

Intel® Ethernet Connection I217

Datasheet v2.3

Product Features

■ General

- 10 BASE-T IEEE 802.3 specification conformance
- 100 BASE-TX IEEE 802.3 specification conformance
- 1000 BASE-T IEEE 802.3 specification conformance
- Energy Efficient Ethernet (EEE) IEEE 802.3az support [Low Power Idle (LPI) mode]
- IEEE 802.3u auto-negotiation conformance
- Supports carrier extension (half duplex)
- Loopback modes for diagnostics
- Advanced digital baseline wander correction
- Automatic MDI/MDIX crossover at all speeds of operation
- Automatic polarity correction
- MDC/MDIO management interface
- Flexible filters in PHY to reduce integrated LAN controller power
- Smart speed operation for automatic speed reduction on faulty cable plants
- PMA loopback capable (no echo cancel)
- 802.1as/1588 conformance
- Power Optimizer Support
- Intel® Stable Image Platform Program (SIPP)
- iSCSI Boot
- Network proxy/ARP Offload support

■ Security & Manageability

- Intel® vPro support with appropriate Intel chipset components

■ Performance

- Jumbo Frames (up to 9 kB)
- 802.1Q & 802.1p
- Receive Side Scaling (RSS)
- Two Queues (Tx & Rx)

■ Power

- Reduced power consumption during normal operation and power down modes
- Integrated Intel® Auto Connect Battery Saver (ACBS)
- Single-pin LAN Disable for easier BIOS implementation
- Fully integrated Switching Voltage Regulator (iSVR)
- Low Power LinkUp (LPLU)

■ MAC/PHY Interconnect

- PCIe-based interface for active state operation (S0 state)
- SMBus-based interface for host and management traffic (Sx low power state)

■ Package/Design

- 48-pin package, 6 x 6 mm with a 0.4 mm lead pitch and an Exposed Pad* for ground
- Three configurable LED outputs
- Integrated MDI interface termination resistors to reduce BOM costs
- Reduced BOM cost by sharing SPI flash with PCH



Legal Information

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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document 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. Intel and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2009-2013, Intel Corporation. All Rights Reserved.



Revision	Date	Description
0.25	April 2011	Initial release.
0.26	May 2011	<ul style="list-style-type: none"> Changed signal name "VDD1P0" to "VDD0P9" Changed signal name "CTRL1P0" to "CTRL0P9"
0.8	September 2011	Added: Features list, Introduction, Interconnects, Initialization, Power Management and Delivery, Device Functionality, Phy Functionality, Programmers' Visible State, Electrical and Timing Specifications, 1588 Information, Reference Schematics information, Schematic and Board Layout Checklist Information, and Mobile and Non-Mobile Design Considerations and Guidelines.
0.9	December 2011	Updated: name for Power Optimizer, Figure1 Block Diagram, Table 1 (Power Optimizer information), Status Packet from PHY information, Figure 3 and notes, VDD0P9 signal names, VDD3P3 signal names, Pinout diagram, Ball Mapping table (signal names), Table 26 Address Map, Table 30 PHY Identifier Register 1, Table 83 LANWAKE #, Figure 12 Power Delivery Schematic, Interconnects diagram Chapter 15 and Chapter 16, Power Delivery schematic Chapters 15 and 16.
0.93	April 2012	Added: Reference to <i>Intel® Ethernet Controller I218 Datasheet</i> ; Notice that only the internal Voltage Regulator will be supported Updated: Legal Information front matter; Table 1 SKUs and Features; Table 15-10 Crystal Manufacturers and Part Numbers; Table 15-13 Inductors and Manufacturers; Table 16-10 Crystal Manufacturers and Part Numbers; Table 16-13 Inductors and Manufacturers
0.94	July 2012	<ul style="list-style-type: none"> Included up-to-date power table Added note to inform readers that an external voltage regulator is not to be used Removed all information and references regarding the use of an external voltage regulator
0.95	August 2012	<ul style="list-style-type: none"> Updated Power Delivery Diagram Updated Power and Power Control Pins table Updated Ball Mapping table Updated Pinout diagram Updated SMBus Interface Pins table
0.96	November 2012	<ul style="list-style-type: none"> Updated e-PAD information Removed all references to shared SVR
2.0	April 2013	<ul style="list-style-type: none"> Initial release (Intel public).¹
2.1	August 2013	<ul style="list-style-type: none"> Added section 6.3.1.9.
2.2	September 2013	<ul style="list-style-type: none"> Revised section 6.3.1.4.2 (EEE LPI Unsupported Features; added network proxy offload bullet).
2.3	October 2013	<ul style="list-style-type: none"> Revised section 7.5 (Network Proxy Functionality; added a note to section 7.5.1).

1. Revision number changed to 2.0 at product release. No other versions have been released between revisions 0.96 and 2.0.



NOTE: *This page intentionally left blank.*



Contents

Contents	v
1 Introduction	1
1.1 Overview	1
1.2 Main Flows	2
1.3 References	2
1.4 Product SKUs and Device IDs	3
2 Interconnects	5
2.1 Introduction.....	5
2.2 PCIe-Based	5
2.2.1 PCIe Interface Signals.....	5
2.2.2 PCIe Operation and Channel Behavior	6
2.3 SMBus.....	9
2.3.1 Overview.....	9
2.3.2 SMBus Pass Trough (PT).....	13
2.3.3 Slave Transactions	15
2.3.4 Master Transactions	18
2.4 Transitions between SMBus and PCIe interfaces	21
2.4.1 Switching from SMBus to PCIe	21
2.4.2 Switching from PCIe to SMBus	21
2.5 SMBus/PCIe Interconnects	21
3 Pin Interface	23
3.1 Pin Assignment	23
3.1.1 Signal Type Definitions.....	23
3.1.2 PCIe Interface Pins (8).....	23
3.1.3 SMBus Interface Pins (2).....	23
3.1.4 Miscellaneous Pins (3)	24
3.1.5 PHY Pins (14).....	24
3.1.6 Testability Pins (5).....	25
3.1.7 Power and Power Control Pins (16)	25
3.2 Pinout	26
4 Package	27
4.1 Package Type and Mechanical	27
4.2 Package Electrical and Thermal Characteristics	28
4.3 Power and Ground Requirements	28
4.4 Ball Mapping.....	28
5 Initialization	29
5.1 Power Up.....	29
5.2 Reset Operation	30
5.3 Timing Parameters	31
5.3.1 Timing Requirements	31
5.3.2 Timing Guarantees.....	31
6 Power Management and Delivery	33
6.1 Power Information	33



- 6.2 Power Delivery 34
 - 6.2.1 Core Vdc Supply 34
- 6.3 Power Management 34
 - 6.3.1 Global Power States 34
- 7 Device Functionality 41**
 - 7.1 Tx Flow 41
 - 7.2 Rx Flow 41
 - 7.3 Flow Control 42
 - 7.3.1 MAC Control Frames and Reception of Flow Control Packets 42
 - 7.3.2 Transmitting PAUSE Frames 44
 - 7.4 Wake Up 44
 - 7.4.1 Host Wake Up 45
 - 7.5 Network Proxy Functionality 55
 - 7.5.1 Introduction 55
 - 7.5.2 Network Proxy Activation 55
 - 7.5.3 IPv4 Proxy—ARP 56
 - 7.5.4 IPv6 Proxy—Neighbor Discovery 57
 - 7.5.5 Multicast Listener Discovery Support 60
 - 7.6 Loopback 62
- 8 Programmers' Visible State 65**
 - 8.1 Terminology 65
 - 8.2 MDIO Access 66
 - 8.3 Addressing 66
 - 8.3.1 Registers Byte Ordering 67
 - 8.4 Address Map 68
 - 8.5 PHY Registers (Page 0) 71
 - 8.5.1 Loopback Mode Settings 79
 - 8.6 Port Control Registers (Page 769) 90
 - 8.7 Statistics Registers 93
 - 8.8 PCIe Registers 95
 - 8.9 General Registers 96
 - 8.9.1 Interrupts 102
 - 8.10 Wake Up Registers 102
 - 8.10.1 Accessing Wake Up Registers Using MDIO 102
 - 8.10.2 Host Wake Up Control Status Register Description 103
 - 8.10.3 ME Wake Up Control Status Register Description 115
 - 8.10.4 LPI MMD PHY Registers 126
- 9 Non-Volatile Memory (NVM) 127**
 - 9.1 Introduction 127
 - 9.2 NVM Programming Procedure Overview 127
 - 9.3 LAN NVM Format and Contents 129
 - 9.3.1 Hardware Accessed Words 130
 - 9.3.2 Software Accessed Words 137
 - 9.3.3 Basic Configuration Software Words 141
- 10 Electrical & Timing Specifications 143**
 - 10.1 Introduction 143



10.2	Operating Conditions.....	143
10.2.1	Absolute Maximum Ratings.....	143
10.2.2	Recommended Operating Conditions	143
10.3	Power Delivery	144
10.3.1	Voltage Regulator Power Supply Specifications	144
10.3.2	SVR Specification (Internal)	144
10.3.3	Power On/Off Sequence.....	144
10.3.4	Power Delivery Schematic Drawing	145
10.4	I/O DC Parameter.....	146
10.4.1	3.3 Vdc I/O (Open Drain).....	146
10.4.2	3.3 Vdc I/O.....	146
10.4.3	3.3 Vdc I/O Pull-up & Pull-down Information	146
10.4.4	Input Buffer Only	147
10.4.5	Input Buffer Vdc I/O Pull-up & Pull-down.....	147
10.4.6	PCIe DC/AC Specifications	147
10.5	Discrete/Integrated Magnetics Specifications	150
10.6	Mechanical	150
10.7	Oscillator/Crystal Specifications	151
11	Driver Programming Interface	153
11.1	Registers Byte Ordering	153
11.1.1	Register Conventions	154
11.1.2	PCI Configuration and Status Registers—CSR Space	154
11.1.3	PCIm Configuration and Status Registers—CSR Space	237
12	Time Synch	
	(IEEE1588 and 802.1AS)	281
12.1	Overview	281
12.1.1	Flow and HW/SW responsibilities.....	281
12.1.2	HW Time sync elements.....	283
12.1.3	PTP packet structure.....	285
13	Mobile Design	
	Considerations & Guidelines	289
13.1	PHY Overview	290
13.1.1	PHY Interconnects.....	290
13.1.2	PCIe-Based Interface	291
13.1.3	SMBus Interface	291
13.1.4	PCIe and SMBus Modes	292
13.1.5	Transitions between PCIe and SMBus Interfaces.....	292
13.2	Platform LAN Design Guidelines.....	293
13.2.1	General Design Considerations for PHYs	293
13.2.2	NVM for PHY Implementations.....	295
13.2.3	LAN Switch	296
13.2.4	LED.....	296
13.2.5	Exposed Pad* (e-PAD) Design and SMT Assembly Guide.....	297
13.3	PCH—SMBus/PCIe LOM Design Guidelines.....	302
13.4	SMBus Design Considerations.....	303
13.5	General Layout Guidelines	304
13.6	Layout Considerations.....	304
13.7	Guidelines for Component Placement.....	304
13.7.1	PHY Placement Recommendations	305



- 13.8 MDI Differential-Pair Trace Routing for LAN Design 307
- 13.9 Signal Trace Geometry 307
- 13.10 Trace Length and Symmetry 309
- 13.11 Impedance Discontinuities 310
- 13.12 Reducing Circuit Inductance 310
- 13.13 Signal Isolation 310
- 13.14 Power and Ground Planes 311
- 13.15 Traces for Decoupling Capacitors 313
- 13.16 Ground Planes under a Magnetics Module 313
- 13.17 Light Emitting Diodes 315
- 13.18 Considerations for Layout 315
- 13.19 Frequency Control Device Design Considerations 316
- 13.20 Crystals and Oscillators 316
- 13.21 Quartz Crystal 316
- 13.22 Fixed Crystal Oscillator 317
- 13.23 Crystal Selection Parameters 317
- 13.24 Vibrational Mode 317
- 13.25 Nominal Frequency 317
- 13.26 Frequency Tolerance 318
- 13.27 Temperature Stability and Environmental Requirements 318
- 13.28 Calibration Mode 318
- 13.29 Load Capacitance 319
- 13.30 Shunt Capacitance 319
- 13.31 Equivalent Series Resistance 319
- 13.32 Drive Level 320
- 13.33 Aging 320
- 13.34 Reference Crystal 320
 - 13.34.1 Reference Crystal Selection 320
 - 13.34.2 Circuit Board 321
 - 13.34.3 Temperature Changes 321
- 13.35 Oscillator Support 321
- 13.36 Oscillator Placement and Layout Recommendations 322
- 13.37 LAN Switch 322
- 13.38 Troubleshooting Common Physical Layout Issues 323
- 13.39 Power Delivery 323
- 13.40 Power Sequencing 325
- 14 Non-Mobile Design Considerations & Guidelines 327**
 - 14.1 PHY Overview 328
 - 14.1.1 PHY Interconnects 328
 - 14.1.2 PCIe-Based Interface 329
 - 14.1.3 SMBus Interface 329
 - 14.1.4 PCIe and SMBus Modes 330
 - 14.1.5 Transitions Between PCIe and SMBus Interfaces 330
 - 14.2 Platform LAN Design Guidelines 331
 - 14.2.1 General Design Considerations for PHYs 331
 - 14.2.2 NVM for PHY Implementations 333
 - 14.2.3 LED 333
 - 14.2.4 Exposed Pad* (e-PAD) Design and SMT Assembly Guide 334



14.3	PCH—SMBus/PCIe LOM Design Guidelines.....	339
14.4	SMBus Design Considerations.....	340
14.5	General Layout Guidelines.....	341
14.6	Layout Considerations.....	341
14.7	Guidelines for Component Placement.....	341
14.7.1	PHY Placement Recommendations.....	342
14.8	MDI Differential-Pair Trace Routing for LAN Design.....	343
14.9	Signal Trace Geometry.....	344
14.10	Trace Length and Symmetry.....	346
14.11	Impedance Discontinuities.....	347
14.12	Reducing Circuit Inductance.....	347
14.13	Signal Isolation.....	347
14.14	Power and Ground Planes.....	348
14.15	Traces for Decoupling Capacitors.....	350
14.16	Ground Planes Under a Magnetics Module.....	350
14.17	Light Emitting Diodes.....	352
14.18	Frequency Control Device Design Considerations.....	352
14.19	Crystals and Oscillators.....	353
14.20	Quartz Crystal.....	353
14.21	Fixed Crystal Oscillator.....	353
14.22	Crystal Selection Parameters.....	354
14.23	Vibrational Mode.....	354
14.24	Nominal Frequency.....	354
14.25	Frequency Tolerance.....	354
14.26	Temperature Stability and Environmental Requirements.....	354
14.27	Calibration Mode.....	355
14.28	Load Capacitance.....	355
14.29	Shunt Capacitance.....	356
14.30	Equivalent Series Resistance.....	356
14.31	Drive Level.....	356
14.32	Aging.....	356
14.33	Reference Crystal.....	356
14.33.1	Reference Crystal Selection.....	357
14.33.2	Circuit Board.....	357
14.33.3	Temperature Changes.....	357
14.34	Oscillator Support.....	358
14.35	Oscillator Placement and Layout Recommendations.....	358
14.36	Troubleshooting Common Physical Layout Issues.....	359
14.37	Power Delivery.....	360
14.38	Intel® Ethernet Controller I217 Power Sequencing.....	361
15	Checklists.....	363
16	Models.....	365
17	Reference Schematics.....	367



Note: This page intentionally left blank.



1 Introduction

1.1 Overview

The Intel® Ethernet Connection I217 is a single-port Gigabit Ethernet Physical Layer Transceiver (PHY). It connects to an integrated Media Access Controller (MAC) through a dedicated interconnect. The I217 supports operation at 1000/100/10 Mb/s data rates. The PHY circuitry provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3, 802.3u, and 802.3ab). The I217 also supports the Energy Efficient Ethernet (EEE) 802.az specification.

Note: 1 Gb/s half-duplex operation is not supported.

The I217 is packaged in a small footprint QFN package. Package size is 6 x 6mm with a 0.4mm lead pitch and a height of 0.85mm, making it very attractive for small form-factor platforms.

The I217 interfaces with its MAC through two interfaces: PCIe*-based and SMBus. The PCIe (main) interface is used for all link speeds when the system is in an active state (S0) while the SMBus is used only when the system is in a low power state (Sx). In SMBus mode, the link speed is reduced to 10 Mb/s (dependent on low power options). The PCIe interface incorporates two aspects: a PCIe SerDes (electrically) and a custom logic protocol.

Note: The I217 PCIe interface is not PCIe compliant. It operates at half of the PCI Express* (PCIe*) Specification v1.1 (2.5 GT/s) speed. In this datasheet the term PCIe-based is interchangeable with PCIe. There are no design layout differences between normal PCIe and the Intel® Ethernet Controller I217's PCIe-based interface.

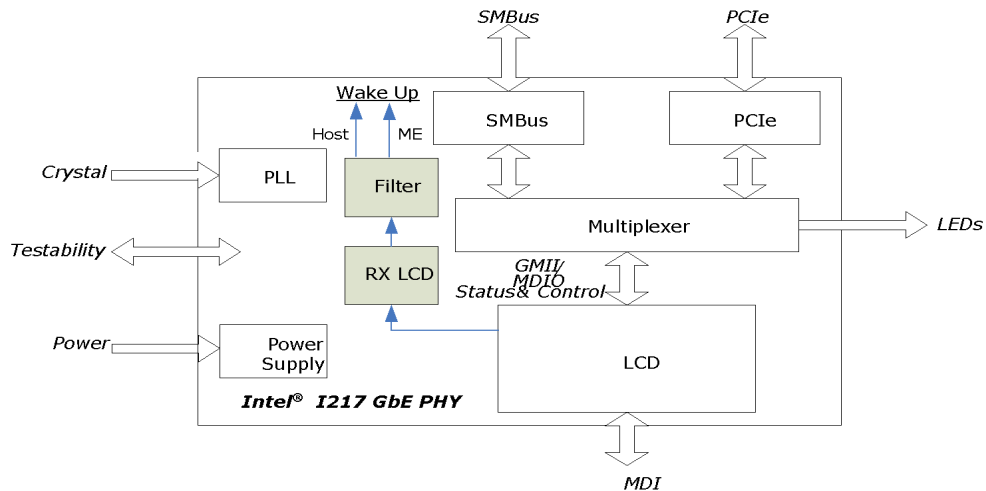


Figure 1-1. I217 Block Diagram



1.2 Main Flows

The main Intel® I217 GbE PHY interfaces are PCIe and SMBus on the host side and the MDI interface on the link side. Transmit traffic is received from the MAC device through either PCIe or SMBus on the host interconnect and then transmitted on the MDI link. Receive traffic arrives on the MDI link and transferred to the integrated LAN controller through either the PCIe or SMBus interconnects.

The integrated LAN controller and system software control the Intel® Ethernet Controller I217's functionality through two mechanisms:

- The I217 configuration registers are mapped into the MDIO space and can be accessed by the integrated LAN controller through the PCIe or SMBus interconnects.
- The MDIO traffic is embedded in specific fields in SMBus packets or carried by special packets over the PCIe encoded interconnect as defined by the custom protocol.

Specific flows are described in other sections of this document:

The power-up flow is described in [Section 5.1](#) .

Reset flows are described in [Section 5.2](#) .

Power delivery options are described in [Section 6.2](#) . Power management is described in [Section 6.3](#) .

1.3 References

- Information Technology—Telecommunication & Information Exchange Between Systems—LAN/MAN—Specific Requirements—Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, IEEE Standard No.: 802.3-2008.
- Energy Efficient Ethernet (EEE) 802.az specification.
- SMBus specification revision 2.0.



1.4 Product SKUs and Device IDs

Table 1 lists the product SKUs and Device IDs for the Intel® Ethernet Connection I217 GbE controller. Refer to the *Intel® I217 Specification Update* for device ordering information.

Table 1. Product SKU and Feature Matrix

					Performance				Extended Power Savings				Advanced Features							
Method of enabling/disabling features in SKUs =>					Driver	Driver	MAC	MAC	PHY	Driver	Software	PHY	Platform	HW	HW	Driver	Platform	BIOS	FW	FW
Link Speed	Platform	Segment	Product/Code Name	Device ID	Jumbo Frames (up to 9k)	802.1Q & 802.1p	Receive Side Scaling (RSS)	2 Tx & 2 Rx Queues	Low 'No-Link' Power (ACBS)	Link Speed Battery Saver	Server Operating System Support	Low Power Linkup (LPLU)	Power Optimizer	802.1as	EEE - 802.3az	Ability to Initiate a Team	Intel® SIPP	iSCSI Boot	Standard Manageability	Intel® vPro*
Gigabit	Ultrabook Mobile	Corporate	I217-LM (Lynx Point-H+I217LM)	153A	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	Desktop Embedded	Consumer	I217-V (Lynx Point-H+I217V)	153B	X	X	X	X	X	X		X	X	X	X	X			X	
Gigabit	Entry Server Workstation	Corporate	I217-LM (Lynx Point-H+I217LM)	153A	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Notes for Table 1:

- An "X" indicates that the specified feature is to be enabled (that is, available to the user) for the associated SKU.
- The list of features does not include all of the features for the product; the features listed are those enabled/disabled at the SKU level.
- For the list of supported Operating Systems, contact your Intel Field Service Representative.



NOTE: *This page intentionally left blank.*



2 Interconnects

2.1 Introduction

The I217 implements two interconnects to the integrated LAN controller:

- PCIe—A high-speed SerDes interface using PCIe electrical signaling at half speed while keeping the custom logical protocol for active state operation mode.
- System Management Bus (SMBus)—A very low speed connection for low power state mode for manageability communication only. At this low power state mode the Ethernet link speed is reduced to 10 Mb/s.

System	LAN Connected Device	
	SMBus	PCIe
S0 and PHY Power Down	Not used	Idle
S0 and Idle or Link Disc	Not used	Idle
S0 and Link in Low Power Idle (LPI)	Not used	Idle
S0 and active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

The I217 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

2.2 PCIe-Based

Note: The I217 PCIe interface is not PCIe compliant. It operates at half of the PCI Express* (PCIe*) Specification v1.1 (2.5 GT/s) speed. In this document the term PCIe-based is interchangeable with PCIe. There are no design layout differences between normal PCIe and the LCD's PCIe-based interface. Standard PCIe validation tools cannot be used to validate this interface. Refer to [Section 13.1.2](#) for PCIe-based specifications.

2.2.1 PCIe Interface Signals

The signals used to connect between the integrated LAN Controller and the PHY in this mode are:

- Serial differential pair running at 1.25 Gb/s for Rx
- Serial differential pair running at 1.25 Gb/s for Tx
- 100 MHz differential clock input to the PHY running at 100 MHz
- Power and clock good indication to the PHY PE_RSTn
- Clock control through CLKREQn



2.2.2 PCIe Operation and Channel Behavior

The I217 runs at 1250 Mb/s KX speed, which is 1/2 of the gen1 2.5 Gb/s PCIe frequency. To operate with KX only devices, each of the PCIe root ports in the PCH-integrated MAC have the ability to run at the KX rate. There is no need to implement a mechanism to detect that a KX only device is attached. The port configuration (if any), attached to a KX only device, is pre-loaded from NVM. The selected port adjusts the transmitter to run at the KX rate and does not need to be PCIe compliant.

Packets transmitted and received over the PCIe interface are full Ethernet packets and not PCIe transaction/link/physical layer packets.

After the PCIe power-up sequence completes, each transmitter starts transmitting idle symbols and the receiver acquires synchronization as specified in 802.3z.

2.2.2.1 PCIe In-Band Messages

In-band messages are used to transfer control information between the Intel® Ethernet Controller I217 and the integrated LAN Controller. The Intel® Ethernet Controller I217 only initiates PHY status in-bands and then waits for an acknowledge. For each in-band message on PCIe, there is an equivalent message on the SMBus. As a result, if an interface switch took place before an acknowledgement was received, the equivalent message is sent on the other interface.

2.2.2.1.1 MDIO Access Packet Transmitted by integrated LAN Controller

This in-band message is equivalent to the Mtp Configuration command on the SMBus.

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address and Controls	Bit 7 0b => MDIO register. Bit 6 0b => Command. Bit 5 0b => Read access. 1b => Write access. Bits 4:0 Register address.
1st Data Byte	MSB data byte if write. Must be 0b if read.
2nd Data Byte	LSB data byte if write. Must be 0b if read.
PHY Address	Bits 7:5 Reserved (0b). Bits 4:0 PHY address.
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.



2.2.2.1.2 MDIO Access Acknowledge/Response Packet Transmitted by PHY

This in-band message is equivalent to the PtM Configuration Acknowledge command on the SMBus.

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address and Controls	Bit 7 0b => MDIO register. Bit 6 1b => Acknowledge. Bit 5: 0b => Read access. 1b => Write access. Bits 4:0 Register address.
1st Data Byte	Data MSB.
2nd Data Byte	Data LSB.
Reserved Byte	Reserved byte (must be 0b).
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.

2.2.2.1.3 Status Packet from PHY

This in-band message is equivalent to the PtM PHY Status command on the SMBus. The status command is sent by the Intel® Ethernet Controller I217 on every status change at the PHY side or when a time out defined in the In-Band Control register has expired. The 217 re-transmits the Status command if no acknowledge arrived after a time out defined in the PCIe Diagnostic register has expired.

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address and & Controls	Bit 7 1b => Status. Bit 6 0b Bit 5:0 Reserved (must be 0x0).
1st Data Byte	Bits 7:6 Reserved (must be 0x0). Bit 5 Inband Host WoL indication Bit 4 Reserved. Bit 3 Reserved (must be 0b). Bit 2 EI entry request. Bit 1 TX Off (MAC back pressure required). Bit 0 K1 entry request.
2nd Data Byte	Bit 7 Reset complete. Bit 6 Interrupt request (Not used) Bit 5 GLCI link status: 1b => Link up. 0b => Link down. Bit 4 PHY cable disconnected. Bit 3 PHY Link status: 1b => Link up. 0b => Link down. Bit 2 Duplex mode: 1b => Full duplex. 0b => Half duplex. Bits 1:0 Speed: 11b => Reserved. 10b => 1000 Mb/s. 01b => 100 Mb/s. 00b => 10 Mb/s.
Reserved Byte	Reserved byte (must be 0b).



Byte	Description
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.

2.2.2.1.4 Acknowledge Packet From the integrated LAN Controller

This in-band message is equivalent to the MtP PHY Status Acknowledge command on the SMBus.

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address and Controls	Bit 7 1b => Status. Bit 6 1b => Acknowledge. Bits 5:0 Reserved (must be 0x0).
1st Data Byte	Bits 7:6 Reserved (must be 0x0). Bit 5 Inband Host WoL indication Bit 4 Reserved. 1b =>RX in EEE LPI (Low Power Idle) state indication. 0b =>RX Active indication. Bit 3 Reserved (must be 0b). Bit 2 EI entry request. Bit 1 TX Off (MAC back pressure required). Bit 0 K1 entry request.
2nd Data Byte	Bit 7 Reset complete. Bit 6 Interrupt request (Not used) Bit 5 GLCI link status: 1b => Link up. 0b => Link down. Bit 4 PHY cable disconnected. Bit 3 PHY Link status: 1b => Link up. 0b => Link down. Bit 2 Duplex mode: 1b => Full duplex. 0b => Half duplex. Bits 1:0 Speed: 11b => Reserved. 10b => 1000 Mb/s. 01b => 100 Mb/s. 00b => 10 Mb/s.
Reserved Byte	Reserved Byte (must be 0x0).
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.



2.2.2.1.5 Status Packet From the Integrated LAN Controller

This in-band message is equivalent to the MtP Control command on the SMBus.

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address & Controls	Bit 7 1b Status. Bit 6 0b Bits 5:0 Reserved (must be 0x0).
1st Data Byte	Bits 7:6 Reserved (must be 0x0). Bit 5 Reserved. Bit 4 Reserved. Bit 3 XON request. Bit 2 EI entry request. Bit 1 XOFF request. Bit 0 K1 entry request.
2nd Data Byte	Bit 7 Port Reset. Bit 6 Power down. Bits 5:4 Power saving D-states. Bits 3:0 MAC LED Status (not used).
Reserved Byte	Reserved Byte (must be 0b).
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.

2.2.2.1.6 Acknowledge Packet From the PHY

Byte	Description
Special MDIO Symbol	/K28.1/
Register Address and Controls	Bit 7 1b => Status. Bit 6 1b => Acknowledge. Bits 5:0 Reserved (must be 0x0).
1st Data Byte	Bits 7:3 Reserved (must be 0x0). Bit 2:0 Reserved (must be 0x0)
2nd Data Byte	Bits 7:0 Reserved (must be 0x0).
Reserved Byte	Reserved byte (must be 0b).
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.

2.3 SMBus

2.3.1 Overview

SMBus is a low speed (100 kHz/400 kHz/1000 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the Intel® Ethernet Controller I217 and the integrated LAN Controller when the system is in a low power Sx state. The interface is also used to enable the integrated LAN Controller to configure the Intel® Ethernet Controller I217 as well as passing in-band information between them.

The SMBus uses two primary signals: SMBCLK and SMBDAT, to communicate. Both of these signals are open drain with board-level pull-ups.



The SMBus protocol includes various types of message protocols composed of individual bytes. The message protocols supported by the Intel® Ethernet Controller I217 are described in the relevant sections.

For further details on SMBus behavior, see the SMBus specification.

2.3.1.1 SMBus Channel Behavior

The SMBus specification defines the maximum frequency of the SMBus as 100 kHz or 1000 kHz. When operating at 1000 kHz the SMBus specification parameters are defined by the I2C specification. To change the Intel® Ethernet Controller I217's SMBus frequency to 1000 kHz, bit 12,8 in the SMBus Address register (register 26, address 01, page 0) should be set to 10b.

2.3.1.2 SMBus Addressing

The 217's address is assigned using SMBus ARP protocol. The default SMBus address is 0xC8.

2.3.1.3 Bus Time Outs

The 217 can detect (as a master or a slave) an SMBCLK time out on the main SMBus. If the SMBus clock line is held low for less than 25 ms, the Intel® Ethernet Controller I217 does not abort the transaction. If the SMBus clock line is held low for 25 ms or longer, the Intel® Ethernet Controller I217 aborts the transaction.

As a slave, the I217 detects the time out and goes into an idle state. In idle, the slave releases the SMB_CLK and SMB_DATA lines. Any data that was received before the time out might have been processed depending on the transaction.

As a master, the I217 detects a time out and issues a STOP on the SMBus at the next convenient opportunity and then brings the SMBus back to idle (releases SMB_CLK and SMB_DATA). Any master transaction on which the I217 detects a time out is aborted.

2.3.1.4 Bus Hangs

Although uncommon, SMBus bus hangs can happen in a system. The catalyst for the hang is typically an unexpected, asynchronous reset or noise coupled onto the SMBus. Slaves can contribute to SMBus hangs by not implementing the SMBus time outs as specified in SMBus 2.0 specification. Masters or host masters can contribute to SMBus hangs by not detecting the failures and by not attempting to correct the bus hangs.

Because of the potential bus hang scenario, the I217 has the capability of detecting a hung bus. If SMB_CLK or SMB_DATA are stuck low for more than 35 ms, the I217 forces the bus to idle (both SMB_CLK and SMB_DATA set), if it is the cause of the bus hang.

2.3.1.5 Packet Error Code (PEC) Support

PEC is defined in the SMBus 2.0 specification. It is an extra byte at the end of the SMBus transaction, which is a CRC-8 calculated on all of the preceding bytes (not including ACKs, NACKs, STARTs, or STOPs) in the SMBus transaction. The polynomial for this CRC-8 is:

$$x^8 + x^2 + x + 1$$



The PEC calculation is reset when any of the following occurs:

- A STOP condition is detected on the host SMBus
- An SMBus hang is detected on the host SMBus
- The SMBCLK is detected high for $\sim 50 \mu\text{s}$

2.3.1.6 SMBus ARP Functionality

The I217 supports the SMBus ARP protocol as defined in the SMBus 2.0 specification. The I217 is a persistent slave address device, meaning its SMBus address is valid after power up (constant 0xC8). The I217 supports all SMBus ARP commands defined in the SMBus specification, both general and directed.

2.3.1.6.1 SMBus ARP Flow

SMBus ARP flow is based on the status of two flags:

- AV-Address Valid—This flag is set when the I217 has a valid SMBus address.
- AR-Address Resolved—This flag is set when the I217 SMBus address is resolved (SMBus address was assigned by the SMBus ARP process).

Note: These flags are internal Intel® Ethernet Controller I217 flags and are not shown to external SMBus devices.

Since the I217 is a Persistent SMBus Address (PSA) device, the AV flag is always set, while the AR flag is cleared after power up until the SMBus ARP process completes. Since AV is always set, it means that the I217 always has a valid SMBus address.

When the SMBus master wants to start a SMBus ARP process, it resets (in terms of ARP functionality) all the devices on the SMBus by issuing either Prepare to ARP or Reset Device commands. When the I217 accepts one of these commands, it clears its AR flag (if set from previous SMBus ARP process), but not its AV flag; the current SMBus address remains valid until the end of the SMBus ARP process.

With the AR flag cleared, the I217 answers the following SMBus ARP transactions that are issued by the master. The SMBus master then issues a Get UDID command (General or Directed) to identify the devices on the SMBus. The I217 responds to the Directed command every time and to the General command only if its AR flag is not set. After a Get UDID command, the master assigns the I217 an SMBus address by issuing an Assign Address command. The I217 checks whether the UDID matches its own UDID and if matched, switches its SMBus address to the address assigned by the command (byte 17). After accepting the Assign Address command, the AR flag is set, and from this point on (as long as the AR flag is set) the I217 does not respond to the Get UDID General command, while all other commands should be processed even if the AR flag is set. [Figure 2-1](#) shows the SMBus ARP behavior of the I217.




Figure 2-1. SMBus ARP Flow


2.3.1.6.2 SMBus ARP UDID Content

The Unique Device Identifier (UDID) provides a mechanism to isolate each device for the purpose of address assignment. Each device has a unique identifier. The 128-bit number is comprised of the following fields:

1 Byte	1 Byte	2 Bytes	2 Bytes	2 Bytes	2 Bytes	2 Bytes	4 Bytes
Device Capabilities	Version / Revision	Vendor ID	Device ID	Interface	Subsystem Vendor ID	Subsystem Device ID	Vendor Specific ID
See below	See below	0x8086	0x10D4	0x0004	0x0000	0x0000	See below
MSB							LSB

Where:

- Vendor ID—The device manufacturer's ID as assigned by the SBS Implementers' Forum or the PCI SIG. Constant value: 0x8086.
- Device ID—The device ID as assigned by the device manufacturer (identified by the Vendor ID field). The value for the LAN Connected device is constant: 0x10D4.
- Interface the protocol layer interfaces supported over the SMBus connection by the device. In this case, SMBus Version 2.0. Constant value: 0x0004.



- Subsystem Fields—These fields are not supported and return zeros.
- Device Capabilities: Dynamic and persistent address, PEC support bit:

7		6		5		4		3		2		1		0	
Address Type		Reserved		Reserved		Reserved		Reserved		Reserved		Reserved		PEC Supported	
0b		1b		0b		0b		0b		0b		0b		1b	
MSB														LSB	

- Version/Revision: UDID Version 1, Silicon Revision:

7		6		5		4		3		2		1		0	
Reserved		Reserved		UDID Version				Silicon Revision ID							
0b		0b		001b				See below							
MSB														LSB	

- Silicon Revision ID:

Silicon Version								Revision ID							
A0								000b							

Vendor Specific ID: Four bytes—constant 0x44332211:

1 Byte				1 Byte				1 Byte				1 Byte			
44				33				22				1x ¹			
MSB												LSB			

Notes:

1. Bit 0 value is defined by the value sampled at reset on GPIO[0].

2.3.1.7 SMBus ARP Transactions

Note: All SMBus ARP transactions include a PEC byte. For the layout of these transactions refer to the SMBus 2.0 specification.

Supported SMBus ARP transactions:

- Prepare to ARP
- Reset Device (General and Directed)
- Assign Address
- Get UDID (General and Directed)

2.3.2 SMBus Pass Trough (PT)

This section describes the SMBus transactions supported as PT between the integrated LAN Controller and the I217. All traffic between the I217 and the integrated LAN Controller is in network order (the MSB is sent first).



2.3.2.1 Receive Flow

The maximum SMBus fragment length is configurable and can either be 32 bytes or 64 bytes. The default value is 32 bytes. The incoming packets are divided into fragments, where the I217 uses the maximum fragment size allowed in each fragment. The data of the packet is transferred using the Receive TCO packet transaction as described in [Figure 2.3.4.1.1](#)

Any time out on the SMBus results in discarding the entire packet. Any NACK by the integrated LAN Controller on one of the I217 receive bytes causes the packet to be re-transmitted up to four times. If after four times the packet fails to be transmitted, it is silently discarded.

The maximum size of the received packet is limited by the I217 hardware to 1522 bytes. Packets larger than 1522 bytes are silently discarded.

2.3.2.2 Transmit Flow

The I217 is used as a channel for transmitting packets from the integrated LAN Controller to the network link. The network packet is transferred from the integrated LAN Controller over the SMBus (starting with a preamble), and then, when fully received by the I217, it is transmitted over the network link.

The I217 supports packets up to an Ethernet packet length of 1522 bytes. SMBus transactions are configurable up to a 64-byte length, which means that packets can be transferred over the SMBus in more than one fragment. Fragments within a packet are marked with the F and L flags. The I217 does not change any field in the transmitted packet. A packet bigger than 1522 bytes is silently discarded by the integrated LAN Controller and not sent to the I217.

If the network link is down when the I217 is receiving SMBus fragments of the packet, it silently discards the packet.

The transmit SMBus transaction is described in [Section 2.3.1.6.1](#) .

2.3.2.2.1 Transmit Errors in Sequence Handling

Once a packet is transferred over the SMBus from the integrated LAN Controller to the I217 the F and L flag should follow specific rules. The F flag defines that this is the first fragment of the packet, and the L flag defines that the transaction contains the last fragment of the packet.

Table 1. Flag options in transmit packet transactions:

Previous	Current	Action/Notes
Last	First	Accept both.
Last	Not First	Error for current transaction. All fragments, until one with the F flag set, are discarded, unless the current fragment is a Single.
Not Last	First	Error for the previous transaction. Previous packet is sent with a bad CRC. No abort status is asserted, unless the previous fragment is a Single.
Not Last	Not First	Process the current transaction. In case of Single after Middle error for the previous transaction and process the current transaction.



Note that since every other Block Write command in the messaging protocol has both first and last flags off, they cause flushing any pending transmit fragments that were previously received. In other words, when running the transmit flow, no other Block Write transactions are allowed in between the fragments.

2.3.2.3 Concurrent SMBus Transactions

Concurrent SMBus transactions (receive, transmit and configuration read/write) on opposite directions are allowed and there is no limitation for it. Transmit fragments can be sent between receive fragments and configuration Read/Write commands. Acknowledges can also be issued between receive and transmit fragments in the opposite direction.

2.3.3 Slave Transactions

2.3.3.1 SMBus Transactions (Integrated LAN Controller to the I217)

The following table lists the slave SMBus transactions supported by the I217.

Table 2. SMBus Slave Transactions

Command	Transaction	Command		Fragmentation	Section
Transmit Packet	Block Write	First	0x84	Multiple	2.3.3.1.1
		Middle	0x04	Multiple	
		Last	0x44	Multiple	
		Single	0xC4	Single	
MtP Control	Block Write	Single	0xC3	Single	2.3.3.1.2
MtP Status Acknowledge	Block Write	Single	0xC3	Single	2.3.3.1.3
MtP configuration	Block Write	Single	0xC3	Single	2.3.3.1.4

2.3.3.1.1 Transmit Packet Command

The transmit packet fragments have the following format:

Function	Command	Byte Count	Data 1	...	Data N
Transmit first fragment	0x84	N	Packet data MSB	...	Packet data LSB
Transmit middle fragment	0x04				
Transmit last fragment	0x44				
Transmit single fragment	0xC4				

If the overall packet length is bigger than 1522 bytes, the packet is silently discarded by the integrated LAN Controller.



2.3.3.1.2 MtP Control Command

The MtP Control command is a single fragment command enabling the integrated LAN Controller to send messages to the I217 informing status changes or sending directed control commands (not through registers).

MtP Control command format:

Function	Command	Byte Count	Data 1	Data 2	Data 3	Data 4
Control	0xC3	4	Address and Control	Command 1st Byte	Command 2nd Byte	Reserved

Table 3. Address and control

Field	Bit(s)	Description
CFG_CTL	7	Configuration/control indication—This bit should be set to 1b indicating control/status.
CMD_ACK	6	Command/acknowledge indication—This bit should be set to 0b indicating command.
Reserved	5:0	Reserved, must be set to 0x0.

Table 4. Command 1st Byte

Field	Bit(s)	Description
Reserved.	7:6	Reserved.
Reserved.	5	Reserved.
Reserved.	4	Reserved.
Reserved	3:0	Reserved.

Table 5. Command 2nd Byte

Field	Bit(s)	Description
RST	7	Port reset.
PWDN	6	Power down.
Reserved	5:1	Reserved.
PINSTOP	0	Clear the LANWAKE# pin indication

2.3.3.1.3 MtP PHY Status Acknowledge Command

The MtP PHY Status Acknowledge command is a single fragment command sent by the integrated LAN Controller as an acknowledge to the I217's PtM Status Command.

MtP Status command format:

Function	Command	Byte Count	Data 1	Data 2	Data 3	Data 4
Status	0xC3	4	Address and Control	PHY Status 1st Byte	PHY Status 2nd Byte	Reserved

**Table 6. Address and control**

Field	Bit(s)	Description
CFG_CTL	7	Configuration or control/status indication. This bit should be set to 1b indicating control/status.
CMD_ACK	6	Command/acknowledge indication. This bit should be set to 1b indicating acknowledge.
Reserved	5:0	Reserved, must be set to 0x0.

Table 7. PHY Status 1st Byte

Field	Bit(s)	Description
Reserved	7:6	Reserved.
Host WoL	5	Inband Host WoL indication
Reserved.	4	Reserved.
Reserved	3:0	Reserved.

Table 8. PHY Status 2nd Byte

Field	Bit(s)	Description
RSTC	7	Reset complete.
Reserved (INT)	6	Interrupt request (not used).
Reserved (KLINK)	5	PCIe link status (not used).
CDIS	4	Ethernet cable disconnected.
ELINK	3	Ethernet Link Status 1b = Link up. 0b = Link down.
DPX	2	Duplex Mode 1b = Full duplex. 0b = Half duplex.
SPD	1:0	Ethernet Link Speed 11b = Reserved. 10b = 1000 Mb/s. 01b = 100 Mb/s. 00b = 10 Mb/s.

2.3.3.1.4 MtP Configuration Command

The Configuration command is a single fragment command enabling the integrated LAN Controller access to all the I217-registers over the SMBus link.

Configuration command format:

Function	Command	Byte Count	Data 1	Data 2	Data 3	Data 4	
Configuration	0xC3	4	Address and Control	MSB Data	LSB Data	Reserved ¹	PHY Address

1. Bits 7:5



Table 9. Address and control

Field	Bit(s)	Description
CFG_CTL	7	Configuration/control indication. This bit should be set to 0b indicating configuration.
CMD_ACK	6	Command/acknowledge indication. This bit should be set to 0b indicating Command.
RW	5	Read/write indication. 1b = Indicates write access. 0b = Indicates read access.
Add	4:0	Register address.

2.3.4 Master Transactions

2.3.4.1 SMBus Transactions (I217 to the integrated LAN Controller)

In order to avoid starvation on the SMBus in the opposite direction, the I217-adds a byte-time idle between any two fragments it sends.

The following table lists the master SMBus transactions supported by the Intel® Ethernet Controller I217.

Table 10. SMBus Master Transactions

Command	Transaction	Command	Fragmentation	Section	
Receive Packet	Block Write	First Middle Last Single	0x90 0x10 0x50 0xD0	Multiple Multiple Multiple Single	2.3.4.1.1
PtM Status	Block Write	Single	0xD3	Single	2.3.4.1.2
PtM Configuration Acknowledge	Block Write	Single	0xD3	Single	2.3.4.1.3
PtM WU	Block Write	Single	0xD5	Single	2.3.4.1.4

2.3.4.1.1 Receive Packet Transaction

When the Intel® Ethernet Controller I217-has a packet to deliver to the integrated LAN Controller, it should begin issuing Receive packet transaction commands using the Block Write protocol. The packet can be delivered in more than one SMBus fragment, and the integrated LAN Controller should follow the fragments order.

The opcode can have these values:

- 0x90—First fragment
- 0x10—Middle fragment
- 0x50—Last fragment of the packet
- 0xD0—Single fragment packet



Function	Byte Count	Data 1 (Op-Code)	Data 2	...	Data N
Receive TCO First Fragment	N	90	Packet Data Byte	...	Packet Data Byte
Receive TCO Middle Fragment		10			
Receive TCO Last Fragment		50			
Receive TCO Single Fragment		D0			

2.3.4.1.2 PtM Status Command

The PtM Status command is a single fragment command sent by the Intel® Ethernet Controller I217 on every status change at the PHY side or when a time out defined in the In-Band Control register has expired. The I217 re-transmits the Status command if no acknowledge arrived after a time out defined in the PCIe Diagnostic register has expired. If after 32 retries no acknowledge arrived, the Intel® Ethernet Controller I217 aborts this command and continues with the next operation.

MtP Status command format:

Function	Command	Byte Count	Data 1	Data 2	Data 3	Data 4
Status	0xD3	4	Address and Control	PHY Status 1st Byte	PHY Status 2nd Byte	Reserved

Table 11. Address and control

Field	Bit(s)	Description
CFG_CTL	7	Configuration or control/status indication. This bit should be set to 1b indicating control/status.
CMD_ACK	6	Command/acknowledge indication. This bit should be set to 0b indicating command.
Reserved	5:0	Reserved, must be set to 0x0.

Table 12. PHY Status 1st Byte

Field	Bit(s)	Description
Reserved	7:6	Reserved (must be 0x0).
Host WoL	5	Host WoL indication
Reserved.	4	Reserved.
Reserved	3:2	Reserved (must be 0x0).
Reserved.	1	Reserved
Reserved	0	Reserved, must be set to 0b.

Table 13. PHY Status 2nd Byte

Field	Bit(s)	Description
RSTC	7	Reset complete.
Reserved (INT)	6	Interrupt request (not used).
Reserved (KLINK)	5	PCIe Link Status (not used).
CDIS	4	Ethernet cable disconnected.



Field	Bit(s)	Description
ELINK	3	Ethernet Link Status 1b = Link up. 0b = Link down.
DPX	2	Duplex Mode 1b = Full duplex. 0b = Half duplex.
SPD	1:0	Ethernet Link Speed 11b = Reserved. 10b = 1000 Mb/s. 01b = 100 Mb/s. 00b = 10 Mb/s.

2.3.4.1.3 PtM Configuration Acknowledge Command

The Configuration Acknowledge command is a single fragment command enabling the Intel® Ethernet Controller I217 to acknowledge the integrated LAN Controller access to its registers over the SMBus link.

Configuration command format:

Function	Command	Byte Count	Data 1	Data 2	Data 3	Data 4	
Configuration	0xD3	4	Address and Control	MSB Data	LSB Data	Reserved ¹	PHY Address

Notes:

1. Bits 7:5.

Table 14. Address and control

Field	Bit(s)	Description
CFG_CTL	7	Configuration/control indication. This bit should be set to 0b indicating configuration.
CMD_ACK	6	Command/acknowledge indication. This bit should be set to 1b indicating Acknowledge.
RW	5	Read/Write indication 1b = Indicates write acknowledge. 0b = Indicates read acknowledge.
Add	4:0	Register address.

2.3.4.1.4 PtM Wake Up Message

The host and Manageability Engine (ME) wake up message is a single fragment message indicating to the integrated LAN Controller that the PHY received a wake up packet/event that should cause both the host and ME to wake up.

Host and ME Wake Up message format:

Function	Command	Byte Count	Data 1
Wake Up	0xD5	1	Wake Up Control



Table 15. Wake up control

Field	Bit(s)	Description
Host Wake Up	0	Host wakeup indication.
ME Wake Up	1	ME wakeup indication.
Reserved	7:2	Reserved.

2.4 Transitions between SMBus and PCIe interfaces

2.4.1 Switching from SMBus to PCIe

Communication between the integrated LAN Controller and the Intel® Ethernet Controller I217 is done through the SMBus each time the system is in a low power state (Sx); PE_RST_N signal is low. The integrated LAN Controller/PHY interface is needed while the ME is still active or to enable host wake up from the Intel® Ethernet Controller I217.

Possible states for activity over the SMBus:

1. After power on (G3 to S5).
2. On system standby (Sx).

While in this state, the SMBus is used to transfer traffic, configuration, control and status between the ME through the integrated LAN Controller and the Intel® Ethernet Controller I217.

The switching from the SMBus to PCIe is done when the PE_RSTn signal is high.

- Any transmit/receive packet that is not completed when PE_RSTn is asserted is discarded.
- Any in-band message that was sent over the SMBus and was not acknowledged is re-transmitted over PCIe.

2.4.2 Switching from PCIe to SMBus

The communication between the integrated LAN Controller and the Intel® Ethernet Controller I217 is done through PCIe each time the system is in active power state (S0); PE_RSTn signal is high. Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME/PMC.

The switching from PCIe to SMBus is done when the PE_RSTn signal is low.

- Any transmit/receive packet that is not completed when PE_RSTn goes to 0b is discarded.
- Any in-band message that was sent over PCIe and was not acknowledged is re-transmitted over SMBus.

2.5 SMBus/PCIe Interconnects

The I217 can be connected to any x1 PCIe port in the chipset. The PCIe port that connects to the I217 LCD is selected by PCHSTRP9, bits [11:8] in the SPI Flash descriptor region. For more information on this setting, please refer to the chipset External Datasheet Specification. The chipset-to-I217 PCIe port connection in the

reference schematic must match the previously mentioned chipset SPI strap setting. Choosing another port can result in unexpected system behavior. The SMBus/PCIe interface can be configured in as shown Figure 2-2.




Figure 2-2. PCIe/SMBus Interface

Notes:

1. For latest PHY schematic connection recommendations, refer to the *Intel® I217 GbE PHY Reference Schematic*, and the *Intel® I217 GbE PHY Schematic and Layout Checklist*. For access to documentation, contact your Intel Field Service Representative.
2. Any free PCIe lanes [8:1] on the PCH can be used to connect the PCIe interface on the PHY.
3. Any free CLKOUT_PCIE[7:0] and PCIECLKRQ[7:0]# from PCH can be used to connect to PE_CLK[P:N] and CLK_REQ_N on the PHY respectively.
4. Any CLKOUT_SRC[7:0] and SRC[7:0]CLKRQ# can be used to connect to PE_CLK and CLK_REQ_N on the PHY.
5. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.



3 Pin Interface

3.1 Pin Assignment

The I217 is packaged in a 48-pin package, 6 x 6mm with a 0.4mm lead pitch. There are 48 pins on the periphery and a die pad (Exposed Pad*) for ground.

3.1.1 Signal Type Definitions

Signal Type	Definition
In	Input is a standard input-only signal.
I	A standard input-only signal.
Out (O)	Totem pole output is a standard active driver.
T/s	Tri-state is a bi-directional, tri-state input/output pin.
S/t/s	Sustained tri-state is an active low tri-state signal owned and driven by one and only one agent at a time. The agent that drives an s/t/s pin low must drive it high for at least one clock before letting it float. A new agent cannot start driving an s/t/s signal any sooner than one clock after the previous owner tri-states it.
O/d	Open drain enables multiple devices to share as a wire-OR.
Analog	Analog input/output signal.
A-in	Analog input signal.
A-out	Analog output signal.
B	Input bias

3.1.2 PCIe Interface Pins (8)

Pin Name	Pin #	Type	Op Mode	Name and Function
PE_RST_N	36	I	Input	PCIe reset.
PETp PETn	38 39	A-out	Output	PCIe Tx.
PERp PERn	41 42	A-in	Input	PCIe Rx.
PE_CLKP PE_CLKN	44 45	A-in	Input	PCIe clock.
CLK_REQ_N	48	O/d		Clock request.

3.1.3 SMBus Interface Pins (2)

Pin Name	Pin #	Type	Op Mode	Name and Function
SMB_CLK	28	O/d	BI-dir	SMBus clock. Pull this signal up to 3.3 V dc (auxiliary supply) through a 499 Ω resistor (while in Sx mode).
SMB_DATA	31	O/d	BI-dir	SMBus data. Pull this signal up to 3.3 V dc (auxiliary supply) through a 499 Ω resistor (while in Sx mode).



3.1.4 Miscellaneous Pins (3)

Pin Name	Pin #	Type	Op Mode	Name and Function
RSVD1_VCC3P3	1	T/s		
LANWAKE_N	2	O/d		Wake signal.
LAN_DISABLE_N	3	I		When this pin is set to 0b, the Intel® Ethernet Controller I217 is disabled.

3.1.5 PHY Pins (14)

3.1.5.1 LEDs (3)

Pin Name	Pin #	Type	Op Mode	Name and Function
LED0	26	O	Output	This signal is used for the programmable LED.
LED1	27	O	Output	This signal is used for the programmable LED.
LED2	25	O	Output	This signal is used for the programmable LED.

3.1.5.2 Analog Pins (11)

Pin Name	Pin#	Type	Op Mode	Name and Function
MDI_PLUS0 MDI_MINUS0	13 14	Analog	Bi-dir	Media Dependent Interface[0] 100BASE-T: In MDI configuration, MDI[0]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[0]+/- corresponds to BI_DB+/-. 100BASE-TX: In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair. 10BASE-T: In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair.
MDI_PLUS1 MDI_MINUS1	17 18	Analog	Bi-dir	Media Dependent Interface[1] 100BASE-T: In MDI configuration, MDI[1]+/- corresponds to BI_DB+/- and in MDI-X configuration MDI[1]+/- corresponds to BI_DA+/-. 100BASE-TX: In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair. 10BASE-T: In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair.
MDI_PLUS2 MDI_MINUS2 MDI_PLUS3 MDI_MINUS3	20 21 23 24	Analog	Bi-dir	Media Dependent Interface[3:2] 100BASE-T: In MDI configuration, MDI[3:2]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[3:2]+/- corresponds to BI_DB+/-. 100BASE-TX: Unused. 10BASE-T: Unused.
XTAL_OUT	9	O		Output crystal.
XTAL_IN	10	I		Input crystal.
RBIAS	12	Analog		Connect to ground through a 3.01 K Ω +/-1%.



3.1.6 Testability Pins (5)

Pin Name	Pin #	Type	Op Mode	Name and Function
JTAG_TCK	35	In	Input	JTAG clock input.
JTAG_TDI	32	In PU	Input	JTAG TDI input.
JTAG_TDO	34	T/s	Output	JTAG TDO output.
JTAG_TMS	33	In PU	Input	JTAG TMS input.
TEST_EN	30	In	Input	Should be connected to ground through a 1 kΩ resistor, when connected to logic 1b and test mode is enabled.


Note: The I217 uses the JTAG interface to support XOR files for manufacturing test. BSDL is not supported.

3.1.7 Power and Power Control Pins (16)

Pin Name	Pin #	Type	Name and Function
VDD0P9_8	8	Power	Core Vdc supply
VDD0P9_11	11		
VDD0P9_16	16		
VDD0P9_22	22		
VDD0P9_37	37		
VDD0P9_40	40		
VDD0P9_43	43		
VDD0P9_46	46		
VDD0P9_47	47		
VDD3P3_4	4	Power	3.3 Vdc supply
VDD3P3_5	5		
VDD3P3_15	15		
VDD3P3_19	19		
VDD3P3_29	29		
CTRL0P9	7	Analog	Internal SVR control pin. Connect to a 4.7 μH inductor and to the core rail.
SVR_EN_N	6	Input	SVR Enable pin: Connect to GND for internal SVR use. Connect to VDD3 when using external power.



3.2 Pinout





4 Package

4.1 Package Type and Mechanical

The I217 package is a 6mm x 6mm, 48-pin QFN Halogen Free and Pb Free package with e-PAD size of 3mm x 3mm.




Figure 4-1. Package Illustration



4.2 Package Electrical and Thermal Characteristics

The thermal resistance from junction to case, q_{JC} , is $15.1 \times C/Watt$. The thermal resistance from junction to ambient, q_{JA} , is as follows, 4-layer PCB, $85 \text{ }^\circ C_{ambient}$.

Air Flow (m/s)	Maximum T_j	$q_{JA} (\times C/Watt)$
0	119	34
1	118	33
2	116	31

No heat sink is required.


4.3 Power and Ground Requirements

All the grounds for the package is down-bonded to an Exposed Pad* e-PAD*. Core power supply provided by internal SVR controlled through pin 7 using external inductor and capacitor.

4.4 Ball Mapping

Pin Name	Side	Pin Number
RSVD1_VCC3P3	Left	1
LANWAKE_N	Left	2
LAN_DISABLE_N	Left	3
VDD3P3_4	Left	4
VDD3P3_5	Left	5
SVR_EN_N	Left	6
CTRL_0P9	Left	7
VDD0P9_8	Left	8
XTAL_OUT	Left	9
XTAL_IN	Left	10
VDD0P9_11	Left	11
RBIAS	Left	12
LED2	Right	25
LED0	Right	26
LED1	Right	27
SMB_CLK	Right	28
VDD3P3_29	Right	29
TEST_EN	Right	30
SMB_DATA	Right	31
JTAG_TDI	Right	32
JTAG_TMS	Right	33
JTAG_TDO	Right	34
JTAG_TCK	Right	35
PE_RST_N	Right	36
GND_EPAD	Epad (Center)	49

Pin Name	Side	Pin Number
MDI_PLUS0	Bottom	13
MDI_MINUS0	Bottom	14
VDD3P3_15	Bottom	15
VDD0P9_16	Bottom	16
MDI_PLUS1	Bottom	17
MDI_MINUS1	Bottom	18
VDD3P3_19	Bottom	19
MDI_PLUS2	Bottom	20
MDI_MINUS2	Bottom	21
VDD0P9_22	Bottom	22
MDI_PLUS3	Bottom	23
MDI_MINUS3	Bottom	24
VDD0P9_37	Top	37
PETp	Top	38
PETn	Top	39
VDD0P9_40	Top	40
PERp	Top	41
PERn	Top	42
VDD0P9_43	Top	43
PE_CLKP	Top	44
PE_CLKN	Top	45
VDD0P9_46	Top	46
VDD0P9_47	Top	47
CLK_REQ_N	Top	48



5 Initialization

5.1 Power Up

Initialization begins with power up.

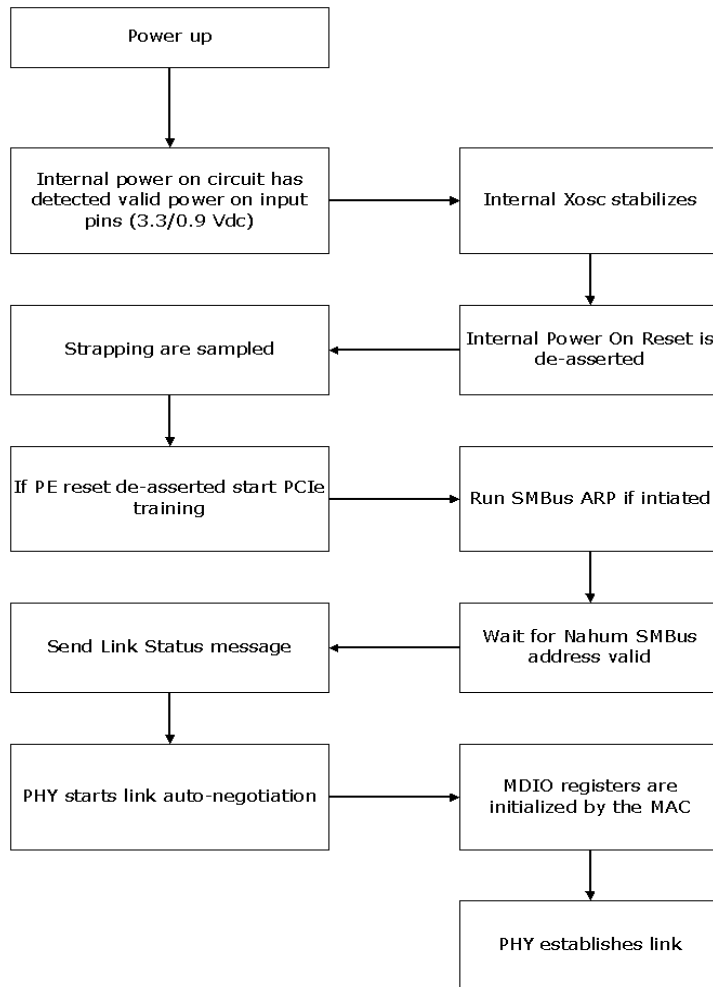


Figure 5-1. Power-Up Sequence



Note: Internal Power On Reset is an active low signal.

Note	
1	Platform power ramps up (3.3 V dc/Core V dc)
2	XTAL is stable after T_{XTAL} sec.
3	Internal Power On Reset triggers T_{POR} after XTAL is stable. Strapping options are latched.
4	PCIe training if PE reset is de-asserted.
5	SMBus ARP if initiated.
6	Wait for the integrated LAN Controller SMBus address valid.
7	Send Link Status message.
8	MAC configures the Intel® Ethernet Controller I217.
9	PHY goes through auto-negotiation to acquire link.

Power requirements during the power-up sequence are described in [Section 6.3.1.1](#) .

5.2 Reset Operation

The reset sources for the Intel® Ethernet Controller I217 are as follows:

- **Internal Power On Reset (POR)**—The I217 has an internal mechanism for sensing the power pins. Until power is up and stable, it generates an internal active low reset. This reset acts as a master reset for Intel® Ethernet Controller I217. While the internal reset is 0b, all registers in Intel® Ethernet Controller I217 are reset to their default values. Strapping values are latched after Internal POR is de-asserted.
- **PHY Soft Reset**—A PHY reset caused by writing to bit 15 in MDIO register 0. Setting the bit resets the PHY, but does not reset non-PHY parts. The PHY registers are reset, but other Intel® Ethernet Controller I217 registers are not.

Note: The integrated LAN controller configures the LCD registers. Other I217 GbE LCD registers do not need to be configured.

- **PCIe Reset** from pin—After asserting a PCIe reset, Intel® Ethernet Controller I217 stops the SerDes and if in the middle of transmitting a packet it should be dropped. De-asserting PCIe reset resets the internal FIFO unless wake-up is activated and causes a switch from SMBus to PCIe.
- **In-Band Reset**—An in-band message causing complete reset of Intel® Ethernet Controller I217 except the wake up filters content.

Note: LAN_DISABLE_N is the only external signal that can reset the PHY. Refer to [Section 6.3.1.3](#) for more details.

Some of the bonding option registers are writable and can be loaded from the NVM or written by the integrated LAN Controller (SKU register). The effect of the various reset options on these and other registers is listed in the following table.



Effects/ Sources	PCIe- Based Interface	Non-PHY Registers and State	PHY Registers and State	Reset Complete Indication ¹	Strapping Options	Fuse Registers	Move Out of Power Down Mode	Wake Up Register
Internal POR ²	X	X	X	X	X	X		X
PHY Soft Reset ³			X	X				
PCIe Reset	X							
In-Band Reset	X	X	X	X		X	X	

Notes:

1. Once the PHY completes its internal reset a reset complete indication is sent to the integrated LAN controller over the interconnect. The integrated LAN controller then configures the PHY.
2. Asserting a 3.3 Vdc power on reset should move the PHY out of power down mode.
3. PHY registers (page 0 in MDIO space and any aliases to page 0) are reset during a PHY soft reset. The rest of Intel® Ethernet Controller I217’s MDIO space is not reset.

5.3 Timing Parameters

5.3.1 Timing Requirements

The I217 requires the following start-up and power-state transitions.

Parameter	Description	Min.	Max.	Notes
T _{r2init}	Completing a PHY configuration following a reset complete indication.		0.5 s	

5.3.2 Timing Guarantees

The I217 guarantees the following start-up and power state transition related timing parameters.

Note: For platform power sequencing requirements for the integrated LAN controller, refer to the chipset EDS.

Parameter	Description	Min.	Max.	Notes
T _{PHY_Reset}	Reset de-assertion to PHY reset complete		10 ms	PHY configuration should be delayed until PHY completes its reset.
T _{c2an}	Cable connect at start of auto-negotiation	1.2 s	1.3 s	Per 802.3 specification.
T _{XTAL}	XTAL frequency stable after platform power ramp up		45 ms	
T _{POR}	Internal POR trigger after XTAL stable		40 ms	



NOTE: *This page intentionally left blank.*



6 Power Management and Delivery

This section describes power management in Intel® Ethernet Controller I217.

6.1 Power Information

The following table lists the targets for device power for the I217. Note that power is reduced according to link speed and link activity.

Power Consumption Target

System State		Link State	3.3 V Rail (mA) ¹	Core Rail (mA) ² -external-	Device Power (mW)		
S0 (Max)		Thermal Design Power (TDP)	132	110	535		
S0 (Typ)		1000 Mbps Active	132	110	535		
		1000 Mbps Idle	120	58	448		
		1000 Mbps LPI (EEE link partner only)	11	10	45		
		100 Mbps Active	62	61	260		
		100 Mbps Idle	49	21	180		
		100 Mbps LPI (EEE link partner only)	10	11	43		
		10 Mbps Active	77	56	304		
		10 Mbps Idle	18	15	73		
				No Link (No circuitry)	3	8	17
				LAN disabled using driver	2.5	2	10
Sx	WoL Enabled	100 Mbps WOL enabled	48	19	176		
		10 Mbps WOL enabled	17	13	68		
		Cable Disconnect	3	8	17		
	WoL Disabled	WOL disabled using driver	2.5	2	10		
		WOL disabled using BIOS ³	0	0	0		

1. Power was supplied by an external power supply.

2. Device Power was calculated using this formula: $3.3 \times I_{3.3} + (0.9 \times I_{0.9})$

3. Use SLP_LAN# to gate PHY power and the ME must be off in Sx state.

Note: Measured power could be higher or lower based on lab setup.



The following sections describe requirements in specific power states.

6.2 Power Delivery

The I217 operates from a 3.3 Vdc external power rail.

6.2.1 Core Vdc Supply

The Core Vdc rail is supplied by the integrated SVR, which requires an external inductor and capacitor.

6.3 Power Management

6.3.1 Global Power States

The I217 transitions between power states based on a status packet received over the interconnect and based on the Ethernet link state. The following power states are defined:

- **Power Up**—Defined as the period from the time power is applied to Intel® Ethernet Controller I217 and until the I217 powers up its PHY. The I217 needs to consume less than 40 mA during this period.
- **Active 10/100/1000 Mb/s**—Ethernet link is established with a link partner at any of 10/100/1000 Mb/s speed. The I217 is either transmitting/receiving data or is capable of doing so without delay (for example, no clock gating that requires lengthy wake).
- **Idle 10/100/1000 Mb/s**—Ethernet link is established with a link partner at any of 10/100/1000 Mb/s speed. The I217 is not actively transmitting or receiving data and might enter a lower power state (for example, an interface can be in electrical idle).
- **IEEE Power Down**—The standard IEEE power-down initiated by the Host setting the POWER_DOWN bit (bit 11) of the PHY Control Register to 1b (see [section 9.4](#)).
- **LPI**—IEEE802.3az [Energy Efficient Ethernet (EEE)] defines an optional Low Power Idle (LPI) mode for 1000BASE-T, 100BASE-TX and other interfaces. LPI enables power saving by switching off part of the I217 functionality when no data needs to be transmitted or/and received. When LPI support is enabled the I217 will shut off RX circuitry and send an inband RX LPI Indication on detection that link the partner's TX moved into LPI state. The I217 will move TX into LPI state and power-down transmit circuitry when receiving an Inband TX LPI request from the integrated LAN controller.

6.3.1.1 Power Up

Defined as the period from the time power is applied to the I217 and until the I217 powers up its PHY. The I217 should consume less than ~40 mA during this period. Following the I217 LCD entering reset, the power-up sequence is considered done and the requirement is removed. See [Section 5.1](#) for a description of the power-up sequence.



6.3.1.2 Cable Disconnect State

The I217 enters a cable disconnect state if it detects a cable disconnect condition on the Ethernet link. Power is reduced during cable disconnect mode by several means:

- The PHY enters energy detect mode.
- The PCIe link enters power down.

An exit from cable disconnect happens when the I217 detects energy on the MDI link, and starts the following exit sequence:

- The I217 signals the integrated LAN controller that link energy was detected by clearing the *Cable Disconnect* bit in the PCIe or SMBus interface.
- The PHY waits until the auto-negotiation break link timer expires (T_{c2an} time) and then starts to advertise data on the line.

Note: When a network cable is removed from an active system, the link LED remains lit approximately two to five seconds. This is expected behavior for a link-down.

6.3.1.3 Power Down State

The I217 enters a power-down state when the LAN_POWER_GOODn pin is set to zero. Exiting this mode requires setting the LAN_DISABLE_N pin to a logic one.

Note: Following a power up or reset, the power-down bit must not be set until the configuration cycle completes.

The *Device Power Down Mode* field in the MDIO register space defines the response to a power-down command. The I217 takes one of two possible actions:

- Device stays active—No change in functionality and no power reduction.
- Device power down—The PHY enters power down, clocks are gated, PCIe enters Electrical Idle (EI).

Figure 6 shows the power-down sequence in the two later cases.




Figure 6-1. Power-Down Sequence

Note: In cases where the LAN_DISABLE_N pin cannot be used, use a power-down inband. When used, the power savings are lower since not all logic is turned off in this mode.

Note	Description
2	Once the I217 detects the LAN_DISABLE_N transitions to a logic zero, the PHY enters a power-down state.
3	The PCIe link (if enabled) enters electrical idle state.
4	PCIe/SMBus exits a reset state and performs link initialization.

5	The integrated LAN controller configures the I217 through the MDIO interface.
6	PHY goes through auto-negotiation to acquire link.

6.3.1.4 EEE LPI State

EEE (Energy Efficient Ethernet) Low Power Idle (LPI) mode defined in IEEE802.3az optionally allows power saving by switching off part of the integrated LAN controller and the I217 functionality when no data needs to be transmitted or/and received. Decision on whether the I217 transmit path should enter Low Power Idle mode or exit Low Power Idle mode is done on the integrated LAN controller level and communicated to the I217 in order to allow power saving in the transmit circuitry. Information on whether Link Partner has entered Low Power Idle mode is detected by the I217 and communicated to the integrated LAN controller to allow for power saving in the receive circuitry.

The following table and figure illustrate general principles of an EEE LPI operation on the Ethernet Link.

Table 1. LPI Parameters

Parameter	Description
Sleep Time (Ts)	Duration PHY sends Sleep symbols before going Quiet.
Quiet Duration (Tq)	Duration PHY remains Quiet before it must wake for Refresh period.
Refresh Duration (Tr)	Duration PHY sends Refresh symbols for timing recovery and coefficient synchronization.
PHY Wake Time (Tw_PHY)	Minimum duration PHY takes to resume to Active state after decision to Wake.
Receive System Wake Time (Tw_System_rx)	Wait period where no data is expected to be received to give the local receiving system time to wake up.
Transmit System Wake Time (Tw_System_tx)	Wait period where no data is transmitted to give the remote receiving system time to wake up.




Figure 6-2. EEE LPI Compliant Operation

In the transmit direction entrance to Low Power Idle mode of operation is triggered by the reception of LPI TX Request from the integrated LAN controller. Following reception of the LPI TX in band Request, PHY transmits special Sleep symbols to communicate to the link partner that the local system is entering Low Power Idle mode. In 100BASE-TX LPI mode PHY enters low power operation in an asymmetric manner. After Sleep symbols transmission, the transmit function of the local PHY immediately enters a low power quiet mode. In 1000BASE-T LPI mode, PHY entry into low power mode is symmetric. Only after the local PHY transmits and receives sleep symbols from the remote PHY does the transmit function of the local PHY enter the quiet mode. Periodically the local PHY transmits Refresh symbols that are used by the link partner to



update adaptive filters and timing circuits in order to maintain link integrity. This quiet-refresh cycle continues until the local integrated LAN controller sends an in band message with a clear (0) LPI TX Request, which signals to the I217 that Low Power Idle mode should end. The I217 communicates this to the link partner by sending special Wake symbols for a pre-defined period of time. Then the PHY enters Active state and resumes normal operation. Data can be transmitted after a `Tw_System_tx` duration.

6.3.1.4.1 EEE Capabilities Auto-Negotiation

EEE support is advertised during Auto-Negotiation stage. Auto-Negotiation provides the capability to detect the abilities supported by the device at the other end of the link, determine common abilities, and configure for joint operation. Auto-Negotiation is performed at power up, on command from integrated LAN controller, upon detection of a PHY error, or following Ethernet cable re-connection.

During the link establishment process, both link partners indicate their EEE capabilities. If EEE is supported by both link partners for the negotiated PHY type then the EEE function may be used independently in either direction. The Auto-Negotiation process uses next page messages as defined in IEEE802.3az clauses 28C.12 and 28C.13.

6.3.1.4.2 EEE LPI Unsupported Features

EEE LPI does not support:

- Half-duplex operation
- 10 Mb/s speed
- When network proxy offload is enabled.

Note: These features should not be enabled while EEE is enabled.

6.3.1.5 Intel® Auto Connect Battery Saver (ACBS)

Intel® Auto Connect Battery Saver for the Intel® Ethernet Connection I217 is a hardware-only feature that automatically reduces the LCD to a lower power state when the power cable is disconnected. When the power cable is reconnected, it renegotiates the line speed following IEEE specifications for auto negotiation. By default, auto negotiation starts at 1 Gb/s, then 100 Mb/s full duplex/half duplex, then 10 Mb/s full duplex/half duplex.

- ACBS is only supported during auto negotiation. If link is forced, the I217 does not enter ACBS mode.

The I217 ACBS works in both S0 and Sx states. Since the I217 ACBS has no driver control, the feature is always enabled, allowing power savings by default.

- The crystal clock drivers are intermittently disabled when the network cable is unplugged and the I217 is in ACBS mode.

6.3.1.6 Automatic Link Downshift

Automatic link downshift is a collection of power saving features that enable a link downshift from 1000 Mb/s to a lower speed to save power under different conditions like the AC cable plugged in, monitor idle, or entering Sx states.



6.3.1.7 Link Speed Battery Saver

Link speed battery saver is a power saving feature that negotiates to the lowest speed possible when the I217 operates in battery mode to save power. When in AC mode, where performance is more important than power, it negotiates to the highest speed possible. The Windows NDIS drivers (Windows XP and later) monitor the AC-to-battery transition on the system to make the LCD negotiate to the lowest connection speed supported by the link partner (usually 10 Mb/s) when the power cable is unplugged (switches from AC to battery power). When the AC cable is plugged in, the speed negotiates back to the fastest LAN speed. This feature can be enabled/disabled directly from DMiX or through the advanced settings of the Window's driver.

When transferring packets at 1000/100 Mb/s speed, if there is an AC-to-battery mode transition, the speed renegotiates to the lower speed. Any packet that was in process is re-transmitted by the protocol layer. If the link partner is hard-set to only advertise a certain speed, then the driver negotiates to the advertised speed. Note that since the feature is driver based, it is available in S0 state only.

Link speed battery saver handles duplex mismatches/errors on link seamlessly by re-initiating auto negotiation while changing speed. Link speed battery saver also supports spanning tree protocol.

Note: Packets are re-transmitted for any protocol other than TCP as well.

6.3.1.8 System Idle Power Saver (SIPS)

SIPS is a software-based power saving feature that is enabled only with Microsoft* Windows* Vista* or Windows* 7. This feature is only supported in the S0 state and can be enabled/disabled using the advanced tab of the Windows driver or through DMiX. The power savings from this feature is dependent on the link speed.



6.3.1.9 Low Power Link Up (LPLU)

LPLU is a firmware/hardware-based feature that enables the designer to make the PHY negotiate to the lowest connection speed first and then to the next higher speed and so on. This power saving setting can be used when power is more important than performance.

When speed negotiation starts, the PHY tries to negotiate for a 10 Mb/s link, independent of speed advertisement. If link establishment fails, the PHY tries to negotiate with different speeds. It enables all speeds up to the lowest speed supported by the partner. For example, if the I217 advertises 10 Mb/s only and the link partner supports 1000/100 Mb/s only, a 100 Mb/s link is established.

LPLU is controlled through the LPLU bit in the PHY Power Management register. The integrated LAN controller sets and clears the bit according to hardware/software settings. The I217 auto negotiates with the updated LPLU setting on the following auto-negotiation operation. The I217 does not automatically auto-negotiate after a change in the LPLU value. LPLU is not dependent on whether the system is in Vac or Vdc mode. In S0 state, link speed battery saver overrides the LPLU functionality.


LPLU is enabled for non-D0a states by GbE NVM image word 0x17 (bit 10):

- 0b = LPLU is disabled.
- 1b = LPLU is enabled in all non-D0a states.

LPLU power consumption depends on what speed it negotiates at. [Section 6.1](#) includes all of the power numbers for the I217 in the various speeds.



NOTE: *This page intentionally left blank.*



7 Device Functionality

7.1 Tx Flow

When packets are ready for transmission in the integrated LAN controller it transfers them to the I217 through the PCIe or the SMBus (depends on system state). I217 starts transmitting the arrived packet over the wire after it gathers 8 bytes of data if the PCIe interface is active or after all packet data is received if it was transferred over the SMBus; this behavior has no dependency on the link speed. I217 design is based on the assumption that the integrated LAN controller has the full packet ready for transmission.

In several cases the I217 has to stop transmission over the wire while still accepting data from the upper connection (PCIe or SMBus). For those cases, the I217 maintains a 3 KB FIFO. The cases where the I217 needs to stop Tx are:

- PAUSE packet was received on the Rx side while flow control is enabled. For full support of flow control, the *Receive Flow Control Enable (RFCE)* bit (bit 7) in the LCD Receive Control register should be set in addition to the configuration in the integrated LAN controller.
- In half-duplex mode while the I217 is in the middle of a receive (DEFER).
- In half-duplex mode while a collision was detected on the wire.

In addition to stop transmission, the I217 sends an in-band message to the integrated LAN controller with the Tx *OFF* bit set. This in-band message must be sent at the first gap between received packets if (at the same time) the event caused the stop transmit is not valid and transmission over the wire is activated, the I217 might avoid sending the in-band message. An in-band message with the Tx *OFF* bit cleared is sent when the collided packet was successfully transmitted or dropped after 16 retries (see [Section 7.3.1](#)).

In-band messages from the integrated LAN controller to the I217 always come in between packets during the IPG. I217 does not accept in-band messages in the middle of a packet.

7.2 Rx Flow

I217 maintains a 3 KB FIFO on the receive side in order not to lose packets when PCIe is active but at K1 power save mode. In this case the Intel® Ethernet Controller I217 initiates recovery of the PCIe when a reception has started. If the link is at 1 Gb/s, the transmission of the packet over the PCIe bus starts immediately after recovery. If the link speed is lower, the I217 starts the transmission after the entire packet is received. I217 assumes maximum recovery time (from the K1 state) of 10 μ s on both sides of the PCIe side. Higher recovery time causes a packet drop on the receive side.

I217 identifies PAUSE packets, stop transmission, and a send in-band message as described in the previous section.

In-band messages from the I217 to the integrated LAN controller always come in between packets during the IPG.



When the PCIe is not active, packet drop is not avoidable due to the big difference in line rate between the MDI and the SMBus.

7.3 Flow Control

Flow control as defined in 802.3x, as well as the specific operation of asymmetrical flow control defined by 802.3z, is supported in the integrated LAN controller during PCIe mode and in the LAN Connected Device during SMBus mode. Some of the flow control functionality has moved to the I217. The following registers are duplicated to the I217 for the implementation of flow control:

- Flow Control Address is: 0x01, 0x80, 0xC2, 0x00, 0x00, 0x01; where 01 is the first byte on the wire, 0x80 is the second, etc.
- Flow Control Type (FCT): a 16-bit field to indicate flow control type
- Flow Control Transmit Timer Value (FCTTV): a 16-bit timer value to include in transmitted PAUSE frame
- Flow Control Refresh Threshold Value (FCRTV): a 16 bit PAUSE refresh threshold value

Flow control is implemented as a mean of reducing the possibility of receive buffer overflows, which result in the dropping of received packets, and allows for local controlling of network congestion levels. This can be accomplished by sending an indication to a transmitting station of a nearly full receive buffer condition at a receiving station. The implementation of asymmetric flow control allows for one link partner to send flow control packets while being allowed to ignore their reception. For example, not required to respond to PAUSE frames.

7.3.1 MAC Control Frames and Reception of Flow Control Packets

Three comparisons are used to determine the validity of a flow control frame:

1. A match on the six-byte multicast address for MAC control frames or to the station address of the device (Receive Address Register 0).
2. A match on the type field
3. A comparison of the *MAC Control Opcode* field.

The 802.3x standard defines the MAC control frame multicast address as 01-80-C2-00-00-01. The *Flow Control Packet's Type* field is checked to determine if it is a valid flow control packet: XON or XOFF. 802.3x reserves this as 0x8808. The final check for a valid PAUSE frame is the *MAC Control Opcode* field. At this time only the PAUSE control frame opcode is defined. It has a value of 0x0001. Frame based flow control differentiates XOFF from XON based on the value of the PAUSE *Timer* field. Non-zero values constitute XOFF frames while a value of zero constitutes an XON frame. Values in the *Timer* field are in units of slot time. A slot time is hard-wired to 64 byte times.

Note: An XON frame signals cancelling the pause from being initiated by an XOFF frame (Pause for zero slot times).

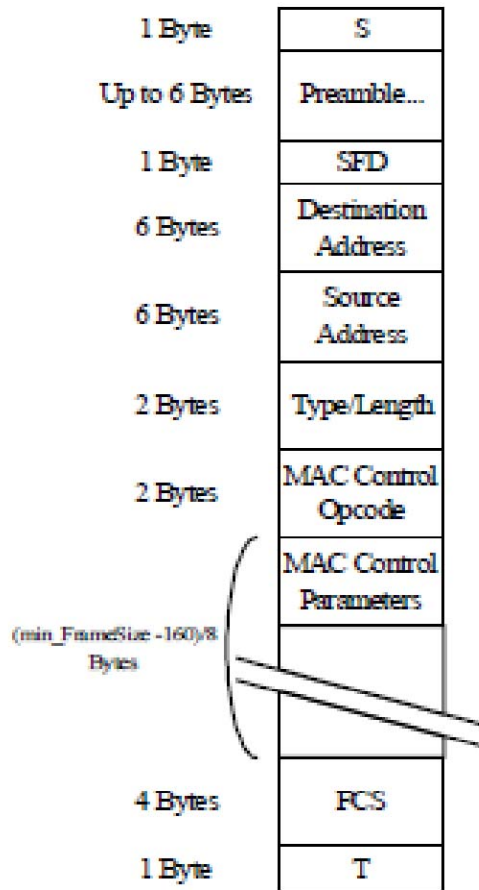


Figure 7-1. 802.3x MAC Control Frame Format

Where S is the start-of-packet delimiter and T is the first part of the end-of-packet delimiter for 802.3z encapsulation. The receiver is enabled to receive flow control frames if flow control is enabled via the *RFCE* bit in the Device Control (CTRL) register.

Note: Flow control capability must be negotiated between link partners via the auto-negotiation process. The auto-negotiation process might modify the value of these bits based on the resolved capability between the local device and the link partner.

Once the I217 has validated the reception of an XOFF, or PAUSE frame, it performs the following:

- Initializes the pause timer based on the packet's PAUSE *Timer* field
- Disables packet transmission or schedules the disabling of transmission after the current packet completes.
- Sends an in-band status command with the TX OFF bit set.
- Forward the XOFF or PAUSE frame to integrated LAN controller.



Resuming transmission might occur under the following conditions:

- Expiration of the PAUSE timer.
- Reception of an XON frame (a frame with its PAUSE timer set to zero).¹

Once the I217 has validated the reception of an XON frame, it performs the following:

- Enables packet transmission.
- Sends an in-band status command with the Tx *OFF* bit cleared.
- Forwards the XON frame to the integrated LAN controller.

7.3.2 Transmitting PAUSE Frames

During PCIe mode transmitting PAUSE frames is done as a result of an In-Band Control command from the integrated LAN controller. The integrated LAN controller initiates an in-band message if it is enabled by software by writing a 1b to the *TFCE* bit in the Device Control register.

Note: Similar to receiving flow control packets previously mentioned, XOFF packets can be transmitted only if this configuration has been negotiated between the link partners via the auto-negotiation process. In other words, the setting of this bit indicates the desired configuration.

When the in-band message from the integrated LAN controller is received, the I217 sends a PAUSE frame with its *PAUSE Timer* field equal to *FCTTV*. Once the receive buffer fullness reaches the low water mark, the integrated LAN controller sends an in-band message indicating to send an XON message (a PAUSE frame with a timer value of zero).

During SMBus mode transmitting PAUSE frames is done as a result of the I217 receive and transmit FIFO status. If the sum of data in both FIFOs is greater than the configured *FCTH.HTHRSH* the I217 sends a PAUSE frame with its *PAUSE Timer* field equal to *FCTTV*. Once the sum of data in the receive and transmit FIFOs is lower than *FCTH.LTHRSH* the I217 sends a PAUSE frame with a timer value of zero (XON). The I217 will send an XOFF refresh message with the PAUSE Timer set to *FCTTV* if after *FCRTV* time from the previous XOFF message the transmit and receive buffer fullness is still above the low threshold value.

Note: Transmitting flow control frames should only be enabled in full-duplex mode per the IEEE 802.3 standard. Software should make sure that the transmission of flow control packets is disabled when the I217 is operating in half-duplex mode.

7.4 Wake Up

The I217 supports host wake up.

The host can enable host wake up from the I217 by setting the *Host_WU_Active* bit. When this bit is set

Setting the I217's wake up:

1. The XON frame is also forwarded to integrated LAN controller.



1. Verify Host_WU_Active bit (bit 4) in the Port General Configuration register (page 769, register 17) is clear, this is needed to allow the configuration of the filters to wake up mode.
2. Set bit 2 (*MACPD_enable*) of the Port Control register (page 769, register 17) to enable the I217 wake up capability and software accesses to page 800.
3. Set the *Slave Access Enable* bit (bit 2) in the Receive Control register (page 800, register 0) to enable access to the Flex Filter register, if setting those bits is needed in the next stage. The registers affected are:
 - a. Flexible Filter Value Table LSB—FFVT_L (filters 01)
 - b. Flexible Filter Value Table MSBs—FFVT_H (filters 23)
 - c. Flexible Filter Value Table—FFVT_45 (filters 45)
 - d. Flexible Filter Value Table—FFVT_67 (filters 67)
 - e. Flexible TCO Filter Value/Mask Table LSBs—FTFT_L
 - f. Flexible TCO Filter Value/Mask Table MSBs—FTFT_H
 - g. Extended Flexible Filter Value Table—EFFVT (Filters 8-15)
4. Configure the I217's wake up registers per ACPI/APM wake up needs.
5. Clear the *Slave Access Enable* bit (bit 2) in the Receive Control register (page 800, register 0) to enable the flex filters.
6. Set the *Host_WU_Active* bit (bit 4) in the Port General Configuration register (page 769, register 17) to activate the I217's wake up functionality.

On a detection of a Host wake up event the I217 will:

1. Assert the LANWAKE# pin indicating wake to PMC.
2. If in DeepSx, PMC will power up the SUS well.
3. PMC will send a PINSTOP message to the I217 through SMBus
4. On a reception of a PINSTOP message, the I217 will stop asserting the LANWAKE# pin and send a WU message over SMBus indicating the WU source.
 - a. I217 will send a WU message on every PINSTOP message reception
5. I217 will not assert the LANWAKE# pin again until a WU message was sent and acknowledged power-up PMC. In case of host wake up the integrated LAN controller wakes the host.
6. Host should issue a LCD reset to the I217 before clearing the *Host_WU_Active* bit.
7. Host reads the Wake Up Status register (WUS); wake up status from the I217.

When a wake up packet is identified, the wake up in-band message is sent and the host should clear the *Host_WU_Active* bit (bit 4) in the Port General Configuration register (page 769, register 17). up to the LAN Controller, and clears the PME_Status bit in the WUC register

While in wake up active mode new wake up packets received will not overwrite the packet in the FIFO. I217 re-transmits the wake up in-band message after 50 ms if no change in the *Host_WU_Active* bits occurred.

7.4.1 Host Wake Up

I217 supports two types of wake up mechanisms:

- Advanced Power Management (APM) wake up
- ACPI Power Management wake up



7.4.1.1 Advanced Power Management Wake Up

Advanced Power Management Wakeup or APM Wakeup was previously known as Wake on LAN (WoL). The basic premise is to receive a broadcast or unicast packet with an explicit data pattern, and then to assert a signal to wake up the system or issue an in-band PM_PME message (if configured to).

At power up, if the I217's wake up functionality is enabled, the *APM Enable* bits from the NVM are written to the I217 by the integrated LAN controller to the *APM Enable* (APME) bits of the Wakeup Control (WUC) register. These bits control the enabling of APM wake up.

When APM wake up is enabled, the I217 checks all incoming packets for Magic Packets. See [Section 7.4.1.3.1.4](#) for a definition of Magic Packets.

To enable APM wake up, programmers should write a 1b to bit 10 in register 26 on page 0 PHY address 01, and then the station address to registers 27, 28, 29 at page 0 PHY address 01. The order is mandatory since registers RALO[31:0] and RAHO[15:0] are updated with a corresponding value from registers 27, 28, 29, if the *APM WoL Enable* bit is set in register 26. The Address Valid bit (bit 31 in RAHO) is automatically set with a write to register 29, if the *APM WoL Enable* bit is set in register 26. The *APM Enable* bit (bit 0 in the WUC) is automatically set with a write to register 29, if the *APM WoL Enable* bit is set in register 26.

Once the I217 receives a matching magic packet, it:

- Sets the *Magic Packet Received* bit in the WUS register.
- Initiates the integrated LAN controller wake up event through an in-band message.

APM wake up is supported in all power states and only disabled if a subsequent NVM read results in the *APM Wake Up* bit being cleared or software explicitly writes a 0b to the *APM Wake Up* (APM) bit of the WUC register.

7.4.1.1.1 Link Status Change

When the *LSCWO* bit (bit 5 in the WUC register) is set, wake up is generated if all of the following conditions are met:

- APM wake up is enabled (*APME* bit is set in the WUC register)
- The *LSCWE* bit (bit 4) is set in the WUC register
- Link status change is detected

When the I217 detects a link status change it:

- Sets the *Link Status Changed* (LNKC) bit (bit 0) in the WUS register.
- Initiates the integrated LAN controller wake up event.

When the *LSCWO* bit is set, wake up is never generated on link status change if either APM wake up is disabled or the *LSCWE* bit is cleared. In this case, the *LNKC* bit (bit 0) in the Wake up Filter Control (WUFC) register is read as zero, independent of the value written to it.



7.4.1.2 ACPI Power Management Wake Up

I217 supports ACPI Power Management based wake ups and can generate system wake up events from three sources:

- Reception of a Magic Packet
- Reception of a ACPI wake up packet
- Detection of a link change of state

Activating ACPI Power Management wake up requires the following steps:

- Programming of the WUFC register to indicate the packets it needs to wake up and supply the necessary data to the IPv4 Address Table (IP4AT) and the Flexible Filter Mask Table (FFMT), Flexible Filter Length Table (FFLT), and the Flexible Filter Value Table (FFVT). It can also set the *Link Status Change Wake up Enable* (LNKC) bit (bit 0) in the WUFC register to cause wake up when the link changes state.
- Setting bit 2 (*MACPD_enable*) of the Port Control register (page 769, register 17) to put the I217 in wake up mode.

Once wake up is enabled, the I217 monitors incoming packets by first filtering them according to its standard address filtering method and then by filtering them with all enabled wake up filters. If a packet passes both the standard address filtering and at least one of the enabled wake up filters, the I217:

- Initiates an integrated LAN controller wake up event.
- Sets one or more of the *Received* bits in the WUS register. Note that more than one bit is set if a packet matches more than one filter.

If enabled, a link state change wake up causes similar results.

7.4.1.3 Wake Up Packet Storage

A packet initiating Host wake up can be stored in the I217 by setting 772.25[4] prior to WoL entry.

Post wakeup the host clears the Host WoL setting and can read the content of the WoL packet from 803.0-63, before that the driver should verify if a packet was stored by checking at 803.64 to get the packet length—a value of zero means no packet was captured. By writing to 803.66 the driver will clear the wake up packet indication in preparation for a subsequent low power state entry.

7.4.1.4 Wake Up Packets

I217 supports various wake up packets using two types of filters:

- Pre-defined filters
- Flexible filters

Each of these filters are enabled if the corresponding bit in the WUFC register is set to 1b. If the wake up packet passes one of the manageability filters enabled in the Management Control (MANC) register, system wake up also depends on the *NoTCO* bit (11) in the WUFC register being inactive or the MANC.NoHostWoLonMETraffic bit that enable Host WoL only if a packet passed ME filter but was also directed to the Host using the MANC2H and MANC2H2 registers.



7.4.1.4.1 Pre-Defined Filters

The following packets are supported by the I217's pre-defined filters:

- Directed Packet (including exact, multicast indexed, and broadcast)
- Magic Packet
- ARP/IPv4 Request Packet
- Directed IPv4 Packet
- Directed IPv6 Packet
- Flexible UDP/TCP and IP filters packets

Each of the filters are enabled if the corresponding bit in the WUFC register is set to 1b.

The explanation of each filter includes a table showing which bytes at which offsets are compared to determine if the packet passes the filter. Note that both VLAN frames and LLC/Snap can increase the given offsets if they are present.

7.4.1.4.1.2 Directed Exact Packet

I217 generates a wake up event after receiving any packet whose destination address matches one of the valid programmed receive addresses if the *Directed Exact Wake Up Enable* bit (bit 2) is set in the WUFC register.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	Match any pre-programmed address as defined in the receive address

7.4.1.4.1.3 Directed Multicast Packet

For multicast packets, the upper bits of the incoming packet's destination address indexes a bit vector and the Multicast Table Array indicates whether to accept the packet. If the *Directed Multicast Wake Up Enable* bit (bit 3) is set in the WUFC register and the indexed bit in the vector is one, the I217 generates a wake up event. The exact bits used in the comparison are programmed by software in the *Multicast Offset* field (bits 4:3) of the RCTL register.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	See above paragraph.

7.4.1.4.1.4 Broadcast

If the *Broadcast Wake Up Enable* bit (bit 4) in the WUFC register is set, the I217 generates a wake up event when it receives a broadcast packet.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address	FF*6	Compare	

7.4.1.4.1.5 Magic Packet

Magic packets are defined as follows:

- **Magic Packet Technology Details**—Once the I217 has been put into Magic Packet mode, it scans all incoming frames addressed to the node for a specific data sequence, which indicates to the MAC that this is a Magic Packet frame. A Magic Packet frame must also meet the basic requirements for the LAN technology



chosen, such as Source address, Destination Address (which might be the receiving station's IEEE address or a Multicast address that includes the Broadcast address) and CRC. The specific data sequence consists of 16 duplications of the IEEE address of this node with no breaks or interruptions. This sequence can be located anywhere within the packet, but must be preceded by a synchronization stream. The synchronization stream enables the scanning state machine to be much simpler. The synchronization stream is defined as 6 bytes of 0xFF. The device also accepts a Broadcast frame, as long as the 16 duplications of the IEEE address match the address of the system that needs to wake up.

I217 expects the destination address to either:

1. Be the broadcast address (FF.FF.FF.FF.FF.FF)
2. Match the value in Receive Address (RAH0/RAL0) register 0. This is initially loaded from the NVM but can be changed by the software device driver.
3. Match any other address filtering enabled by the software device driver.

If the packet destination address met one of the three criteria previously listed, the I217 searches for 16 repetitions of the same destination address in the packet's data field. Those 16 repetitions must be preceded by (in the data field) at least 6 bytes of 0xFF, which act as a synchronization stream. If the destination address is NOT the broadcast address (FF.FF.FF.FF.FF.FF), the I217 assumes that the first non-0xFF byte following at least 6 0xFF bytes is the first byte of the possible matching destination address. If the 96 bytes following the last 0xFF are 16 repetitions of the destination address, the I217 accepts the packet as a valid wake up Magic Packet. Note that this definition precludes the first byte of the destination address from being 0xFF.

A Magic Packet's destination address must match the address filtering enabled in the configuration registers with the exception that broadcast packets are considered to match even if the *Broadcast Accept* bit (bit 5) of the RCTL register is 0b. If APM wake up is enabled in the NVM, the I217 starts up with the RAH0/RAL0 register 0 loaded from the NVM. This enables the I217 to accept packets with the matching IEEE address before the software device driver comes up.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Len/LLC/SNAP Header		Skip	
12 + S + D	2	Type		Skip	
Any	6	Synchronizing Stream	FF*6+	Compare	
any+6	96	16 copies of Node Address	A*16	Compare	Compared to RAH0/RAL0 register

7.4.1.4.1.6 ARP/IPv4 Request Packet

I217 supports receiving ARP Request packets for wake up if the ARP bit (bit 5) is set in the WUFC register. Three IPv4 addresses are supported, which are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must contain a broadcast MAC address, a protocol type of 0x0806, an ARP opcode of 0x01, and one of the three host programmed IPv4 addresses. I217 also handles ARP Request packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.



To enable broadcast ARP wakeup RCTL.BAM should be set to accept broadcast packets.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Possible Len/LLC/SNAP Header		Skip	
12 + S + D	2	Type	0x0806	Compare	ARP
14 + S + D	2	Hardware Type	0x0001	Compare	
16 + S + D	2	Protocol Type	0x0800	Compare	
18 + S + D	1	Hardware Size	0x06	Compare	
19 + S + D	1	Protocol Address Length	0x04	Compare	
20 + S + D	2	Operation	0x0001	Compare	
22 + S + D	6	Sender Hardware Address	-	Ignore	
28 + S + D	4	Sender IP Address	-	Ignore	
32 + S + D	6	Target Hardware Address	-	Ignore	
38 + S + D	4	Target IP Address	IP4AT	Compare	Might match any of three host values in IP4AT

7.4.1.4.2 Directed IPv4 Packet

I217 supports receiving Directed IPv4 packets for wake up if the *IPV4* bit (bit 6) is set in the WUFC register. Three IPv4 addresses are supported, which are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must contain the station's MAC address, a Protocol Type of 0x0800, and one of the three host programmed IPv4 addresses. I217 also handles Directed IPv4 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Possible Len/LLC/SNAP Header		Skip	
12 + S + D	2	Type	0x0800	Compare	IP
14 + S + D	1	Version/ HDR length	0x4X	Compare	Check IPv4
15 + S + D	1	Type of Service	-	Ignore	
16 + S + D	2	Packet Length	-	Ignore	
18 + S + D	2	Identification	-	Ignore	
20 + S + D	2	Fragment Info	-	Ignore	
22 + S + D	1	Time to live	-	Ignore	
23 + S + D	1	Protocol	-	Ignore	
24 + S + D	2	Header Checksum	-	Ignore	
26 + S + D	4	Source IP Address	-	Ignore	
30 + S + D	4	Destination IP Address	IP4AT	Compare	Might match any of three host values in IP4AT



7.4.1.4.2.7 Directed IPv6 Packet

The I217 supports receiving Directed IPv6 packets for wake up if the *IPV6* bit (bit 7) is set in the WUFC register. One IPv6 address is supported, which is programmed in the IPv6 Address Table (IP6AT). A successfully matched packet must contain the station's MAC address, a protocol type of 0x0800, and the programmed IPv6 address. The I217 also handles Directed IPv6 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Possible Len/LLC/SNAP Header		Skip	
12 + S + D	2	Type	0x0800	Compare	IP
14 + S + D	1	Version/ Priority	0x6X	Compare	Check IPv6
15 + S + D	3	Flow Label	-	Ignore	
18 + S + D	2	Payload Length	-	Ignore	
20 + S + D	1	Next Header	-	Ignore	
21 + S + D	1	Hop Limit	-	Ignore	
22 + S + D	16	Source IP Address	-	Ignore	
38 + S + D	16	Destination IP Address	IP6AT	Compare	Match value in IP6AT

7.4.1.4.3 Flexible Filter

I217 supports a total of 16 flexible filters. Each filter can be configured to recognize any arbitrary pattern within the first 128 bytes of the packet. To configure the flexible filter, software programs the mask values into the Flexible Filter Mask Table (FFMT) and the required values into the Flexible Filter Value Table (FFVT), and the minimum packet length into the Flexible Filter Length Table (FFLT). These contain separate values for each filter. Software must also enable the filter in the WUFC register, and enable the overall wake up functionality must be enabled by setting *PME_En* in the Power Management Control Status Register (PMCSR) or the WUC register.

Once enabled, the flexible filters scan incoming packets for a match. If the filter encounters any byte in the packet where the mask bit is one and the byte doesn't match the byte programmed in the Flexible Filter Value Table (FFVT) then the filter fails that packet. If the filter reaches the required length without failing the packet, it passes the packet and generates a wake up event. It ignores any mask bits set to one beyond the required length.

7.4.1.4.3.8 IPv6 Neighbor Solicitation Message Filter

In IPv6, a Neighbor Solicitation Message packet (type 135) is used for address resolution. A flexible filter can be used to check for a Neighborhood Solicitation Message packet (type 135).

Note: The fields checked for detection of a Neighbor Solicitation Message packet (type 135) are type, code and target addresses.

The I217 has the ability to wake up the ME. Any packet that can cause ME wake up must first match the MAC address filtering this includes Exact Unicast/Multicast filtering; hash based Multicast filtering any unicast MAC address if Promiscuous is enabled and Broadcast packets if enabled. Other filters can be enabled based on the MANC register setting.

Figure 10 shows a top level diagram of the ME wake up filters. The following sections describe these filters.




Figure 7-9. ME Receive Wake Up Filters Top Level



7.4.1.4.4 Manageability Registers

The Manageability filtering configuration is done by the manageability firmware. The following is a list of registers in Intel® Ethernet Controller I217 ME control space that are used to setup the manageability filtering:

- Management VLAN filters—MAVTV[x], x=03
- Flexible TCO filters—FTFT table
- Flexible UDP/TCP and IP filters—MFUTP[11:0], MIP4AT, L24IPV6
- Management Filter control registers—MANC, MANC2, MANC3
- Shared MAC filters control—SHRAFER
- IP Binding Control registers—IPV4MBC, IPV4HBC, IPV6BC

These registers are only reset on an Internal Power On Reset.

7.4.1.4.5 Manageability Packets Types

This section describes the different types of packets that can be configured in the receive filters.

Packets cause an ME wake up if they match a specific filter that it is enabled in the MANC register and the *RCV_TCO_EN* bit (bit 17) is set.

A Unicast packet causes an ME wake up if (1) it matches a unicast MAC filter or (2) unicast promiscuous mode is enabled (these are L2 conditions), and it passes any of the enabled manageability filters as defined in the MANC register.

A Multicast packet causes an ME wake up if (1) it passes the L2 multicast filter or (2) multicast promiscuous mode is enabled (these are L2 conditions), and it passes any of the enabled manageability filters as defined in the MANC register.

A Broadcast packet cause an ME wake up if it passes *any* of the enabled manageability filters as defined in the MANC register.

In addition, unicast or multicast packets that match any of the previous conditions with a VLAN header causes an ME wake up if it passes one of the VLAN filters as defined by the MAVTV(x) registers.

7.4.1.4.6 ARP Packets Filtering

I217 has the ability to wake the ME on ARP Request packets, ARP Response packets or both. ARP Request packets wake up the ME if the *ARP_REQ_EN* bit (bit 13) is set in the MANC register. ARP Response packets wake up the ME if *ARP_RES_EN* bit (bit 15) is set in the MANC register. Note that the hardware filter does not check the IP address for ARP Response packets.

There is also a support in Neighbor Solicitation Message packet (type 135) filtering. If the *NEIGHBOR_EN* bit (bit 14) is set in the MANC register, then Neighbor Solicitation Message packets (type 135) wake up the ME.



7.4.1.5 Flexible TCP UDP port and IP address filtering

I217 has the ability to direct packets to the ME if their L4 (TCP/UDP) destination port matches a specific value. There are 12 flexible TCP or UDP port numbers depending on the FLEXPOR(x) bits in the MANC/MANC2 registers and the corresponding MFUTP(x) registers.

The flexible port filters may be further dependant on IP address filtering as defined by the MIP4AT and IP6AT if the EN_IPFLEX(x) bits in the MANC register are set.

Six of the 12 flexible port filters can be configured to filter source port instead of destination port numbers.

To enhance the IPv6 filters four L3 filters were added each filter can be configured to match the 24 LSB of the L3 address and can also be combined to match one of the UDP/TCP flexible ports or ICMPv6 next header.

Note: For incoming tunneled packets TCP/UDP port filtering is done only for IPv6 packets in an IPv4 tunnel.

7.4.1.6 Dedicated TCP UDP port and IP address filtering

I217 has the ability to direct packets to the ME if their L4 (TCP/UDP) destination port matches a specific value. There are two constant port numbers: 0x26F depends on the RMCP_EN bit in the MANC register and 0x298 depends on the EN_0298 bit in the MANC register). Additional dedicated TCP or UDP port filters are controlled by the MANC3 register.

The first two constant L4 filters may be further dependant on IP address filtering as defined by the MIP4AT and IP6AT if the EN_IP_ADDRFILTER bit in the MANC register is set.

Some of the dedicated port filters introduced in MANC3 have the ability to be further dependant on a specific IP address and to pass filtering the packet must pass IP filtering as well as port and protocol filtering.

7.4.1.6.1 L2 Filtering

The 217 has the ability to wake up the ME according to the packet L2 parameters:

Broadcast Filtering—if bit *BR_EN* (bit 24) is set in the MANC register then all broadcast packets wake up the ME.

Unicast Filtering—There are up to eleven MAC addresses that can be allocated to manageability MAC address filtering defined by the Shared Receive Address Low (SHRAL) register and the Shared Receive Address High (SHRAH) register. The MAC address filtering to manageability is controlled by the MNG MAC Must match bit (bit 25) in the MANC register. Packets that match the MAC address filtering must match any of the manageability filters as defined in the SHRAFER register to wake up the ME.

VLAN Filtering—There are four dedicated filters for VLAN addresses, which can be configured in registers MAVTV(x). Each register has a 12-bit field that represents the VID field of the incoming VLAN header and an enable bit. If the enable bit is set, HW compares the VID field to the VID field of the incoming packets. If it matches, Intel® Ethernet Controller I217 wakes up the ME (the packet must pass any of the MAC filters as well).



7.4.1.6.2 Flexible TCO Filtering

The I217 includes two flexible filters as defined by the Flexible TCO Filter Table (FTFT). These filters compare an exact pattern match within the 1st 128 bytes of the packet. Enabling these filters is done by setting *FLEX_TCO1/0_FILTER_EN* bits (bits 6 and 7) in the MANC register.

7.4.1.7 Accessing the I217's Wake Up Register Using MDIC

When software needs to configure the wake up state (either read or write to these registers) the MDIO page should be set to 800 (for host accesses) or 801 (for ME accesses) until the page is not changed to a different value wake up register access is enabled. For more details on wake up configuration using MDIC see [Section 9.10.1](#).

7.5 Network Proxy Functionality

7.5.1 Introduction

In prior operating system releases, ARP and IPv6 neighbor discovery messages were one of the possible wakeup types for the platform. ARP and IPv6 neighbor discovery packets are required to enable other network devices to discover the link layer address used by the PC. Supporting these protocols while the host is in low power state is fundamental to maintain remote network accessibility to the sleeping host. If the host does not respond, other devices in the network will eventually not be able to send routable network traffic (such as IPV4 and IPV6) to the sleeping host.

Prior to network proxy, devices that wanted to maintain their network presence would have configured the ARP and neighbor discovery messages as wake up patterns to the system. Analysis shows that many of these ARP wake ups are unnecessary as they are generated by automated processes whose sole purpose is to verify that the system is alive on the network (Microsoft* has stated in their testing 95% of the wake ups are due entirely to ARP wake ups).

Ethernet devices that implement ARP offload must implement it as defined in the Power Management specification on the NDIS Program Connect site. Specifically, the offload must respond to an ARP Request (operation = 1) by responding with an ARP Reply (operation = 2) as defined in RFC 826.

Note: ARP offload is only supported using signed device drivers in Windows® operating systems.

Ethernet devices that implement IPv6 NS offload must implement it as defined in Power Management specification on the NDIS Program, Connect site. Specifically, the offload must respond to a Neighbor Solicitation (operation = 135) by responding with an NS Advertisement (operation = 136) as defined in RFC 2461. Devices must support at least two NS offloads, each with up to two target IPv6 addresses.

7.5.2 Network Proxy Activation

As part of the system sleep flow and after receiving from the OS the network proxy and WoL patterns the SW driver should follow the following steps to activate network proxy in the I217:

1. Program the WoL patterns according to the WoL flow with the addition of the network proxy specific configuration as described in the following steps.
2. Program the appropriate IPv4/IPv6 addresses in IP4AT and IP6AT registers.



3. Program the relevant L2 MAC addresses or broadcast reception.
4. Enable ARP/NS proxy through IP4V[15:14].

Note: A packet matching both a proxy filter and a WoL filter should only cause WoL.

Also, the I217 should not respond to illegal network proxy packets with CRC or checksum errors.

7.5.3 IPv4 Proxy—ARP

In IPv4 networks, ARP provides the address mapping of the IP address to a corresponding MAC address. ARP is a key protocol for remaining responsive on the network.

The delay time between repeated packets is undefined but may be relatively short. As a consequence it is possible for the transition between the proxy and host to miss packets and for a brief time appear off the network (no ARP response). Since ARP is an unreliable protocol there are no specific requirements for proxies.

The sending node generates an ARP Request as a MAC broadcast datagram. The endpoint with the requested IP address must generate a MAC unicast or MAC broadcast datagram ARP Response informing the sending node of its presence. In order to be fully responsive on the network, the Proxy of a sleeping host must respond to ARP requests by generating the necessary responses. Response packet timings and ARP cache timeout values are undefined in the RFCs 826 and 1122.

The I217 supports responding to ARP Request packets (proxy) if enabled through PRXC register. Three IPv4 addresses are supported, which are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must contain a broadcast MAC address or one of the pre-programmed unicast MAC addresses, a protocol type of 0x0806, an ARP opcode of 0x01, and one of the three programmed IPv4 addresses. The I217 also handles ARP Request packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.



7.5.3.1 ARP Request Packet

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Possible LEN/LLC/SNAP Header		Skip	
12 + S + D	2	Type	0x0806	Compare	ARP
14 + S + D	2	Hardware Type	0x0001	Compare	
16 + S + D	2	Protocol Type	0x0800	Compare	
18 + S + D	1	Hardware Size	0x06	Compare	
19 + S + D	1	Protocol Address Length	0x04	Compare	
20 + S + D	2	Operation	0x0001	Compare	
22 + S + D	6	Sender Hardware Address	-	Ignore	
28 + S + D	4	Sender IP Address	-	Ignore	
32 + S + D	6	Target Hardware Address	-	Ignore	
38 + S + D	4	Target IP Address	IP4AT	Compare	match IP4AT values or zero
42 + S + D	18 - S - D	Padding	0x00	Ignore	Padding to 64bytes
60	4	CRC		Check	

7.5.3.2 ARP Response Packet

Offset	# of Bytes	Field	Value	Action
0	6	Destination Address		Copy from ARP Request Source Address
6	6	Source Address		Station address
12	S = (0/4)	Possible VLAN Tag		Copy from ARP Request
12 + S	D = (0/8)	Possible LLC/SNAP Header		Copy from ARP Request
12 + S + D	2	Type	0x0806	Constant (Copy from ARP Request)
14 + S + D	2	Hardware Type	0x0001	Constant (Copy from ARP Request)
16 + S + D	2	Protocol Type	0x0800	Constant (Copy from ARP Request)
18 + S + D	1	Hardware Size	0x06	Constant (Copy from ARP Request)
19 + S + D	1	Protocol Address Length	0x04	Constant (Copy from ARP Request)
20 + S + D	2	Operation	0x0002	Constant
22 + S + D	6	Sender Hardware Address		Station Address
28 + S + D	4	Sender IP Address		Target IP address from ARP Request or valid IP address if Target IP was zero
32 + S + D	6	Target Hardware Address		Sender MAC address from ARP Request
38 + S + D	4	Target IP Address		Sender IP address from ARP Request
42 + S + D	18 - S - D	Padding	0x00	Padding to 64 bytes
60	4	CRC		Calculate

7.5.4 IPv6 Proxy—Neighbor Discovery

In IPv6 networks, ICMPv6 Neighbor solicitation and Neighbor advertisement provides the address mapping of the IP address to a corresponding MAC address.




Neighbor Discovery is a set of 5 message types that are implemented on ICMPv6. The message types are

- Router Solicitation
- Router Advertisement
- Neighbor Solicitation
- Neighbor Advertisement
- Redirect

Only two of these messages that are significant for resolving IPv6 addresses to the MAC address Neighbor Solicitation and Neighbor Advertisement.

Machines that operate in IPv6 networks are sent an ICMPv6 Neighbor Solicitation and must respond with their link-layer (MAC) address in their ICMPv6 Neighbor Advertisement response. The solicitation may be for either the link-local, global, or a temporary IPv6 addresses.

Neighbor discovery messages have both an IPv6 header and the ICMPv6 header. The IPv6 header is a standard one, including the source and destination IP addresses. The Network proxy offload does not support IPv6 Neighbor discovery messages that also have IPv6 header extensions these packets will be silently discarded with no reply.





7.5.4.1 Ipv6 Neighbor Solicitation Packet

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S=(0/4)	Possible VLAN Tag		Skip	
12+S	D=(0/8)	Possible LLC/SNAP Header		Skip	
IPv6 header					
12+D+S	2	Type	0x86DD	Compare	IPv6
14+D+S	1	Version/ Traffic Class	0x6	Compare	Check IPv6
15+D+S	3	Traffic Class/Flow Label		Ignore	
18+D+S	2	Payload Length		Ignore	
20+D+S	1	Next Header	0x3A	Check	ICMPv6
21+D+S	1	Hop Limit	0xFF	Compare	
22+D+S	16	Source Address		Ignore	Check if source address is undefined
38+D+S	16	Destination Address		Ignore	
ICMPv6 header					
54+D+S	1	Type	0x87	Compare	Neighbor Solicitation
55+D+S	1	Code	0x0	Compare	
56+D+S	2	Checksum		Check	
58+D+S	4	Reserved	0x0000	Ignore	
62+D+S	16	Target IP Address	IP6AT	Compare	
78+D+S	1	Type	0x1	Compare	Possible Source Link Layer Address option (Should not appear if source address is undefined)
79+D+S	1	Length	0x1	Compare	
80+D+S	6	Link Layer Address		Skip	
86+D+S	4	CRC		Check	

7.5.4.2 Ipv6 Neighbor Advertisement Packet

Offset	# of bytes	Field	Value (hex)	Action
0	6	Destination Address		Copy from ND packet
6	6	Source Address		Station Address
12	S=(0/4)	Possible VLAN Tag		Copy from ND packet
12+S	D=(0/8)	Possible LLC/SNAP Header		Copy from ND packet
IPv6 header				
12+D+S	2	Type	0x86DD	Constant (Copy from ND packet)
14+D+S	1	Version/ Traffic Class	0x6	Constant (Copy from ND packet)
15+D+S	3	Traffic Class/Flow Label		Constant (Copy from ND packet)
18+D+S	2	Payload Length		
20+D+S	1	Next Header	0x3A	Constant
21+D+S	1	Hop Limit	0xFF	Constant



Offset	# of bytes	Field	Value (hex)	Action
22+D+S	16	Source Address		relevant IPv6AT entry (ND target address)
38+D+S	16	Destination Address		Copy from ND packet Source address If source address was undefined, send to All Nodes Multicast (FF02:1)
ICMPv6 header				
54+D+S	1	Type	0x88	Constant
55+D+S	1	Code	0x0	
56+D+S	2	Checksum		Calculate
58+D+S	4	Flags	0x60000000	Constant (Solicited, Override) if the source address was defined
			0x20000000	Constant (Override) if the source address was undefined
62+D+S	16	Target IP Address	IP6AT	Same as source address
78+D+S	1	Type	0x2	Target Link Layer Address option
79+D+S	1	Length	0x1	
80+D+S	6	Link Layer Address	From ND	
86+D+S	4	CRC		Calculate

7.5.5 Multicast Listener Discovery Support

Microsoft requires that any device that claims support of NS (IPv6 Neighbor Solicitation) Proxying will also support protocol offload of the Multicast Listener Discovery (MLD) protocol for the solicited node addresses of the device IPv6 addresses.

The purpose of Multicast Listener Discovery (MLD) is to enable each IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. This information is then provided to whichever IPv6 multicast routing protocol such as Neighbor Solicitation (NS) is being used by the router, in order to ensure that multicast packets are delivered to all links where there are interested receivers. If no responses are received on a specific link where MLD queries for a specific multicast address then a packet with this multicast address will not be forwarded to this Link.

MLD is a sub-protocol of ICMPv6, MLD message types are a subset of the set of ICMPv6 messages, and MLD messages are identified in IPv6 packets by a preceding Next Header value of 58 (i.e. the ICMPv6 Next Header value).

Two versions of MLD messaging (MLDv1 and MLDv2) are defined in:

- RFC 2710, Multicast Listener Discovery (MLDv1) for IPv6, 1999
- RFC 3810, Multicast Listener Discovery Version 2 (MLDv2) for IPv6, 2004

MLDv2 is designed to be interoperable with MLDv1.

MLD messages supported by the I217

As a receiver:

- Multicast Listener Query (ICMPv6 Type = decimal 130).



There are two subtypes of Multicast Listener Query messages supported by the I217:

- General Query used to learn which multicast addresses have listeners on an attached link.
- Multicast-Address-Specific Query, used to learn if a particular multicast address has any listeners on an attached link.

As a transmitter:

- Multicast Listener Report (ICMPv6 Type = decimal 131).

7.5.5.1 MLD Query Packet

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Compare	MAC Header—processed by main address filter
6	6	Source Address		Skip	
12	S=(0/4)	Possible VLAN Tag		Skip	
12+S	D=(0/8)	Possible LLC/SNAP Header		Skip	
IPv6 header					
12+D+S	2	Type	0x86DD	Compare	IPv6
14+D+S	1	Version/ Traffic Class	0x6	Compare	Check IPv6
15+D+S	3	Traffic Class/Flow Label		Ignore	
18+D+S	2	Payload Length		Ignore	
20+D+S	1	Next Header	0x00	Compare	Hop by Hop
21+D+S	1	Hop Limit	0x01	Compare	
22+D+S	16	Source Address		Ignore	
38+D+S	16	Destination Address		Compare	
Hop by Hop header					
54+D+S	1	Next Header	0x3A	Compare	ICMPv6
55+D+S	1	Header Extended Length	0x00	Compare	
56+D+S	1	Type	0x05	Ignore	Router Alert
57+D+S	1	Length	0x02	Ignore	
58+D+S	2	MLD	0x0000	Ignore	MLD packet
60+D+S	1	PadN Option	0x01	Ignore	
61+D+S	1	PadN Length	0x00	Ignore	
ICMPv6 header					
62+D+S	1	Type	0x82	Compare	MLD Query
63+D+S	1	Code	0x0	Ignore	
64+D+S	2	Checksum		Check	
66+D+S	2	Maximum Respond Delay	0x0000	Ignore	
68+D+S	2	Reserved	0x0000	Ignore	
70+D+S	16	Multicast IP Address		Compare	104 MSB FF02::0001:FFxx:xxxx assume lower 24 bits match



7.5.5.2 MLD Report Packet

Offset	# of bytes	Field	Value (hex)	Action
0	6	Destination Address		Copy from MLD query packet
6	6	Source Address		Station Address
12	S=(0/4)	Possible VLAN Tag		Copy from MLD packet
12+S	D=(0/8)	Possible LLC/SNAP Header		Copy from MLD packet
IPv6 header				
12+D+S	2	Type	0x86DD	Constant (Copy from MLD packet)
14+D+S	1	Version/ Traffic Class	0x6	Constant (Copy from MLD packet)
15+D+S	3	Traffic Class/Flow Label		Constant (Copy from MLD packet)
18+D+S	2	Payload Length		
20+D+S	1	Next Header	0x00	Constant (Copy from MLD packet)
21+D+S	1	Hop Limit	0x01	Constant (Copy from MLD packet)
22+D+S	16	Source Address		Link Local Address
38+D+S	16	Destination Address		Multicast Address being reported
Hop by Hop header				
54+D+S	1	Next Header	0x3A	Constant (Copy from MLD packet)
55+D+S	1	Header Extended Length	0x00	Constant (Copy from MLD packet)
56+D+S	1	Type	0x05	Constant (Copy from MLD packet)
57+D+S	1	Length	0x02	Constant (Copy from MLD packet)
58+D+S	2	MLD	0x0000	Constant (Copy from MLD packet)
60+D+S	1	PadN Option	0x01	Constant (Copy from MLD packet)
61+D+S	1	PadN Length	0x00	Constant (Copy from MLD packet)
ICMPv6 header				
62+D+S	1	Type	0x83	Constant
63+D+S	1	Code	0x00	Constant
64+D+S	2	Checksum		Calculate
66+D+S	2	Maximum Respond Delay	0x0000	
68+D+S	2	Reserved	0x0000	
70+D+S	16	Multicast Address		Copy from MLD query packet, if undefined send a single packet for each of the I217 IP addresses

7.6 Loopback

PHY loopback is supported in the LAN Connect Device. Software or Firmware should set the LAN Connected Device to the loopback mode via MDIC register writing to the PHY Control Register (Page 0 Register 00). The PHY supports a number of loopback modes configured through the Loopback Control Register (Page 0 Register 19).

For more information on the different loopback modes, see the tables starting on [page 81](#).



The LAN Controller must be in forced link and in full duplex mode for PHY loopback to operate. The following bits must be configured in the LAN Controller to enable PHY loopback:

- CTRL.FRCDPLX = 1b: // Force duplex mode by the integrated LAN controller
- CTRL.FD = 1b: // Set Full Duplex mode

Note: Host driver needs to disable energy detect prior to configuring the LAN Connected Device into loopback mode.



Note: *This page intentionally left blank.*



8 Programmers' Visible State

8.1 Terminology

Shorthand	Description
R/W	Read/Write. A register with this attribute can be read and written. If written since reset, the value read reflects the value written.
R/W S	Read/Write Status. A register with this attribute can be read and written. This bit represents status of some sort, so the value read might not reflect the value written.
RO	Read Only. If a register is read only, writes to this register have no effect.
WO	Write Only. Reading this register might not return a meaningful value.
R/WC	Read/Write Clear. A register bit with this attribute can be read and written. However, a write of 1b clears (sets to 0b) the corresponding bit and a write of 0b has no effect.
R/W SC	Read/Write Self Clearing. When written to 1b the bit causes an action to be initiated. Once the action is complete the bit return to 0b.
RO/LH	Read Only, Latch High. The bit records an event or the occurrence of a condition to be recorded. When the event occurs the bit is set to 1b. After the bit is read, it returns to 0b unless the event is still occurring.
RO/LL	Read Only, Latch Low. The bit records an event. When the event occurs the bit is set to 0b. After the bit is read, it reflects the current status.
RO/SC	Read Only, Self Clear. Writes to this register have no effect. Reading the register clears (set to 0b) the corresponding bits.
RW0	Ignore Read, Write Zero. The bit is a reserved bit. Any values read should be ignored. When writing to this bit always write as 0b.
RWP	Ignore Read, Write Preserving. This bit is a reserved bit. Any values read should be ignored. However, they must be saved. When writing the register the value read out must be written back. (There are currently no bits that have this definition.)

This document names registers as follows.

- By register number
 - Registers 0-15 are independent of the page and can be designated by their register number.
 - When a register number is used for registers 16-21, or 23-28, it refers to the register in page 0.
 - Register 31 in PHY address 01, is the page register itself and doesn't belong to any page. It is always written as register 31.
- By page and register number
 - This can be written out as page x, register y, but is often abbreviated x.y
- By name
 - Most functional registers also have a name.

Register bits are designated by a dot followed by a number after the register address. Thus, bit 4.16.2 is page 4, register 16 and bit 2. Multi-bit fields follow the MSB, colon, LSB convention and so bits 4.16.5:4 is page 4, register 16, bits 5:4. All fields in a register have a name.



Register bits with default values marked with an asterisk * are loaded by the integrated LAN controller during the I217 power up and following reset. Other fields in the same 16-bit register must be loaded with their default values.

8.2 MDIO Access

After LCD reset to the I217 a delay of 10 ms is required before attempting to access MDIO registers.

Access using MDIO should be done only when bit 10 in page 769 register 16 is set.

8.3 Addressing

Addressing is based on the IEEE 802.3 MII Management Interface specification defined in clause 22 of 802.3, particularly section 22.2.4.

The 217 registers are spread over two PHY addresses 01, 02, where general registers are located under PHY address 01 and the PHY specific registers are at PHY address 02. The IEEE specification allows five bits for the register access. Registers 0 to 15 are defined by the specification, while registers 16 to 31 are left available to the vendor. The PHY implements many registers for diagnostic purposes. In addition, the I217 contains registers controlling the custom interface as well as other the I217 functions. The total number of registers implemented far exceeds the 16 registers available to the vendor. When this occurs, a common technique is to use paging. The 217 registers in PHY address 01, are divided into pages. Each page has 32 registers. Registers 0-15 are identical in all the pages and are the IEEE defined registers. Register 31 is the page register in all pages of PHY address 01. All other registers are page specific.

In order to read or write a register, software should define the appropriate PHY address. For PHY address 01, in order to access registers other than 0-15, software should first set the page register to map to the appropriate page. Software can then read or write any register in that page. Setting the page is done by writing page_num x 32 to Register 31. This is because only the 11 MSB's of register 31 are used for defining the page. During write to the page register, the five LSB's are ignored.

In pages 800 and 801, the register address space is more than 32. See [Section 8.10](#) for a description of register addressing in these pages.

Accessing more than 32 registers, in PHY address 02, is done without using pages. Instead, two registers from register address 16 to 31 are used as Address Offset port and Data port for extended set of registers. See [Section 8.5](#) for details about these registers.



8.3.1 Registers Byte Ordering

This section defines the structure of registers that contain fields carried over the network. Some examples are L2, L3, L4 fields. The following example is used to describe byte ordering over the wire (hex notation):

Last First

...,06, 05, 04, 03, 02, 01, 00

where each byte is sent with the LSbit first. That is, the bit order over-the-wire for this example is:

Last First

..., 0000 0011, 0000 0010, 0000 0001, 0000 0000

The general rule for register ordering is to use Host Ordering (also called little endian). Using the above example, a 6-byte fields (e.g. MAC address) is stored in a CSR in the following manner:

Byte 1 Byte0

DW address (N) 0x01 0x00

DW address (N+1) 0x03 0x02

DW address (N+2) 0x05 0x04

The exceptions listed below use network ordering (also called big endian). Using the above example, a 16-bit field (e.g. EtherType) is stored in a register in the following manner:

Byte 1 Byte0

(Word aligned) 0x00 0x01

The following exceptions use network ordering:

- All EtherType fields

Note: The “normal” notation as it appears in text books, etc., is to use network ordering.

Example: Suppose a MAC address of 00-A0-C9-00-00-00. The order on the network is 00, then A0, then C9, etc. However, the host ordering presentation would be

Byte 1 Byte0

DW address (N) A0 00

DW address (N+1) 00 C9

DW address (N+2) 00 00



8.4 Address Map

PHY Address	Page	Register	Name	Page
02	Any	0	Control	71
02	Any	1	Status	72
02	Any	2	PHY Identifier 1	72
02	Any	3	PHY Identifier 2	73
02	Any	4	Auto-Negotiation Advertisement	73
02	Any	5	Auto-Negotiation Link Partner Ability	73
02	Any	6	Auto-Negotiation Expansion	74
02	Any	7	Auto-Negotiation Next Page Transmit	75
02	Any	8	Link Partner Next Page	75
02	Any	9	1000BASE-T Control	75
02	Any	10	1000BASE-T Status	76
02	Any	14:11	Reserved	
02	Any	15	Extended Status	77
02	0	17:16	Reserved	
02	0	18	PHY Control 2	77
02	0	19	Loopback Control	79
02	0	20	Rx Error Counter	79
02	0	21	Management Interface (MI) Control	80
02	0	22	PHY Configuration	80
02	0	23	PHY Control	80
02	0	24	Interrupt Mask	82
02	0	25	Interrupt Status	82
02	0	26	PHY Status	83-
02	0	27	LED Control 1	84
02	0	28	LED Control 2	85
02	0	29	LED Control 3	86
02	0	30	Diagnostics Control (Linking Disabled)	88
02	0	31	Diagnostics Status	89
Page 769—Port Control Registers				
01	769	16	Custom Mode Control	90
01	769	17	Port General Configuration	90
01	769	21	Power Management Control	91
01	769	23	SMBus Control	91
01	769	25	Rate Adaptation Control	91
01	769	27	Flow Control Transmit Timer Value	92
01	769	28	System Low Power Control	92
Page 778—Statistics Registers				
01	778	16-17	Single Collision Count	93
01	778	18-19	Excessive Collisions Count	93
01	778	20-21	Multiple Collisions Count	93
01	778	23-24	Late Collision Count	93
01	778	25-26	Collision Count	94
01	778	27-28	Defer Count	94
01	778	29-30	Transmit with No CRS—TNCRS	94



PHY Address	Page	Register	Name	Page
PCIe Registers				
01	770	16	PCIe FIFOs Control/Status	95
01	770	17	PCIe Power Management Control	95
01	770	18	In-Band Control	96
01	770	20	PCIe Diagnostics	96
01	770	21	Time Outs	96
01	770	23	PCIe K-State Minimum Duration Timeout	96
01	774	23-24	SERDES MDI Control Register	92
LPI Registers				
01	772	18	LowPower Idle GPIO Control	96
01	772	20	Low Power Idle Control	97
01	772	23	Flow Control Refresh Threshold Value	97
01	772	24	Flow Control Thresholds	97
01	772	25	LANWAKE# Control	98
01	772	26	Memories Power	98
01	772	29	Config Register	98
General Registers				
01	776	18	I217 Detect Mode	98
01	776	19	I217 Capability Register	99
01	0	25	OEM Bits	99
01	0	26	SMBus Address	100
01	0	27-28	Shadow Register for RAL0[31:0].	100
01	0	29	Shadow Register for RAH0[15:0].	101
01	0	30	LED Configuration	101
Page 800—Wake Up Registers				
01	800	0	Receive Control Register	103
01	800	1	Wake Up Control Register	104
01	800	2	Wake Up Filter Control Register	105
01	800	3	Wake Up Status Register	105
01	800	16	Receive Address Low 0	106
01	800	18	Receive Address High 0	106
01	800	20-21 + 4*n (n=0...10)	Shared Receive Address Low	106
01	800	22-23 + 4*n (n=0...8,10)	Shared Receive Address High	106
01	800	58 - 59	Shared Receive Address High 9	107
01	800	64	IP Address Valid—IPAV	107
01	800	70	Proxy Control	108
01	800	71	Proxy Code Checksum	108
01	800	72	Proxy Control 2	108
01	800	75	Flex Filters Proxy Control	108
01	800	82-83	IPv4 Address Table—IP4AT 0	107
01	800	88-89	IPv6 Address Table—IP6AT 0	109
01	800	128-191	Multicast Table Array—MTA[31:0]	109
01	800	256 + 2*n (n = 0-127)	Flexible Filter Value Table LSB—FFVT_01	111
01	800	257 + 2*n (n = 0-127)	Flexible Filter Value Table MSB—FFVT_23	111
01	800	512 + 2*n (n = 0-127)	Flexible Filter Value Table—FFVT_45	112
01	800	1024 + 2*n (n = 0-127)	Flexible Filter Value Table—FFVT_67	112
01	800	768 + n (n = 0-127)	Flexible Filter Mask Table—FFMT	112



PHY Address	Page	Register	Name	Page
01	800	896 + n (n = 0 - 3)	Flexible Filter Length Table—FFLT03	113
01	800	904 + n (n = 0 - 1)	Flexible Filter Length Table—FFLT45	113
01	800	908 + n (n = 0 - 1)	Flexible Filter Length Table—FFLT67	113
01	800	2304 + 2*n (n = 0 - 127)	Flexible Filter Value Table—FFVT_89	113
01	800	2305 + 2*n (n = 0 - 127)	Flexible Filter Value Table—FFVT_1011	114
01	800	2560 + 2*n (n = 0 - 127)	Flexible Filter Value Table—FFVT_1213	114
01	800	3072 + 2*n (n = 0 - 127)	Flexible Filter Value Table—FFVT_1415	114
01	800	2816 + n (n = 0 - 127)	Flexible Filter Mask Table 2- FFMT2	114
01	800	2944 + n (n = 0 - 3)	Flexible Filter Length Table—FFLT891011	114
01	800	2952 + n (n = 0 - 1)	Flexible Filter Length Table—FFLT1213	115
01	800	2956 + n (n = 0 - 1)	Flexible Filter Length Table—FFLT1415	114



8.5 PHY Registers (Page 0)

Note: The PHY registers were copied directly from the PHY vendor document.

Table 1. Control Register—Address 0

Bits	Field	Type	Default	Description
15	Reset	R/W, SC	0b	Writing a 1b to this bit causes immediate PHY reset. Once the operation completes, this bit clears to 0b automatically. 1b = PHY reset. 0b = Normal operation.
14	Loopback	R/W	0b	This is the master enable for digital and analog loopback as defined by the IEEE standard. The exact type of loopback is determined by the Loopback Control register (19). 1b = Enables loopback. 0b = Disables loopback.
13	Speed Select (LSB)	R/W	0b	The speed selection address 0 (bits 13 and 6) might be used to configure the link manually. Setting these bits has no effect unless bit 0.12 (AN En) is cleared. 11b = Reserved. 10b = 1000 Mb/s. 01b = 100 Mb/s. 00b = 10 Mb/s.
12	Auto-Negotiation Enable	R/W	1b	When this bit is cleared, the link configuration is determined manually. 1b = Enables auto-negotiation process. 0b = Disables auto-negotiation process.
11	Power Down	R/W	0b	1b = Power down. 0b = Normal operation.
10	Isolate	R/W	0b	Setting this bit isolates the PHY from the MII or GMII interfaces. 1b = Isolates the PHY from MII. 0b = Normal operation.
9	Restart Auto-Negotiation	R/W,SC	0b	1b = Restarts auto-negotiation process. 0b = Normal operation.
8	Duplex Mode	R/W	1b	This bit might be used to configure the link manually. Setting this bit has no effect unless bit 0.12 (AN En) is cleared. 1b = Full-duplex. 0b = Half-duplex.
7	Collision Test	R/W	0b	Enables IEEE 22.2.4.1.9 collision test. 1b = Enable collision test. 0b = Disable collision test.
6	Speed Selection (MSB)	R/W	1b	See description in bit 13.
5:0	Reserved	RO	Always 0x0	Reserved, always set to 0x0.



Table 2. Status Register—Address 1

Bits	Field	Type	Default	Description
15	100BASE-T4	RO	0b	100BASE-T4. This protocol is not supported. This register bit is always set to 0b. 0b = Not 100BASE-T4 capable.
14	100BASE-TX Full-Duplex	RO	1b	1b = 100BASE-TX full duplex capable. 0b = Not 100BASE-TX full duplex capable.
13	100BASE-TX Half-Duplex	RO	1b	1b = 100BASE-TX half duplex capable. 0b = Not 100BASE-TX half duplex capable.
12	10 Mb/s Full-Duplex	RO	1b	1b = 10BASE-T full duplex capable. 0b = Not 10BASE-T full duplex capable.
11	10 Mb/s Half-Duplex	RO	1b	1b = 10BASE-T half duplex capable. 0b = Not 10BASE-T half duplex capable.
10	100BASE-T2 Full-Duplex	RO	0b	Not able to perform 100BASE-T2.
9	100BASE-T2 Half-Duplex	RO	0b	Not able to perform 100BASE-T2.
8	Extended Status	RO	1b	Extended status information in the register Extended Status (0xF).
7	Reserved	RO	0b	Must always be set to 0b.
6	MF Preamble Suppression	RO	1b	1b = PHY accepts management frames with preamble suppressed.
5	Auto-Negotiation Complete	RO	0b	This bit is set after auto-negotiation completes. 1b = Auto-negotiation process complete. 0b = Auto-negotiation process not complete.
4	Remote Fault	RO,LH	0b	This bit indicates that a remote fault has been detected. Once set, it remains set until it is cleared by reading register 1 via the management interface or by PHY reset. 1b = Remote fault condition detected. 0b = Remote fault condition not detected.
3	Auto-Negotiation Ability	RO	1b	1b = PHY able to perform auto-negotiation. 0b = PHY not able to perform auto-negotiation.
2	Link Status	RO,LL	0b	This bit indicates that a valid link has been established. Once cleared, due to link failure, this bit remains cleared until register 1 is read via the management interface. 1b = Link is up. 0b = Link is down.
1	Jabber Detect	RO,LH	0b	1b = Jabber condition detected. 0 = Jabber condition not detected.
0	Extended Capability	RO	1b	Indicates that the PHY provides an extended set of capabilities that might be accessed through the extended register set. For a PHY that incorporates a GMII/RGMII, the extended register set consists of all management registers except registers 0, 1, and 15. 1b = Extended register capabilities.

Table 3. PHY Identifier Register 1—Address 282

Bits	Field	Type	Default	Description
15:0	PHY ID Number ¹	RO	0x0154	The PHY identifier composed of bits 3 through 18 of the Organizationally Unique Identifier (OUI)

Notes:

1. PHY ID Number based on Intel assigned OUI number of 00-AA-00 following bit reversal.

**Table 4. PHY Identifier Register 2—Address 3**

Bits	Field	Type	Default	Description
15:10	PHY ID Number ¹	RO	0x0	The PHY identifier composed of bits 19 through 24 of the OUI.
9:4	Model Number	RO	0x0A	The value is part of the PHY identifier and represents the Device Model Number.
3:0	Revision Number	RO	0x1	The value is part of the PHY identifier and represents the Device Revision Number.

Notes:

- PHY ID Number based on Intel assigned OUI number of 00-AA-00 following bit reversal.

Note:

Any write to the Auto-Negotiation Advertisement register, prior to auto-negotiation completion, is followed by a restart of auto-negotiation. Also note that this register is not updated following auto-negotiation.

Table 5. Auto-Negotiation Advertisement Register—Address 4

Bits	Field	Type	Default	Description
15	Next Page	R/W	0b	1b = Advertises next page ability supported. 0b = Advertises next page ability not supported.
14	Reserved	RO	Always 0b	Must be 0b.
13	Remote Fault	R/W	0b	1b = Advertises remote fault detected. 0b = Advertises no remote fault detected.
12	Reserved	RO	0b	Reserved
11	Asymmetric Pause	R/W	0b	1b = Advertises asymmetric pause ability. 0b = Advertises no asymmetric pause ability.
10	Pause Capable	R/W	0b	1b = Capable of full duplex pause operation. 0b = Not capable of pause operation.
9	100BASE-T4 Capability	RO	0b	The PHY does not support 100BASE-T4. The default value of this register bit is 0b. 1b = 100BASE-T4 capable. 0b = Not 100BASE-T4 capable.
8	100BASE-TX Full-Duplex Capable	R/W	1b	1b = 100BASE-TX full duplex capable. 0b = Not 100BASE-TX full duplex capable.
7	100BASE-TX Half-Duplex Capable	R/W	1b	1b = 100BASE-TX half duplex capable. 0b = Not 100BASE-TX half duplex capable.
6	10BASE-TX Full-Duplex Capable	R/W	1b	1b = 10BASE-TX full duplex capable. 0b = Not 10BASE-TX full duplex capable.
5	10BASE-TX Half-Duplex Capable	R/W	1b	1b = 10BASE-TX half duplex capable. 0b = Not 10BASE-TX half duplex capable.
4:0	Selector Field	R/W	00001b	00001b = IEEE 802.3 CSMA/CD.

Table 6. Auto-Negotiation Link Partner Ability—Address 5

Bits	Field	Type	Default	Description
15	Next Page	RO	0b	1b = Link partner has next page ability. 0b = Link partner does not have next page ability.
14	Acknowledge	RO	0b	1b = Link partner has received link code word. 0b = Link partner has not received link code word.



Bits	Field	Type	Default	Description
13	Remote Fault	RO	0b	1b = Link partner has detected remote fault. 0b = Link partner has not detected remote fault.
12	Reserved	RO	0b	Reserved.
11	Asymmetric Pause	RO	0b	1b = Link partner requests asymmetric pause. 0b = Link partner does not request asymmetric pause.
10	Pause Capable	RO	0b	1b = Link partner is capable of full duplex pause operation. 0b = Link partner is not capable of pause operation.
9	100BASE-T4 Capability	RO	0b	1b = Link partner is 100BASE-T4 capable. 0b = Link partner is not 100BASE-T4 capable.
8	100BASE-TX Full-Duplex Capability	RO	0b	1b = Link partner is 100BASE-TX full-duplex capable. 0b = Link partner is not 100BASE-TX full-duplex capable.
7	100BASE-TX Half-Duplex Capability	RO	0b	1b = Link partner is 100BASE-TX half-duplex capable. 0b = Link partner is not 100BASE-TX half-duplex capable.
6	10BASE-T Full-Duplex Capability	RO	0b	1b = Link partner is 10BASE-T full-duplex capable. 0b = Link partner is not 10BASE-T full-duplex capable.
5	10BASE-T Half-Duplex Capability	RO	0b	1b = Link partner is 10BASE-T half-duplex capable. 0b = Link partner is not 10BASE-T half-duplex capable.
4:0	Protocol Selector Field	RO	0x00	Link partner protocol selector field.

Table 7. Auto-Negotiation Expansion

Bits	Field	Type	Default	Description
15:5	Reserved	RO	0x00	Reserved, must be set to 0x00.
4	Parallel Detection Fault	RO,LH	0b	1b = Parallel link fault detected. 0b = Parallel link fault not detected.
3	Link Partner Next Page Ability	RO	0b	1b = Link partner has next page capability. 0b = Link partner does not have next page capability.
2	Next Page Capability	RO, LH	1b	1b = Local device has next page capability. 0b = Local device does not have next page capability.
1	Page Received	RO, LH	0b	1b = A new page has been received from a link partner. 0b = A new page has not been received from a link partner.
0	Link Partner Auto-Negotiation Ability	RO	0b	1b = Link partner has auto-negotiation capability. 0b = Link partner does not have auto-negotiation capability.

**Table 8. Auto-Negotiation Next Page Transmit**

Bits	Field	Type	Default	Description
15	Next Page	R/W	0b	1b = Additional next pages to follow. 0b = Sending last next page.
14	Reserved	RO	0b	Reserved.
13	Message Page	R/W	1b	1b = Formatted page. 0b = Unformatted page.
12	Acknowledge 2	R/W	0b	1b = Complies with message. 0b = Cannot comply with message.
11	Toggle	RO	0b	1b = Previous value of transmitted link code word was a logic zero. 0b = Previous value of transmitted link code word was a logic one.
10:0	Message/ Unformatted Field	R/W	0x3FF	Next page message code or unformatted data.

Table 9. Link Partner Next Page

Bits	Field	Type	Default	Description
15	Next Page	RO	0b	1b = Additional next pages to follow. 0b = Sending last next page.
14	Acknowledge	RO	0b	1b = Acknowledge. 0b = No acknowledge.
13	Message Page	RO	0b	1b = Formatted page. 0b = Unformatted page.
12	Acknowledge2	RO	0b	1b = Complies with message. 0b = Cannot comply with message.
11	Toggle	RO	0b	1b = Previous value of transmitted link code word was a logic zero. 0b = Previous value of transmitted link code word was a logic one.
10:0	Message/ Unformatted Code Field	RO	0x00	Next page message code or unformatted data.

Table 10. 1000BASE-T Control

Bits	Field	Type	Default	Description
15:13	Test Mode	R/W	000b	000b = Normal mode. 001b = Test Mode 1—Transmit waveform test. 010b = Test Mode 2—Master transmit jitter test. 011b = Test Mode 3—Slave transmit jitter test. 100b = Test Mode 4—Transmit distortion test. 101b, 110b, 111b = Reserved.
12	Master/Slave Manual Configuration Enable	R/W	0b	1b = Enables master/slave configuration. 0b = Automatic master/slave configuration.
11	Master/Slave Configuration Value	R/W	0b	Setting this bit has no effect unless address 9, bit 12 is set. 1b = configures PHY as a master. 0b = Configures PHY as a slave.



Bits	Field	Type	Default	Description
10	Port Type	R/W	0b	1b = Multi-port device (prefer master). 0b = Single port device (prefer slave).
9	Advertise 1000BASE-T Full-Duplex Capability	R/W	0b	1b = Advertises 1000BASE-T full-duplex capability. 0b = Advertises no 1000BASE-T full-duplex capability.
8	Advertise 1000BASE-T Half-Duplex Capability	R/W	0b	1b = Advertises 1000BASE-T half-duplex capability. 0b = Advertises no 1000BASE-T half-duplex capability. Note: 1000BASE-T half-duplex not supported.
7:0	Reserved	RO	0x00	Set these bits to 0x00.

Note: Logically, bits 12:8 can be regarded as an extension of the *Technology Ability* field in Register 4.

Table 11. 1000BASE-T Status

Bits	Field	Type	Default	Description
15	Master/Slave Configuration Fault	RO,LH, SC	0b	Once set, this bit remains set until cleared by the following actions: <ul style="list-style-type: none"> • Read of Register 10 via the management interface. • Reset. • Auto-negotiation completed. • Auto-negotiation enabled. 1b = Master/slave configuration fault detected. 0b = No master/slave configuration fault detected.
14	Master/Slave Configuration Resolution	RO	0b	This bit is not valid when bit 15 is set. 1b = Local PHY resolved to master. 0b = Local PHY resolved to slave.
13	Local Receiver Status	RO	0b	1b = Local receiver is correct. 0b = Local receiver is incorrect.
12	Remote Receiver Status	RO	0b	1b = Remote receiver is correct. 0b = Remote receiver is incorrect.
11	Link Partner 1000BASE-T Full-Duplex Capability	RO	0b	1b = Link partner 1000BASE-T full-duplex capable. 0b = Link partner not 1000BASE-T full-duplex capable. Note: Logically, bits 11:10 might be regarded as an extension of the Technology Ability field in Register 5.
10	Link Partner 1000BASE-T Half-Duplex Capability	RO	0b	1b = Link partner 1000BASE-T half-duplex capable. 0b = Link partner not 1000BASE-T half-duplex capable. Note: Logically, bits 11:10 might be regarded as an extension of the Technology Ability field in Register 5.
9:8	Reserved	RO	00b	Reserved
7:0	Idle Error Count	RO	0x00	These bits contain a cumulative count of the errors detected when the receiver is receiving idles and both local and remote receiver status are operating correctly. The count is held at 255 in the event of overflow and is reset to zero by reading Register 10 via the management interface or by reset. MSB of idle error count.



Table 12. Extended Status

Bits	Field	Type	Default	Description
15	1000BASE-X Full-Duplex	RO	0b	0b = Not 1000BASE-X full-duplex capable.
14	1000BASE-X Half-Duplex	RO	0b	0b = Not 1000BASE-X half-duplex capable.
13	1000BASE-T Full-Duplex	RO	1b	1b = 1000BASE-T full-duplex capable. 0b = Not 1000BASE-T full-duplex capable.
12	1000BASE-T Half-Duplex	RO	1b	1b = 1000BASE-T half-duplex capable. 0b = Not 1000BASE-T half-duplex capable.
11:0	Reserved	RO	0x00	Reserved

Table 13. PHY Control 2

Bits	Field	Type	Default	Description
15	Resolve MDI/MDI-X Before Forced Speed	R/W	1b	1b = Resolves MDI/MDI-X configuration before forcing speed. 0b = Does not resolve MDI/MDI-X configuration before forcing speed.
14	Count False Carrier Events	R/W	0b	Count symbol errors (bit 13) and count false carrier events (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (refer to Table 40). The default is to count CRC errors. 1b = Rx error counter counts false carrier events. 0b = Rx error counter does not count false carrier events.
13	Count Symbol Errors	R/W	0b	Count symbol errors (bit 13) and count false carrier events (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (refer to Table 40). The default is to count CRC errors. 1b = Rx error counter counts symbol errors. 0b = Rx error counter counts CRC errors.
12:11	Reserved			Reserved.
10	Automatic MDI/MDI-X	R/W	1b	1b = Enables automatic MDI/MDI-X configuration. 0b = Disables automatic MDI/MDI-X configuration.
9	MDI-MDI-X Configuration	R/W	0b	1b = Manual MDI-X configuration. 0b = Manual MDI configuration. Refer to Table 41 for further information.
8:3	Reserved		0x0	Reserved.
2	Enable Diagnostics	R/W	0b	This bit enables PHY diagnostics, which include IP phone detection and TDR cable diagnostics. It is not recommended to enable this bit in normal operation (when the link is active). This bit does not need to be set for link analysis cable diagnostics. 1b = Enables diagnostics. 0b = Disables diagnostics.
1:0	Reserved		0x0	Reserved.



Table 14. Rx Error Counter Characteristics

Count False Carrier Events	Count Symbol Errors	Rx Error Counter
1	1	Counts symbol errors and false carrier events.
1	0	Counts CRC errors and false carrier events.
0	1	Counts symbol errors.
0	0	Counts CRC errors.

Bit 9 of the PHY Control register manually sets the MDI/MDI-X configuration if automatic MDI-X is disabled (refer to [Table 41](#)).

Table 15. MDI/MDI-X Configuration Parameters

Automatic MDI/MDI-X	MDI/MDI-X Configuration	MDI/MDI-X Mode
1	X	Automatic MDI/MDI-X detection.
0	0	MDI configuration (NIC/DTE).
0	1	MDI-X configuration (switch).

The mapping of the transmitter and receiver to pins for MDI and MDI-X configuration for 10BASE-T, 100BASE-TX, and 1000BASE-T is listed in. Note that even in manual MDI/MDI-X configuration, the PHY automatically detects and corrects for C and D pair swaps.

Table 16. MDI/MDI-X Pin Mapping

Pin	MDI Pin Mapping			MDI-X Pin Mapping		
	10BASE-T	100BASE-TX	1000BASE-T	10BASE-T	100BASE-TX	1000BASE-T
TRD[0]+/-	Tx +/-	Tx +/-	Tx A+/- Rx B+/-	Rx +/-	Rx +/-	Tx B+/- Rx A+/-
TRD[1]+/-	Rx +/-	Rx +/-	Tx B+/- Rx A+/-	Tx +/-	Tx +/-	Tx A+/- Rx B+/-
TRD[2]+/-			Tx C+/- Rx D+/-			Tx D+/- Rx C+/-
TRD[3]+/-			Tx D+/- Rx C+/-			Tx C+/- Rx D+/-



Table 17. Loopback Control Register—Address 19

Bits	Field	Type	Default	Description
15	MII	R/W	0b	1b = MII loopback selected. 0b = MII loopback not selected.
14:13	Reserved			Reserved
12	All Digital	R/W	1b	1b = All digital loopback enabled. 0b = All digital loopback disabled.
11	Reserved			Reserved
10	Line Driver	R/W	0b	1b = Line driver loopback enabled. 0b = Line driver loopback disabled.
9	Remote	R/W	0b	1b = Remote loopback enabled. 0b = Remote loopback disabled.
8	Reserved	R/W		Reserved
7	External Cable	R/W	0b	1b = External cable loopback enabled. 0b = External cable loopback disabled.
6	Tx Suppression	R/W	1b	1b = Suppress Tx during all digital loopback. 0b = Do not suppress Tx during all digital loopback.
5:1	Reserved			Reserved
0	Force Link Status	R/W	1b	This bit can be used to force link status operational during MII loopback. In MII loopback, the link status bit is not set unless force link status is used. In all other loopback mode, the link status bit is set when the link comes up. 1b = Forces link status operational in MII loopback. 0b = Forces link status not operational in MII loopback.

8.5.1 Loopback Mode Settings

Table 18 lists how the loopback bit (Register 0, bit 14) and the *Link Enable* bit (Register 23, bit 13) should be set for each loopback mode. It also indicates whether the loopback mode sets the *Link Status* bit and when the PHY is ready to receive data.

Table 18. Loopback Bit (Register 0, Bit 14) Settings for Loopback Mode

Loopback	Register 0, Bit 14 = 1b	Register 26, Bit 6 (Link Status Set)	PHY Ready for Data
MII	Yes	Register 19, bit 0	After a few ms
All Digital	Yes	Yes	Link Status
Line Driver	Yes	Yes	Link Status
Ext Cable	No	Yes	Link Status
Remote	No	Yes	Never

Table 19. Rx Error Counter Register—Address 20

Bits	Field	Type	Default	Description
15:0	Rx Error Counter	RO, SC	0x00	16-bit Rx error counter. Note: Refer to Register 18, bits 13 and 14 for error type descriptions.



Table 20. Management Interface (MI) Register—Address 21

Bits	Field	Type	Default	Description
15:4	Reserved		0x0	Reserved
3	Energy Detect Power Down Enable	R/W	1b	1b = Enables energy detect power down. 0b = Disables energy detect power down.
2	Energy-Detect Powerdown Mode Transmit Enable	R/W	1b	1: Enables NLP transmission during energy-detect powerdown. 0: Disables NLP transmission during energy-detect powerdown.
1:0	Reserved		0x0	Reserved

Table 21. PHY Control Register—Address 22

Bits	Field	Type	Default	Description
15	CRS Transmit Enable	R/W	0b	1b = Enables CRS on transmit in half-duplex mode. 0b = Disables CRS on transmit.
14:12	Reserved	R/O	0x0	Reserved
11:10	Automatic Speed Downshift Mode	R/W	11b	If automatic downshift is enabled and the PHY fails to auto-negotiate at 1000BASE-T, the PHY falls back to attempt connection at 100BASE-TX and, subsequently, 10BASE-T. This cycle repeats. If the link is broken at any speed, the PHY restarts this process by re-attempting connection at the highest possible speed (1000BASE-T). 00b = Automatic speed downshift disabled. 01b = 10BASE-T downshift enabled. 10b = 100BASE-TX downshift enabled. 11b = 100BASE-TX and 10BASE-T enabled.
9:8	Reserved			Reserved
7	Alternate Next Page	RO	0b	1b = Enables manual control of 1000BASE-T next pages only. 0b = Normal operation of 1000BASE-T next page exchange.
6	Group MDIO Mode Enable	R/W	0b	1b = Enables group MDIO mode. 0b = Disables group MDIO mode.
5	Transmit Clock Enable	R/W	0b	When this bit is set, the transmit test clock is available on pin TX_TCLK. 1b = Enables output of mixer clock (transmit clock in 1000BASE-T). 0b = Disables output.
4:0	Reserved		0x0	Reserved

Table 22. PHY Control Register—Address 23

Bits	Field	Type	Default	Description
15:14	Reserved	R/O	0x0	Reserved



Table 22. PHY Control Register—Address 23

Bits	Field	Type	Default	Description
13	LNK_EN (Link Enable)	R/W	1b	If LNK_EN is set, the PHY attempts to bring up a link with a remote partner and monitors the MDI for link pulses. If LNK_EN is cleared, the PHY takes down any active link, goes into standby, and does not respond to link pulses from a remote link partner. In standby, IP phone detect and TDR functions are available. 1b = Enables linking. 0b = Disables linking.
12:10	Automatic Speed Downshift Attempts Before Downshift	R/W	100b	000b = 1. 001b = 2. 010b = 3. 011b = 4. 100b = 5. 101b = 6. 110b = 7. 111b = 8.
9:8	Reserved		0x0	Reserved
7	Link Partner Detected	RO, LH	0b	When linking is disabled, the PHY automatically monitors for the appearance of a link partner and sets this bit if detected. Linking is disabled when LNK_EN is cleared (bit 13 = 0b). 1b = Link partner detected. 0b = Link partner not detected.
6	Jabber (10BASE-T)	R/W	1b	1b = Disables jabber. 0b = Normal operation.
5	SQE (10BASE-T)	R/W	0b	1b = Enables heartbeat. 0b = Disables heartbeat.
4	TP_LOOPBACK (10BASE-T)	R/W	0b	1b = Disables TP loopback during half duplex. 0b = Normal operation.
3:2	10BASE-T Preamble Length	R/W	10b	00b = 10BASE-T preamble length of zero bytes sent. 01b = 10BASE-T preamble length of one byte sent. 10b = 10BASE-T preamble length of two bytes sent. 11b = 10BASE-T preamble length of seven bytes sent.
1	Reserved		0b	Reserved
0	Force Interrupt	R/W	0b	1b = Asserts MDINT_N pin. 0b = Deasserts MDINT_N pin.

Table 23. Interrupt Mask Register—Address 24

Bits	Field	Type	Default	Description
15:11	Reserved		0x0	Reserved.
10	TDR/IP Phone	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
9	MDIO Sync Lost	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
8	Auto-Negotiation Status Change	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
7	CRC Errors	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
6	Next Page Received	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
5	Error Count Full	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
4	FIFO Overflow/ Underflow	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
3	Receive Status Change	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
2	Link Status Change	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
1	Automatic Speed Downshift	R/W	0b	1b = Interrupt enabled. 0b = Interrupt disabled.
0	MDINT_N Enable	R/W	0b	1b = MDINT_N enabled. ¹ 0b = MDINT_N disabled.

Notes:

1. MDINT_N is asserted (active low) if MII interrupt pending = 1b.

Table 24. Interrupt Status Register—Address 25

Bits	Field	Type	Default	Description
15:11	Reserved		0x0	Reserved.
10	TDR/IP Phone	RO, LH	0b	1b = Event completed. 0b = Event has not completed.
9	MDIO Sync Lost	RO, LH	0b	If the management frame preamble is suppressed (MF preamble suppression, Register 0, bit 6), it is possible for the PHY to lose synchronization if there is a glitch at the interface. The PHY can recover if a single frame with a preamble is sent to the PHY. The MDIO sync lost interrupt can be used to detect loss of synchronization and, thus, enable recovery. 1b = Event has occurred. 0b = Event has not occurred.
8	Auto-Negotiation Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
7	CRC Errors	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
6	Next Page Received	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.



Bits	Field	Type	Default	Description
5	Error Count Full	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
4	FIFO Overflow/ Underflow	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
3	Receive Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
2	Link Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
1	Automatic Speed Downshift	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
0	MII Interrupt Pending	RO, LH	0b	An event has occurred and the corresponding interrupt mask bit is enabled (set to 1b). 1b = Interrupt pending. 0b = No interrupt pending.

Table 25. PHY Status Register—Address 26

Bits	Field	Type	Default	Description
15	PHY in Standby	RO	0b	This bit indicates that the PHY is in standby mode and is ready to perform IP phone detection or TDR cable diagnostics. The PHY enters standby mode when LNK_EN is cleared (Register 23, bit 13 = 0b) and exits standby mode and attempts to auto- negotiate a link when LNK-EN is set (Register 23, bit 13 = 1b). 1b = PHY in standby mode. 0b = PHY not in standby mode.
14:13	Auto-Negotiation Fault Status	RO	00b	11b = Reserved. 10b = Master/slave auto-negotiation fault. 01b = Parallel detect auto-negotiation fault. 00b = No auto-negotiation fault.
12	Auto-Negotiation Status	RO	0b	1b = Auto-negotiation complete. 0b = Auto-negotiation not complete.
11	Pair Swap on Pairs A and B	RO	0b	1b = Pairs A and B swapped. 0b = Pairs A and B not swapped.
10	Polarity Status	RO	1b	1b = Polarity inverted (10BASE-T only). 0b = Polarity normal (10BASE-T only).
9:8	Speed Status	RO	11b	11b = Undetermined. 10b = 1000BASE-T. 01 = 100BASE-TX. 00b = 10BASE-T.
7	Duplex Status	RO	0b	1b = Full duplex. 0b = Half duplex.
6	Link Status	RO	0b	1b = Link up. 0b = Link down.
5	Transmit Status	RO	0b	1b = PHY transmitting a packet. 0b = PHY not transmitting a packet.
4	Receive Status	RO	0b	1b = PHY receiving a packet. 0b = PHY not receiving a packet.
3	Collision Status	RO	0b	1b = Collision occurring. 0b = Collision not occurring.



Bits	Field	Type	Default	Description
2	Auto-Negotiation Enabled	RO	0b	1b = Both partners have auto-negotiation enabled. 0b = Both partners do not have auto-negotiation enabled.
1	Link Partner Advertised PAUSE	RO	0b	1b = Link partner advertised PAUSE. 0b = Link partner did not advertise PAUSE.
0	Link Partner Advertised Asymmetric PAUSE	RO	0b	1b = Link partner advertised asymmetric PAUSE. 0b = Link partner did not advertise asymmetric PAUSE.

Table 26. LED Control Register 1—Address 27

Bits	Field	Type	Default	Description
15	Two-Color Mode LED_100/LED_10	R/W	0b	If two-color mode is enabled for pair LED_LNK/ACT and LED_1000, the signal output for LED_LNK/ACT is equal to LED_LNK/ACT and LED_1000. When LED_LNK/ACT and LED_1000 are not mutually exclusive (such as duplex and collision), this mode can simplify the external circuitry because it ensures either LED_LNK/ACT and LED_1000 is on, and not both at the same time. The same rule applies to pair LED_100 and LED_10. 1b = Two-color mode for LED_100 and LED_10. 0b = Normal mode for LED_100 and LED_10.
14	Two-Color Mode LED_LNK/ACT/LED_1000	R/W	0b	If two-color mode is enabled for pair LED_LNK/ACT and LED_1000, the signal output for LED_LNK/ACT is equal to LED_LNK/ACT and LED_1000. When LED_LNK/ACT and LED_1000 are not mutually exclusive (such as duplex and collision), this mode can simplify the external circuitry because it ensures either LED_LNK/ACT and LED_1000 is on, and not both at the same time. The same rule applies to pair LED_100 and LED_10. 1b = Two-color mode for LED_LNK/ACT and LED_1000. 0b = Normal mode for LED_LNK/ACT and LED_1000.
13	LED_10 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. The default value of this bit is set by the LED_CFG pin at reset. 1b = Extended modes for LED_10. 0b = Standard modes for LED_10.
12	LED_100 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_100. 0b = Standard modes for LED_100.
11	LED_1000 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_1000. 0b = Standard modes for LED_1000.
10	LED_LNK/ACT Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_LNK/ACT. 0b = Standard modes for LED_LNK/ACT.



Bits	Field	Type	Default	Description
9:8	Reserved		0x0	Reserved
7:4	LED Blink Pattern Pause	R/W	0x0	LED blink pattern pause cycles.
3:2	LED Pause Duration	R/W	00b	The pulse duration for the setting, Register 27, bits 3:2 = 11b, can be programmed in the range 0 ms to 2 s, in steps of 4 ms using the extended register set. 00b = Stretch LED events to 32 ms. 01b = Stretch LED events to 64 ms. 10b = Stretch LED events to 104 ms. 11b = Reserved.
1	LED Output Disable	R/W	0b	The default value of this bit is set by the LED_CFG pin at reset. 1b = Disables LED outputs. 0b = Enables LED outputs.
0	Pulse Stretch 0	R/W	1b	1b = Enables pulse stretching of LED functions: transmit activity, receive activity, and collision. 0b = Disables pulse stretching of LED functions: transmit activity, receive activity, and collision.

Table 27. LED Control Register 2—Address 28

Bits	Field	Type	Default	Description
15:12	LED_10	R/W	LED_CFG	See description for bits 3:0.
11:8	LED_100	R/W	LED_CFG	See description for bits 3:0.
7:4	LED_1000	R/W	LED_CFG	See description for bits 3:0.
3:0	LED_LNK/ ACT	R/W	LED_CFG	Standard modes: 0000 = 1000BASE-T. 0001 = 100BASE-TX. 0010 = 10BASE-T. 0011 = 1000BASE-T on; 100BASE-TX blink. 0100 = Link established. 0101 = Transmit. 0110 = Receive. 0111 = Transmit or receive activity. 1000 = Full duplex. 1001 = Collision. 1010 = Link established (on) and activity (blink). 1011 = Link established (on) and receive (blink). 1100 = Full duplex (on) and collision (blink). 1101 = Blink. 1110 = On. 1111 = Off. Extended modes: 0000 = 10BASE-T or 100BASE-TX. 0001 = 100BASE-TX or 1000BASE-T. 0010 = 10BASE-T (on) and activity (blink). 0011 = 100BASE-TX (on) and activity (blink). 0100 = 1000BASE-T (on) and activity (blink). 0101 = 10BASE-T or 100BASE-TX on and activity (blink). 0110 = 100BASE-TX or 1000BASE-T on and activity (blink). 0111 = 10BASE-T or 1000BASE-T. 1000 = 10BASE-T or 1000BASE-T on and activity (blink). 1xxx = Reserved.

Table 28. LED Control Register 3—Address 29


Bits	Field	Type	Default	Description
15:14	LED Blink Pattern Address	R/W	00b	Select LED blink pattern register set. 00b = Select register set for LED_LNK/ACT. 01b = Select register set for LED_1000. 10b = Select register set for LED_100. 11b = Select register set for LED_10.
13:8	LED Blink Pattern Frequency	R/W	0x1F	LED blink pattern clock frequency divide ratio. The default pattern is 512 ms blink.
7:0	LED Blink Pattern	R/W	0x55	LED blink pattern. The default pattern is 512 ms blink.

Table 29. Late Collision Counter—Address 30

Bits	Field	Type	Default	Description
7:0	LateColCnt02	RO/SC	0x00	When the PHY is operating in half duplex mode collisions may occur during bytes 0 to 63 of a transmit frame. Collisions occurring after that are counted by the late collision counters as follows: Late col cnt 0: Collisions during bytes 64 - 95 Late col cnt 1: Collisions during bytes 96 - 127 Late col cnt 2: Collisions during bytes 128 - 191 Late col cnt 3: Collisions during bytes 192 - 319 If any counter saturates then all of the counters stop incrementing. To use the late collision counters the LateColCntEn register in the indirect access address map must be set.
15:0	LateColCnt13	RO/SC	0x00	Each of the late collision counters is 8 bits. They can be read two at a time. The LateColCnt02 field corresponds either to counter 0 or 2 and the LateColCnt13 field corresponds either to counter 1 or 3 depending on the value of LateColCnt23Sel. When LateColCnt23Sel is clear counters 0 and 1 may be read. After reading counters 0 and 1 all of the counters stop incrementing and LateColCnt23Sel is set so as to allow counters 2 and 3 to be read. After reading counters 2 and 3 all of the counters are cleared and start incrementing once again and LateColCnt23Sel is cleared.

Table 30. Link Polarity and Length Status—Address 31

Bits	Field	Type	Default	Description
15	LateColCnt23Sel	RO	0b	This bit indicates that late collision counters 2 and 3 are selected for read. See LateColCnt02 and LateColCnt13 for more information.
14	Pair23Swap	RO	0b	The pairs 2 and 3 swapped bit indicates that the PHY has detected that dimensions 2 and 3 are swapped. This bit is asserted when the PHY determines that it is receiving on pair 2 the signal that is being transmitted by the link partner on pair 2 whereas this signal is expected to be received on pair 3.



Bits	Field	Type	Default	Description
13	Pair3PolInv	RO	0b	The pair 3 polarity inverted bit indicates that the PHY has detected that the polarity of the signal that it is receiving on pair 3 is inverted. The PHY can only detect this condition in the case of a 1000BASE-T link.
12	Pair2PolInv	RO	0b	The pair 2 polarity inverted bit indicates that the PHY has detected that the polarity of the signal that it is receiving on pair 2 is inverted. The PHY can only detect this condition in the case of a 1000BASE-T link.
11	Pair1PolInv	RO	0b	The pair 1 polarity inverted bit indicates that the PHY has detected that the polarity of the signal that it is receiving on pair 1 is inverted. The PHY can detect this condition in the case of a 10BASE-T link when the Pair01Swap bit is clear or in the case of a 1000BASE-T link.
10	Pair0PolInv			The pair 0 polarity inverted bit indicates that the PHY has detected that the polarity of the signal that it is receiving on pair 0 is inverted. The PHY can detect this condition in the case of a 10BASE-T link when the Pair01Swap bit is set or in the case of a 1000BASE-T link.
9:2	DspCblLen	RO	0xFF	This field allows the cable length estimate determined by the DSP engine to be observed. This estimate is computed during establishment of a 100BASE-TX or 1000BASE-T link. The reported cable length is in meters with the value 8'd255 indicating indeterminate.
0	B1000DscrAcqErr	RO	0b	The 1000BASE-T descrambler acquisition error bit indicates that 1000BASE-T link establishment failed at the descrambler acquisition stage. The most likely explanation for this failure is excessive pair skew. Once this bit goes high it latches high until it is cleared by a subsequent successful 1000BASE-T descrambler acquisition



Table 31. Diagnostics Status Register (Linking Disabled)—Address 31

Bits	Field	Type	Default	Description
15:14	TDR Fault Type Pair X	R/W, SC	11b	<p>The first time this register is read after automatic TDR analysis completed, it indicates the fault type for pair A. The second time it is read, it indicates the fault type for pair B, the third for pair C, and the fourth for pair D. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported.</p> <p>A value of 01b indicates either an open or a short. If bits 13:10 of Register 31 equal 0000b, it is an open. For all other values of bits 13:10 in Register 31, each bit indicates a short to pair A, B, C, and D.</p> <p>A value of 11b indicates that the results for this pair are invalid. An invalid result usually occurs when unexpected pulses are received during the TDR operation. For example, from a remote PHY that is also doing TDR or trying to bring up a link. When an invalid result is indicated, the distance in bits 9:2 of Register 31 is 0xFF and should be ignored.</p> <p>11b = Result invalid 10b = Open or short found on pair X. 01b = Strong impedance mismatch found on pair X. 00b = Good termination found on pair X.</p>
13	Short Between Pairs X and D	RO	0b	<p>The first time these bits are read after automatic TDR analysis has completed, indicate a short between pair A and pair A, B, C, and D, respectively. The second time they are read, indicate a short between pair B and pair A, B, C, and D, respectively. The third time with pair C and the fourth time with pair D. It then cycles back to pair A. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported.</p> <p>1b = Short between pairs X and D. 0b = No short between pairs X and D.</p>
12	Short Between Pairs X and C	RO	0b	<p>The first time these bits are read after automatic TDR analysis has completed, indicate a short between pair A and pair A, B, C, and D, respectively. The second time they are read, indicate a short between pair B and pair A, B, C, and D, respectively. The third time with pair C and the fourth time with pair D. It then cycles back to pair A. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported.</p> <p>1b = Short between pairs X and C. 0b = No short between pairs X and C.</p>



Bits	Field	Type	Default	Description
11	Short Between Pairs X and B	RO	0b	The first time these bits are read after automatic TDR analysis has completed, indicate a short between pair A and pair A, B, C, and D, respectively. The second time they are read, indicate a short between pair B and pair A, B, C, and D, respectively. The third time with pair C and the fourth time with pair D. It then cycles back to pair A. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported. 1b = Short between pairs X and B. 0b = No short between pairs X and B.
10	Short Between Pairs X and A	RO	0b	The first time these bits are read after automatic TDR analysis has completed, indicate a short between pair A and pair A, B, C, and D, respectively. The second time they are read, indicate a short between pair B and pair A, B, C, and D, respectively. The third time with pair C and the fourth time with pair D. It then cycles back to pair A. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported. 1b = Short between pairs X and A. 0b = No short between pairs X and A.
9:2	Distance to Fault	RO	0x0	Distance to first open, short, or SIM fault on pair X. The first time this register is read, after automatic TDR analysis has completed, it indicates the distance to the first fault on pair A. The second time it is read, it indicates the distance to the first fault on pair B, the third time on pair C, and the fourth time on pair D. It then cycles back to pair A. The pair indication bits (Register 31, bits 1:0) indicate to which pair the results correspond to. Bits 13:12 of Register 30 can be used to specify a pair other than pair A as the first dimension that is reported. This 8-bit integer value is the distance in meters. The value 0xFF indicates an unknown result.
1:0	Pair Indication	RO	00b	These bits indicate the pair to which the results in bits 15:2 of Register 31 correspond to. 00b = results are for pair A. 01b = results are for pair B. 10b = results are for pair C. 11b = results are for pair D.

Table 32. Diagnostics Status Register (Linking Enabled)—Address 31

Bits	Field	Type	Default	Description
15	Reserved			Reserved.
14	Pair Swap on Pairs C and D	RO	0b	If this bit is set, the PHY has detected that received pair 2 (RJ-45 pins 4 and 5) and pair 3 (RJ-45 pins 7 and 8) have crossed over. 1b = Pairs C and D are swapped (1000BASE-T only). 0b = Pairs C and D are not swapped (1000BASE-T only).



Bits	Field	Type	Default	Description
13	Polarity on Pair D	RO	0b	1b = Polarity on pair D is inverted (1000BASE-T only). 0b = Polarity on pair D is normal (1000BASE-T only).
12	Polarity on Pair C	RO	0b	1b = Polarity on pair C is inverted (1000BASE-T only). 0b = Polarity on pair C is normal (1000BASE-T only).
11	Polarity on Pair B	RO	0b	1b = Polarity on pair B is inverted (10BASE-T or 1000BASE-T only). 0b = Polarity on pair B is normal (10BASE-T or 1000BASE-T only).
10	Polarity on Pair A	RO	0b	1b = Polarity on pair A is inverted (10BASE-T or 1000BASE-T only). 0b = Polarity on pair A is normal (10BASE-T or 1000BASE-T only).
9:2	Cable Length	RO	0b	This 8-bit integer value is the cable length in meters when the link is active. The value 0xFF indicates an unknown result. Cable length when the link is active.
1	Reserved			Reserved.
0	Excessive Pair Skew	RO	0b	Excessive pair skew in 1000BASE-T is detected by detecting that the scrambler has not acquired a 1000BASE-T link and cannot be brought up. In this case, the PHY usually falls back to 100BASE-TX or 10BASE-T. It is possible for other scrambler acquisition errors to be mistaken for excessive pair skew. 1b = Excessive pair skew (1000BASE-T only). 0b = No excessive pair skew (1000BASE-T only).

8.6 Port Control Registers (Page 769)

Table 33. Custom Mode Control PHY—Address 01, Page 769, Register 16

Name	Default	Bits	Description	Type
Reserved	0x04	15:11	Reserved	R/W
MDIO frequency access	0b	10	0b = normal MDIO frequency access 1b = reduced MDIO frequency access (required for read during cable disconnect)	R/W
Reserved	0x180	9:0	Reserved	R/W

Table 34. PHY—Address 01, Page 769, Register 17

Bits	Name	Type	Default	Description
15:11	Tx Gate Wait IFS	RW	00111b	Determines the size (in nibbles) of non-deferring window from CRS de-assertion.
10:8	BP extension Wait	R/W	000b	Additional waiting byte times after TX Gate Wait IPG expires until the <i>Back Pressure In-band</i> bit is cleared.
7	Reserved	R/W	0b	Reserved
6	Active_PD_enable	R/W	0b	Active Power Down Enable (sD3 Enable) When set to 1b, The MAC needs to enter integrated LAN controller power down mode.

**Table 34. PHY—Address 01, Page 769, Register 17**

5	Reserved	R/W	1b	Reserved
4	Host_WU_Active	R/W	0b	Enables host wake up from the I217. This bit is reset by power on reset only.
3	Wakeup clocks stop	R/W	1b	Wake-up clocks are stopped while wake up is disabled.
2	MACPD_enable	R/W	1b	Written as 1b when the integrated LAN controller needs to globally enable the MAC power down feature while the I217 supports WoL. When set to 1b, pages 800 and 801 are enabled for configuration and <i>Host_WU_Active</i> is not blocked for writes.
1:0	Reserved	RO	00b	Reserved

Register resets on Power Good only.

Table 35. DFT Control Register PHY Address 01, Page 769, Register 20

Name	Default	Bits	Description	Type
Reserved	0x140	15:7	Reserved	R/W
Gate PHY Reset on MAC Reset	0x0	6	Restricted/Reserved not for external documentation: When set to 1b integrated LAN controller reset does not automatically cause a PHY reset.	R/W
Reserved	0x0	5:0	Reserved	R/W

Table 36. Power Management Control Register PHY Address 01, Page 769, Register 21

Name	Default	Bits	Description	Type
Reserved	0x00	15:9	Reserved, write to 0x00	RO
Collision threshold	0x0F	8:1	Number of retries for a collided packet.	R/W
Retry late collision	0b	0	Retry late collision.	R/W

Table 37. SMBus Control Register PHY—Address 01, Page 769, Register 23

Name	Default	Bits	Description	Type
Reserved	0x400	15:5	Reserved	RO
Use LANWAKE#	0b	4	Use LANWAKE#, reset on power good	RW
Reserved	0b	3:2	Reserved	RW
dis_SMB_filtering	0b	1	When set, disables filtering of RX packets for the SMBus. In wake up mode, this configuration is ignored and the filters are enabled.	R/W
Force SMBus	0b	0	Force SMBus, reset on PCI reset de-assertion	RO

Table 38. Rate Adaptation Control Register PHY—Address 01, Page 769, Register 25

Name	Default	Bits	Description	Type
Reserved	0100010b	15:9	Reserved, write as read.	RWP



rx_en_rxdv_preamble	1	8	Enable generation of early preamble based on RX_DV in the receive path.	R/W
rx_en_crs_preamble	0	7	Enable generation of early preamble based on CRS in the receive path.	R/W
reserved	0	6	Reserved, write as read.	RWP
rx_flip_bad_sfd	1	5	Align the packet's start of frame delimiter to a byte boundary in the receive path.	R/W
read_delay_fd	10001b	4:0	Reserved, write as read.	RWP

Table 39. Flow Control Transmit Timer Value PHY—Address 01, Page 769, Register 27

Name	Default	Bits	Description	Type
Flow Control Transmit Timer Value	0x0000	15:0	The TTV field is inserted into a transmitted frame (either XOFF frames or any pause frame value in any software transmitted packets). It counts in units of slot time. If software needs to send an XON frame, it must set TTV to 0x0000 prior to initiating the pause frame.	RW

Note: Register resets on Power Good only.

Table 40. System Low Power Control—SxCTRL PHY Address 01, Page 769, Register 28

Attribute	Bit(s)	Initial Value	Description
RO	1:0	0b	Reserved
RW	2	0b	Enable ICMPv6 filtering to proxy controller
RW	3	0b	Enable Flow Control in SMBus
RW	4	0b	Enable new indicating for Flow Control
RO	11:5	1100000b	Reserved
RW	12	0b	Enable LPI enable to reset only on power good
RO	15:13	111b	Reserved

Table 41. Register resets on Power Good only. This register is used by software to access the SERDES MDI Control Register—SMDIC PHY Address 01, Page 774, Registers 23-24

Bits	Type	Reset	Description
15:0	RW	X	Data (DATA). In a Write command, software places the data bits and the MAC shifts them out to the LAN Connected Device. In a Read command, the MAC reads these bits serially from the LAN Connected Device and software can read them from this location.
20:16	RW	0	LAN Connected Device Register address (REGADD). i.e., Reg 0, 1, 2, ... 31.
25:21	RW	0	LAN Connected Device Address (PHYADD).
27:26	RW	0	Op-code (OP). 00 reserved 01 for MDI Write 10 for MDI Read. 11 reserved.



Table 41. SERDES MDI Control Register—SMDIC PHY Address 01, Page 774, Registers 23-24

28	RW	1	Ready bit (R) . Set to 1 by LAN Controller at the end of the MDI transaction (i.e., indicates a Read or Write has been completed). It should be reset to 0 by software at the same time the command is written.
30:29	RW	0	Start of Frame for the MDIO access (ST) 00b—MDIO access compliant to IEEE 802.3 clause 45 01b—MDIO access compliant to IEEE 802.3 clause 22 10b:11b—Reserved
31	RO	0	Reserved

SERDES registers in the LAN Connected Device.

8.7 Statistics Registers

Note: Each statistics register is constructed out of a pairs of two 16 bit registers. The lower 16 bits of the register are mapped to the higher numbered register and the higher 16 bits of the register are mapped to the lower numbered register.

Table 42. Single Collision Count—SCC PHY Address 01, Page 778, Registers 16—17

Bit	Type	Reset	Description
31:0	RO/V	0x00	SCC Number of times a transmit encountered a single collision.

This register counts the number of times that a successfully transmitted packet encountered a single collision. This register only increments if transmits are enabled and the I217 is in half-duplex mode.

Table 43. Excessive Collisions Count—ECOL PHY Address 01, Page 778, Register 18—19

Bit	Type	Reset	Description
31:0	RO/V	0x00	ECC Number of packets with more than 16 collisions.

When 16 or more collisions have occurred on a packet, this register increments, regardless of the value of collision threshold. If collision threshold is set below 16, this counter won't increment. This register only increments if transmits are enabled and the I217 is in half-duplex mode.

Table 44. Multiple Collision Count—MCC PHY Address 01, Page 778, Register 20—21

Bit	Type	Reset	Description
31:0	RO/V	0x00	MCC Number of times a successful transmit encountered multiple collisions.

This register counts the number of times that a transmit encountered more than one collision but less than 16. This register only increments if transmits are enabled and the I217 is in half-duplex mode.

Table 45. Late Collisions Count—LATECOL PHY Address 01, Page 778, Register 23—24

Bit	Type	Reset	Description
-----	------	-------	-------------



Table 45. Late Collisions Count—LATECOL PHY Address 01, Page 778, Register 23–24

31:0	RO/V	0x00	LCC Number of packets with late collisions.
------	------	------	--

Late collisions are collisions that occur after one slot time. This register only increments if transmits are enabled and the I217 is in half-duplex mode.

Table 46. Collision Count—COLC PHY Address 01, Page 778, Register 25–26

Bit	Type	Reset	Description
31:0	RO/V	0x00	COLC Total number of collisions experienced by the transmitter.

This register counts the total number of collisions seen by the transmitter. This register only increments if transmits are enabled and the I217 is in half-duplex mode. This register applies to clear as well as secure traffic.

Table 47. Defer Count—DC PHY Address 01, Page 778, Register 27–28

Bit	Type	Reset	Description
31:0	RO/V	0x00	CDC Number of defer events.

This register counts defer events. A defer event occurs when the transmitter cannot immediately send a packet due to the medium busy either because another device is transmitting, the IPG timer has not expired, half-duplex deferral events, reception of XOFF frames, or the link is not up. This register only increment if transmits are enabled. The counter does not increment for streaming transmits that are deferred due to TX IPG.

Table 48. Transmit with No CRS—TNCRS PHY Address 01, Page 778, Register 29–30

Bit	Type	Reset	Description
31:0	RO/V	0x00	TNCRS Number of transmissions without a CRS assertion from the I217.

This register counts the number of successful packet transmission in which the CRS input from the I217 was not asserted within one slot time of start of transmission from the integrated LAN controller. Start of transmission is defined as the assertion of TX_EN to the I217.

The I217 should assert CRS during every transmission. Failure to do so might indicate that the link has failed, or the I217 has an incorrect link configuration. This register only increments if transmits are enabled. This register is only valid when the I217 is operating at half duplex.



8.8 PCIe Registers

Table 49. PCIe FIFOs Control/Status PHY Address 01, Page 770, Register 16)

Name	Default	Bits	Description	Type
Reserved	0000001b	15:9	Reserved	RO
Rx FIFO overflow	0b	8	Rx FIFO overflow occurred.	RO/SC
Reserved	0b	7	Reserved	RO
Tx FIFO overflow	0b	6	Tx FIFO overflow occurred.	RO/SC
Reserved	000000b	5:0	Reserved	RO

Table 50. PCIe Power Management Control PHY Address 01, Page 770, Register 17

Name	Default	Bits	Description	Type
Reserved	1b	15	Reserved.	RO
K1 enable ¹	0b	14	Enable K1 Power Save Mode 1b = Enable. 0b = Disable.	RW
Giga_K1_disable	0b	13	When set, the I217 does not enter K1 while link speed is at 1000Mb/s.	RW
Reserved	100b	12:10	Reserved.	RO
Request a PCIe clock in K1	1b	9	Use CLK_REQ to request PCIe clock in K1	R/W
PLL stop in K1 giga	0b	8	Enables stopping the SERDES PLL in K1 state (in 1Gbps) 1b = Enable. 0b = Disable.	R/W
PLL stop in K1	1b	7	Enables stopping SerDes PLL in K1 state (in 10Mbps and 100Mbps). 1b = Disable. 0b = Enable.	R/W
Reserved	0b	6	Reserved	R/W
Reserved	0b	5	Reserved	R/W
Reserved ¹	0010b	4:1	Reserved	R/W
Enable Electrical Idle in Cable Disconnect	0b	0	Consider reserved (no entry to Electrical Idle due to Cable Disconnect)	R/W

Notes:

1. While in SMBus mode, this bit is cleared. To re-enable K1 after switching back to PCIe, this register needs to be re-configured.



Table 51. In-Band Control PHY Address 01, Page 770, Register 18¹

Name	Default	Bits	Description	Type
Reserved	0x0	15:14	Reserved	R/W
Link status transmit timeout	0x5	13:8	Link status retransmission period in tens of microseconds.	R/W
kum_pad_use_dis	0b	7	Disables 1000 Mb/s in-band messages during packets in 10/100 Mb/s mode.	R/W
Max retries	0x7	6:0	Maximum retries when not receiving an acknowledge to an in-band message.	R/W

1. All in-band time outs are multiplied by 1000 while in SMBus mode.

Table 52. PCIe Diagnostic PHY Address 01, Page 770, Register 20¹

Name	Default	Bits	Description	Type
Inband MDIO acknowledge timeout	0x55	15:8	Timeout in microseconds for receiving acknowledge for an inband MDIO message	R/W
Reserved	0x0	76	Reserved	RW
In-band status acknowledge timeout	0x04	5:0	Timeout in microseconds for receiving an acknowledge for an in-band status message.	R/W

1. All in-band timeouts are multiplied by 1000 while in SMBus mode.

Table 53. Timeouts PHY Address 01, Page 770, Register 21¹

Name	Default	Bits	Description	Type
Reserved	0000b	15:12	Reserved, write as read.	RWP
K1 exit timeout	010100b	11:6	These bits define how much time IDLE symbols are sent on the TX pair after exiting from K1 state before the I217 starts sending data to the integrated LAN controller (each bit represents 80ns).	R/W
Reserved	010100b	5:0	Reserved	R/W

1. All in-band time outs are multiplied by 1000 while in SMBus mode.

8.9 General Registers

Table 54. PCIe Kstate Minimum Duration Timeout PHY Address 01, Page 770, Register 23¹

Name	Default	Bits	Description	Type
Reserved	0x00	15:5	Reserved, write as read.	RWP
EI_min_dur timeout	0x10	4:0	These bits define the minimum time the I217 stays in electrical idle state once entered (each bit represents 80 ns).	R/W

1. All in-band time outs are multiplied by 1000 while in SMBus mode.

Table 55. Low Power Idle GPIO Control PHY Address 01, Page 772, Register 18

Name	Default	Bits	Description	Type
Reserved	0x40	15:10	Reserved	R/W
Auto EN LPI	0x0	11	Auto Enable LPI after link up. When set to 0x1 772.20[14:13] will be automatically set by HW after link up.	R/W

**Table 55. Low Power Idle GPIO Control PHY Address 01, Page 772, Register 18**

Name	Default	Bits	Description	Type
TX_LPI_GPIO0	0x0	10	Route Tx LPI indication to GPIO 0	R/W
Reserved	0xC9	9:0	Reserved	R/W

Note: Register resets on Power Good only.

Table 56. Low Power Idle Control PHY Address 01, Page 772, Register 20

Name	Default	Bits	Description	Type
Reserved	0x0	15	Reserved	R/W
1000Enable	0x0	14	1000Enable—Enable EEE on 1 Gb/s link speed. This bit auto clears on link down.	R/W
100Enable	0x0	13	100Enable—Enable EEE on 100 Mb/s link speed. This bit auto clears on link down.	R/W
ForceLPI	0x0	12	Force LPI entry—When set to 1b by software the PHY enters LPI mode even when not in K1.	R/W
PostLPICount	0x1	11:9	Post LPI Counter—When in LPI active and an Ethernet packet of an inband XOFF message is received from the LC count: 000b = 8 μ s 001b = 16 μ s 010b = 24 μ s 011b = 32 μ s 100b = 40 μ s 101b = 48 μ s 110b = 56 μ s 111b = 64 μ s Before transmitting XOFF or a valid packet.	R/W
PLLLockCnt	0x2	6:4	PLL Lock Counter—when LPI <i>100Enable</i> or <i>1000Enable</i> is asserted, this counter forces the PLL Lock count to be in the range of 10.02 μ s to 81.88 μ s in steps of 10.02 μ s.	
Reserved	0x0	3:0	Reserved	RO

Table 57. Flow Control Refresh Threshold Value FCRTV PHY Address 01, Page 772, Register 23

Name	Default	Bits	Description	Type
FCRTV	0x0000	15:0	Flow Control Refresh Threshold Value. This value indicates the threshold value of the flow control shadow counter. When the counter reaches this value, and the transmit and receive buffer fullness is still above the low threshold value), a pause (XOFF) frame is sent to the link partner. The FCRTV timer count interval counts at slot times of 64 byte times. If this field is set to zero, Flow Control Refresh feature is disabled.	RW

Note: Register resets on Power Good only.

Table 58. Flow Control Thresholds FCTH PHY Address 01, Page 772, Register 24

Name	Default	Bits	Description	Type
HTHRSH	0x00	7:0	High threshold for sending XOFF (in units of 32 bytes)	RW



Table 58. Flow Control Thresholds FCTH PHY Address 01, Page 772, Register 24

LTHRSR	0x00	15:8	Low threshold for sending XON (in units of 32 bytes)	RW
--------	------	------	--	----

Note: Register resets on Power Good only.

Table 59. LANWAKE# Control LANWAKEC PHY Address 01, Page 772, Register 25

Name	Default	Bits	Description	Type
Rx FIFO empty threshold	0xA	3:0	For flow control in SMB	RW
Store Host WoL packet	0x1	4	Store the Host WoL packet in the Proxy RAM	RW
Use LANWAKE pin	0x1	5	Always use LANWAKE pin to indicate WoL	RW
Clear LANWAKE pin	0x0	6	When set to 1 will clear the LANWAKE pin, this bit is auto cleared to 0	RW
Use Legacy wake	0x0	7	When set to 1 use I217 WoL flows	RW
Reserved	0x0	8	Reserved	RW
Filter on 1Gbps	0x0	9	Filter packets in 1G for jumbo fix	RW
Crop ANM DA Jumbo	0x0	10	Crop ANM DA in Jumbo fix	RW
Crop BCST DA in Jumbo	0x0	11	Crop BCST DA in Jumbo fix	RW
Enable WoL if no packet capture	0x1	12	When set to 1 WoL is enabled even if the WoL packet is not captured, when set to 0 no capture means no WoL	RW
Reserved	0x6	15:13	Reserved	RW

Note: Register resets on Power Good only.

Table 60. Memories Power PHY Address 01, Page 772, Register 26

Name	Default	Bits	Description	Type
Reserved	0x1	15:13	Reserved	R/W
MOEM	0b	12	Mask OEM bits/Gig Disable/restart AN bits impact	R/W
Reserved	0x7F	11:5	Reserved	R/W
DISSMABREKONLCD	1b	4	Disable the SMB release on LCD reset	R/W
Reserved	0xF	3:0	Reserved	R/W

Table 61. Configuration Register PHY Address 01, Page 772, Register 20

Name	Default	Bits	Description	Type
Reserved	0x3FC0	15:2	Reserved	R/W
ENMTAONPWRGD	0b	1	Enable MTA to reset only on power good	R/W
Reserved	0b	0	Reserved	R/W

Table 62. Off Board LAN Connected Device Control PHY Address 01, Page 779, Register 20

Name	Default	Bits	Description	Type
OBLCD_En	0b	15	OBLCD_En—Enables the beacon feature for OBLCD connection indication.	R/W
	0x0	14:12	Reserved	R/W



Table 62. Off Board LAN Connected Device Control PHY Address 01, Page 779, Register 20

Name	Default	Bits	Description	Type
REF_DLY	0x1	11:8	REF_DLY—Reference clock delay, defines the time K1 exit is delayed, waiting for reference clock. Granularity in 1uSec	R/W
BCN_INTER	0x4	7:4	BCN_INTER—Beacon Interval, defines the time interval between beacons. Granularity in 50mSec 0x0 is not a valid setting	R/W
BCN_DUR	0x2	3:0	BCN_DUR—Beacon Duration, defines the time of a single beacon cycle. Granularity in 0.5uSec 0x0 is not a valid setting	R/W

Energy Detect Mode

Table 63. Intel® Ethernet Controller I217 Capability PHY Address 01, Page 776, Register 19

Name	Default	Bits	Description	Type
Reserved	000000b	15:10	Reserved for future capabilities.	RO
Intel® Active Management Technology (Intel® AMT) and Circuit Breaker	0b	9	Intel® AMT and Circuit Breaker Enables Intel® AMT and circuit breaker capability.	RO
802.1Q & 802.1p	0b	8	802.1Q & 802.1p Enables support for VLAN per 802.1Q & 802.1p.	RO
Receive Side Scaling	0b	7	Receive Side Scaling (RSS) Enables RSS.	RO
2 Tx and 2 Rx Queues	0b	6	Two Tx and 2 Rx Queues When set, enables dual transmit and dual receive queues. When cleared, a single receive and a single transmit queue are enabled.	RO
Energy Detect	0b	5	Energy Detect Enables energy detect capability.	RO
AC/DC Auto Link Speed Connect	0b	4	AC/DC Auto Link Speed Connect Enables different power management policy in AC and battery modes.	RO
Reserved	0b	3	Reserved	RO
ASF	0b	2	Alert Standard Format (ASF) Enables ASF support	RO
WfM	0b	1	Wired-for-Manageability (WfM) Enables WfM, including ACPI, WoL, and PXE.	RO
Ability to initiate a team	0b	0	Ability to initiate a team; enables teaming capability.	RO

Table 64. OEM Bits PHY Address 01, Page 0, Register 25

Bits	Field	Mode	HW Rst	Description
15:11	Reserved	R/W	00000b	
10	Aneg_now	R/W	0b	Restart auto-negotiation. This bit is self clearing.
9:7	Reserved	R/W	000b	
6	a1000_dis	R/W	0b'	When set to 1b, 1000 Mb/s speed is disabled.



Table 64. OEM Bits PHY Address 01, Page 0, Register 25

Bits	Field	Mode	HW Rst	Description
5:3	Reserved	R/W	000b	
2	rev_aneg	R/W	0b	Low Power Link Up Mechanism. Enables a link to come up at the lowest possible speed in cases where power is more important than performance.
1:0	Reserved	R/W	00b	

Notes:

1.

Note: 0b is the default value after power on reset. When PE_RST_N goes low (switches to SMBus), its value becomes 1b.

Table 65. SMBus Address PHY Address 01, Page 0, Register 26¹

Name	Default	Bits	Description	Type
Reserved	0x00	15:12	Reserved	RO
SMBus Frequency High	0b	12	SMBus frequency high bit, see bit description in SMBus frequency low	RW
SMB fragments size	0b	11	Select SMBus Fragments Size When set to 1b, the fragment size is 64 bytes, otherwise 32 bytes.	RW
APM Enable	0b	10	APM WoL enable.	RW
PEC Enable	1b	9	Defines if the I217 supports PEC on the SMBus.	RW
SMBus Frequency Low	0b	8	Together with SMBus Frequency High defines the SMBus frequency: High, Low-Frequency 0,0—100Khz 0,1—400Khz 1,0—1000Khz 1,1—Reserved	RW
SMBus Address Valid	0b	7	0b = Address not valid. 1b = SMBus address valid. This bit is written by the integrated LAN controller when the SMBus Address field is updated. The 217 cannot send SMBus transactions to the integrated LAN controller unless this bit is set.	RW
SMBus Address	0x00	6:0	This is the integrated LAN controller SMBus address. The 217 uses it for master functionality.	RW

1. This register is reset only on internal power on reset.

Table 66. Shadow Receive Address Low0—SRAL0 PHY Address 01, Page 0, Registers

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	Receive Address Low (RAL) The lower 32 bits of the 48-bit Ethernet address n (n=0, 1...6). RAL 0 is loaded from words 0x0 and 0x1 in the NVM.

**Table 67. Shadow Receive Address High0—RAH0 PHY Address 01, Page 0, Register 29**

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address n (n=0, 1...6). RAH 0 is loaded from word 0x2 in the NVM.

Table 68. LED Configuration PHY Address 01, Page 0, Register 30

Name	Default	Bits	Description	Type
Blink rate	0b	15	Specifies the blink mode of the LEDs. 0b = Blinks at 200 ms on and 200 ms off. 1b = Blinks at 83 ms on and 83 ms off.	RW
LED2 Blink	0b	14	LED2_BLINK Field 0b = No blinking. 1b = Blinking.	RW
LED2 Invert	0b	13	LED2_IVRT Field 0b = Active low output. 1b = Active high output.	RW
LED2 Mode	110b	12:10	Mode specifying what event/state/pattern is displayed on LED2.	RW
LED1 Blink	0b	9	LED1_BLINK Field 0b = No blinking. 1b = Blinking.	RW
LED1 Invert	0b	8	LED1_IVRT Field 0b = Active low output. 1b = Active high output.	RW
LED1 Mode	111b	7:5	Mode specifying what event/state/pattern is displayed on LED1.	RW
LED0 Blink	1b	4	LED0_BLINK Field 0b = No blinking. 1b = Blinking.	RW
LED0 Invert	0b	3	LED0_IVRT Field 0b = Active low output. 1b = Active high output.	RW
LED0 Mode	100b	2:0	Mode specifying what event/state/pattern is displayed on LED0.	RW

NOTES:

- When LED Blink mode is enabled the appropriate Led Invert bit should be set to zero.
- The dynamic LED's modes (LINK/ACTIVITY and ACTIVITY) should be used with LED Blink mode enabled.

Table 69. LED Modes

Mode	Selected Mode	Source Indication
000	Link 10/1000	Asserted when either 10 or 1000Mbps link is established and maintained
001	Link 100/1000	Asserted when either 100 or 1000Mbps link is established and maintained
010	Link Up	Asserted when any speed link is established and maintained.
011	Activity	Asserted when link is established and packets are being transmitted or received
100	Link/Activity	Asserted when link is established AND when there is NO transmit or receive activity



Table 69. LED Modes

Mode	Selected Mode	Source Indication
101	Link 10	Asserted when a 10Mbps link is established and maintained.
110	Link 100	Asserted when a 100Mbps link is established and maintained
111	Link 1000	Asserted when a 1000Mbps link is established and maintained

8.9.1 Interrupts

The I217 maintains status bits (per interrupt cause) to reflect the source of the interrupt request. System software is expected to clear these status bits once the interrupt is being handled.

8.10 Wake Up Registers

8.10.1 Accessing Wake Up Registers Using MDIC

When software needs to configure the wake up state (either read or write to these registers) the MDIO page should be set to 800 (for host accesses) until the page is not changed to a different value wake up register access is enabled. After the page was set to the wake up page, the address field is no longer translated as *reg_addr* (register address) but as an instruction. If the given address is in the [0..15] range, meaning PHY registers, the functionality remains unchanged. There are two valid instructions:

1. Address Set—0x11—Wake up space address is set for either reading or writing.
2. Data cycle—0x12—Wake up space accesses read or write cycle.

For the I217 the wake area read cycle sequence of events is as follows:

1. Setting page 800 The software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b
 - b. Op-Code = 01b (write)
 - c. PHYADD = The I217's address from the MDI register
 - d. REGADD = Page setting
 - e. DATA = 800 (wake up page)
2. Address setting; the software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b
 - b. Op-Code = 01b (write)
 - c. PHYADD = The I217's address from the MDI register
 - d. REGADD = 0x11 (address set)
 - e. DATA = XXXX (address of the register to be read)
3. Reading a register; the software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b



- b. Op-Code = 10b (read)
- c. PHYADD = The 217's address from the MDI register
- d. REGADD = 0x12 (data cycle for read)
- e. DATA = YYYY (data is valid when the ready bit is set)

For the I217, the wake area write cycle sequence of events is as follows:

1. Setting page 800; the software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b
 - b. Op-Code = 01b (write)
 - c. PHYADD = The 217's address from the MDI register
 - d. REGADD = Page setting
 - e. DATA = 800 (wake up page)
2. Address setting; The software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b
 - b. Op-Code = 01b (write)
 - c. PHYADD = The 217's address from the MDI register
 - d. REGADD = 0x11 (address set)
 - e. DATA = XXXX (address of the register to be read)
3. Writing a register; the software device driver performs a write cycle to the MDI register with:
 - a. Ready = 0b
 - b. Op-Code = 01b (write)
 - c. PHYADD = The 217's address from the MDI register
 - d. REGADD = 0x12 (data cycle for write)
 - e. DATA = YYYY (data to be written to the register)

8.10.2 Host Wake Up Control Status Register Description

Note: All Wake-Up registers (Pg. 800-801) are not cleared when PHY reset is asserted. These registers are only cleared when internal power on reset is de-asserted or when cleared by the software device driver.

Table 70. Receive Control—RCTL PHY Address 01, Page 800, Register 0

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	Unicast Promiscuous Enable (UPE) 0b = Disabled. 1b = Enabled.
RW	1	0b	Multicast Promiscuous Enable (MPE) 0b = Disabled. 1b = Enabled.

Table 70. Receive Control—RCTL PHY Address 01, Page 800, Register 0

RW	2	1b	Slave Access Enable 0b = Access disabled, the filters are active. 1b = Access enabled, the filters are not active.
RW	4:3	00b	Multicast Offset (MO) This determines which bits of the incoming multicast address are used in looking up the bit vector. 00b = [47:38]. 01b = [46:37]. 10b = [45:36]. 11b = [43:34].
RW	5	0b	Broadcast Accept Mode (BAM) 0b = Ignore broadcast (unless it matches through exact or imperfect filters) 1b = Accept broadcast packets.
RW	6	0b	Pass MAC Control Frames. PMCF 0b = Do not (specially) pass iMAC control frames. 1b = Pass any iMAC control frame (type field value of 0x8808).
RW	7	0b	Receive Flow Control Enable (RFCE) Indicates that the I217 responds to the reception of flow control packets. If auto-negotiation is enabled, this bit is set to the negotiated duplex value.
RW	8	0b	Reserved
Reserved.	15:9	0x00	Reserved

PMCF controls the usage of MAC control frames (including flow control). A MAC control frame in this context must be addressed to the flow control multicast address 0x0100_00C2_8001 and match the type field (0x8808). If PMCF=1b, then frames meeting this criteria participate in wake up filtering.

Table 71. Wake Up Control—WUC PHY Address 01, Page 800, Register 1

Attribute	Bit(s)	Initial Value	Description
RW/SN	0	0b	Advance Power Management Enable (APME) If set to 1b, APM wake up is enabled.
RW/V	1	0b	PME_En If set to 1b, ACPI wake up is enabled.
RWC	2	0b	PME_Status This bit is set when the I217 receives a wake up event. This bit is cleared by writing 1b to clear or by clearing the Host_WU_Active/ME_WU_Active bits.
RW	3	0b	Link Status change on Energy detect. When this bit is set enable wake in energy on the lines (instead of actual link status change). By default this bit is in 0b.
RW/SN	4	0b	Link Status Change Wake Enable (LSCWE) Enables wake on link status change as part of APM wake capabilities.
RW/SN	5	0b	Link Status Change Wake Override (LSCWO) If set to 1b, wake on link status change does not depend on the LNKC bit in the WUFC register. Instead, it is determined by the APM settings in the WUC register.
RO	13:6	0x00	Reserved
RW	14	0b	FLX6 Flexible filter 6enable.
RW	15	0b	FLX7 Flexible filter 7enable.


Table 72. Wake Up Filter Control—WUFC PHY Address 01, Page 800, Register 2

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	LNKC Link status change wake up enable.
RW	1	0b	MAG Magic packet wake up enable.
RW	2	0b	EX Directed exact wake up enable.
RW	3	0b	MC Directed multicast wake up enable.
RW	4	0b	BC Broadcast wake up enable.
RW	5	0b	ARP ARP/IPv4 request packet wake up enable. IPv4 filtering applies only to the 3 host IPv4 addresses in IP4AT
RW	6	0b	IPV4 Directed IPv4 packet wake up enable. IPv4 filtering applies only to the 3 host IPv4 addresses in IP4AT
RW	7	0b	IPV6 Directed IPv6 packet wake up enable.
RO	8	0b	Reserved.
RW	9	0	FLX4 Flexible filter 4 enable.
RW	10	0b	FLX5 Flexible filter 5 enable.
RW	11	0b	NoTCO Ignore TCO packets for host wake up. If the <i>NoTCO</i> bit is set, then any packet that passes the manageability packet filtering does not cause a host wake up event even if it passes one of the host wake up filters.
RW	12	0b	FLX0 Flexible filter 0 enable
RW	13	0b	FLX1 Flexible filter 1 enable
RW	14	0b	FLX2 Flexible filter 2 enable
RW	15	0b	FLX3 Flexible filter 3 enable

This register is used to enable each of the pre-defined and flexible filters for wake up support. A value of 1b means the filter is turned on. A value of 0b means the filter is turned off.

Table 73. Wake Up Status—WUS PHY Address 01, Page 800, Register 3

Attribute	Bit(s)	Initial Value	Description
RWC	0	0b	LNKC Link status changed
RWC	1	0b	MAG Magic packet received
RWC	2	0b	EX Directed exact packet received. The packet's address matched one of the 7 pre-programmed exact values in the Receive Address registers.
RWC	3	0b	MC Directed multicast packet received. The packet was a multicast packet that was hashed to a value that corresponded to a 1-bit in the multicast table array.
RWC	4	0b	BC Broadcast packet received.
RWC	5	0b	ARP ARP/IPv4 request packet received.

Table 73. Wake Up Status—WUS PHY Address 01, Page 800, Register 3

RWC	6	0b	IPV4 Directed IPv4 packet received.
RWC	7	0b	IPV6 Directed IPv6 packet received.
RWC	8	0b	FLX4 Flexible filter 4 match.
RWC	9	0b	FLX5 Flexible filter 5 match.
RWC	10	0b	FLX6 Flexible filter 6 match.
RWC	11	0b	FLX7 Flexible filter 7 match.
RWC	12	0b	FLX0 Flexible filter 0 match.
RWC	13	0b	FLX1 Flexible filter 1 match.
RWC	14	0b	FLX2 Flexible filter 2 match.
RWC	15	0b	FLX3 Flexible filter 3 match.

This register is used to record statistics about all wake up packets received. Note that packets that match multiple criteria might set multiple bits. Writing a 1b to any bit clears that bit.

This register is not cleared when PHY reset is asserted. It is only cleared when internal power on reset is de-asserted or when cleared by the software device driver.

Table 74. Receive Address Low—RAL PHY Address 01, Page 800, Registers 16-17

Attribute	Bit(s)	Initial Value	Description
RW	31:0	0	Receive Address Low (RAL) The lower 32 bits of the 48-bit Ethernet address

Table 75. Receive Address High—RAH PHY Address 01, Page 800, Registers 18-19

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address
Reserved.	17:16	0x0	Reserved.
RO	30:18	0x00	Reserved, reads as 0b and ignored on writes.
RW	31	0	Address valid (AV) When this bit is set, the relevant RAL and RAH are valid (compared against the incoming packet).

Table 76. Shared Receive Address Low—SHRAL PHY Address 01, Page 800, Registers 20-21 + 4*n (n=0...10)

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	Receive Address Low (RAL) The lower 32 bits of the 48-bit Ethernet address n (n=0...10).

Table 77. Shared Receive Address High—SHRAH PHY Address 01, Page 800, Registers 22-23 + 4*n (n=0...8,10)

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address n (n=0...8,10).



Table 77. Shared Receive Address High—SHRAH PHY Address 01, Page 800, Registers 22-23 + 4*n (n=0...8,10)

Reserved.	17:16	0x0	Reserved.
RO	30:18	0x00	Reserved, reads as 0b and is ignored on writes.
RW	31	0b	Address valid (AV) When this bit is set, the relevant RAL and RAH are valid (compared against the incoming packet).

Table 78. Shared Receive Address High 9—SHRAH[9] PHY Address 01, Page 800, Registers 58-59

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address 9
Reserved.	17:16	0x0	Reserved.
RO	29:18	0x00	Reserved, reads as 0x00 and is ignored on writes.
RW	30	0b	All Nodes Multicast Address valid (MAV) The all nodes multicast address (33:33:00:00:00:01) is valid when this bit is set. Note that 0x33 is the first byte on the wire.
RW	31	0b	Address valid (AV) When this bit is set, the relevant address 3 is valid (compared against the incoming packet).

Table 79. IP Address Valid—IPAV1 PHY Address 01, Page 800, Register 64

Attribute	Bit(s)	Initial Value	Description
RO	0	0b	Reserved
RW	1	0b	IP4AT1 address valid IPv4 address 1 valid.
RW	2	0b	IP4AT1 address valid IPv4 address 2 valid.
RW	3	0b	IP4AT1 address valid IPv4 address 3 valid.
RW	4	0b	IP6AT3 address valid.
RW	5	0b	IP6AT2 address valid.
RW	6	0b	IP6AT1 address valid.
RW	7	0b	IP6AT0 address valid (Duplicate of bit 15).
RW	8	0b	Enable L2 for IPv6 multicast according to IP6AT0
RW	9	0b	Enable L2 for IPv6 multicast according to IP6AT1
RW	10	0b	Enable L2 for IPv6 multicast according to IP6AT2
RW	11	0b	Enable L2 for IPv6 multicast according to IP6AT3
RO	12:14	0x00	Reserved
RW	15	0b	V60 IPv6 address valid.

Table 80. IPv4 Address Table—IP4AT1 PHY Address 01, Page 800, Registers 82-83 + 2*n (n=0, 1, 2)

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	IPADD IP address n (n = 0, 1, 2).

1. The IPv4 address table is used to store the three IPv4 addresses for ARP/IPv4 request packets and directed



IPv4 packet wake ups.

Table 81. Proxy Control—PRXC PHY Address 01, Page 800, Register 70

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	Proxy_mode. Should be set in the end of Proxy configuration. That's the last MDIO access to Phy, unless we put arbitration on MDIO accesses between Proxy uCtl and regular MDIO accesses through MAC). As long as this bit is clear, Proxy logic is under reset
RW	1	0b	Code_loaded. Set either by SW after loading the uCode through MDIO, or by HW after reception of the Code packet from MAC (the first packet after setting PRXC. nxt_pkt_is_code)
RW	2	0b	Code_from_MDIO. Should be set by SW before it starts loading the instruction code memory through MDIO accesses to page 802
RW	3	0b	Nxt_pkt_is_code. Should be set by SW before it transmits uCode packet
RW	4	0b	Auto disable proxying after link-down deactivation period
RW	5	0b	ARP PROxy Enable
RW	6	0b	ND Proxy Enable
Reserved.	7	0b	Reserved.
RW	13:8	0b	Link Down deactivation period in 1 seconds granularity
RW	14	0b	Reserved.
RW	15	0b	Ready_for_Code. this bit is set by HW post setting of Code_from_MDIO or Nxt_pkt_is_code indicating to SW Intel I217 is ready for the code load.

Table 82. Proxy Code Checksum—PRCC PHY Address 01, Page 800, Register 71

Attribute	Bit(s)	Initial Value	Description
RO	15:0	0000h	Proxy Code Checksum. This register holds the checksum calculation for the proxy code loaded to the transmit FIFO as a packet starting from the start of frame till the end of frame

Table 83. Proxy Control 2—PRXC2 PHY Address 01, Page 800, Register 72

Attribute	Bit(s)	Initial Value	Description
Reserved.	13:0	0b	Reserved
RW	14	0b	MLD Proxy Enable
Reserved.	15	0b	Reserved

Table 84. Flex Filters Proxy Control—FFPRXC PHY Address 01, Page 800, Register 75

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	Route Flex filter 0 to the proxy uController
RW	1	0b	Route Flex filter 1 to the proxy uController
RW	2	0b	Route Flex filter 2 to the proxy uController
RW	3	0b	Route Flex filter 3 to the proxy uController
RW	4	0b	Route Flex filter 4 to the proxy uController
RW	5	0b	Route Flex filter 5 to the proxy uController
RW	6	0b	Route Flex filter 6 to the proxy uController



RW	7	0b	Route Flex filter 7 to the proxy uController
RWC	15:8	0x00	Flex filter match status An incoming packet matching one of the flex filters will set a bit in this status, the bits are cleared on write of 1 Bit 8 Flex filter 0 match Bit 9 Flex filter 1 match Bit 10 Flex filter 2 match Bit 11 Flex filter 3 match Bit 12 Flex filter 4 match Bit 13 Flex filter 5 match Bit 14 Flex filter 6 match Bit 15 Flex filter 7 match

Table 85. Wake Up Filter Control 2—WUS2 PHY Address 01, Page 800, Register 76

Attribute	Bit(s)	Initial Value	Description
RW	0	0	FLX8, Flexible filter 8 enable.
RW	1	0b	FLX9, Flexible filter 9 enable.
RW	2	0b	FLX10, Flexible filter 10 enable
RW	3	0b	FLX11, Flexible filter 11 enable
RW	4	0b	FLX12, Flexible filter 12 enable
RW	5	0b	FLX13, Flexible filter 13 enable
RW	6	0b	FLX14, Flexible filter 14 enable
RW	7	0b	FLX15, Flexible filter 15 enable
RO	15:8	0x0	Reserved

Table 86. Wake Up Filter Status 2—WUS2 PHY Address 01, Page 800, Register 77

Attribute	Bit(s)	Initial Value	Description
RW	0	0	FLX8, Flexible filter 8 matched.
RW	1	0b	FLX9, Flexible filter 9 matched.
RW	2	0b	FLX10, Flexible filter 10 matched
RW	3	0b	FLX11, Flexible filter 11 matched
RW	4	0b	FLX12, Flexible filter 12 matched
RW	5	0b	FLX13, Flexible filter 13 matched
RW	6	0b	FLX14, Flexible filter 14 matched
RW	7	0b	FLX15, Flexible filter 15 matched
RO	15:8	0x0	Reserved

**Table 87. IPv6 Address Table—IP6AT[3:0]
PHY Address 01, Page 800, Registers 88-89 + 2*n (n=0...3)
PHY Address 01, Page 800, Registers 96-97 + 2*n (n=0...3)
PHY Address 01, Page 800, Registers 104-105 + 2*n (n=0...3)
PHY Address 01, Page 800, Registers 112-113 + 2*n (n=0...3)**

Attribute	Bit(s)	Initial Value	Description
RW	31:0	0x0	IPv6 Address IPv6 address bytes n*4...n*4+3 (n=0, 1, 2, 3) while byte 0 is first on the wire and byte 15 is last.



The IPv6 address table is used to store the IPv6 addresses for directed IPv6 packet wake up (only using the first IPv6 address) and network proxy filtering.

Configuration example for IPv6 address: fe80:0:0:0:200:1ff:fe30:100

- 01.800.88 - 0x80fe
- 01.800.89 - 0x0000
- 01.800.90 - 0x0000
- 01.800.91 - 0x0000
- 01.800.92 - 0x0002
- 01.800.93 - 0xff01
- 01.800.94 - 0x30fe
- 01.800.95 - 0x0001

Table 88. Multicast Table Array—MTA[31:0] PHY Address 01, Page 800, Registers 128-191

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	Bit Vector. Word-wide bit vector specifying 32 bits in the multicast address filter table.

There is one register per 32 bits of the multicast address table for a total of 32 registers (thus the MTA[31:0] designation). The size of the word array depends on the number of bits implemented in the multicast address table. Software must mask to the desired bit on reads and supply a 32-bit word on writes.

Note: All accesses to this table must be 32-bit.

Figure 8-1 shows the multicast lookup algorithm. The destination address shown represents the internally stored ordering of the received destination address. Note that Byte 1 bit 0 shown in Figure 8-1 is the first on the wire. The bits that are directed to the multicast table array in this diagram match a multicast offset in the CTRL register equals 00b. The complete multicast offset options are:

Multicast Offset	Bits Directed to the Multicast Table Array
00b	DA[47:38] = Byte 6 bits 7:0, Byte 5 bits 1:0
01b	DA[46:37] = Byte 6 bits 6:0, Byte 5 bits 2:0
10b	DA[45:36] = Byte 6 bits 5:0, Byte 5 bits 3:0
11b	DA[43:34] = Byte 6 bits 3:0, Byte 5 bits 5:0

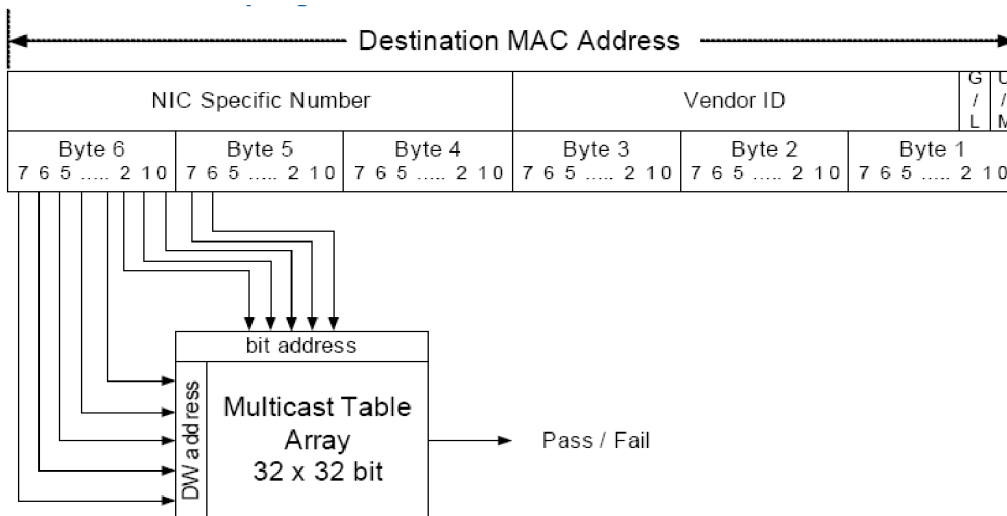


Figure 8-1. Multicast Table Array Algorithm

Table 89. Flexible Filter Value Table LSB—FFVT_01 PHY Address 01, Page 800, Registers 256 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value 0 Value of filter 0 byte n (n=0, 1... 127).
RW	15:8	X	Value 1 Value of filter 1 byte n (n=0, 1... 127).

There are 128 filter values. The flexible filter value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is one, the flexible filter compares the incoming data byte to the values stored in this table.

In the I217, since each address contains 16 bits, only the least significant bytes are stored in those addresses.

Table 90. Flexible Filter Value Table MSBs—FFVT_23 PHY Address 01, Page 800, Registers 257 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value 2 Value of filter 2 byte n (n=0, 1... 127).
RW	15:8	X	Value 3 Value of filter 3 byte n (n=0, 1... 127).

There are 128 filter values. The flexible filter value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is one, the flexible filter then compares the incoming data byte to the values stored in this table.

In the I217, since each address contains 16 bits, only the most significant bytes are stored in those addresses.



Before writing to the flexible filter value table, the software device driver must first disable the flexible filters by writing zeros to the *Flexible Filter Enable* bits of the WUFC register (WUFC.FLXn).

Table 91. Flexible Filter Value Table—FFVT_45 PHY Address 01, Page 800, Registers 512 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value 4 Value of filter 4 byte n (n=0, 1... 127).
RW	15:8	X	Value 5 Value of filter 5 byte n (n=0, 1... 127).

Table 92. Flexible Filter Value Table—FFVT_67 PHY Address 01, Page 800, Registers 1024 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value 6 Value of filter 6 byte n (n=0, 1... 127).
RW	15:8	X	Value 7 Value of filter 7 byte n (n=0, 1... 127).

Table 93. Flexible Filter Mask Table—FFMT PHY Address 01, Page 800, Registers 768 + n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	0	X	Mask 0 Mask for filter 0 byte n (n=0, 1... 127).
RW	1	X	Mask 1 Mask for filter 1 byte n (n=0, 1... 127).
RW	2	X	Mask 2 Mask for filter 2 byte n (n=0, 1... 127).
RW	3	X	Mask 3 Mask for filter 3 byte n (n=0, 1... 127).
RW	4	X	Mask 4 Mask for filter 4 byte n (n=0, 1... 127).
RW	5	X	Mask 5 Mask for filter 5 byte n (n=0, 1... 127).
RW	6	X	Mask 6 Mask for filter 6 byte n (n=0, 1... 127).
RW	7	X	Mask 7 Mask for filter 7 byte n (n=0, 1... 127).
RO	15:8	X	Reserved.

There are 128 mask entries. The flexible filter mask and table is used to store the four 1-bit masks for each of the first 128 data bytes in a packet, one for each flexible filter. If the mask bit is one, the corresponding flexible filter compares the incoming data byte at the index of the mask bit to the data byte stored in the flexible filter value table.

Note: Before writing to the flexible filter mask table the software device driver must first disable the flexible filters by writing zeros to the *Flexible Filter Enable* bits of the WUFC register (WUFC.FLXn).



Table 94. Flexible Filter Length Table—FFLT03 PHY Address 01, Page 800, Registers 896 + n (n=0...3)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1... 3).
RO	15:11	X	Reserved.

All reserved fields read as zeros and are ignored on writes.

There are eight flexible filters lengths covered by FFLT03, FFLT45, FFLT67 registers. The flexible filter length table stores the minimum packet lengths required to pass each of the flexible filters. Any packets that are shorter than the programmed length won't pass that filter. Each flexible filter considers a packet that doesn't have any mismatches up to that point to have passed the flexible filter when it reaches the required length. It does not check any bytes past that point.

Note: Before writing to the flexible filter length table the software device driver must first disable the flexible filters by writing zeros to the *Flexible Filter Enable* bits of the WUFC register (WUFC.FLXn).

Table 95. Flexible Filter Length Table—FFLT45 PHY Address 01, Page 800, Registers 904 + n (n=0...1)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1).
RO	15:11	X	Reserved.

Table 96. Flexible Filter Length Table—FFLT67 PHY Address 01, Page 800, Registers 908 + n (n=0...1)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1).
RO	15:11	X	Reserved.

Table 97. Flexible Filter Value Table 89—FFVT_89 PHY Address 01, Page 800, Registers 2304 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value of filter 8 byte n (n=0, 1... 127).
RW	15:8	X	Value of filter 9 byte n (n=0, 1... 127).

There are 128 filter values. The flexible filter value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is one, then the flexible filter compares the incoming data byte to the values stored in this table.



Table 98. Flexible Filter Value Table 1011—EFFVT_1011 PHY Address 01, Page 800, Registers 2305 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value of filter 10 byte n (n=0, 1... 127).
RW	15:8	X	Value of filter 11 byte n (n=0, 1... 127).

There are 128 filter values. The flexible filter value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is one, then the flexible filter compares the incoming data byte to the values stored in this table.

Table 99. Flexible Filter Value Table 1213—FFVT_1213 PHY Address 01, Page 800, Registers 2560 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value of filter 12 byte n (n=0, 1... 127).
RW	15:8	X	Value of filter 13 byte n (n=0, 1... 127).

Table 100. Flexible Filter Value Table 1415—FFVT_1415 PHY Address 01, Page 800, Registers 3072 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	Value of filter 14 byte n (n=0, 1... 127).
RW	15:8	X	Value of filter 15 byte n (n=0, 1... 127).

Table 101. Flexible Filter Mask Table 2—FFMT2 PHY Address 01, Page 800, Registers 2816 + n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	0	X	Mask for filter 8 byte n (n=0, 1... 127).
RW	1	X	Mask for filter 9 byte n (n=0, 1... 127).
RW	2	X	Mask for filter 10 byte n (n=0, 1... 127).
RW	3	X	Mask for filter 11 byte n (n=0, 1... 127).
RW	4	X	Mask for filter 12 byte n (n=0, 1... 127).
RW	5	X	Mask for filter 13 byte n (n=0, 1... 127).
RW	6	X	Mask for filter 14 byte n (n=0, 1... 127).
RW	7	X	Mask for filter 15 byte n (n=0, 1... 127).
RO	15:8	X	Reserved.

There are 128 mask entries. The flexible filter mask and table is used to store the four 1-bit masks for each of the first 128 data bytes in a packet, one for each flexible filter. If the mask bit is one, the corresponding flexible filter compares the incoming data byte at the index of the mask bit to the data byte stored in the flexible filter value table.

Table 102. Flexible Filter Length Table 891011—FFLT891011 PHY Address 01, Page 800, Registers 2944 + n (n=0...3)

Attribute	Bit(s)	Initial Value	Description
-----------	--------	---------------	-------------



Table 102. Flexible Filter Length Table 891011—FFLT891011 PHY Address 01, Page 800, Registers 2944 + n (n=0...3)

RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1... 3) for filters (8, 9, 10, 11).
RO	15:11	X	Reserved.

Table 103. Flexible Filter Length Table 1213—FFLT1213 PHY Address 01, Page 800, Registers 2952 + n (n=0...1)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1). For filters 12, 13.
RO	15:11	X	Reserved.

Table 104. Flexible Filter Length Table 1415—FFLT1415 PHY Address 01, Page 800, Registers 2956 + n (n=0...1)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN Minimum length for flexible filter n (n=0, 1). For filters 14, 15.
RO	15:11	X	Reserved.

8.10.3 ME Wake Up Control Status Register Description

Table 105. Management Wake Up Status—MWUS PHY Address 01, Page 801, Register 0

Attribute	Bit(s)	Initial Value	Description
RWC	0	0b	LNKC Link status changed
RO	15:1	0x0	Reserved



Table 106. Management 2 Host Control Register—MANC2H PHY Address 01, Page 801, Registers 30-31

Bits	Type	Reset	Description
0	RW	0	Flex Port 0 (FP0). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
1	RW	0	Flex Port 1 (FP1). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
2	RW	0	Flex Port 2 (FP2). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
3	RW	0	Flex TCO 0 (FT0). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
4	RW	0	Flex TCO 1 (FT1). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
5	RW	0	026F (FLT_026F). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
6	RW	0	0298 (FLT_0298). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
7	RW	0	ARP_REQuest (ARP_REQ). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
8	RW	0	ARP_RESponse (ARP_RES). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
9	RW	0	Broadcast (BR). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
10	RW	0	Neighbor (NE). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
11	RW	0	VLAN 0 (VLAN0). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
12	RW	0	VLAN 1 (VLAN1). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
13	RW	0	VLAN 2 (VLAN2). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
14	RW	0	VLAN 3 (VLAN3). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
15	RW	0	Manageability MAC (MNG_MAC). When set indicates that packets that are routed to the ME due to a match of the destination MAC address to any of the Shared Receive Addresses, are sent to the HOST as well.
16	RW	0	Flex Port 3 (FP3). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
17	RW	0	Manageability All Nodes Multicast MAC (MNG_ANM). When set to '1' packets that are routed to the ME due to a match of the destination MAC address to 33:33:00:00:00:01 are sent to the HOST as well.
18	RW	0	Low IPv6 address 0 (L24IPV60). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
19	RW	0	Low IPv6 address 1 (L24IPV61). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.



Table 106. Management 2 Host Control Register—MANC2H PHY Address 01, Page 801, Registers 30-31

20	RW	0	Low IPv6 address 2 (L24IPV62). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
21	RW	0	Low IPv6 address 3 (L24IPV63). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
22	RW	0	Flex Port 4 (FP4). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
23	RW	0	Flex Port 5 (FP5). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
24	RW	0	ICMP_IPV4 When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
25	RW	0	IPV4 When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
26	RW	0	Flex Port 6 (FP6). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
27	RW	0	Flex Port 7 (FP7). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
28	RW	0	Flex Port 8 (FP8). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
29	RW	0	Flex Port 9 (FP9). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
30	RW	0	Flex Port 10 (FP10). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
31	RW	0	Flex Port 11 (FP11). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.

Table 107. Management 2 Host Control Register 2—MANC2H2 PHY Address 01, Page 801, Registers 32-33

Bits	Type	Reset	Description
0	RW	0	TCPPORT0. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
1	RW	0	TCPPORT1. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
2	RW	0	TCPPORT2. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
3	RW	0	TCPPORT3. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
4	RW	0	TCPPORT4. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
5	RW	0	TCPPORT5. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
6	RW	0	TCPPORT6. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.



Table 107. Management 2 Host Control Register 2—MANC2H2 PHY Address 01, Page 801, Registers 32-33

7	RW	0	TCPPORT7. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
8	RW	0	TCPPORT8. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
9	RW	0	TCPPORT9. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
10	RW	00	TCPPORT10. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
11	RO	0	Reserved
12	RW	0	UDPPORT0. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
13	RW	0	DHCPv6. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
14	RW	0	EAPoUDP When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
15	RW	0	DNS. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
16	RW	0	UDPIPPORT0. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
17	RO	0	Reserved
18	RW	0	UDPIPPORT1. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
19	RO	0	Reserved
20	RW	0	UDPIPPORT2. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
21	RO	0	Reserved
22	RW	0	UDPIPPORT3. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
23	RO	0	Reserved
24	RW	0	MRFUTPF. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
25	RW	0	Host MTA. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
26	RW	0	ICMPv6. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
27	RO	0	Reserved
28	RW	0	ETHERTYPE0. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
29	RW	0	ETHERTYPE1. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
30	RW	0	ETHERTYPE2. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
31	RO	0	Reserved



Table 108. IPv4 ME Binding Control—IPV4MBC PHY Address 01, Page 801, Registers 40-41

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCSPORT0). When set, only packets that match the ME IPv4 filter may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCSPORT1). When set, only packets that match the ME IPv4 filter may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCSPORT2). When set, only packets that match the ME IPv4 filter may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCSPORT3). When set, only packets that match the ME IPv4 filter may pass TCPSPORT3 filtering.
4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCSPORT4). When set, only packets that match the ME IPv4 filter may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCSPORT5). When set, only packets that match the ME IPv4 filter may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCSPORT6). When set, only packets that match the ME IPv4 filter may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCSPORT7). When set, only packets that match the ME IPv4 filter may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCSPORT8). When set, only packets that match the ME IPv4 filter may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCSPORT9). When set, only packets that match the ME IPv4 filter may pass TCPSPORT9 filtering.
10	RW	00	Enable IP filter for TCPSPORT10 (EN_IPTCSPORT10). When set, only packets that match the ME IPv4 filter may pass TCPSPORT10 filtering.
11	RO	0	Reserved
12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the ME IPv4 filter may pass UDPPORT0 filtering.
13	RO	0	Reserved
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the ME IPv4 filter may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the ME IPv4 filter may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the ME IPv4 filter may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the ME IPv4 filter may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the ME IPv4 filter may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the ME IPv4 filter may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the ME IPv4 filter may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the ME IPv4 filter may pass flex port 5 filtering.



Table 108. IPv4 ME Binding Control—IPV4MBC PHY Address 01, Page 801, Registers 40-41

22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the ME IPv4 filter may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the ME IPv4 filter may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the ME IPv4 filter may pass flex port 8 filtering.
25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the ME IPv4 filter may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the ME IPv4 filter may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the ME IPv4 filter may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFUTPF (EN_IPMRFUTPF). When set, only packets that match the ME IPv4 filter may pass MRFUTPF range port filtering.
29	RW	0	Enable IP filter for ICMPv4. When set, only packets that match the ME IPv4 filter may pass ICMPv4 filtering.
30	RW	0	Enable IP filter for ARP. When set, only packets that match the ME IPv4 filter may pass ARP request filtering.
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the Host IPv4 filters may pass RMCP filtering

Table 109. IPv4 Host Binding Control—IPV4HBC PHY Address 01, Page 801, Registers 42-43

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCSPORT0). When set, only packets that match the Host IPv4 filters may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCSPORT1). When set, only packets that match the Host IPv4 filters may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCSPORT2). When set, only packets that match the Host IPv4 filters may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCSPORT3). When set, only packets that match the Host IPv4 filters may pass TCPSPORT3 filtering.
4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCSPORT4). When set, only packets that match the Host IPv4 filters may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCSPORT5). When set, only packets that match the Host IPv4 filters may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCSPORT6). When set, only packets that match the Host IPv4 filters may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCSPORT7). When set, only packets that match the Host IPv4 filters may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCSPORT8). When set, only packets that match the Host IPv4 filters may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCSPORT9). When set, only packets that match the Host IPv4 filters may pass TCPSPORT9 filtering.



Table 109. IPv4 Host Binding Control—IPV4HBC PHY Address 01, Page 801, Registers 42-43

10	RW	00	Enable IP filter for TCPPORT10 (EN_IPTCPPORT10). When set, only packets that match the Host IPv4 filters may pass TCPPORT10 filtering.
11	RO	0	Reserved
12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the Host IPv4 filters may pass UDPPORT0 filtering.
13	RO	0	Reserved
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the Host IPv4 filters may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the Host IPv4 filters may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the Host IPv4 filters may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the Host IPv4 filters may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the Host IPv4 filters may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the Host IPv4 filters may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the Host IPv4 filters may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the Host IPv4 filters may pass flex port 5 filtering.
22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the Host IPv4 filters may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the Host IPv4 filters may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the Host IPv4 filters may pass flex port 8 filtering.
25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the Host IPv4 filters may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the Host IPv4 filters may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the Host IPv4 filters may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFUTPF (EN_IPMRFUTPF). When set, only packets that match the Host IPv4 filters may pass MRFUTPF range port filtering.
29	RW	0	Enable IP filter for ICMPv4. When set, only packets that match the Host IPv4 filters may pass ICMPv4 filtering.
30	RW	0	Enable IP filter for ARP. When set, only packets that match the Host IPv4 filters may pass ARP request/response filtering.
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the Host IPv4 filters may pass RMCP filtering.



Table 110. IPv6 Binding Control—IPV6BC PHY Address 01, Page 801, Registers 50-51

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCSPORT0). When set, only packets that match the L24IPV6 filters may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCSPORT1). When set, only packets that match the L24IPV6 filters may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCSPORT2). When set, only packets that match the L24IPV6 filters may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCSPORT3). When set, only packets that match the L24IPV6 filters may pass TCPSPORT3 filtering.
4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCSPORT4). When set, only packets that match the L24IPV6 filters may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCSPORT5). When set, only packets that match the L24IPV6 filters may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCSPORT6). When set, only packets that match the L24IPV6 filters may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCSPORT7). When set, only packets that match the L24IPV6 filters may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCSPORT8). When set, only packets that match the L24IPV6 filters may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCSPORT9). When set, only packets that match the L24IPV6 filters may pass TCPSPORT9 filtering.
10	RW	00	Enable IP filter for TCPSPORT10 (EN_IPTCSPORT10). When set, only packets that match the L24IPV6 filters may pass TCPSPORT10 filtering.
11	RO	0	Reserved
12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the L24IPV6 filters may pass UDPPORT0 filtering.
13	RO	0	Enable IP filter for DHCPv6 (EN_IPDHCPv6). When set, only packets that match the L24IPV6 filters may pass DHCPv6 filtering.
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the L24IPV6 filters may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the L24IPV6 filters may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the L24IPV6 filters may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the L24IPV6 filters may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the L24IPV6 filters may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the L24IPV6 filters may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the L24IPV6 filters may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the L24IPV6 filters may pass flex port 5 filtering.


Table 110. IPv6 Binding Control—IPV6BC PHY Address 01, Page 801, Registers 50-51

22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the L24IPV6 filters may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the L24IPV6 filters may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the L24IPV6 filters may pass flex port 8 filtering.
25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the L24IPV6 filters may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the L24IPV6 filters may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the L24IPV6 filters may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFUTPF (EN_IPMRFUTPF). When set, only packets that match the L24IPV6 filters may pass MRFUTPF range port filtering.
29	RW	0	Enable IP filter for ICMPv6. When set, only packets that match the L24IPV6 filters may pass ICMPv6 filtering.
30	RO	0	Reserved
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the L24IPV6 filters may pass RMCP filtering

Table 111. SHRA Filter Enable Register—SHRAFER PHY Address 01, Page 801, Register 52

Bits	Type	Reset	Description
0	RW	0	EN_SHRA0_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[0] register and the Rx packet matches SHRA[0] the packet will be routed to the ME.
1	RW	0	EN_SHRA1_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[1] register and the Rx packet matches SHRA[1] the packet will be routed to the ME.
2	RW	0	EN_SHRA2_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[2] register and the Rx packet matches SHRA[2] the packet will be routed to the ME.
3	RW	0	EN_SHRA3_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[3] register and the Rx packet matches SHRA[3] the packet will be routed to the ME.
4	RW	0	EN_SHRA4_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[4] register and the Rx packet matches SHRA[4] the packet will be routed to the ME.
5	RW	0	EN_SHRA5_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[5] register and the Rx packet matches SHRA[5] the packet will be routed to the ME.
6	RW	0	EN_SHRA6_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[6] register and the Rx packet matches SHRA[6] the packet will be routed to the ME.
7	RW	0	EN_SHRA7_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[7] register and the Rx packet matches SHRA[7] the packet will be routed to the ME.



Table 111. SHRA Filter Enable Register—SHRAFER PHY Address 01, Page 801, Register 52

Bits	Type	Reset	Description
8	RW	0	EN_SHRA8_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[8] register and the Rx packet matches SHRA[8] the packet will be routed to the ME.
9	RW	0	EN_SHRA9_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[9] register and the Rx packet matches SHRA[9] the packet will be routed to the ME.
10	RW	0	EN_SHRA10_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[10] register and the Rx packet matches SHRA[10] the packet will be routed to the ME.
15:11	RO	0x0	Reserved

Table 112. Management IPv4 Address Table¹—MIP4AT PHY Address 01, Page 801, Registers 80-81

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	MIP4ADD Management IPv4 address filter (L.S. byte is first on the wire).

¹ The management IPv4 address table is used to store the management IPv4 address filter for ARP/IPv4 request packets and directed IPv4 packet wake ups.

Table 113. Management VLAN TAG Value Table—MAVTV PHY Address 01, Page 801, Registers 96 + n (n=0...3)

Attribute	Bit(s)	Initial Value	Description
RW	11:0	0x00	VLAN ID Contains the VLAN ID that should be compared with the incoming packet if bit 31 is set.
RO	14:12	0x00	Reserved
RW	15	0b	En Enable VID filtering.

Table 114. Management IPv6—MIPV6 PHY Address 01, Page 801, Registers 44-45 + 2*n (n = 0,1) and Registers 56-57 + 2*n (n = 0,1)

Attribute	Bit(s)	Initial Value	Description
RW	23:0	0x00	L24IPv6 This field holds the lower 24 bits of the IPv6 address to be compared with any IPv6 incoming packet.



RW	30:24	0x00	Those bits control whether a packet passing the IPv6 lower 24 bits should also be matched to additional filters as pass criteria of this filter. The supported values are: 0000000b = No additional match required. 0000001b = MFUTP0_EN match to MFUTP0 UDP/TCP flexible port 0. 0000010b = MFUTP1_EN match to MFUTP0 UDP/TCP flexible port 1. 0000100b = MFUTP2_EN match to MFUTP1 UDP/TCP flexible port 2. 0001000b = MFUTP3_EN match to MFUTP1 UDP/TCP flexible port 3. 0010000b = MFUTP4_EN match to MFUTP2 UDP/TCP flexible port 4. 0100000b = MFUTP5_EN match to MFUTP2 UDP/TCP flexible port 5. 1000000b = ICMP_EN match ICMP filtering.
RO	31	0b	Reserved

Table 115. Flexible TCO Filter Value/Mask Table LSBs—FTFT_L PHY Address 01, Page 801, Registers 512 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	TCO0 Byte n Flexible TCO filter 0, Byte n (n=0, 1... 127).
RW	8	X	TCO0 Mask Flexible TCO filter 0, Mask n (n=0, 1... 127).
RW	15:9	X	Reserved

Table 116. ME WoL Information—MEWOLI Address 01, Page 800, Register 74

Attribute	Bit(s)	Initial Value	Description
RW	2:0	X	Reserved
RW	15:3	X	ME WoL Information. Reserved for ME FW usage.

Table 117. Flexible TCO Filter Value¹/Mask Table MSBs—FTFT_H PHY Address 01, Page 801, Registers 513 + 2*n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	7:0	X	TCO1 Byte n Flexible TCO filter 1, Byte n (n=0, 1... 127).
RW	8	X	TCO1 Mask Flexible TCO filter 1, Mask n (n=0, 1... 127).
RW	15:9	X	Reserved

1. Values are used for flex TCO filter 1.

Table 118. TCO Flexible Filter Length Table—TCOFFLT PHY Address 01, Page 801, Registers 900 + n (n=0...1)

Attribute	Bit(s)	Initial Value	Description
RW	10:0	X	LEN TCO Minimum length for flexible TCO filter n (n=0, 1).
RW	15:11	X	Reserved. ¹

Notes:

1. All reserved fields read as zero's and are ignored on writes.



8.10.4 LPI MMD PHY Registers

LPI MMD PHY registers are part of the I217's EMI registers. These registers are accessed via MDIO by programming the EMI address to register MI16 and reading/writing the data from/to register MI17.

Table 119. Intel® Ethernet Controller I217 EMI Registers PHY Address 02, Page 0, Registers 16/17

IEEE MMD	MMD Bits	EMI Address	EMI Bits	Description	Type
3.0	10	9400	10	Clock stoppable	
3.1	11	9401	11	Tx LP idle received	
3.1	10	9401	10	Rx LP idle received	
3.1	9	9401	9	Tx LP idle indication	
3.1	8	9401	8	RX LP idle indication	
3.20	15:0	8000	15:0	EEE capability register	
3.22	15:0	A000	15:0	EEE wake error counter in 100BASE-TX mode	
3.22	15:0	A000	15:0	EEE wake error counter in 1000BASE-T mode	
7.60	15:0	8001	15:0	EEE advertisement	
7.61	15:0	8002	15:0	EEE LP advertisement	



9 Non-Volatile Memory (NVM)

9.1 Introduction

This section is intended for designs using a 10/100/1000 Mb/s Intel® C220 Series Chipset integrated LAN controller in conjunction with the I217.

There are several LAN clients that might access the NVM such as hardware, LAN driver, and BIOS. Refer to the *Intel® C220 Series Chipset External Design Specification (Intel® C220 Series Chipset EDS)* and the *Intel® C220 Series Chipset SPI Programming Guide* for more details.

Unless otherwise specified, all numbers in this section use the following numbering convention:

- Numbers that do not have a suffix are decimal (base 10).
- Numbers with a prefix of "0x" are hexadecimal (base 16).
- Numbers with a suffix of "b" are binary (base 2).

9.2 NVM Programming Procedure Overview

The LAN NVM shares space on an SPI Flash device (or devices) along with the BIOS, Manageability Firmware, and a Flash Descriptor Region. It is programmed through the Intel® C220 Series Chipset. This combined image is shown in [Figure 9-1](#). The Flash Descriptor Region is used to define vendor specific information and the location, allocated space, and read and write permissions for each region. The Manageability (ME) Region contains the code and configuration data for ME functions such as Intel® Active Management Technology. The system BIOS is contained in the BIOS Region. The ME Region and BIOS Region are beyond the scope of this document and a more detailed explanation of these areas can be found in the *Intel® C220 Series Chipset Family External Design Specification (Intel® C220 Series Chipset EDS)*. This document describes the LAN image contained in the Gigabit Ethernet (GbE) region.




Figure 9-1. LAN NVM Regions

To access the NVM, it is essential to correctly setup the following:

1. A valid Flash Descriptor Region must be present. Details for the Flash Descriptor Region are contained in the *Intel® C220 Series Chipset EDS*. This process is described in detail in the *Intel® Active Management Technology OEM Bring-Up Guide*.

The *Intel® Active Management Technology OEM Bring-Up Guide* can be obtained by contacting your local Intel Field Service Representative.

2. The GbE region must be part of the original image flashed onto the part.
3. For Intel LAN tools and drivers to work correctly, the BIOS must set the VSCC register(s) correctly. There are two sets of VSCC registers, the upper (UVSCC) and lower (LVSCC). Note that the LVSCC register is only used if the NVM attributes change. For example, the use of a second flash component, a change in erase size between segments, etc. Due to the architecture of the Intel® C220 Series Chipset, if these registers are not set correctly, the LAN tools might not report an error message even though the NVM contents remain unchanged. Refer to the *Intel® C220 Series Chipset EDS* for more information



4. The GbE region of the NVM must be accessible. To keep this region accessible, the Protected Range register of the GbE LAN Memory Mapped Configuration registers must be set to their default value of 0x0000 0000. (The GbE Protected Range registers are described in the Intel® C220 Series Chipset EDS).
5. The sector size of the NVM must equal 256 bytes, 4 KB, or 64 KB. When a Flash device that uses a 64 KB sector erase is used, the GbE region size must equal 128 KB. If the Flash part uses a 4 KB or 256-byte sector erase, then the GbE region size must be set to 8 KB.

The NVM image contains both static and dynamic data. The static data is the basic platform configuration, and includes OEM specific configuration bits as well as the unique Printed Circuit Board Assembly (PBA). The dynamic data holds the product's Ethernet Individual Address (IA) and Checksum. This file can be created using a text editor.

9.3 LAN NVM Format and Contents

Table 1 lists the NVM maps for the LAN region. Each word listed is described in detail in the following sections.

Table 1. LAN NVM Address Map

LAN Word Offset	NVM Byte Offset	Used By	15	0	Image Value
0x00	0x00	HW-Shared	Ethernet Address Byte 2, 1		IA (2, 1)
0x01	0x02	HW-Shared	Ethernet Address Byte 4, 3		IA (4, 3)
0x02	0x04	HW-Shared	Ethernet Address Byte 6, 5		IA (6, 5)
0x03	0x06	SW	Reserved		0x0800
0x04	0x08	SW	Reserved		0xFFFF
0x05	0x0A	SW	Image Version Information 1		
0x06	0x0C	SW	Reserved		0xFFFF
0x07	0x0E	SW	Reserved		0xFFFF
0x08	0x10	SW	PBA Low		
0x09	0x12	SW	PBA High		
0x0A	0x14	HW-PCI	PCI Init Control Word		
0x0B	0x16	HW-PCI	Subsystem ID		
0x0C	0x18	HW-PCI	Subsystem Vendor ID		
0x0D	0x1A	HW-PCI	Device ID		0x153A 0x153B
0x0E	0x1C	HW-PCI	Reserved		
0x0F	0x1E	HW-PCI	Reserved		
0x10	0x20	HW-PCI	LAN Power Consumption		
0x11	0x22	HW	Reserved		
0x12	0x24		Reserved		
0x13	0x26	HW-Shared	Shared Init Control Word		
0x14	0x28	HW-Shared	Extended Configuration Word 1		



0x15	0x2A	HW-Shared	Extended Configuration Word 2	
0x16	0x2C	HW-Shared	Extended Configuration Word 3	
LAN Word Offset	NVM Byte Offset	Used By	15	0
0x17	0x2E	HW-Shared	OEM Configuration Defaults	
0x18	0x30 (See note below.)	HW-Shared	LED 0 - 2	
0x19:0x2F	0x32:0x5E	HW-Shared	Reserved	0x0000
0x30:0x3E	0x60:0x7C	PXE	PXE Software Region	
0x3F	0x7E	SW	Software Checksum (Bytes 0x00 through 0x7D)	
0x40:0x4A	0x80:0x94	HW	G3 -> S5 PHY Configuration	

Table notes:

- SW = Software: This is access from the network configuration tools and drivers.
- PXE = PXE Boot Agent: This is access from the PXE option ROM code in BIOS.
- HW-Shared = Hardware-Shared: This is read when the shared configuration is reset.
- HW-PCI = Hardware-PCI: This is read when the PCI Configuration is reset.
- Word 0x30: For more information, see *Intel® iSCSI Remote Boot Application Notes for BIOS Engineers, Reference Number 322328*.

9.3.1 Hardware Accessed Words

This section describes the NVM words that are loaded by the integrated LAN controller hardware.

9.3.1.1 Ethernet Address (Words 0x00-0x02)

The Ethernet Individual Address (IA) is a 6-byte field that must be unique for each Network Interface Card (NIC) or LAN on Motherboard (LOM), and thus unique for each copy of the NVM image. The first three bytes are vendor specific—for example, the IA is equal to [00 AA 00] or [00 A0 C9] for Intel products. The value from this field is loaded into the Receive Address Register 0 (RAL0/RAH0).

For the purpose of this section, the IA byte numbering convention is indicated as follows; byte 1, bit 0 is first on the wire and byte 6, bit 7 is last. Note that byte 1, bit 0 is the unicast/multicast address indication while zero means unicast address. Byte 1, bit 1 identifies the global/local indication while zero means a global address.

	IA Byte/Value					
Vendor	1	2	3	4	5	6
Intel Original	00	AA	00	variable	variable	variable
Intel New	00	A0	C9	variable	variable	variable

9.3.1.2 PCI Init Control Word (Word 0x0A)

This word contains initialization values that:



- Set defaults for some internal registers
- Enable/disable specific features
- Determines which PCI configuration space values are loaded from the NVM

Bit	Name	Default	Description
15:8	Reserved	0x10	Reserved
7	AUX PWR	1b	Auxiliary Power Indication If set and if PM Ena is set, D3cold wake-up is advertised in the PMC register of the PCI function. 0b = No AUX power. 1b = AUX power.
6	PM Enable	1b	Power Management Enable (PME-WoL) Enables asserting PME in the PCI function at any power state. This bit affects the advertised PME_Support indication in the PMC register of the PCI function. 0b = Disable. 1b = Enable.
5:3	Reserved	0x0	These bits are reserved and must be set to 0x0.
2	Reserved	0b	Reserved, set to 0b.
1	Load Subsystem IDs	1b	Load Subsystem IDs from NVM When set to 1b, indicates that the device is to load its PCI Subsystem ID and Subsystem Vendor ID from the NVM (words 0x0B and 0x0C).
0	Load Device IDs	1b	Load Device ID from NVM When set to 1b, indicates that the device is to load its PCI Device ID from the NVM (word 0x0D).

9.3.1.3 Subsystem ID (Word 0x0B)

If the Load Subsystem ID in word 0x0A is set, this word is read-in to initialize the Subsystem ID. Default value is 0x0000.

9.3.1.4 Subsystem Vendor ID (Word 0x0C)

If the Load Subsystem ID in word 0x0A is set, this word is read-in to initialize the Subsystem Vendor ID. Default value is 0x8086.

9.3.1.5 Device ID (Word 0x0D)

If the Load Device ID in word 0x0A is set, this word is read-in to initialize the Device ID of the I217LM PHY. Default value is 0x153A.

Note: When the I217V SKU is used in combination with certain chipset SKUs, the default value for this word is 0x153B.

9.3.1.6 Words 0x0E and 0x0F Are Reserved

Default value is 0x0.

Note: In some OEM custom images these words are used for adding the track ID.



9.3.1.7 LAN Power Consumption (Word 0x10)

This word is meaningful only if the power management is enabled. The default value is 0x0702.

Bits	Name	Default	Description
15:8	LAN D0 Power	0x7	The value in this field is reflected in the PCI Power Management Data register for D0 power consumption and dissipation (<i>Data_Select</i> = 0 or 4). Power is defined in 100 mW units. The power also includes the external logic required for the LAN function.
7:5	Reserved	000b	Reserved, set to 000b.
4:0	LAN D3 Power	0x2	The value in this field is reflected in the PCI Power Management Data register for D3 power consumption and dissipation (<i>Data_Select</i> = 3 or 7). Power is defined in 100 mW units. The power also includes the external logic required for the LAN function. The most significant bits in the Data register that reflects the power values are padded with zeros.

9.3.1.8 Word 0x12 and Word 0x11 Are Reserved

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved, set to 0x0000.

9.3.1.9 Shared Init Control Word (Word 0x13)

This word controls general initialization values.

Bits	Name	Default	Description
15:14	Sign	10b	Valid Indication A 2-bit valid indication field indicates to the device that there is a valid NVM present. If the valid field does not equal 10b the integrated LAN controller does not read the rest of the NVM data and default values are used for the device configuration.
13	LinkSec Disable	1b	This bit enables the Linksec logic when set to 0b. It is loaded to the <i>LSecCK</i> bit in the <i>CTRL_EXT</i> register. When set, the Linksec logic is disabled and its clocks are gated. When cleared, the Linksec logic is enabled. Note: Refer to MACsec doc.
12:10	Reserved	001b	Reserved, set to 001b.
9	PHY PD Ena	0b	Enable PHY Power Down When set, enables PHY power down at DMoff/D3 or Dr and no WoL. This bit is loaded to the <i>PHY Power Down Enable</i> bit in the Extended Device Control (<i>CTRL_EXT</i>) register. 1b = Enable PHY power down. 0b = PHY always powered up.
8	Reserved	1b	Reserved, should be set to 1b.
7:6	PHYT	00b	PHY Device Type Indicates that the PHY is connected to the integrated LAN controller and resulted mode of operation of the integrated LAN controller/PHY link buses. 00b = I217 01b = Reserved. 10b = Reserved. 11b = Reserved.
5	Reserved	0b	Reserved, should be set to 0b.
4	FRCSPEED	0b	Default setting for the <i>Force Speed</i> bit in the Device Control register (<i>CTRL[11]</i>).



Bits	Name	Default	Description
3	FD	0b	Default setting for the <i>Full Duplex</i> bit in the Device Control register (CTRL[0]). The hardware default value is 1b.
2	Reserved	1b	Reserved, set to 0b.
1	CLK_CNT_1_4	0b	When set, automatically reduces DMA frequency. Mapped to the Device Status register (STATUS[31]).
0	Dynamic Clock gating	1b	When set, enables dynamic clock gating of the DMA and integrated LAN controller units. This bit is loaded to the <i>DynCK</i> bit in the CTRL_EXT register.

9.3.1.10 Extended Configuration Word 1 (Word 0x14)

Bits	Name	Default	Description
15:14	Reserved	00b	Reserved, set to 00b.
13	PHY Write Enable	1b	When set, enables loading of the extended PHY configuration area in the Lan Controller. When disabled, the extended PHY configuration area is ignored. Loaded to the EXTCNF_CTRL register.
12	OEM Write Enable	1b	When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[3]). 1b = OEM bits written to the Lan Controller 0b = No OEM bits configuration.
11:0	Extended Configuration Pointer	0x30	Defines the base address (in Dwords) of the Extended Configuration area in the NVM. The base address defines an offset value relative to the beginning of the LAN space in the NVM. A value of 0x00 is not supported when operating with the Lan Controller. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[27:16]).

9.3.1.11 Extended Configuration Word 2 (Word 0x15)

Bits	Name	Default	Description
15:8	Extended PHY Length	0x00	Size (in Dwords) of the Extended PHY configuration area loaded to the Extended Configuration Size register (EXTCNF_SIZE[23:16]). If an extended configuration area is disabled by bit 13 in word 0x14, its length must be set to zero.
7:0	Reserved	0x00	Reserved, must be set to 0x00.

Note: This field is dependent upon the length of the extended configuration area. The default value above is for mobile images to be used on platforms with a LAN switch. Refer to the image relevant to the platform for the appropriate default value.

9.3.1.12 Extended Configuration Word 3 (Word 0x16)

Bits	Name	Default	Description
15:0	Reserved	0x00	Reserved, set to 0x00.



9.3.1.13 OEM Configuration Defaults (Word 0x17)

This word defines the OEM fields for the PHY power management parameters loaded to the PHY Control (PHY_CTRL) register.

Bits	Name	Default	Description
15	B2B Enable	1b	Enable SPD in Back-to-Back Link setup
14	GbE Disable	0b	When set, GbE operation is disabled in all power states (including D0a).
13:12	Reserved	00b	Reserved, set to 00b.
11	GbE Disable in non-D0a	1b	Disables GbE operation in non-D0a states. This bit must be set if <i>GbE Disable</i> (bit 14) is set.
10	LPLU Enable in non-D0a	1b	Low Power Link Up Enables a reduction in link speed in non-D0a states when power policy and power management states are set to do so. This bit must be set if LPLU Enable in D0a bit is set.
9	LPLU Enable in D0a	0b	Low Power Link Up Enables a reduction in link speed in all power states.
8:0	Reserved	0x000	Reserved, set to 0x000.

9.3.1.14 LED 0–2 Configuration Defaults (Word 0x18)

This NVM word specifies the hardware defaults for the LED Control (LEDCTL) register fields controlling the LED1 (LINK_1000), LED0 (LINK/ACTIVITY) and LED2 (LINK_100) output behaviors. Refer to the *Intel® C220 Series Chipset Family PDG* and the *I217 Reference Schematics* for LED connection details. Also, [Table 2](#) lists mode encodings for LED outputs.

Note: In all system states, the LEDs function as defined in Word 0x18 of the GbE region of the NVM after the software driver loads.



Bits	Name	Default	Description
15	Blink Rate	0b	Blink Rate 0b = Blink at 200 ms on and 200 ms off. 1b = Blink at 83 ms on and 83 ms off.
14	LED2 Blink	0b	Initial Value of LED2_BLINK Field 0b = Non-blinking. 1b = Blinking.
13	LED2 Invert	0b	Initial Value of LED2_IVRT Field 0b = Active-low output.
12:10	LED2 Mode	110b	LED2 Mode Specifies what event/state/pattern is displayed on the LED2 output. 0110b = 100 Mb/s link_up.
9	LED1 Blink	0b	Initial Value of LED1_BLINK Field 0b = Non-blinking. 1b = Blinking.
8	LED1 Invert	0b	Initial Value of LED1_IVRT Field 0b = Active-low output.
7:5	LED1 Mode	111b	LED1 Mode Specifies what event/state/pattern is displayed on the LED1 output. 0111b = 1000 Mb/s link_up.
4	LED0 Blink	1b	Initial Value of LED0_BLINK Field 0b = Non-blinking. 1b = Blinking.
3	LED0 Invert	0b	Initial Value of LED0_IVRT Field 0b = Active-low output.
2:0	LED0 Mode	100b	LED0 Mode Specifies what event/state/pattern is displayed on the LED0 output. 100b = Filter activity on.

Table 2. Mode Encodings for LED Outputs

Mode	Mnemonic	State / Event Indicated
000b	LINK_10/1000	Asserted when either 10 or 1000 Mb/s link is established and maintained.
001b	LINK_100/1000	Asserted when either 100 or 1000 Mb/s link is established and maintained.
010b	LINK_UP	Asserted when any speed link is established and maintained.
011b	ACTIVITY	Asserted when link is established and packets are being transmitted or received.
100b	LINK/ACTIVITY	Asserted when link is established and when there is no transmit or receive activity.
101b	LINK_10	Asserted when a 10 Mb/s link is established and maintained.
110b	LINK_100	Asserted when a 100 Mb/s link is established and maintained.
111b	LINK_1000	Asserted when a 1000 Mb/s link is established and maintained.



9.3.1.15 Reserved (Word 0x19)

Bits	Name	Default	Description
15:0	Reserved	0x0A00	Reserved, set to 0x0A00.

9.3.1.16 Reserved (Word 0x1A)

Bits	Name	Default	Description
15:1	Reserved	0x0000	Reserved, set to 0x0000.
0	APM Enable	1b	APM Enable Initial value of Advanced Power Management Wake Up Enable in the Wake Up Control (WUC.APME) register. 1b = Advanced power management enabled. 0b = Advanced power management disabled.

9.3.1.17 Reserved (Word 0x1B)

Bits	Name	Default	Description
15:0	Reserved	0x0113	Reserved, set to 0x0113.

9.3.1.18 Reserved (Word 0x1C)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

9.3.1.19 Reserved (Word 0x1D)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

9.3.1.20 Reserved (Word 0x1E)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

9.3.1.21 Reserved (Word 0x1F)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

9.3.1.22 Reserved (Word 0x20)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved



9.3.1.23 Reserved (Word 0x21)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

9.3.1.24 Reserved (Word 0x22)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

9.3.1.25 Reserved (Word 0x23)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

9.3.1.26 Reserved (Word 0x24)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved, set to 0x0000.

9.3.1.27 Reserved (Word 0x25)

Bits	Name	Default	Description
15:0	Reserved	0x8080	Reserved, set to 0x8080

9.3.1.28 Reserved (Word 0x26)

Bits	Name	Default	Description
15:0	Reserved	0x4E00	Reserved, set to 0x4E00

9.3.1.29 Reserved (Word 0x27)

Bits	Name	Default	Description
15:0	Reserved	0x0886	Reserved

9.3.2 Software Accessed Words

9.3.2.1 PXE Words (Words 0x30 Through 0x3E)

Words 0x30 through 0x3E (bytes 0x60 through 0x7D) have been reserved for configuration and version values to be used by PXE code.



9.3.2.1.1 Boot Agent Main Setup Options (Word 0x30)

The boot agent software configuration is controlled by the NVM with the main setup options stored in word 0x30. These options are those that can be changed by using the Control-S setup menu or by using the IBA Intel Boot Agent utility. Note that these settings only apply to Boot Agent software.

Table 3. Boot Agent Main Setup Options

Bit	Name	Default	Description
15:14	Reserved	00b	Reserved, set to 00b.
13	Reserved	0b	Reserved, must be set to 0b.
12	FDP	0b	Force Full Duplex. Set this bit to 0b for half duplex and 1b for full duplex. Note that this bit is a don't care unless bits 10 and 11 are set.
11:10	FSP	00b	Force Speed. These bits determine speed. 01b = 10 Mb/s. 10b = 100 Mb/s. 11b = Not allowed. All zeros indicate auto-negotiate (the current bit state). Note that bit 12 is a don't care unless these bits are set.
9	Reserved	0b	Reserved Set this bit to 0b.
8	DSM	1b	Display Setup Message. If this bit is set to 1b, the "Press Control-S" message appears after the title message. The default for this bit is 1b.
7:6	PT	00b	Prompt Time. These bits control how long the "Press Control-S" setup prompt message appears, if enabled by DIM. 00b = 2 seconds (default). 01b = 3 seconds. 10b = 5 seconds. 11b = 0 seconds. Note that the Ctrl-S message does not appear if 0 seconds prompt time is selected.
5	Reserved	0b	Reserved
4:3	DBS	00b	Default Boot Selection. These bits select which device is the default boot device. These bits are only used if the agent detects that the BIOS does not support boot order selection or if the MODE field of word 0x31 is set to MODE_LEGACY. 00b = Network boot, then local boot. 01b = Local boot, then network boot. 10b = Network boot only. 11b = Local boot only.
2	Reserved	0b	Reserved
1:0	PS	00b	Protocol Select. These bits select the boot protocol. 00b = PXE (default value). 01b = Reserved. Other values are undefined.



9.3.2.1.2 Boot Agent Configuration Customization Options (Word 0x31)

Word 0x31 contains settings that can be programmed by an OEM or network administrator to customize the operation of the software. These settings cannot be changed from within the Control-S setup menu or the IBA Intel Boot Agent utility. The lower byte contains settings that would typically be configured by a network administrator using the Intel Boot Agent utility; these settings generally control which setup menu options are changeable. The upper byte are generally settings that would be used by an OEM to control the operation of the agent in a LOM environment, although there is nothing in the agent to prevent their use on a NIC implementation.

Table 4. Boot Agent Configuration Customization Options (Word 0x31)

Bit	Name	Default	Description
15:14	SIG	01b	Signature Set these bits to 11b to indicate valid data.
13:12	Reserved	00b	Reserved, must be set to 00b.
11		0b	Continuous Retry Disabled (0b default).
10:8	MODE	0x0	Selects the agent's boot order setup mode. This field changes the agent's default behavior in order to make it compatible with systems that do not completely support the BBS and PnP Expansion ROM standards. Valid values and their meanings are: 000b = Normal behavior. The agent attempts to detect BBS and PnP Expansion ROM support as it normally does. 001b = Force Legacy mode. The agent does not attempt to detect BBS or PnP Expansion ROM supports in the BIOS and assumes the BIOS is not compliant. The BIOS boot order can be changed in the Setup Menu. 010b = Force BBS mode. The agent assumes the BIOS is BBS-compliant, even though it might not be detected as such by the agent's detection code. The BIOS boot order CANNOT be changed in the Setup Menu. 011b = Force PnP Int18 mode. The agent assumes the BIOS allows boot order setup for PnP Expansion ROMs and hooks interrupt 18h (to inform the BIOS that the agent is a bootable device) in addition to registering as a BBS IPL device. The BIOS boot order CANNOT be changed in the Setup Menu. 100b = Force PnP Int19 mode. The agent assumes the BIOS allows boot order setup for PnP Expansion ROMs and hooks interrupt 0x19 (to inform the BIOS that the agent is a bootable device) in addition to registering as a BBS IPL device. The BIOS boot order CANNOT be changed in the Setup Menu. 101b = Reserved for future use. If specified, treated as value 000b. 110b = Reserved for future use. If specified, treated as value 000b. 111b = Reserved for future use. If specified, treated as value 000b.
7:6	Reserved	00b	Reserved, must be set to 00b.
5	DFU	0b	Disable Flash Update If set to 1b, no updates to the Flash image using PROSet is allowed. The default for this bit is 0b; allow Flash image updates using PROSet.



Bit	Name	Default	Description
4	DLWS	0b	Disable Legacy Wakeup Support If set to 1b, no changes to the Legacy OS Wakeup Support menu option is allowed. The default for this bit is 0b; allow Legacy OS Wakeup Support menu option changes.
3	DBS	0b	Disable Boot Selection If set to 1b, no changes to the boot order menu option is allowed. The default for this bit is 0b; allow boot order menu option changes.
2	DPS	0b	Disable Protocol Select If set to 1b, no changes to the boot protocol is allowed. The default for this bit is 0b; allow changes to the boot protocol.
1	DTM	0b	Disable Title Message If set to 1b, the title message displaying the version of the boot agent is suppressed; the Control-S message is also suppressed. This is for OEMs who do not want the boot agent to display any messages at system boot. The default for this bit is 0b; allow the title message that displays the version of the boot agent and the Control-S message.
0	DSM	0b	Disable Setup Menu If set to 1b, no invoking the setup menu by pressing Control-S is allowed. In this case, the EEPROM can only be changed via an external program. The default for this bit is 0b; allow invoking the setup menu by pressing Control-S.

9.3.2.1.3 Boot Agent Configuration Customization Options (Word 0x32)

Word 0x32 is used to store the version of the boot agent that is stored in the Flash image. When the Boot Agent loads, it can check this value to determine if any first-time configuration needs to be performed. The agent then updates this word with its version. Some diagnostic tools to report the version of the Boot Agent in the Flash also read this word. This word is only valid if the PPB is set to 0b. Otherwise the contents might be undefined.

Table 5. Boot Agent Configuration Customization Options (Word 0x32)

Bit	Name	Default	Description
15:12	MAJOR	0x1	PXE boot agent major version. The default for these bits is 0x1.
11:8	MINOR	0x2	PXE boot agent minor version. The default for these bits is 0x2.
7:0	BUILD	0x28	PXE boot agent build number. The default for these bits is 0x28.

9.3.2.1.4 IBA Capabilities (Word 0x33)

Word 0x33 is used to enumerate the boot technologies that have been programmed into the Flash. It is updated by IBA configuration tools and is not updated or read by IBA.



Table 6. IBA Capabilities

Bit	Name	Default	Description
15:14	SIG	01b	Signature These bits must be set to 01b to indicate that this word has been programmed by the agent or other configuration software.
13:5	Reserved	0x00	Reserved, must be set to 0x00.
4	iSCSI Boot Capability not present	0b	iSCSI boot capability not present (0b default).
3	EFI	0b	EFI EBC capability is present in Flash. 0b = The EFI code is not present (default). 1b = The EFI code is present.
2	Reserved	1b	Reserved, set to 1b.
1	UNDI	1b	PXE/UNDI capability is present in Flash. 1b = The PXE base code is present (default). 0b = The PXE base code is not present.
0	BC	1b	PXE base code is present in Flash. 0b = The PXE base code is not present. 1b = The PXE base code is present (default).

9.3.2.2 Checksum Word Calculation (Word 0x3F)

The Checksum word (Word 0x3F, NVM bytes 0x7E and 0x7F) is used to ensure that the base NVM image is a valid image. The value of this word should be calculated such that after adding all the words (0x00-0x3F) / bytes (0x00-0x7F), including the Checksum word itself, the sum should be 0xBABA. The initial value in the 16 bit summing register should be 0x0000 and the carry bit should be ignored after each addition.

Note: Hardware does not calculate the word 0x3F checksum during NVM write; it must be calculated by software independently and included in the NVM write data. Hardware does not compute a checksum over words 0x00-0x3F during NVM reads in order to determine validity of the NVM image; this field is provided strictly for software verification of NVM validity. All hardware configuration based on word 0x00-0x3F content is based on the validity of the Signature field of the NVM.

9.3.3 Basic Configuration Software Words

This section describes the meaningful NVM words in the basic configuration space that are used by software at word addresses 0x03-0x09.

9.3.3.1 Reserved (Word 0x3)

Bits	Name	Default	Description
15:12	Reserved	0x0	Reserved, set to 0x0.
11	LOM	1b	LOM Set to 1b.
10:0	Reserved	0x00	Reserved, set to 0x00.

Note: When software calculates the checksum, bit 1 of this word is set to 1b to indicate that the checksum is valid after the image is successfully programmed.



9.3.3.2 Reserved (Word 0x04)

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved

9.3.3.3 Image Version Information (Word 0x05)

0x03 denotes tuning to support a design that includes a LAN switch. 0x04 denotes tuning for all other designs.

9.3.3.4 PBA Low and PBA High (Words 0x08 and 0x09)

Bits	Word	Default	Description
15:0	0x08	0xFFFF	PBA low.
15:0	0x09	0xFFFF	PBA high.

The nine-digit Printed Board Assembly (PBA) number used for Intel manufactured Network Interface Cards (NICs) and Lan on Motherboard (LOMs) are stored in a four-byte field. The dash itself is not stored, neither is the first digit of the 3-digit suffix, as it is always zero for the affected products. Note that through the course of hardware ECOs, the suffix field (byte 4) is incremented. The purpose of this information is to allow customer support (or any user) to identify the exact revision level of a product.

Note: Network driver software should not rely on this field to identify the product or its capabilities.

Example: PBA number = 123456-003 to Word 0x08 = 0x1234; Word 0x09 = 0x5603.





10 Electrical & Timing Specifications

10.1 Introduction

This section describes the I217's recommended operating conditions, power delivery, DC electrical characteristics, power sequencing and reset requirements, PCIe specifications, reference clock, and packaging information.

10.2 Operating Conditions

10.2.1 Absolute Maximum Ratings

Symbol	Parameter	Min	Max	Units
T _{case}	Case Temperature Under Bias	0	106	°C
T _{storage}	Storage Temperature Range	-40	125	°C
V _i /V _o	3.3 Vdc I/O Voltage	-0.3	3.7	Vdc
VCC	3.3 Vdc Periphery DC Supply Voltage	-0.3	3.7	Vdc
VCC1p0	Core Vdc Supply Voltage	-0.3	1.2	Vdc

Notes:

1. Ratings in this table are those beyond which permanent device damage is likely to occur. These values should not be used as the limits for normal device operation. Exposure to absolute maximum rating conditions for extended periods might affect device reliability.
2. Recommended operation conditions require accuracy of power supply of +/-5% relative to the nominal voltage.
3. Maximum ratings are referenced to ground (VSS).

10.2.2 Recommended Operating Conditions

Symbol	Parameter	Min	Max	Units
T _a	Operating Temperature Range Commercial (Ambient; 0 CFS airflow)	0	85 ¹	°C
T _j	Junction Temperature		120	°C

Notes:

1. For normal device operation, adhere to the limits in this table. Sustained operations of a device at conditions exceeding these values, even if they are within the absolute maximum rating limits, can result in permanent device damage or impaired device reliability. Device functionality to stated Vdc and Vac limits is not guaranteed if conditions exceed recommended operating conditions.



10.3 Power Delivery

10.3.1 Voltage Regulator Power Supply Specifications

10.3.1.1 3.3 Vdc Rail

Title	Description	Min	Max	Units
Rise Time	Time from 10% to 90% mark	0.1	100	mS
Monotonicity	Voltage dip allowed in ramp	N/A	0	mV
Slope	Ramp rate at any given time between 10% and 90% Min: 0.8*V(min)/Rise time (max) Max: 0.8*V(max)/Rise time (min)	24	28800	V/S
Operational Range	Voltage range for normal operating conditions	3.13	3.46	V
Ripple	Maximum voltage ripple (peak to peak)	N/A	70	mV
Overshoot	Maximum overshoot allowed	N/A	100	mV

10.3.2 SVR Specification (Internal)

Parameter	Specifications			Units	Comments
	Min	Typ	Max		
Regulator Output Voltage ("Core" Voltage)	0.8		1.2	V dc	The default voltage is set to 0.93 Vdc
Output Voltage Accuracy	-3		+3	%	Not including line and load regulation errors.
Input Voltage Range	2.9	3.3	3.7	Vdc	Supply voltage range.
Load Current	0.01	0.3	0.5	A	Average value.
Output Voltage Under/Over Shoot	-10		+10	%	For min-to-max average load current change.
Transient Settling Time		100		µs	Duration of overshoot or undershoot.
Conversion Efficiency	80	85	90	%	
Switching Frequency		1.5625		MHz	
Output Filter Inductor	3.9	4.7		µH	
Output Filter Inductor DCR		0.1	0.318	Ω	+/-20%, values higher than the typical DCR value will lower the SVR conversion efficiency
Output Filter Capacitor	20			µF	
Output Filter Capacitor ESR		5	50	mΩ	
Input Capacitor	22			µF	

10.3.3 Power On/Off Sequence

- There is no power sequencing requirement for the I217.



10.3.4 Power Delivery Schematic Drawing

The following diagram shows the power delivery schematic at a high level. For the latest PHY schematic connection recommendations, refer to the Intel I217 GbE PHY reference schematic. For access to documentation, contact your Intel Field Service Representative.




Figure 10-1. Power Delivery Schematic

Table 1. Power Detection Threshold

Symbol	Parameter	Specifications			Units
		Min	Typ	Max	
V1a	High-threshold for 3.3 Vdc supply	2.35	2.45	2.6	Vdc
V2a	Low-threshold for 3.3 Vdc supply	2.1	2.45	2.6	Vdc
V1b	High-threshold for Core Vdc supply	0.6	0.75	0.85	Vdc
V2b	Low-threshold for Core Vdc supply	0.45	0.65	0.75	Vdc



10.4 I/O DC Parameter

10.4.1 3.3 Vdc I/O (Open Drain)

Parameter	Minimum	Typical	Maximum	Unit
VIL	-0.4	0	0.8	Vdc
VIH	2	3.3	3.6	Vdc
VOL	-0.4	0	0.4	Vdc
VOH	2.4	3.3	3.6	Vdc
Ipullup	30	50	75	μA
Ileakage			10	μA
Ci		2	4	pF

Pin Name	Bus Size	Description
CLK_REQ_N	1	Open-drain I/O.
SMB_CLK	1	Open-drain I(H)/O with snap back NMOS ESD cell.
SMB_DATA	1	Open-drain I(H)/O with snap back NMOS ESD cell.

Note: SMBus leakage current when the I217 is off is <180uA.

10.4.2 3.3 Vdc I/O

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.3	0	0.4	Vdc
VIH		2	3.3	3.6	Vdc
VOL	I _{OL} = 9 mA VCC = Min	-0.4	0	0.4	Vdc
VOH	I _{OH} = -9 mA VCC = Min	2	2.6	2.8	Vdc
Ipullup		30	50	75	μA
Ileakage		15 (pull down)	25 (pull down)	35 (pull down)	μA
Ci			2	4	pF
PU			50		KΩ
PD			50		KΩ

10.4.3 3.3 Vdc I/O Pull-up & Pull-down Information

Pin Name	Bus Size	Description
RSVD1_VCC3P3	1	I/O, PU
LED0, LED1, LED2	3	I/O, PU
JTAG_TDI	1	I/O, PU
JTAG_TMS	1	I/O, PU
JTAG_TDO	1	I/O, PU
JTAG_TCK	1	I/O, PU



10.4.4 Input Buffer Only

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.3	0	0.8	Vdc
VIH		2	3.3	3.6	Vdc
Ipullup		30	50	75	μA
Ileakage				10	μA
Ci			2	4	pF

10.4.5 Input Buffer Vdc I/O Pull-up & Pull-down

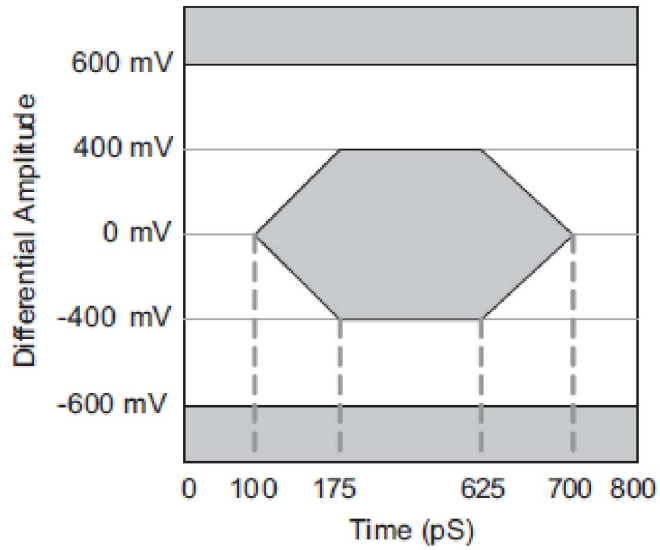
Pin Name	Bus Size	Description
LAN_DISABLE_N	1	I(H), PU
TEST_EN	1	I (no PU, no PD)
PE_RST_N	1	I(H), PU

10.4.6 PCIe DC/AC Specifications

10.4.6.1 PCIe Specifications (Transmitter)


Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
UI	Unit interval	799.92	800.08	ps	Each UI is 800 pS +/- 100 ppm
V _{tx-diff-pp}	Differential peak-to-peak Tx voltage swing	0.8	1.2	V dc	
T _{tx-eye}	Transmitter eye including all jitter sources	0.75		UI	
T _{tx-eye-median-to-max-jitter}	Maximum time between the jitter median and maximum deviation from the median		0.125	UI	
RL _{tx-diff}	Tx package plus silicon differential return loss	7		db	
RL _{tx-cm}	Tx package plus silicon common mode return loss	6		db	
Z _{tx-diff-dc}	DC differential Tx impedance	75	120	W	
V _{tx-cm-ac-p}	Tx V ac common mode voltage (2.5 GT/s)		20	mV	
I _{tx-short}	Transmitter short-circuit current limit		90	mA	
V _{tx-dc-cm}	Transmitter DC common mode voltage	0	3.6	V dc	
V _{tx-cm-dc-active-idle-delta}	Absolute delta of DC common mode voltage during L0 and electrical idle	0	100	mV	
V _{tx-cm-dc-line-delta}	Absolute delta of DC common mode voltage between D+ and D-	0	25	mV	

$V_{tx-idle-diff-ac-p}$	Electrical idle differential peak output voltage	0	20	mV	
$T_{tx-idle-set-to-idle}$	Maximum time to transition to a valid electrical idle after sending an EIOS		35	ns	
$T_{tx-idle-to-diff-data}$	Maximum time to transition to valid differential signaling after leaving electrical idle		35	ns	



Note: Not To Scale


Figure 10-2. Transmitter Eye Diagram



10.4.6.2 PCIe Specifications (Receiver)

Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
UI	Unit interval	799.92	800.08	ps	Each UI is 800 ps +/- 100 ppm
$V_{rx-diff-pp-cc}$	Differential peak-to-peak Rx voltage swing for common clock	0.175	1.2	V dc	
$V_{rx-diff-pp-dc}$	Differential peak-to-peak Rx voltage swing for data clock	0.175	1.2	V dc	
T_{rx-eye}	Receiver minimum eye time opening	0.4	N/A	UI	
$T_{rx-eye-median2maxjitter}$	Maximum time delta between median and deviation from median	N/A	0.3	UI	
$RL_{rx-diff}$	Rx differential return loss	6	N/A	dB	
RL_{rx-cm}	Rx CM return loss	5	N/A	dB	
$Z_{rx-diff-dc}$	Rx differential Vdc impedance	80	120	W	
$V_{rx-cm-ac-p}$	Rx Vac CM voltage	N/A	150	mVp	
$Z_{rx-high-imp-dc-pos}$	DC input CM impedance for V>0	50 K	N/A	W	
$Z_{rx-high-imp-dc-neg}$	DC input CM impedance for V<0	1 K	N/A	W	
$V_{rx-idle-det-diff-p}$	Electrical idle detect threshold	65	175	mV	

Note: The I217 has integrated PCIe termination that results in attenuating the voltage swing of the PCIe clock supplied by the platform. This is in compliance with the PCIe CEM 1.1 specification. For more detail, check the platform design guide.



Note: Not To Scale

Figure 10-3. Receiver Eye Diagram



10.4.6.3 PCIe Clock Specifications

The PCIe clock specification can be found in the PCI Express Card Electromechanical Specification 1.1, section 2.1.

10.5 Discrete/Integrated Magnetics Specifications

Criteria	Condition	Values (Min/Max)
Voltage Isolation	At 50 to 60 Hz for 60 seconds	1500 Vrms (min)
	For 60 seconds	2250 V dc (min)
Open Circuit Inductance (OCL) or OCL (alternate)	With 8 mA DC bias at 25 °C	400 μH (min)
	With 8 mA DC bias at 0 °C to 70 °C	350 μH (min)
Insertion Loss	100 kHz through 999 kHz	1 dB (max)
	1.0 MHz through 60 MHz	0.6 dB (max)
	60.1 MHz through 80 MHz	0.8 dB (max)
	80.1 MHz through 100 MHz	1.0 dB (max)
	100.1 MHz through 125 MHz	2.4 dB (max)
Return Loss	1.0 MHz through 40 MHz	18 dB (min)
	40.1 MHz through 100 MHz	12 to 20 * LOG (frequency in MHz / 80) dB (min)
	When reference impedance is 85 Ω, 100 Ω, and 115 Ω	
	Note that return loss values might vary with MDI trace lengths. The LAN magnetics might need to be measured in the platform where it is used.	
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	-50.3+(8.8*(freq in MHz / 30)) dB (max)
	30 MHz through 250 MHz	-26-(16.8*(LOG(freq in MHz / 250))) dB (max)
	250.1 MHz through 375 MHz	-26 dB (max)
Crosstalk Isolation Integrated Modules	1.0 MHz through 10 MHz	-50.8+(8.8*(freq in MHz / 10)) dB (max)
	10.1 MHz through 100 MHz	-26-(16.8*(LOG(freq in MHz / 100))) dB (max)
	100.1 MHz through 375 MHz	-26 dB (max)
Diff to CMR	1.0 MHz through 29.9 MHz	-40.2+(5.3*((freq in MHz / 30)) dB (max)
	30 MHz through 500 MHz	-22-(14*(LOG((freq in MHz / 250))) dB (max)
CM to CMR	1.0 MHz through 270 MHz	-57+(38*((freq in MHz / 270)) dB (max)
	270.1 MHz through 300 MHz	-17-2*((300-(freq in MHz) / 30) dB (max)
	300.1 MHz through 500 MHz	-17 dB (max)

10.6 Mechanical

Body Size (mm)	Ball Count	Ball Pitch	Ball Matrix	Center Matrix	Substrate
6x6 mm	48	0.4 mm	N/A, Peripheral	N/A, Exposed Pad	N/A Lead frame-Based Package



10.7 Oscillator/Crystal Specifications

Table 2 lists required parameters. See the reference schematic for connection information.

Table 2. External Crystal Specifications

Parameter Name	Symbol	Recommended Value	Max/Min Range	Conditions
Frequency	f_o	25 [MHz]		@25 [°C]
Vibration Mode		Fundamental		
Frequency Tolerance @25 °C	$Df/f_o @25^{\circ}C$	±30 [ppm]		@25 [°C]
Temperature Tolerance	Df/f_o	±30 [ppm]		0 to +70 [°C]
Series Resistance (ESR)	R_s		50 [Ω] max	@25 [MHz]
Crystal Load Capacitance	C_{load}	18 [pF]		
Shunt Capacitance	C_o		6 [pF] max	
Drive Level	D_L		200 [μ W] max	
Aging	Df/f_o	±5 ppm per year	±5 ppm per year max	
Calibration Mode		Parallel		
Insulation Resistance			500 [$M\Omega$] min	@ 100 Vdc

Note: Crystal must meet or exceed the specified drive Level (D_L). Refer to the crystal design guidelines in the Intel® 5 Series Family PDG.

Table 3. Clock Oscillator Specifications

Parameter Name	Symbol/Parameter	Conditions	Min	Typ	Max	Unit
Frequency	f_o	@25 [°C]		25.0		MHz
Clock Amplitude	V_{max}		0.8		1.8	Vdc
Clock Amplitude	V_{min}				0	Vdc
Frequency Tolerance	f/f_o	20 to +70		±50		[ppm]
Operating Temperature	T_{opr}	-20 to +70				°C
Aging	f/f_o			±5 ppm per year		[ppm]
TH_XTAL_IN	XTAL_IN High Time		13	20		nS
TL_XTAL_IN	XTAL_IN Low Time		13	20		nS
TR_XTAL_IN	XTAL_IN Rise	10% to 90%			5	nS
TF_XTAL_IN	XTAL_IN Fall	10% to 90%			5	nS
TJ_XTAL_IN	XTAL_IN Total Jitter				200 ¹	pS

Notes:

1. Broadband peak-to-peak = 200 pS, Broadband rms = 3 pS, 12 KHz to 20 MHz rms = 1 ps

XTAL_IN/XTAL_OUT Timing




Figure 10-4. XTAL Timing Diagram



NOTE: *This page intentionally left blank.*



11 Driver Programming Interface

11.1 Registers Byte Ordering

This section defines the structure of registers that contain fields carried over the network. Some examples are L2, L3, L4 fields and LinkSec fields.

The following example is used to describe byte ordering over the wire (hex notation):

Last	First
...,06, 05, 04, 03, 02, 01, 00	

where each byte is sent with the LSbit first. That is, the bit order over the wire for this example is

Last	First
..., 0000 0011, 0000 0010, 0000 0001, 0000 0000	

The general rule for register ordering is to use Host Ordering (also called little endian). Using the above example, a 6-byte fields (e.g. MAC address) is stored in a CSR in the following manner:

	Byte 3	Byte 2	Byte 1	Byte0
DW address (N)	0x03	0x02	0x01	0x00
DW address (N+4)	0x05	0x04

The exceptions listed below use network ordering (also called big endian). Using the above example, a 16-bit field (e.g. EtherType) is stored in a CSR in the following manner:

	Byte 3	Byte 2	Byte 1	Byte0
(DW aligned)	0x00	0x01

or

(Word aligned)	0x00	0x01
----------------	------	------	-----	-----

The following exceptions use network ordering:

- All ETHERType fields

Note: The “normal” notation as it appears in text books, etc. is to use network ordering. Example: Suppose a MAC address of 00-A0-C9-00-00-00. The order on the network is 00, then A0, then C9, etc. However, the host ordering presentation would be

	Byte 3	Byte 2	Byte 1	Byte0
DW address (N)	00	C9	A0	00
DW address (N+4)	00	00



11.1.1 Register Conventions

All registers in the LAN Controller are defined to be 32 bits, so write cycles should be accessed as 32 bit double-words, There are some exceptions to this rule:

- Register pairs where two 32 bit registers make up a larger logical size

Reserved bit positions: Some registers contain certain bits that are marked as “reserved.” These bits should never be set to a value of “one” by software. Reads from registers containing reserved bits may return indeterminate values in the reserved bit-positions unless read values are explicitly stated. When read, these reserved bits should be ignored by software.

Reserved and/or undefined addresses: any register address not explicitly declared in this specification should be considered to be reserved, and should not be written to. Writing to reserved or undefined register addresses may cause indeterminate behavior. Reads from reserved or undefined configuration register addresses may return indeterminate values unless read values are explicitly stated for specific addresses. Reserved fields within defined registers are defined as Read-Only (RO). When writing to these registers the RO fields should be set to their init value. Reading from reserved fields may return indeterminate values.

Initial values: most registers define the initial hardware values prior to being programmed. In some cases, hardware initial values are undefined and will be listed as such via the text “undefined”, “unknown”, or “X”. Some such configuration values may need to be set via NVM configuration or via software in order for proper operation to occur; this need is dependent on the function of the bit. Other registers may cite a hardware default which is overridden by a higher-precedence operation. Operations which may supersede hardware defaults may include a valid NVM load, completion of a hardware operation (such as hardware auto-negotiation), or writing of a different register whose value is then reflected in another bit.

For registers that should be accessed as 32 bit double words, partial writes (less than a 32 bit double word) will not take effect (i.e. the write is ignored). Partial reads will return all 32 bits of data regardless of the byte enables.

Note: Partial reads to read-on-clear registers (e.g. ICR) can have unexpected results since all 32 bits are actually read regardless of the byte enables. Partial reads should not be done.

Note: All statistics registers are implemented as 32 bit registers. Though some logical statistics registers represent counters in excess of 32-bits in width, registers must be accessed using 32-bit operations (e.g. independent access to each 32-bit field).

Note: The LAN Controller supports a single memory access at a time.

See special notes for Multicast Table Arrays in their specific register definitions.

11.1.2 PCI Configuration and Status Registers—CSR Space

11.1.2.1 PCI Register Map

All configuration registers are listed in the table below. These registers are ordered by grouping and are not necessarily listed in order that they appear in the address space.

Register Based Legend:



RW—Read Write register.

RO—Read Only Register.

RO/CR—Read Only Register, Clear on Read.

RO/V—Read Only Register, Read status is not constant.

RW/RO—Read write by FW, Read only by SW.

R/WC—Read Write Clear registers. Writing '0's has no affect. Writing '1's clears the appropriate fields (see detailed description of the specific registers).

RW/V—Read Write register This bit self-clears immediately.

RW/SN—Read Write register initial value loaded from NVM.

RC/WC—Read Clear/ Write Clear registers. Writing '0's has no affect. Writing '1's clears the appropriate fields. Read may also clear the register depending on enablement (see specific registers).

RWC/CR/V—Read Write register clear on read, clear on write.

WO—Write only registers. Reading from these registers does not reflect any meaningful data. Mostly it would be all zero's (see detailed description of the specific registers).

Table 11-1. Register Summary

Offset	Abbreviation	Name	RW	Link
General Register Descriptions				
0x00000	CTRL	Device Control Register	RW	160
0x00008	STATUS	Device Status Register	RO	162
0x0000C	STRAP	Strapping Option Register	RO	164
0x00010	FEXTNVM6	Future Extended6 NVM Register	RW	173
0x00014	FEXTNVM5	Future Extended5 NVM Register	RW	172
0x00018	CTRL_EXT	Extended Device Control Register	RW	164
0x00020	MDIC	MDI Control Register	RW	165
0x00024	FEXTNVM4	Future Extended4 NVM Register	RW	171
0x00028	FEXTNVM	Future Extended NVM Register	RW	169
0x0002C	FEXT	Future Extended Register	RW	170
0x00030	FEXTNVM2	Future Extended2 NVM Register	RW	170
0x00034	KUMCTRLSTA	Kumeran control and status registers	RW	171
0x00038	BUSNUM	Device and Bus Number	RO	174
0x000F8	LTRV	Latency Tolerance Reporting Value	RW	174
0x000FC	LPIC	Low Power Idle Control	RW	175
0x00170	FCTTV	Flow Control Transmit Timer Value	RW	176
0x05F40	FCRTV	Flow Control Refresh Threshold Value	RW	176
0x00F00	EXTCNF_CTRL	Extended Configuration Control	RW	176
0x00F08	EXTCNF_SIZE	Extended Configuration Size	RW	177
0x00F10	PHY_CTRL	PHY Control Register	RW	177



Offset	Abbreviation	Name	RW	Link
0x00F18	PCIEANACFG	PCIE Analog Configuration	RW	178
0x01000	PBA	Packet Buffer Allocation	RW	178
0x01008	PBS	Packet Buffer Size	RW	178
0x05B00	DCR	DMA Control Register	RW	178
Interrupt Register Descriptions				
0x000C0	ICR	Interrupt Cause Read Register	RC/WC	178
0x000C4	ITR	Interrupt Throttling Register	RW	180
0x000C8	ICS	Interrupt Cause Set Register	WO	181
0x000D0	IMS	Interrupt Mask Set/Read Register	RW	181
0x000D8	IMC	Interrupt Mask Clear Register	WO	182
0x000E0	Mask—IAM	Interrupt Acknowledge Auto	RW	184
Receive Register Descriptions				
0x00100	RCTL	Receive Control Register	RW	184
0x00104	RCTL1	Receive Control Register 1	RW	186
0x02008	ERT	Early Receive Threshold	RW	187
0x02170 + n*0x4 [n=0..1]	PSRCTL	Packet Split Receive Control Register	RW	188
0x02160	FCRTL	Flow Control Receive Threshold Low	RW	188
0x02168	FCRTH	Flow Control Receive Threshold High	RW	189
0x02800 + n*0x100[n=0..1]	RDBAL	Receive Descriptor Base Address Low queue	RW	189
0x02804 + n*0x100[n=0..1]	RDBAH	Receive Descriptor Base Address High queue	RW	189
0x02808 + n*0x100[n=0..1]	RDLEN	Receive Descriptor Length queue	RW	189
0x02810 + n*0x100[n=0..1]	RDH	Receive Descriptor Head queue	RW	190
0x02818 + n*0x100[n=0..1]	RDT	Receive Descriptor Tail queue	RW	190
0x02820 + n*0x100[n=0..1]	RDTR	Interrupt Delay Timer (Packet Timer)	RW	190
0x02828 + n*0x100[n=0..1]	RXDCTL	Receive Descriptor Control	RW	191
0x0282C	RADV	Receive Interrupt Absolute Delay Timer	RW	192
0x02C00	RSRPD	Receive Small Packet Detect Interrupt	RW	192
0x02C08	RAID	Receive ACK Interrupt Delay Register	RW	193
0x02C10	CPUVEC	CPU Vector Register	RW	193
0x05000	RXCSUM	Receive Checksum Control	RW	193
0x05008	RFCTL	Receive Filter Control Register	RW	194
0x05200-0x0527C	MTA[31:0]	Multicast Table Array	RW	195
0x05400	RAL	Receive Address Low	RW	196
0x05404	RAH	Receive Address High	RW	196
0x05408 + 8*n (n=0...10)	SRAL	Shared Receive Address Low	RW	197
0x0540C + 8*n (n=0...10)	SRAH	Shared Receive Address High 0...10	RW	197
0x05454	SHRAH[9]	Shared Receive Address High 9	RW	198
0x05818	MRQC	Multiple Receive Queues Command Register	RW	198
0x05864	RSSIM	RSS Interrupt Mask Register	RW	199
0x05868	RSSIR	RSS Interrupt Request Register	RW	199
0x05C00 + 4*n (n=0...31)	RETA	Redirection Table	RW	199



Offset	Abbreviation	Name	RW	Link
0x05C80 + 4*n (n=0..9)	RSSRK	Random Key Register	RW	199
Transmit Register Descriptions				
0x00400	TCTL	Transmit Control Register	RW	200
0x00410	TIPG	Transmit IPG Register	RW	201
0x00458	AIT	Adaptive IFS Throttle	RW	202
0x03800 + n*0x100[n=0..1]	TDBAL	Transmit Descriptor Base Address Low	RW	202
0x03804 + n*0x100[n=0..1]	TDBAH	Transmit Descriptor Base Address High	RW	202
0x03808 + n*0x100[n=0..1]	TDLEN	Transmit Descriptor Length	RW	203
0x03810 + n*0x100[n=0..1]	TDH	Transmit Descriptor Head	RW	203
0x03818 + n*0x100[n=0..1]	TDT	Transmit Descriptor Tail	RW	203
0x03840 + n*0x100[n=0..1]	TARC	Transmit Arbitration Count	RW	203
0x03820	TIDV	Transmit Interrupt Delay Value	RW	204
0x03828 + n*0x100[n=0..1]	TXDCTL	Transmit Descriptor Control	RW	205
0x0382C	TADV	Transmit Absolute Interrupt Delay Value	RW	207
Statistic Register Descriptions				
0x04000	CRCERRS	CRC Error Count	RO	207
0x04004	ALGNERRC	Alignment Error Count	RO	208
0x0400C	RXERRC	RX Error Count	RO	208
0x04010	MPC	Missed Packets Count	RO	208
0x0403C	CEXTERR	Carrier Extension Error Count	RO	208
0x04040	RLEC	Receive Length Error Count	RO	208
0x04048	XONRXC	XON Received Count	RO	209
0x0404C	XONTXC	XON Transmitted Count	RO	209
0x04050	XOFFRXC	XOFF Received Count	RO	209
0x04054	XOFFTXC	XOFF Transmitted Count	RO	209
0x04058	FCRUC	FC Received Unsupported Count	RO	209
0x04074	GPRC	Good Packets Received Count	RO	210
0x04078	BPRC	Broadcast Packets Received Count	RO	210
0x0407C	MPRC	Multicast Packets Received Count	RO	210
0x04080	GPTC	Good Packets Transmitted Count	RO	210
0x04088	GORCL	Good Octets Received Count Low	RO	210
0x0408C	GORCH	Good Octets Received Count High	RO	210
0x04090	GOTCL	Good Octets Transmitted Count Low	RO	211
0x04094	GOTCH	Good Octets Transmitted Count High	RO	211
0x040A0	RNBC	Receive No Buffers Count	RO	211
0x040A4	RUC	Receive Undersize Count	RO	211
0x040A8	RFC	Receive Fragment Count	RO	212
0x040AC	ROC	Receive Oversize Count	RO	212
0x040B0	RJC	Receive Jabber Count	RO	212
0x040B4	MNGPRC	Management Packets Received Count	RO	212
0x040B8	MNGPDC	Management Packets Dropped Count	RO	213



Offset	Abbreviation	Name	RW	Link
0x040BC	MNGPTC	Management Packets Transmitted Count	RO	213
0x040D8	TCBPD	Tx Circuit Breaker Packets Dropped	RO	213
0x040C0	TORL	Total Octets Received Low	RO	213
0x040C4	TORH	Total Octets Received High	RO	213
0x040C8	TOTL	Total Octets Transmitted	RO	214
0x040CC	TOTH	Total Octets Transmitted	RO	214
0x040D0	TPR	Total Packets Received	RO	214
0x040D4	TPT	Total Packets Transmitted	RO	214
0x040F0	MPTC	Multicast Packets Transmitted Count	RO	215
0x040F4	BPTC	Broadcast Packets Transmitted Count	RO	215
0x040F8	TSCTC	TCP Segmentation Context Transmitted Count	RO	215
0x04100	IAC	Interrupt Assertion Count	RO	215
Management Register Descriptions				
0x05800	WUC	Wake Up Control Register	RW	216
0x05808	WUFC	Wake Up Filter Control Register	RW	216
0x05810	WUS	Wake Up Status Register	RW	217
0x05820	MANC_S	Management Control Register Shadow	RO	218
0x05824	MANC2_S	Management Control 2 Register Shadow	RO	116
0x05828	MANC3_S	Management Control 3 Register Shadow	RO	218
0x0582C	SHRAFER_S	SHRA Filter Enable Register Shadow	RO	218
0x05860	MANC2H_S	Management Control To Host Register Shadow	RO	218
0x05870	MANC2H2_S	Management Control To Host Register 2 Shadow	RO	218
0x5838	IPAV	IP Address Valid	RW	218
0x05840 + 8*n (n=1...3)	IP4AT	IPv4 Address Table	RW	218
0x05880 + 4*n (n=0...3)	IP6AT	IPv6 Address Table	RW	218
0x05B50	H2ME	Host to ME Register	RW	219
0x05B54	FWSM_S	Firmware Semaphore Register Shadow	RO	219
0x05B58	EXFWSM_S	Extended Firmware Semaphore Register Shadow	RO	219
0x05F00 + 8*n (n=0...7)	FFLT	Flexible Filter Length Table	RW	219
0x09000 + 8*n (n=0...127)	FFMT	Flexible Filter Mask Table	RW	220
0x09800 + 8*n (n=0...127)	FFVT	Flexible Filter Value Table	RW	220
0x09804 + 8*n (n=0...127)	FFVT2	Flexible Filter Value Table	RW	220
0x0B620	TSYNCRXCTL	RX Time Sync Control register	RW	221
0x0B624	RXSTMP_L	RX timestamp Low	RO	221
0x0B628	RXSTMP_H	RX timestamp High	RO	221
0x0B62C	RXSATRL	RX timestamp attributes low	RO	221
0x0B630	RXSATRH	RX timestamp attributes high	RO	221
0x0B634	RXMTRL	RX message type register low	RW	222
0x0B638	RXUDP	RX UDP port	RW	222
0x0B614	TSYNCTXCTL	TX Time Sync Control register	RW	222
0x0B618	TXSTMP_L	TX timestamp value Low	RO	222



Offset	Abbreviation	Name	RW	Link
0x0B61C	TXSTMPH	TX timestamp value High	RO	222
0x0B600	SYSTIML	System time register Low	RO	222
0x0B604	SYSTIMH	System time register High	RO	223
0x0B608	TIMINCA	Increment attributes register	RW	223
0x0B60C	TIMADJL	Time adjustment offset register low	RW	223
0x0B610	TIMADJH	Time adjustment offset register high	RW	223
Diagnostic Register Descriptions				
0x02410	RDFH	Receive Data FIFO Head Register	RW	223
0x02418	RDFT	Receive Data FIFO Tail Register	RW	223
0x02420	RDFHS	Receive Data FIFO Head Saved Register	RW	224
0x02428	RDFTS	Receive Data FIFO Tail Saved Register	RW	224
0x02430	RDFPC	Receive Data FIFO Packet Count	RW	224
0x03410	TDFH	Transmit Data FIFO Head Register	RW	224
0x03418	TDFT	Transmit Data FIFO Tail Register	RW	225
0x03420	TDFHS	Transmit Data FIFO Head Saved Register	RW	225
0x03428	TDFTS	Transmit Data FIFO Tail Saved Register	RW	225
0x03430	TDFPC	Transmit Data FIFO Packet Count	RW	225
0x10000 - 0x15FFC	PBM	Packet Buffer Memory	RW	226
0x0C000 - 0x0C3FC	GMD	Ghost Memory Data	RW	226
0x0C400 - 0x0C5FC	DMD	Descriptor Memory Data	RW	226
Hidden Testability Register Descriptions				
0x05B60	MNGCCR	MNG CSR Control register	RW	226
0x05B64	MNGCAR	MNG CSR Address Register	RW	227
0x05B68	MNGCDO	MNG CSR Data Out Register	RO	227
0x05B6C	MNGCDI	MNG CSR Data In Register	RW	227
LinkSec Register Descriptions				
0x0B000	LSECTXCAP	LinkSec TX Capabilities register	RW	228
0x0B300	LSECRXCAP	LinkSec RX Capabilities register	RW	228
0x0B004	LSECTXCTRL	LinkSec TX Control register	RW	228
0x0B304	LSECRXCTRL	LinkSec RX Control register	RW	229
0x0B008	LSECTXSCL	LinkSec TX SCI Low	RW	229
0x0B00C	LSECTXSCH	LinkSec TX SCI High	RW	229
0x0B010	LSECTXSA	LinkSec TX SA	RW	230
0x0B018	LSECTXPN0	LinkSec TX SA PN 0	RW	230
0x0B01C	LSECTXPN1	LinkSec TX SA PN 1	RW	230
0x0B020	LSECTXKEY0 [n]	LinkSec TX Key 0 0 + 4*n (n=0...3) 0x0B02	WO	231
0x0B030	LSECTXKEY1 [n]	LinkSec TX Key 1 0 + 4*n (n=0...3) 0x0B03	WO	231
0x0B3D0 + 4*n (n=0...3)	LSECRXSCL[n]	LinkSec RX SCI Low	RW	231
0x0B3E0 + 4*n (n=0...3)	LSECRXSCH[n]	LinkSec RX SCI High	RW	231
0x0B310 + 4*n (n=0...7)	LSECRXSA[n]	LinkSec RX SA	RW	232
0x0B330 + 4*n (n=0...7)	LSECRXSAPN	LinkSec RX SA PN	RW	232



Offset	Abbreviation	Name	RW	Link
0x0B350 + 0x10*n (n=0...7) + 4*m (m=0..3)	LSECRXKEY[n,m]	LinkSec RX Key	WO	233
0x04300	LSECTXUT	Tx Untagged Packet Counter	RC	233
0x04304	LSECTXPKTE	Encrypted Tx Packets	RC	233
0x04308	LSECTXPKTP	Protected Tx Packets	RC	233
0x0430C	LSECTXOCTE	Encrypted Tx Octets	RC	234
0x04310	LSECTXOCTP	Protected Tx Octets	RC	234
0x04314	LSECRXUT	LinkSec Untagged RX Packet	RC	234
0x0431C	LSECRXOCTE	LinkSec RX Octets Decrypted	RC	234
0x04320	LSECRXOCTP	LinkSec RX Octets Validated	RC	234
0x04324	LSECRXBAD	LinkSec RX Packet with Bad Tag	RC	235
0x04328	LSECRXNOSCI	LinkSec RX Packet No SCI	RC	235
0x0432C	LSECRXUNSCI	LinkSec RX Packet Unknown SCI count	RC	235
0x04330	LSECRXUNCH	LinkSec RX Unchecked Packets	RC	235
0x04340 + 4*n (n=0...3)	LSECRXDELAY[n]	LinkSec RX Delayed Packets	RC	235
0x04350 + 4*n (n=0...3)	LSECRXLATE[n]	LinkSec RX Late Packets	RC	235
0x04360 + 4*n (n=0...7)	LSECRXOK[n]	LinkSec RX Packet OK	RC	236
0x043A0 + 4*n (n=0...7)	LSECRXINV[n]	LinkSec Check RX Invalid	RC	236
0x04380 + 4*n [n=0...7]	LSECRXNV[n]	LinkSec RX Not valid count	RC	236
0x043C0 + 4*n (n=0...3)	LSECRXNUSA[n]	LinkSec RX Not Using SA	RC	236
0x043D0 + 4*n (n=0...3)	LSECRXUNSA[n]	LinkSec RX Unused SA	RC	236

11.1.2.2 General Register Descriptions

11.1.2.2.1 Device Control Register—CTRL (0x00000; RW)

Bit	Type	Reset	Description
0	RW/SN	1	Full Duplex (FD). 0 – half duplex 1 – full duplex. Controls the MAC duplex setting when explicitly setting by software. Loaded from the NVM word 13h.
1	RO	0	Reserved. Write as 0 for future compatibility
2	RW	0	Master Disable. When set, the LAN Controller blocks new master requests on the PCI device. Once no master requests are pending by this function, the <i>Master Enable Status</i> bit is cleared.
6:3	RO	0	Reserved. Write as 0 for future compatibility.
7	RO	0	Reserved. Must be set to '0' (was Invert Loss-of-Signal in predecessors).
9:8	RW	10b	Speed selection (SPEED). These bits may determine the speed configuration and are written by software after reading the PHY configuration through the MDIO interface. These signals are ignored when Auto-Speed Detection is enabled. (00)b – 10Mb/s (01)b – 100Mb/s (10)b – 1000Mb/s (11)b – not used



10	RO	0	Reserved. Write as 0 for future compatibility.
11	RW/SN	0	Force Speed (FRCSPLD). This bit is set when software wants to manually configure the MAC speed settings according to the SPEED bits above. When using a PHY device, note that the PHY device must resolve to the same speed configuration or software must manually set it to the same speed as the MAC. The value is loaded from word 13h in the NVM. Note that this bit is superseded by the CTRL_EXT.SPD_BYPS bit which has a similar function.
12	RW	0	Force Duplex (FRCDPLX). When set to 1, software may override the duplex indication from the PHY that is indicated in the FDX to the MAC. Otherwise, the duplex setting is sampled from the PHY FDX indication into the MAC on the asserting edge of the PHY LINK signal. When asserted, the CTRL.FD bit sets duplex.
14:13	RO	0	Reserved.
15	RO	0	Reserved. Reads as 0.
16	RW	0	LANPHYPC override When set to 1 this bit provides the SW driver the ability to control the LANPHYPC pin value.
17	RW	0	LANPHYPC Value When LANPHYPC override is set to 1 this bit will define the value of the LANPHYPC pin.
18	RO	0	Reserved.
19	RW	0	
23:20	RO	0	Reserved.
24	RW	0	LCD Power Down (LCDPD). When the bit is cleared to '0', the LCD power down setting is controlled by the internal logic of the LAN controller. When set to '1' and the CTRL_EXT.PHYPDEN is set as well, the LAN controller sets the external LCD to power down mode using the LANPHYPC.
25	RW	0	Host to ME Interrupt (H2MEINT). Setting this bit asserts the Host interrupt to ME. This bit is self-clearing.
26	RW/V	0	Host Software Reset (SWRST). This bit performs a reset to the PCI data path and the relevant shared logic (see). Writing 1 initiates the reset. This bit is self-clearing.
27	RW	0	Receive Flow Control Enable (RFCE). Indicates the device will respond to the reception of flow control packets. If Auto-Negotiation is enabled, this bit is set to the negotiated duplex value.
28	RW	0	Transmit Flow Control Enable (TFCE). Indicates the device will transmit flow control packets (XON & XOFF frames) based on receiver fullness. If Auto-Negotiation is enabled, this bit is set to the negotiated duplex value.
29	RO	0	Reserved
30	RW	0	VLAN Mode Enable (VME). When set to 1, all packets transmitted from LAN Controller that have VLE set are sent with an 802.1Q header added to the packet. The contents of the header come from the transmit descriptor and from the VLAN type register. On receive, VLAN information is stripped from 802.1Q packets. See Table 63 for more details.
31	RW/V	0	LAN Connected Device Reset (LCD_RST). Controls a 0 – normal (operational) 1 – reset to PHY is asserted. The LCD_RST functionality is gated by the FWSM.RSPCIPHY bit. If the FWSM.RSPCIPHY bit is not set to '1', then setting the LCD_RST has no impact. For proper operation Software or Firmware must also set the SWRST bit in the register at the same time. This bit is self-clearing.

Note: Fields loaded from the NVM are set by the NVM only if the signature bits of the NVM's Initialization Control Word match (01)b.



This register, as well as the Extended Device Control register (CTRL_EXT), controls the major operational modes for the device. While software writes to this register to control device settings, several bits (such as FD and SPEED) may be overridden depending on other bit settings and the resultant link configuration is determined by the PHY's Auto-Negotiation resolution.

The FD (duplex) and SPEED configurations of the device are normally determined from the link configuration process. Software may specifically override/set these MAC settings via certain bits in a forced-link scenario; if so, the values used to configure the MAC must be consistent with the PHY settings.

Manual link configuration is controlled through the PHY's MII management interface.

Host Software Reset (bit 26), may be used to globally reset the entire host data path and shared logic with the PCIm. This register is provided primarily as a last-ditch software mechanism to recover from an indeterminate or suspected hung hardware state. Most registers (receive, transmit, interrupt, statistics, etc.), and state machines will be set to their power-on reset values, approximating the state following a power-on or PCI reset. One internal configuration register, the Packet Buffer Allocation (PBA) register, retains its value through a Software reset.

Note: To ensure that global device reset has fully completed and that the controller will respond to subsequent accesses, one must wait approximately 1 microsecond after setting before attempting to check to see if the bit has cleared or to access (read or write) any other device register.

Note: This register's address is reflected also at address 0x00004 for legacy reasons. Neither the SW driver nor the FW should use it since it may be unsupported in next generations.

11.1.2.2.2 Device Status Register—STATUS (0x00008; RO)

Bits	Attribute	Reset	Description
0	RO/V	X	Full Duplex (FD). 0 – half duplex 1 – full duplex Reflects duplex setting of the MAC and/or link.
1	RO/V	X	Link up (LU). 0 – no link established 1 – link established For this to be valid, the <i>Set Link Up</i> bit of the <i>Device Control Register</i> (CTRL.SU) must be set.
3,2	RO/V	00b	PHY Type Indication (PHYTYPE). Indicates the LAN Connected Device attached to Intel® Ethernet Controller I217 and resulted mode of operation of the MAC/LAN Connected Device Link buses. PHYTYPE PHY-Device 00 I217 01 Reserved 10 11 Reserved This field is loaded from the Shared Init control word in the NVM.
4	RO/V	X	Transmission Paused (TXOFF). Indication of pause state of the transmit function when symmetrical flow control is enabled.



5	RO/V	1	PHY Power Up not (PHYPWR). RO bit that indicates the power state of the PHY. 0 – The PHY is powered on in the active state. 1 – The PHY is in the power down state The PHYPWR bit is valid only after PHY reset is asserted. Note: The PHY Power Up indication reflects the status of the LANPHYPC signaling to the LCD.
7:6	RO/V	X	Link speed setting (SPEED). This bit reflects the speed setting of the MAC and/or link. (00)b – 10Mb/s (01)b – 100Mb/s (10)b – 1000Mb/s (11)b – 1000Mb/s
8	RO/V	X	Master Read Completions Blocked. This bit is set when the device receives a completion with an error (EP = 1 or status!= successful); It is cleared on PCI reset.
9	RW/V/C	0	LAN Init Done. This bit is Asserted following completion of the LAN initialization from the FLASH. See “LAN Init Done Event” section for a complete description. Software is expected to clear this field to make it usable for the next Init done event.
10	RW/V/C	1	PHY Reset Asserted (PHYRA). This bit is R/W. The HW sets this bit following the assertion of LCD reset (either HW or in-band). The bit is cleared on writing '0' to it. This bit may be used by SW as an indication that the ME initiated an LCD_RST.
18:1	RO	0	Reserved.
19	RO/V	1	Master Enable Status. Cleared by LAN Controller when the <i>Master Disable</i> bit is set and no master requests are pending by this function, otherwise this bit is set. This bit indicates that no master requests will be issued by this function as long as the <i>Master Disable</i> bit is set.
29:20	RO	0	Reserved. Reads as 0.
30	RO	0	PCI _m Function state indication 0b DMon 1b DMoff
31	RO/SN	1	Clock Control ¼ (CLK_CNT_1_4). This bit is loaded from the NVM word 13h and indicates the device supports lowering its DMA clock to ¼ of its value.

FD reflects the actual MAC duplex configuration. This normally reflects the duplex setting for the entire link, as it normally reflects the duplex configuration negotiated between the PHY and link partner (copper link) or MAC and link partner (fiber link).

Link up provides a useful indication of whether something is attached to the port. Successful negotiation of features/link parameters results in link activity. The link startup process (and consequently the duration for this activity after reset) may be several 100's of ms. It reflects whether the PHY's LINK indication is present. Refer to section for more details.

TXOFF indicates the state of the transmit function when symmetrical flow control has been enabled and negotiated with the link partner. This bit is set to 1 when transmission is paused due to the reception of an XOFF frame. It is cleared upon expiration of the pause timer or the receipt of an XON frame.

SPEED indicates the actual MAC speed configuration. These bits normally reflect the speed of the actual link, negotiated by the PHY and link partner, and reflected internally from the PHY to the MAC. These bits may represent the speed configuration of the MAC only, if the MAC speed setting has been forced via software (CTRL.SPEED). Speed indications are mapped as shown below:



(00)b – 10Mb/s

(01)b – 100Mb/s

(10)b – 1000Mb/s

(11)b – 1000Mb/s

11.1.2.2.3 Strapping Option Register—STRAP (0x0000C; RO)

This register reflects the values of the soft strapping options fetched from the NVM descriptor in the CH space. These signals are sampled by the LAN controller following PCIIm reset de-assertion that follows LAN power good or global reset (assertion of both PCI and PCIIm Resets).

Attribute	Bit(s)	Initial Value	Description
RO	0	1	Reserved.
RO	5:1	0	LAN NVM Size (NVMS). LAN NVM Space size is indicated in multiples of 4KB. LAN NVM size may vary from 4KB to 128KB while 0 value means 4KB.
RO	13:12	0	SMT Frequency
RO	15:14	0	LCD SMT Connection PHYCON
RO	16	0	LC SMBus address enable (LCSMBADDEN)
RO	23:17	0	LC SMBus address (LCSMBADD)
RO	24	0	LCD SMBus address enable (LCDSMBADDEN)
RO	31:25	0	LCD SMBus address (LCDSMBADD)

11.1.2.2.4 Extended Device Control Register—CTRL_EXT (0x00018; RW)

Bits	Type	Reset	Description
1:0	RO	0	NVM Valid Sector (NVMVS). These two bits reflects which of the LAN Controller NVM sectors is valid: 00b: No valid sector 01b: Reserved 10b: Sector 0 is valid 11b: Sector 1 is valid
2	RO	0	LCD Power Cycle Done (LPCD). This bit indicates whether LCD power cycle is done—the bit is set 50/100mSec after LANPHYPC pin assertion. SW should not override LANPHYPC when the CTRL_EXT.LPCD is zero. The bit is cleared on power good or LANPHYPC de-assertion
1:3	RO	0	Reserved.
12	RW/V	1	LinkSec Clock Gate (LSecCK). When cleared, the LinkSec logic gets its clocks. When the LSecCK is set the LinkSec logic (including all its CSR registers) do not get any clocks. This bit is loaded from NVM word 13h. Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '0' as an output from this register. BIOS reading this register should always return the correct value.
14:13	RO	0	Reserved.
15	RW	0	Speed Select Bypass (SPD_BYPS). When set to 1, all speed detection mechanisms are bypassed, and the device is immediately set to the speed indicated by CTRL.SPEED. This provides a method for software to have full control of the speed settings of the device when the change takes place by overriding the hardware clock switching circuitry.
18:16	RO	0	



19	RW/SN	0	Dynamic Clock Gating (DynCK) . When set, this bit enables dynamic clock gating of the DMA and MAC units. Also see the description of the DynWakeCK in this register. The bit is loaded from NVM word 13h. Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '0' as an output from this register. BIOS reading this register should always return the correct value.
20	RW/SN	1	PHY Power Down Enable (PHYPDEN) . When set, this bit enables the LCD to enter a low-power state when the LAN controller is at the DMoff / D3 or Dr with no WoL. This bit is loaded from word 13h in the NVM.
24:21	RO	0	Reserved.
25	RW	0	DMA Clock Control (DMACKCTL) . Controls the DMA clock source in none GbE mode (10/100 and no Link). In GbE mode the DMA clock source is always Kumeran PLL divided by 2. In nominal operation this bit should be in the default state in which the DMA clock source in none GbE is mosc_clk. In test mode the DMACKCTL and PLLGateDis should be set to 1 and CLK_CNT_1_4 in the NVM should not be set. In this mode the DMA clock source is Kumeran PLL divided by 2.
26	RW	0	Disable Static Kumeran PLL Gating (PLLGateDis) . By default the PLL is functional only when Kumeran link is required, and inactive when it is not required (at non GbE mode if Jordan is available). When set to 1 the Kumeran PLL is always active.
27	RW	0	IAME . When the IAM (interrupt acknowledge auto-mask enable) bit is set, a read or write to the ICR register will have the side effect of writing the value in the IAM register to the IMC register. When this bit is 0, the feature is disabled.
28	RW	0	Driver loaded (DRV_LOAD) . This bit should be set by the driver after it was loaded, Cleared when the driver unloads or soft reset. The MNG controller loads this bit to indicate to the manageability controller that the driver has loaded.
29	RW	0	INT_TIMERS_CLEAR_ENA . When set this bit enables the clear of the interrupt timers following an IMS clear. In this state, successive interrupts will occur only after the timers will expire again. When clear, successive interrupts following IMS clear may happen immediately.
30	RO	0	Reserved.
31	RO	0	Reserved. Reads as 0.

This register provides extended control of device functionality beyond that provided by the Device Control register (CTRL).

Note: If software uses the EE_RST function and desires to retain current configuration information, the contents of the control registers should be read and stored by software. Control register values are changed by a read of the NVM which occurs upon assertion of the EE_RST bit.

Note: The EEPROM reset function may read configuration information out of the NVM which affects the configuration of PCI configuration space BAR settings. The changes to the BAR's are not visible unless the system is rebooted and the BIOS is allowed to re-map them.

Note: The SPD_BYPS bit performs a similar function as the CTRL.FRCSPPD bit in that the device's speed settings are determined by the value software writes to the CTRL.SPEED bits. However, with the SPD_BYPS bit asserted, the settings in CTRL.SPEED take effect immediately rather than waiting until after the device's clock switching circuitry performs the change.

11.1.2.2.5 MDI Control Register—MDIC (0x00020; RW)

Bits	Type	Reset	Description
------	------	-------	-------------



15:0	RW/V	X	Data (DATA). In a Write command, software places the data bits and the MAC shifts them out to the LAN Connected Device. In a Read command, the MAC reads these bits serially from the LAN Connected Device and software can read them from this location.
20:16	RW/V	0	LAN Connected Device Register address (REGADD). i.e., Reg 0, 1, 2, ... 31.
25:21	RW/V	0	LAN Connected Device Address (PHYADD).
27:26	RW/V	0	Op-code (OP). 01 for MDI Write 10 for MDI Read. Other values are reserved.
28	RW/V	1	Ready bit (R). Set to 1 by LAN Controller at the end of the MDI transaction (i.e., indicates a Read or Write has been completed). It should be reset to 0 by software at the same time the command is written.
29	RW/V	0	Interrupt Enable (I). When set to 1 by software, it will cause an Interrupt to be asserted to indicate the end of an MDI cycle.
30	RW/V	0	Error (E). This bit set is to 1 by the HW when it fails to complete an MDI read. Software should make sure this bit is clear (0) before making a MDI read or write command.
31	RO	0	Wait. This bit is set to 1 by HW to indicate a PCIe-SMBus transition is taking place. ME/Host should not issue new MDIC transactions while this bit is set to 1. The bit is auto cleared by HW after the transition had occurred.

This register is used by software to read or write MDI (Management Data Interface) registers in a GMII/MII LAN Connected Device.

Note:

Internal logic uses the MDIC to communicate with the LCD. All fields in these registers are indicated as "/V" since the internal logic may use them to access the LCD. Since the HW uses this register, all HW, SW and FW must use semaphore logic (the Ownership flags) before accessing the MDIC.

For an MDI Read cycle the sequence of events is as follows:

1. the CPU performs a write cycle to the MII register with:
 - Ready = 0
 - Interrupt Enable bit set to 1 or 0.
 - Op-Code = 10b (read)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = the register address of the specific register to be accessed (0 through 31)
2. the MAC applies the following sequence on the MDIO signal to the LAN Connected Device:
 - <PREAMBLE><01><10><PHYADD><REGADD><Z>

where the Z stands for the MAC tri-stating the MDIO signal.

3. the LAN Connected Device returns the following sequence on the MDIO signal:
 - <0><DATA><IDLE>
4. the MAC discards the leading bit and places the following 16 data bits in the MII register.
5. LAN Controller asserts an Interrupt indicating MDI "Done", if the Interrupt Enable bit was set.



6. LAN Controller sets the Ready bit in the MII register indicating the Read is complete.
7. the CPU may read the data from the MII register and issue a new MDI command.

For an MDI Write cycle the sequence of events is as follows:

1. the CPU performs a write cycle to the MII register with:
 - Ready = 0
 - Interrupt Enable bit set to 1 or 0.
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = the register address of the specific register to be accessed (0 through 31)
 - Data = specific Data for Desired Control of LAN Connected Device
2. the MAC applies the following sequence on the MDIO signal to the LAN Connected Device:
 - <PREAMBLE><01><01><PHYADD><REGADD><10><DATA><IDLE>
3. the LAN Controller asserts an Interrupt indicating MDI “Done” if the Interrupt Enable bit was set.
4. the LAN Controller sets the Ready bit in the MII register to indicate Step 2. has been completed.
5. the CPU may issue a new MDI command.

Note: An MDI Read or Write may take as long as 64 microseconds from the CPU Write to the Ready bit assertion.

If an invalid opcode is written by software, the MAC will not execute any accesses to the LAN Connected Device registers.

If the LAN Connected Device does not generate a zero as the second bit of the turnaround cycle for reads, the MAC will abort the access, set the E (error) bit, write 0xFFFF to the data field to indicate an error condition, and set the ready bit.

11.1.2.2.5.1 Accessing LCD Wakeup register using MDIC

A new page is defined in the LAN Connected Device to hold the wakeup register space.

When SW wants to configure the wakeup state (either read or write to these registers) the MDIO page should be set to 800 (for host accesses) or 801 (for ME accesses). While the page remains the same Wakeup register access is enabled.

While the page is set to the Wakeup page the address field is no longer translated as reg_addr (register address) but as an instruction. If the given address is in [0..15] range meaning PHY registers, the functionality remains unchanged.

There are 2 valid instructions:

Address Set – 0x11 – Wakeup space address is set for either reading or writing.

Data cycle – 0x12 – Wakeup space accesses read or write cycle.



For the LAN Connected Device, in the wake area Read cycle, the sequence of events is as follows:

Setting page 0x0800

1. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = page setting
 - DATA = 0x0800 (Wakeup page)

Address setting

2. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x11 (Address set)
 - DATA = XXXX (Address of the register to be read)

Reading a register

1. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 10b (read)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x12 (data cycle for read)
 - DATA = YYYY (Data will be valid when the ready bit is set)

For the LAN Connected Device, in the wake area Write cycle, the sequence of events is as follows:

Setting page 0x0800

1. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = page setting
 - DATA = 0x0800 (Wakeup page)

Address setting

2. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x11 (Address set)
 - DATA = XXXX (Address of the register to be read)



Writing a register

3. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x12 (data cycle for write)
 - DATA = YYYY (Data to be written to the register)

11.1.2.2.6 Future Extended NVM Register—FEXTNVM (0x00028; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 19h and bits 31:16 are loaded from the NVM word 1Ah.

Bits	Type	Reset	Description
0	RW/SN	0	Reserved
1	RW/SN	0	dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
2	RW/SN	0	wake_dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
3	RW/SN	0	gpt_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
4	RW/SN	0	mac_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM). Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '1' as an output from this register. BIOS reading this register should always return the correct value.
5	RW/SN	0	m2k_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
6	RW/SN	0	Invalid image CSUM. When cleared this bit indicates to the LAD NVM programming tools (eeupdate..) that the Image CSUM needs to be corrected. When set the CSUM is assumed to be correct.
9:7	RW/SN	0	Reserved.
10	RW/SN	0	Enable MDIO Watchdog Timer (MDIOWatchEna). When set to '0', the 100msec MDIO watchdog timer is enabled. (In Nahum1 this was EXTCNF_CTRL[1] loaded from word 0x14) Default NVM setting is '1'.
11	RW/SN	0	RESERVED
12	RW/SN	0	At '1', When in giga mode, state machine for mt_x_mac_in does not need to wait for sync between clock domain (the clock domain is the same) and the sync stage is skipped. At '0' the sync stage is not skipped. When operating in 10/100 Mbps the sync is still needed, therefore it is never skipped. Default NVM setting is '0'.
13	RW/SN	0	Reserved.
14	RW/SN	0	Reserved.
15	RW/SN	0	Reserved.
16	RW/SN	0	Reserved.
19:17	RW/SN	0	Reserved.



Bits	Type	Reset	Description
20	RW/SN	0	Disable CLK gate Enable Due to D3hot. When set it disables assertion of bb_clkgaten due to D3hot. Default NVM setting is '0'.
21	RW/SN	0	LAN Disable Mode. When set to '1', legacy flow managed by BIOS routine should be performed to disable the Lan. Otherwise, the whole flow will be managed by hardware when Lan-Disable RTC well bit is set to '1'in ICH9. Default NVM setting is '0'.
29:22	RW/SN	0	Reserved.
30	RW/SN	0	Enables assertion of "gbe_pmcmsus_powerdown_rdy_mosc" in WoL enabled configuration.
31	RW/SN	0	Enables assertion of "gbe_pmcmsus_powerdown_rdy_mosc" in D3/DMoff/no_wake configuration.

11.1.2.2.7 Future Extended Register—FEXT (0x0002C; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values.

Bits	Type	Reset	Description
2:0	RO	0	Reserved.
3	RO/V	0	LAN PHY Power Control (PHYPC). Set to 1 indicates external power to the LAN Controller is On. 0,external power is off.
4	RW	0	Reserved
5	RW	0	Reserved
6	RW	0	Reserved
7	RW	0	Reserved.
8	RW	0	HW SW CRC mismatch trigger —when set to 1, the LC generates a trigger signal whenever there is a mismatch between the software calculated CRC and the hardware calculated CRC. This feature is ignored when CRC calculation is off-loaded to HW.
9	RW	0	Write disable Ghost and DMA RAMs on CRC mismatch —when set to 1: disable any writes to the following rams in the event of CRC mismatch until reset: Ghost read PCI descriptor Ghost read PCI data The four rams in the descriptor engine The packet buffer
10	RW	0	When set to 1: Enables the data visibility of the Ghost read PCI descriptor and PCI data RAMs to the NOA.
11	RW	0	Visibility in/out read data select (1 = in) FEXT.10 must be set to 1.
12	RW	0	Visibility data/desc read ram select (1 = data) FEXT.10 must be set to 1
13	RW	0	When set to 1, the Ghost read RAMs are readable by the slave bus.
14	RW	0	Reserved
15	RW	0	Reserved
16	RW	0	Reserved
17	RW	0	Reserved
29:18	RW	0x00	Future Extended. Reserved for future setting.



Bits	Type	Reset	Description
30	RW	0	WoL Enable setting by BIOS —indicating ME FW to enable ME WoL for the case Host WoL is disabled in BIOS. 1—WoL Enabled 0—WoL Disabled
31	RW	0	WoL indication Valid —set to 1 by the BIOS to indicate the WoL indication setting in FEXT[30] is valid

11.1.2.2.8 Future Extended NVM 2—FEXTNVM2 (0x00030; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 24h and bits 31:16 are loaded from the NVM word 25h.

Bits	Type	Reset	Description
31:0	RW/SN	0	Reserved.

11.1.2.2.9 Future Extended NVM 4- FEXTNVM4 (0x00024; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 2Ah and bits 31:16 are loaded from the NVM word 2Bh.

Bits	Type	Reset	Description
2:0	RW/SN	000	Beacon duration period: 000: 10 us 001: 12 us 010: 14 us 011: 16 us 100: 18 us 101: 20 us 110: 22 us 111: 08 us
3	RW/SN	0	Wait a minimum of k0s_min_time before exiting K1 0=Wait
5:4	RW/SN	00	Delay the assertion of gotorxeleidle by: 00: 128ns 01: 256ns 10: 512ns 11: 1024ns
6	RW/SN	0	Drive rxsqlchdisable to the UAFE
7	RW/SN	0	Force the rx_valid that goes to the GBE (from the UAFE)
14:8	RW/SN	0x0	Mask the errors coming form the UAFE Bit 0 masking RxStatus = 001 Bit 1 masking RxStatus = 010 Bit 6 masking RxStatus = 111
15	RW/SN	0	Forces cppxp_lg_ckwaktat
23:16	RW/SN	0x0	Txeleidle count value



Bits	Type	Reset	Description
25:24	RW/SN	00	Delay the K1 exit: 00 = no delay 01 = 2ms 10 = 4ms 11 = 8ms
26	RW/SN	0	Forces gotorxelecidle to the UAFE
27	RW/SN	0	Enable the assertion of gotorxelecidle on K1
28	RW/SN	0	Force txdetectrx_or_loopback to the UAFE
29	RW/SN	0	Enable PLL shut on 1000Mbps link up. When set to 0 and 1000Mbps link is up the LAN Controller will not approve PLL shut in K1, when set to 1 PLL shut on K1 will not be gated in 1000Mbps speed. This bit is loaded from NVM word 2Bh bit 13
31:	RW/SN	000	Reserved

11.1.2.2.10 Future Extended NVM 5—FEXTNVM5 (0x00014; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 2Ch and bits 31:16 are loaded from the NVM word 2Dh.

Bits	Type	Reset	Description
0	RW/SN	0	Disable CB CLK gate. When set it disables the extended gating of Tx/Rx circuit breaker clocks. Default NVM setting is '0'. Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '1' as an output from this register. BIOS reading this register should always return the correct value.
1	RW/SN	0	Disable Ghost PCI CLK gate. When set it disables the extended gating of Ghost PCI clock. Default NVM setting is '0'. Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '1' as an output from this register. BIOS reading this register should always return the correct value.
2	RW/SN	0	Enable Ghost PCIm CLK gate. When set it enables the extended gating of Ghost PCIm clock. Default NVM setting is '0'. Note: If Dynamic Clock Gating Disable Fuse is '1', hardware will always see '0' as an output from this register. BIOS reading this register should always return the correct value.
3	RW/SN	0	Disable MAC CLK gate. When set it disables the extended gating of the MAC clock. Default NVM setting is '0'.
6:4	RW/SN	0	LANPHYC-down duration in use. Whenever LANPHYC is de-asserted by GbE, it will remain down for at least this period of time. Itimer granularity is 1usec, default settings 0x2
7	RW/SN	0	This bit revert the lanphyc logic to use the internal GBE counter and not the PMC counter. 1—PMC counter will be used 0—GBE counter will be used
14:8	RW/SN	0	Reserved
15	RW/SN	0	PRELIMINARY_DR_EN. When set, reception of PSCHS with NS_PLTRSTB ==0 during S0 will influence just like going to DR power state, e.g. LANPHYC might be de-asserted if all the rest of the conditions are met.
16	RW/SN	0	IGNORE_PPW_AND_OEM. Ignore PPW and OEM flows while waiting for MAC-Phy flows to complete during GbE-PMC handshake.
17	RW/SN	0	IGNORE_SKU. Ignore SKU-read flow while waiting for MAC-Phy flows to complete during GbE-PMC handshake.



Bits	Type	Reset	Description
18	RW/SN	0	IGNORE_LCD_RST. Ignore in-band LCD reset sending while waiting for MAC-Phy flows to complete during GbE-PMC handshake.
19	RW/SN	0	IGNORE_MDIO. Ignore OEM completion while waiting for MAC-Phy flows to complete during GbE-PMC handshake.
20	RW/SN	0	DONT_GATE_PPW. Don't gate PPW during PLTRST#/ASW_POK state transitions.
21	RW/SN	0	DONT_GATE_OEM. Don't gate OEM-bits writing during PLTRST#/ASW_POK state transitions.
22	RW/SN	0	DONT_GATE_SKU. Don't gate SKU-read during PLTRST#/ASW_POK state transitions.
23	RW/SN	0	DONT_GATE_LCD_RST. Don't gate in-band LCD reset sending during PLTRST#/ASW_POK state transitions.
24	RW/SN	0	FULL_PPW. If GbE-PMC handshake happens during PPW, allow to complete all of PPW configurations even if "Last DW of basic Phy configuration" is used.
25	RW/SN	0	PPW_IN_BASIC. If PPW is not gated, always perform PPW (after LCD reset or LANPHYPC de-assertion) before acknowledging state-transition.
26	RW/SN	0	OEM_BITS_IN_BASIC. If PPW is not gated, always perform OEM-bits writing (after LCD reset or LANPHYPC de-assertion) before acknowledging state-transition.
27	RW/SN	0	SKU_IN_BASIC. If SKU is not gated, always perform SKU-read before acknowledging state-transition.
28	RW/SN	0	SKU_ON_PCIE_ONLY. Allow SKU-read only through PCIe.
29	RW/SN	0	DIS_SEND_ACK2_PMC. Never send acknowledge to PMC—let its watch dog timer to expire.
30	RW/SN	0	SEND_ACK2_PMC_IMMEDIATELY. Send acknowledge to PMC immediately without gating any MAC-Phy flow.
31	RW/SN	0	IGNORE_ASW_POK_STATUS. Don't consider ASW_POK de-assertion as a change that should be prepared for.

11.1.2.2.11 Future Extended NVM 6—FEXTNVM6 (0x00010; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 2Eh and bits 31:16 are loaded from the NVM word 2Fh.

Bits	Type	Reset	Description
0	RW/SN	0	lcd_master_rst enable on FLR—when set, MAC's internal lcd_master_rst_n will be triggered as result of FLR. (fallback option. Setting this bit will cause MAC Kumeran i/f to be reset, without resetting Kumeran side of Phy).
1	RW/SN	0	Enable backup/forced switching of func_clk switch in S0. If set, forced switching (by strobe) is enabled. (fallback option. Setting might cause func_clk to be sourced by mbb_clk for a couple of strobes, while SPXB i/f is already active, if switch-select toggled back and forth too quickly).
2	RW/SN	0	When set the fix in m2s_tx_in is disabled. The fix prevents m2s_tx_in parameters sent to m2s_tx_out to be reset as result of bb_rst de-assertion and it allows the MBB VDM that started to be sent to SMT to be accomplished properly without being chopped on Sx'S0 transition.
3	RW/SN	0	Disable sending VMD request to SMT (over MBB) in S0. (Setting this bit might cause gbe_mbu_req3 on MBB to be de-asserted as soon as PCI reset is de-asserted).



Bits	Type	Reset	Description
5:4	RW/SN	0	Selection of NOA signals sent to PMC cluster. Can be used for MAC debug for the cases its signals need to be observable right after GbE power ungating)
7:6	RW/SN	0	Reserved
8	RW/SN	0	Enable assertion of cppxp_bbc_ckreq. In LPT-H it is not needed because bb_clk is already controlled by bbclk_gateen.
9	RW/SN	0	Enable K1 entry condition to get into account the current link speed. In 100Mbps/10Mbps k0s_entry_ltncy is multiplied by 16/256 respectively. (In LPT-H A0 there is a bug and this bit should not be set) HSD 3715410
12:10	RW/SN	0	Reserved.
13	RW/SN	0	Invert polarity of PCIe lane. (driven as gbe_sp_rxpolarity) HSD 3715158
31:14	RW/SN	0	Reserved


K FSM Control

11.1.2.2.12 Device and Bus Number—BUSNUM (0x00038; RO)

Bit	Type	Reset	Description
7:0	RO	0x00	Reserved.
10:8	RO	000b	Function Number. The LAN controller is a single PCI function being function 0.
15:11	RO	0x19	Device Number. During nominal operation the LAN controller has a predefined Device number equal to 25 (0x19).
23:16	RO	0x00	Bus Number. The LAN controller captures its bus number during host configuration write cycles type 0 aimed at the device. This field is initialized by LAN power good reset, PCI reset and D3 to D0 transition.
31:24	RO	0x00	Reserved.

11.1.2.2.13 Latency Tolerance Reporting Values—LTRV (0x000F8; RW)

Bit Location	Attribute	Default Value	Description
9:0	RW	0x0	GbE Snoop Latency Value (GBE_SNOOP_VAL) GbE snoop latency tolerance is this value multiplied by GBE_SNOOP_SCALE
12:10	RW	0x0	GbE Snoop Latency Scale (GBE_SNOOP_SCALE). Specifies the scale for the value reported in GBE_SNOOP_VAL field Encoding: 000 – Value times 1 ns 001 – Value times 32 ns 010 – Value times 1,024 ns 011 – Value times 32,768 ns 100 – Value times 1,048,576 ns 101 – Value times 33,554,432 ns 110-111 – Not Permitted
14:13	RW	0x0	Reserved



Bit Location	Attribute	Default Value	Description
15	RW	0x0	GbE Snoop Latency Requirement (GBE_SNOOP_REQ) If this bit is set to 1 than GbE's snoop latency tolerance is GBE_SNOOP_VAL multiplied by GBE_SNOOP_SCALE. If this bit is 0 than GbE has no snoop latency requirement (i.e. infinite snoop latency)
25:16	RW	0x0	GbE Non-Snoop Latency Value (GBE_NONSNOOP_VAL) GbE non snoop latency tolerance is this value multiplied by GBE_NON_SNOOP_SCALE
28:26	RW	0x0	GbE Non-Snoop Latency Scale (GBE_NONSNOOP_SCALE). Specifies the scale for the value reported in GBE_NON_SNOOP_VAL field Encoding: 000 – Value times 1 ns 001 – Value times 32 ns 010 – Value times 1,024 ns 011 – Value times 32,768 ns 100 – Value times 1,048,576 ns 101 – Value times 33,554,432 ns 110-111 – Not Permitted
29	RW	0x0	Reserved
30	RW	0x0	Send message (Send). When set to 1 bit the driver the LAN Controller will send a message to PMC through SB IOSF with the latency tolerance values as indicated in LTRV This bit is auto cleared by the LAN Controller after the SB IOSF message was sent
31	RW	0x0	GbE Non-Snoop Latency Requirement (GBE_NONSNOOP_REQ) If this bit is set to 1 than GbE's non snoop latency tolerance is GBE_NONSNOOP_VAL multiplied by GBE_NONSNOOP_SCALE. If this bit is 0 than GbE has no non snoop latency requirement (i.e. infinite non snoop latency)

11.1.2.2.14 Low Power IDLE Control—LPIC (0x000FC; RW)

Bit	Type	Reset	Description
15:0	RW	0x11	Low Power IDLE Tw_stsyem_tx Value (LPITXV) . Value to be used to back pressure the transmit start after low power idle exit in 1uSec granularity.
16	RW	0	1000Enable —Enable EEE on 1Gbps link speed. This bit auto clears on link down.
17	RW	0	100Enable —Enable EEE on 100Mbps link speed. This bit auto clears on link down.
18	RW	0	Force LPI entry —When set to 1 by SW the LAN Controller will send the LPI entry in-band message to the LAN Connected Device
23:19	RO	0x0	Reserved
31:24	RW	0x11	Low Power IDLE Entry Time (LPIET) . Value to be used to count the minimum time in between Tx LPI entry request to Tx LPI exit request counted in 1uSec granularity

Note: According to the EEE specification the device should not request LPI entry for 1 second after link up indication. This implies that the 100/1000Enable bits are self cleared on link down indication and should be set to enable by the driver/ME after link up indication.



11.1.2.2.15 Flow Control Transmit Timer Value—FCTTV (0x00170; RW)

Bit	Type	Reset	Description
15:0	RW	X	Transmit Timer Value (TTV) . to be included in XOFF frame.
31:16	RO	0	Reserved . Read as 0. Should be written to 0 for future compatibility.

The 16-bit value in the TTV field is inserted into a transmitted frame (either XOFF frames or any PAUSE frame value in any software transmitted packets). It counts in units of slot time. If software wishes to send an XON frame, it must set TTV to 0 prior to initiating the PAUSE frame.

Note: The LAN Controller uses a fixed slot time value of 64 byte times.

11.1.2.2.16 Flow Control Refresh Threshold Value—FCRTV (0x05F40; RW)

Bit	Type	Reset	Description
15:0	RW	X	Flow Control Refresh Threshold (FCRT) . This value indicates the threshold value of the flow control shadow counter. When the counter reaches this value, and the conditions for a pause state are still valid (buffer fullness above low threshold value), a pause (XOFF) frame is sent to the link partner. The FCRTV timer count interval is the same as other flow control timers and counts at slot times of 64 byte times. If this field contains a zero value, the Flow Control Refresh is disabled.
31:16	RO	0	Reserved . Read as 0. Should be written to 0 for future compatibility.

11.1.2.2.17 Extended Configuration Control—EXTCNF_CTRL (0x00F00; RW)

Bit	Type	Reset	Description
0	RW/SN	0	LCD Write Enable . When set, enables the Extended LAN Connected Device Configuration area in the LAN Controller. When disabled, the Extended LAN Connected Device Configuration area is ignored. Loaded from NVM word 14h.
1	RW/SN	0	Reserved
2	RW/SN	0	Reserved
3	RW/SN	1	OEM Write Enable . When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. Loaded from NVM word 14h.
4	RO	0	Reserved .
5	RW/V	0	SW Semaphore FLAG (SWFLAG) . This bit is set by the device driver to gain access permission to shared CSR registers with the firmware and hardware The bit is initialized on power-up PCI reset and software reset.
6	RO/V	0	MDIO HW Ownership . HW requests access to MDIO. Part of the arbitration scheme for MDIO access (see Section 2.2.2.1.1). This is a RO bit.
7	RW/V	0	Gate Phy Configuration (PPW, SKU read, OEM configuration).
15:8	RO	0x0	Reserved .
27:16	RW/SN	0x001	Extended Configuration Pointer . Defines the base address (in DW) of the Extended Configuration area in the NVM.
31:28	RW	0	Reserved .



11.1.2.2.18 Extended Configuration Size—EXTCNF_SIZE (0x00F08; RW)

Bit	Type	Reset	Description
31:24	RO	0x0	Reserved.
23:16	RW/SN	0x0	Extended LCD Length. Size (in DW) of the Extended LAN Connected Device Configuration area loaded from Extended Configuration word 2 in the NVM. If an extended configuration area is disabled by "LCD Write Enable" field in word 14h in the NVM, this length must be set to zero.
15:8	RW/SN	0x0	Reserved
7:0	RW/SN	0x0	Reserved.

11.1.2.2.19 PHY Control Register—PHY_CTRL (0x00F10; RW) This register is initialized to the HW default at LAN PWR Good reset.

Bit	Type	Reset	Description
31:29	RO	0x0	Reserved
28:25	RO	0x0	SKU read data. These four bits contain the SKU value read from the LAN Connected Device SKU register. Using these bits, the SKU mechanism determines the Device ID.
24	RO	0x0	Reserved.
23	RO	0x0	SKU done. This bit indicates the termination of SKU read.
22	RW	0x0	Reserved.
21	RW	0	Reserved
20	RW	0	Reserved.
19:17	RW	0x2	Reserved. Used to be: Periodic Power Down Time (PD Time). Defines the power down time in the periodic power down scheme as follows: Delay Time = 100msec x (PD Time + 1), Default = 300msec.
16	RW	0	Reserved.
15:8	RO	0x0	Reserved
7	RW/SN	0	B2B Ena. Enables SPD in Back To Back link setup. Bit is initialized by word 17h bit 15 in the NVM.
6	RW/SN	0	Global GbE Disable. Prevents PHY auto negotiating 1000 Mb/s link in all power states (including D0a). Bit is initialized by word 17h bit 14 in the NVM.
5:4	RO	00b	Reserved.
3	RW/SN	1	GbE Disable at non D0a. Prevents PHY from auto negotiating 1000Mb/s link in all power states except D0a (DR, D0u and D3). Bit is initialized by word 17h bit 11 in the NVM. This bit must be set since GbE is not supported in Sx by the platform.
2	RW/SN	1	LPLU in non D0a. Enables PHY to negotiate for slowest possible link (Reverse AN) in all power states except D0a (DR, D0u and D3). Bit is initialized by word 17h bit 10 in the NVM.
1	RW/SN	0	LPLU in D0a. Enables PHY to negotiate for slowest possible link (Reverse AN) in all power states (including D0a). This bit overrides the LPLU in non D0a bit. Bit is initialized by word 17h bit 9 in the NVM.
0	RW/SN	0	SPD Ena. Enables PHY Smart Power Down mode. Bit is initialized by word 17h bit 8 in the NVM.



11.1.2.2.20 PCIE Analog Configuration—PCIEANACFG (0x00F18; RW)

Bit	Type	Reset	Description
3:0	RW/SN	0x0	Reserved.
4	RW/SN	0	Reserved.
5	RW/SN	0	Reserved.
31:6	RO	0	Reserved. Read as 0 ignore on write.

Packet Buffer Allocation—PBA (0x01000; RW) This register sets the on-chip receive and transmit storage allocation ratio.

Note: Programming this register does not automatically re-load or initialize internal packet-buffer RAM pointers. The software must reset both transmit and receive operation (using the global device reset CTRL.SWRST bit) after changing this register in order for it to take effect. The PBA register itself will not be reset by assertion of the Software reset, but will only be reset upon initial hardware power-on.

Note: If Early Receive functionality is not enabled (indicate field/register), the Receive packet buffer should be larger than the max expected received packet + 32B.

Note: For best performance the transmit buffer allocation should be set to accept two full sized packets.

Note: Transmit Packet Buffer size should be configured to be more than 4KB.

Packet Buffer Size—PBS (0x01008; RW) This register sets the on-chip receive and transmit storage allocation size, The allocation value is read/write for the lower 6 bits. The division between transmit and receive is done according to the PBA register.

Note: Programming this register does not automatically re-load or initialize internal packet-buffer RAM pointers. The software must reset both transmit and receive operation (using the global device reset CTRL.SWRST bit) after changing this register in order for it to take effect. The PBS register itself will not be reset by assertion of the Software reset, but will only be reset upon initial hardware power-on.

Note: Programming this register should be aligned with programming the PBA register HW operation, if PBA and PBS are not coordinated is not determined.

11.1.2.2.21 DMA Control Register—DCR (0x05B00; RW)

Reserved—Interrupt Register Descriptions

11.1.2.2.22 Interrupt Cause Read Register—ICR (0x000C0; RC/WC)

This register is Read-Clear or Write-Clear (see details after the tables). If enabled, read access also clears the ICR content after it is posted to the SW. Otherwise a Write cycle is required to clear the relevant bit fields. Write a '1' clears the written bit while writing '0' has no affect (with the exception of the INT_ASSERTED bit as detailed below).



Bit	Type	Reset	Description
0	RWC/CR/V	0	Transmit Descriptor Written Back (TXDW) . Set when hardware processes a descriptor with either RS set. If using delayed interrupts (IDE set), the interrupt is delayed until after one of the delayed-timers (TIDV or TADV) expires.
1	RWC/CR/V	0	Transmit Queue Empty (TXQE) . Set when, the last descriptor block for a transmit queue has been used. When configured to use more than one transmit queue this interrupt indication will be issued if one of the queues is empty and will not be cleared until all the queues have valid descriptors.
2	RWC/CR/V	0	Link Status Change (LSC) . This bit is set whenever the link status changes (either from up to down, or from down to up). This bit is affected by the LINK indication from the PHY.
3	RO	0	Reserved.
4	RWC/CR/V	0	Receive Descriptor Minimum Threshold hit (RXDMTO) . Indicates that the minimum number of receive descriptors RCTL.RDMTS are available and software should load more receive descriptors.
5	RWC/CR/V	0	Disable SW Write Access (DSW) . The DSW bit indicates that the FW changed the status of the DISSW or the DISSWLNK bits in the FWSM register.
6	RWC/CR/V	0	Receiver Overrun (RXO) . Set on receive data FIFO overrun. Could be caused either because there are no available buffers or because receive bandwidth is inadequate.
7	RWC/CR/V	0	Receiver Timer Interrupt (RXT0) . Set when the timer expires.
8		0	Reserved.
9	RWC/CR/V	0	MDIO Access Complete (MDAC) . Set when the MDIO access is completed. See Section 11.1.2.2.5 for details.
11:10	RO	0	Reserved.
12	RWC/CR/V	0	PHY Interrupt (PHYINT) . Set when the LAN Connected Device generates an interrupt.
13	RO	0	Reserved.
14	RWC/CR/V	0	LinkSec Packet Number (LSECPN) . The Tx Packet Number hit the "PN exhaustion threshold" as defined in the LSECTXCTRL register and the host is the KaY.
15	RWC/CR/V	0	Transmit Descriptor Low Threshold hit (TXD_LOW) . Indicates that the descriptor ring has reached the threshold specified in the Transmit Descriptor Control register.
16	RWC/CR/V	0	Small Receive Packet Detected (SRPD) . Indicates that a packet of size < RSRPD.SIZE register has been detected and transferred to host memory. The interrupt is only asserted if RSRPD.SIZE register has a non-zero value.
17	RWC/CR/V	0	Receive ACK Frame Detected (ACK) . Indicates that an ACK frame has been received and the timer in RAID.ACK_DELAY has expired.
18	RWC/CR/V	0	Manageability Event Detected (MNG) . Indicates that a Manageability read the H2ME_S register. Can be set also by setting the MNG bit in the ICS register.
19	RWC/CR/V	0	Reserved.
20	RWC/CR/V	0	ME Reset (EPRST) . This bit is set when PCIm Reset or PCIm SW Reset has occurred.
21	RWC/CR/V	0	Reserved.
22	RWC/CR/V	0	ECC Error (ECCER) . Indicates an uncorrectable EEC error had occurred
30:23	RO	0	Reserved. Reads as 0.
31	RWC/CR/V	0	Interrupt Asserted (INT_ASSERTED) . This bit is set when the LAN port has a pending interrupt. If the Interrupt is enabled in the PCI configuration space an Interrupt is asserted.



This register contains all interrupt conditions for LAN Controller. Whenever an interrupt causing event occurs, the corresponding interrupt bit is set in this register. An interrupt is generated whenever one of the bits in this register is set, and the corresponding interrupt is enabled via the Interrupt Mask Set/Read Register (see Section 11.1.2.3.5).

Whenever an interrupt causing event occurs, all timers of delayed interrupts are cleared and their cause event is set in the ICR.

- Read ICR register is affected differently in the following cases:
 - Case 1—Interrupt Mask Register equals 0x0000 (mask all). ICR content will be cleared.
 - Case 2—Interrupt was asserted (ICR.INT_ASSERTED=1). ICR content will be cleared and Auto Mask is active, meaning, the IAM register is written to the IMC register.
 - Case 3—Interrupt was not asserted (ICR.INT_ASSERTED=0). Read has no side affect.

Writing a 1 to any bit in the register will also clear that bit. Writing a 0 to any bit will have no effect on that bit. The INT_ASSERTED bit is a special case. Writing a 1 or 0 to this bit has no affect. It is cleared only when all interrupt sources are cleared.

11.1.2.2.23 Interrupt Throttling Register—ITR (0x000C4; RW)

Bit	Type	Reset	Description
19:0	RW	0	INTERVAL. Minimum inter-interrupt interval. The interval is specified in 256ns units. Zero disables interrupt throttling logic.
31:20	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Software can use this register to pace (or even out) the delivery of interrupts to the host CPU. This register provides a guaranteed inter-interrupt delay between interrupts asserted by the network controller, regardless of network traffic conditions. To independently validate configuration settings, software can use the following algorithm to convert the inter-interrupt interval value to the common 'interrupts/sec' performance metric:

$$\text{Interrupts/sec} = (256 \times 10^{-9}\text{sec} \times \text{interval})^{-1}$$

For example, if the interval is programmed to 500d, the network controller guarantees the CPU will not be interrupted by the network controller for 128 microseconds from the last interrupt.

Inversely, inter-interrupt interval value can be calculated as:

$$\text{Inter-interrupt interval} = (256 \times 10^{-9}\text{sec} \times \text{interrupts/sec})^{-1}$$

The optimal performance setting for this register is very system and configuration specific. An initial suggested range for the interval value is 65-5580 (28B-15CC).

Note: When working at 10/100Mbps and running at ¼ clock the interval time is multiplied by four.



11.1.2.2.24 Interrupt Cause Set Register—ICS (0x000C8; WO)

Bit	Type	Reset	Description
0	WO	X	TXDW. Sets Transmit Descriptor Written Back.
1	WO	X	TXQE. Sets Transmit Queue Empty.
2	WO	X	LSC. Sets Link Status Change.
3	RO	X	Reserved.
4	WO	X	RXDMT. Sets Receive Descriptor Minimum Threshold hit.
5	WO	X	DSW. Sets Block SW Write accesses.
6	WO	X	RXO. Sets Receiver Overrun. Set on receive data FIFO overrun.
7	WO	X	RXT. Sets Receiver Timer Interrupt.
8		X	Reserved.
9	WO	X	MDAC. Sets MDIO Access Complete Interrupt.
11:10	RO	X	Reserved.
12	WO	X	PHYINT. Sets PHY Interrupt.
13	RO	X	Reserved.
14	WO	X	LSECPN. Sets LinkSec Packet Number Interrupt.
15	WO	X	TXD_LOW. Transmit Descriptor Low Threshold Hit.
16	WO	X	SRPD. Small Receive Packet Detected and Transferred.
17	WO	X	ACK. Set Receive ACK frame detected.
18	WO	X	MNG. Set the Manageability Event Interrupt.
19	WO	X	Reserved.
20	WO	X	EPRST. Set ME Reset event.
21	RO	X	Reserved.
22	WO	X	ECCER Set uncorrectable EEC error
31:23	RO	X	Reserved. Should be written with 0 to ensure future compatibility.

Software uses this register to set an interrupt condition. Any bit written with a “1” sets the corresponding interrupt. This results in the corresponding bit being set in the Interrupt Cause Read Register (see [Section 11.1.2.2.22](#)), and an interrupt is generated if one of the bits in this register is set, and the corresponding interrupt is enabled via the Interrupt Mask Set/Read Register (see [Section 11.1.2.2.25](#)).

Bits written with “0” are unchanged.

11.1.2.2.25 Interrupt Mask Set/Read Register—IMS (0x000D0; RW)

Bit	Type	Reset	Description
0	RWS	0	TXDW. Sets mask for Transmit Descriptor Written Back.
1	RWS	0	TXQE. Sets mask for Transmit Queue Empty.
2	RWS	0	LSC. Sets mask for Link Status Change.
3	RO	0	Reserved.
4	RWS	0	RXDMT0. Sets mask for Receive Descriptor Minimum Threshold hit.
5	RWS	0	DSW. Sets mask for Block SW Write accesses.
6	RWS	0	RXO. Sets mask for Receiver Overrun. Set on receive data FIFO overrun.



Bit	Type	Reset	Description
7	RWS	0	RXT0. Sets mask for Receiver Timer Interrupt.
8	R	0	Reserved.
9	RWS	0	MDAC. Sets mask for MDIO Access Complete Interrupt.
11:10	RO	0	Reserved.
12	RWS	0	PHYINT. Sets mask for PHY Interrupt.
13	RO	0	Reserved.
14	RWS	0	LSECPN. Sets the mask for LinkSec Packet Number Int.
15	RWS	0	TXD_LOW. Sets the mask for Transmit Descriptor Low Threshold hit.
16	RWS	0	SRPD. Sets mask for Small Receive Packet Detection.
17	RWS	0	ACK. Sets the mask for Receive ACK frame detection.
18	RWS	0	MNG. Sets mask for Manageability Event Interrupt.
19	RWS	0	Reserved.
20	RWS	0	EPRST. Sets mask for ME Reset event.
21	RO	0	Reserved.
22	RWS	0	ECCER Sets mask for uncorrectable EEC error
31:23	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Reading this register returns which bits have an interrupt mask set. An interrupt is enabled if its corresponding mask bit is set to 1, and disabled if its corresponding mask bit is set to 0. An interrupt is generated whenever one of the bits in this register is set, and the corresponding interrupt condition occurs. The occurrence of an interrupt condition is reflected by having a bit set in the Interrupt Cause Read Register (see [Section 11.1.2.2.22](#)).

A particular interrupt may be enabled by writing a 1 to the corresponding mask bit in this register. Any bits written with a 0 are unchanged.

Note: If software desires to disable a particular interrupt condition that had been previously enabled, it must write to the Interrupt Mask Clear Register (see [Section 11.1.2.2.26](#)), rather than writing a 0 to a bit in this register.

When the CTRL_EXT.INT_TIMERS_CLEAR_ENA bit is set, then following writing all 1's to the IMS register (enable all Interrupts) all interrupt timers are cleared to their initial value. This auto clear provides the required latency before the next INT event.

11.1.2.2.26 Interrupt Mask Clear Register—IMC (0x000D8; WO)

Bit	Type	Reset	Description
0	WO	0	TXDW. Clears mask for Transmit Descriptor Written Back.
1	WO	0	TXQE. Clears mask for Transmit Queue Empty.
2	WO	0	LSC. Clears mask for Link Status Change.
3	RO	0	Reserved.
4	WO	0	RXDMT0. Clears mask for Receive Descriptor Minimum Threshold hit.
5	WO	0	DSW. Clears mask for Block SW Write accesses.
6	WO	0	RXO. Clears mask for Receiver Overrun.
7	WO	0	RXT0. Clears mask for Receiver Timer Interrupt.



Bit	Type	Reset	Description
8	O	0	Reserved.
9	WO	0	MDAC. Clears mask for MDIO Access Complete Interrupt.
11:10	RO	0	Reserved. Reads as 0.
12	WO	0	PHYINT. Clears PHY Interrupt.
13	RO	0	Reserved.
14	WO	0	LSECPN. Clears the LinkSec Packet Number Interrupt.
15	WO	0	TXD_LOW. Clears the mask for Transmit Descriptor Low Threshold hit.
16	WO	0	SRPD. Clears mask for Small Receive Packet Detect Interrupt.
17	WO	0	ACK. Clears the mask for Receive ACK frame detect Interrupt.
18	WO	0	MNG. Clears mask for the Manageability Event Interrupt.
19	WO	0	Reserved.
20	WO	0	EPRST. Clears mask for the ME Reset event.
21	RO	0	Reserved.
22	WO	0	ECCER Clears the mask for uncorrectable EEC error
31:23	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Software uses this register to disable an interrupt. Interrupts are presented to the bus interface only when the mask bit is a one and the cause bit is a one. The status of the mask bit is reflected in the Interrupt Mask Set/Read Register, and the status of the cause bit is reflected in the Interrupt Cause Read register (see [Section 11.1.2.2.22](#)).

Software blocks interrupts by clearing the corresponding mask bit. This is accomplished by writing a 1 to the corresponding bit in this register. Bits written with 0 are unchanged (i.e. their mask status does not change).

In summary, the sole purpose of this register is to allow software a way to disable certain, or all, interrupts. Software disables a given interrupt by writing a 1 to the corresponding bit in this register.



11.1.2.2.27 Interrupt Acknowledge Auto-Mask—IAM (0x000E0; RW)

Bit	Type	Reset	Description
0-31	RW	0	IAM_VALUE. When the CTRL_EXT.IAME bit is set and the ICR.INT_ASSERTED=1, an ICR read or write will have the side effect of writing the contents of this register to the IMC register.

11.1.2.3 Receive Register Descriptions

11.1.2.3.1 Receive Control Register—RCTL (0x00100; RW)

Bit	Type	Reset	Description
0	RO	0	Reserved. This bit represented a hardware reset of the receive-related portion of the device in previous controllers, but is no longer applicable. Only a full device reset CTRL.SWRST is supported. Write as 0 for future compatibility.
1	RW	0	Enable (EN). The receiver is enabled when this bit is 1. Writing this bit to 0, stops reception after receipt of any in progress packets. All subsequent packets are then immediately dropped until this bit is set to 1. Note that this bit controls only DMA functionality to the host. Packets are counted by the statistics even when this bit is cleared.
2	RW	0	Store bad packets (SBP). 0 – do not store bad packets 1 – store bad packets Note that CRC errors before the SFD are ignored. Any packet must have a valid SFD in order to be recognized by the device (even bad packets). Note: Packet errors will not be routed to the MNG even if this bit is set.
3	RW	0	Unicast promiscuous enable (UPE). 0 – disabled 1 – enabled
4	RW	0	Multicast promiscuous enable (MPE). 0 – disabled 1 – enabled
5	RW	0	Long packet enable (LPE). 0 – disabled 1 – enabled
7:6	RW	00b	Reserved.
9:8	RW	0	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set whenever the fractional number of free descriptors becomes equal to RDMTS. Table 11-2 below lists which fractional values correspond to RDMTS values. See Section 11.1.2.3.9 for details regarding RDLEN.
11:10	RW	0	Descriptor Type (DTYP). 00 – Legacy or Extended descriptor type 01 – Packet Split descriptor type 10 and 11 – Reserved
13:12	RW	0	Multicast Offset (MO). This determines which bits of the incoming multicast address are used in looking up the bit vector. 00 – [47:38] 01 – [46:37] 10 – [45:36] 11 – [43:34]
14	RW	0	Reserved.
15	RW	0	Broadcast Accept Mode (BAM). 0 – ignore broadcast (unless it matches through exact or imperfect filters) 1 – accept broadcast packets.



Bit	Type	Reset	Description
17:16	RW	0	Receive Buffer Size (BSIZE). RCTL.BSEX – 0 00 – 2048 Bytes 01 – 1024 Bytes 10 – 512 Bytes 11 – 256 Bytes RCTL.BSEX – 1 00 – reserved 01 – 16384 Bytes 10 – 8192 Bytes 11 – 4096 Bytes BSIZE is only used when DTYP – 00. When DTYP – 01, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than 0, in that case, FLXBUF determines the buffer size.
21:18	RO	0	Reserved. Should be written with 0.
22	RW	0	Reserved.
23	RW	0	Pass MAC Control Frames (PMCF). 0 – do not (specially) pass MAC control frames. 1 – pass any MAC control frame (type field value of 0x8808) that does not contain the pause opcode of 0x0001.
24	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
25	RW	0	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE above). 0 – Buffer size is as defined in BSIZE 1 – Original BSIZE values are multiplied by 16.
26	RW	0	Strip Ethernet CRC from incoming packet (SECR). 0 – does not strip CRC 1 – Strips CRC The stripped CRC is not DMA'd to host memory and is not included in the length reported in the descriptor
30:27	RW	0	FLXBUF. Determines a flexible buffer size. When this field is "0000", the buffer size is determined by BSIZE. If this field is different from "0000", the receive buffer size is the number represented in Kbytes: i.e. "0001" = 1KB (1024 Bytes).
31	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

LPE controls whether long packet reception is permitted. Hardware discards long packets if LPE is 0. A long packet is one longer than 1522 bytes. If LPE is 1, the maximum packet size that the device can receive is 9018 bytes.

RDMTS{1,0} determines the threshold value for free receive descriptors according to the following table:

Table 11-2. RDMTS Values

RDMTS	Free Buffer Threshold
00	1/2
01	1/4
10	1/8
11	Reserved



BSIZE controls the size of the receive buffers and permits software to trade-off descriptor performance versus required storage space. Buffers that are 2048 bytes require only one descriptor per receive packet maximizing descriptor efficiency. Buffers that are 256 bytes maximize memory efficiency at a cost of multiple descriptors for packets longer than 256 bytes.

PMCF controls the DMA function of MAC control frames (other than flow control). A MAC control frame in this context must be addressed to either the MAC control frame multicast address or the station address, match the type field and NOT match the PAUSE opcode of 0x0001. If PMCF = 1 then frames meeting this criteria will be DMA'd to host memory.

The SECRC bit controls whether the hardware strips the Ethernet CRC from the received packet. This stripping occurs prior to any checksum calculations. The stripped CRC is not DMA'd to host memory and is not included in the length reported in the descriptor.

11.1.2.3.2 Receive Control Register 1—RCTL1 (0x00104; RW)

Bit	Type	Reset	Description
7:0	RO	0	Reserved. This bit represented a hardware reset of the receive-related portion of the device in previous controllers, but is no longer applicable. Only a full device reset CTRL.SWRST is supported. Write as 0 for future compatibility.
9:8	RW	0	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set whenever the fractional number of free descriptors becomes equal to RDMTS. Table 11-2 lists which fractional values correspond to RDMTS values. See Section 11.1.2.3.9 for details regarding RDLEN.
11:10	RW	0	Descriptor Type (DTYP). 00 – Legacy or Extended descriptor type 01 – Packet Split descriptor type 10 and 11 – Reserved The value of RCTL1.DTYP should be the same as RCTL.DTYP (same descriptor types used in both descriptor queues).
15:12	RO	0	Reserved.
17:16	RW	0	Receive Buffer Size (BSIZE). RCTL.BSEX – 0 00 – 2048 Bytes 01 – 1024 Bytes 10 – 512 Bytes 11 – 256 Bytes RCTL.BSEX – 1 00 – reserved 01 – 16384 Bytes 10 – 8192 Bytes 11 – 4096 Bytes BSIZE is only used when DTYP – 00. When DTYP – 01, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than 0, in that case, FLXBUF determines the buffer size.
24:18	RO	0	Reserved. Should be written with 0.
25	RW	0	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE above). 0 – Buffer size is as defined in BSIZE 1 – Original BSIZE values are multiplied by 16.
26	RW	0	Reserved. Should be written with 0.



30:27	RW	0	FLXBUF. Determine a flexible buffer size. When this field is "0000", the buffer size is determined by BSIZE. If this field is different from "0000", the receive buffer size is the number represented in Kbytes: i.e. "0001" – 1KB (1024 Bytes).
31	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

This register is used to configure queue1 registers when working in VMDq mode.

11.1.2.3.3 Early Receive Threshold—ERT (0x02008; RW)

Bit	Type	Reset	Description
12:0	RW	0x0000	Receive Threshold Value (RxThreshold). This threshold is in units of 8 bytes.
13	RW	0	Reserved. When SW activates the "Early Receive" mechanism (by setting the RxThreshold field to a non-zero value) it must set this bit as well.
14	RO	0	Reserved
15	RW	0	Reserved.
16	RW	0	Reserved.
17	RW	0	Reserved.
25:22	RO	0	Reserved.
26	RW	0	Reserved.
31:7	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the RxThreshold value. This threshold determines how many bytes of a given packet should be in the LAN Controller's on-chip receive packet buffer before it attempts to begin transmission of the frame on the host bus. This register allows software to configure the "early receive" mode.

This field has a granularity of 8 bytes. So, if this field is written to 0x20, which corresponds to a threshold of 256 (decimal) bytes. If the size of a given packet is smaller than the threshold value, or if this register is set to 0, then LAN Controller will start the PCI transfer only after the entire packet is contained in LAN Controller's receive packet buffer. LAN Controller examines this register on a cycle by cycle basis to determine if there is enough data to start a transfer for the given frame over the PCI bus.

Once the device acquires the bus, it will attempt to DMA all of the data collected in the internal receive packet buffer so far.

The only negative affect of setting this value too low is that it will cause additional PCI bursts for the packet. In other words, this register allows software to trade-off latency versus bus utilization. Too high a value will effectively eliminate the early receive benefits (at least for short packets) and too low a value will deteriorate PCI bus performance due to a large number of small bursts for each packet. The RUTEC statistic counts certain cases where the ERT has been set too low, and thus provides software a feedback mechanism to better tune the value of the ERT.

It should also be noted that this register will have an effect only when the receive packet buffer is nearly empty (the only data in the packet buffer is from the packet that is currently on the wire).

Note: When Early receive is used in parallel to the Packet split feature the minimum value of the ERT register should be bigger than the header size to enable the actual packet split.



Note: Early Receive should be enabled only when working in a Jumbo frames enabled environment and the ERT.RxThreshold should be set to 0xFA so that only packets bigger than 2K bytes would trigger the early receive mechanism.

11.1.2.3.4 Packet Split Receive Control Register— PSRCTL (0x02170 + n*0x4 [n=0..1]; RW)

Bit	Type	Reset	Description
6:0	RW	2	Receive Buffer Size for Buffer 0 (BSIZE0). The value is in 128 byte resolution. Value can be from 128 bytes to 16256 bytes (15.875 Kbytes). Default buffer size is 256B. SW should not program this field to a zero value.
7	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
13:8	RW	4	Receive Buffer Size for Buffer 1 (BSIZE1). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 4KB. SW should not program this field to a zero value.
15:14	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
21:16	RW	4	Receive Buffer Size for Buffer 2 (BSIZE2). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 4KB. SW may program this field to any value.
23:22	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
29:24	RW	0	Receive Buffer Size for Buffer 3 (BSIZE3). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 0KB. SW may program this field to any value.
31:30	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Note: If SW sets a buffer size to zero, all buffers following that one must be set to zero as well. Pointers in the Receive descriptors to buffers with a zero size should be set to anything but NULL pointers.

11.1.2.3.5 Flow Control Receive Threshold Low—FCRTL (0x02160; RW)

Bit	Type	Reset	Description
2:0	RO	0	Reserved. The underlying bits might not be implemented in all versions of the chip. Must be written with 0.
15:3	RW	0	Receive Threshold Low (RTL). FIFO low water mark for flow control transmission.
30:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
31	RW	0	XON Enable (XONE). 0 – disabled 1 – enabled.

This register contains the receive threshold used to determine when to send an XON packet. It counts in units of bytes. The lower 3 bits must be programmed to 0 (8 byte granularity). Software must set XONE to enable the transmission of XON frames. Whenever hardware crosses the receive high threshold (becoming more full), and then crosses the receive low threshold and XONE is enabled (= 1), hardware transmits an XON frame.

Note that flow control reception/transmission are negotiated capabilities by the Auto-Negotiation process. When the device is manually configured, flow control operation is determined by the RFCE and TFCE bits of the device control register.



11.1.2.3.6 Flow Control Receive Threshold High—FCRTH (0x02168; RW)

Bit	Type	Reset	Description
2:0	RO	0	Reserved. The underlying bits might not be implemented in all versions of the chip. Must be written with 0.
15:3	RW	0	Receive Threshold High (RTH). FIFO high water mark for flow control transmission.
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the receive threshold used to determine when to send an XOFF packet. It counts in units of bytes. This value must be at least 8 bytes less than the maximum number of bytes allocated to the Receive Packet Buffer (PBA, RXA), and the lower 3 bits must be programmed to 0 (8 byte granularity). Whenever the receive FIFO reaches the fullness indicated by RTH, hardware transmits a PAUSE frame if the transmission of flow control frames is enabled.

Note that flow control reception/transmission are negotiated capabilities by the Auto-Negotiation process. When the device is manually configured, flow control operation is determined by the RFCE and TFCE bits of the device control register.

11.1.2.3.7 Receive Descriptor Base Address Low queue—RDBAL (0x02800 + n*0x100[n=0..1]; RW)

Bit	Type	Reset	Description
3:0	RO	0	Reserved. Ignored on writes. Returns 0 on reads.
31:4	RW	X	Receive Descriptor Base Address Low (RDBAL).

This register contains the lower bits of the 64 bit descriptor base address. The lower 4 bits are always ignored. The Receive Descriptor Base Address must point to a 16B aligned block of data.

11.1.2.3.8 Receive Descriptor Base Address High queue—RDBAH (0x02804 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
31:0	RW	X	Receive Descriptor Base Address [63:32] (RDBAH).

This register contains the upper 32 bits of the 64 bit Descriptor base address.

11.1.2.3.9 Receive Descriptor Length queue—RDLEN (0x02808 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
6:0	RO	0	Reserved. Ignore on write. Reads back as 0.
19:7	RW	0	Descriptor Length (LEN)
31:20	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register sets the number of bytes allocated for descriptors in the circular descriptor buffer. It must be 128B aligned.

Note: The descriptor ring must be equal to or larger than 8 descriptors.

11.1.2.3.10 Receive Descriptor Head queue—RDH (0x02810 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
15:0	RW/V	0	Receive Descriptor Head (RDH).
31:16	RO	0	Reserved. Should be written with 0

This register contains the head pointer for the receive descriptor buffer. The register points to a 16B datum. Hardware controls the pointer. The only time that software should write to this register is after a reset (hardware reset or CTRL.SWRST) and before enabling the receive function (RCTL.EN). If software were to write to this register while the receive function was enabled, the on-chip descriptor buffers may be invalidated and the hardware could become confused.

11.1.2.3.11 Receive Descriptor Tail queue—RDT (0x02818 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
15:0	RW	0	Receive Descriptor Tail (RDT).
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the tail pointer for the receive descriptor buffer. The register points to a 16B datum. Software writes the tail register to add receive descriptors for the hardware to process.

11.1.2.3.12 Interrupt Delay Timer (Packet Timer)—RDTR (0x02820 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
15:0	RW	0	Receive Delay Timer. Receive Packet delay timer measured in increments of 1.024 microseconds.
30:16	RO	0	Reserved. Reads as 0.
31	WO	0	Flush Partial Descriptor Block (FPD). when set to 1; ignored otherwise. Reads 0.

This register is used to delay interrupt notification for the receive descriptor ring by coalescing interrupts for multiple received packets. Delaying interrupt notification helps maximize the number of receive packets serviced by a single interrupt.

This feature operates by initiating a countdown timer upon successfully receiving each packet to system memory. If a subsequent packet is received *before* the timer expires, the timer is re-initialized to the programmed value and re-starts its countdown. If the timer expires due to *not* having received a subsequent packet within the programmed interval, pending receive descriptor write backs are flushed and a receive timer interrupt is generated.

Setting the value to 0 represents no delay from a receive packet to the interrupt notification, and results in immediate interrupt notification for each received packet.

Writing this register with FPD set initiates an immediate expiration of the timer, causing a write back of any consumed receive descriptors pending write back, and results in a receive timer interrupt in the ICR.



Receive interrupts due to a Receive Absolute Timer (RADV) expiration will cancel a pending RDTR interrupt. The RDTR countdown timer is reloaded but halted, so as to avoid generation of a spurious second interrupt after the RADV has been noted, but may be restarted by a subsequent received packet.

Note: FPD is self-clearing.

11.1.2.3.13 Receive Descriptor Control—RXDCTL (0x02828 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
14	RW	0	Reserved.
15	RW	0	Reserved.
21:16	RW	0x01	Write-back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0x0	Granularity (GRAN). Units for the thresholds in this register. 0 – cache lines 1 – descriptors
31:25	RO	0x00	Reserved.

Note: This register was not fully validated. SW should set it to 0x0000 during nominal operation.

This register controls the fetching and write-back of receive descriptors. The three threshold values are used to determine when descriptors will be read from and written to host memory. The values may be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag. If GRAN=0 (specifications are in cache-line granularity), the thresholds specified (based on the cache line size specified in the PCI configuration space CLS field) must not represent greater than 31 descriptors.

Note: When (WTHRESH = 0) or (WTHRESH = 1 and GRAN = 1) only descriptors with the 'RS' bit set will be written back.

PTHRESH is used to control when a prefetch of descriptors will be considered. This threshold refers to the number of valid, unprocessed receive descriptors the chip has in its on-chip buffer. If this number drops below PTHRESH, the algorithm will consider pre-fetching descriptors from host memory. This fetch will not happen however unless there are at least HTHRESH valid descriptors in host memory to fetch.

Note: HTHRESH should be given a non zero value when ever PTHRESH is used.

WTHRESH controls the write-back of processed receive descriptors. This threshold refers to the number of receive descriptors in the on-chip buffer which are ready to be written back to host memory. In the absence of external events (explicit flushes), the write-back will occur only after at least WTHRESH descriptors are available for write-back.

Note: Possible values:

GRAN = 1 (descriptor granularity):



PTHRESH = 0...31

WTHRESH = 0...31

HTHRESH = 0...31

GRAN = 0 (cache line granularity):

PTHRESH = 0...3 (for 16 descriptors cache line—256 bytes)

WTHRESH = 0...3

HTHRESH = 0...4

Note: For any WTHRESH value other than 0—The packet and absolute timers must get a non zero value for WTHRESH feature to take affect.

Note: Since the default value for write-back threshold is 1, the descriptors are normally written back as soon as one cache line is available. WTHRESH must contain a non-zero value to take advantage of the write-back bursting capabilities of LAN Controller.

Note: RXDCTL1 is only accessible when VMDq is enabled (MRQC.MRxQueue = 10).

11.1.2.3.14 Receive Interrupt Absolute Delay Timer- RADV (0x0282C; RW)

Bits	Type	Reset	Description
15:0	RW	0	Receive Absolute Delay Timer. Receive Absolute delay timer measured in increments of 1.024 microseconds (0 = disabled)
31:16	RO	0	Reserved. Reads as 0

If the packet delay timer is used to coalesce receive interrupts, it will ensure that when receive traffic abates, an interrupt will be generated within a specified interval of no receives. During times when receive traffic is continuous, it may be necessary to ensure that no receive remains unnoticed for too long an interval. This register may be used to *ensure* that a receive interrupt occurs at some predefined interval after the first packet is received.

When this timer is enabled, a separate absolute countdown timer is initiated upon successfully receiving each packet to system memory. When this absolute timer expires, pending receive descriptor write backs are flushed and a receive timer interrupt is generated.

Setting this register to 0 disables the absolute timer mechanism (the RDTR register should be used with a value of 0 to cause immediate interrupts for all receive packets).

Receive interrupts due to a Receive Packet Timer (RDTR) expiration will cancel a pending RADV interrupt. If enabled, the RADV countdown timer is reloaded but halted, so as to avoid generation of a spurious second interrupt after the RDTR has been noted.

11.1.2.3.15 Receive Small Packet Detect Interrupt- RSRPD (0x02C00; RW)

Bits	Type	Reset	Description
11:0	RW	0	SIZE. If the interrupt is enabled any receive packet of size <= SIZE will assert an Interrupt. SIZE is specified in bytes and includes the headers and the CRC. It does not include the VLAN header in size calculation if it is stripped.



Bits	Type	Reset	Description
31:12	RO	X	Reserved.

11.1.2.3.16 Receive ACK Interrupt Delay Register—RAID (0x02C08; RW)

Bits	Type	Reset	Description
15:0	RW	0	ACK_DELAY. ACK delay timer measured in increments of 1.024 microseconds. When the Receive ACK frame detect Interrupt is enabled in the IMS register, ACK packets being received will use a unique delay timer to generate an interrupt. When an ACK is received, an absolute timer will load to the value of ACK_DELAY. The interrupt signal will be set only when the timer expires. If another ACK packet is received while the timer is counting down, the timer will not be reloaded to ACK_DELAY.
31:16	RO		Reserved.

If an immediate (non-scheduled) interrupt is desired for any received Ack frame, the ACK_DELAY should be set to 0.

11.1.2.3.17 CPU Vector Register—CPUVEC (0x02C10; RW)

Bits	Type	Reset	Description
31:0	Reserved.	0x00	Reserved.

11.1.2.3.18 Receive Checksum Control—RXCSUM (0x05000; RW)

Bits	Type	Reset	Description
7:0	RW	0x00	Packet Checksum Start (PCSS).
8	RW	1	IP Checksum Offload Enable (IPOFL).
9	RW	1	TCP/UDP Checksum Offload Enable (TUOFL).
11:10	RO	0	Reserved.
12	RW	0	IP Payload Checksum Enable (IPPCSE).
13	RW	0	Packet Checksum Disable (PCSD).
15	RW	0	Reserved.
31:16	RO	0	Reserved.

The Receive Checksum Control register controls the receive checksum off loading features of LAN Controller. LAN Controller supports the off loading of three receive checksum calculations: the Packet Checksum, the IP Header Checksum, and the TCP/UDP Checksum.

PCSD: The Packet Checksum and IP Identification fields are mutually exclusive with the RSS hash. Only one of the two options is reported in the Rx descriptor. The RXCSUM.PCSD effect is shown in the table below:

RXCSUM.PCSD	0 (Checksum Enable)	1 (Checksum Disable)
Legacy Rx Descriptor (RCTL.DTYP = 00b)	Packet Checksum is reported in the Rx Descriptor	Forbidden Configuration
Extended or Header Split Rx Descriptor (RCTL.DTYP = 01b)	Packet Checksum and IP Identification are reported in the Rx Descriptor	RSS Hash value is reported in the Rx Descriptor



PCSS IPPCSE: The PCSS and the IPPCSE control the Packet Checksum calculation. As noted above, the packet checksum shares the same location as the RSS field. The Packet checksum is reported in the Receive descriptor when the RXCSUM.PCSD bit is cleared.

If RXCSUM.IPPCSE cleared (the default value), the checksum calculation that is reported in the Rx Packet checksum field is the unadjusted "16 bit ones complement" of the packet. The Packet Checksum starts from the byte indicated by RXCSUM.PCSS (0 corresponds to the first byte of the packet), after VLAN stripping if enabled (by CTRL.VME). For example, for an Ethernet II frame encapsulated as an 802.3ac VLAN packet and with RXCSUM.PCSS set to 14, the packet checksum would include the entire encapsulated frame, excluding the 14-byte Ethernet header (DA, SA, Type/Length) and the 4-byte VLAN tag. The Packet Checksum will not include the Ethernet CRC if the RCTL.SECRC bit is set. Software must make the required offsetting computation (to back out the bytes that should not have been included and to include the pseudo-header) prior to comparing the Packet Checksum against the TCP checksum stored in the packet.

If the RXCSUM.IPPCSE is set, the Packet checksum is aimed to accelerate checksum calculation of fragmented UDP packets.

The PCSS value should not exceed a pointer to IP header start or else it will erroneously calculate IP header checksum or TCP/UDP checksum.

RXCSUM.IPOFLD is used to enable the IP Checksum off-loading feature. If RXCSUM.IPOFLD is set to one, LAN Controller will calculate the IP checksum and indicate a pass/fail indication to software via the IP Checksum Error bit (IPE) in the ERROR field of the receive descriptor. Similarly, if RXCSUM.TUOFLD is set to one, LAN Controller will calculate the TCP or UDP checksum and indicate a pass/fail indication to software via the TCP/UDP Checksum Error bit (TCPE). Similarly, if RFCTL.IPv6_DIS and RFCTL.IP6Xsum_DIS are cleared to zero and RXCSUM.TUOFLD is set to one, LAN Controller will calculate the TCP or UDP checksum for IPv6 packets. It then indicates a pass/fail condition in the TCP/UDP Checksum Error bit (RDESC.TCPE).

This applies to checksum off loading only. Supported Frame Types:

- Ethernet II
- Ethernet SNAP

This register should only be initialized (written) when the receiver is not enabled (e.g. only write this register when RCTL.EN = 0).

11.1.2.3.19 Receive Filter Control Register—RFCTL (0x05008; RW)

Bits	Type	Reset	Description
0	RW	0	iSCSI Disable (ISCSI_DIS) . Disable the iSCSI filtering for header split functionality.
1:5	RW	0	iSCSI DWord count (ISCSI_DWC) . This field indicated the Dword count of the iSCSI header, which is used for packet split mechanism.
6	RW	0	NFS Write disable (NFSW_DIS) . Disable filtering of NFS write request headers for header split functionality.
7	RW	0	NFS Read disable (NFSR_DIS) . Disable filtering of NFS read reply headers for header split functionality.



Bits	Type	Reset	Description
9:8	RW	00	NFS Version (NFS_VER). 00 – NFS version 2 01 – NFS version 3 10 – NFS version 4 11 – Reserved for future use
10	RW	0	Reserved.
11	RW	0	Reserved.
12	RW	0	ACK accelerate disable (ACKDIS). When this bit is set LAN Controller will not accelerate interrupt on TCP ACK packets.
13	RW	0	ACK data Disable (ACKD_DIS). 1 – LAN Controller will recognize ACK packets according to the ACK bit in the TCP header + No –CP data 0 – LAN Controller will recognize ACK packets according to the ACK bit only. This bit is relevant only if the ACKDIS bit is not set.
14	RW	0	IP Fragment Split Disable (IPFRSP_DIS). When this bit is set the header of IP fragmented packets will not be set.
15	RW	0	Extended status Enable (EXSTEN). When the EXSTEN bit is set or when the Packet Split receive descriptor is used, LAN Controller writes the extended status to the Rx descriptor.
16	RO RW	0	Reserved.
17	RO RW	0	Reserved.
31:18	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

11.1.2.3.20 Multicast Table Array—MTA[31:0] (0x05200-0x0527C; RW)

Bits	Type	Reset	Description
31:0	RW	X	Bit Vector. Word wide bit vector specifying 32 bits in the multicast address filter table.

There is one register per 32 bits of the Multicast Address Table for a total of 32 registers (thus the MTA[31:0] designation). The size of the word array depends on the number of bits implemented in the multicast address table. Software must mask to the desired bit on reads and supply a 32-bit word on writes.

Note: All accesses to this table must be 32-bit.

The figure below diagrams the multicast lookup algorithm. The destination address shown represents the internally stored ordering of the received DA. Note that Byte 1 bit 0 indicated in this diagram is the first on the wire. The bits that are directed to the multicast table array in this diagram match a Multicast offset in the CTRL equals 00b. The complete multicast offset options are:

Multicast Offset	Bits directed to the Multicast Table Array
00b	DA[47:38] = Byte 6 bits 7:0, Byte 5 bits 7:6

Multicast Offset	Bits directed to the Multicast Table Array
01b	DA[46:37] = Byte 6 bits 6:0, Byte 5 bits 7:5
10b	DA[45:36] = Byte 6 bits 5:0, Byte 5 bits 7:4
11b	DA[43:34] = Byte 6 bits 3:0, Byte 5 bits 7:2

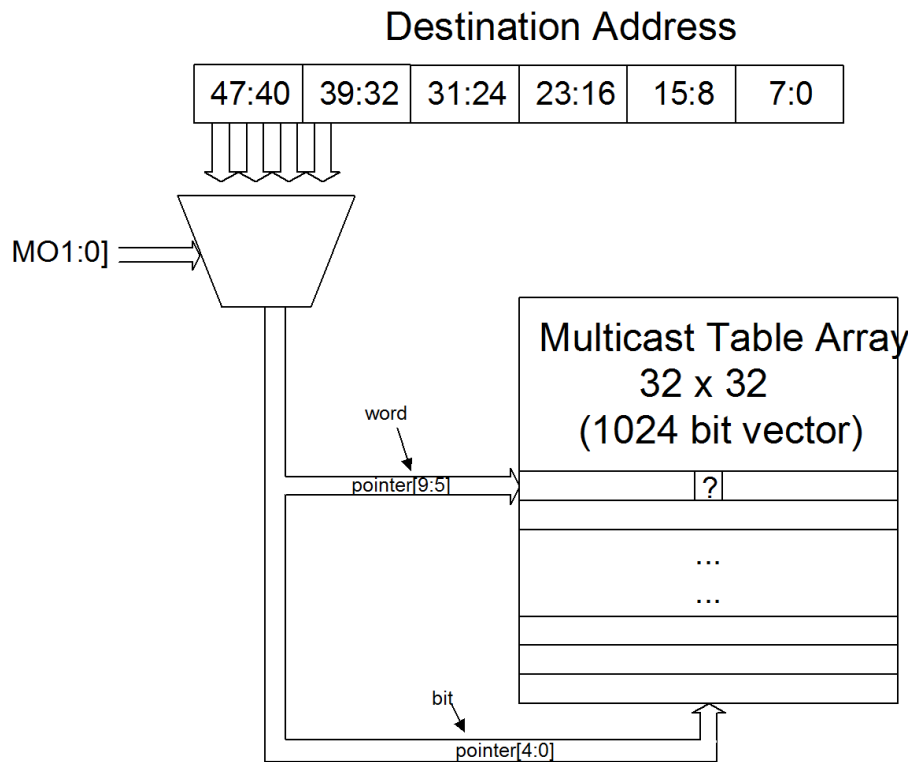


Figure 11-2. Multicast Table Array Algorithm

11.1.2.3.21 Receive Address Low—RAL (0x05400; RW)

Bits	Type	Reset	Description
31:0	RW	X	Receive Address Low (RAL) . The lower 32 bits of the 48 bit Ethernet address. RAL 0 is loaded from words 0 and 1 in the NVM.

11.1.2.3.22 Receive Address High—RAH (0x05404; RW)

Bits	Type	Reset	Description
15:0	RW	X	Receive Address High (RAH) . The upper 16 bits of the 48 bit Ethernet address. RAH 0 is loaded from word 2 in the NVM.



Bits	Type	Reset	Description
17:16	RW	X	Address Select (ASEL) . Selects how the address is to be used. Decoded as follows: 00 – Destination address (must be set to this in normal mode) 01 – Source address 10 – Reserved 11 – Reserved
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
30:19	RO	0	Reserved . Reads as 0. Ignored on write.
31	RW	see description	Address valid (AV) . Cleared after master reset. If the NVM is present, the <i>Address Valid</i> field of <i>Receive Address Register 0</i> will be set to 1 after a software or PCI reset or NVM read. This bit is cleared by master (software) reset.

AV determines whether this address is compared against the incoming packet. AV is cleared by a master (software) reset.

ASEL enables the device to perform special filtering on receive packets.

Note: RAR0 is also used for exact match pause frame checking (DA matches the first register). Therefore RAR0 should always be used to store the individual Ethernet MAC address of the adapter.

After reset, if the NVM is present, Receive Address Register 0 will be loaded from the IA field in the NVM, its *Address Select* field will be 00, and its *Address Valid* field will be 1. If no NVM is present the *Address Valid* field will be 0. The *Address Valid* field for all of the other registers will be 0.

11.1.2.3.23 Shared Receive Address Low—SHRAL[n] (0x05408 + 8*n (n=0...10); RW)

Bits	Type	Reset	Description
31:0	RW	X	Receive Address Low (RAL) . The lower 32 bits of the 48 bit Ethernet address n (n=0...10).

These registers may be WR locked by the LockMAC field in the FWSM register. The ME locks these registers when it uses them. The Host SW may use these registers that are not locked by ME.

11.1.2.3.24 Shared Receive Address High 0...10—SHRAH[n] (0x0540C + 8*n (n=0...10); RW)

Bits	Type	Reset	Description
15:0	RW	X	Receive Address High (RAH) . The upper 16 bits of the 48 bit Ethernet address n (n=0...10).
17:16	RO	00	Address Select (ASEL) . Selects how the address is to be used. 00b means that it is used to decode the Destination MAC address.
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
30:19	RO	0	Reserved . Reads as 0. Ignored on write.
31	RW	0	Address valid (AV) . When this bit is set, the relevant RAL,RAH are valid (compared against the incoming packet). Init trigger of this field depends on the state of the LockMAC state (see below).



These registers may be WR locked by the LockMAC field in the FWSM register. The ME locks these registers when it uses them. The Host SW may use those registers that are not locked by ME. The Address Valid bit of locked register is initialized with the PCI_m Data Path. If not locked by ME, the Address Valid bit is initialized with the PCI Data Path.

The description of SHRAH[9] is different and is described in [Section 11.1.2.3.25](#).

Note: The first receive address register (RAR0) is also used for exact match pause frame checking (DA matches the first register). Therefore RAR0 should always be used to store the individual Ethernet MAC address of the adapter.

After reset, the first register (Receive Address Register 0) will be loaded from the IA field in the NVM, its *Address Valid* field will be 1.

11.1.2.3.25 Shared Receive Address High 9—SHRAH[9] (0x05454; RW)

Bits	Type	Reset	Description
15:0	RW	X	Receive Address High (RAH). The upper 16 bits of the 48 bit Shared Ethernet address 9.
17:16	RO	00	Address Select (ASEL). Selects how the address is to be used. 00b means that it is used to decode the Destination MAC address.
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
29:19	RO	0	Reserved. Reads as 0. Ignored on write.
30	RW	0	All Nodes Multicast Address valid (MAV). The All Nodes Multicast address (33:33:00:00:00:01) is valid when this bit is set. Init trigger of this bit depends on the state of the LockMAC state (see below). Note that 0x33 is the first byte on the wire.
31	RW	0	Address valid (AV). When this bit is set the relevant Address 3 is valid (compared against the incoming packet). Init trigger of this field depends on the state of the LockMAC state (see below).

These registers may be WR locked by the LockMAC field in the FWSM register. The ME locks these registers when it uses them. The Host SW may use those registers that are not locked by ME. The Address Valid bit of locked register is initialized with the PCI_m Data Path. If not locked by ME, the Address Valid bit is initialized with the PCI Data Path.

[n]+ 4*n (n=0...11)

11.1.2.3.26 Multiple Receive Queues Command register—MRQC (0x05818; RW)

Bits	Type	Reset	Description
1:0	RW	0x00	Multiple Receive Queues Enable (MRxQueue). Enables support for Multiple Receive Queues and defines the mechanism that controls queue allocation. This field can be modified only when receive to host is not enabled (RCTL.EN = 0). 00b – Multiple Receive Queues are disabled 01b – Multiple Receive Queues as defined by MSFT RSS. The RSS Field Enable bits define the header fields used by the hash function. 10b – VMDq enable, enables VMDq operation as defined in section Receive Queuing for Virtual Machine Devices 11b – Reserved



Bits	Type	Reset	Description
2	Reserved. RW	0x0	Reserved.
15:3	RO	0x0	Reserved.
21:16	RW	0x0	Reserved.
31:22	RO	0x0	Reserved.

11.1.2.3.27 RSS Interrupt Mask Register—RSSIM (0x05864; RW)

Bits	Type	Reset	Description
31:0	Reserved.	0x00	Reserved.

11.1.2.3.28 RSS Interrupt Request Register—RSSIR (0x5868; RW)

Bits	Type	Reset	Description
31:0	Reserved.	0x00	Reserved.

11.1.2.3.29 Redirection Table—RETA (0x05C00 + 4*n (n=0...31); RW)

The redirection table is a 32 entry table. Each entry is composed of 4 Tags each 8-bits wide. Only the first or last 6 bits of each Tag are used (5 bits for the CPU index and 1 bit for Queue index).

Offset	31:24	23:16	15:8	7:0
0x05C00 + n*4	Tag 4*n+3	Tag 4*n+2	Tag 4*n+1	Tag 4*n

Bits	Type	Reset	Description
4:0	RW	X	CPU INDX 0. CPU index for Tag 4*n (n=0,1,...31)
6:5	RO	X	Reserved.
7	RW	X	QUE INDX 0. Queue Index for Tag 4*n (n=0,1,...31)
12:8	RW	X	CPU INDX 1. CPU index for Tag 4*n+1 (n=0,1,...31)
14:13	RO	X	Reserved.
15	RW	X	QUE INDX 1. Queue Index for Tag 4*n+1 (n=0,1,...31)
20:16	RW	X	CPU INDX 2. CPU index for Tag 4*n+2 (n=0,1,...31)
22:21	RO	X	Reserved.
23	RW	X	QUE INDX 2. Queue Index for Tag 4*n+2 (n=0,1,...31)
28:24	RW	X	CPU INDX 3. CPU index for Tag 4*n+3 (n=0,1,...31)
30:29	RO	X	Reserved.
31	RW	X	QUE INDX 3. Queue Index for Tag 4*n+3 (n=0,1,...31)

Note: RETA cannot be read when RSS is enabled.

11.1.2.3.30 Random Key Register—RSSRK (0x05C80 + 4*n (n=0...9); RW)

The RSS Random Key Register stores a 40 byte key (10 Dword entry table) used by the RSS hash function.



Bits	Type	Reset	Description
7:0	RW	0x00	K0. Byte n*4 of the RSS random key (n=0,1,...9)
15:8	RW	0x00	K1. Byte n*4+1 of the RSS random key (n=0,1,...9)
23:16	RW	0x00	K2. Byte n*4+2 of the RSS random key (n=0,1,...9)
31:24	RW	0x00	K3. Byte n*4+3 of the RSS random key (n=0,1,...9)

11.1.2.4 Transmit Register Descriptions

11.1.2.4.1 Transmit Control Register—TCTL (0x00400; RW)

Bits	Type	Reset	Description
0	RW	0	IP Identification 15 bit (IPID15). When '1', the IP Identification field will be incremented and wrapped around on 15-bit base. For example, if IP ID is equal to 0x7FFF then the next value will be 0x0000; if IP ID is equal to 0xFFFF then the next value will be 0x8000. When '0', the IP Identification field will be incremented and wrapped around on 16-bit base. In this case, the value following 0x7FFF is 0x8000, and the value following 0xFFFF is 0x0000. The purpose of this feature is to enable the software to manage 2 subgroups of connections.
1	RW	0	Enable (EN). The transmitter is enabled when this bit is 1. Writing this bit to 0 will stop transmission after any in progress packets are sent. Data remains in the transmit FIFO until the device is re-enabled. Software should combine this with reset if the packets in the FIFO should be flushed.
2	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
3	RW	1	Pad short packets (PSP). (with valid data, NOT padding symbols). 0 – do not pad; 1 – pad. Padding makes the packet 64B. This is not the same as the minimum collision distance. If Padding of short packets is allowed, the value in TX descriptor length field should be not less than 17 bytes.
11:4	RW	0x0F	Collision Threshold (CT). This determines the number of attempts at retransmission prior to giving up on the packet (not including the first transmission attempt). While this can be varied, it should be set to a value of 15 in order to comply with the IEEE specification requiring a total of 16 attempts. The Ethernet back-off algorithm is implemented and clamps to the maximum number of slot-times after 10 retries. This field only has meaning when in half-duplex operation.
21:12	RW	0x3F	Collision Distance (COLD). Specifies the minimum number of byte times which must elapse for proper CSMA/CD operation. Packets are padded with special symbols, not valid data bytes. Hardware checks and pads to this value plus one byte even in full-duplex operation. Default value is 64B – 512B times.
22	RW/V	0	Software XOFF Transmission (SWXOFF). When set to a 1 the device will schedule the transmission of an XOFF (PAUSE) frame using the current value of the PAUSE timer. This bit self clears upon transmission of the XOFF frame.
23	RW	0	Reserved.
24	RW	0	Re-transmit on Late Collision (RTLCL). Enables the device to retransmit on a late collision event.
27:25	RW	0x0	Reserved. Used to be UNORTX and TXDSCMT in predecessors.
28	RO	1	Reserved.



Bits	Type	Reset	Description
30:29	RW	01	Read Request Threshold (RRTHRESH). These bits will define the threshold size for the intermediate buffer to determine when to send the read command to the Packet buffer. Threshold is defined as follow: RRTHRESH – 00b Threshold – 2 lines of 16 bytes RRTHRESH – 01b Threshold – 4 lines of 16 bytes RRTHRESH – 10b Threshold – 8 lines of 16 bytes RRTHRESH – 11b Threshold – No threshold (transfer data after all of the request is in the RFIFO)
31	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Two fields deserve special mention: CT and COLD. Software may choose to abort packet transmission in less than the Ethernet mandated 16 collisions. For this reason, hardware provides CT.

Wire speeds of 1000 Mbps result in a very short collision radius with traditional minimum packet sizes. COLD specifies the minimum number of bytes in the packet to satisfy the desired collision distance. It is important to note that the resulting packet has special characters appended to the end. These are NOT regular data characters. Hardware strips special characters for packets that go from 1000 Mbps environments to 100 Mbps environments. Note that the hardware evaluates this field against the packet size in Full Duplex as well.

Note: While 802.3x flow control is only defined during full duplex operation, the sending of PAUSE frames via the SWXOFF bit is not gated by the duplex settings within the device. Software should not write a 1 to this bit while the device is configured for half duplex operation.

RTLIC configures the LAN Controller to perform retransmission of packets when a late collision is detected. Note that the collision window is speed dependent: 64 bytes for 10/100 Mb/s and 512 bytes for 1000Mb/s operation. If a late collision is detected when this bit is disabled, the transmit function assumes the packet is successfully transmitted. This bit is ignored in full-duplex mode.

11.1.2.4.2 Transmit IPG Register—TIPG (0x00410; RW)

Bits	Type	Reset	Description
9:0	RW	0x8	IPG Transmit Time (IPGT). Specifies the IPG length for back-to-back transmissions equal to $[(IPGT+4) \times 8]$ bit time.
19:10	RW	0x8	IPG Receive Time 1 (IPGR1). Specifies the defer IPG part 1 (during which carrier sense is monitored). Equal to $(IPGR1 \times 8)$ when DJHDX=0 and equals to $(IPGR1+2) \times 8$ when DJHDX=1.
29:20	RW	0x9	IPG Receive Time 2 (IPGR2). Specifies the defer IPG. Equal to $(IPGR2+3) \times 8$ when DJHDX=0 and equal to $(IPGR2+5) \times 8$ when DJHDX=1.
31:30	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register controls the IPG (Inter Packet Gap) timer. IPGT specifies the IPG length for back-to-back transmissions in both full and half duplex. Note that an offset of 4 byte times is added to the programmed value to determine the total IPG. Therefore, a value of 8 is recommended to achieve a 12 byte time IPG.

IPGR1 specifies the portion of the IPG in which the transmitter will defer to receive events. This should be set to 2/3 of the total effective IPG, or 8.



IPGR2 specifies the total IPG time for non back-to-back transmissions (transmission following deferral) in half duplex.

An offset of 5 byte times is added to the programmed value to determine the total IPG after a defer event. Therefore, a value of 7 is recommended to achieve a 12 byte time effective IPG for this case. Note the IPGR should never be set to a value greater than IPGT. If IPGR is set to a value equal to or larger than IPGT, it will override the IPGT IPG setting in half duplex, resulting in inter packet gaps that are larger than intended by IPGT in that case. Full Duplex will be unaffected by this, and will always rely on IPGT only.

In summary, the recommended TIPG value to achieve 802.3 compliant minimum transmit IPG values in full and half duplex is 0x00702008.

11.1.2.4.3 Adaptive IFS Throttle—AIT (0x00458; RW)

Bits	Type	Reset	Description
15:0	RW	0x0000	Adaptive IFS value (AIFS) . This value is in units of 8 nanoseconds.
31:16	RO	0x0000	Reserved . This field should be written with 0.

Adaptive IFS throttles back-to-back transmissions in the transmit packet buffer and delays their transfer to the CSMA/CD transmit function, and thus can be used to delay the transmission of back-to-back packets on the wire. Normally, this register should be set to 0. However, if additional delay is desired between back-to-back transmits, then this register may be set with a value greater than zero.

The Adaptive IFS field provides a similar function to the IPGT field in the TIPG register (see Section 11.1.2.5.2). However, it only affects the initial transmission timing, not re-transmission timing.

Note: If the value of the AdaptiveIFS field is less than the IPGTransmitTime field in the Transmit IPG registers then it will have no effect, as the chip will select the maximum of the two values.

11.1.2.4.4 Transmit Descriptor Base Address Low—TDBAL (0x03800 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
3:0	RO	0	Reserved . Ignored on writes. Returns 0 on reads
31:4	RW	X	Transmit Descriptor Base Address Low (TDBAL)

This register contains the lower bits of the 64 bit descriptor base address. The lower 4 bits are ignored. The Transmit Descriptor Base Address must point to a 16B aligned block of data.

11.1.2.4.5 Transmit Descriptor Base Address High—TDBAH (0x03804 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
31:0	RW	X	Transmit Descriptor Base Address [63:32] (TDBAH).

This register contains the upper 32 bits of the 64 bit Descriptor base address.



11.1.2.4.6 Transmit Descriptor Length—TDLEN (0x03808 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
6:0	RO	0	Reserved. Ignore on write. Reads back as 0.
19:7	RW	0	Descriptor Length (LEN).
31:20	RO	0	Reserved. Reads as 0. Should be written to 0.

This register contains the descriptor length and must be 128B aligned.

Note: The descriptor ring must be equal to or larger than 8 descriptors.

11.1.2.4.7 Transmit Descriptor Head—TDH (0x03810 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
15:0	RW/V	0	Transmit Descriptor Head (TDH).
31:16	RO	0	Reserved. Should be written with 0.

This register contains the head pointer for the transmit descriptor ring. It points to a 16B datum. Hardware controls this pointer. The only time that software should write to this register is after a reset (hardware reset or CTRL.SWRST) and before enabling the transmit function (TCTL.EN). If software were to write to this register while the transmit function was enabled, the on-chip descriptor buffers may be invalidated and the hardware could be become confused.

11.1.2.4.8 Transmit Descriptor Tail—TDT (0x03818 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
15:0	RW	0	Transmit Descriptor Tail (TDT).
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Note: This register contains the tail pointer for the transmit descriptor ring. It points to a 16B datum. Software writes the tail pointer to add more descriptors to the transmit ready queue. Hardware attempts to transmit all packets referenced by descriptors between head and tail.

11.1.2.4.9 Transmit Arbitration Count—TARC (0x03840 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
6:0	RW	3	Transmit Arbitration Count (COUNT). number of packets that can be sent from queue 0 to make the N over M arbitration between the queues. Writing 0 to this register is forbidden.
7	RW	0	Compensation mode (COMP). when set to 1 the HW will compensate this queue according to the compensation ratio, if the number of packets in a TCP segmentation in queue 1 caused the counter in queue 1 to go below zero
9:8	RW	00	Compensation Ratio (RATIO). this value will determine the ratio between the number of packets transmitted on queue1 in a TCP segmentation offload to the number of compensated packets transmitted from queue 0 00 – 1/1 compensation ratio 01 – 1/2 compensation ratio 10 – 1/4 compensation ratio 11 – 1/8 compensation ratio



Bits	Type	Reset	Description
10	RW	1	Descriptor enable (ENABLE). The ENABLE bit of the Transmit queue 0 should always be set.
26:11	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
27	RW	0	Reserved. Reserved for Multiple Tx request disable. This bit should not be modified by SW.
31:28	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

The default hardware value for TARC0.COUNT is **3** (this value is also reflected after reset).

The counter is subtracted as a part of the transmit arbitration.

It is reloaded to its high (last written) value when it decreases below zero.

- Upon read, the hardware returns the current counter value.
- Upon write, the counter will update the high value in the **next** counter-reload.
- The counter may be decreased in chunks (when transmitting TCP segmentation packets). It should never roll because of that.
The size of chunks is determined according to the TCP segmentation (number of packets sent).

When the counter reaches zero, other TX queues should be selected for transmission as soon as possible (usually after current transmission).

COMP is the enable bit to compensate between the two queues, when enabled (set to 1) the HW will compensate between the two queues if one of the queues is transmitting TCP segmentation packets and its counter went below zero, the HW will compensate the other queue according to the ratio in the TARC1.RATIO.

For example if the TARC0.COUNT reached (-5) after sending TCP segmentation packets and both TARC0.COMP and TARC1.COMP are enabled (set to 1) and TARC1.RATIO is 01 (1/2 compensation) TARC1.COUNT will be adjusted by adding $5/2=2$ to the current count.

RATIO is the multiplier to compensate between the two queues. The compensation method is described in the explanation above.

For DHG 802.3p using qWAVE API the following configuration will be used:

TARC0: COUNT = 1, COMP = 0, RATIO = 00.

TARC1: COUNT = 4, COMP = 1, RATIO = 00.

11.1.2.4.10 Transmit Interrupt Delay Value—TIDV (0x03820; RW)

Bits	Type	Reset	Description
15:0	RW	0	Interrupt Delay Value (IDV). Counts in units of 1.024 microseconds. A value of 0 is not allowed.
30:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
31	WO	0	Flush Partial Descriptor Block (FPD). when set to 1; ignored otherwise. Reads 0.



This register is used to delay interrupt notification for transmit operations by coalescing interrupts for multiple transmitted buffers. Delaying interrupt notification helps maximize the amount of transmit buffers reclaimed by a single interrupt. This feature *only* applies to transmit descriptor operations where (a) interrupt-based reporting is requested (RS set) and (b) the use of the timer function is requested (IDE is set).

This feature operates by initiating a countdown timer upon successfully transmitting the buffer. If a subsequent transmit delayed-interrupt is scheduled *before* the timer expires, the timer is re-initialized to the programmed value and re-starts its countdown. When the timer expires, a transmit-complete interrupt (ICR.TXDW) is generated.

Setting the value to 0 is not allowed. If an immediate (non-scheduled) interrupt is desired for any transmit descriptor, the descriptor IDE should be set to 0.

The occurrence of either an immediate (non-scheduled) or absolute transmit timer interrupt will halt the TIDV timer and eliminate any spurious second interrupts.

Transmit interrupts due to a Transmit Absolute Timer (TADV) expiration or an immediate interrupt (RS/RSP=1, IDE=0) will cancel a pending TIDV interrupt. The TIDV countdown timer is reloaded but halted, though it may be restarted by a processing a subsequent transmit descriptor.

Writing this register with FPD set initiates an immediate expiration of the timer, causing a write back of any consumed transmit descriptors pending write back, and results in a transmit timer interrupt in the ICR.

Note: FPD is self-clearing.

11.1.2.4.11 Transmit Descriptor Control—TXDCTL (0x03828 + n*0x100[n=0..1]; RW)

Note: This register was not fully validated. SW should set it to 0x0000 during nominal operation.

Bits	Type	Reset	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
15:14	RO	0x00	Reserved.
21:16	RW	0x00	Write-back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0x0	Granularity (GRAN) . Units for the thresholds in this register. 0 – cache lines 1 – descriptors
31:25	RW	0x0	Transmit descriptor Low Threshold (LWTHRESH). Interrupt asserted when the number of descriptors pending service in the transmit descriptor queue (processing distance from the TDT) drops below this threshold.

This register controls the fetching and write-back of transmit descriptors. The three threshold values are used to determine when descriptors will be read from and written to host memory. The values may be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag.

Note: When GRAN = 1 all descriptors will be written back (even if not requested).



PTHRESH is used to control when a prefetch of descriptors will be considered. This threshold refers to the number of valid, unprocessed transmit descriptors the chip has in its on-chip buffer. If this number drops below PTHRESH, the algorithm will consider pre-fetching descriptors from host memory. This fetch will not happen however, unless there are at least HTHRESH valid descriptors in host memory to fetch.

Note: HTHRESH should be given a non zero value when ever PTHRESH is used.

WTHRESH controls the write-back of processed transmit descriptors. This threshold refers to the number of transmit descriptors in the on-chip buffer which are ready to be written back to host memory. In the absence of external events (explicit flushes), the write-back will occur only after at least WTHRESH descriptors are available for write-back.

Possible values:

GRAN = 1 (descriptor granularity):

PTHRESH = 0..31

WTHRESH = 0..31

HTHRESH = 0..31

GRAN = 0 (cacheline granularity):

PTHRESH = 0..3 (for 16 descriptors cacheline—256 bytes)

WTHRESH = 0..3

HTHRESH = 0..4

Note: For any WTHRESH value other than 0—The packet and absolute timers must get a non zero value for the WTHRESH feature to take affect.

Note: Since the default value for write-back threshold is 0, descriptors are normally written back as soon as they are processed. WTHRESH must be written to a non-zero value to take advantage of the write-back bursting capabilities of the LAN Controller. If the WTHRESH is written to a non-zero value then all of the descriptors are written back consecutively no matter the setting of the RS bit.

Since write-back of transmit descriptors is optional (under the control of RS bit in the descriptor), not all processed descriptors are counted with respect to WTHRESH. Descriptors start accumulating after a descriptor with RS is set. Furthermore, with transmit descriptor bursting enabled, all of the descriptors are written back consecutively no matter the setting of the RS bit.

LWTHRESH controls the number of pre-fetched transmit descriptors at which a transmit descriptor-low interrupt (ICR.TXD_LOW) is reported. This may allow software to operate more efficiently by maintaining a continuous addition of transmit work, interrupting only when the hardware nears completion of all submitted work. LWTHRESH specifies a multiple of 8 descriptors. An interrupt is asserted when the number of descriptors available transitions from (threshold level=8*LWTHRESH)+1 to (threshold level=8*LWTHRESH). Setting this value to 0 will disable this feature.



11.1.2.4.12 Transmit Absolute Interrupt Delay Value-TADV (0x0382C; RW)

Bits	Type	Reset	Description
15:0	RW	0	Interrupt Delay Value (IDV) . Counts in units of 1.024 microseconds. (0 – disabled)
31:16	RO	0	Reserved . Reads as 0. Should be written to 0 for future compatibility.

The transmit interrupt delay timer (TIDV) may be used to coalesce transmit interrupts. However, it may be necessary to ensure that no completed transmit remains unnoticed for too long an interval in order ensure timely release of transmit buffers. This register may be used to *ensure* that a transmit interrupt occurs at some predefined interval after a transmit is completed. Like the delayed-transmit timer, the absolute transmit timer *only* applies to transmit descriptor operations where (a) interrupt-based reporting is requested (RS set) and (b) the use of the timer function is requested (IDE is set).

This feature operates by initiating a countdown timer upon successfully transmitting the buffer. When the timer expires, a transmit-complete interrupt (ICR.TXDW) is generated. The occurrence of either an immediate (non-scheduled) or delayed transmit timer (TIDV) expiration interrupt will halt the TADV timer and eliminate any spurious second interrupts.

Setting the value to 0 disables the transmit absolute delay function. If an immediate (non-scheduled) interrupt is desired for any transmit descriptor, the descriptor IDE should be set to 0.

11.1.2.5 Statistic Register Descriptions

Note: All Statistics registers reset when read. In addition, they stick at 0xFFFF_FFFF when the maximum value is reached.

Note: For the receive statistics it should be noted that a packet is indicated as “received” if it passes the device’s filters and is placed into the packet buffer memory. A packet does not have to be DMA’d to host memory in order to be counted as “received”.

Note: Due to divergent paths between interrupt-generation and logging of relevant statistics counts, it may be possible to generate an interrupt to the system for a noteworthy event prior to the associated statistics count actually being incremented. This is extremely unlikely due to expected delays associated with the system interrupt-collection and ISR delay, but might be observed as an interrupt for which statistics values do not quite make sense. Hardware guarantees that any event noteworthy of inclusion in a statistics count will be reflected in the appropriate count within 1 microsecond; a small time-delay prior to read of statistics may be necessary to avoid the potential for receiving an interrupt and observing an inconsistent statistics count as part of the ISR.

11.1.2.5.1 CRC Error Count—CRCERRS (0x04000; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	CEC . CRC error count.

Counts the number of receive packets with CRC errors. In order for a packet to be counted in this register, it must pass address filtering and must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. If receives are not enabled, then this register will not increment.



11.1.2.5.2 RX Error Count—RXERRC (0x0400C; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RXEC. RX error count.

Counts the number of packets received in which RX_ER was asserted by the PHY. In order for a packet to be counted in this register, it must pass address filtering and must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. If receives are not enabled, then this register will not increment.

11.1.2.5.3 Missed Packets Count—MPC (0x04010; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	MPC. Missed Packets Count.

Counts the number of missed packets. Packets are missed when the receive FIFO has insufficient space to store the incoming packet. This could be caused by too few buffers allocated, or because there is insufficient bandwidth on the IO bus. Events setting this counter cause RXO, the receiver overrun interrupt, to be set. This register does not increment if receives are not enabled.

Note: Note that these packets will also be counted in the Total Packets Received register as well as in Total Octets Received.

11.1.2.5.4 Carrier Extension Error Count—CEXTERR (0x0403C; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	CEXTERR. Number of packets received with a carrier extension error.

This register counts the number of packets received in which the carrier extension error was signaled across the GMII interface. The PHY propagates carrier extension errors to the MAC when an error is detected during the carrier extended time of a packet reception. An extension error is signaled by the PHY by the encoding of 0x1F on the receive data inputs while RX_ER is asserted to the MAC. This register will only increment if receives are enabled and the device is operating at 1000Mb/s.

11.1.2.5.5 Receive Length Error Count—RLEC (0x04040; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RLEC. Number of packets with receive length errors.

This register counts receive length error events. A length error occurs if an incoming packet passes the filter criteria but is undersized or oversized. Packets less than 64 bytes are undersized. Packets over 1522 bytes are oversized if LongPacketEnable is 0. If LongPacketEnable (LPE) is 1, then an incoming packet is considered oversized if it exceeds 16384 bytes.

If receives are not enabled, this register will not increment. These lengths are based on bytes in the received packet from <Destination Address> through <CRC>, inclusively.



11.1.2.5.6 XON Received Count—XONRXC (0x04048; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	XONRXC . Number of XON packets received.

This register counts the number of XON packets received. XON packets can use the global address, or the station address. This register will only increment if receives are enabled.

11.1.2.5.7 XON Transmitted Count—XONTXC (0x0404C; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	XONTXC . Number of XON packets transmitted.

This register counts the number of XON packets transmitted. These can be either due to queue fullness, or due to software initiated action (using SWXOFF). This register will only increment if transmits are enabled.

11.1.2.5.8 XOFF Received Count—XOFFRXC (0x04050; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	XOFFRXC . Number of XOFF packets received.

This register counts the number of XOFF packets received. XOFF packets can use the global address, or the station address. This register will only increment if receives are enabled.

11.1.2.5.9 XOFF Transmitted Count—XOFFTXC (0x04054; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	XOFFTXC . Number of XOFF packets transmitted.

This register counts the number of XOFF packets transmitted. These can be either due to queue fullness, or due to software initiated action (using SWXOFF). This register will only increment if transmits are enabled.

11.1.2.5.10 FC Received Unsupported Count—FCRUC (0x04058; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	FCRUC . Number of unsupported flow control frames received.

This register counts the number of unsupported flow control frames that are received.

The FCRUC counter is incremented when a flow control packet is received which matches either the reserved flow control multicast address (in FCAH/L) or the MAC station address, and has a matching flow control type field match (to the value in FCT), but has an incorrect opcode field. This register will only increment if receives are enabled.



11.1.2.5.11 Good Packets Received Count—GPRC (0x04074; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	GPRC. Number of good packets received (of any legal length).

This register counts the number of good (non-erred) packets received of any legal length. The legal length for the received packet is defined by the value of LongPacketEnable (see 10.2.7.13 Receive Length Error Count). This register does not include received flow control packets and only counts packets that pass filtering. This register will only increment if receives are enabled. This register does not count packets counted by the *Missed Packet Count (MPC)* register.

11.1.2.5.12 Broadcast Packets Received Count—BPRC (0x04078; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	BPRC. Number of broadcast packets received.

This register counts the number of good (non-erred) broadcast packets received. This register counts every broadcast packets received that passed filtering. This register does not count packets counted by the *Missed Packet Count (MPC)* register.

11.1.2.5.13 Multicast Packets Received Count—MPRC (0x0407C; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	MPRC. Number of multicast packets received.

This register counts the number of good (non-erred) multicast packets received. This register does not count multicast packets received that fail to pass address filtering nor does it count received flow control packets. This register does not count packets counted by the *Missed Packet Count (MPC)* register.

11.1.2.5.14 Good Packets Transmitted Count—GPTC (0x04080; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	GPTC. Number of good packets transmitted.

This register counts the number of good (non-erred) packets transmitted. A good transmit packet is considered one that is 64 or more bytes in length (from <Destination Address> through <CRC>, inclusively) in length. This does not include transmitted flow control packets. This register will only increment if transmits are enabled. This register does not count packets counted by the *Missed Packet Count (MPC)* register. The register counts clear as well as secure packets.

11.1.2.5.15 Good Octets Received Count—GORCL (0x04088; RO)

11.1.2.5.16 Good Octets Received Count—GORCH (0x0408C; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	GORCL. Number of good octets received – lower 4 bytes.
31:0	RO/V	0	GORCH. Number of good octets received – upper 4 bytes.



These registers make up a logical 64-bit register which counts the number of good (non-erred) octets received. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (GORCL) at this point the HW will snapshot the upper 32 bits to be read (GORCH).

In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached. Only packets that pass address filtering are counted in this register. This register will only increment if receives are enabled.

These octets do not include octets in received flow control packets.

11.1.2.5.17 Good Octets Transmitted Count—GOTCL (0x04090; RO);

11.1.2.5.18 Good Octets Transmitted Count—GOTCH (0x04094; RO);

Bits	Type	Reset	Description
31:0	RO/V	0	GOTCL. Number of good octets transmitted – lower 4 bytes.
31:0	RO/V	0	GOTCH. Number of good octets transmitted – upper 4 bytes.

These registers make up a logical 64-bit register which counts the number of good (non-erred) packets transmitted. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (GOTCL) at this point the HW will snapshot the upper 32 bits to be read (GOTCH).

In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached. This register includes bytes transmitted in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register counts octets in successfully transmitted packets which are 64 or more bytes in length. This register will only increment if transmits are enabled. The register counts clear as well as secure octets.

These octets do not include octets in transmitted flow control packets.

11.1.2.5.19 Receive No Buffers Count—RNBC (0x040A0; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RNBC. Number of receive no buffer conditions.

This register counts the number of times that frames were received when there were no available buffers in host memory to store those frames (receive descriptor head and tail pointers were equal). The packet will still be received if there is space in the FIFO. This register will only increment if receives are enabled.

This register does not increment when flow control packets are received.

11.1.2.5.20 Receive Undersize Count—RUC (0x040A4; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RUC. Number of receive undersize errors.



This register counts the number of received frames that passed address filtering, and were less than the minimum size (64 bytes from <Destination Address> through <CRC>, inclusively), and had a valid CRC. This register will only increment if receives are enabled.

11.1.2.5.21 Receive Fragment Count—RFC (0x040A8; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RFC. Number of receive fragment errors.

This register counts the number of received frames that passed address filtering, and were less than the minimum size (64 bytes from <Destination Address> through <CRC>, inclusively), but had a bad CRC (this is slightly different from the Receive Undersize Count register). This register will only increment if receives are enabled.

11.1.2.5.22 Receive Oversize Count—ROC (0x040AC; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	ROC. Number of receive oversize errors.

This register counts the number of received frames that passed address filtering, and were greater than maximum size. Packets over 1522 bytes are oversized if LongPacketEnable is 0. If LongPacketEnable (LPE) is 1, then an incoming, packet is considered oversized if it exceeds 16384 bytes.

If receives are not enabled, this register will not increment. These lengths are based on bytes in the received packet from <Destination Address> through <CRC>, inclusively.

11.1.2.5.23 Receive Jabber Count—RJC (0x040B0; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	RJC. Number of receive jabber errors.

This register counts the number of received frames that passed address filtering, and were greater than maximum size and had a bad CRC (this is slightly different from the Receive Oversize Count register).

Packets over 1522 bytes are oversized if LongPacketEnable is 0. If LongPacketEnable (LPE) is 1, then an incoming packet is considered oversized if it exceeds 16384 bytes.

If receives are not enabled, this register will not increment. These lengths are based on bytes in the received packet from <Destination Address> through <CRC>, inclusively.

11.1.2.5.24 Management Packets Received Count—MNGPRC (0x040B4; RO)

Bits	Type	Reset	Description
15:0	RO/V	0	MPRC. Number of management packets that were sent to the ME due to ME filtering or Receive circuit breaker redirection.
31:16	RO	0	Reserved.



This register counts the total number of packets received that pass the management filters or receive circuit breaker redirection. Any packets with errors are not counted, except for packets dropped because the management receive FIFO is full.

11.1.2.5.25 Management Packets Dropped Count—MNGPDC (0x040B8; RO)

Bits	Type	Reset	Description
15:0	RO/V	0	MPDC. Number of management packets dropped.
31:16	RO	0	Reserved.

This register counts the total number of packets received that pass the management filters and then are dropped because the management receive FIFO is full.

11.1.2.5.26 Management Packets Transmitted Count—MNGPTC (0x040BC; RO)

Bits	Type	Reset	Description
15:0	RO/V	0	MPTC. Number of management packets transmitted.
31:16	RO	0	Reserved.

This register counts the total number of packets that are transmitted that are received over the LAN Controller PCI-M interface.

11.1.2.5.27 Tx Circuit Breaker Packets Dropped -TCBPD (0x040D8; RO)

Bits	Type	Reset	Description
15:0	RO/V	0	TCBD. Dropped redirected transmit CB packets. Packets are dropped due to lack of room in the redirection MNG FIFO. It could be either due to M-Link BW or no memory resources.
31:16	RO	0	Reserved.

11.1.2.5.28 Total Octets Received—TORL (0x040C0; RO);

11.1.2.5.29 Total Octets Received—TORH (0x040C4; RO);

Bits	Type	Reset	Description
31:0	RO/V	0	TORL. Number of total octets received – lower 4 bytes.
31:0	RO/V	0	TORH. Number of total octets received – upper 4 bytes.

These registers make up a logical 64-bit register which count the total number of octets received. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (TORL) at this point the HW will snapshot the upper 32 bits to be read (TORH). In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached.

All packets received that pass address filtering will have their octets summed into this register, regardless of their length, whether they are errored, or whether they are flow control packets. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register will only increment if receives are enabled.



Note: Broadcast rejected packets will be counted in this counter (in contradiction to all other rejected packets that are not counted).

11.1.2.5.30 Total Octets Transmitted—TOTL (0x040C8; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	TOTL. Number of total octets transmitted – lower 4 bytes.

The TOTL and TOTH registers make up a logical 64-bit register which count the total number of octets transmitted. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (TORL) at this point the HW will snapshot the upper 32 bits to be read (TOTH). In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached.

All transmitted packets will have their octets summed into this register, regardless of their length or whether they are flow control packets. This register includes bytes transmitted in a packet from the <Destination Address> field through the <CRC> field, inclusively.

Octets transmitted as part of partial packet transmissions (e.g. collisions in half-duplex mode) are not included in this register. This register will only increment if transmits are enabled.

11.1.2.5.31 Total Octets Transmitted—TOTH (0x040CC; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	TOTH. Number of total octets transmitted – upper 4 bytes

See explanation of the TOTL above.

11.1.2.5.32 Total Packets Received—TPR (0x040D0; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	TPR. Number of all packets received.

This register counts the total number of all packets received. All packets received will be counted in this register, regardless of their length, whether they are erred, or whether they are flow control packets. This register will only increment if receives are enabled.

Note: Broadcast rejected packets will be counted in this counter (in contradiction to all other rejected packets that are not counted).

11.1.2.5.33 Total Packets Transmitted—TPT (0x040D4; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	TPT. Number of all packets transmitted.

This register counts the total number of all packets transmitted. All packets transmitted will be counted in this register, regardless of their length, or whether they are flow control packets.



Partial packet transmissions (e.g. collisions in half-duplex mode) are not included in this register. This register will only increment if transmits are enabled. This register counts all packets, including standard packets, secure packets, packets received over the SMB, and packets generated by the ASF function.

11.1.2.5.34 Multicast Packets Transmitted Count—MPTC (0x040F0; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	MPTC . Number of multicast packets transmitted.

This register counts the number of multicast packets transmitted. This register does not include flow control packets and increments only if transmits are enabled. Counts clear as well as secure traffic.

11.1.2.5.35 Broadcast Packets Transmitted Count—BPTC (0x040F4; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	BPTC . Number of broadcast packets transmitted count.

This register counts the number of broadcast packets transmitted. This register will only increment if transmits are enabled. This register counts all packets, including standard and secure packets. (Management packets will never be more than 200 bytes).

11.1.2.5.36 TCP Segmentation Context Transmitted Count—TSCTC (0x040F8; RO)

Bits	Type	Reset	Description
31:0	RO/V	0	TSCTC . Number of TCP Segmentation contexts transmitted count.

This register counts the number of TCP segmentation offload transmissions and increments once the last portion of the TCP segmentation context payload is segmented and loaded as a packet into the on-chip transmit buffer. Note that it is not a measurement of the number of packets sent out (covered by other registers). This register will only increment if transmits and TCP Segmentation offload are enabled.

11.1.2.5.37 Interrupt Assertion Count—IAC (0x04100; RO)

Bits	Type	Reset	Description
0-31	RO/V	0	IAC . This is a count of the interrupt assertions that have occurred. It counts the total number of interrupts generated in the system.



11.1.2.6 Management Register Descriptions

11.1.2.6.1 Wake Up Control Register—WUC (0x05800; RW)

Bits	Type	Reset	Description
0	RW/ SN	0	Advance Power Management Enable (APME). 1 – APM Wakeup is enabled 0 – APM Wakeup is disabled (See Section 7.4.1.1) Loaded from the NVM word 1Ah.
1	RW/V	0	PME_En. This read/write bit is used by the driver to access the PME_En bit of the <i>Power Management Control / Status Register</i> (PMCSR) without writing to PCI configuration space.
2	RWC	0	PME_Status. This bit is set when LAN Controller receives a wakeup event. It is the same as the PME_Status bit in the <i>Power Management Control / Status Register</i> (PMCSR). Writing a “1” to this bit will clear it, and also clear the PME_Status bit in the PMCSR.
3	RW	1	Assert PME On APM Wakeup (APMPME). If it is 1, LAN Controller will set the <i>PME_Status</i> bit in the <i>Power Management Control / Status Register</i> (PMCSR) and assert Host_Wake when APM Wakeup is enabled and LAN Controller receives a matching magic packet. (See Section 7.4.1.1).
4	RW/ SN	0	Link Status Change Wake Enable (LSCWE). Enables wake on link status change as part of APM wake capabilities. (See Section 7.4.1.1).
5	RW/ SN	0	Link Status Change Wake Override (LSCWO). If “1”, wake on Link Status Change does not depend on the LNKC bit in the Wake Up Filter Control Register (WUFC). Instead, it is determined by the APM settings in the WUC register. (See Section 7.4.1.1).
6	RO	0	Reserved. Was “APM Flexible Filter Allocation (APMFFA)”.
7	RO	0	Reserved. Was “Flexible APM filter Enable (FLEX_APM_FILTER_EN)”.
8	RW/ SN	0	Phy_Wake. This bit indicates if the Phy connected to the LAN controller supports wakeup. This bit is loaded from NVM word 13h bit 8.
29:9	RO	0	Reserved. Reads as 0.
31:30	RO	0	Reserved.

The *PME_Status* bits are cleared in the following conditions:

- If there is VAUX then the PME Status bits should be cleared by:
 - PWR Good (or PCI + PCIm reset)
 - Explicit Software Clear
- If there is NO VAUX then the PME Status bits should be cleared by:
 - PWR Good (or PCI + PCIm reset)
 - PCI Reset de-assertion
 - Explicit Software Clear

11.1.2.6.2 Wake Up Filter Control Register—WUFC (0x05808; RW)

Bits	Type	Reset	Description
0	RW	0	LNKC. Link Status Change Wake Up Enable
1	RW	0	MAG. Magic Packet Wake Up Enable
2	RW	0	EX. Directed Exact Wake Up Enable
3	RW	0	MC. Directed Multicast Wake Up Enable
4	RW	0	BC. Broadcast Wake Up Enable



Bits	Type	Reset	Description
5	RW	0	ARP. ARP/IPv4 Request Packet Wake Up Enable, IPv4 filtering applies only to the 3 host IPv4 addresses in IP4AT
6	RW	0	IPV4. Directed IPv4 Packet Wake Up Enable, IPv4 filtering applies only to the 3 host IPv4 addresses in IP4AT
7	RW	0	IPV6. Directed IPv6 Packet Wake Up Enable
8	RW	0	LCDINTWAKE. Enable wakeup based on LCD interrupt request (not used)
9	RW	0	LCD. Wakeup Done from the LAN Connected Device only
10:14	RO	0	Reserved.
15	RW	0	NoTCO. Ignore TCO Packets for TCO. If the NoTCO bit is set, then any packet that passes the manageability packet filtering will not cause a Wake Up event even if it passes one of the Wake Up Filters.
16	RW	0	FLX0. Flexible Filter 0 Enable
17	RW	0	FLX1. Flexible Filter 1 Enable
18	RW	0	FLX2. Flexible Filter 2 Enable
19	RW	0	FLX3. Flexible Filter 3 Enable
22	RW	0	FLX6. Flexible Filter 6 Enable
23	RW	0	FLX7. Flexible Filter 7 Enable

This register is used to enable each of the pre-defined and flexible filters for wake up support. A value of 1 means the filter is turned on, and a value of 0 means the filter is turned off.

11.1.2.6.3 Wake Up Status Register—WUS (0x05810; RW)

Bits	Type	Reset	Description
0	RW	0	LNKC. Link Status Changed
1	RW	0	MAG. Magic Packet Received
2	RW	0	EX. Directed Exact Packet Received. The packet's address matched one of the 7 pre-programmed exact values in the <i>Receive Address</i> registers.
3	RW	0	MC. Directed Multicast Packet Received. The packet was a multicast packet that hashed to a value corresponding to a 1 bit in the <i>Multicast Table Array</i> .
4	RW	0	BC. Broadcast Packet Received
5	RW	0	ARP. ARP/IPv4 Request Packet Received
6	RW	0	IPV4. Directed IPv4 Packet Received
7	RW	0	IPV6. Directed IPv6 Packet Received
15:8	RO	0	Reserved. Read as 0
16	RW	0	FLX0. Flexible Filter 0 Match
17	RW	0	FLX1. Flexible Filter 1 Match
18	RW	0	FLX2. Flexible Filter 2 Match
19	RW	0	FLX3. Flexible Filter 3 Match
22	RW	0	FLX6. Flexible Filter 6 Match
23	RW	0	FLX7. Flexible Filter 7 Match

This register is used to record statistics about all Wake Up packets received. A packet that matches multiple criteria may set multiple bits. Writing a 1 to any bit will clear that bit.



This register will not be cleared when PCI_RST_N is asserted. It will only be cleared when LAN PWR Good is de-asserted or when cleared by the driver.

11.1.2.6.4 Management Control Register Shadow—MANC_S (0x05820; RO)

This register is a RO shadow image of the PCIIm CSR at offset 0x00018.

11.1.2.6.5 Management Control 3 Register Shadow—MANC3_S (0x05828; RO)

This register is a RO shadow image of the PCIIm CSR at offset 0x00024.

11.1.2.6.6 SHRA Filter Enable Register Shadow—SHRAFER_S (0x0582C; RO)

This register is a RO shadow image of the PCIIm CSR at offset 0x00034.

11.1.2.6.7 Management Control To Host Register Shadow—MANC2H_S (0x05860; RO)

This register is a RO shadow image of the PCIIm CSR at offset 0x00020.

11.1.2.6.8 Management Control To Host 2 Register Shadow—MANC2H_S (0x05870; RO)

This register is a RO shadow image of the PCIIm CSR at offset 0x00028.

11.1.2.6.9 IP Address Valid—IPAV (0x5838; RW)

The IP Address Valid indicates whether the IP addresses in the IP Address Table are valid:

Bits	Type	Reset	Description
0	RO	0	Reserved.
1	RW	0	V41. IPv4 Address 1 Valid
2	RW	0	V42. IPv4 Address 2 Valid
3	RW	0	V43. IPv4 Address 3 Valid
4:15	RO	0x00	Reserved.
16	RW	0	V60. IPv6 Address Valid
31:17	RO	0x00	Reserved.

11.1.2.6.10 IPv4 Address Table—IP4AT (0x05840 + 8*n (n=1...3); RW)

The IPv4 Address Table is used to store the three IPv4 addresses for ARP/IPv4 Request packet and Directed IPv4 packet wake up. It is a 4 entry table with the following format:

Bits	Type	Reset	Description
31:0	RW	X	IPADD. IP Address n (n=1, 2, 3)

The register at address 0x5840 (n=0) was used in predecessors and reserved in the LAN Controller.

11.1.2.6.11 IPv6 Address Table—IP6AT (0x05880 + 4*n (n=0...3); RW)

The IPv6 Address Table is used to store the IPv6 address for Directed IPv6 packet wake up and Manageability traffic filtering. The IP6AT has the following format:



Bits	Type	Reset	Description
31:0	RW	X	IPv6 Address. IPv6 Address bytes n*4...n*4+3 (n=0, 1, 2, 3) while byte 0 is first on the wire and byte 15 is last.

11.1.2.6.12 Host to ME Register—H2ME (0x05B50; RW)

This register is used for message passing from Host to ME.

Bits	Type	Reset	Description
0	RW	0x0	LinkSec Request (LSECREQ). The host driver may assert this bit to request control from ME over LinkSec logic.
1	RW	0x0	Host LinkSec Connection Active (LSECA). Host sets this bit to '1' when it owns LinkSec logic and the host SW established a LinkSec channel. The host clears this bit when LinkSec channel is inactive.
2	RW	0x0	LinkSec Setup Failed (LSECSF). The host driver can assert this bit to notify the ME it had failed to setup a secured link.
3	RW	0x0	LinkSec Disabled (LSECD). The host driver sets this bit to 1 to indicate to the ME the LinkSec functionality is disabled.
4	RW	0x0	Reserved.
5	RW	0x0	IPV4_ARP_EN. The host driver sets this bit to 1 to indicate to the ME it should respond the IPv4 ARP requests as configured by the network proxy.
6	RW	0x0	IPV6_NS_EN. The host driver sets this bit to 1 to indicate to the ME it should respond the IPv6 NS requests as configured by the network proxy.
7	RW	0x0	Proxy_UC_Loaded. The host driver sets this bit to 1 to indicate the ME that the proxy UC was loaded to the PB
8	RW	0x0	Host Messaging Available (HMA). Host supports mailbox and there is a host application registered to receive messages
9	RW	0x0	Host Service Request (HSR). Host is requesting that the ME relinquish ownership of the mailbox. Cleared by the Host once ME reads H2ME_S (or once host assumes ownership of the mailbox)
10	RW	0x0	Host Message Ready (HMR). Set by the host to trigger interrupt to ME when a new message is ready. Cleared by the host when ME reads H2ME_S
31:11	RW	0x0	Host to ME Data (H2MED). Message passing content from host to ME. See message passing flow at Section 8.8.9.

11.1.2.6.13 Firmware Semaphore Register Shadow—FWSM_S (0x05B54; RO)

This register is a shadow RO image of the PCIm CSR at offset 0x00010.

11.1.2.6.14 Extended Firmware Semaphore Register Shadow—EXFWSM_S (0x05B58; RO)

This register is a shadow RO image of the PCIm CSR at offset 0x00060.

11.1.2.6.15 Flexible Filter Length Table—FFLT (0x05F00 + 8*n (n=0...7); RW)

There are 8 flexible filters Lengths. The Flexible Filter Length Table stores the minimum packet lengths required to pass each of the Flexible Filters. Any packets that are shorter than the programmed length will not pass that filter. Each Flexible Filter will consider a packet that does not have any mismatches up to that point to have passed the Flexible Filter when it reaches the required length. It will not check any bytes past that point.

Bits	Type	Reset	Description
10:0	RW	X	LEN. Minimum Length for Flexible Filter n.



Bits	Type	Reset	Description
31:11	RO	X	Reserved.

All reserved fields read as 0's and ignore writes.

Note: Before writing to the Flexible Filter Length Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

11.1.2.6.16 Flexible Filter Mask Table—FFMT (0x09000 + 8*n (n=0...127); RW)

There are 128 mask entries. The Flexible Filter Mask and Table is used to store the four 1-bit masks for each of the first 128 data bytes in a packet, one for each Flexible Filter. If the mask bit is 1, the corresponding Flexible Filter will compare the incoming data byte at the index of the mask bit to the data byte stored in the Flexible Filter Value Table.

Bits	Type	Reset	Description
0	RW	X	Mask 0. Mask for filter 0 byte n (n=0, 1... 127)
1	RW	X	Mask 1. Mask for filter 1 byte n (n=0, 1... 127)
2	RW	X	Mask 2. Mask for filter 2 byte n (n=0, 1... 127)
3	RW	X	Mask 3. Mask for filter 3 byte n (n=0, 1... 127)
6	RW	X	Mask 6. Mask for filter 6 byte n (n=0, 1... 127)
7	RW	X	Mask 7. Mask for filter 7 byte n (n=0, 1... 127)

Note: The table is organized to permit expansion to 8 (or more) filters and 256 bytes in a future product without changing the address map.

Note: Before writing to the Flexible Filter Mask Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

11.1.2.6.17 Flexible Filter Value Table—FFVT (0x09800 + 8*n (n=0...127); RW)

There are 128 filter values. The Flexible Filter Value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is 1, the Flexible Filter will compare the incoming data byte to the values stored in this table.

Bits	Type	Reset	Description
7:0	RW	X	Value 0. Value of filter 0 byte n (n=0, 1... 127)
15:8	RW	X	Value 1. Value of filter 1 byte n (n=0, 1... 127)
23:16	RW	X	Value 2. Value of filter 2 byte n (n=0, 1... 127)
31:24	RW	X	Value 3. Value of filter 3 byte n (n=0, 1... 127)

Before writing to the Flexible Filter Value Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

11.1.2.7 Time Sync Register Description

Note: All the registers bellow are for logical needs only, upon implementation their offset or structure can be changed according to the project specific needs.



11.1.2.7.1 RX Time Sync Control register—TSYNCRXCTL (0xB620; RW)

Bits	Type	Reset	Description
0	(RO/V)	0	RXTT. Rx timestamp valid equals '1' when a valid value for Rx timestamp is captured in the Rx timestamp register, clear by read of Rx timestamp register RXSTMPH.
3:1	RW	0x0	Type. Type of packets to timestamp 000b—Time stamp L2 (V2) packets only (Sync or Delay_req depends on message type in Section 11.1.2.7.6 and packets with message ID 2 and 3) 001b—Time stamp L4 (V1) packets only (Sync or Delay_req depends on message type in Section 11.1.2.7.6) 010b—Time stamp V2 (L2 and L4) packets (Sync or Delay_req depends on message type in Section 11.1.2.7.6 and packets with message ID 2 and 3) 100b—Time stamp all packets (in this mode no locking is done to the value in the timestamp registers and no indications in receive descriptors will be transferred) 101b—Time stamp all packets which message id bit 3 is zero, which means timestamp all event packets. This is applicable for V2 packets only. 011b, 110b and 111b – reserved
4	RW	0x0	En. Enable RX timestamp 0x0—time stamping disabled. 0x1—time stamping enabled.
5	RO	0x0	System Clock Frequency Indication (SYSCFI). This bit is used to indicate the host driver what is the clock frequency that is used for the system clock. 0b: 25Mhz 1b: 96Mhz
31:6	RO	0x0	Reserved

11.1.2.7.2 RX timestamp Low—RXSTMPL (0x0B624; RO)

Bits	Type	Reset	Description
31:0	RO	0x0	RXSTMPL. Rx timestamp LSB value

11.1.2.7.3 RX timestamp High—RXSTMPH (0x0B628; RO)

Bits	Type	Reset	Description
31:0	RO	0x0	RXSTMPH. Rx timestamp MSB value

11.1.2.7.4 RX timestamp attributes low—RXSATRL (0x0B62C; RO)

Bits	Type	Reset	Description
31:0	RO	0x0	SourceIDL. Sourceuuid low The value of this register is in host order.

11.1.2.7.5 RX timestamp attributes high—RXSATRH (0x0B630; RO)

Bits	Type	Reset	Description
15:0	RO	0x0	SourceIDH. Sourceuuid high The value of this register is in host order.
31:16	RO	0x0	SequenceID. Sequence ID The value of this register is in host order.

**11.1.2.7.6 RX message type register low—RXMTRL (0x0B634; RW)**

Bits	Type	Reset	Description
15:0	RW	0x88F7	PTP L2 EtherType to timestamp The value of this register is programmed/read in host order.
23:16	RW	0x0	V1 control to timestamp
31:24	RW	0x0	V2 messageId to timestamp

11.1.2.7.7 RX UDP port—RXUDP (0x0B638; RW)

Bits	Type	Reset	Description
15:0	RW	0x013F	UDP PORT. UDP port number to time stamp The value of this register is programmed/read in network order.
31:16	RO	0x0	Reserved.

11.1.2.7.8 TX Time Sync Control register—TSYNCTXCTL (0x0B614; RW)

Bits	Type	Reset	Description
0	RO/V	0	TXTT. Tx timestamp valid equals '1' when a valid value for Rx timestamp is captured in the Rx timestamp register, clear by read of Tx timestamp register TXSTMPH.
3:1	RO	0	Reserved.
4	RW	0	EN. Enable TX timestamp 0x0—time stamping disabled. 0x1—time stamping enabled.
5	RW	0	Packet mode. In this mode (deterministic SYSTIM) the SYSTIM register is incremented by TIMINCA.IV (increment value) every time a packet is time stamped. There is no significance whatsoever to the TIMINCA.IP (increment period) value. TIMADJ (time adjustment) registers work as they do in normal mode: every time they're set to a certain value, that value is immediately (next clock, not packet) added to or subtracted from the SYSTIM value. 0x0—Packet mode disabled 0x1—Packet mode enabled
31:6	RO	0	Reserved

11.1.2.7.9 TX timestamp value Low—TXSTMPL (0x0B618; RO)

Bits	Type	Reset	Description
31:0	RO	0x0	TXSTMPL. Tx timestamp LSB value

11.1.2.7.10 TX timestamp value High—TXSTMPH (0x0B61C; RO)

Bits	Type	Reset	Description
31:0	RO	0x0	TXSTMPH. Tx timestamp MSB value

11.1.2.7.11 System time register Low—SYSTIML (0x0B600; RO)

Bits	Type	Reset	Description
31:0	RW	0x0	STL. System time LSB register



11.1.2.7.12 System time register High—SYSTIMH (0x0B604; RO)

Bits	Type	Reset	Description
31:0	RW	0x0	STH. System time MSB register

11.1.2.7.13 Increment attributes register—TIMINCA (0x0B608; RW)

Bits	Type	Reset	Description
23:0	RW	0x0	IV. Increment value – <i>incvalue</i>
31:24	RW	0x0	IP. Increment period – <i>inperiod</i>

11.1.2.7.14 Time adjustment offset register low—TIMADJL (0x0B60C; RW)

Bits	Type	Reset	Description
31:0	RW	0x00	TADJL. Time adjustment value – Low

11.1.2.7.15 Time adjustment offset register high—TIMADJH (0x0B610;RW)

Bits	Type	Reset	Description
30:0	RW	0x00	TADJH. Time adjustment value—High
31	RW	0x0	Sign. Sign ("0"="+", "1"="-")

11.1.2.8 Diagnostic Register Descriptions

LAN Controller contains several diagnostic registers. These registers allow software to directly access the contents of the LAN Controller's internal Packet Buffer Memory (PBM), also referred to as FIFO space. These registers also give software visibility into what locations in the PBM that the HW currently considers to be the "head" and "tail" for both transmit and receive operations.

11.1.2.8.1 Receive Data FIFO Head Register—RDFH (0x02410; RW)

Bits	Type	Reset	Description
12:0	RW	0	Receive FIFO Head pointer (FIFO Head).
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores the head pointer of the on-chip receive data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Receive FIFO Head. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Receive FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.2 Receive Data FIFO Tail Register—RDFT (0x02418; RW)

Bits	Type	Reset	Description
12:0	RW	0	Receive FIFO Tail pointer (FIFO Tail).
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.



This register stores the tail pointer of the on-chip receive data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Receive FIFO Tail. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Receive FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.3 Receive Data FIFO Head Saved Register—RDFHS (0x02420; RW)

Bits	Type	Reset	Description
12:0	RW	0	FIFO Head. A "saved" value of the Receive FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores a copy of the Receive Data FIFO Head register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.4 Receive Data FIFO Tail Saved Register—RDFTS (0x02428; RW)

Bits	Type	Reset	Description
12:0	RW	0	FIFO Tail. A "saved" value of the Receive FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores a copy of the Receive Data FIFO Tail register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.5 Receive Data FIFO Packet Count—RDFPC (0x02430; RW)

Bits	Type	Reset	Description
12:0	RW	0	RX FIFO Packet Count. The number of received packets currently in the RX FIFO.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register reflects the number of receive packets that are currently in the Receive FIFO. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.6 Transmit Data FIFO Head Register—TDFH (0x03410; RW)

Bits	Type	Reset	Description
12:0	RW/V	0x0900 ¹	FIFO Head. Transmit FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Notes:

1. The initial value equals PBA.RXA times 128.

This register stores the head pointer of the on-chip transmit data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Transmit FIFO Head. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Transmit FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.



11.1.2.8.7 Transmit Data FIFO Tail Register—TDFT (0x03418; RW)

Bits	Type	Reset	Description
12:0	RW/V	0x0900 ¹	FIFO Tail. Transmit FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Notes:

1. The initial value equals PBA.RXA times 128.

This register stores the head pointer of the on-chip transmit data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Transmit FIFO Tail. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Transmit FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.8 Transmit Data FIFO Head Saved Register—TDFHS (0x03420; RW)

Bits	Type	Reset	Description
12:0	RW/V	0x000 ¹	FIFO Head. A "saved" value of the Transmit FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Notes:

1. The initial value equals PBA.RXA times 128.

This register stores a copy of the Transmit Data FIFO Head register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.9 Transmit Data FIFO Tail Saved Register—TDFTS (0x03428; RW)

Bits	Type	Reset	Description
12:0	RW/V	0x000 ¹	FIFO Tail. A "saved" value of the Transmit FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Notes:

1. The initial value equals PBA.RXA times 128.

This register stores a copy of the Transmit Data FIFO Tail register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

11.1.2.8.10 Transmit Data FIFO Packet Count—TDFPC (0x03430; RW)

Bits	Type	Reset	Description
12:0	RW	0	TX FIFO Packet Count. The number of packets to be transmitted that are currently in the TX FIFO.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register reflects the number of packets to be transmitted that are currently in the Transmit FIFO. This register is available for diagnostic purposes only, and should not be written during normal operation.

**11.1.2.8.11 Packet Buffer Memory—PBM (0x10000 - 0x15FFC; RW)**

Bits	Type	Reset	Description
31:0	RW	X	FIFO Data. Packet Buffer Data.

All PBM (FIFO) data is available to diagnostics. Locations can be accessed as 32 bit or 64 bit words. The internal PBM is 32KB in size. The PB size is set by the PBS register while RX and TX FIFO sizes are set by the PBA register. Regardless of the individual FIFO sizes that SW configures, the RX FIFO is located first in the memory mapped PBM space. So for the default FIFO configuration, the RX FIFO starts at offsets 0x10000 of the memory mapped space, while the TX FIFO follows the Rx FIFO.

Note: Packet buffer memory is only available when the device is in Testability Mode (PCI2PCIM).

11.1.2.8.12 Ghost Memory Data—GMD (0x0C000 - 0x0C3FC; RW)

Bits	Type	Reset	Description
31:0	RW	X	Data. Ghost Memory Data.

Ghost memory data is available to diagnostics. Locations can be accessed as 32 bit words.

11.1.2.8.13 Descriptor Memory Data—DMD (0x0C400 - 0x0C5FC; RW)

Bits	Type	Reset	Description
31:0	RW	X	Data. Descriptor Memory Data.

Descriptor memory data is available to diagnostics. Locations can be accessed as 32 bit words.

11.1.2.8.14 MNG CSR Control register—MNGCCR (0x05B60; RW)

Bits	Type	Reset	Description
0	RW/V	0x0	Valid (V). The valid bit is set by the SW for read or write transaction on the MNG CSR interface. The 'V' bit is cleared by HW at the end of the transaction.
1	RW	0x0	Write (Wr). When set indicates a write transaction. Valid only when the V bit is set.
2	RW/V	0x1	Complete (C). HW indication that the transaction was completed OK. This bit is valid only when the "V" bit is cleared.
4:3	RO	0x0	Destination (DST). Indicates the destination of the transaction. 0x0 PCIm CSR space. Else reserved
31:5	RW	X	Reserved.

Note: If the "V" bit is cleared by HW, and the "C" bit is not set, the SW should wait 2000 cycles before accessing the MNG CSR interface again.

Setting the "V" bit results in redirected access to the PCIm space but only in a testability mode as defined in [Strapping Options](#) section.



11.1.2.8.15 MNG CSR Address Register—MNGCAR (0x05B64; RW)

Bits	Type	Reset	Description
15:0	WR	X	CSR Address (CA). Address for the write/read transaction on the MNG CSR Interface. Valid when the V bit on the control register is set.
31:16	RO	X	Reserved.

11.1.2.8.16 MNG CSR Data Out Register—MNGCDO (0x05B68; RO)

This register will contain the data that was read by the MNG CSR interface. It is valid when the HW sets the C bit in the MNGMCR on a read transaction until the next time the 'V' bit is set by SW.

Bits	Type	Reset	Description
0-31	RO	0x0	DATAO. MNG CSR Data Out. Output data on MNG CSR read cycles.

11.1.2.8.17 MNG CSR Data In Register—MNGCDI (0x05B6C; RW)

This register contains the data that should be written by the MNG CSR interface. This register should be valid before the SW sets the 'V' bit on write transaction and must not be changed until the complete bit is set.

Bits	Type	Reset	Description
0-31	RW	0x0	DATAI. MNG CSR Data In. Input data for the MNG CSR write cycles.

Bits	Type	Reset	Description
11:0	RO	0	Reserved.
12	RW	0	Don't gate wake_dma_clk if DA removal is enabled 0b: Fix is disabled
15:13	RO	0	Reserved
18	RW	1	Reserved.
19	RW	1	Eco cb. Controls 4 changes in XTX: mask false irdy on dtaken—prevents false dtaken reduces # of dma->xtx wait states remove unneeded state from output fsm enables retx on prefetch packets
20	RO/ CR	0	HW SW CRC mismatch —Sticky bit clear on read This bit is asserted in case of mismatch between the software calculated CRC and the hardware calculated CRC. FEXT.8 must be set to enable this indication.
21	RW	0	When set to 0 no packet number counter update replay_protect feature is disabled, HSD2743902 bug fix (0 = fix enabled)
28:22	RO	0x0	Reserved.
29	RW	1	Reserved.
30	RW	1	Reserved.
31	RW	1	Reserved.



11.1.2.9 LinkSec Register Descriptions

11.1.2.9.1 LinkSec TX Capabilities register—LSECTXCAP (0x0B000; RW)

Bits	Type	Reset	Description
2:0	RO	1b	TX CA-supported Number of CA's supported by the device.
6:3	RO	1b	TX SC Capable. Number of SC's supported by the device on the transmit data path. The LAN controller supports twice the number of SA's as the TX SC for seamless re-keying, i.e., 2 SA's.
15:7	RO	0x0	Reserved.
23:16	RO/V	0x0	Tx LSEC Key SUM (LSECTXSUM). A bit wise XOR of the LSECTXKEY 0 bytes and LSECTXKEY 1 bytes. This register may be used by KaY (the programming entity) to validate key programming.
31:24	RO/V	0x0	Reserved.

11.1.2.9.2 LinkSec RX Capabilities register—LSECRXCAP (0x0B300; RW)

Bits	Type	Reset	Description
2:0	RO	1b	RX CA-supported Number of CA's supported by the device.
6:3	RO	4b1b	RX SC Capable. Number of SC's supported by the device on the receive data path. The LAN controller supports twice SA's as the RX SC for seamless re-keying, i.e., 8 2 SA's.
15:7	RO	0x0	Reserved.
23:16	RO/V	0x0	Rx LSEC Key SUM (LSECRXSUM). A bit wise XOR of the Rx LinkSec keys 0..7 as defined in registers LSECRXKEY [n, m]. Each byte is XORed with the respective byte of the other keys. This register may be used by KaY (the programming entity) to validate key programming.
31:24	RO	0x0	Reserved.

11.1.2.9.3 LinkSec TX Control register—LSECTXCTRL (0x0B004; RW)

Bits	Type	Reset	Description
1:0	RW	00b	Enable Tx LinkSec. Enable Tx LinkSec off loading. 00b – Disable Tx LinkSec (Tx all packets w/o LinkSec offload). 01b – Add integrity signature. 10b – Encrypt and add integrity signature. 11b – Reserved. When this field equals 00b (LinkSec offload is disabled). The "Tx Untagged Packet" register is not incremented for transmitted packets when the "Enable Tx LinkSec" equals 00b.
2	RW	0	PNID PN Increase Disable 0—Normal operation 1—PN is not incremented, used for testability mode only.
3	RW	0	(Testability feature) when set to 1 the HW ignores the ILSEC bit in the TX descriptor and transmits the packet as if the ILSEC bit was set.
4	RW	0b	Reserved



Bits	Type	Reset	Description
5	RW	1b	Always Include SCI. This field controls whether SCI is explicitly included in the transmitted SecTag. 0b – False 1b – True, SCI is explicitly included
6	RW	0b	Reserved
7	RW	1	Reserved
31:8	RW	11..1b	PN exhaustion threshold. MSB of the threshold over which HW needs to interrupt the KaY to warn of TX SA PN exhaustion and will trigger a new SA renegotiation. Bits 7:0 of the threshold are all 1's.

11.1.2.9.4 LinkSec RX Control register—LSECRXCTRL (0x0B304; RW)

Bits	Type	Reset	Description
1:0	RW	00b	Reserved
3:2	RW	00b	Enable Rx LinkSec. Controls the level of LinkSec packet filtering. 00b – Disable Rx LinkSec (pass all packets to host without LinkSec processing and no LinkSec header strip). 01b – Check (execute LinkSec offload and post frame to host and ME even when it fails LinkSec operation unless failed ICV and C bit was set). 10b – Strict (execute LinkSec offload and post frame to host and ME only if it does not fail LinkSec operation) 11b – Disable Rx LinkSec (drop all packets that include LinkSec header).
4	RO	1b	Reserved
5	RO	1b	Reserved
6	RW	0b	Reserved.
7	RW	1	Replay Protect. Enable replay protection.
8:8	RO	0x0	Reserved
31	RW	0	Reserved.

11.1.2.9.5 LinkSec TX SCI Low—LSECTXSCL (0x0B008; RW)

Bits	Type	Reset	Description
31:0	RW	0b	MAC Address SecY Low. The 4 LS bytes of the MAC address copied to the SCI field in the LinkSec header. The value of this register is programmed/read in host order.

11.1.2.9.6 LinkSec TX SCI High—LSECTXSCH (0x0B00C; RW)

Bits	Type	Reset	Description
15:0	RW	0b	MAC Address SecY High. The 2 MS bytes of the MAC address copied to the SCI field in the LinkSec header. The value of this register is programmed/read in host order.
31:16	RW	0b	Port Identifier. Always zero for transmitted packets.



11.1.2.9.7 LinkSec TX SA—LSECTXSA (0x0B010; RW)

Bits	Type	Reset	Description
1:0	RW	0b	AN0 – Association Number 0. This 2 bit field is posted to the AN field in the transmitted LinkSec header when SA 0 is active.
3:2	RW	0b	AN1 – Association Number 1. This 2 bit field is posted to the AN field in the transmitted LinkSec header when SA 1 is active.
4	RW	0b	SA Select (SelSA). This bit selects between SA 0 or SA 1 smoothly, i.e., on a packet boundary. A value of '0' selects SA 0 and a value of '1' selects SA 1.
5	RO/V	0b	Active SA (ActSA). This bit indicates the active SA. The ActSA follows the value of the SelSA on a packet boundary. The KaY (the programming entity) may use this indication to retire the old SA.
6	RW	0	Reserved.
7	RW	0	Reserved.
31:8	RW	0x0	Reserved.

11.1.2.9.8 LinkSec TX SA PN 0—LSECTXPNO (0x0B018; RW)

Bits	Type	Reset	Description
31:0	RW	0b	PN – Packet number. This field is posted to the PN field in the transmitted LinkSec header when SA 0 is active. It is initialized by the KaY at SA creation and then increments by 1 for each transmitted packet using this SA. Packets should never be transmitted if the PN repeats itself. In order to protect against such an event, the HW generates an LSECPN interrupt to KaY when the PN reaches the exhaustion threshold as defined in the LSECTXCTRL register. There is additional level of defense against repeating the PN. The HW will never transmit packets after the PN reaches a value of 0xFF..FF. In order to guarantee this, the HW clears the "Enable Tx LinkSec" field in the LSECTXCTRL register to 00b once a packet is transmitted with a PN that equals to 0xFF..F0. The value of this register is programmed/read in host order.

11.1.2.9.9 LinkSec TX SA PN 1—LSECTXPNI (0x0B01C; RW)

Bits	Type	Reset	Description
31:0	RW	0b	PN – Packet number. This field is posted to the PN field in the transmitted LinkSec header when SA 1 is active. It is initialized by the KaY at SA creation and then increments by 1 for each transmitted packet using this SA. Packets should never be transmitted if the PN repeats itself. In order to protect against such an event the HW generates an LSECPN interrupt to KaY when the PN reaches the exhaustion threshold as defined in the LSECTXCTRL register. There is additional level of defense against repeating the PN. The HW will never transmit packets after the PN reaches a value of 0xFF..FF. In order to guarantee this, the HW clears the "Enable Tx LinkSec" field in the LSECTXCTRL register to 00b once a packet is transmitted with a PN that equals to 0xFF..F0. The value of this register is programmed/read in host order.



11.1.2.9.10 LinkSec TX Key 0—LSECTXKEY0 [n] (0x0B020 + 4*n (n=0...3); WO)

Bits	Type	Reset	Description
31:0	WO	0x0	<p>LSEC Key 0. Transmit LinkSec key of SA 0.</p> <p>n – 0LSEC Key defines bits 31:0 of the Tx LinkSec Key n – 1LSEC Key defines bits 63:32 of the Tx LinkSec Key n – 2LSEC Key defines bits 95:64 of the Tx LinkSec Key n – 3LSEC Key defines bits 127:96 of the Tx LinkSec Key</p> <p>This field is WO for confidentiality protection. For data integrity check, the hash value may read the LSECTXSUM field in the LSECCAP register. If for some reason a read request is aimed to this register a value of all zeros will be returned.</p> <p>The value of this register is programmed/read in host order.</p>

11.1.2.9.11 LinkSec TX Key 1—LSECTXKEY1 [n] (0x0B030 + 4*n (n=0...3); WO)

Bits	Type	Reset	Description
31:0	WO	0x0	<p>LSEC Key 1. Transmit LinkSec key of SA 1.</p> <p>n – 0LSEC Key defines bits 31:0 of the Tx LinkSec Key n – 1LSEC Key defines bits 63:32 of the Tx LinkSec Key n – 2LSEC Key defines bits 95:64 of the Tx LinkSec Key n – 3LSEC Key defines bits 127:96 of the Tx LinkSec Key</p> <p>This field is WO for confidentiality protection. For data integrity check, the hash value may read the LSECTXSUM field in the LSECCAP register. If for some reason a read request is aimed to this register a value of all zeros will be returned.</p> <p>The value of this register is programmed/read in host order.</p>

11.1.2.10 LinkSec Rx SC Registers

The registers in this section relate to LinkSec Receive SC context. There are 4 SC's in the receive data path defined as SC0, SC1, SC2 and SC3. The registers below with index n relates to the SC index, while n=0, 1, 2, 3.

11.1.2.10.1 LinkSec RX SCI Low—LSECRXSCL (0x0B3D0 + 4*n (n=0...3); RW)

Bits	Type	Reset	Description
31:0	RW	0b	<p>MAC Address SecY low. The 4 LS bytes of the MAC address in the SCI field in the incoming packet that are compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set.</p> <p>Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively.</p> <p>The value of this register is programmed/read in host order.</p>

11.1.2.10.2 LinkSec RX SCI High—LSECRXSCH (0x0B3E0 + 4*n (n=0...3); RW)

Bits	Type	Reset	Description
15:0	RW	0b	<p>MAC Address SecY High. The 2 MS bytes of the MAC address in the SCI field in the incoming packet that are compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set.</p> <p>Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively.</p> <p>The value of this register is programmed/read in host order.</p>



Bits	Type	Reset	Description
31:16	RW	0b	Port Identifier. The Port Number in the SCI field in the incoming packet that is compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set. Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively. The value of this register is programmed/read in host order.

11.1.2.11 LinkSec Rx SA Registers

The registers in this section relates to LinkSec Receive SA context. There are 8 SA's in the receive data path defined as SA0 and SA1... SA7. The registers below with index n relates to the SA index, while n=0, 1... 7. SA0 and SA1 are assigned to SCI0, SA2 and SA3 are assigned to SCI1, SA4 and SA5 are assigned to SCI2 and SA6 and SA7 are assigned to SCI3.

11.1.2.11.1 LinkSec RX SA—LSECRXSA (0x0B310 + 4*n (n=0...7); RW)

Bits	Type	Reset	Description
1:0	RW	0b	AN – Association Number. This field is compared with the AN field in the TCI field in the incoming packet for match.
2	RW	0b	SA Valid. This bit is set or cleared by the KaY to validate or invalidate the SA.
3	RO/V	0b	Frame received. This bit is cleared when the SA Valid (bit 2) transitions from 0->1, and is set when a frame is received with this SA. When the Frame received bit is set the Retired bit of the other SA of the same SC is set. Note that a single frame reception with the new SA is sufficient to retire the old SA since we assume the Replay Window is 0.
4	RO/V	0b	Retired. When this bit is set the SA is invalid (retired). This bit is cleared when a new SA is configured by the KaY (SA Valid transition to 1). It is set to '1' when a packet is received with the other SA of the same SC. Note that a single frame reception with the new SA is sufficient to retire the old SA since we assume the Replay Window is 0.
31:5	RO	0	Reserved.

11.1.2.11.2 LinkSec RX SA PN—LSECRXSAPN (0x0B330 + 4*n (n=0...7); RW)

Bits	Type	Reset	Description
31:0	RW	0b	PN – Packet number. This register holds the PN field of the next incoming packet that uses this SA. The PN field in the incoming packet must be greater or equal to the PN register. The PN register is set by KaY at SA creation. It is updated by the hardware for each received packet using this SA to be Received PN + 1. The value of this register is programmed/read in host order.



11.1.2.11.3 LinkSec RX Key—LSECRXKEY[n,m] (0x0B350 + 0x10*n (n=0...7) + 4*m (m=0...3); WO)

Bits	Type	Reset	Description
31:0	WO	0x0	<p>LSEC Key. Receive LinkSec key of SA n, while n=0,1,2...7.</p> <p>m—0LSEC Key defines bits 31:0 of the Rx LinkSec Key</p> <p>m—LSEC Key defines bits 63:32 of the Rx LinkSec Key</p> <p>m—2LSEC Key defines bits 95:64 of the Rx LinkSec Key</p> <p>m—3LSEC Key defines bits 127:96 of the Rx LinkSec Key</p> <p>This field is WO for confidentiality protection. For data integrity check, the KaY hash value may read the LSECRXSUM field in the LSECCAP registers. If for some reason a read request is aimed to this register a value of all zeros will be returned. The value of this register is programmed/read in host order.</p>

11.1.2.12 LinkSec Tx Port Statistics

These counters are defined by spec as 64bit while implementing only 32 bit in the hardware. The KaY must implement the 64 bit counter in SW by regularly polling the hardware statistic counters. The HW section of the statistics counter is cleared upon read action.

11.1.2.12.1 Tx Untagged Packet Counter—LSECTXUT (0x04300; RC)

This statistic implements the SecyStatsRxUntaggedPkts statistic of the 802.1ae MIB while in non strict mode. In strict mode, this implements the secyStatsRxNoTagPkts statistic of the 802.1ae MIB.

Bits	Type	Reset	Description
31:0	RC	0x0	<p>Untagged Packet CNT. Increments for each transmitted packet that is transmitted with the ILSec bit cleared in the packet descriptor while "Enable Tx LinkSec" field in the LSECTXCTRL register is either 01b or 10b. The KaY must implement a 64 bit counter. It can do that by reading the LSECTXUT register regularly.</p>

11.1.2.12.2 Encrypted Tx Packets—LSECTXPKTE (0x04304; RC)

Bits	Type	Reset	Description
31:0	RC	0x0	<p>Encrypted Packet CNT. Increments for each transmitted packet through the controlled port with E bit set (i.e. confidentiality was prescribed for this packet by SW/FW).</p>

11.1.2.12.3 Protected Tx Packets—LSECTXPKTP (0x04308; RC)

Bits	Type	Reset	Description
31:0	RC	0x0	<p>Protected Packet CNT. Increments for each transmitted packet through the controlled port with E bit cleared (i.e. integrity only was prescribed for this packet by SW/FW).</p>



11.1.2.12.4 Encrypted Tx Octets—LSECTXOCTE (0x0430C; RC)

Bits	Type	Reset	Description
31:0	RC	0x0	Encrypted Octet CNT. Increments for each byte of user data through the controlled port with E bit set (i.e. confidentiality was prescribed for this packet by SW/FW).

11.1.2.12.5 Protected Tx Octets—LSECTXOCTP (0x04310; RC)

Bits	Type	Reset	Description
31:0	RC	0x0	Protected Octet CNT. Increments for each byte of user data through the controlled port with E bit (i.e. integrity only was prescribed for this packet by SW/FW).

11.1.2.13 LinkSec Rx Port Statistic Counters

These counters are defined by spec as 64bit while implementing only 32 bit in the hardware. The KaY must implement the 64 bit counter in SW by regularly polling the hardware statistic counters.

11.1.2.13.1 LinkSec Untagged RX Packet—LSECRXUT (0x04314; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Untagged Packet CNT. Increments for each packet received having no tag. Increments only when "Enable Rx LinkSec" field in the LSECRXCTRL register is either 01b or 10b.

11.1.2.13.2 LinkSec RX Octets Decrypted—LSECRXOCTD (0x0431C; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Decrypted Rx Octet CNT. The number of octets of User Data recovered from received frames that were both integrity protected and encrypted. This includes the octets from SecTag to ICV not inclusive. These counts are incremented even if the User Data recovered failed the integrity check or could not be recovered.

11.1.2.13.3 LinkSec RX Octets Validated—LSECRXOCTV (0x04320; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Validated Rx Octet CNT. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted. This includes the octets from SecTag to ICV not inclusive. These counts are incremented even if the User Data recovered failed the integrity check or could not be recovered.



11.1.2.13.4 LinkSec RX Packet with Bad Tag—LSECRXBAD (0x04324; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Bad Rx Packet CNT. Number of packets received having an invalid tag.

11.1.2.13.5 LinkSec RX Packet No SCI—LSECRXNOSCI (0x04328; RC)

Bits	Type	Reset	Description
31:0	RC	0b	No SCI Rx Packet CNT. Number of packets received having unrecognized SCI and dropped due to that condition.

11.1.2.13.6 LinkSec RX Packet Unknown SCI count—LSECRXUNSCI (0x432C; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Unknown SCI Rx Packet CNT. Number of packets received with an unrecognized SCI but still forwarded to the host.

11.1.2.14 LinkSec Rx SC Statistic Counters

11.1.2.14.1 LinkSec RX Unchecked Packets—LSECRXUNCH (0x04330; RC)

SW/FW needs to maintain the full sized register.

Bits	Type	Reset	Description
31:0	RC	0b	Unchecked Rx Packet CNT. Rx Packet CNT. Number of packets received with LinkSec encapsulation (SecTag) while ValidateFrames is disabled (LSECRXCTRL bits 3:2 equal 00b)."

11.1.2.14.2 LinkSec RX Delayed Packets—LSECRXDELAY (0x04340 + 4*n (n=0...3); RC)

SW/FW needs to maintain the full sized register.

Bits	Type	Reset	Description
31:0	RC	0b	Delayed Rx Packet CNT. Number of packets received and accepted for validation having failed replay-protection and ReplayProtect is false (LSECRXCTRL bit 7 is zero).

11.1.2.14.3 LinkSec RX Late Packets—LSECRXLATE (0x04350 + 4*n (n=0...3); RC)

SW/FW needs to maintain the full sized register.



Bits	Type	Reset	Description
31:0	RC	0b	Late Rx Packet CNT. Number of packets received and accepted for validation having failed replay-protection and ReplayProtect is true (LSECRXCTRL bit 7 is '1').

11.1.2.15 LinkSec Rx SA Statistic Counters

11.1.2.15.1 LinkSec RX Packet OK—LSECRXOK[n] (0x04360 + 4*n (n=0...7); RC)

Bits	Type	Reset	Description
31:0	RC	0b	OK Rx Packet CNT. Number of packets received that were valid (authenticated) and passed replay protection.

11.1.2.15.2 LinkSec Check RX Invalid—LSECRXINV[n] (0x43A0 + 4*n (n=0...7); RC)

Bits	Type	Reset	Description
31:0	RC	0b	Invalid Rx Packet CNT. Number of packets received that were not valid (authentication failed) and were forwarded to host.

11.1.2.15.3 LinkSec RX Not valid count—LSECRXNV[n] (0x04380 + 4*n [n=0...7]; RC)

Bits	Type	Reset	Description
31:0	RC	0b	Invalid Rx Packet CNT. Number of packets received that were not valid (authentication failed) and were dropped.

11.1.2.15.4 LinkSec RX Not using SA—LSECRXNUSA[n] (0x043C0 + 4*n (n=0...3); RC)

Bits	Type	Reset	Description
31:0	RC	0b	Invalid SA Rx Packet CNT. Number of packets received that were associated with an SA that is not "inUse" (No match on AN or not valid or retired) and were dropped. ¹

Notes:

1. The implementation maintains one such counter per SC.

11.1.2.15.5 LinkSec RX Unused SA—LSECRXUNSA[n] (0x043D0 + 4*n (n=0...3); RC)

Bits	Type	Reset	Description
31:0	RC	0b	Invalid SA Rx Packet CNT. Number of packets received that were associated with an SA that is not "inUse" (No match on AN or not valid or retired) and where forwarded to host. ¹

**Notes:**

1. The implementation maintains one such counter per SC.

11.1.3 PCI_m Configuration and Status Registers—CSR Space

11.1.3.1 PCI_m Register Map

All configuration registers are listed in the table below. These registers are ordered by grouping and are not necessarily listed in order that they appear in the address space. Note that the registers should be handled as DWORD entities therefore write accesses should be limited to whole 32bit DWORD cycles.

Register based Legend:

RW—Read Write register.

RO—Read Only Register.

R/WC—Read Write Clear register. Writing '0's has no affect. Writing '1's clears the appropriate fields (see detailed description of the specific registers).

RC/WC—Read Clear/ Write Clear register. Writing '0's has no affect. Writing '1's clears the appropriate fields (see detailed description of the specific registers). Read also clears the register.

Table 11-3. Register Summary

General Management Register Descriptions				
0x0000	MGCR	Management General Control Register	RW	240
0x0008	MDCR	Management DMA Control Register	RW	178
0x00010	FWSM	Firmware Semaphore Register	RW	241
0x00014	H2ME_S	Host to ME Register Shadow	RO	242
0x00018	MANC	Management Control Register	RW	242
0x0001C	MANC2	Management Control Register 2	RW	244
0x00020	MANC2H	Management to Host Control Register	RW	245
0x00024	MANC3	Management Control Register 3	RW	246
0x00028	MANC2H2	Management to Host Control Register 2	RW	248
0x0002C	IPV4MBC	IPv4 ME Binding Control	RW	249
0x00030	IPV6BC	IPv6 Binding Control	RW	547
0x00034	SHRAFER	SHRA Filter Enable Register	RW	550
0x0003C	IPV4HBC	IPv4 Host Binding Control	RW	547
0x00040	M2SCTRL	MAC To Smbus Control Register	RW	550
0x00044	M2SCFG	MAC To Smbus Configuration Register	RW	254
0x00048	M2STX	MAC To Smbus Transmit Control Register	RW	254
0x00050	H2MEM	Host To ME Mask	RW	254
0x00060	EXFWSM	Extended Firmware Semaphore Register	RW	255
0x000A0	CTRL_EXT_S	Extended Device Control Shadow	RO	254



0x000A4	MDIC_S	MDI Control Shadow	RO	255
0x000A8	RAL0_S	Receive Address Low 0 Shadow	RO	255
0x000AC	RAH0_S	Receive Address High 0 Shadow	RO	255
0x000B0	GENERAL_S	General bits Shadow	RO	255
0x000C0	FLXWINTC	Flex Write to Interrupt Control	RW	256
0x000C4	FLXCWINT01	Flex CSR Write to interrupt 01	RW	256
0x000C8	FLXCWINT23	Flex CSR Write to Interrupt 23	RW	256
0x000CC	FLXPWINT01	Flex Phy Write to Interrupt 01	RW	256
0x000D0	FLXPWINT23	Flex Phy Write to Interrupt 23	RW	256
Interrupt Register Descriptions				
0x0080	MICR	Management Interrupt Cause Register	RC/WC	257
0x0088	MIER	Management Interrupt Enable Register	RW	258
Receive Register Descriptions				
0x0100	MRBCR	Management Receive Buffer Control Register	RW	258
0x0108	MRBBA	Management Receive Buffer Base Address	RW	259
0x0110	MRBLA	Management Receive Buffer Last Address	RW	259
0x0120	MRBHR	Management Receive Buffer Head Register	RO	259
0x0128	MRBTR	Management Receive Buffer Tail Register	RW	260
0x0144	MRFUTPF	Management Range Flex UDP/TCP Port Filter	RW	260
0x0148	MFUTP01	Management Flex UDP/TCP Ports 0/1	RW	260
0x0150	MFUTP23	Management Flex UDP/TCP Port 2/3	RW	260
0x0158	MFUTP45	Management Flex UDP/TCP Port 4/5	RW	261
0x0160 + 0x4*n (n=0..3)	MIPV6	Management IPv6	RW	261
0x0180 + 8*n (n=0...1)	TCOFFLT	TCO Flexible Filter Length Table	RW	261
0x0190	TCOFFDAI	TCO Flexible Filter DA Index	RW	261
0x0800 + 8*n (n=0...127)	FTFT	Flexible TCO Filter Table	RW	262
0x01A0+4*n (n=0...3)	MAVTV[n]	Management VLAN TAG Value Table	RW	262
0x01C0	MIP4AT	Management IPv4 Address Table	RW	262
Transmit Register Descriptions				
0x0200	MTBCR	Management Transmit Buffer Control Register	RW	262
0x0208	MTBBA	Management Transmit Buffer Base Address	RW	263
0x0210	MTBLA	Management Transmit Buffer Last Address	RW	263
0x0220	MTBHR	Management Transmit Buffer Head Register	RW	263
0x0228	MTBTR	Management Transmit Buffer Tail Register	RO	264
PCI CSR Access Register Descriptions				
0x0300	MCCR	Management to MAC CSR Control register	RW	264
0x0310	MCAR	Management to MAC CSR Address Register	RW	265
0x0318	MCDO	Management to MAC CSR Data Out Register	RO	265
0x0320	MCDI	Management to MAC CSR Data In Register	RW	265
0x0340	FACTPS	PCI Function Power State to MNG	RO	265



Circuit Breaker Register Descriptions				
0x1000	CBCR	Circuit Breaker Configuration Register	RW	265
0x1010	CBSEED	CB Hash Seed Register	RW	266
0x1100	CBCS	Circuit Breaker Counter Status	R/WC	267
0x1200 + 8*n (n=0...31)	CBCTC[n]	Counter/Threshold Configuration	RW	267
0x1300 + 8*n (n=0...31)	CBCTV[n]	Counter/Threshold Value	RW	267
0x2000 + 8*n (n=0...31)	CBTC[n]	Transmit Filter Configuration Registers	RW	267
0x2100	CBTFS	Circuit Breaker Transmit Filter Status	R/WC	269
0x2200 + 8*n (n=0...30)	CBTIPV0[n]	Transmit Filter IP Address Value 0	RW	269
0x2300 + 8*n (n=0...30)	CBTIPM[n]	Transmit Filter IP Address Mask	RW	270
0x2400 + 8*n (n=0...30)	CBTPTV[n]	Transmit Filter Port / Type Value	RW	271
0x2500 + 8*n (n=0...30)	CBTNHFV[n]	Transmit Filter IP Next Header/Flags Value	RW	271
0x2600 + 8*n (n=0...30)	CBTTFM[n]	Transmit Filter TCP Flags Mask	RW	272
0x2700 + 8*n (n=0...30)	CBTIPV1[n]	Transmit Filter IP Address Value 1	RW	272
0x2800 + 8*n (n=0...30)	CBTIPV2[n]	Transmit Filter IP Address Value 2	RW	272
0x2900 + 8*n (n=0...30)	CBTIPV3[n]	Transmit Filter IP Address Value 3	RW	272
0x3000 + 8*n (n=0...31)	CBRC[n]	Receive Filter Configuration Registers	RW	273
0x3100	CBRFS	Circuit Breaker Receive Filter Status	R/WC	274
0x3200 + 8*n (n=0...30)	CBRIPV0[n]	Receive Filter IP Address Value 0	RW	274
0x3300 + 8*n (n=0...30)	CBRIPM[n]	Receive Filter IP Address Mask	RW	275
0x3400 + 8*n (n=0...30)	CBRPTV[n]	Receive Filter Port / Type Value	RW	276
0x3500 + 8*n (n=0...30)	CBRNHFV[n]	Receive Filter IP Next Header/Flags Value	RW	276
0x3600 + 8*n (n=0...30)	CBRTFM[n]	Receive Filter TCP Flags Mask	RW	277
0x3700 + 8*n (n=0...30)	CBRIPV1[n]	Receive Filter IP Address Value 1	RW	277
0x3800 + 8*n (n=0...30)	CBRIPV2[n]	Receive Filter IP Address Value 2	RW	277
0x3900 + 8*n (n=0...30)	CBRIPV3[n]	Receive Filter IP Address Value 3	RW	277
Header Redirection Register Descriptions				
0x0400	TXHCTL	Tx Header Buffer Control Register	RW	278
0x0408	TXHUTIL	Tx Header Utilization Control Register	RW	278
0x0410	MTXHBBA	Transmit Header Buffer Base Address	RW	279



0x0418	MTXHBLA	Transmit Header Buffer Last Address	RW	279
0x0420	MTXHBHR	Transmit Header Buffer Head Register	RO	279
0x0428	MTXHBTR	Transmit Header Buffer Tail Register	RW	280

11.1.3.2 General Management Register Descriptions

11.1.3.2.1 Management General Control Register—MGCR (0x0000; RW)

This register is used to control and setup general PCI-M capabilities.

Bits	Type	Reset	Description
0	RW/V	0	PCI-M Reset (RESET). When set, this bit resets all PCI-M resources, except for PCI-M configuration space. The bit is self-clearing.
1	RW/V/C	1	PHY Reset Asserted (PHYRA). This bit is R/W. The HW sets this bit following the assertion of LCD_RST. The bit is cleared on writing '0' to it. This bit may be used by FW as an indication that the SW initiated an LCD_RST.
2	RO	1	PHY Power Up not (PHYPWR). RO bit that indicates the power state of the PHY. It is a shadow bit of the PHYPWR in the STATUS register in the PCI space. 1—The PHY is powered on; fully active state. 0—The PHY is in the power down state. The PHYPWR bit is valid only after PHY reset is asserted. Note: The PHY Power Up indication reflects well the status of the LANPHYPC signaling to the LCD.
3	RO/V	0	MDIO HW Ownership. HW request for access to MDIO. Part of the arbitration scheme for MDIO access (see Section 2.2.2.1.1). This is a RO bit. It is a shadow bit of the same field in the EXTCNF_CTRL register in the PCI space.
4	RW/V	0	MDIO FW Ownership. FW request for access to MDIO. Part of the arbitration scheme for MDIO access (see Section 2.2.2.1.1). This is a copy of the FWFLAG bit field in the FWSM register.
5	RO/V	0	LAN Init Done. Asserted following completion of the LAN initialization from the FLASH. See "LAN Init Done Event" section for a complete description. Software is expected to clear this field to make it usable for the next Init done event.
6	RO/V	0	Master Read Completions Blocked: This bit is set when the device receives a completion with error (EP = 1 or status = successful). It is cleared on PCI-M reset.
7	RW	0	ME Wake on Link Status Change (WoLS). When set, the WoL to ME due to Link status change is enabled.
8	RW	0	Reserved.
10	RW	0	MRST Warn. When set to 1 indicates the ME is about to issue MRST and wants to make sure there is no LAN Controller SPI accesses in progress, once this bit is set the LAN Controller will find the next possible opportunity to stop any further SPI transactions and set MGCR.MRST_OK to indicate no SPI transactions will occur until the after MRST.
11	RO	0	MRST OK. When set to 1 by the LAN Controller indicates no SPI transactions will occur until after MRST. This bit will be set to 1 only after MGCR.MRST_Warn was set to 1 by FW and there are no pending SPI transactions.
31:12	RO	0x0	Reserved.

The RESET bit should only be set when firmware has a problem such that it must reset the PCI-M resources. Setting this bit will clear PCI-M resources, except for the PCI-M configuration space. All other PCI-M resources will default PCI-M to its initial hardware reset state. It will not, however, affect PCI space or shared resources (MAC, PHY, etc.). Firmware should make sure all outstanding transmit requests are completed before issuing a reset to avoid deadlock among the shared resources. Firmware should also disable bus mastering support before issuing the reset.



11.1.3.2.2 Firmware Semaphore Register—FWSM (0x00010; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05B54. The FWSM register is initialized only by LAN Power Good.

Bit	Type	Reset	Description
0	RW/V	0x0	FWFLAG. FW Semaphore FLAG. This bit is set by the FW (Read only to the SW). The bit is set by the HW only after the SWFLAG in the EXTCNF_CTRL register is cleared and the HW does not access the MDIO. The FW should set this bit and then wait till it is set. When it is set, it means that the FW can read/write from/to the shared resources to FW and LAN on the PCI. The FW should clear this bit when finishing its access.
1	RW	0x0	MEHostWakeEvent. When set to 1 indicates that the Host wake event was caused by ME. This bit should be cleared by FW prior to Sx entry.
2	RW	0x0	AMT_mode. Indicates FW support for AMT, AMT_mode is meaningful only when the FW_Val_bit is set
3	RW	0x0	PROXY_mode. Indicates FW support for Network Proxy, PROXY_mode is meaningful only when the FW_Val_bit is set
4	RW	0x0	ME Messaging Capable (MEMC). Indication that ME supports mailbox and there is a ME application registered to receive messages.
5	RW	0x0	ME Service Request (MESR). Indication that ME is requesting that the Host LAN driver relinquish ownership of the mailbox.
6	RW	1	RSPCIPHY. Reset PHY on PCI Reset. When this bit is set The LAN Connected Device is initialized by PCI Reset and SW LCD_RST. This bit should be cleared if manageability functionality is required over the link across host resets. Remote IDE is just one example for such requirement.
9:7	RW	0	Write Lock MAC Addresses (LockMAC). ME sets this field to lock MAC addresses that it uses against further WR accesses. The ME is expected to set this field after the MAC addresses that it uses are configured. The SW driver does not use MAC addresses that are locked by ME. LockMAC Functionality 000 None of the MAC addresses are locked 001 Shared Rx MAC Address 0,1,2,3,4,5,6,7,8,9,10 is WR protected 010 Shared Rx MAC Address 2,3,4,5,6,7,8,9,10 are WR protected 011 Shared Rx MAC Address 3,4,5,6,7,8,9,10 are WR protected 100 Shared Rx MAC Address 4,5,6,7,8,9,10 are WR protected 101 Shared Rx MAC Address 5,6,7,8,9,10 are WR protected 110 Shared Rx MAC Address 6,7,8,9,10 are WR protected 111 Shared Rx MAC Address 7,8,9,10 are WR protected
10	RW	0x0	Reserved
11	RW	0x0	LinkSec Ownership (LSECO). ME sets this bit to '1' when ME owns LinkSec logic. The ME may give-up ownership of the LinkSec logic if the host requests it by setting the LSECREQ bit in the H2ME register. By hardware default the host CPU owns the LinkSec logic. When the LSECO bit is set, the LinkSec logic directs all interrupts to ME. When the LSECO bit is cleared the LinkSec logic directs all interrupts to the host.
12	RW	0x0	Lock LinkSec Logic (LLL). ME sets this bit to '1' to lock any Read and Write accesses to the LinkSec CSR registers by the host. When LinkSec is locked, any Read cycles by the host return unexpected values and write accesses do not impact the context of the device. Both Read and Write accesses are completed w/o PCI error messages.
13	RW	0x0	ME LinkSec Connection Active (LSECA). ME sets this bit to '1' when ME owns LinkSec logic and the firmware established LinkSec channel. ME clears this bit when it closed the LinkSec channel.



14	RW	0x0	Proxy uCode load Support. Indicates FW support for loading the proxy uCode to the LCD.
15	RW	0x0	FW_Val_bit. FW Valid bit. HW clears this bit following PCIm reset or PCIm SW reset. FW set it to 1 when it is ready (end of boot sequence).
18:16	RW	0x0	Reset_cnt. Reset counter. FW increments it upon every reset
24:19	RW	0x0	Ext_err_ind. External error indication. FW writes here the reason that the FW has reset /clock gated (i.e. EEPROM, flash, patch corruption). Possible values: 0x00: No Error 0x01: Invalid EEPROM checksum 0x02: Unlocked secured EEPROM 0x03: Clock Off host command 0x04: Invalid FLASH checksum 0x05: C0 checksum failed 0x06: C1 checksum failed 0x07: C2 checksum failed 0x08: C3 checksum failed 0x09: TLB table exceeded 0x0A: DMA load failed 0x0B: Bad hardware version in patch load 0x0C: FLASH device not supported in the 82573 0x0D: Unspecified Error 0x3F: Reserved—max error value
25	RW	0x0	Reserved.
26	RW	0x0	Bits are valid in Extended FWSM.
27	RW	0	DISSWLNK. Disable SW Write Access from the Link Control registers. When set SW has no write access rights to the following registers: CTRL, CTRL_EXT, MDIC, PHY_CTRL. SW can read any of the CSR registers and NVM. A change in the state of this bit asserts the DSW interrupt bit in the PCI-ICR. See "Tamper Resistance" section for thorough explanation of the DISSWLNK usage.
28	RW	0	DISSW. Disable SW Write Access. When set SW has no write access rights to the LAN Controller. SW can read any of the CSR registers and NVM. Reading the ICR will clear it the same as it functions during nominal operation enabling the driver to clear the Interrupt source. A change in the state of this bit asserts the DSW interrupt bit in the PCI-ICR. See "Tamper Resistance" section for thorough explanation of the DISSW usage. Setting the DISSW also clears the IAM register in the PCI space or the CTRL_EXT.IAME bit. It is the SW driver's responsibility to restore the IAM register once it gains back control on the LAN controller.
30:29	RW	11b	Reserved.
31	RW	0	Gate HW Phy Configuration cycle (PPW, SKU read, OEM configuration).

11.1.3.2.3 Host to ME Register Shadow—H2ME_S (0x00014; RO)

This register is a RO Shadow register of the H2ME register in the PCI CSR space.

11.1.3.2.4 Management Control Register—MANC (0x00018; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05820.

Bits	Type	Reset	Description
1:0	RW	00	FLEXPOR0. This field enables the flexible UDP/TCP filter in MFUTP01 register as follows: 00 – MFUTP0 Filter is disabled 01 – MFUTP0 defines a UDP port number 10 – MFUTP0 defines a TCP port number 11 – MFUTP0 defines a UDP or TCP port



Bits	Type	Reset	Description
3:2	RW	00	FLEXPOR1. This field enables the flexible UDP/TCP filter in MFUTP01 register as follows: 00 – MFUTP1 Filter is disabled 01 – MFUTP1 defines a UDP port number 10 – MFUTP1 defines a TCP port number 11 – MFUTP1 defines a UDP or TCP port
5:4	RW	00	FLEXPOR2. This field enables the flexible UDP/TCP filter in MFUTP23 register as follows: 00 – MFUTP2 Filter is disabled 01 – MFUTP2 defines a UDP port number 10 – MFUTP2 defines a TCP port number 11 – MFUTP2 defines a UDP or TCP port
6	RW	0	FLEX_TCO1_FILTER_EN. Flex TCO filter 1 filtering enable. When this bit is set the flex TCO filtering mechanism is enabled for filter 1.
7	RW	0	FLEX_TCO0_FILTER_EN. Flex TCO filter 0 filtering enable. When this bit is set the flex TCO filtering mechanism is enabled for filter 0.
8	RW	0	RMCP_EN. Enable RMCP 026Fh Filtering. Sends RMCP packets with a destination port of 026Fh that pass RMCP filtering to the SMB or internal ASF controller.
9	RW	0	EN_0298. Enables RMCP 0298h Filtering. Sends RMCP packets with a destination port of 0298h that pass RMCP filtering to the SMB or internal ASF controller.
11:10	RW	00	FLEXPOR3. This field enables the flexible UDP/TCP filter in MFUTP23 register as follows: 00 – MFUTP3 Filter is disabled 01 – MFUTP3 defines a UDP port number 10 – MFUTP3 defines a TCP port number 11 – MFUTP3 defines a UDP or TCP port
12	RW	0	Reserved
13	RW	0	ARP_REQ_EN. Enable ARP Request Filtering. Sends ARP request packets that pass ARP filtering to the SMB or internal ASF controller.
14	RW	0	NEIGHBOR_EN. Enables Neighbor Solicitation Message Filtering. Sends packets that pass Neighbor Solicitation Message filtering to the ME.
15	RW	0	ARP_RES_EN. Enables ARP Response Filtering. Sends ARP response packets that pass ARP filtering to the SMB or internal ASF controller.
16	RW	0	EN_ANM_FILTER. Enables the manageability receive all nodes multicast filter for IPV6. When this bit is set and the MAV bit is set in the SHRAH[3] register and the incoming packet matches the All Node Multicast Address filter the packet will be routed to the ME.
17	RW	0	RCV_TCO_EN. Receive TCO Packets Enabled. When this bit is set it enables the receive flow from the wire to ME.
18	RW	0	Explicitly add Host MTA filter to ME Filter. When Add Host MTA filter is set to 1 setting this bit will explicitly add the host MTA table filter to the filters that can pass to ME (not bound to MNG MAC Must match and no need to pass a higher level filter)
19	RW	0	RCV_ALL. Receive All Enable. When set it means that all packets received from the wire will be routed to the ME.
20	RW	0	Add Host MTA filter to MNG MAC Must match. When MNG MAC Must match is set to 1 setting this bit will also add the host MTA table filter to the filters that can pass to ME (with a higher level filter)
21	RW	0	EN_MNG2HOST. Enables MNG packets to host memory. This bit enables the functionality of the MANC2H register. When set, the packets that are specified in the MANC2H registers will be forwarded to the HOST memory too, if they pass manageability filters.
22	RW	0	NoHost WoL on ME traffic. When set to 1 a packet that is routed only to ME will not be candidate for Host WoL. A packet that is routed to ME but also to the Host (MANC2H) will be candidate for Host WoL. When set to 0 all packets are candidates for Host WoL.



Bits	Type	Reset	Description
23	RW	0	EN_XSUM_FILTER. Enable XSUM filtering to MNG. When this bit is set, only packets that pass L3, L4 checksum will be sent to the ME.
24	RW	0	BR_EN. Enables Broadcast filtering. When set, the Filter sends broadcast packets to the ME.
25	RW	0	MNG MAC Must match. At '0' all packets that pass any of the MAC address filters may be directed to the ME, if it matches the MNG filters as well. When set to '1' only packets that pass any of the Shared Receive Addresses dedicated for MNG as defined by the LockMAC field in the FWSM register, or Broadcast packets or MTA matched packets (if Add Host MTA filter is set), can be directed to the ME if they match the MNG filters as well.
30:26	RW	0	Reserved
31	RW	0	Reserved

Note: If RCV_TCO_EN is not set, then FW should not set any of the specific MNG filters.

11.1.3.2.5 Management Control Register 2—MANC2 (0x0001C; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05824.

Bits	Type	Reset	Description
1:0	RW	00	FLEXPOR4. This field enables the flexible UDP/TCP filter in MFUTP45 register as follow: 00 – MFUTP4 Filter is disabled 01 – MFUTP4 defines a UDP port number 10 – MFUTP4 defines a TCP port number 11 – MFUTP4 defines a UDP or TCP port
3:2	RW	00	FLEXPOR5. This field enables the flexible UDP/TCP filter in MFUTP45 register as follow: 00 – MFUTP5 Filter is disabled 01 – MFUTP5 defines a UDP port number 10 – MFUTP5 defines a TCP port number 11 – MFUTP5 defines a UDP or TCP port
8	RW	0	DSTSRC4. This field controls if the MFUTP4 is filtering for a destination or a source port number. When set to: 0 – destination port is compared 1 – source port is compared
9	RW	0	DSTSRC5. This field controls if the MFUTP5 is filtering for a destination or a source port number. When set to: 0 – destination port is compared 1 – source port is compared
12	RW	0	L24IPV6_0. This field enables the L24IPV6[0] filter.
13	RW	0	L24IPV6_1. This field enables the L24IPV6[1] filter.
14	RW	0	L24IPV6_2. This field enables the L24IPV6[2] filter.
15	RW	0	L24IPV6_3. This field enables the L24IPV6[3] filter.
16	RW	0	Reserved
17	RW	0	ICMPv6. This field enables the ICMPv6 filter. When set to 1 ICMPv6 packets received will be routed to the ME.
18	RW	0	ICMP_IPV4. This field enables the IPV4 ICMP filter. When set it means that IPV4 ICMP packets received from the wire will be routed to the ME.



11.1.3.2.6 Management Control Register—MANC2H (0x00020; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05860.

The MANC2H register allows routing of manageability packets to the host based on the filter type that routed it to the manageability micro-controller. Each manageability filter has a corresponding bit in the MANC2H register. When a legacy filter to the manageability micro-controller routes a packet, it is also routed to the host if the respective filter bit is set in the MANC2H register and if the EN_MNG2HOST bit is set. The EN_MNG2HOST bit serves as a global enable for the MANC2H bits.

Bits	Type	Reset	Description
0	RW	0	Flex Port 0 (FP0) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
1	RW	0	Flex Port 1 (FP1) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
2	RW	0	Flex Port 2 (FP2) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
3	RW	0	Flex TCO 0 (FT0) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
4	RW	0	Flex TCO 1 (FT1) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
5	RW	0	026F (FLT_026F) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
6	RW	0	0298 (FLT_0298) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
7	RW	0	ARP_REQuest (ARP_REQ) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
8	RW	0	ARP_RESponse (ARP_RES) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
9	RW	0	Broadcast (BR) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
10	RW	0	Neighbor (NE) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
11	RW	0	VLAN 0 (VLAN0) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
12	RW	0	VLAN 1 (VLAN1) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
13	RW	0	VLAN 2 (VLAN2) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
14	RW	0	VLAN 3 (VLAN3) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
15	RW	0	Manageability MAC (MNG_MAC) . When set indicates that packets that are routed to the ME due to a match of the destination MAC address to any of the Shared Receive Addresses, are sent to the HOST as well.
16	RW	0	Flex Port 3 (FP3) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
17	RW	0	Manageability All Nodes Multicast MAC (MNG_ANM) . When set to '1' packets that are routed to the ME due to a match of the destination MAC address to 33:33:00:00:00:01 are sent to the HOST as well.
18	RW	0	Low IPv6 address 0 (L24IPV60) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
19	RW	0	Low IPv6 address 1 (L24IPV61) . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.



20	RW	0	Low IPv6 address 2 (L24IPV62). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
21	RW	0	Low IPv6 address 3 (L24IPV63). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
22	RW	0	Flex Port 4 (FP4). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
23	RW	0	Flex Port 5 (FP5). When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
24	RW	0	ICMP_IPV4 When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.

11.1.3.2.7 Management Control Register 3- MANC3 (0x00024; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05828.

Bits	Type	Reset	Description
0	RW	0	TCPPORT0. This field enables the TCP port filter 16992 (decimal) 0 - TCPPORT0 Filter is disabled 1 - TCPPORT0 Filter is enabled
1	RW	0	TCPPORT1. This field enables the TCP port filter 16993 (decimal) 0 - TCPPORT1 Filter is disabled 1 - TCPPORT1 Filter is enabled
2	RW	0	TCPPORT2. This field enables the TCP port filter 16994 (decimal) 0 - TCPPORT2 Filter is disabled 1 - TCPPORT2 Filter is enabled
3	RW	0	TCPPORT3. This field enables the TCP port filter 16995 (decimal) 0 - TCPPORT3 Filter is disabled 1 - TCPPORT3 Filter is enabled
4	RW	0	TCPPORT4. This field enables the TCP port filter 5900 (decimal) 0 - TCPPORT4 Filter is disabled 1 - TCPPORT4 Filter is enabled
5	RW	0	TCPPORT5. This field enables the TCP port filter 5223 (decimal) 0 - TCPPORT5 Filter is disabled 1 - TCPPORT5 Filter is enabled
6	RW	0	TCPPORT6. This field enables the TCP port filter 139 (decimal) 0 - TCPPORT6 Filter is disabled 1 - TCPPORT6 Filter is enabled
7	RW	0	TCPPORT7. This field enables the TCP port filter 445 (decimal) 0 - TCPPORT7 Filter is disabled 1 - TCPPORT7 Filter is enabled
8	RW	0	TCPPORT8. This field enables the TCP port filter 5060 (decimal) 0 - TCPPORT8 Filter is disabled 1 - TCPPORT8 Filter is enabled
9	RW	0	TCPPORT9. This field enables the TCP port filter 5061 (decimal) 0 - TCPPORT9 Filter is disabled 1 - TCPPORT9 Filter is enabled
10	RW	00	TCPPORT10. This field enables the TCP port filter 3544 (decimal) 0 - TCPPORT10 Filter is disabled 1 - TCPPORT10 Filter is enabled
11	RO	0	Reserved



12	RW	0	UDPPORT0. This field enables the UDP port filter 68 (decimal) 0 – UDPPORT0 Filter is disabled 1 – UDPPORT0 Filter is enabled
13	RW	0	DHCPv6. This field enables the UDP port filter 546 (decimal). 0 – DHCPv6 Filter is disabled 1 – DHCPv6 Filter is enabled
14	RW	0	EAPoUDP This field enables the UDP port filter 21862 (decimal). 0 – EAPoUDP Filter is disabled 1 – EAPoUDP Filter is enabled
15	RW	0	DNS. This field enables the TCP/UDP source port filter 53 (decimal). 0 – DNS Filter is disabled 1 – DNS Filter is enabled
16	RW	0	UDPIPPORT0. This field enables the UDP port filter 1900 (decimal) 0 – UDPIPPORT0 Filter is disabled 1 – UDPIPPORT0 Filter is enabled
17	RW	0	IPPORT0. This field enables the UDPIPPORT0 port filter to be restricted with IPv4 239.255.255.250 or IPv6 FF02::C 0 – IPPORT0 Filter is disabled 1 – IPPORT0 Filter is enabled
18	RW	0	UDPIPPORT1. This field enables the UDP port filter 3702 (decimal) 0 – UDPIPPORT1 Filter is disabled 1 – UDPIPPORT1 Filter is enabled
19	RW	0	IPPORT1. This field enables the UDPIPPORT1 port filter to be restricted with IPv4 239.255.255.250 or IPv6 FF02::C 0 – IPPORT1 Filter is disabled 1 – IPPORT1 Filter is enabled
20	RW	0	UDPIPPORT2. This field enables the UDP port filter 5355 (decimal) 0 – UDPIPPORT2 Filter is disabled 1 – UDPIPPORT2 Filter is enabled
21	RW	0	IPPORT2. This field enables the UDPIPPORT2 port filter to be restricted with IPv4 224.0.0.252 or IPv6 FF02:0:0:0:0:0:1:3 0 – IPPORT2 Filter is disabled 1 – IPPORT2 Filter is enabled
22	RW	0	UDPIPPORT3. This field enables the UDP port filter 5353 (decimal) 0 – UDPPORT3 Filter is disabled 1 – UDPPORT3 Filter is enabled
23	RW	0	IPPORT3. This field enables the UDPIPPORT3 port filter to be restricted with IPv4 224.0.0.251 or IPv6 FF02::FB 0 – IPPORT3 Filter is disabled 1 – IPPORT3 Filter is enabled
25:24	RW	00	ENMRFUTPF. This field enables the flexible range UDP/TCP port filter: 00 – MRFUTPF Filter is disabled 01 – MRFUTPF defines a UDP range port numbers 10 – MRFUTPF defines a TCP range port numbers 11 – MRFUTPF defines a UDP or TCP range port numbers
27:26	RO	0	Reserved
28	RW	0	ETHERTYPE0. This field enables the Ethertype filter 0x8863 0 – ETHERTYPE0 Filter is disabled 1 – ETHERTYPE0 Filter is enabled
29	RW	0	ETHERTYPE1. This field enables the Ethertype filter 0x8864 0 – ETHERTYPE1 Filter is disabled 1 – ETHERTYPE1 Filter is enabled



30	RW	0	ETHERTYPE2 . This field enables the Ethertype filter 0x88D9 0 - ETHERTYPE2 Filter is disabled 1 - ETHERTYPE2 Filter is enabled
31	RO	0	Reserved

11.1.3.2.8 Management Control Register—MANC2H2 (0x00028; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05870.

The MANC2H2 register allows routing of manageability packets to the host based on the filter type that routed it to the manageability micro-controller. Each manageability filter defined in MANC3 has a corresponding bit in the MANC2H2 register. When a filter in MANC3 routes a packet to the ME, the same packet is also routed to the host if the respective filter bit is set in the MANC2H2 register and if the EN_MNG2HOST bit is set. The EN_MNG2HOST bit serves as a global enable for the MANC2H/MANC2H2 bits.

Bits	Type	Reset	Description
0	RW	0	TCPPORT0 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
1	RW	0	TCPPORT1 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
2	RW	0	TCPPORT2 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
3	RW	0	TCPPORT3 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
4	RW	0	TCPPORT4 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
5	RW	0	TCPPORT5 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
6	RW	0	TCPPORT6 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
7	RW	0	TCPPORT7 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
8	RW	0	TCPPORT8 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
9	RW	0	TCPPORT9 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
10	RW	00	TCPPORT10 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
11	RO	0	Reserved
12	RW	0	UDPPORT0 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
13	RW	0	DHCPv6 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
14	RW	0	EAPoUDP When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
15	RW	0	DNS . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
16	RW	0	UDPIPPORT0 . When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
17	RO	0	Reserved



18	RW	0	UDPIPPORT1. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
19	RO	0	Reserved
20	RW	0	UDPIPPORT2. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
21	RO	0	Reserved
22	RW	0	UDPIPPORT3. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
23	RO	0	Reserved
24	RW	0	MRFUTPF. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
25	RW	0	Host MTA. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
26	RW	0	ICMPv6. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
27	RO	0	Reserved
28	RW	0	ETHERTYPE0. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
29	RW	0	ETHERTYPE1. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
30	RW	0	ETHERTYPE2. When set indicates that packets that are routed to the ME due to this filter will be sent to the HOST as well.
31	RO	0	Reserved

11.1.3.2.9 IPv4 ME Binding Control—IPV4MBC (0x0002C; RW)

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCPSPORT0). When set, only packets that match the ME IPv4 filter may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCPSPORT1). When set, only packets that match the ME IPv4 filter may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCPSPORT2). When set, only packets that match the ME IPv4 filter may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCPSPORT3). When set, only packets that match the ME IPv4 filter may pass TCPSPORT3 filtering.
4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCPSPORT4). When set, only packets that match the ME IPv4 filter may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCPSPORT5). When set, only packets that match the ME IPv4 filter may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCPSPORT6). When set, only packets that match the ME IPv4 filter may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCPSPORT7). When set, only packets that match the ME IPv4 filter may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCPSPORT8). When set, only packets that match the ME IPv4 filter may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCPSPORT9). When set, only packets that match the ME IPv4 filter may pass TCPSPORT9 filtering.
10	RW	00	Enable IP filter for TCPSPORT10 (EN_IPTCPSPORT10). When set, only packets that match the ME IPv4 filter may pass TCPSPORT10 filtering.
11	RO	0	Reserved



12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the ME IPv4 filter may pass UDPPORT0 filtering.
13	RO	0	Reserved
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the ME IPv4 filter may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the ME IPv4 filter may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the ME IPv4 filter may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the ME IPv4 filter may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the ME IPv4 filter may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the ME IPv4 filter may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the ME IPv4 filter may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the ME IPv4 filter may pass flex port 5 filtering.
22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the ME IPv4 filter may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the ME IPv4 filter may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the ME IPv4 filter may pass flex port 8 filtering.
25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the ME IPv4 filter may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the ME IPv4 filter may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the ME IPv4 filter may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFTP (EN_IPMRFTP). When set, only packets that match the ME IPv4 filter may pass MRFTP range port filtering.
29	RW	0	Enable IP filter for ICMPv4. When set, only packets that match the ME IPv4 filter may pass ICMPv4 filtering.
30	RW	0	Enable IP filter for ARP. When set, only packets that match the ME IPv4 filter may pass ARP request filtering.
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the ME IPv4 filter may pass RMCP filtering.

11.1.3.2.10 IPv4 Host Binding Control—IPV4HBC (0x0003C; RW)

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCPSPORT0). When set, only packets that match the Host IPv4 filters may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCPSPORT1). When set, only packets that match the Host IPv4 filters may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCPSPORT2). When set, only packets that match the Host IPv4 filters may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCPSPORT3). When set, only packets that match the Host IPv4 filters may pass TCPSPORT3 filtering.



4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCPSPORT4). When set, only packets that match the Host IPv4 filters may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCPSPORT5). When set, only packets that match the Host IPv4 filters may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCPSPORT6). When set, only packets that match the Host IPv4 filters may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCPSPORT7). When set, only packets that match the Host IPv4 filters may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCPSPORT8). When set, only packets that match the Host IPv4 filters may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCPSPORT9). When set, only packets that match the Host IPv4 filters may pass TCPSPORT9 filtering.
10	RW	00	Enable IP filter for TCPSPORT10 (EN_IPTCPSPORT10). When set, only packets that match the Host IPv4 filters may pass TCPSPORT10 filtering.
11	RO	0	Reserved
12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the Host IPv4 filters may pass UDPPORT0 filtering.
13	RO	0	Reserved
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the Host IPv4 filters may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the Host IPv4 filters may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the Host IPv4 filters may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the Host IPv4 filters may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the Host IPv4 filters may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the Host IPv4 filters may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the Host IPv4 filters may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the Host IPv4 filters may pass flex port 5 filtering.
22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the Host IPv4 filters may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the Host IPv4 filters may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the Host IPv4 filters may pass flex port 8 filtering.
25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the Host IPv4 filters may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the Host IPv4 filters may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the Host IPv4 filters may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFUTPF (EN_IPMRFUTPF). When set, only packets that match the Host IPv4 filters may pass MRFUTPF range port filtering.
29	RW	0	Enable IP filter for ICMPv4. When set, only packets that match the Host IPv4 filters may pass ICMPv4 filtering.
30	RW	0	Enable IP filter for ARP. When set, only packets that match the Host IPv4 filters may pass ARP request filtering.
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the Host IPv4 filters may pass RMCP filtering.



11.1.3.2.11 IPv6 Binding Control—IPV6BC (0x00030; RW)

Bits	Type	Reset	Description
0	RW	0	Enable IP filter for TCPSPORT0 (EN_IPTCPSPORT0). When set, only packets that match the L24IPV6 filters may pass TCPSPORT0 filtering.
1	RW	0	Enable IP filter for TCPSPORT1 (EN_IPTCPSPORT1). When set, only packets that match the L24IPV6 filters may pass TCPSPORT1 filtering.
2	RW	0	Enable IP filter for TCPSPORT2 (EN_IPTCPSPORT2). When set, only packets that match the L24IPV6 filters may pass TCPSPORT2 filtering.
3	RW	0	Enable IP filter for TCPSPORT3 (EN_IPTCPSPORT3). When set, only packets that match the L24IPV6 filters may pass TCPSPORT3 filtering.
4	RW	0	Enable IP filter for TCPSPORT4 (EN_IPTCPSPORT4). When set, only packets that match the L24IPV6 filters may pass TCPSPORT4 filtering.
5	RW	0	Enable IP filter for TCPSPORT5 (EN_IPTCPSPORT5). When set, only packets that match the L24IPV6 filters may pass TCPSPORT5 filtering.
6	RW	0	Enable IP filter for TCPSPORT6 (EN_IPTCPSPORT6). When set, only packets that match the L24IPV6 filters may pass TCPSPORT6 filtering.
7	RW	0	Enable IP filter for TCPSPORT7 (EN_IPTCPSPORT7). When set, only packets that match the L24IPV6 filters may pass TCPSPORT7 filtering.
8	RW	0	Enable IP filter for TCPSPORT8 (EN_IPTCPSPORT8). When set, only packets that match the L24IPV6 filters may pass TCPSPORT8 filtering.
9	RW	0	Enable IP filter for TCPSPORT9 (EN_IPTCPSPORT9). When set, only packets that match the L24IPV6 filters may pass TCPSPORT9 filtering.
10	RW	00	Enable IP filter for TCPSPORT10 (EN_IPTCPSPORT10). When set, only packets that match the L24IPV6 filters may pass TCPSPORT10 filtering.
11	RO	0	Reserved
12	RW	0	Enable IP filter for UDPPORT0 (EN_IPUDPPORT0). When set, only packets that match the L24IPV6 filters may pass UDPPORT0 filtering.
13	RO	0	Enable IP filter for DHCPv6 (EN_IPDHCPv6). When set, only packets that match the L24IPV6 filters may pass DHCPv6 filtering.
14	RW	0	Enable IP filter for EAPoUDP (EN_IPEAPoUDP). When set, only packets that match the L24IPV6 filters may pass EAPoUDP filtering.
15	RW	0	Enable IP filter for DNS (EN_IPDNS). When set, only packets that match the L24IPV6 filters may pass DNS filtering.
16	RW	0	Enable IP filter for Flex port 0 (EN_IPFLEX0). When set, only packets that match the L24IPV6 filters may pass flex port 0 filtering.
17	RW	0	Enable IP filter for Flex port 1 (EN_IPFLEX1). When set, only packets that match the L24IPV6 filters may pass flex port 1 filtering.
18	RW	0	Enable IP filter for Flex port 2 (EN_IPFLEX2). When set, only packets that match the L24IPV6 filters may pass flex port 2 filtering.
19	RW	0	Enable IP filter for Flex port 3 (EN_IPFLEX3). When set, only packets that match the L24IPV6 filters may pass flex port 3 filtering.
20	RW	0	Enable IP filter for Flex port 4 (EN_IPFLEX4). When set, only packets that match the L24IPV6 filters may pass flex port 4 filtering.
21	RW	0	Enable IP filter for Flex port 5 (EN_IPFLEX5). When set, only packets that match the L24IPV6 filters may pass flex port 5 filtering.
22	RW	0	Enable IP filter for Flex port 6 (EN_IPFLEX6). When set, only packets that match the L24IPV6 filters may pass flex port 6 filtering.
23	RW	0	Enable IP filter for Flex port 7 (EN_IPFLEX7). When set, only packets that match the L24IPV6 filters may pass flex port 7 filtering.
24	RW	0	Enable IP filter for Flex port 8 (EN_IPFLEX8). When set, only packets that match the L24IPV6 filters may pass flex port 8 filtering.



25	RW	0	Enable IP filter for Flex port 9 (EN_IPFLEX9). When set, only packets that match the L24IPV6 filters may pass flex port 9 filtering.
26	RW	0	Enable IP filter for Flex port 10 (EN_IPFLEX10). When set, only packets that match the L24IPV6 filters may pass flex port 10 filtering.
27	RW	0	Enable IP filter for Flex port 11 (EN_IPFLEX11). When set, only packets that match the L24IPV6 filters may pass flex port 11 filtering.
28	RW	0	Enable IP filter for MRFUTPF (EN_IPMRFUTPF). When set, only packets that match the L24IPV6 filters may pass MRFUTPF range port filtering.
29	RW	0	Enable IP filter for ICMPv6. When set, only packets that match the L24IPV6 filters may pass ICMPv6 filtering.
30	RO	0	Reserved
31	RW	0	Enable IP filter for RMCP. When set, only packets that match the L24IPV6 filters may pass RMCP filtering

11.1.3.2.12 SHRA Filter Enable Register—SHRAFER (0x00034; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x0582C.

Bits	Type	Reset	Description
0	RW	0	EN_SHRA0_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[0] register and the Rx packet matches SHRA[0] the packet will be routed to the ME.
1	RW	0	EN_SHRA1_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[1] register and the Rx packet matches SHRA[1] the packet will be routed to the ME.
2	RW	0	EN_SHRA2_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[2] register and the Rx packet matches SHRA[2] the packet will be routed to the ME.
3	RW	0	EN_SHRA3_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[3] register and the Rx packet matches SHRA[3] the packet will be routed to the ME.
4	RW	0	EN_SHRA4_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[4] register and the Rx packet matches SHRA[4] the packet will be routed to the ME.
5	RW	0	EN_SHRA5_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[5] register and the Rx packet matches SHRA[5] the packet will be routed to the ME.
6	RW	0	EN_SHRA6_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[6] register and the Rx packet matches SHRA[6] the packet will be routed to the ME.
7	RW	0	EN_SHRA7_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[7] register and the Rx packet matches SHRA[7] the packet will be routed to the ME.
8	RW	0	EN_SHRA8_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[8] register and the Rx packet matches SHRA[8] the packet will be routed to the ME.
9	RW	0	EN_SHRA9_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[9] register and the Rx packet matches SHRA[9] the packet will be routed to the ME.
10	RW	0	EN_SHRA10_FILTER. Enables Shared MAC address filtering. When this bit set and the AV bit is set in the SHRAH[10] register and the Rx packet matches SHRA[10] the packet will be routed to the ME.
31:11	RO	0x0	Reserved



11.1.3.2.13 MAC To SMBus Control Register—M2SCTRL (0x00040; RW)

Bits	Type	Reset	Description
0	RW	1	LCD Interface Control (LCDIC). When clear to 0 overrides the LC-LCD interface to the SMBus (Force SMBus), although PE_RST_N is set to 1, asserts on PE_RST_N.PE_RST_N, LCDIC - LCD Interface 0, 0 - SMBus 0, 1 - SMBus 1, 0 - SMBus 1, 1 - PCIe
2:1	RW	1	Selects the target ESMT to be used for MAC/PHY SMBus communication 00 - Reserved 01 - ESMT1 10 - ESMT2 11 - ESMT3
3	RO	0	Reserved
7:4	RW	0xC	IPG. Amount of IPG (inter-packet Gap) (idle clocks) inserted into Rx burst FIFO from M2S.
31:2	RO	0	Reserved

11.1.3.2.14 MAC To SMBus Configuration Register—M2SCFG (0x00044; RW)

Bits	Type	Reset	Description
15:0	RW	0x0141	SMT Target ID (SMTTI). Indicates the target ID to be used by the LAN Controller on MBB transactions to the LCD through SMBus
16	RW	0	LCD SMBus Address Valid
23:17	RW	0x64	LCD SMBus Address (LCDSMBA)
24	RW	0	LC SMBus Address Valid
31:25	RW	0	LC SMBus Address (LCSMBA)

11.1.3.2.15 MAC To SMBus Transmit Control Register—M2STX (0x00048; RW)

Bits	Type	Reset	Description
6:0	RW	0x20	SMBus Data Length (SMBDL). Indicates the maximum data length (bytes) that can be used in SMBus transactions from LC to LCD. Allowed configurations: 0x40 - 64 bytes of data 0x30 - 48 bytes of data 0x20 - 32 bytes of data 0x10 - 16 bytes of data
31:7	RO	0	Reserved

11.1.3.2.16 H2ME Mask—H2MEM (0x00050; RW)

Bits	Type	Reset	Description
31:0	RW	0	Interrupt mask for the relevant H2ME bit ME FW need to set a bit to mask the matching H2ME bit interrupt generation



11.1.3.2.17 Extended Firmware Semaphore Register—EXFWSM (0x00060; RW)

This register is RW to ME. It has a shadow RO image in the Host CSR space at offset 0x05B58. The EXFWSM register is initialized only by LAN Power Good.

Bit	Type	Reset	Description
31:0	RW	0x0	Reserved for future use

11.1.3.2.18 Extended Device Control Shadow—CTRL_EXT_S (0x000A0; RO)

This register is a RO Shadow register of the CTRL_EXT register in the PCI CSR space.

11.1.3.2.19 MDI Control Shadow—MDIC_S (0x000A4; RO)

This register is a RO Shadow register of the MDIC register in the PCI CSR space.

11.1.3.2.20 Receive Address Low 0 Shadow—RAL0_S (0x000A8; RO)

This register is a RO Shadow register of the RAL0 register in the PCI CSR space.

11.1.3.2.21 Receive Address High 0 Shadow—RAH0_S (0x000AC; RO)

This register is a RO Shadow register of the RAH0 register in the PCI CSR space.

11.1.3.2.22 Kumeran Control and Status Shadow—GENERAL_S (0x000B0; RO)

This register is a RO Shadow register of some bits from the KUMERAN control registers in the PCI CSR space.

Bits	Type	Reset	Description
0	RO	0b	Shadow of: Far end Loopback. Digital far-end loopback mode. Connects the Rx GMII to the Tx GMII.
1	RO	0b	Shadow of: Near end Loopback (NELPBK). Digital near-end loopback mode. Connects the 8B/10B encoder (Tx) to the 8B/10B decoder (Rx).
2	RO	0b	Shadow of: Driver loaded (DRV_LOAD) bit from CTRL_EXT register
31:3	RO	0	Reserved



11.1.3.2.23 Flex Write to Interrupt Control—FLXWINTC (0x000C0; RW)

Bits	Type	Reset	Description
0	RW	0x0	Enable Capture CSR Register 0
1	RW	0x0	Enable Capture CSR Register 1
2	RW	0x0	Enable Capture CSR Register 2
3	RW	0x0	Enable Capture CSR Register 3
4	RW	0x0	Enable Capture Phy Register 0
5	RW	0x0	Enable Capture Phy Register 1
6	RW	0x0	Enable Capture Phy Register 2
7	RW	0x0	Enable Capture Phy Register 3
15:8	RO	0x00	Reserved
16	RW	0x0	Enable Capture to CTRL_EXT Register
17	RW	0x0	Enable Capture to MDIC Register
18	RW	0x0	Enable Capture to RAL/H Register
19	RW	0x0	Enable Capture to GENERAL_S Register
31:20	RW	0x00	Reserved

11.1.3.2.24 Flex CSR Write to Interrupt 01—FLXCWINT01 (0x000C4; RW)

Bits	Type	Reset	Description
15:0	RW	0x00	CSR Capture Register 0
31:16	RW	0x00	CSR Capture Register 1

11.1.3.2.25 Flex CSR Write to Interrupt 23—FLXCWINT23 (0x000C8; RW)

Bits	Type	Reset	Description
15:0	RW	0x00	CSR Capture Register 2
31:16	RW	0x00	CSR Capture Register 3

11.1.3.2.26 Flex Phy Write to Interrupt 01—FLXPWINT01 (0x000CC; RW)

Bits	Type	Reset	Description
4:0	RW	0x00	Phy Capture REGADD 0
9:5	RW	0x00	Phy Capture PHYADD 0
11:10	RW	0x0	Phy Capture OP 0
15:12	RO	0x0	Reserved
20:16	RW	0x00	Phy Capture REGADD 1
25:21	RW	0x00	Phy Capture PHYADD 1
27:26	RW	0x0	Phy Capture OP 1
31:28	RO	0x0	Reserved

11.1.3.2.27 Flex Phy Write to Interrupt 23—FLXPWINT23 (0x000D0; RW)

Bits	Type	Reset	Description
4:0	RW	0x00	Phy Capture REGADD 2



10:6	RW	0x00	Phy Capture PHYADD 2
12:11	RW	0x0	Phy Capture OP 2
15:13	RO	0x0	Reserved
20:16	RW	0x00	Phy Capture REGADD 3
25:21	RW	0x00	Phy Capture PHYADD 3
27:26	RW	0x0	Phy Capture OP 3
31:28	RO	0x0	Reserved

11.1.3.3 Interrupt Register Descriptions

11.1.3.3.1 Management Interrupt Cause Register—MICR (0x0080; RC/WC)

This interrupt cause register reflects various interrupt sources related to the LAN interface. This register is Write-Clear. Writing a '1' clears the written bit while writing '0' has no affect. An active high level on each bit field triggers an MSI signaling to the ME (if it is enabled in the MIER). Once an MSI message is sent to the ME, the relevant bit in the MICR is auto-cleared to '0'. In case a bit is cleared and a new event is sent at the same time, the bit stays active at '1' and an additional MSI will be sent. When multiple bits in the MICR are set at the same time, the LAN controller will initiate multiple MSI messages in a round robin prioritization.

An MSI message includes a Data (operand) associated with the specific interrupt cause. The LAN controller captures the data once it sets the interrupt bit in the MICR. In case there is an additional interrupt event before an MSI was sent, the most recent event will be captured as the data to be reflected in the MSI operand.

Bits	Type	Reset	Description
0	RWC/ V	0	Receive Frame In (RFI). New manageability received frame from LAN is in the ME receive circular queue (i.e. DMA complete and the Head pointer is updated).
1	RWC/ V	0	XMT Frame Out (XFO). Full XMT frame was DMAed from ME memory space to the internal XMT FIFO.
2	RWC/ V	0	LinkSec Packet Number. The Tx Packet Number hit the "PN exhaustion threshold" as defined in the LSECTXCTRL register and the ME is the KaY.
3	RWC/ V	0	Link Status changed (LSC). Asserted when the Link status is changed.
4		0	FMSW flag. FMSW flag is set
5	RWC/ V	0	SWRST Event. Asserted following setting of the SWRST bit in the CTRL register.
6	RWC/ V	0	Host Write Event. Asserted following host write to pre defined set of registers.
7	RWC/ V	0	LCD Reset Event. Asserted following setting of the LCD_RST bit in the CTRL register.
8	RWC/ V	0	LAN Init Done. Asserted following completion of the LAN initialization from the FLASH. See "LAN Init Done Event" section for a complete description.
9	RWC/ V	0	PCI State Change. Asserted following a change of the PCI power state as reflected in the FACTPS register.
10	RWC/ V	0	Host Interrupt. Asserted following setting of the H2MEINT bit in the CTRL register.



11	RWC/ V	0	Host NVM Cycle Event (NVMCYC). Asserted when the host initiates a Program (Write) or Erase cycle to the FLASH on the LAN host interface.
12	RWC/ V	0	Tx Circuit Breaker (TxCB). Asserted when a transmitted packet hits a circuit breaker filter that is enabled to generate an interrupt.
13	RWC/ V	0	Rx Circuit Breaker (RxCB). Asserted when a received packet hits a circuit breaker filter that is enabled to generate an interrupt.
14	RWC/ V	0	Tx Header (TxHCB). Asserted when a transmitted packet's header is posted to ME host memory and enabled to generate an interrupt.
15	RWC/ V	0	H2ME Interrupt. Asserted when an un-masked bit in the H2ME changes (set or clear)
31:16	RO	0x0	Reserved.

11.1.3.3.2 Management Interrupt Enable Register—MIER (0x0088; RW)

Each bit that is set in this register enables the relevant bit in the ICR to set an interrupt.

Default value of the IER is 0x0.

Bits	Type	Reset	Description
0	RW	0	RFI. When set to '1' this Interrupt is enabled.
1	RW	0	XFO. When set to '1' this Interrupt is enabled.
2	RW	0	LinkSec Packet Number. When set to '1' this Interrupt is enabled.
3	RW	0	LSC. When set to '1' this Interrupt is enabled.
4	R	0	FMSW flag Event
5	RW	0	SWRST Event. When set to '1' this Interrupt is enabled.
6	RW	0	Host Write Event. When set to '1' this Interrupt is enabled.
7	RW	0	LCD Reset Event. When set to '1' this Interrupt is enabled.
8	RW	0	LAN Init Done. When set to '1' this Interrupt is enabled.
9	RW	0	PCI State Change. When set to '1' this Interrupt is enabled.
10	RW	0	Host Interrupt. When set to '1' this Interrupt is enabled.
11	RW	0	NVMCYC. When set to '1' this Interrupt is enabled.
12	RW	0	TxCB. When set to '1' this Interrupt is enabled.
13	RW	0	RxCB. When set to '1' this Interrupt is enabled.
14	RW	0	TxHCB. When set to '1' this Interrupt is enabled.
15	RW	0	H2ME Interrupt. When set to '1' this interrupt is enabled
31:16	RO	0x0	Reserved.

11.1.3.4 Receive Register Descriptions

11.1.3.4.1 Management Receive Buffer Control Register—MRBCR (0x0100; RW)

Bits	Type	Reset	Description
0	RW	0x0	Enable (En). When set it enables the Receive operation. <i>Note:</i> Buffer Base and Last should be changed only when the enable bit is not set.



1	RW/V	0	Reset Buffer (RB) . When set, it resets the current Receive operation and resets the head and tail pointers to the value of the base address—MRBBA. Self cleared by HW. This bit should be set by Firmware together with the RF bit in this register.
2	RW/V	0	Reset FIFO (RF) . When set, this resets the management receive FIFO pointers and the transmit redirection FIFO pointers, thus flushing the FIFO(s). It also aborts any reception from the wire or Circuit Breaker redirection to the FIFO(s). Self cleared by HW. This bit should be set by Firmware together with the RB bit in this register.
30:3	RO	0x0	Reserved.
31	RW	1	GBT CB Priority . Selects the host access arbitration priority of the receive path between normal RX management packets and redirected CB packets. When = 0 a 1:1 ratio is selected; when = 1 a 4:1 ratio (Mng to CB) is selected.

Note: When the enable bit is not set than the HW will drop all the packets that are received from the Rx-Filter.

Note: When the FW wants to abort reception it should follow the steps described in [Section 8.9.2.3](#).

11.1.3.4.2 Management Receive Buffer Base Address—MRBBA (0x0108; RW)

Bits	Type	Reset	Description
3:0	RO	0x0	Reserved.
31:4	RW	0x0	Buffer Base Address (BBA) . The base address of the receive buffer in the ME memory space. The Head and Tail registers get the base register value when the base register is being written. The pointer is defined in 16 bytes granularity.

11.1.3.4.3 Management Receive Buffer Last Address—MRBLA (0x0110; RW)

Bits	Type	Reset	Description
5:0	RO	0x0	Reserved.
31:6	RW/V	0x0	Buffer Last Address (BLA) . The last memory line of the receive buffer. The Buffer size must be 64 byte aligned therefore the 6 LS bits in this register are all zeroes. The register may be initialized to MRBBA when the MRBBA is set.

11.1.3.4.4 Management Receive Buffer Head Register—MRBHR (0x0120; RO)

This register is read only by FW. The head is incremented by the HW when a new line is stored in the ARC memory, and not in packet granularity.

Bits	Type	Reset	Description
3:0	RO	0x0	Reserved.
31:4	RO/V	0x00	Buffer Head (BH) . This value points to the next memory line that the HW will write the data into. The pointer is 16 bytes line aligned (Head pointer is an absolute pointer). The register may be initialized to MRBBA when the MRBBA is set or by Reset Buffer bit in the MRBCR register.



11.1.3.4.5 Management Receive Buffer Tail Register—MRBTR (0x0128; RW)

This register is controlled by FW. The HW can write data until (Tail-1). It is recommended that the FW change the value of the tail register in packet granularity.

Bits	Type	Reset	Description
3:0	RO	0x0	Reserved.
31:4	RW/V	0x00	Buffer Tail (BT). This value points to the next memory line that the FW will read. The HW can write to the buffer until address BT-1. The pointer must be 16 bytes line aligned (Tail pointer is an absolute pointer). The register may be initialized to MRBBA when the MRBBA is set or by Reset Buffer bit in the MRBCR register.

11.1.3.4.6 Management Range Flex UDP/TCP Ports Filter—MRFUTPF (0x0144; RW)

Bits	Type	Reset	Description
15:0	RW	0	MRFUTPFH. Management Range Flex UDP/TCP ports Filter High This port filter marks the highest port value for the range port filter
31:16	RW	0	MRFUTPFL. Management Range Flex UDP/TCP ports Filter Low This port filter marks the Lowest port value for the range port filter

11.1.3.4.7 Management Flex UDP/TCP Ports 0/1—MFUTP01 (0x0148; RW)

Bits	Type	Reset	Description
15:0	RW	0	MFUTP0. Management Flex UDP/TCP port This filter is enabled by the FLEXP0 field in the MANC register.
31:16	RW	0	MFUTP1. Management Flex UDP/TCP port This filter is enabled by the FLEXP1 field in the MANC register.

11.1.3.4.8 Management Flex UDP/TCP Port 2/3—MFUTP23 (0x0150; RW)

Bits	Type	Reset	Description
15:0	RW	0	MFUTP2. Management Flex UDP/TCP port This filter is enabled by the FLEXP2 field in the MANC register.
31:16	RW	0	MFUTP3. Management Flex UDP/TCP port This filter is enabled by the FLEXP3 field in the MANC register.



11.1.3.4.9 Management Flex UDP/TCP Port 4/5—MFUTP45 (0x0158; RW)

Bits	Type	Reset	Description
15:0	RW	0	MFUTP4. Management Flex UDP/TCP port This filter is enabled by the FLEXPOR4 field in the MANC2 register.
31:16	RW	0	MFUTP5. Management Flex UDP/TCP port This filter is enabled by the FLEXPOR5 field in the MANC2 register.

11.1.3.4.10 Management IPv6—MIPV6 (0x0160 + 0x4*n (n=0..3); RW)

Bits	Type	Reset	Description
23:0	RW	0	L24IPV6. This field holds the lower 24 bits of the IPv6 address to be compared with any IPv6 incoming packet, field should be set in network order.
31:24	RO	0	Reserved

11.1.3.4.11 TCO Flexible Filter Length Table—TCOFFLT (0x0180 + 8*n (n=0...1); RW)

The TCO Flexible Filter Length Table stores the minimum packet lengths required to pass each of the TCO Flexible Filters. Any packets that are shorter than the programmed length won't pass that filter. Each Flexible Filter will consider a packet that doesn't have any mismatches up to that point to have passed the Flexible Filter when it reaches the required length. It will not check any bytes past that point.

Bits	Type	Reset	Description
10:0	RW	X	LEN TCO. Min Length for flexible TCO filter n (n=0, 1)
31:11	RW	X	Reserved.

All reserved fields read as 0's and ignore writes.

11.1.3.4.12 TCO Flexible Filter DA Index- TCOFFDAI (0x0190; RW)

Bits	Type	Reset	Description
3:0	RW	0x0	TCO0_DA_IDX—Index of RAL/RAH that should match in addition to TCO_0 flex filter if the correspondent mask bit is cleared
7:4	RO	0x0	Reserved
8	RW	0x0	Mask bit for TCO0_DA_IDX
15:9	RO	0x0	Reserved
19:16	RW	0x0	TCO1_DA_IDX—Index of RAL/RAH that should match in addition to TCO_1 flex filter if the correspondent mask bit is cleared
23:20	RO	0x0	Reserved
24	RW	0x0	Mask bit for TCO1_DA_IDX
31:25	RO	0x0	Reserved



11.1.3.4.13 Flexible TCO Filter Table—FTFT (0x0800 + 8*n (n=0...127); RW)

There are 128 entries to the TCO filter table. These registers can be used by SW to update the flex-TCO filter bytes that should be compared. As opposed to the wake-up table this structure contains the byte value and the bit mask in the same address. The enable bits for the engines are in the MANC register.

Bits 7:0 and 8 are used for flex TCO filter 0 and bits 16:9 and 17 are used for flex TCO filter 1.

Bits	Type	Reset	Description
7:0	RW	X	TCO0 Byte n. Flexible TCO filter 0, Byte n (n=0, 1... 127)
8	RW	X	TCO0 MASK. Flexible TCO filter 0, Mask n (n=0, 1... 127)
16:9	RW	X	TCO1 Byte n. Flexible TCO filter 1, Byte n (n=0, 1... 127)
17	RW	X	TCO1 MASK. Flexible TCO filter 1, Mask n (n=0, 1... 127)
31:18	RO	X	Reserved.

11.1.3.4.14 Management VLAN TAG Value Table—MAVTV[n] (0x01A0+4*n (n=0...3); RW)

Bits	Type	Reset	Description
11:0	RW	0x0	VLAN ID. Contains the VLAN ID that should be compared with the incoming packet if bit 31 is set.
30:12	RO	0x0	Reserved.
31	RW	0x0	En. Enable VID filtering

11.1.3.4.15 Management IPv4 Address Table—MIP4AT (0x01C0; RW)

The Management IPv4 Address Table is used to store the Management IPv4 address filter for ARP/IPv4 Request packets and Directed IPv4 packet wake up.

Bits	Type	Reset	Description
31:0	RW	X	MIP4ADD. Management IPv4 Address filter (L.S. byte is first on the wire)

11.1.3.5 Transmit Register Descriptions

11.1.3.5.1 Management Transmit Buffer Control Register—MTBCR (0x0200; RW)

Bits	Type	Reset	Description
0	RW	0x0	Enable (En). When set it enables the Transmit operation. Note: Buffer Base and Last should be changed only when the enable bit is not set.
1	RW/V	0	Reset Buffer (RB). When set it resets the current transmit operation and resets the head and tail pointer to the value of the base address—MTBBA. Self cleared by HW. This bit should be set by Firmware together with the RF bit in this register.



2	RW/V	0	Reset FIFO (RF). When set, this resets the management transmit path FIFO pointers / logic; flushing all data in the internal management memories. The transmit path should be disabled before asserting this bit otherwise results may be indeterminate. Self cleared by HW. This bit should be set by Firmware together with the RB bit in this register.
31:3	RO	0x0	Reserved.

Note: When the FW wants to abort a transmission, it should follow the steps described in [Section 8.9.2.3](#). To de-assert the enable bit after it was set, the FW should first set the RB bit. The MMS HW should indicate to the NW XMT unit that the current packet is aborted (the XMT unit will send the current packet with bad CRC value).

11.1.3.5.2 Management Transmit Buffer Base Address—MTBBA (0x0208; RW)

Bits	Type	Reset	Description
3:0	RO	0x0	Reserved.
31:4	RW	0x0	Buffer Base Address (BBA). The base address of the Transmit buffer in the ME memory space. The Head and Tail registers get the base register value when the base register is being written. The address is defined in 16-byte granularity.

Note: Following a PCIm Reset or PCIm SW Reset(MGCR.RESET) or clearing the Tx Enable (MTBCR.En) or clearing the Tx FIFO (MTBCR.RF or MTBCR.RB) the FW should avoid changing the setting of the Transmit buffer address registers for 2msec.

11.1.3.5.3 Management Transmit Buffer Last Address—MTBLA (0x0210; RW)

Bits	Type	Reset	Description
5:0	RO	0x0	Reserved.
31:6	RW/V	0x0	Buffer Last Address (BLA). The last memory line of the transmit buffer. The Buffer size must be 64 byte aligned, therefore, the 6 LS bits in this register are all zeroes. The register may be initialized to MTBBA when the MTBBA is set.

Note: Following a PCIm Reset or PCIm SW Reset(MGCR.RESET) or clearing the Tx Enable (MTBCR.En) or clearing the Tx FIFO (MTBCR.RF or MTBCR.RB) the FW should avoid changing the setting of the Transmit buffer address registers for 2msec.

11.1.3.5.4 Management Transmit Buffer Head Register—MTBHR (0x0220; RW)

This register is controlled by FW. The FW should increment the head value in a packet granularity, which means that only when a full packet is written into the buffer, and ready to be sent, then the head register is incremented.

The Head register gets the base register value when the base register is being written or when the "RB" bit is set in the TBCR.

Bits	Type	Reset	Description
3:0	RO	0x00	Reserved.



Bits	Type	Reset	Description
31:4	RW/V	0x00	Buffer Head (BH). This value points to the next memory line that the FW will write data into. The FW can increment the head up to one line after the tail register so that the head does not point inside the area owned by the HW. As a result, the transmit buffer can be utilized by the FW up to its size minus one line. The pointer must be 16 bytes granularity (head pointer is an absolute pointer). The register may be initialized to MTBBA when the MTBBA is set or by Reset Buffer bit in the MTBCR register.

Note: Following a PCIm Reset or PCIm SW Reset(MGCR.RESET) or clearing the Tx Enable (MTBCR.En) or clearing the Tx FIFO (MTBCR.RF or MTBCR.RB) the FW should avoid changing the setting of the Transmit buffer address registers for 2msec.

11.1.3.5.5 Management Transmit Buffer Tail Register—MTBTR (0x0228; RO)

This register is read only by FW. The FW can write data into the FIFO until (Tail-1). The HW changes the value of the tail register in packet granularity.

The Tail register gets the base register value when the base register is being written or when the "RB" bit is set in the TBCR.

Bits	Type	Reset	Description
3:0	RO	0x00	Reserved.
31:4	RO/V	0x00	Buffer Tail (BT). This value points to the next memory line that the HW will read. The FW can write to the buffer until address BT-1. The pointer is in 16 bytes granularity (Tail pointer is an absolute pointer). The register may be initialized to MTBBA when the MTBBA is set or by Reset Buffer bit in the MTBCR register.

Note: Following a PCIm Reset or PCIm SW Reset(MGCR.RESET) or clearing the Tx Enable (MTBCR.En) or clearing the Tx FIFO (MTBCR.RF or MTBCR.RB) the FW should avoid changing the setting of the Transmit buffer address registers for 2msec.

11.1.3.6 PCI CSR Access Register Descriptions

11.1.3.6.1 Management to MAC CSR Control register—MCCR (0x0300; RW)

Bits	Type	Reset	Description
0	RW/V	0x0	Valid (V). Valid bit for read or write transaction on the MAC CSR interface. The 'V' bit is cleared by HW at the end of the transaction.
1	RW	0x0	Write (Wr). When set indicates a write transaction. Valid only when the V bit is set.
2	RW/V	0x1	Complete (C). HW indication that the transaction was completed OK. This bit is valid when the 'V' bit is not set.
31:3	RO	0x0	Reserved.

Note: If the "V" bit is cleared by HW, and the "C" bit is not set, the FW should wait 2000 cycles before accessing the MAC-CSR interface again.



11.1.3.6.2 Management to MAC CSR Address Register—MCAR (0x0310; RW)

Bits	Type	Reset	Description
16:0	RW	0x0	CSR Address (CA) . Address for the write/read transaction on the MAC CSR Interface. Valid when the V bit on the control register is set.
31:17	RO	0x0	Reserved.

11.1.3.6.3 Management to MAC CSR Data Out Register—MCDO (0x0318; RO)

This register will contain the data that was read by the MAC CSR interface. It is valid when the HW sets the C bit in the MMCR on read transaction until the next time the 'V' bit is set by FW.

Bits	Type	Reset	Description
31:0	RO	0x0	MAC CSR Data Out (MCDO) . Output data on MAC CSR read cycles.

11.1.3.6.4 Management to MAC CSR Data In Register—MCDI (0x0320; RW)

This register will contain the data that should be written by the MAC CSR interface. This register should be valid before the FW sets the 'V' bit on write transaction and must not be changed until the complete bit is set.

Bits	Type	Reset	Description
31:0	RW	0x0	MAC CSR Data In (MCDI) . Input data for the MAC CSR write cycles.

11.1.3.6.5 PCI Function Power State to MNG—FACTPS (0x0340; RO)

Register for use by the device FW for configuration.

Bits	Type	Reset	Description
31:2	RO	0	Reserved.
1:0	RO/V	00	PCI Power State (PCIPSTATE) . Power state indication of the LAN PCI function 00 - DR 01 - D0u 10 - D0a 11 - D3

11.1.3.7 Circuit Breaker Register Descriptions

11.1.3.7.1 Circuit Breaker Configuration Register—CBCR (0x1000; RW)

This register configures the global configuration mechanisms of the Circuit Breaker.

Bits	Type	Reset	Description
0	RW/SN	0	Rx Circuit Breaker Enable (RXCBE) . Enables the Rx circuit breaker filters.
1	RW/SN	0	Tx Circuit Breaker Enable (TXCBE) . Enables the Tx circuit breaker filters.



Bits	Type	Reset	Description
2	RO RW	0	Reserved.
3	RW	0	Receive Filter Priority (RFP). If set, receive PASS filters will override BLOCK filters in the case of an overlap. If cleared, receive BLOCK filters will override PASS filters.
4	RW	0	Transmit Filter Priority (TFP). If set, transmit PASS filters will override BLOCK filters in the case of an overlap. If cleared, transmit BLOCK filters will override PASS filters.
6:5	RW	0	Reserved.
7	RW	0	Tx Fragmented Header Redirection Enable (TXFRHRE). This bit controls header redirection action on the transmit data path of a first fragment of fragmented IP packets. When set to 1 the header redirection is enabled. If cleared header redirection is not enabled. Please refer also to the FTXFR field in this register. Note that filters that do not check L4 fields may match fragmented as well as non fragmented packets. In case of match, the header of packets that match these filters can be redirected regardless if the packets are fragmented and regardless of the value of this field (TXFRHRE).
8	RW	0	Force Tx Fragmented Header Redirection (FTXFR). When it is cleared, then the Tx fragmented header redirection follows the HSKIP field in the TXHUTIL register. When the FRXFR is set then header redirection ignores (override) the 1/N rule. Note: This bit may be active only if the TXFRHRE is active.
9	RO	0	Reserved.
10	RO	0	Reserved.
11	RW	0	Tx Fragments Block (TXFRBLK). This bit controls the action for fragmented packets. When set to 1 all transmitted fragmented packets will be blocked. If cleared (0) fragmented packets will follow the regular transmit CB filter behavior.
12	RW	0	Rx Fragments Block (RXFRBLK). This bit controls the action for fragmented packets. When set to 1 all received fragmented packets will be blocked. If cleared (0) fragmented packets will follow the regular receive CB filter behavior.
23:13	RO	0x0	Reserved.
28:24	RW	00000	Anti-Spoofing Filter Select (ASFS). Defines the transmit filters assigned for Anti-Spoofing. Each bit of the ASFS field enables a CB filter as Anti-Spoofing when set to 1. Bit 24 enables CB filter 0, bit 25 enables CB filter 1... bit 28 enables CB filter 4. A packet is considered as Spoof if its source IP address does not match any of the IP addresses defined by filters 0 through 4, that are enabled as Anti-Spoofing.
29	RO	00	Reserved.

11.1.3.7.2 CB Hash Seed Register—CBSEED (0x1010; RW)

Bits	Type	Reset	Description
31:0	RW	0xF..F	CB SEED. Seed value that is used by the CRC hash engine for IP addresses. The hash value is part of a packet's header that can be posted to ME. See "Header Redirection" section for more details. Note: In order to avoid inconsistency the FW should avoid CBSEED modification when header redirection may be active. That is, the HEN bits in the TXHCTL as well as the RXHCTL registers are cleared.



11.1.3.7.3 Circuit Breaker Counter Status—CBCS (0x1100; R/WC)

This register identifies which of the counters have reached their thresholds. Once the counter threshold is reached, the bit is set and is then only cleared by firmware writing a 1b to it.

Bits	Type	Reset	Description
31:0	RWC/C	0x00	Counter Status (CS). Each bit indicates the corresponding counter (i.e. bit 0 = counter #0). If the bit is set, the counter has reached the threshold value and gets back to zero at least once. This bit is cleared by writing a 1b to it.

11.1.3.7.4 Counter/Threshold Configuration—CBCTC[n] (0x1200 + 8*n (n=0...31); RW)

There is one instance of this register for each of the 32 counters. It is used to enable and configure the operation of the filter.

Bits	Type	Reset	Description
0	RW	0	Enable (EN). Enables this counter for use.
1	RW	0	Filter Type (FT). 1 – transmit filter 0 – receive filter
11:2	RO	0x0	Reserved.
12	RW	0x0	Filter Number (FN). Specifies the filter number that this counter/threshold is associated with as shown below. Filter Number = (n && 111110b) + FN (while && represent a logic AND operation). A counter can be associated to filters that are set to enable host traffic with the CBRC/CBTC.HEN bit set or redirecting packets to ME with the CBRC/CBTC.TOME bit set.
31:13	RO	0x0	Reserved.

11.1.3.7.5 Counter/Threshold Value—CBCTV[n] (0x1300 + 8*n (n=0...31); RW)

This register contains the threshold and counter values that can be read or written by the ME.

Bits	Type	Reset	Description
31:16	RW	0x00	Threshold (TH). The value that the incrementing counter will be compared against.
15:0	RW	0x00	Counter (CO). The running counter that is incremented at every time that the filter associated with it has a packet match.

11.1.3.7.6 Transmit Filter Configuration Registers—CBTC[n] (0x2000 + 8*n (n=0...31); RW)

There is one instance of this register for each of the 31 Transmit Filters. The 32nd entry is used to describe/configure the transmit “else” case. It is used to enable and configure the operation of the filter.

Configuration filters 0-30;



Bits	Type	Reset	Description
0	RW	0	Enable Host Filter (ENH). Enables this filter for host filtering.
1	RW	0	Interrupt Enable (IEN). Enables matching of this filter (or of the counter/threshold) associated with this filter to cause an interrupt to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will cause an interrupt. The IEN bit can be set only if the ENH bit is set as well for the same filter.
2	RW	0	Block (BLK). Indicates that this filter specifies packets that should be blocked for transmission from the host to the LAN. If cleared, it identifies a "pass" filter. This configuration applies to each and every packet.
3	RW	0	To ME (TOME). If set, the packet that passes this filter is sent to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will be sent. Note that a packet blocked by the anti spoofing filters is not passed to the ME even if the TOME bit is set in a passing filter.
4	RW	0	Header To ME (HTOME). If set, the header of the packet that matches this filter is sent to the ME.
5	RW	0	Force Header To ME (FHTOME). When FHTOME is set the header redirection to ME overrides the HSKIP mechanism. This field may be set only when the HTOME field is set as well. Note in the case that a packet matches multiple filters with active HTOME bit but only some of them have active FHTOME bit. The header redirection to ME overrides the HSKIP mechanism only if the filter with the lowest index (0...30) has an active FHTOME bit.
6	RW	0	QOSSET. This field determines whether the HW will override the DSCP value of an outgoing packet. 0b – Keep current QoS (DSCP) value 1b – Set QoS (DSCP) value as defined by QOSVALUE
7	RO	0	Reserved.
9:8	RW	00	Packet Type (PKTTYPE). Determines what type of filter this is: 00 – IP / TCP 01 – IP / UDP 10 – IP / Other 11 – Non-IP NOTE: This field is Reserved 11b for the "else" case.
15:10	RO	0x0	Reserved.
17:16	RW	00	IP Address Type (IPATYPE). Determines how the filter looks at the IP Address for IP/TCP, IP/UDP, and IP/Other packet types. 00 – No IP Address Compare 01 – Compare IP Address against Source Address 10 – Compare IP Address against Destination Address 11 – Reserved NOTE: This field is Reserved 11b for the "else" case.
23:18	RO	0x0	Reserved.
25:24	RW	00	Port Type (PORTTYPE). Determines how the filter looks at the Source and/or Destination Ports for IP/TCP or IP/UDP packet types. 00 – Ports ignored and not compared 01 – Source Port Range 10 – Destination Port Range 11 – Source and Destination Port Match NOTE: This field is Reserved 00b for the "else" case.
31:26	RW	0x0	QOSVALUE. This field determines the value that will be set in the DSCP field of an outgoing packet when the QOSSET bit is set to 1 and this CB filter matched the outgoing packet.



11.1.3.7.6.3 Transmit Filter Configuration—CBTC[31] (0x20F8; RW) The else configuration filter (filter number 31).

Bits	Type	Reset	Description
0	RO	1	Reserved. This field is used as the "Enables Host filter" in all other filters. The else filter is always enabled.
1	RW	0	Interrupt Enable (IEN). Enables matching of this filter (or of the counter/threshold) associated with this filter to cause an interrupt to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will cause an interrupt. The IEN bit can be set only if the ENH bit is set as well for the same filter.
2	RW/SN	0	Block (BLK). Indicates that this filter specifies packets that should be blocked for transmission. If cleared, it identifies a "pass" filter. This configuration applies to each and every packet that does not match any of the filters 0...30 with active ENH bit.
3	RW	0	To ME (TOME). If set, the packet that does not match any of the filters 0...30 with active TOME bit is sent to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will be sent to ME.
4	RW	0	Header To ME (HTOME). If set, the header of the packet that does not match any of the filters 0...30 with active HTOME bit is sent to the ME.
5	RW	0	Force Header To ME (FHTOME). When it is cleared, then the header follows the HSKIP field in the TXHUTIL register. When FHTOME is set the header is posted to the Header FIFO and ignores the HSKIP mechanism. Note in the case that a packet matches multiple filters with active HTOME bit but only some of them have active FHTOME bit. The header redirection to ME overrides the HSKIP mechanism only if the filter with the lowest index (0...30) has an active FHTOME bit.
6	RW	0	QOSSET. This field determines whether the HW will override the DSCP value of an outgoing packet. 0b – Keep current QoS (DSCP) value 1b – Set QoS (DSCP) value as defined by QOSVALUE
25:7	RO	0x0	Reserved.
31:26	RW	0x0	QOSVALUE. This field determines the value that will be set in the DSCP field of an outgoing packet when the QOSSET bit is set to 1 and this CB filter matched the outgoing packet.

11.1.3.7.7 Circuit Breaker Transmit Filter Status—CBTFS (0x2100; R/WC)

Bits	Type	Reset	Description
31:0	RWC/V	0x00	Transmit Filter Status (TFS). Each bit indicates the corresponding transmit filter (i.e. bit 0 = filter #0). If the bit is set, at least one packet has been transmitted from the host interface (or attempted transmitted) that hit this filter while either the "Enable Host Filter" or "To ME" bits are set in the CBTC register. This bit is cleared by writing a 1b to it. Bit 31 (else filter, index #31) is set for packets that do not match any other filter with active "Enable Host Filter" bit.

11.1.3.7.8 Transmit Filter IP Address Value 0—CBTIPV0[n] (0x2200 + 8*n (n=0...30); RW)

This register contains the match value for the IP Address field.

Bits	Type	Reset	Description
------	------	-------	-------------



31:0	RW	0x00	IP Address (IP). For IPv4 filters these registers contain the match value for the IP Address field. For IPv6 filters these registers contain the match value for bits 31:0 of the IP address field. IPv4 vs. IPv6 type is set by the TYPEIPV6 field in the CBTIPM[n] register. The IPATYPE field in the CBTC[n] register determines if this field is compared against the source address, the destination address, or none. The IP filters are defined in little endian notation so for IPv4 filters the MS Byte is compared against the first Byte on the wire and so on. For IPv6 filters, the MS Byte of CBTIPV3 goes first on the wire while LS Byte of CBTIPV0 goes last on the wire.
------	----	------	---

11.1.3.7.9 Transmit Filter IP Address Mask—CBTIPM[n] (0x2300 + 8*n (n=0...30); RW)

This register contains the bit mask configuration as the complement enable for the IP Address match value field.

Bits	Type	Reset	Description
7:0	RW	0x00	<p>IP Address Mask (IPM). The mask bits of the IP address in the filter. The bits are encoded to provide the following masks:</p> <p>IPM IPv4 32 bit Mask Functionality</p> <p>0x00 0x00...00 ignore bits 31:0 0x01 0x80...00 ignore bits 30:0 0x02 0xC0...00 ignore bits 29:0 0x03 0xE0...00 ignore bits 28:0 ... 0x1D 0xFF..F8 ignore bits 2:0 0x1E 0xFF..FC ignore bits 1:0 0x1F 0xFF..FE ignore bit 0 0x20 0xFF..FF match complete value other reserved undefined</p> <p>IPM IPv6 128 bit Mask Functionality</p> <p>0x00 0x00...00 ignore bits 127:0 0x01 0x80...00 ignore bits 126:0 0x02 0xC0...00 ignore bits 125:0 0x03 0xE0...00 ignore bits 124:0 ... 0x5E ignore bits 33:0 0x5F ignore bits 32:0 0x60 ignore bits 31:0 0x61 ignore bits 30:0 ... 0x7D 0xFF..F8 ignore bits 2:0 0x7E 0xFF..FC ignore bits 1:0 0x7F 0xFF..FE ignore bit 0 0x80 0xFF..FF complete address match other reserved undefined</p>
29:8	RO	0x00	Reserved.
30	RW	0	Select IPv6 Address type (TYPEIPV6).
31	RW	0	Complement (COMP). If set, this IP Address match result will be inverted. This enables a filter match on any IP address except the one that is specified in the IP Address Value field.



11.1.3.7.10 Transmit Filter Port / Type Value—CBTPTV[n] (0x2400 + 8*n (n=0...30); RW)

This register contains the TCP/UDP Port match values or the Ethernet Type match value, depending on packet type.

Bits	Type	Reset	Description
15:0	RW	0x00	Port Value #2 (PV2). The value of the second port in the filter. The filter is defined in little endian notation so MS Byte is compared against the first Byte on the wire.
31:16	RW	0x00	Port / Type Value #1 (PTV1). The value of the first port / type in the filter. The filter is defined in little endian notation so MS Byte is compared against the first Byte on the wire.

There are 2 port values identified in this field. They are applied according to the PKTTYPE and PORTYPE fields corresponding to the filter as such:

PKT TYPE	PORT TYPE	Description	PTV1 Field	PV2 Field
00 or 01	00	TCP or UDP packet, but no port check.	N/A	N/A
	01	TCP or UDP packet, source port range check.	Low Source Port	High Source Port
	10	TCP or UDP packet, destination port range check.	Low Destination Port	High Destination Port
	11	TCP or UDP packet, source & destination port check.	Source Port	Destination Port
10	00	No data match in the next protocol fields.	N/A	N/A
	01	Illegal setting. Reserved	Low Source Port	High Source Port
	10	Illegal setting. Reserved	Low Destination Port	High Destination Port
	11	Illegal setting. Reserved	Source Port	Destination Port
11	00	Non-IP packet, but no Ethernet Type check.	N/A	N/A
	01	Illegal setting.	N/A	N/A
	10	Illegal setting.	N/A	N/A
	11	Non-IP packet, Ethernet Type check.	Ethernet Type	N/A

11.1.3.7.11 Transmit Filter IP Next Header/Flags Value—CBTNHFV[n] (0x2500 + 8*n (n=0...30); RW)

This register contains the IP Next Header match value or the TCP Flags match value, depending on the packet type.

Bits	Type	Reset	Description
7:0	RW	0x00	Next Header / Flags (NHF). This field holds the compare value of the IP Next Header when the PKTTYPE is IP/Other, or it holds the compare value for the TCP Flags value in bits 5:0 if the PKTTYPE is IP/TCP. Otherwise, the value is ignored. TCP Flags are: bit 7:6 5 4 3 2 1 0 Flag Resrv URG ACK PSH RST SYN FIN



Bits	Type	Reset	Description
31:8	RO	0x00	Reserved.

11.1.3.7.12 Transmit Filter TCP Flags Mask—CBTTFM[n] (0x2600 + 8*n (n=0...30); RW)

This register contains the TCP Flags bit mask.

Bits	Type	Reset	Description
0	RW	0	TCP Flags Mask 0 / Next Header Mask (TFM0_NHM). For TCP packet this bit masks TCP Flag 0 – FIN flag. While 1b enables comparing the relevant TCP flag and 0b ignores that bit field. For IP/Other packets this field masks the next Header. While 1b enables comparing the next header and 0b means that filtering is done only on the IP address.
5:1	RW	0x00	TCP Flags Mask (TFM). The mask bits of the TCP Flags bits 5:1 value in the filter. While 1b enables comparing the relevant TCP flag and 0b ignores that bit field. The field is ignored if the PKTTYPE is other than IP/TCP.
31:6	RO	0x00	Reserved.

11.1.3.7.13 Transmit Filter IP Address Value 1—CBTIPV1[n] (0x2700 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
31:0	RW	0x00	IP Address (IP). Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBTIPM[n] register. The IP field contains the match value for the IP Address field bits 63:32. The IP filters are defined in little endian notation so the MS Byte of CBTIPV3 goes first on the wire while LS Byte of CBTIPV0 goes last on the wire.

11.1.3.7.14 Transmit Filter IP Address Value 2—CBTIPV2[n] (0x2800 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
31:0	RW	0x00	IP Address (IP). Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBTIPM[n] register. The IP field contains the match value for the IP Address field bits 95:64. The IP filters are defined in little endian notation so the MS Byte of CBTIPV3 goes first on the wire while LS Byte of CBTIPV0 goes last on the wire.

11.1.3.7.15 Transmit Filter IP Address Value 3—CBTIPV3[n] (0x2900 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
------	------	-------	-------------



31:0	RW	0x00	IP Address (IP) . Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBTIPM[n] register. The IP field contains the match value for the IP Address field bits 127:96. The IP filters are defined in little endian notation so the MS Byte of CBTIPV3 goes first on the wire while LS Byte of CBTIPV0 goes last on the wire.
------	----	------	--

11.1.3.7.16 Receive Filter Configuration Registers— CBRC[n] (0x3000 + 8*n (n=0...30); RW)

There is one instance of this register for each of the Receive Filters. The 32nd entry is used to describe/configure the receive “else” case. It is used to enable and configure the operation of the filter.

Configuration filters 0-30;

Bits	Type	Reset	Description
0	RW	0	Enable Host Filter (ENH) . Enables this filter for host filtering.
1	RW	0	Interrupt Enable (IEN) . Enables matching of this filter (or of the counter/threshold) associated with this filter to cause an interrupt to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will cause an interrupt. The IEN bit can be set only if the ENH bit is set as well for the same filter.
2	RW	0	Block (BLK) . Indicates that this filter specifies packets that should be blocked for reception from the LAN to the host. If cleared, it identifies a “pass” filter. This configuration applies to each and every packet.
3	RW	0	To ME (TOME) . If set, the packet that passes this filter is sent to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will be sent.
4	RW	0	Header To ME (HTOME) . If set, the header of the packet that matches this filter is sent to the ME.
5	RW	0	Force Header To ME (FHTOME) . When FHTOME is set the header redirection to ME overrides the HSKIP mechanism. This field may be set only when the HTOME field is set as well. Note in the case that a packet matches multiple filters with active HTOME bit but only some of them have active FHTOME bit. The header redirection to ME overrides the HSKIP mechanism only if the filter with the lowest index (0...30) has an active FHTOME bit.
7:6	RO	0x0	Reserved.
9:8	RW	00	Packet Type (PKTTYPE) . Determines what type of filter this is: 00 – IP / TCP 01 – IP / UDP 10 – IP / Other 11 – Non-IP NOTE: This field is Reserved 11b for the “else” case.
15:10	RO	0x0	Reserved.
17:16	RW	00	IP Address Type (IPATYPE) . Determines how the filter looks at the IP Address for IP/TCP, IP/UDP, and IP/Other packet types. 00 – No IP Address Compare 01 – Compare IP Address against Source Address 10 – Compare IP Address against Destination Address 11 – Reserved NOTE: This field is Reserved 11b for the “else” case.
23:18	RO	0x0	Reserved.



Bits	Type	Reset	Description
25:24	RW	00	Port Type (PORTTYPE). Determines how the filter looks at the Source and/or Destination Ports for IP/TCP or IP/UDP packet types. 00 – Ports ignored and not compared 01 – Source Port Range 10 – Destination Port Range 11 – Source and Destination Port Match NOTE: This field is Reserved 00b for the “else” case.
31:26	RO	0x0	Reserved.

**11.1.3.7.16.4 Receive Filter Configuration—CBRC[31] (0x30F8; RW)
The else configuration filter (filter number 31);**

Bits	Type	Reset	Description
0	RO	1	Reserved. This field is used as the “Enables Host filter” in all other filters. The else filter is always enabled.
1	RW	0	Interrupt Enable (IEN). Enables matching of this filter (or of the counter/threshold) associated with this filter to cause an interrupt to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will cause an interrupt. The IEN bit can be set only if the ENH bit is set as well for the same filter.
2	RW/ SN	0	Block (BLK). Indicates that this filter specifies packets that should be blocked for reception to the host. If cleared, it identifies a “pass” filter. This configuration applies to each and every packet that does not match any of the filters 0...30 with active ENH bit.
3	RW	0	To ME (TOME). If set, the packet that does not match any of the filters 0...30 with active TOME bit is sent to the ME. If a counter/threshold is associated with this filter, then only the packet that occurs on the threshold value will be sent to ME.
4	RW	0	Reserved.
5	RW	0	Reserved.
31:6	RO	0x0	Reserved.

11.1.3.7.17 Circuit Breaker Receive Filter Status—CBRFS (0x3100; R/WC)

Bits	Type	Reset	Description
31:0	RWC/ V	0x00	Receive Filter Status (RFS). Each bit indicates the corresponding receive filter (i.e. bit 0 = filter #0). If the bit is set, at least one packet has been received that hit this filter while either the “Enable Host Filter” or “To ME” bits are set in the CBRC register. This bit is cleared by writing a 1b to it. Bit 31 (else filter, index #31) is set for packets that do not match any other filter with active “Enable Host Filter” bit.

11.1.3.7.18 Receive Filter IP Address Value 0—CBRIPV0[n] (0x3200 + 8*n (n=0...30); RW)

This register contains the match value for the IP Address field.

Bits	Type	Reset	Description
------	------	-------	-------------



31:0	RW	0x00	IP Address (IP). For IPv4 filters these registers contain the match value for the IP Address field. For IPv6 filters these registers contain the match value for bits 31:0 of the IP address field. IPv4 vs. IPv6 type is set by the TYPEIPV6 field in the CBRIPM[n] register. The IPATYPE field in the CBRC[n] register determines if this field is compared against the source address, the destination address, or none. The IP filters are defined in little endian notation so for IPv4 filters the MS Byte is compared against the first Byte on the wire and so on. For IPv6 filters, the MS Byte of CBRIPV3 goes first on the wire while LS Byte of CBRIPV0 goes last on the wire.
------	----	------	---

11.1.3.7.19 Receive Filter IP Address Mask—CBRIPM[n] (0x3300 + 8*n (n=0...30); RW)

This register contains the bit mask configuration as the complement enable for the IP Address match value field.

Bits	Type	Reset	Description
7:0	RW	0x00	<p>IP Address Mask (IPM). The mask bits of the IP address in the filter. The bits are encoded to provide the following masks:</p> <p>IPM IPv4 32 bit Mask Functionality</p> <p>0x00 0x00...00 ignore bits 31:0 0x01 0x80...00 ignore bits 30:0 0x02 0xC0...00 ignore bits 29:0 0x03 0xE0...00 ignore bits 28:0</p> <p>...</p> <p>0x1D 0xFF...F8 ignore bit 2:0 0x1E 0xFF...FC ignore bit 1:0 0x1F 0xFF...FE ignore bit 0 0x20 0xFF...FF match complete value other reserved undefined</p> <p>IPM IPv6 128 bit Mask Functionality</p> <p>0x00 0x00...00 ignore bits 127:0 0x01 0x80...00 ignore bits 126:0 0x02 0xC0...00 ignore bits 125:0 0x03 0xE0...00 ignore bits 124:0</p> <p>...</p> <p>0x5E ignore bits 33:0 0x5F ignore bits 32:0 0x60 ignore bits 31:0 0x61 ignore bits 30:0</p> <p>...</p> <p>0x7D 0xFF...F8 ignore bit 2:0 0x7E 0xFF...FC ignore bit 1:0 0x7F 0xFF...FE ignore bit 0 0x80 0xFF...FF match complete value other reserved—undefined</p>
29:8	RO	0x00	Reserved.
30	RW	0	Select IPv6 Address type (TYPEIPV6).
31	RW	0	Complement (COMP). If set, this IP Address match result will be inverted. This enables a filter match on any IP address except the one that is specified in the IP Address Value field.



11.1.3.7.20 Receive Filter Port / Type Value—CBRPTV[n] (0x3400 + 8*n (n=0...30); RW)

This register contains the TCP/UDP Port match values or the Ethernet Type match value, depending on packet type.

Bits	Type	Reset	Description
15:0	RW	0x00	Port Value #2 (PV2). The value of the second port in the filter. The filter is defined in little endian notation so MS Byte is compared against the first Byte on the wire.
31:16	RW	0x00	Port / Type Value #1 (PTV1). The value of the first port in the filter. The filter is defined in little endian notation so MS Byte is compared against the first Byte on the wire.

There are 2 port values identified in this field. They are applied according to the PKTTYPE and PORTYPE fields corresponding to the filter as such:

PKT TYPE	PORT TYPE	Description	PTV1 Field	PV2 Field
00 or 01	00	TCP or UDP packet, but no port check.	N/A	N/A
	01	TCP or UDP packet, source port range check.	Low Source Port	High Source Port
	10	TCP or UDP packet, destination port check.	Low Destination Port	High Destination Port
	11	TCP or UDP packet, source & destination port check.	Source Port	Destination Port
10	00	No data match in the next protocol fields.	N/A	N/A
	01	Check next header in the same location and manner as the TCP or UDP packet, source port range check.	Low Source Port	High Source Port
	10	Check next header in the same location and manner as the TCP or UDP packet, destination port range check.	Low Destination Port	High Destination Port
	11	Check next header in the same location and manner as the TCP or UDP packet, source & destination port check.	Source Port	Destination Port
11	00	Non-IP packet, but no Ethernet Type check.	N/A	N/A
	01	Illegal setting.	N/A	N/A
	10	Illegal setting.	N/A	N/A
	11	Non-IP packet, Ethernet Type check.	Ethernet Type	N/A

11.1.3.7.21 Receive Filter IP Next Header/Flags Value—CBRNHFV[n] (0x3500 + 8*n (n=0...30); RW)

This register contains the IP Next Header match value or the TCP Flags match value, depending on packet type.

Bits	Type	Reset	Description
7:0	RW	0x00	Nest Header/Flags (NHF). This field holds the compare value of the IP Next Header when the PKTTYPE is IP/Other, or it holds the compare value for the TCP Flags value in bits 5:0 if the PKTTYPE is IP/TCP. Otherwise, the value is ignored. TCP Flags are: bit 7:6 5 4 3 2 1 0 Flag Resrv URG ACK PSH RST SYN FIN



Bits	Type	Reset	Description
31:8	RO	0x00	Reserved.

11.1.3.7.22 Receive Filter TCP Flags Mask—CBRTFM[n] (0x3600 + 8*n (n=0...30); RW)

This register contains the TCP Flags bit mask.

Bits	Type	Reset	Description
0	RW	0	TCP Flags Mask 0 / Next Header Mask (TFM0_NHM). For TCP packet this bit masks TCP Flag 0 – FIN flag. While 1b enables comparing the relevant TCP flag and 0b ignores that bit field. For IP/Other packets this field masks the next Header. While 1b enables comparing the next header and 0b means that filtering is done only on the IP address.
5:1	RW	0x00	TCP Flags Mask (TFM). The mask bits of the TCP Flags bits 5:1 value in the filter. While 1b enables comparing the relevant TCP flag and 0b ignores that bit field. The field is ignored if the PKTTYPE is other than IP/TCP.
31:6	RO	0x00	Reserved.

11.1.3.7.23 Receive Filter IP Address Value 1—CBRIPV1[n] (0x3700 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
31:0	RW	0x00	IP Address (IP). Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBRIPM[n] register. The IP field contains the match value for the IP Address field bits 63:32. The IP filters are defined in little endian notation so the MS Byte of CBRIPV3 goes first on the wire while LS Byte of CBRIPV0 goes last on the wire.

11.1.3.7.24 Receive Filter IP Address Value 2—CBRIPV2[n] (0x3800 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
31:0	RW	0x00	IP Address (IP). Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBRIPM[n] register. The IP field contains the match value for the IP Address field bits 95:64. The IP filters are defined in little endian notation so the MS Byte of CBRIPV3 goes first on the wire while LS Byte of CBRIPV0 goes last on the wire.

11.1.3.7.25 Receive Filter IP Address Value 3—CBRIPV3[n] (0x3900 + 8*n (n=0...30); RW)

Bits	Type	Reset	Description
------	------	-------	-------------



31:0	RW	0x00	IP Address (IP). Meaningful only for IPv6 address filters as selected by the TYPEIPV6 field in the CBRIPM[n] register. The IP field contains the match value for the IP Address field bits 127:96. The IP filters are defined in little endian notation so the MS Byte of CBRIPV3 goes first on the wire while LS Byte of CBRIPV0 goes last on the wire.
------	----	------	--

11.1.3.8 Header Redirection Register Descriptions

11.1.3.8.1 Tx Header Buffer Control Register—TXHCTL (0x0400; RW)

Bits	Type	Reset	Description
0	RW	0x0	Tx Header DMA Enable (DEn) . When set it enables the Tx Header DMA operation. When it is cleared then the LAN controller does not initiate new DMA requests to ME memory. Note: Tx Header Buffer Base and Last should be changed only when the enable bit is not set.
1	RW/V	0	Reset Tx Buffer (RB) . When set it resets the current Tx Header operation and resets the head and tail pointers to the value of the base address. Self cleared by HW. This bit should be set by Firmware together with the RF bit in this register.
2	RW/V	0	Reset Tx FIFO (RF) . When set, it resets the Tx Header FIFO pointers thus flushing the FIFO. It also aborts any reception to the FIFO. Self cleared by HW. This bit should be set by Firmware together with the RB bit in this register.
3	RW	0	Tx DMA Moderation (DMAMOD) . When set, the Tx Headers are posted from the Tx Header FIFO to host memory only after there are at least 64 valid bytes in the FIFO. When cleared, DMA to host memory is requested whenever there is a valid data in the Tx Header FIFO. This field is static and may be set only at buffer initialization when the Tx Enable bit in this register is not set.
4	RW	0	Tx Header Redirect Enable (HEn) . When set it enables the Tx Header posting to the internal Header FIFO. When it is cleared then new headers are not posted to the internal FIFO. Note: The Tx Header FIFO can be cleared (setting the RF bit in this register) only <i>after</i> the HEn bit is cleared as well.
31:5	RO	0x0	Reserved.

Note: When the HEn bit is not set, the HW will drop all the packets received from the Tx-Filter.

Note: When the FW wants to abort reception, it should follow the steps described in the “Tx Header Abort Flow” section.

11.1.3.8.2 Tx Header Utilization Control Register—TXHUTIL (0x0408; RW)

Bits	Type	Reset	Description
7:0	RW	0x0	Tx Header Skip (HSKIP) . The HSKIP defines the number of headers between adjacent packets that the LAN controller skips and does not post to ME. A value of 0 means that all headers are posted to ME. A value of 1 means that every other header is posted to ME and so on. This field may be changed on the fly while reception to the memory is enabled. Please see the note below for HSKIP programming.
15:8	RO	0x0	Reserved.



Bits	Type	Reset	Description
23:16	RW	0x0	Tx Interrupt Throttling (INTT) . The INTT defines the number of Tx Header DMA events between CB Header Interrupts. A value of 0x00 defines no throttling mode where each DMA causes an interrupt. A value of 1 means that every other DMA causes an interrupt and so on. This field may be changed on the fly while reception to the memory is enabled.
31:24	RO	0x0	Reserved.

Note: The HSKIP is designed to be a dynamic field. Yet, in order to minimize validation effort its dynamic nature is not covered. Therefore, the FW should avoid its programming while dynamic operation. Before any programming of the HSKIP field all filters that uses the HSKIP mechanism should be disabled by clearing their HTOME bit. Only after the HSKIP is programmed, the relevant HTOME bits can be set to their active state.

11.1.3.8.3 Transmit Header Buffer Base Address—MTXHBBA (0x0410; RW)

Bits	Type	Reset	Description
5:0	RO	0x0	Reserved. The Address is 64 byte aligned.
31:6	RW	0x0	Buffer Base Address (BBA) . The base address of the Tx Header buffer in the ME memory space. The Head and Tail registers get the base register value when the base register is being written.

11.1.3.8.4 Transmit Header Buffer Last Address—MTXHBLA (0x0418; RW)

Bits	Type	Reset	Description
5:0	RO	0x0	Reserved. The Address is 64 byte aligned.
31:6	RW/V	0x0	Buffer Last Address (BLA) . The last memory line of the transmit buffer. The register is initialized to MTBBA when the MTBBA is set. The size of the ME header ring buffer should be greater than 64Bytes (note that a 64byte ring is considered full before any headers are written).

11.1.3.8.5 Transmit Header Buffer Head Register—MTXHBHR (0x0420; RO)

Bits	Type	Reset	Description
2:0	RO/V	0x0	Overflow Counter (OVRCNT) . This field reports the number of Tx headers dropped due to FIFO full since the last DMA. The counter does not wraparound. If it reaches a value of 0x7 it stops counting. This field is cleared when this register is read by ME or reported to ME in the MSI. Note that the Buffer head pointer is always 16 bytes aligned so the OVRCNT field does not hide any data of the pointer.
3	RO/V	0x0	Overflow Event (OVRFL) . This bit is set to 1 when a header is dropped due to FIFO full while the ring buffer in ME memory is full as well. This bit is cleared when this register is read by ME or reported to ME in the MSI. Note that the Buffer head pointer is always 16 bytes aligned so the OVRFL field does not hide any data of the pointer.



Bits	Type	Reset	Description
31:4	RO/V	0x00	Buffer Head (BH). This value points to the next memory address in the Tx Header ring buffer that the HW will write the data into. The register is initialized to MTXHBBA when the MTXHBBA is set or by Reset Tx Buffer bit in the TXHCTL register. When the Buffer Head reaches as close as 4 headers (64 bytes) from the Buffer Tail the ring buffer is considered full and the HW stops writing to the ring buffer.

Reads of this register produce the same behavior as that of sending a **Tx** Header (**TxH**) MSI in the sense of clearing the OVRCNT and OVRFL fields and initializing the interrupt throttling mechanism.

11.1.3.8.6 Transmit Header Buffer Tail Register—MTXHBTR (0x0428; RW)

Bits	Type	Reset	Description
5:0	RO	0x0	Reserved. The FW always updates a 64 byte aligned address. It simplifies the HW in the case that it DMAs data to memory, 64 bytes at a time.
31:6	RW/V	0x00	Buffer Tail (BT). This value points to the next memory that the FW will read. The HW can write to the buffer until address BT-1. The register may be initialized to MTXHBBA when the MTXHBBA is set or by Reset Tx Buffer bit in the TXHCTL register.



12 Time Synch (IEEE1588 and 802.1AS)

12.1 Overview

IEEE 1588 addresses the clock synchronization requirements of measurement and control systems. The protocol supports system-wide synchronization accuracy in the sub-microsecond range with minimal network and local clock computing resources. The protocol is spatially localized and allows simple systems to be installed and operated without requiring the administrative attention of users.

The IEEE802.1AS standard specifies the protocol used to ensure that synchronization requirements are met for time sensitive applications, such as audio and video, across Bridged and Virtual Bridged Local Area Networks consisting of LAN media where the transmission delays are fixed and symmetrical. For example, IEEE 802.3 full duplex links. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration. It specifies the use of IEEE 1588 specifications where applicable.

Activation of the LAN Controller Time Sync mechanism is possible in full duplex mode only. No limitations on wire speed exist, although wire speed might affect the accuracy. Time Sync protocol is tolerant of dropping packets as well as missing timestamps.

12.1.1 Flow and HW/SW responsibilities

The operation of a PTP (precision time protocol) enabled network is divided into two stages, Initialization and time synchronization.

At the initialization stage every master enabled node starts by sending Sync packets that include the clock parameters of its clock. Upon receipt of a Sync packet a node compares the received clock parameters to its own and if the received parameters are better, then this node moves to Slave state and stops sending Sync packets. When in slave state the node continuously compares the incoming Sync packets to its currently chosen master and if the new clock parameters are better then the master selection is transferred to this better master clock. Eventually the best master clock is chosen. Every node has a defined time-out interval after which if no Sync packet is received from its chosen master clock it moves back to master state and starts sending Sync packets until a new best master clock (BMC) is chosen.

The time synchronization stage is different between master and slave nodes. If a node is at a master state it should periodically send a Sync packet which is time stamped by hardware on the TX path (as close as possible to the PHY). After the Sync packet a Follow_Up packet is sent which includes the value of the timestamp kept from the Sync packet. In addition the master should timestamp Delay_Req packets on its RX path and return, to the slave that sent it, the timestamp value using a Delay_Response packet. A node in Slave state should timestamp every incoming Sync packet and, if it came from its selected master, software will use this value for time offset calculation. In addition it should periodically send Delay_Req packets in order to calculate the path delay from its master. Every sent Delay_Req packet sent by the slave is time stamped and kept. With the timestamp value received from the master with the Delay_Response packet,

the slave can now calculate the path delay from the master to the slave. The synchronization protocol flow and the offset calculation are described in the following figure.




Figure 12-1. Sync flow and offset calculation

The HW responsibilities are:

- Identify the packets that require time stamping.
- Time stamp the packets on both RX and TX paths.
- Store the time stamp value for SW.
- Keep the system time in HW and give a time adjustment service to the SW.

The SW is responsible for:

- Best Master Clock (BMC) protocol execution which determines which clock is the highest quality clock within the network. As a result of the protocol the SW will set the node state (master or slave) and selection of the master clock if the node is in slave state.
- Generate PTP packets, consume PTP packets.
- Calculate the time offset and adjust the system time using HW mechanism.

Table 12-1. Chronological order of events for Sync and path delay

Action	Responsibility	Node Role
Generate a Sync packet with timestamp notification in descriptor.	SW	Master
Timestamp the packet and store the value in registers (T1).	HW	Master
Timestamp incoming Sync packet, store the value in register and store the sourceID and sequenceID in registers (T2).	HW	Slave
Read the timestamp from register T1 and put it in a Follow_Up packet and send.	SW	Master
Once the Follow_Up packet arrives, store T2 from registers and T1 from the Follow_up packet.	SW	Slave



Action	Responsibility	Node Role
Generate a Delay_Req packet with timestamp notification in descriptor.	SW	Slave
Timestamp the packet and store the value in registers (T3).	HW	Slave
Timestamp incoming Delay_Req packet, store the value in register and store the sourceID and sequenceID in registers (T4).	HW	Master
Read the timestamp from the register and send back to the Slave using a Delay_Response packet.	SW	Master
Once the Delay_Response packet arrives, calculate the offset using T1, T2, T3 and T4 values.	SW	Slave

12.1.1.1 TimeSync indications in RX and TX packet descriptors

Some indications need to be transferred between SW and HW regarding PTP packets.

On the transmit path the SW should set the TST bit in the ExtCMD field in the TX advanced descriptor to indicate that the transmit packet time stamp should be captured.

On the RX path, the HW has two indications to indicate to the SW in the receive descriptor:

1. An indication that this packet is a PTP packet (no matter if a timestamp was taken or not), this is also for other types of PTP packets needed for management of the protocol, this bit is set only for the L2 type of packets (the PTP packet is identified according to its Ethertype). PTP packets have the PACKETTYPE field set to 0xE to indicate that the Etype matches the filter number set by the SW to filter PTP packets. The UDP type of PTP packets don't need such an indication since the port number (319 for event and 320 all the rest PTP packets) will direct the packets toward the time sync application.
2. An indication using the TST bit in the Extended Status field of the Rx descriptor to indicate to the SW that a time stamp was taken for this packet. SW needs to access the time stamp registers to get the time stamp values.

12.1.2 HW Time sync elements

All time sync HW elements are reset to their initial values as defined in the registers section upon MAC reset.

12.1.2.1 System time structure and mode of operation

The time sync logic contains an up counter to maintain the system time value. This is a 64 bit counter that is built using the **SYSTEML** and **SYSTEMH** registers. When in Master state the **SYSTEMH** and **SYSTEML** registers should be set once by the software according to the general system, when in slave state software should update the system time on every sync event as described in [Section 12.1.2.3](#). Setting the system time is done by a direct write to the **SYSTEMH** register and fine tune setting of the **SYSTEM** register using the adjustment mechanism described in [Section 12.1.2.3](#).

Read access to the **SYSTEMH** and **SYSTEML** registers should be executed in the following manner:

1. SW read register **SYSTEML**, at this stage the HW should latch the value of **SYSTEMH**.



2. SW read register **SYSTMH**, the latched (from last read from **SYSTM**) value should be returned by HW.

The SYSTM timer value in **SYSTM** and **SYSTMH** registers is updated periodically every **TIMINCA.incperiod** (If **TIMINCA.incperiod** is one, then an increment event should occur on every clock cycle).

Upon increment event the system time value should increment its value by the value stored in **TIMINCA.incvalue**. The *incvalue* defines the granularity in which the time is represented by the **SYSTMH/L** registers. For example if the cycle time is 16ns and the *incperiod* is 1 then if the *incvalue* is 16 then the time is represented in nanoseconds if the *incvalue* is 160 then the time is represented in 0.1ns units and so on. The *incperiod* helps to avoid inaccuracy in cases where T value can not be represented as a simple integer and should be multiplied to get to an integer representation. The *incperiod* value should be as small as possible to achieve best accuracy possible.

The clock used has stable 96MHz or 25Mhz frequency that is always active at S0 state, independent of the link speed. The Host driver is indicated which clock is used for the system clock by reading TSYNCRXCTL.SYSCFI.

Note: System time registers should be implemented on a free running clock to make sure the system time is kept valid on traffic idle times (dynamic clock gating).

12.1.2.2 Time stamping mechanism

The time stamping logic is located on TX and RX paths at a location as close as possible to the PHY. This is to reduce delay uncertainties originating from implementation differences. The operation of this logic is slightly different on TX and on RX.

The TX part decides to timestamp a packet if the TX timestamp is enabled and the time stamp bit in the packet descriptor is set. On the TX side only the time is captured.

On the RX side the logic parses the traversing frame and if it is matching the message type defined in register described in [Section 12.1.2.3](#) the time, sourceId and sequenceId are latched in the timestamp registers. In addition two bits in the RX descriptor are set, one to identify that this is a PTP packet (this bit is set only for L2 packets since on the UDP packets the port number will direct the packet to the application) and the second (TS) to identify that a time stamp was taken for this packet. If this PTP packet is not Sync or Delay_Req or for some reason time stamp was not taken only the first bit is set.

For more details please refer to the timestamp registers section. The following figure defines the exact point where the time value should be captured.

On both sides the timestamp values is locked in the registers until SW access. This means that if a new PTP packet that requires time stamp has arrived before SW access is not time stamped. In some cases on the RX path a packet that was timestamped might be lost and not get to the host, to avoid lock condition the SW should keep a watch dog timer to clear locking of the time stamp register. The value of such timer should be at least higher then the expected interval between two Sync or Delay_Req packets depends the state (Master or Slave).

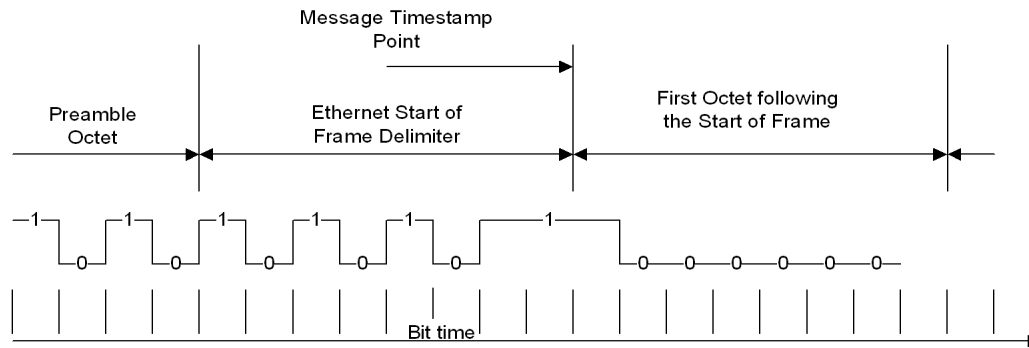



Figure 12-1. Time stamp point

12.1.2.3 Time adjustment mode of operation

A node in the Time Sync network can be in one of two states Master or Slave. When a Time Sync entity is at Master state it synchronizes other entities to its System Clock through the sending out of TimeSync, Follow-up and delay response packets. Master nodes require no time adjustments. Slave nodes adjust their system clocks by using the data arrived with the Follow_Up and Delay_Response packets and to the time stamp values of Sync and Delay_Req packets. When having all the values the SW on the slave node can calculate its offset in the following manner.

After offset calculation the system time register should be updated. This is done by writing the calculated offset to **TIMADJL** and **TIMADJH** registers. The order should be as follows:

1. Write the lower portion of the offset to **TIMADJL**.
2. Write the high portion of the offset to **TIMADJH** to the lower 31 bits and the sign to the most significant bit.

After the write cycle to **TIMADJH** the value of **TIMADJH** and **TIMADJL** should be added to the system time.

12.1.3 PTP packet structure

The time sync implementation supports both the 1588 V1 and V2 PTP frame formats. The V1 structure can come only as UDP payload over IPv4 while the V2 can come over L2 with its Ethertype or as a UDP payload over IPv4 or IPv6. The 802.1AS uses only the layer2 V2 format.

Table 12-2. V1 and V2 PTP message structure

Offset in bytes	V1 fields	V2 fields	
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	
0	versionPTP	transportSpecific ¹	messageId
1		Reserved	versionPTP
2	versionNetwork	messageLength	
3			



Offset in bytes	V1 fields	V2 fields
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0
4	Subdomain	SubdomainNumber
5		Reserved
6		flags
7		
8		correctionNs
9		
10		
11		
12		
13		correctionSubNs
14		
15		
16		reserved
17		
18		
19		
20	messageType	Reserved
21	Source communication technology	Source communication technology
22	Sourceuuid	Sourceuuid
23		
24		
25		
26		
27		
28	sourceportid	sourceportid
29		
30	sequenceId	sequenceId
31		
32	control	control
33	reserved	logMessagePeriod
34	flags	N/A
35		

Notes:

1. Should be all zero.

Note: Only the **BOLD** highlighted fields are of interest to the hardware.

**Table 12-3. PTP message over layer 2**

Ethernet (L2)	VLAN (Optional)	PTP Ethertype	PTP message
---------------	-----------------	---------------	-------------

Table 12-4. PTP message over layer 4

Ethernet (L2)	IP (L3)	UDP	PTP message
---------------	---------	-----	-------------

When a PTP packet is recognized (by Ethertype or UDP port address) on the RX side, the version should be checked. If it is V1, then the control field at offset 32 should be compared to the control field in the register described at [Section 12.1.2.3](#). Otherwise the byte at offset 0 (messageId) should be used for comparison to the messageId field.

The rest of the needed fields are at the same location and the same size for both V1 and V2 versions.

Table 12-5. Message decoding for V1 (the control field at offset 32)

Enumeration	Value
PTP_SYNC_MESSAGE	0
PTP_DELAY_REQ_MESSAGE	1
PTP_FOLLOWUP_MESSAGE	2
PTP_DELAY_RESP_MESSAGE	3
PTP_MANAGEMENT_MESSAGE	4
reserved	5-255

Table 12-6. Message decoding for V2 (messageId field at offset 0)

MessageId	Message Type	Value (hex)
PTP_SYNC_MESSAGE	Event	0
PTP_DELAY_REQ_MESSAGE	Event	1
PTP_PATH_DELAY_REQ_MESSAGE	Event	2
PTP_PATH_DELAY_RESP_MESSAGE	Event	3
Unused		4-7
PTP_FOLLOWUP_MESSAGE	General	8
PTP_DELAY_RESP_MESSAGE	General	9
PTP_PATH_DELAY_FOLLOWUP_MESSAGE	General	A
PTP_ANNOUNCE_MESSAGE	General	B
PTP_SIGNALLING_MESSAGE	General	C
PTP_MANAGEMENT_MESSAGE	General	D
Unused		E-F

If V2 mode is configured in the register described in [Section 12.1.2.3](#) then a time stamp should be taken on PTP_PATH_DELAY_REQ_MESSAGE and PTP_PATH_DELAY_RESP_MESSAGE for any value in the message field in register



described at [Section 12.1.2.3](#).

§ §



13 Mobile Design Considerations & Guidelines

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel® Ethernet Controller I217 (PHY) shown in [Figure 13-1](#). Its bus master capabilities enable the component to process high-level commands and perform multiple operations, which lowers processor use by off loading communication tasks from the processor.

The PCH, which hereinafter refers to the integrated MAC within the PCH, supports the SMBus interface for manageability while in an Sx state and PCI Express* (PCIe*) for 10/100/1000 Mbps traffic in an S0 state.

Note: The PCIe interface is not PCIe compliant. It operates at half of the PCI Express* (PCIe*) Specification v1.0 (2.5 GT/s) speed. In this chapter, the term “PCIe-based” is interchangeable with “PCIe.” There are no design layout differences between normal PCIe and the PCIe-based interface.


The PHY interfaces with the integrated MAC through two interfaces: PCIe and SMBus. In SMBus mode, the link speed is reduced to 10 Mbps. The PCIe interface incorporates two aspects: a PCIe-based SerDes (electrically) and a custom logic protocol for messaging between the integrated MAC and the PHY.

Note: Gigabit Ethernet requires an SPI Flash to host firmware and does not work without an SPI Flash on board.

The integrated MAC supports multi-speed operation (10/100/1000 Mbps). The integrated MAC also operates in full-duplex at all supported speeds or half-duplex at 10/100 Mbps as well as adhering to the IEEE 802.3x Flow Control Specification.

Note: References to the AUX power rail means the power rail is available in all power states including G3 to S5 transitions and Sx states with Wake on LAN (WoL) enabled. For example, V3P3_AUX in this chapter refers to a rail that is powered under the conditions previously mentioned.

Figure 13-1. PCH/PHY Interface Connections



**Table 13-1. SMBus Data Signals on the PCH**

Group	PHY Signal Name	PCH Signal Name	Description
Data	SMB_DATA	SMLINK0_DATA	SMBus data

Table 13-2. PCIe Data Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Data	PETp PETn	PETp PETn	PCIe transmit pair
Data	PERp PERn	PERp PERn	PCIe receive pair

Notes:

1. Can be connected to any PCIe port on the integrated MAC. The appropriate NVM descriptor soft strap (PCHSTRP9) should define which PCIe port is configured as GbE LAN.

Table 13-3. Clock and Reset Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Clock	SMB_CLK	SML0_CLK	SMBus clock
Clock	PE_CLKP PE_CLKN	CLKOUT_PCIE[7:0]_P ¹ CLKOUT_PCIE[7:0]_N ¹	PCIe* clock
Clock ²	CLK_REQ_N	PCIECLKRQ[7:0]#	PCIe clock request
Reset	PE_RST_N	PLTRST# ³	PCIe reset

Notes:

1. These signals come from the PCH and drive the PHY.
2. See Fig. 1-9 for connection information.

13.1 PHY Overview

The PHY is a single port compact component designed for 10/100/1000 Mbps operation. It enables a single port Gigabit Ethernet (GbE) implementation in a very small area, easing routing constraints from the PCH chipset to the PHY.

The PHY provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3ab, 802.3u, and 802.3i, respectively).

13.1.1 PHY Interconnects

The main interfaces for either PHY are PCIe and SMBus on the host side and Media Dependent Interface (MDI) on the link side. Transmit traffic is received from the PCH as either PCIe or SMBus packets on the host interconnect and transmitted as Ethernet packets on the MDI link. Receive traffic arrives as Ethernet packets on the MDI link and transferred to the PCH through either the PCIe or SMBus interconnects.

The PHY switches the in-band traffic automatically between PCIe and SMBus based on platform reset. The transition protocol is done through SMBus. The PCIe interface is powered down when the Ethernet link is running in an Sx state.



13.1.2 PCIe-Based Interface

A high-speed SerDes interface that uses PCIe electrical signaling at half speed while utilizing a custom logical protocol for active state operation mode.

Note: PCIe validation tools cannot be used for electrical validation of this interface; however, PCIe layout rules apply for on-board routing.

13.1.2.0.1 PCIe Interface Signals

The signals used to connect between the PCH and the PHY in this mode are:

- Serial differential pair running at 1.25 Gb/s for Rx.
- Serial differential pair running at 1.25 Gb/s for Tx.
- 100-MHz differential clock input to the PHY is generated by the PCH.
- Power and clock good indication to the PHY PE_RST_N.
- Clock control through CLK_REQ_N (refer to [Table 13-3](#)). This PHY output should be tied to the PCH input and pulled up with a 10k Ω resistor connected to 3.3V DC AUX power (present in G3 to S5).

13.1.2.0.2 PCIe Operation and Channel Behavior

The PHY only runs at 1250 Mbps speed, which is 1/2 of the Gen 1 2.5 Gb/s PCIe frequency. Each of the PCIe root ports in the PCH has the ability to run at 1250 Mbps. Configuring a PCH PCIe port that is attached to a PCIe Intel PHY device is pre-loaded from the GbE region of the NVM. The selected port adjusts the transmitter to run at 1/2 the Gen 1 PCIe speed and does not need to be PCIe compliant.

Packets transmitted and received over the PCIe interface are full Ethernet packets and not PCIe transaction/link/physical layer packets.

13.1.2.0.3 PCIe Connectivity

The PHY transmit/receive pins are output/input signals and are connected to the PCH as listed in [Table 13-1](#) through [Table 13-3](#).

13.1.2.0.4 PCIe Reference Clock

The PCIe Interface uses a 100-MHz differential reference clock, denoted PE_CLKP and PE_CLKN. This signal is typically generated on the platform and routed to the PCIe port.

The frequency tolerance for the PCIe reference clock is ± 300 ppm.

13.1.3 SMBus Interface

SMBus is a low speed (100 kHz/400 kHz/1000 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the PHY and the PCH when the platform is in a low power state (Sx). The interface is also used to enable the PCH to configure the PHY as well as passing in-band information between them.

The SMBus uses two primary signals: SMBCLK and SMBDATA, to communicate. Both of these signals float high with board-level 499 $\Omega \pm 5\%$ pull-up resistors.



The SMBus specification has defined various types of message protocols composed of individual bytes. For more details about SMBus, see the SMBus specification.

13.1.3.0.1 SMBus Connectivity

Table 13-1 through Table 13-3 list the relationship between PHY SMBus pins to the PCH LAN SMBus pins.

Note: The SMBus signals (SMB_DATA and SMB_CLK) cannot be connected to any other devices other than the integrated MAC. Connect the SMB_DATA and SMB_CLK pins to the integrated MAC SML0DATA and SML0CLK pins, respectively.

13.1.4 PCIe and SMBus Modes

In GbE operation, PCIe is used to transmit and receive data and for MDIO status and control. The PHY automatically switches the in-band traffic between PCIe and SMBus based on the platform power state. The table below lists the operating modes of PCIe and SMBus.

The Intel® Ethernet Controller I217 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

System/Intel Management Engine State	PHY	
	SMBus	PCIe
S0 and PHY Power Down	Not used	Electrical Idle (EI)
S0 and Idle or Link Disconnect	Not used	EI
S0 and Link in Low Power Idle (LPI)	Not used	EI
S0 and active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

13.1.5 Transitions between PCIe and SMBus Interfaces

13.1.5.0.1 Switching from SMBus to PCIe

Communication between the integrated MAC and the PHY is done through the SMBus each time the system is in a low power state (Sx). The integrated MAC/PHY interface is needed while the Manageability Engine (ME) is still active to transfer traffic, configuration, control and status or to enable host wake up from the PHY.

Possible states for activity over the SMBus:

1. After power on (G3 to S5).
2. On system standby (Sx).

The switching from the SMBus to PCIe is done when the PE_RST_N signal goes high.

- Any transmit/receive packet that is not completed when PE_RST_N is asserted is discarded.
- Any in-band message that was sent over the SMBus and was not acknowledged is re-transmitted over PCIe.



13.1.5.0.2 Switching from PCIe to SMBus

The communication between the integrated MAC and the PHY is done through PCIe each time the platform is in active power state (S0). Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME.

The switching from PCIe to SMBus is done when the PE_RST_N signal goes low.

- Any transmit/receive packet that is not completed when PE_RST_N goes to 0b is discarded.
- Any in-band message that was sent over PCIe and was not acknowledged is re-transmitted over SMBus.

13.2 Platform LAN Design Guidelines

These sections provide recommendations for selecting components and connecting special pins. For GbE designs, the main elements are the PCH chipset, the Intel® Ethernet Controller I217, a magnetics module and RJ-45 connector, a GbE region NVM (Non Volatile Memory) image, and a clock source.

13.2.1 General Design Considerations for PHYs

Sound engineering practices must be followed with respect to unused inputs by terminating them with pull-up or pull-down resistors, unless otherwise specified in a datasheet, design guide or reference schematic. Pull-up or pull-down resistors must not be attached to any balls identified as “No Connect.” These devices might have special test modes that could be entered unintentionally.

13.2.1.0.1 Clock Source

All designs require a 25-MHz clock source. The PHY uses the 25-MHz source to generate clocks up to 125 MHz and 1.25 GHz for both the PHY circuits and the PCIe interface. For optimum results with lowest cost, a 25-MHz parallel resonant crystal can be used along with the appropriate load capacitors at the XTAL_OUT (X2) and XTAL_IN (X1) leads. The frequency tolerance of the timing device should equal 30 ppm or better. Further detail is found in [Section 13.19](#) and [Section 13.35](#).

Note: XTAL_OUT and XTAL_IN are the signal names for the PHY.

There are three steps to crystal qualification:

1. Verify that the vendor’s published specifications in the component datasheet meet the required conditions for frequency, frequency tolerance, temperature, oscillation mode and load capacitance as specified in the respective datasheet.
2. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems.
3. Independently measure the component’s electrical parameters in real systems. Measure frequency at a test output to avoid test probe loading effects at the PHY. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications. For crystals, it is also important to examine startup behavior while varying system voltage and temperature.



13.2.1.0.2 Magnetics Module

The magnetics module has a critical effect on overall IEEE and emissions conformance. The device should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Carefully qualifying new magnetics modules prevents problems that might arise because of interactions with other components or the printed circuit board itself.

The steps involved in magnetics module qualification are similar to those for crystal qualification:

1. Verify that the vendor’s published specifications in the component datasheet meet or exceed the required IEEE specifications.
2. Independently measure the component’s electrical parameters on the test bench, checking samples from multiple lots. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications.
3. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems. Vary temperature and voltage while performing system level tests.

Magnetics modules for 1000BASE-T Ethernet as used by the PHY only are similar to those designed solely for 10/100 Mbps, except that there are four differential signal pairs instead of two. Refer to the Intel® Ethernet Controller I217 datasheet for specific electrical requirements that the magnetics need to meet.

The following magnetics modules are not recommended; however, they have been used successfully in previous designs:

Table 13-4. Magnetic Modules and Manufacturers

Manufacturer	Part Number	Note
SpeedTech*	P25BPB4MFRT9	USB stack, 8core
SpeedTech*	P25BFB4-RDW9	USB stack, 12core
Foxconn*	JFM38U1A-21C7-4F	USB stack, 8core
Foxconn*	JFM38U1A-7110-4F	USB stack, 8core
Tyco*	1840023-1	USB stack, 8core

13.2.1.0.3 Criteria for Integrated Magnetics Electrical Qualification

The following table gives the criteria used to qualify integrated magnetics.



Table 13-5. Integrated Magnetics Recommended Qualification Criteria

Open Circuit Inductance (OCL)	w/8 mA DC bias; at 25C	400uH Min
	w/8 mA DC bias; at 0C to 70C	350uH Min
Insertion Loss	100 kHz through 999 kHz	1dB Max
	1.0 MHz through 60.0 MHz	0.6dB Max
	60.1 MHz through 80.0 MHz	0.8dB Max
	80.1 MHz through 100.0 MHz	1.0dB Max
	100.1 MHz through 125.0 MHz	2.4dB Max
Return Loss	1.0 MHz through 40.0 MHz	18.0 dB Min
	40.1 MHz through 100.0 MHz When reference impedance is 85 Ohms, 100 Ohms, and 115 Ohms. Note that R.L. values may vary with MDI trace lengths. The LAN magnetics may need to be measured in the platform where it will be used.	12–20 * LOG (Freq in MHz / 80) dB Min
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	-50.3+(8.8*(Freq in MHz / 30)) dB Max
	30.0 MHz through 250.0 MHz	-(26 -(16.8*(LOG(Freq in MHz / 250 MHz)))) dB Max
	250.1 MHz through 375.0 MHz	-26.0 dB Max
Crosstalk Isolation Integrated Modules (Proposed)	1.0 MHz through 10 MHz	-50.8+(8.8*(Freq in MHz / 10)) dB Max
	10.0 MHz through 100.0 MHz	-(26 -(16.8*(LOG(Freq in MHz / 100 MHz)))) dB Max
	100 MHz through 375.0 MHz	-26.0 dB Max
Diff to CMR	1 MHz through 29.9 MHz	-40.2+(5.3*((Freq in MHz / 30)) dB Max
	30.0 MHz through 500 MHz	-(22-(14*(LOG((Freq in MHz / 250)))) dB Max
CM to CMR	1 MHz through 270 MHz	-57+(38*((Freq in MHz / 270)) dB Max
	270.1 MHz through 300 MHz	-17-2*((300-(Freq in MHz) / 30) dB Max
	300.1 MHz through 500 MHz	-17 dB Max
Hi-Voltage Isolation	1500 Vrms at 50 or 60 Hz for 60 sec. or:2250 Vdc for 60 seconds	Minimum

13.2.2 NVM for PHY Implementations

The LAN only supports an SPI Flash, which is connected to the PCH. Several words of the NVM are accessed automatically by the device after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the NVM space is available to software for storing the MAC address, serial numbers, and additional information. More details may be obtained from the Datasheet.

Intel has an MS-DOS* software utility called EUpdate that is used to program the SPI Flash images in development or production line environments. A copy of this program can be obtained through your Intel representative.

13.2.3 LAN Switch

In order to achieve IEEE conformance for applications that must operate both docked and undocked, a LAN switch is recommended. Note that Intel does not recommend specific switches, but those in the following list have been used successfully in previous designs.

Manufacturer	Part Number
Pericom*	PI3L500-AZ
Texas Instrument*	S3L500AE

13.2.4 LED

The PHY has three LED outputs that can be configured via the NVM. The hardware configuration is shown in Figure 13-2.


Refer to the Intel® Ethernet Controller I217 Reference Schematic for default LED color based on reference design.

Refer to the Datasheet for details regarding the programming of the LED’s and the various modes. The default values for the PHY (based on the LED NVM setting—word 0x18 of the LAN region) are listed in the table below:

Table 13-6. LED Default Values

LED	Mode	Color	Blink	Polarity
LED0	Link Up/Activity	Green	200 ms on/200 ms off	Active low
LED1	Link 1000	Yellow	No	Active low
LED2	Link 100	Green	No	Active low

Figure 13-2. LED Hardware Configuration



13.2.4.0.1 RBIAS

RBIAS requires external resistor connection to bias the internal analog section of the device. The input is sensitive to the resistor value. Resistors of 1% tolerance must be used. Connect RBIAS through a 3.01kΩ 1% pull-down resistor to ground and then place it no more than one half inch (0.5”) away from the PHY.



13.2.4.0.2 LAN Disable

The PHY enters a power-down state when the LAN_DISABLE_N pin is asserted low. Exiting this mode requires setting the LAN_DISABLE_N pin to a logic one. Connect LAN_DISABLE_N to LAN_PHY_PWR_CTRL /GPIO12 on the PCH.

13.2.5 Exposed Pad* (e-PAD) Design and SMT Assembly Guide

13.2.5.0.1 Overview

This section provides general information about e-PAD and SMT assemblies. Chip packages have exposed die pads on the bottom of each package to provide electrical interconnections with the printed circuit board. These ePADs also provide excellent thermal performance through efficient heat paths to the PCB.

Packages with e-PADs are very popular due to their low cost. Note that this section only provides basic information and references in regards to the e-PAD. It is recommended that each customer consult their fab and assembly house to obtain more details on how to implement the e-PAD package design. Each fab and assembly house might need to tune the land pattern/stencil and create a solution that best suits their methodology and process.

13.2.5.0.2 PCB Design Requirements

In order to maximize both heat removal and electrical performance, a land pattern must be incorporated on the PCB within the footprint of the package corresponding to the exposed metal pad or exposed heat slug of the package as shown in the following figures. Refer to the specific product datasheet for actual dimensions.

Note: Due to the package size, a via-in-pad configuration must be used [Figure 13-3](#) and [Figure 13-4](#) are general guidelines see [Figure 13-5](#) for -specific via-in-pad thermal pattern recommendations.

Figure 13-3. Typical ePAD Land Pattern





Figure 13-4. Typical Thermal Pad and Via Recommendations




Note: Encroached and uncapped via configurations have voids less than the maximum allowable void percentage. Uncapped via provides a path for trapped air to escape during the reflow soldering process.

Note: Secondary side solder bumps might be seen in an uncapped via design. This needs to be considered when placing components on the opposite side of the PHY.



Figure 13-5. Recommended Thermal Via Patterns for



13.2.5.0.3 Board Mounting Guidelines

The following are general recommendations for mounting a QFN-48 device on the PCB. This should serve as the starting point in assembly process development and it is recommended that the process should be developed based on past experience in mounting standard, non-thermally/electrically enhanced packages.

13.2.5.0.4 Stencil Design


For maximum thermal/electrical performance, it is required that the exposed pad/slug on the package be soldered to the land pattern on the PCB. This can be achieved by applying solder paste on both the pattern for lead attachment as well as on the land pattern for the exposed pad (e-PAD). While for standard (non-thermally/ -electrically enhanced) lead-frame based packages the stencil thickness depends on the lead pitch and package co-planarity, the package standoff must also be considered for the thermally/electrically enhanced packages to determine the stencil thickness. In this case, a stencil foil thickness in the range of 5-6 mils (or 0.127-0.152 mm) is recommended; likely or practically, a choice of either 5 mils or 6 mils. Tolerance wise, it should not be worse than ± 0.5 mil.

Note: Industry specialists typically use ± 0.1 -mil tolerance on stencil for its feasible precision.

The aperture openings should be the same as the solder mask openings on the land pattern. Since a large stencil opening may result in poor release, the aperture opening should be subdivided into an array of smaller openings, similar to the thermal land pattern shown in the figure below.

Note: Refer to the specific product datasheet for actual dimensions.

Figure 13-6. Stencil Design Recommendation for QFN Package Type Products



Important General Guidelines:

- The Stencil Aperture Openings of the e-PAD must not go outside of the exposed landing area (solder mask opening).
- The Stencil Aperture Openings of the e-PAD should be about 80% of the exposed landing area (solder mask opening).

The Intel® Ethernet Connection I217 e-PAD has D2=E2=3mm, therefore the Stencil Design can only have four aperture openings for the e-PAD. This can be achieved by setting R=S=1.35mm and T=0.3. Using this arrangement, the Apertures/e-PAD area is 81% of the exposed landing area (solder mask opening).


This information is only intended as general guidance. Please consult with the manufacturer to confirm the final design meets requirements.



13.2.5.0.5 Assembly Process Flow

The following figure below shows the typical process flow for mounting packages to the PCB.


Figure 13-7. Assembly Flow



13.2.5.0.6 Reflow Guidelines

The typical reflow profile consists of four sections. In the preheat section, the PCB assembly should be preheated at the rate of 1 to 2 °C/sec to start the solvent evaporation and to avoid thermal shock. The assembly should then be thermally soaked for 60 to 120 seconds to remove solder paste volatiles and for activation of flux. The reflow section of the profile, the time above liquidus should be between 45 to 60 seconds with a peak temperature in the range of 245 to 250 °C, and the duration at the peak should not exceed 30 seconds. Finally, the assembly should undergo cool down in the fourth section of the profile. A typical profile band is provided in the figure below, in which 220 °C is referred to as an approximation of the liquidus point. The actual profile parameters depend upon the solder paste used and specific recommendations from the solder paste manufacturers should be followed.

Figure 13-8. Typical Profile Band



1. Preheat: 125°C–220°C, 150–210s at 0.4k/s to 1.0k/s
2. Time at T>220°C: 60–90s
3. Peak Temperature: 245–250°C
4. Peak time: 10–30s
5. Cooling rate: <=6k/s
6. Time from 25 °C to Peak: 240–360s
7. Intel recommends a maximum solder void of 50% after reflow.

Note: Contact your Intel Field Service Representative for any designs unable to meet the recommended guidance for E-pad implementation.

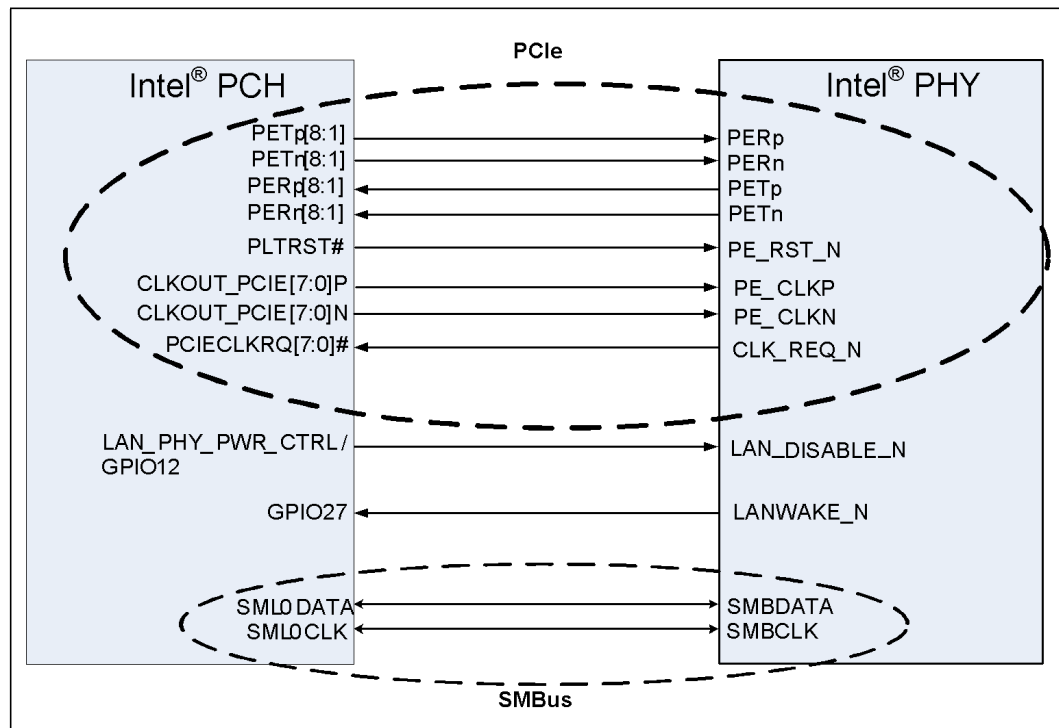
13.3 PCH—SMBus/PCIe LOM Design Guidelines

This section contains guidelines on how to implement a PCH/PHY single solution on a system motherboard. It should not be treated as a specification, and the system designer must ensure through simulations or other techniques that the system meets the specified timings. The following are guidelines for both PCH SMBus and PCIe interfaces. Note that PCIe is only applicable to the PHY.

The SMBus/PCIe Interface between the PCH and PHY is shown at high level in Figure 15-9. For complete design details always refer to the *Intel® I217 GbE Reference Schematic*.

Refer to [Section 13.6](#) for PCI Express Routing Guidelines.

Figure 13-9. Single Solution Interconnect

**Notes:**

1. For latest PHY schematic connection recommendations, refer to the *Intel® I217 GbE PHY Reference Schematic*, and the *Intel® I217 GbE PHY Schematic and Layout Checklist*. For access to documentation, contact your Intel Field Service Representative.
2. Any free PCIe lanes [8:1] on the PCH can be used to connect the PCIe interface on the Intel® Ethernet Controller I217.
3. Any free CLKOUT_PCIE[7:0] and PCIECLKRQ[7:0]# from PCH can be used to connect to PE_CLK[P:N] and CLK_REQ_N on the Intel® Ethernet Controller I217 respectively.
4. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.

13.4 SMBus Design Considerations

No single SMBus design solution works for all platforms. Designers must consider the total bus capacitance and device capabilities when designing SMBus segments. Routing SMBus to the PCI slots makes the design process even more challenging since they add so much capacitance to the bus. This extra capacitance has a large affect on the bus time constant which in turn affects the bus rise and fall times.

Primary considerations in the design process are:

- Amount of V_{CC_SUS3_3} current available, that is, minimizing load of V_{CC_SUS3_3}.
- The pull-up resistor size for the SMBus data and clock signals is dependent on the bus load (this includes all device leakage currents). Generally the SMBus device that can sink the least amount of current is the limiting agent on how small the resistor can be. The pull-up resistor cannot be made so large that the bus time constant (Resistance X Capacitance) does not meet the SMBus rise and time specification.
- The maximum bus capacitance that a physical segment can reach is 400 pF.
- SMBus devices that can operate in STR must be powered by the V_{CC_SUS3_3} supply.



- It is recommended that I²C (Inter-Integrated Circuit) devices be powered by the V_{CC_core} supply. During an SMBus transaction in which the device is sending information to the integrated MAC, the device may not release the SMBus if the integrated MAC receives an asynchronous reset. V_{CC_core} is used to enable the BIOS to reset the device if necessary. SMBus 2.0- compliant devices have a timeout capability that makes them in-susceptible to this I²C issue, enabling flexibility in choosing a voltage supply.
- No other devices (except the integrated MAC and pull-up resistors) should be connected to the SMBus that connects to the PHY.
- **For system LAN on motherboard (LOM) designs:** The traces should be less than 70 inches for stripline and less than 100 inches for Microstrip. These numbers depend on the stackup, dielectric layer thickness, and trace width. The total capacitance on the trace and input buffers should be under 400 pF.
- **For system LAN on daughterboard designs:** Being conservative, the traces should be less than 7 inches for stripline designs and less than 10 inches for Microstrip designs. The lengths depend on the stackup, dielectric layer thickness, and trace width. Longer traces can be used as long as the total capacitance on the trace and input buffers is under 30 pF.

Note: Refer to [Section 13.1.3](#) for additional SMBus design considerations.

13.5 General Layout Guidelines

PHY interface signals must be carefully routed on the motherboard to meet the timing and signal quality requirements of their respective interface specifications. The following are some general guidelines that should be followed in designing a LAN solution. It is recommended that the board designer simulate the board routing to verify that the specifications are met for flight times and skews due to trace mismatch and crosstalk.

13.6 Layout Considerations

Critical signal traces should be kept as short as possible to decrease the likelihood of effects by high frequency noise of other signals, including noise carried on power and ground planes. This can also reduce capacitive loading.

Since the transmission line medium extends onto the printed circuit board, layout and routing of differential signal pairs must be done carefully.

Designing for GbE (1000BASE-T) operation is very similar to designing for 10/100 Mbps. For the PHY, system level tests should be performed at all three speeds.

13.7 Guidelines for Component Placement

Component placement can affect signal quality, emissions, and component operating temperature. Careful component placement can:

Decrease potential problems directly related to electromagnetic interference (EMI), which could cause failure to meet applicable government test specifications. In this case, place the PHY more than one inch from the edge of the board.



Simplify the task of routing traces. To some extent, component orientation affects the complexity of trace routing. The overall objective is to minimize turns and crossovers between traces.


13.7.1 PHY Placement Recommendations

Minimizing the amount of space needed for the PHY is important because other interfaces compete for physical space on a motherboard near the connector. The PHY circuits need to be as close as possible to the connector.

The following figure illustrates some basic placement distance guidelines. To simplify the diagram, it shows only two differential pairs, but the layout can be generalized for a GbE system with four analog pairs. The ideal placement for the PHY (LAN silicon) is approximately one inch behind the magnetics module.

While it is generally a good idea to minimize lengths and distances, this figure also illustrates the need to keep the PHY away from the edge of the board and the magnetics module for best EMI performance.

Figure 13-10. LAN Device Placement: At Least One Inch from Chassis Openings or Unshielded Connectors—Mobile



Note: * this distance is variable and follows the general guidelines.

The PHY, referred to as "LAN Device" in the above figure, must be at least one inch from any chassis openings. To help reduce EMI, the following recommendations should be followed:

- Minimize the length of the MDI interface. See detail in table below: MDI Routing Summary
- Place the MDI traces no closer than 0.5 inch (1.3 cm) from the board edge.
- The Intel® Ethernet Controller I217 must be placed greater than 1" away from any hole to the outside of the chassis larger than 0.125 inches (125 mils) The larger the hole the higher the probability the EMI and ESD immunity will be negatively affected.
- The Intel® Ethernet Controller I217 should be placed greater than 250mils from the board edge.
- If the connector or integrated magnetics module is not shielded, the Intel® Ethernet Controller I217 should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the Intel® Ethernet Controller I217 closer than one inch to unshielded magnetics or connectors will increase the probability of failed EMI and common mode noise. If the LAN switch is too far away it will negatively affect IEEE return loss performance.
- The RBIAS trace length must be less than 1"
- Place the crystal less than one inch (2.54 cm) from the PHY.

Figure 13-11. PLC Placement: At Least One Inch from I/O Backplane





Figure 13-12. Effect of LAN Device Placed Too Close To a Rj-45 Connector or Chassis Opening





13.8 MDI Differential-Pair Trace Routing for LAN Design

Trace routing considerations are important to minimize the effects of crosstalk and propagation delays on sections of the board where high-speed signals exist. Signal traces should be kept as short as possible to decrease interference from other signals, including those propagated through power and ground planes.

13.9 Signal Trace Geometry

One of the key factors in controlling trace EMI radiation are the trace length and the ratio of trace-width to trace-height above the reference plane. To minimize trace inductance, high-speed signals and signal layers that are close to a reference or power plane should be as short and wide as practical. Ideally, the trace-width to trace-height above the ground plane ratio is between 1:1 and 3:1. To maintain trace impedance, the width of the trace should be modified when changing from one board layer to another if the two layers are not equidistant from the neighboring planes.

Each pair of signals should have a differential impedance of $100\ \Omega \pm 15\%$.

A set of trace length calculation tools are available from Intel (via the Intel Business Link (IBL)) to aid with MDI topology design. For access to documentation, contact your Intel Field Service Representative.

When performing a board layout, the automatic router feature of the CAD tool must not route the differential pairs without intervention. In most cases, the differential pairs will require manual routing.

Note: Measuring trace impedance for layout designs targeting $100\ \Omega$ often results in lower actual impedance due to over-etching. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of $105\ \Omega$ to $110\ \Omega$ should compensate for over-etching.

It is necessary to compensate for trace-to-trace edge coupling, which can lower the differential impedance by up to $10\ \Omega$, when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 13-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Signal group	MDI_PLUS[0:3] MDI_MINUS[0:3]		
Microstrip/stripline uncoupled single-ended impedance specification	$50\ \Omega \pm 10\%$		
Microstrip/stripline uncoupled differential impedance specification	$100\ \Omega \pm 15\%$		2,3
Microstrip nominal trace width	Design dependent	Design dependent	4
Microstrip nominal trace space	Design dependent	Design dependent	3,5
Microstrip/stripline trace length	8 in (203 mm) maximum		6,7
Microstrip pair-to-pair space (edge-to-edge)	≥ 7 times the thickness of the thinnest adjacent dielectric layer		Figure 13-13



Table 13-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Stripline pair-to-pair space (edge-to-edge)	≥ 6 times the thickness of the thinnest adjacent dielectric layer		
Microstrip bus-to-bus spacing	≥ 7 times the thickness of the thinnest adjacent dielectric layer		
Stripline bus-to-bus spacing	≥ 6 times the thickness of the thinnest adjacent dielectric layer		

Notes:

1. Pair-to-pair spacing ≥ 3 times the dielectric thickness for a maximum distance of 500 mils from the pin.
2. Board designers should ideally target 100Ω ±15%. If it's not feasible (due to board stack-up) it is recommended that board designers use a 95Ω ±10% target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95Ω. The ±10% tolerance is provided to allow for board manufacturing process variations and not lower target impedances. The minimum value of impedance cannot be lower than 85Ω.
3. Simulation shows 80Ω differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90Ω.
4. Stripline is NOT recommended due to thinner more resistive signal layers.
5. Use a minimum of 21 mil (0.533 mm) pair-to-pair spacing for board designs that use the CRB design stack-up. Using dielectrics that are thicker than the CRB stack-up might require larger pair-to-pair spacing.
6. For applications that require a longer MDI trace length of more than 8 inches (20.32 mm), it is recommended that thicker dielectric or lower Er materials be used. This permits higher differential trace impedance and wider, lower loss traces. Refer to Table 13-8 for examples of microstrip trace geometries for common circuit board materials.
7. If a LAN switch is not used, then the maximum trace length is 4 inches (102 mm). Mobile designs without LAN switch can range up to ~8 inches. Refer to Table 13-8 for trace length information.

Table 13-8. Maximum Trace Lengths Based on Trace Geometry and Board Stack-Up

Dielectric Thickness (mils)	Dielectric Constant (DK) at 1 MHz	Width / Space / Width (mils)	Pair-to-Pair Space (mils)	Nominal Impedance (Ohms)	Impedance Tolerance (±%)	Maximum Trace Length (inches) ¹
2.7	4.05	4/10/4	19	95 ²	17 ²	3.5
2.7	4.05	4/10/4	19	95 ²	15 ²	4
2.7	4.05	4/10/4	19	95	10	5
3.3	4.1	4.2/9/4.2	23	100 ²	17 ²	4
3.3	4.1	4.2/9/4.2	23	100	15	4.6
3.3	4.1	4.2/9/4.2	23	100	10	6
4	4.2	5/9/5	28	100 ²	17 ²	4.5
4	4.2	5/9/5	28	100	15	5.3
4	4.2	5/9/5	28	100	10	7

Notes:

1. Longer MDI trace lengths may be achievable, but may make it more difficult to achieve IEEE conformance. Simulations have shown deviations are possible if traces are kept short. Longer traces are possible; use cost considerations and stack-up tolerance for differential pairs to determine length requirements.
2. Deviations from 100Ω nominal and/or tolerances greater than 15% decrease the maximum length for IEEE conformance.

Note:

Use the MDI Differential Trace Calculator to determine the maximum MDI trace length for your trace geometry and board stack-up. Contact your Intel Field Service Representative for access.


The following factors can limit the maximum MDI differential trace lengths for IEEE conformance:



- Dielectric thickness
- Dielectric constant
- Nominal differential trace impedance
- Trace impedance tolerance
- Copper trace losses
- Additional devices, such as switches, in the MDI path may impact IEEE conformance.

Board geometry should also be factored in when setting trace length.

Figure 13-13.MDI Trace Geometry



13.10 Trace Length and Symmetry

The differential traces should be equal in total length to within 10 mils (0.254 mm) per segment within each pair and as symmetrical as possible. Asymmetrical and unequal length traces in the differential pairs contribute to common mode noise. If a choice has to be made between matching lengths and fixing symmetry, more emphasis should be placed on fixing symmetry. Common mode noise can degrade the receive circuit's performance and contribute to radiated emissions.


The intra-pair length matching on the pairs must be within 10 mils on a segment by segment basis. An MDI segment is defined as any trace within the same layer. For example, transitioning from one layer to another through a via is considered as two separate MDI segments.

The end-to-end total trace lengths within each differential pair must match as shown in the figure titled MDI Trace Geometry. The end to end trace length is defined as the total MDI length from one component to another regardless of layer transitions.

The pair to pair length matching is not as critical as the intra-pair length matching but it should be within 2 inches.

When using Microstrip, the MDI traces should be at least 7x the thinnest adjacent dielectric away from the edge of an adjacent reference plane. When using stripline, the MDI traces should be at least 6x the thinnest adjacent dielectric away from the edge of an adjacent reference plane.

Figure 13-14. MDI Differential Trace Geometry



Note: Similar topology applies to MDI routing from the Intel® Ethernet Controller I217 to the dock RJ45 connector.

13.11 Impedance Discontinuities

Impedance discontinuities cause unwanted signal reflections. Vias (signal through holes) and other transmission line irregularities should be minimized. If vias must be used, a reasonable budget is four or less per differential trace. Unused pads and stub traces should also be avoided.

13.12 Reducing Circuit Inductance

Traces should be routed over a continuous reference plane with no interruptions. If there are vacant areas on a reference or power plane, the signal conductors should not cross the vacant area. This causes impedance mismatches and associated radiated noise levels.

13.13 Signal Isolation

To maintain best signal integrity, keep digital signals far away from the analog traces. Also, keep the MDI traces away from the edge of an adjacent reference plane by a distance that is at least 7x the thickness of the thinnest adjacent dielectric layer (7x when using Microstrip; 6x when using stripline). If digital signals on other board layers cannot be separated by a ground plane, they should be routed perpendicular to the differential pairs. If there is another LAN controller on the board, the differential pairs from that circuit must be kept away.



Other rules to follow for signal isolation include:

- Separate and group signals by function on separate layers if possible. If possible, maintain at least a gap of 30 mils between all differential pairs (Ethernet) and other nets, but group associated differential pairs together.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.
- Isolate I/O signals from high-speed signals to minimize crosstalk, which can increase EMI emission and susceptibility to EMI from other signals.
- Avoid routing high-speed LAN traces near other high-frequency signals associated with a video controller, cache controller, processor, switching power supplies, or other similar devices.

13.14 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and power inputs bypassed to signal return. This will significantly reduce EMI radiation.

The following guidelines help reduce circuit inductance in both backplanes and motherboards:

- Route traces over a continuous plane with no interruptions. Do not route over a split power or ground plane. If there are vacant areas on a ground or power plane, avoid routing signals over the vacant area. This will increase inductance and EMI radiation levels.
- All ground vias should be connected to every ground plane; and every power via, to all power planes at equal potential. This helps reduce circuit inductance.
- Physically locate grounds between a signal path and its return. This will minimize the loop area.
- Split the ground plane beneath a magnetics module. The RJ-45 connector side of the transformer module should have chassis ground beneath it.

Caution: DO NOT do this, if the RJ-45 connector has integrated USB.

Note: All impedance-controlled signals should be routed in reference to a solid plane. If there are plane splits on a reference layer and the signal traces cross those splits then stitching capacitors should be used within 40 mils of where the crossing occurs. See Figure 1-13.

If signals transition from one reference layer to another reference layer then stitching capacitors or connecting vias should be used based on the following:

If the transition is from power-referenced layer to a ground-referenced layer or from one voltage-power referenced layer to a different voltage-power referenced layer, then stitching capacitors should be used within 40 mils of the transition.

If the transition is from one ground-referenced layer to another ground-referenced layer or is from a power-referenced layer to the same net power-referenced layer, then connecting vias should be used within 40 mils of the transition.

Figure 13-15. Trace Transitioning Layers and crossing Plane Splits

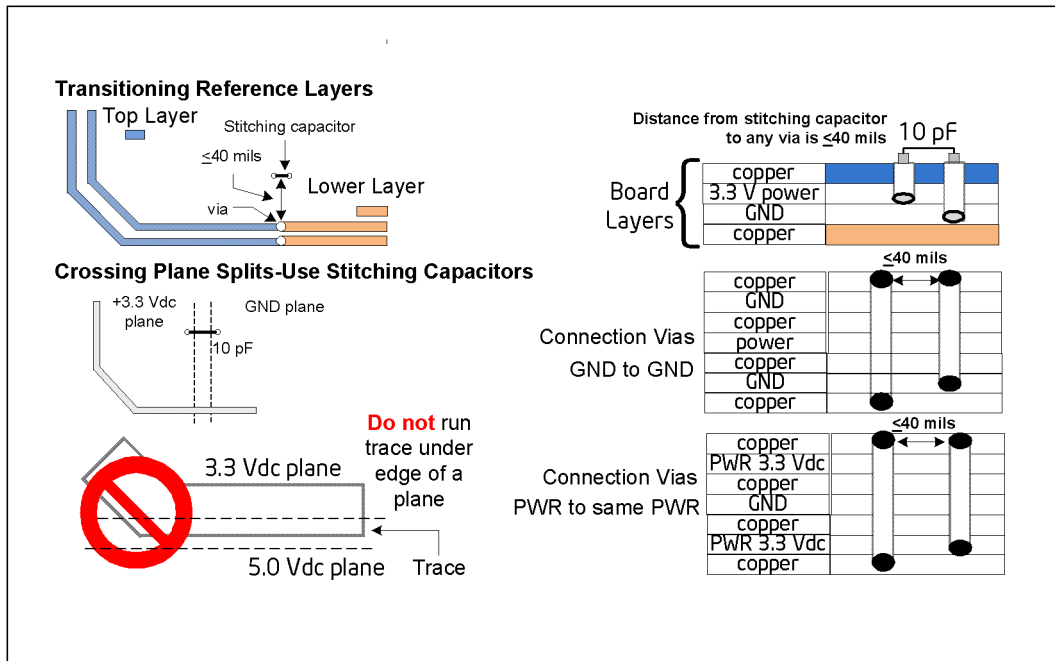


Figure 13-16. Via Connecting GND to GND





Figure 13-17. Stitching Capacitor between Vias Connecting GND to GND



13.15 Traces for Decoupling Capacitors

Traces between decoupling and I/O filter capacitors should be as short and wide as practical. Long and thin traces are more inductive and reduce the intended effect of decoupling capacitors. Also, for similar reasons, traces to I/O signals and signal terminations should be as short as possible. Vias to the decoupling capacitors should be sufficiently large in diameter to decrease series inductance. Refer to the Power Delivery section for the PHY in regards to actual placement requirements of the capacitors.

13.16 Ground Planes under a Magnetics Module

The magnetics module chassis or output ground (secondary side of transformer) should be separated from the digital or input ground (primary side) by a physical separation of 100 mils minimum. Splitting the ground planes beneath the transformer minimizes noise coupling between the primary and secondary sides of the transformer and between the adjacent coils in the magnetics. This arrangement also improves the common mode choke functionality of magnetics module.

Caution: DO NOT do this if the RJ-45 connector has integrated USB.

For more information on split plane layout using a discrete magnetic module, please see I217 Reference Schematic and I217 Schematic/Layout Design Checklist.

Figure 13-18 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.

Figure 13-18. Ideal Ground Split Implementation

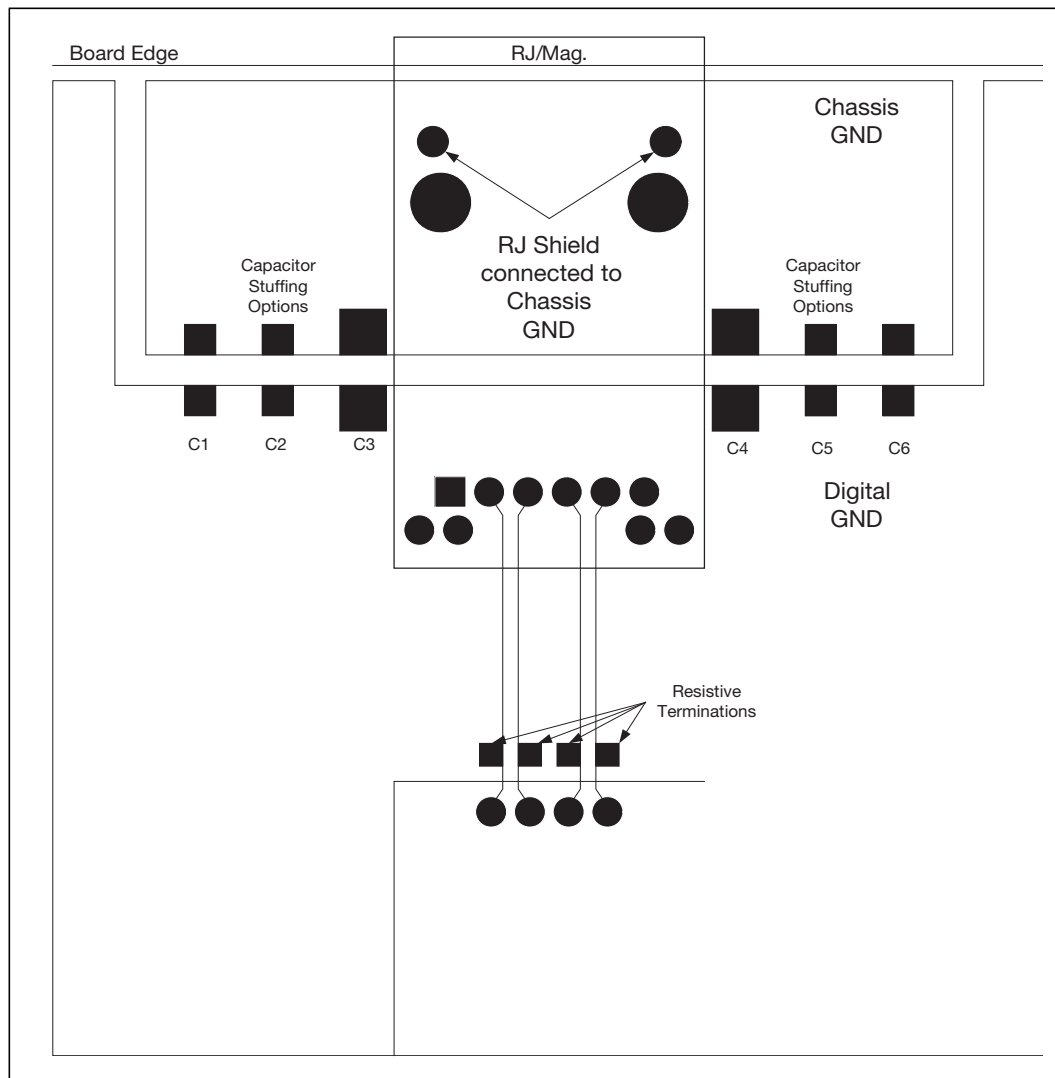


Table 13-9. Capacitor Stuffing Option Recommended Values

Capacitors	Value
C3, C4	4.7 μ F or 10 μ F
C1, C2, C5, C6	470 pF to 0.1 μ F


The placement of C1 through C6 may also differ for each board design (in other words, not all of the capacitors may need to be populated). Also, the capacitors may not be needed on both sides of the magnetics module.

Note: If using an integrated magnetics module without USB, provide a separate chassis ground “island” to ground around the RJ-45 connector. The split in the ground plane should be at least 20 mils wide.



Some integrated magnetics modules/RJ-45 connectors have recently incorporated the USB into the device. For this type of magnetics module, a chassis ground moat may not be feasible due to the digital ground required for the USB pins and their placement relative to the magnetics pins. Thus, a continuous digital ground without any moats or splits must be used. Figure 13-19 provides an example of this.

Figure 13-19. Ground Layout with USB



13.17 Light Emitting Diodes

The device has three high-current outputs to directly drive LEDs for link, activity and speed indication. Since LEDs are likely to be integral to a magnetics module, take care to route the LED traces away from potential sources of EMI noise. In some cases, it may be desirable to attach filter capacitors.

LAN LED traces should be placed at least 6x (side by side separation) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

LAN LED traces should be placed at least 7x (broadside coupling) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

13.18 Considerations for Layout

The PHY MDI routing using microstrip requires a differential impedance of $100 \Omega \pm 15\%$. A 35 mils (0.889 mm) separation is required between pairs. The 35-mil separation can be reduced for 24 mils (0.61 mm) in breakout routing. All MDI traces must be referenced to ground.



13.19 Frequency Control Device Design Considerations

This section provides information regarding frequency control devices, including crystals and oscillators, for use with all Intel Ethernet controllers. Several suitable frequency control devices are available; none of which present any unusual challenges in selection. The concepts documented within this section are applicable to other data communication circuits, including the PHY.

The PHY contains amplifiers that form the basis for feedback oscillators when they are used with the specific external components. These oscillator circuits, which are both economical and reliable, are described in more detail in [Section 13.23](#).

The chosen frequency control device vendor should be consulted early in the design cycle. Crystal and oscillator manufacturers familiar with networking equipment clock requirements may provide assistance in selecting an optimum, low-cost solution.

Several types of third-party frequency reference components are currently available. Descriptions of each type follow in subsequent sections. They are also listed in order of preference.

13.20 Crystals and Oscillators

Clock sources should not be placed near I/O ports or board edges. Radiation from these devices may be coupled onto the I/O ports or out of the system chassis. Crystals should also be kept away from the Ethernet magnetics module to prevent interference.

Crystal and load capacitors should be placed on the printed circuit boards as close to the PHY as possible, which is within 1.0 inch. Traces from XTAL_IN (X1) and XTAL_OUT (X2) should be routed as symmetrically as possible. Do not route X1 and X2 as a differential trace. Doing so increases jitter and degrades LAN performance.

- The crystal trace lengths should be less than 1 inch.
- The crystal load capacitors should be placed less than 1" from the crystal.
- The clock lines must be at least 5 times the height of the thinnest adjacent dielectric layer away from other from other digital traces (especially reset signals), I/O ports, board edge, transformers and differential pairs.
- The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition.
- The clock lines should not cross or run in parallel (within 3x the dielectric thickness of the closest dielectric layer) with any trace (100Mhz signal or higher) on an adjacent layer.

13.21 Quartz Crystal

Quartz crystals are generally considered to be the mainstay of frequency control components due to their low cost and ease of implementation. They are available from numerous vendors in many package types and with various specification options.



13.22 Fixed Crystal Oscillator

A packaged fixed crystal oscillator comprises of an inverter, a quartz crystal, and passive components conveniently packaged together. The device renders a strong, consistent square wave output. Oscillators used with microprocessors are supplied in many configurations and tolerances.

Crystal oscillators should be restricted for use in special situations, such as shared clocking among devices or multiple controllers. Since clock routing can be difficult to accomplish, it is preferable to provide a separate crystal for each device.

Note: Contact your Intel Field Service Representative to obtain the most current device documentation prior to implementing this solution.

13.23 Crystal Selection Parameters

All crystals used with Intel Ethernet controllers are described as “AT-cut,” which refers to the angle at which the unit is sliced with respect to the long axis of the quartz stone.

The following table lists crystals which have been used successfully in past designs. (No particular product is recommended.)

Table 13-10. Crystal Manufacturers and Part Numbers

Manufacturer	Part No.	Note
TXC*	9C25000008	HC-49S SMD
NDK*	41CD25.0F1303018	HC-49S SMD
TXC*	7V25020004	Small SMD 3.2mmx2.5mm
KDS*	DSX321G, 1B/C/N/H225000CC0M	Small SMD 3.2mmx2.5mm
River*	FCX-04-25MJ90141	Small SMD 3.2mmx2.5mm

The datasheet for the PHY lists the crystal electrical parameters and provides suggested values for typical designs. Designers should refer to criteria outlined in their respective PHY datasheet. The parameters are described in the following subsections.

13.24 Vibrational Mode

Crystals in the frequency range referenced above are available in both fundamental and third overtone. Unless there is a special need for third overtone, fundamental mode crystals should be used.

13.25 Nominal Frequency

Intel Ethernet controllers use a crystal frequency of 25.000 MHz. The 25 MHz input is used to generate a 125 MHz transmit clock for 100BASE-TX and 1000BASE-TX operation, and 10 MHz and 20-MHz transmit clocks, for 10BASE-T operation.

13.26 Frequency Tolerance

The frequency tolerance for an Ethernet Platform LAN Connect device is dictated by the IEEE 802.3 specification as ± 50 parts per million (ppm). This measurement is referenced to a standard temperature of 25 °C. Intel recommends a frequency tolerance of ± 30 ppm to ensure for any frequency variance contributed by the PCB.

13.27 Temperature Stability and Environmental Requirements

Temperature stability is a standard measure of how the oscillation frequency varies over the full operational temperature range (and beyond). Several optional temperature ranges are currently available, including -40 °C to +85 °C for industrial environments. Some vendors separate operating temperatures from temperature stability. Manufacturers may also list temperature stability as 50 ppm in their data sheets.

Note: Crystals also carry other specifications for storage temperature, shock resistance, and reflow solder conditions. Crystal vendors should be consulted early in the design cycle to discuss its application and environmental requirements.


13.28 Calibration Mode

The terms “series-resonant” and “parallel-resonant” are often used to describe crystal oscillator circuits. Specifying parallel mode is critical to determining how the crystal frequency is calibrated at the factory.

A crystal specified and tested as series resonant oscillates without problem in a parallel-resonant circuit, but the frequency is higher than nominal by several hundred parts per million. The purpose of adding load capacitors to a crystal oscillator circuit is to establish resonance at a frequency higher than the crystal’s inherent series resonant frequency.

Figure 13-20 illustrates a simplified schematic of the internal oscillator circuit. Pin X1 and X2 refers to XTAL_IN and XTAL_OUT in the Ethernet device, respectively. The crystal and the capacitors form a feedback element for the internal inverting amplifier. This combination is called parallel-resonant, because it has positive reactance at the selected frequency. In other words, the crystal behaves like an inductor in a parallel LC circuit. Oscillators with piezoelectric feedback elements are also known as “Pierce” oscillators.

Figure 13-20. Thermal Oscillator Circuit



13.29 Load Capacitance

The formula for crystal load capacitance is as follows:

$$C_L = \frac{(C1 \cdot C2)}{(C1 + C2)} + C_{stray}$$

where $C1 = C2 = 22 \text{ pF}$ (as suggested in most Intel reference designs) and C_{stray} = allowance for additional capacitance in pads, traces and the chip carrier within the Ethernet device package and C_{damp} .

13.30 Shunt Capacitance

The shunt capacitance parameter is relatively unimportant compared to load capacitance. Shunt capacitance represents the effect of the crystal's mechanical holder and contacts. The shunt capacitance should be a maximum of 6 pF.

13.31 Equivalent Series Resistance

Equivalent Series Resistance (ESR) is the real component of the crystal's impedance at the calibration frequency, which the inverting amplifier's loop gain must overcome. ESR varies inversely with frequency for a given crystal family. The lower the ESR, the faster the crystal starts up. Crystals with an ESR value of 50Ω or better should be used.



13.32 Drive Level

Drive level refers to power dissipation in use. The allowable drive level for a Surface Mounted Technology (SMT) crystal is less than its through-hole counterpart. This is due to the fact that surface mount crystals are typically made from narrow, rectangular AT strips, rather than circular AT quartz blanks.

When selecting a crystal, board designers must ensure that the crystal specification meets at least the drive level specified. For example, if the crystal drive level specification states that the drive level is 200 μ W maximum, then the crystal drive level must be at least 200 μ W. So, a 500 μ W crystal is sufficient, but a 100 μ W crystal is not.

13.33 Aging

Aging is a permanent change in frequency (and resistance) occurring over time. This parameter is most important in its first year because new crystals age faster than old crystals. Crystals with a maximum value of ± 5 ppm per year aging should be used.

13.34 Reference Crystal

The normal tolerances of the discrete crystal components can contribute to small frequency offsets with respect to the target center frequency. To minimize the risk of tolerance-caused frequency offsets causing a small percentage of production line units to be outside of the acceptable frequency range, it is important to account for those shifts while empirically determining the proper values for the discrete loading capacitors, C1 and C2.

Even with a perfect support circuit, most crystals will oscillate slightly higher or lower than the exact center of the target frequency. Therefore, frequency measurements, which determine the correct value for C1 and C2, should be performed with an ideal reference crystal. When the capacitive load is exactly equal to the crystal's load rating, an ideal reference crystal will be perfectly centered at the desired target frequency.

13.34.1 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal:

- If a Saunders and Associates (S&A) crystal network analyzer is available, then discrete crystal components can be tested until one is found with zero or nearly zero ppm deviation (with the appropriate capacitive load). A crystal with zero or near zero ppm deviation will be a good reference crystal to use in subsequent frequency tests to determine the best values for C1 and C2.
- If a crystal analyzer is not available, then the selection of a reference crystal can be done by measuring a statistically valid sample population of crystals, which has units from multiple lots and approved vendors. The crystal, which has an oscillation frequency closest to the center of the distribution, should be the reference crystal used during testing to determine the best values for C1 and C2.
- It may also be possible to ask the approved crystal vendors or manufacturers to provide a reference crystal with zero or nearly zero deviation from the specified frequency when it has the specified C_{Load} capacitance.



When choosing a crystal, designers must keep in mind that to comply with IEEE specifications for 10/100 Mbps operation and 10/100/1000 Mbps operation if applicable, the transmitter reference frequency must be precise within ± 50 ppm. Intel recommends customers use a transmitter reference frequency that is accurate to within ± 30 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance.

13.34.2 Circuit Board

Since the dielectric layers of the circuit board are allowed some reasonable variation in thickness, the stray capacitance from the printed board (to the crystal circuit) will also vary. If the thickness tolerance for the outer layers of dielectric are controlled within $\pm 15\%$ of nominal, then the circuit board should not cause more than ± 2 pF variation to the stray capacitance at the crystal. When tuning crystal frequency, it is recommended that at least three circuit boards are tested for frequency. These boards should be from different production lots of bare circuit boards.

Alternatively, a larger sample population of circuit boards can be used. A larger population will increase the probability of obtaining the full range of possible variations in dielectric thickness and the full range of variation in stray capacitance.

Next, the exact same crystal and discrete load capacitors (C1 and C2) must be soldered onto each board, and the LAN reference frequency should be measured on each circuit board.

The circuit board, which has a LAN reference frequency closest to the center of the frequency distribution, should be used while performing the frequency measurements to select the appropriate value for C1 and C2.

13.34.3 Temperature Changes

Temperature changes can cause the crystal frequency to shift. Therefore, frequency measurements should be done in the final system chassis across the system's rated operating temperature range.

13.35 Oscillator Support

The PHY clock input circuit is optimized for use with an external crystal. However, an oscillator can also be used in place of the crystal with the proper design considerations:

- The clock oscillator has an internal voltage regulator to isolate it from the external noise of other circuits to minimize jitter. If an external clock is used, this imposes a maximum input clock amplitude. For example, if a 3.3V DC oscillator is used, its output signal should be attenuated to a maximum value with a resistive divider circuit.
- The input capacitance introduced by the PHY (approximately 11 to 13pF) is greater than the capacitance specified by a typical oscillator (approximately 15pF).
- The input clock jitter from the oscillator can impact the PHY clock and its performance.

Note: The power consumption of additional circuitry equals about 1.5mW.

The following table lists oscillators that can be used with the PHY. Note that no particular oscillator is recommended):

Table 13-11. Oscillator Manufacturers and Part Numbers

Manufacturer	Part No.
Kyocera*	K30-3C0-SE-25.0000M
MtronPTI*	M214TCN25.0000MHz
TXC*	7C25000230 7X25080001

13.36 Oscillator Placement and Layout Recommendations

Oscillator clock sources should not be placed near I/O ports or board edges. Radiation from these devices can be coupled into the I/O ports and radiate beyond the system chassis. Oscillators should also be kept away from the Ethernet magnetics module to prevent interference.

The oscillator must have its own decoupling capacitors and they must be placed within 0.25 inches. If a power trace is used (not power plane), the trace from the capacitor to the oscillator must not exceed 0.25 inches in length. The decoupling capacitors help to improve the oscillator stability. The oscillator clock trace should be less than two inches from the PHY. If it is greater than 2 inches, then verify the signal quality, jitter, and clock frequency measurements at the PHY.

The clock lines should also target $50\Omega \pm 15\%$ and should have 10Ω series back termination placed close to the series oscillator. To help reduce EMI, the clock lines must be a distance of at least five times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, the board edge, transformers and differential pairs.

The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition. The clock lines should not cross or run in parallel with any trace (100 MHz signal or higher) on an adjacent layer.

There should be a ferrite bead within 250 mils of the oscillator power pin and there must be a 1 uF or greater capacitor within 250 mils of the oscillator, connected to the power trace between the oscillator input and ferrite bead. With a ferrite bead on the power trace for the oscillator, there should be a power pour (or fat trace) to supply power to the oscillator.

Note: For latest PHY schematic connection recommendations, refer to the *Intel® Ethernet Controller I217 Reference Schematic*, and the *Intel® Ethernet Controller I217 Schematic and Layout Checklist*. For access to documentation, contact your Intel Field Service Representative.

13.37 LAN Switch

The following table lists LAN switches that can be used with the Intel® Ethernet Controller I217. Note that no particular LAN switch is recommended:



Table 13-12. LAN Switch Manufacturers and Part Numbers

Manufacturer	Part No.
Pericom*	PI3L500-AZ
Texas Instruments*	TS3L500AE

13.38 Troubleshooting Common Physical Layout Issues

The following is a list of common physical layer design and layout mistakes in LAN on Motherboard (LOM) designs.

1. Lack of symmetry between the two traces within a differential pair. Asymmetry can create common-mode noise and distort the waveforms. For each component and via that one trace encounters, the other trace should encounter the same component or a via at the same distance from the Ethernet silicon.
2. Unequal length of the two traces within a differential pair. Inequalities create common-mode noise and will distort the transmit or receive waveforms.
3. Excessive distance between the Ethernet silicon and the magnetics. Long traces on FR4 fiberglass epoxy substrate will attenuate the analog signals. In addition, any impedance mismatch in the traces will be aggravated if they are longer than the four-inch guideline.
4. Routing any other trace parallel to and close to one of the differential traces. Crosstalk getting onto the receive channel will cause degraded long cable BER. Crosstalk getting onto the transmit channel can cause excessive EMI emissions and can cause poor transmit BER on long cables. At a minimum, for stripline other signals should be kept at least 6x the height of the thinnest adjacent dielectric layer. For microstrip it is 7x. The only possible exceptions are in the vicinities where the traces enter or exit the magnetics, the RJ-45 connector, and the Ethernet silicon.
5. Using a low-quality magnetics module.
6. Reusing an out-of-date physical layer schematic in a Ethernet silicon design. The terminations and decoupling can be different from one PHY to another.
7. Incorrect differential trace impedances. It is important to have about a 100Ω impedance between the two traces within a differential pair. This becomes even more important as the differential traces become longer. To calculate differential impedance, many impedance calculators only multiply the single-ended impedance by two. This does not take into account edge-to-edge capacitive coupling between the two traces. When the two traces within a differential pair are kept close to each other, the edge coupling can lower the effective differential impedance by 5Ω to 20Ω. Short traces will have fewer problems if the differential impedance is slightly off target.

13.39 Power Delivery

The Intel® Ethernet Controller I217 requires a 3.3V power rail and a Core power rail. The internal 3.3V power rail is brought out for decoupling. [Figure 13-21](#) shows a typical power delivery configuration that can be implemented. However, power delivery can be customized based on a specific OEM. In general planes should be used to deliver 3.3 Vdc and the Core voltage. Not using planes can cause resistive voltage drop and/or



inductive voltage drop (due to transient or static currents). Some of the symptoms of these voltage drops can include higher EMI, radiated immunity, radiated emissions, IEEE conformance issues, and register corruption.

Decoupling capacitors (0.1 uF and smaller) should be placed within 250 mils of the LAN device. They also should be distributed around the PHY and some should be in close proximity to the power pins.

The bulk capacitors (1.0 uF or greater) should be placed within 1 inch if using a trace (50 mils wide or wider) or within 1.5 inches if using a plane.

The Core power rail for the Intel® Ethernet Controller I217 is generated using the integrated SVR (iSVR). Using the internal SVR to generate the Core power, the inductor must be placed within 0.5" of the input pin to the PHY and connected with a trace wider than or equal to 20 mil wide. (Please see the reference schematic for further details regarding the Core power rail.)

The following list shows inductors that have been used successfully with Intel® designs:


Table 13-13. Inductors and Manufacturers

Manufacturer	Part Number
muRata*	LQH32PN4R7NN0
muRata*	LQH32CN4R7M53
TDK*	FLF3215T-4R7M
TDK*	VLS252015ET‡
Taiyo Yuden*	NRS2012T4R7MGJ
TDK*	VLS2012MT-4R7M

‡ The height of this part may not be appropriate for some designs.



Figure 13-21. Intel® Ethernet Controller I217 Power Delivery Diagram



Note: For latest PHY schematic connection recommendations, refer to the Intel® Ethernet Controller I217 reference schematic. For access to documentation, contact your Intel Field Service Representative.

13.40 Power Sequencing

The Intel® Ethernet Controller I217 does not require any power sequencing between the 3.3V and Core power rails when the Core voltage is generated from the Intel® Ethernet Controller I217 internal SVR. See the datasheet and reference schematic for details.



NOTE: *This page intentionally left blank.*

§ §



14 Non-Mobile Design Considerations & Guidelines

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel® Ethernet Connection I217 (PHY) shown in [Figure 14-1](#). Its bus master capabilities enable the component to process high-level commands and perform multiple operations, which lowers processor use by off loading communication tasks from the processor.

The PCH, which hereinafter refers to the integrated MAC within the PCH, supports the SMBus interface for manageability while in an Sx state and PCI Express* (PCIe*) for 10/100/1000 Mbps traffic in an S0 state.

Note: The Intel® Ethernet Connection I217 PCIe interface is not PCIe compliant. It operates at half the PCI Express* (PCIe*) Specification v1.0 (2.5 GT/s) speed. In this chapter, the term “PCIe-based” is interchangeable with “PCIe.” There are no design layout differences between normal PCIe and the Intel® Ethernet Controller I217 PCIe-based interface.

The Intel® Ethernet Controller I217 interfaces with the integrated MAC through two interfaces: PCIe and SMBus. In SMBus mode, the link speed is reduced to 10 Mbps. The PCIe interface incorporates two aspects: a PCIe-based SerDes (electrically) and a custom logic protocol for messaging between the integrated MAC and the PHY.

Note: Gigabit Ethernet requires an SPI Flash to host firmware and does not work without an SPI Flash on board.

The integrated MAC supports multi-speed operation (10/100/1000 Mbps). The integrated MAC also operates in full-duplex at all supported speeds or half-duplex at 10/100 Mbps as well as adhering to the IEEE 802.3x Flow Control Specification.

Note: References to the AUX power rail means the power rail is available in all power states including G3 to S5 transitions and Sx states with Wake on LAN (WoL) enabled. For example, V3P3_AUX in this chapter refers to a rail that is powered under the conditions previously mentioned.

Figure 14-1. PCH/PHY Interface Connections






Table 14-1. SMBus Data Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Data	SMB_DATA	SMLINK0_DATA	SMBus data

Table 14-2. PCIe Data Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Data	PETp PETn	PETp PETn	PCIe transmit pair
Data	PERp PERn	PERp PERn	PCIe receive pair

Notes:

1. Can be connected to any PCIe port on the integrated MAC. The appropriate NVM descriptor soft strap (PCHSTRP9) should define which PCIe port is configured as GbE LAN.

Table 14-3. Clock and Reset Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Clock	SMB_CLK	SML0_CLK	SMBus clock
Clock	PE_CLKP PE_CLKN	CLKOUT_PCIE[7:0]_P ¹ CLKOUT_PCIE[7:0]_N ¹	PCIe* clock
Clock ²	CLK_REQ_N	Not Connected	PCIe clock request
Reset	PE_RST_N	PLTRST# ³	PCIe reset

Notes:

1. These signals come from the PCH and drive the PHY.
2. See Fig. 1-9 for connection information.

14.1 PHY Overview

The PHY is a single port compact component designed for 10/100/1000 Mbps operation. It enables a single port Gigabit Ethernet (GbE) implementation in a very small area, easing routing constraints from the PCH chipset to the PHY.

The PHY provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3ab, 802.3u, and 802.3i, respectively).

14.1.1 PHY Interconnects

The main interfaces for either PHY are PCIe and SMBus on the host side and Media Dependent Interface (MDI) on the link side. Transmit traffic is received from the PCH as either PCIe or SMBus packets on the host interconnect and transmitted as Ethernet packets on the MDI link. Receive traffic arrives as Ethernet packets on the MDI link and transferred to the PCH through either the PCIe or SMBus interconnects.

The PHY switches the in-band traffic automatically between PCIe and SMBus based on platform reset. The transition protocol is done through SMBus. The PCIe interface is powered down when the Ethernet link is running in an Sx state.



14.1.2 PCIe-Based Interface

A high-speed SerDes interface that uses PCIe electrical signaling at half speed while utilizing a custom logical protocol for active state operation mode.

Note: PCIe validation tools cannot be used for electrical validation of this interface; however, PCIe layout rules apply for on-board routing.

14.1.2.0.1 PCIe Interface Signals

The signals used to connect between the PCH and the PHY in this mode are:

- Serial differential pair running at 1.25 Gb/s for Rx.
- Serial differential pair running at 1.25 Gb/s for Tx.
- 100-MHz differential clock input to the PHY is generated by the PCH.
- Power and clock good indication to the PHY PE_RST_N.
- Clock control through CLK_REQ_N (refer to [Table 14-3](#)). This PHY output should be pulled up with a 10kΩ resistor connected to 3.3V DC AUX power (present in G3 to S5).

14.1.2.0.2 PCIe Operation and Channel Behavior

The PHY only runs at 1250 Mbps speed, which is 1/2 of the Gen 1 2.5 Gb/s PCIe frequency. Each of the PCIe root ports in the PCH has the ability to run at 1250 Mbps. Configuring a PCH PCIe port that is attached to a PCIe Intel PHY device is pre-loaded from the GbE region of the NVM. The selected port adjusts the transmitter to run at 1/2 the Gen 1 PCIe speed and does not need to be PCIe compliant.

Packets transmitted and received over the PCIe interface are full Ethernet packets and not PCIe transaction/link/physical layer packets.

14.1.2.0.3 PCIe Connectivity

The PHY transmit/receive pins are output/input signals and are connected to the PCH as listed in [Table 14-1](#) through [Table 14-3](#).

14.1.2.0.4 PCIe Reference Clock

The PCIe Interface uses a 100-MHz differential reference clock, denoted PE_CLKP and PE_CLKN. This signal is typically generated on the platform and routed to the PCIe port.

The frequency tolerance for the PCIe reference clock is ± 300 ppm.

14.1.3 SMBus Interface

SMBus is a low speed (100 kHz/400 kHz/1000 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the PHY and the PCH when the platform is in a low power state (Sx). The interface is also used to enable the PCH to configure the PHY as well as passing in-band information between them.

The SMBus uses two primary signals: SMBCLK and SMBDATA, to communicate. Both of these signals float high with board-level $2.2\text{k}\Omega \pm 5\%$ pull-up resistors.



The SMBus specification has defined various types of message protocols composed of individual bytes. For more details about SMBus, see the SMBus specification.

14.1.3.0.1 SMBus Connectivity

Table 14-1 through Table 14-3 list the relationship between PHY SMBus pins to the PCH LAN SMBus pins.

Note: The Intel® Ethernet Controller I217 SMBus signals (SMB_DATA and SMB_CLK) cannot be connected to any other devices other than the integrated MAC. Connect the SMB_DATA and SMB_CLK pins to the integrated MAC SML0DATA and SML0CLK pins, respectively.

14.1.4 PCIe and SMBus Modes

In GbE operation, PCIe is used to transmit and receive data and for MDIO status and control. The PHY automatically switches the in-band traffic between PCIe and SMBus based on the platform power state. The table below lists the operating modes of PCIe and SMBus.

The Intel® Ethernet Controller I217 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

System/Intel Management Engine State	PHY	
	SMBus	PCIe
S0 and PHY Power Down	Not used	Electrical Idle (EI)
S0 and Idle or Link Disconnect	Not used	EI
S0 and Link in Low Power Idle (LPI)	Not used	EI
S0 and active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

14.1.5 Transitions Between PCIe and SMBus Interfaces

14.1.5.0.1 Switching from SMBus to PCIe

Communication between the integrated MAC and the PHY is done through the SMBus each time the system is in a low power state (Sx). The integrated MAC/PHY interface is needed while the Manageability Engine (ME) is still active to transfer traffic, configuration, control and status or to enable host wake up from the PHY.

Possible states for activity over the SMBus:

1. After power on (G3 to S5).
2. On system standby (Sx).

The switching from the SMBus to PCIe is done when the PE_RST_N signal goes high.

- Any transmit/receive packet that is not completed when PE_RST_N is asserted is discarded.
- Any in-band message that was sent over the SMBus and was not acknowledged is re-transmitted over PCIe.



14.1.5.0.2 Switching from PCIe to SMBus

The communication between the integrated MAC and the PHY is done through PCIe each time the platform is in active power state (S0). Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME.

The switching from PCIe to SMBus is done when the PE_RST_N signal goes low.

- Any transmit/receive packet that is not completed when PE_RST_N goes to 0b is discarded.
- Any in-band message that was sent over PCIe and was not acknowledged is re-transmitted over SMBus.

14.2 Platform LAN Design Guidelines

These sections provide recommendations for selecting components and connecting special pins. For GbE designs, the main elements are the PCH chipset, the Intel® Ethernet Controller I217, a magnetics module and RJ-45 connector, a GbE region NVM (Non Volatile Memory) image, and a clock source.

14.2.1 General Design Considerations for PHYs

Sound engineering practices must be followed with respect to unused inputs by terminating them with pull-up or pull-down resistors, unless otherwise specified in a datasheet, design guide or reference schematic. Pull-up or pull-down resistors must not be attached to any balls identified as “No Connect.” These devices might have special test modes that could be entered unintentionally.

14.2.1.0.1 Clock Source

All designs require a 25-MHz clock source. The PHY uses the 25-MHz source to generate clocks up to 125 MHz and 1.25 GHz for both the PHY circuits and the PCIe interface. For optimum results with lowest cost, a 25-MHz parallel resonant crystal can be used along with the appropriate load capacitors at the XTAL_OUT (X2) and XTAL_IN (X1) leads. The frequency tolerance of the timing device should equal 30 ppm or better. Further detail is found in [Section 14.18](#) and [Section 14.34](#).

Note: XTAL_OUT and XTAL_IN are the signal names for the PHY.

There are three steps to crystal qualification:

1. Verify that the vendor’s published specifications in the component datasheet meet the required conditions for frequency, frequency tolerance, temperature, oscillation mode and load capacitance as specified in the respective datasheet.
2. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems.
3. Independently measure the component’s electrical parameters in real systems. Measure frequency at a test output to avoid test probe loading effects at the PHY. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications. For crystals, it is also important to examine startup behavior while varying system voltage and temperature.



14.2.1.0.2 Magnetics Module

The magnetics module has a critical effect on overall IEEE and emissions conformance. The device should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Carefully qualifying new magnetics modules prevents problems that might arise because of interactions with other components or the printed circuit board itself.

The steps involved in magnetics module qualification are similar to those for crystal qualification:

1. Verify that the vendor’s published specifications in the component datasheet meet or exceed the required IEEE specifications.
2. Independently measure the component’s electrical parameters on the test bench, checking samples from multiple lots. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications.
3. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems. Vary temperature and voltage while performing system level tests.

Magnetics modules for 1000BASE-T Ethernet as used by the PHY only are similar to those designed solely for 10/100 Mbps, except that there are four differential signal pairs instead of two. Refer to the Intel® Ethernet Controller I217 datasheet for specific electrical requirements that the magnetics need to meet.

The following magnetics modules are not recommended; however, they have been used successfully in previous designs:

Table 14-4. Magnetic Modules and Manufacturers

Manufacturer	Part Number	Note
SpeedTech*	P25BPP4MFRT9	USB stack, 8core
SpeedTech*	P25BFB4-RDW9	USB stack, 12core
Foxconn*	JFM38U1A-21C7-4F	USB stack, 8core
Foxconn*	JFM38U1A-7110-4F	USB stack, 8core
Tyco*	1840023-1	USB stack, 8core

14.2.1.0.3 Criteria for Integrated Magnetics Electrical Qualification

The following table gives the criteria used to qualify integrated magnetics.



Table 14-5. Integrated Magnetics Recommended Qualification Criteria

Open Circuit Inductance (OCL)	w/8 mA DC bias; at 25C	400uH Min
	w/8 mA DC bias; at 0C to 70C	350uH Min
Insertion Loss	100 kHz through 999 kHz	1dB Max
	1.0 MHz through 60.0 MHz	0.6dB Max
	60.1 MHz through 80.0 MHz	0.8dB Max
	80.1 MHz through 100.0 MHz	1.0dB Max
	100.1 MHz through 125.0 MHz	2.4dB Max
Return Loss	1.0 MHz through 40.0 MHz	18.0 dB Min
	40.1 MHz through 100.0 MHz When reference impedance is 85 Ohms, 100 Ohms, and 115 Ohms. Note that R.L. values may vary with MDI trace lengths. The LAN magnetics may need to be measured in the platform where it will be used.	12–20 * LOG (Freq in MHz / 80) dB Min
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	-50.3+(8.8*(Freq in MHz / 30)) dB Max
	30.0 MHz through 250.0 MHz	-(26 -(16.8*(LOG(Freq in MHz / 250 MHz)))) dB Max
	250.1 MHz through 375.0 MHz	-26.0 dB Max
Crosstalk Isolation Integrated Modules (Proposed)	1.0 MHz through 10 MHz	-50.8+(8.8*(Freq in MHz / 10)) dB Max
	10.0 MHz through 100.0 MHz	-(26 -(16.8*(LOG(Freq in MHz / 100 MHz)))) dB Max
	100 MHz through 375.0 MHz	-26.0 dB Max
Diff to CMR	1 MHz through 29.9 MHz	-40.2+(5.3*((Freq in MHz / 30)) dB Max
	30.0 MHz through 500 MHz	-(22-(14*(LOG((Freq in MHz / 250)))) dB Max
CM to CMR	1 MHz through 270 MHz	-57+(38*((Freq in MHz / 270)) dB Max
	270.1 MHz through 300 MHz	-17-2*((300-(Freq in MHz) / 30) dB Max
	300.1 MHz through 500 MHz	-17 dB Max
Hi-Voltage Isolation	1500 Vrms at 50 or 60 Hz for 60 sec. or:2250 Vdc for 60 seconds	Minimum

14.2.2 NVM for PHY Implementations

The LAN only supports an SPI Flash, which is connected to the PCH. Several words of the NVM are accessed automatically by the device after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the NVM space is available to software for storing the MAC address, serial numbers, and additional information. More details may be obtained from the Intel® Ethernet Controller I217 Datasheet.

Intel has an MS-DOS* software utility called EUpdate that is used to program the SPI Flash images in development or production line environments. A copy of this program can be obtained through your Intel representative.

14.2.3 LED

The PHY has three LED outputs that can be configured via the NVM. The hardware configuration is shown in [Figure 14-2](#).


Refer to the Intel® Ethernet Controller I217 Reference Schematic for default LED color based on reference design.

Refer to the Intel® Ethernet Controller I217 Datasheet for details regarding the programming of the LED’s and the various modes. The default values for the PHY (based on the LED NVM setting—word 0x18 of the LAN region) are listed in the table below:

Table 14-6. LED Default Values

LED	Mode	Color	Blink	Polarity
LED0	Link Up/Activity	Green	200 ms on/200 ms off	Active low
LED1	Link 1000	Yellow	No	Active low
LED2	Link 100	Green	No	Active low

Figure 14-2. LED Hardware Configuration



14.2.3.0.1 RBIAS

RBIAS requires external resistor connection to bias the internal analog section of the device. The input is sensitive to the resistor value. Resistors of 1% tolerance must be used. Connect RBIAS through a 3.01 kΩ 1% pull-down resistor to ground and then place it no more than one half inch (0.5”) away from the PHY.

14.2.3.0.2 LAN Disable

The PHY enters a power-down state when the LAN_DISABLE_N pin is asserted low. Exiting this mode requires setting the LAN_DISABLE_N pin to a logic one. Connect LAN_DISABLE_N to LAN_PHY_PWR_CTRL /GPIO12 on the PCH.

14.2.4 Exposed Pad* (e-PAD) Design and SMT Assembly Guide

14.2.4.0.1 Overview

This section provides general information about e-PAD and SMT assemblies. Chip packages have exposed die pads on the bottom of each package to provide electrical interconnections with the printed circuit board. These e-PADs also provide excellent thermal performance through efficient heat paths to the PCB.



Packages with e-PADs are very popular due to their low cost. Note that this section only provides basic information and references in regards to the e-PAD. It is recommended that each customer consult their fab and assembly house to obtain more details on how to implement the e-PAD package design. Each fab and assembly house might need to tune the land pattern/stencil and create a solution that best suits their methodology and process.

14.2.4.0.2 PCB Design Requirements

In order to maximize both heat removal and electrical performance, a land pattern must be incorporated on the PCB within the footprint of the package corresponding to the exposed metal pad or exposed heat slug of the package as shown in the following figures. Refer to the specific product datasheet for actual dimensions.

Note: Due to the Intel® Ethernet Controller I217 package size, a via-in-pad configuration must be used [Figure 14-3](#) and [Figure 14-4](#) are general guidelines see [Figure 14-5](#) for the Intel® Ethernet Controller I217-specific via-in-pad thermal pattern recommendations.

Figure 14-3. Typical ePAD Land Pattern






Figure 14-4. Typical Thermal Pad and Via Recommendations



Note: Encroached and uncapped via configurations have voids less than the maximum allowable void percentage. Uncapped via provides a path for trapped air to escape during the reflow soldering process.

Note: Secondary side solder bumps might be seen in an uncapped via design. This needs to be considered when placing components on the opposite side of the PHY.

Figure 14-5. Recommended Thermal Via Patterns for the Intel® Ethernet Controller I217



14.2.4.0.3 Board Mounting Guidelines

The following are general recommendations for mounting a QFN-48 device on the PCB. This should serve as the starting point in assembly process development and it is recommended that the process should be developed based on past experience in mounting standard, non-thermally/electrically enhanced packages.



14.2.4.0.4 Stencil Design


For maximum thermal/electrical performance, it is required that the exposed pad/slug on the package be soldered to the land pattern on the PCB. This can be achieved by applying solder paste on both the pattern for lead attachment as well as on the land pattern for the exposed pad. While for standard (non-thermally/-electrically enhanced) lead-frame based packages the stencil thickness depends on the lead pitch and package co-planarity, the package standoff must also be considered for the thermally/ electrically enhanced packages to determine the stencil thickness. In this case, a stencil foil thickness in the range of 5-6 mils (or 0.127-0.152 mm) is recommended; likely or practically, a choice of either 5 mils or 6 mils. Tolerance-wise, it should not be worse than ±0.5 mil.

Note: Industry specialists typically use ±0.1-mil tolerance on stencil for its feasible precision.

The aperture openings should be the same as the solder mask openings on the land pattern. Since a large stencil opening may result in poor release, the aperture opening should be subdivided into an array of smaller openings, similar to the thermal land pattern shown in the figure below.

Note: Refer to the specific product datasheet for actual dimensions.

Figure 14-6. Stencil Design Recommendation for QFN Package Type Products



Important General Guidelines:

- The Stencil Aperture Openings of the e-PAD must not go outside of the exposed landing area (solder mask opening).

- The Stencil Aperture Openings of the e-PAD should be about 80% of the exposed landing area (solder mask opening).


The Intel® Ethernet Connection I217 e-PAD has $D2=E2=3\text{mm}$, therefore the Stencil Design can only have four aperture openings for the e-PAD. This can be achieved by setting $R=S=1.35\text{mm}$ and $T=0.3$. Using this arrangement, the Apertures/e-PAD area is 81% of the exposed landing area (solder mask opening).

This information is only intended as general guidance. Please consult with the manufacturer to confirm the final design meets requirements.

14.2.4.0.5 Assembly Process Flow

The following figure below shows the typical process flow for mounting packages to the PCB.


Figure 14-7. Assembly Flow



14.2.4.0.6 Reflow Guidelines

The typical reflow profile consists of four sections. In the preheat section, the PCB assembly should be preheated at the rate of 1 to 2 °C/sec to start the solvent evaporation and to avoid thermal shock. The assembly should then be thermally soaked for 60 to 120 seconds to remove solder paste volatiles and for activation of flux. The reflow section of the profile, the time above liquidus should be between 45 to 60 seconds with a peak temperature in the range of 245 to 250 °C, and the duration at the peak should not exceed 30 seconds. Finally, the assembly should undergo cool down in the fourth section of the profile. A typical profile band is provided in the figure below, in which 220 °C is referred to as an approximation of the liquidus point. The actual profile parameters depend upon the solder paste used and specific recommendations from the solder paste manufacturers should be followed.

Figure 14-8. Typical Profile Band



1. Preheat: 125°C–220°C, 150–210s at 0.4k/s to 1.0k/s
2. Time at $T > 220^{\circ}\text{C}$: 60–90s
3. Peak Temperature: 245–50°C
4. Peak time: 10–30s
5. Cooling rate: $\leq 6\text{k/s}$
6. Time from 25°C to Peak: 240–360s
7. Intel recommends a maximum solder void of 50% after reflow.

Note: Contact your Intel Field Service Representative for any designs unable to meet the recommended guidance for E-pad implementation.


14.3 PCH—SMBus/PCIe LOM Design Guidelines

This section contains guidelines on how to implement a PCH/PHY single solution on a system motherboard. It should not be treated as a specification, and the system designer must ensure through simulations or other techniques that the system meets the specified timings. The following are guidelines for both PCH SMBus and PCIe interfaces. Note that PCIe is only applicable to the PHY.

The SMBus/PCIe Interface between the PCH and PHY is shown at high level in Figure 15-9. For complete design details always refer to the Intel® Ethernet Controller I217 Reference Schematic.

Refer to [Section 14.6](#) for PCI Express Routing Guidelines.

Figure 14-9. Single Solution Interconnect



Notes:

1. For latest PHY schematic connection recommendations, refer to the Intel® Ethernet Controller I217 Reference Schematic and the Intel® Ethernet Controller I217 Schematic and Layout Checklist. For access to documentation, contact your Intel Field Service Representative.
2. Any free PCIe lanes [8:1] on the PCH can be used to connect the PCIe interface on the Intel® Ethernet Controller I217.
3. Any free CLKOUT_PCIE[7:0] and PCIECLKRQ[7:0]# from PCH can be used to connect to PE_CLK[P:N] and CLK_REQ_N on the Intel® Ethernet Controller I217 respectively.
4. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.

14.4 SMBus Design Considerations

No single SMBus design solution works for all platforms. Designers must consider the total bus capacitance and device capabilities when designing SMBus segments. Routing SMBus to the PCI slots makes the design process even more challenging since they add so much capacitance to the bus. This extra capacitance has a large affect on the bus time constant which in turn affects the bus rise and fall times.

Primary considerations in the design process are:

- Amount of V_{CC_SUS3_3} current available, that is, minimizing load of V_{CC_SUS3_3}.
- The pull-up resistor size for the SMBus data and clock signals is dependent on the bus load (this includes all device leakage currents). Generally the SMBus device that can sink the least amount of current is the limiting agent on how small the resistor can be. The pull-up resistor cannot be made so large that the bus time constant (Resistance X Capacitance) does not meet the SMBus rise and time specification.
- The maximum bus capacitance that a physical segment can reach is 400 pF.
- SMBus devices that can operate in STR must be powered by the V_{CC_SUS3_3} supply.



- It is recommended that I²C (Inter-Integrated Circuit) devices be powered by the V_{CC_core} supply. During an SMBus transaction in which the device is sending information to the integrated MAC, the device may not release the SMBus if the integrated MAC receives an asynchronous reset. V_{CC_core} is used to enable the BIOS to reset the device if necessary. SMBus 2.0- compliant devices have a timeout capability that makes them in-susceptible to this I²C issue, enabling flexibility in choosing a voltage supply.
- No other devices (except the integrated MAC and pull-up resistors) should be connected to the SMBus that connects to the PHY.
- **For system LAN on motherboard (LOM) designs:** The traces should be less than 70 inches for stripline and less than 100 inches for Microstrip. These numbers depend on the stackup, dielectric layer thickness, and trace width. The total capacitance on the trace and input buffers should be under 400 pF.
- **For system LAN on daughterboard designs:** Being conservative, the traces should be less than 7 inches for stripline designs and less than 10 inches for Microstrip designs. The lengths depend on the stackup, dielectric layer thickness, and trace width. Longer traces can be used as long as the total capacitance on the trace and input buffers is under 30 pF.

Note: Refer to [Section 14.1.3](#) for additional SMBus design considerations.

14.5 General Layout Guidelines

PHY interface signals must be carefully routed on the motherboard to meet the timing and signal quality requirements of their respective interface specifications. The following are some general guidelines that should be followed in designing a LAN solution. It is recommended that the board designer simulate the board routing to verify that the specifications are met for flight times and skews due to trace mismatch and crosstalk.

14.6 Layout Considerations

Critical signal traces should be kept as short as possible to decrease the likelihood of effects by high frequency noise of other signals, including noise carried on power and ground planes. This can also reduce capacitive loading.

Since the transmission line medium extends onto the printed circuit board, layout and routing of differential signal pairs must be done carefully.

Designing for GbE (1000BASE-T) operation is very similar to designing for 10/100 Mbps. For the PHY, system level tests should be performed at all three speeds.

14.7 Guidelines for Component Placement

Component placement can affect signal quality, emissions, and component operating temperature. Careful component placement can:

Decrease potential problems directly related to electromagnetic interference (EMI), which could cause failure to meet applicable government test specifications. In this case, place the PHY more than one inch from the edge of the board.

Simplify the task of routing traces. To some extent, component orientation affects the complexity of trace routing. The overall objective is to minimize turns and crossovers between traces.

14.7.1 PHY Placement Recommendations

Minimizing the amount of space needed for the PHY is important because other interfaces compete for physical space on a motherboard near the connector. The PHY circuits need to be as close as possible to the connector.

The following figure illustrates some basic placement distance guidelines. To simplify the diagram, it shows only two differential pairs, but the layout can be generalized for a GbE system with four analog pairs. The ideal placement for the PHY (LAN silicon) is approximately one inch behind the magnetics module.

While it is generally a good idea to minimize lengths and distances, this figure also illustrates the need to keep the PHY away from the edge of the board and the magnetics module for best EMI performance.

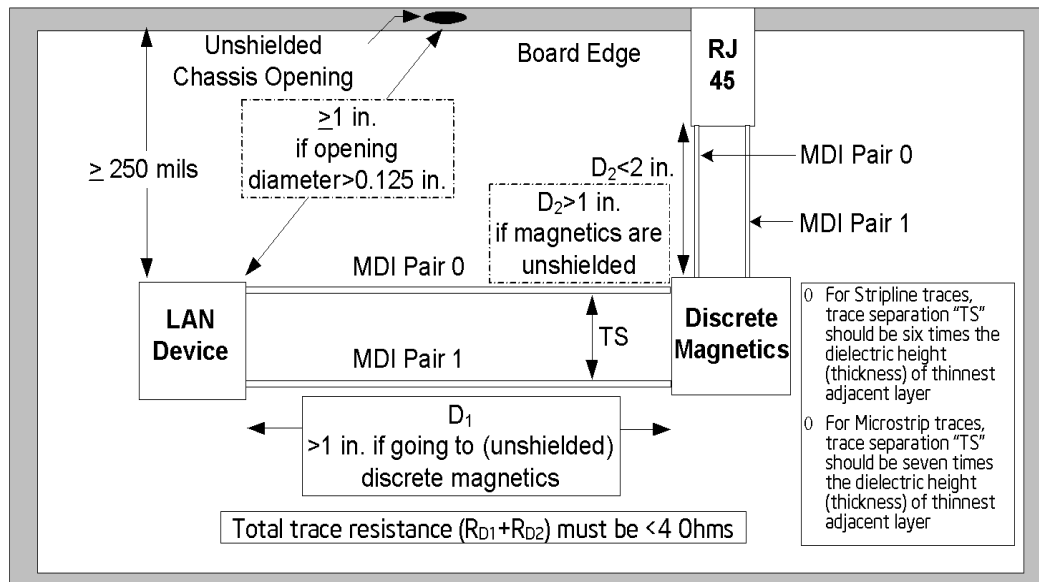


Figure 14-10. LAN Device Placement: At Least One Inch from Chassis Openings or Unshielded Connectors—Desktop

The PHY, referred to as “LAN Device” in the above figure, must be at least one inch from any chassis openings. To help reduce EMI, the following recommendations should be followed:

- Minimize the length of the MDI interface. See detail in table below: MDI Routing Summary
- Place the MDI traces no closer than 0.5 inch (1.3 cm) from the board edge.
- The Intel® Ethernet Controller I217 must be placed greater than 1" away from any hole to the outside of the chassis larger than 0.125 inches (125 mils) The larger the

hole the higher the probability the EMI and ESD immunity will be negatively affected.

- The Intel® Ethernet Controller I217 should be placed greater than 250mils from the board edge.
- If the connector or integrated magnetics module is not shielded, the Intel® Ethernet Controller I217 should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the Intel® Ethernet Controller I217 closer than one inch to Unshielded magnetics or connectors will increase the probability of failed EMI and common mode noise. If the LAN switch is too far away it will negatively affect IEEE return loss performance.
- The RBIAS trace length must be less than 1".
- Place the crystal less than one inch (2.54 cm) from the PHY.

Figure 14-11. PLC Placement: At Least One Inch from I/O Backplane





Figure 14-12. Effect of LAN Device Placed Too Close To an Rj-45 Connector or Chassis Opening



14.8 MDI Differential-Pair Trace Routing for LAN Design

Trace routing considerations are important to minimize the effects of crosstalk and propagation delays on sections of the board where high-speed signals exist. Signal traces should be kept as short as possible to decrease interference from other signals, including those propagated through power and ground planes.

14.9 Signal Trace Geometry

One of the key factors in controlling trace EMI radiation are the trace length and the ratio of trace-width to trace-height above the reference plane. To minimize trace inductance, high-speed signals and signal layers that are close to a reference or power plane should be as short and wide as practical. Ideally, the trace-width to trace-height above the ground plane ratio is between 1:1 and 3:1. To maintain trace impedance, the width of the trace should be modified when changing from one board layer to another if the two layers are not equidistant from the neighboring planes.

Each pair of signals should have a differential impedance of $100\ \Omega \pm 15\%$.

A set of trace length calculation tools are available from Intel (via the Intel Business Link (IBL)) to aid with MDI topology design. For access to documentation, contact your Intel Field Service Representative.

When performing a board layout, the automatic router feature of the CAD tool must not route the differential pairs without intervention. In most cases, the differential pairs will require manual routing.

Note: Measuring trace impedance for layout designs targeting $100\ \Omega$ often results in lower actual impedance due to over-etching. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of $105\ \Omega$ to $110\ \Omega$ should compensate for over-etching.

It is necessary to compensate for trace-to-trace edge coupling, which can lower the differential impedance by up to $10\ \Omega$, when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 14-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Signal group	MDI_PLUS[0:3] MDI_MINUS[0:3]		
Microstrip/stripline uncoupled single-ended impedance specification	$50\ \Omega \pm 10\%$		
Microstrip/stripline uncoupled differential impedance specification	$100\ \Omega \pm 15\%$		2,3
Microstrip nominal trace width	Design dependent	Design dependent	4
Microstrip nominal trace space	Design dependent	Design dependent	3,5
Microstrip/stripline trace length	4 in (102 mm)		6,7
Microstrip pair-to-pair space (edge-to-edge)	≥ 7 times the thickness of the thinnest adjacent dielectric layer		Figure 14-13
Stripline pair-to-pair space (edge-to-edge)	≥ 6 times the thickness of the thinnest adjacent dielectric layer		
Microstrip bus-to-bus spacing	≥ 7 times the thickness of the thinnest adjacent dielectric layer		
Stripline bus-to-bus spacing	≥ 6 times the thickness of the thinnest adjacent dielectric layer		

Notes:

1. Pair-to-pair spacing ≥ 3 times the dielectric thickness for a maximum distance of 500 mils from the pin.



2. Board designers should ideally target $100\Omega \pm 15\%$. If it's not feasible (due to board stack-up) it is recommended that board designers use a $95\Omega \pm 10\%$ target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95Ω . The $\pm 10\%$ tolerance is provided to allow for board manufacturing process variations and not lower target impedances. The minimum value of impedance cannot be lower than 85Ω .
3. Simulation shows 80Ω differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90Ω .
4. Stripline is NOT recommended due to thinner more resistive signal layers.
5. Use a minimum of 21 mil (0.533 mm) pair-to-pair spacing for board designs that use the CRB design stack-up. Using dielectrics that are thicker than the CRB stack-up might require larger pair-to-pair spacing.
6. Mobile designs without LAN switch can range up to ~8 inches. Refer to Table 14-8 for trace length information

Table 14-8. Maximum Trace Lengths Based on Trace Geometry and Board Stack-Up

Dielectric Thickness (mils)	Dielectric Constant (DK) at 1 MHz	Width / Space / Width (mils)	Pair-to-Pair Space (mils)	Nominal Impedance (Ohms)	Impedance Tolerance ($\pm\%$)	Maximum Trace Length (inches) ¹
2.7	4.05	4/10/4	19	95^2	17^2	3.5
2.7	4.05	4/10/4	19	95^2	15^2	4
2.7	4.05	4/10/4	19	95	10	5
3.3	4.1	4.2/9/4.2	23	100^2	17^2	4
3.3	4.1	4.2/9/4.2	23	100	15	4.6
3.3	4.1	4.2/9/4.2	23	100	10	6
4	4.2	5/9/5	28	100^2	17^2	4.5
4	4.2	5/9/5	28	100	15	5.3
4	4.2	5/9/5	28	100	10	7

Notes:

1. Longer MDI trace lengths may be achievable, but may make it more difficult to achieve IEEE conformance. Simulations have shown deviations are possible if traces are kept short. Longer traces are possible; use cost considerations and stack-up tolerance for differential pairs to determine length requirements.
2. Deviations from 100Ω nominal and/or tolerances greater than 15% decrease the maximum length for IEEE conformance.

Note:


Use the MDI Differential Trace Calculator to determine the maximum MDI trace length for your trace geometry and board stack-up. Contact your Intel Field Service Representative for access.

The following factors can limit the maximum MDI differential trace lengths for IEEE conformance:

- Dielectric thickness
- Dielectric constant
- Nominal differential trace impedance
- Trace impedance tolerance
- Copper trace losses
- Additional devices, such as switches, in the MDI path may impact IEEE conformance.

Board geometry should also be factored in when setting trace length.

Figure 14-13.MDI Trace Geometry



14.10 Trace Length and Symmetry

The differential traces should be equal in total length to within 10 mils (0.254 mm) per segment within each pair and as symmetrical as possible. Asymmetrical and unequal length traces in the differential pairs contribute to common mode noise. If a choice has to be made between matching lengths and fixing symmetry, more emphasis should be placed on fixing symmetry. Common mode noise can degrade the receive circuit's performance and contribute to radiated emissions.


The intra-pair length matching on the pairs must be within 10 mils on a segment by segment basis. An MDI segment is defined as any trace within the same layer. For example, transitioning from one layer to another through a via is considered as two separate MDI segments.

The end-to-end total trace lengths within each differential pair must match as shown in the figure titled MDI Trace Geometry. The end to end trace length is defined as the total MDI length from one component to another regardless of layer transitions.

The pair to pair length matching is not as critical as the intra-pair length matching but it should be within 2 inches.

When using Microstrip, the MDI traces should be at least 7x the thinnest adjacent dielectric away from the edge of an adjacent reference plane. When using stripline, the MDI traces should be at least 6x the thinnest adjacent dielectric away from the edge of an adjacent reference plane.

Figure 14-14. MDI Differential Trace Geometry



14.11 Impedance Discontinuities

Impedance discontinuities cause unwanted signal reflections. Vias (signal through holes) and other transmission line irregularities should be minimized. If vias must be used, a reasonable budget is four or less per differential trace. Unused pads and stub traces should also be avoided.

14.12 Reducing Circuit Inductance

Traces should be routed over a continuous reference plane with no interruptions. If there are vacant areas on a reference or power plane, the signal conductors should not cross the vacant area. This causes impedance mismatches and associated radiated noise levels.

14.13 Signal Isolation

To maintain best signal integrity, keep digital signals far away from the analog traces. Also, keep the MDI traces away from the edge of an adjacent reference plane by a distance that is at least 7x the thickness of the thinnest adjacent dielectric layer (7x when using Microstrip; 6x when using stripline). If digital signals on other board layers cannot be separated by a ground plane, they should be routed perpendicular to the differential pairs. If there is another LAN controller on the board, the differential pairs from that circuit must be kept away.

Other rules to follow for signal isolation include:



- Separate and group signals by function on separate layers if possible. If possible, maintain at least a gap of 30 mils between all differential pairs (Ethernet) and other nets, but group associated differential pairs together.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.
- Isolate I/O signals from high-speed signals to minimize crosstalk, which can increase EMI emission and susceptibility to EMI from other signals.
- Avoid routing high-speed LAN traces near other high-frequency signals associated with a video controller, cache controller, processor, switching power supplies, or other similar devices.

14.14 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and power inputs bypassed to signal return. This will significantly reduce EMI radiation.

The following guidelines help reduce circuit inductance in both backplanes and motherboards:

- Route traces over a continuous plane with no interruptions. Do not route over a split power or ground plane. If there are vacant areas on a ground or power plane, avoid routing signals over the vacant area. This will increase inductance and EMI radiation levels.
- All ground vias should be connected to every ground plane; and every power via, to all power planes at equal potential. This helps reduce circuit inductance.
- Physically locate grounds between a signal path and its return. This will minimize the loop area.
- Split the ground plane beneath a magnetics module. The RJ-45 connector side of the transformer module should have chassis ground beneath it.

Caution: DO NOT do this, if the RJ-45 connector has integrated USB.

Note: All impedance-controlled signals should be routed in reference to a solid plane. If there are plane splits on a reference layer, and the signal traces cross those splits, stitching capacitors should be used within 40 mils of where the crossing occurs. See Figure 1-13.

If signals transition from one reference layer to another reference layer, stitching capacitors or connecting vias should be used based on the following:

If the transition is from power-referenced layer to a ground-referenced layer or from one voltage-power referenced layer to a different voltage-power referenced layer, stitching capacitors should be used within 40 mils of the transition.

If the transition is from one ground-referenced layer to another ground-referenced layer or is from a power-referenced layer to the same net power-referenced layer, connecting vias should be used within 40 mils of the transition.

Figure 14-15. Trace Transitioning Layers and crossing Plane Splits




Figure 14-16. Via Connecting GND to GND





Figure 14-17. Stitching Capacitor between Vias Connecting GND to GND



14.15 Traces for Decoupling Capacitors

Traces between decoupling and I/O filter capacitors should be as short and wide as practical. Long and thin traces are more inductive and reduce the intended effect of decoupling capacitors. Also, for similar reasons, traces to I/O signals and signal terminations should be as short as possible. Vias to the decoupling capacitors should be sufficiently large in diameter to decrease series inductance. Refer to the Power Delivery section for the PHY in regards to actual placement requirements of the capacitors.

14.16 Ground Planes Under a Magnetics Module

The magnetics module chassis or output ground (secondary side of transformer) should be separated from the digital or input ground (primary side) by a physical separation of 100 mils minimum. Splitting the ground planes beneath the transformer minimizes noise coupling between the primary and secondary sides of the transformer and between the adjacent coils in the magnetics. This arrangement also improves the common mode choke functionality of magnetics module.

Caution: DO NOT do this if the RJ-45 connector has integrated USB.

Figure 14-17 illustrates the split plane layout for a discrete magnetics module. Capacitors are used to interconnect chassis ground and signal ground.

Figure 14-18 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.

Figure 14-18. Ideal Ground Split Implementation

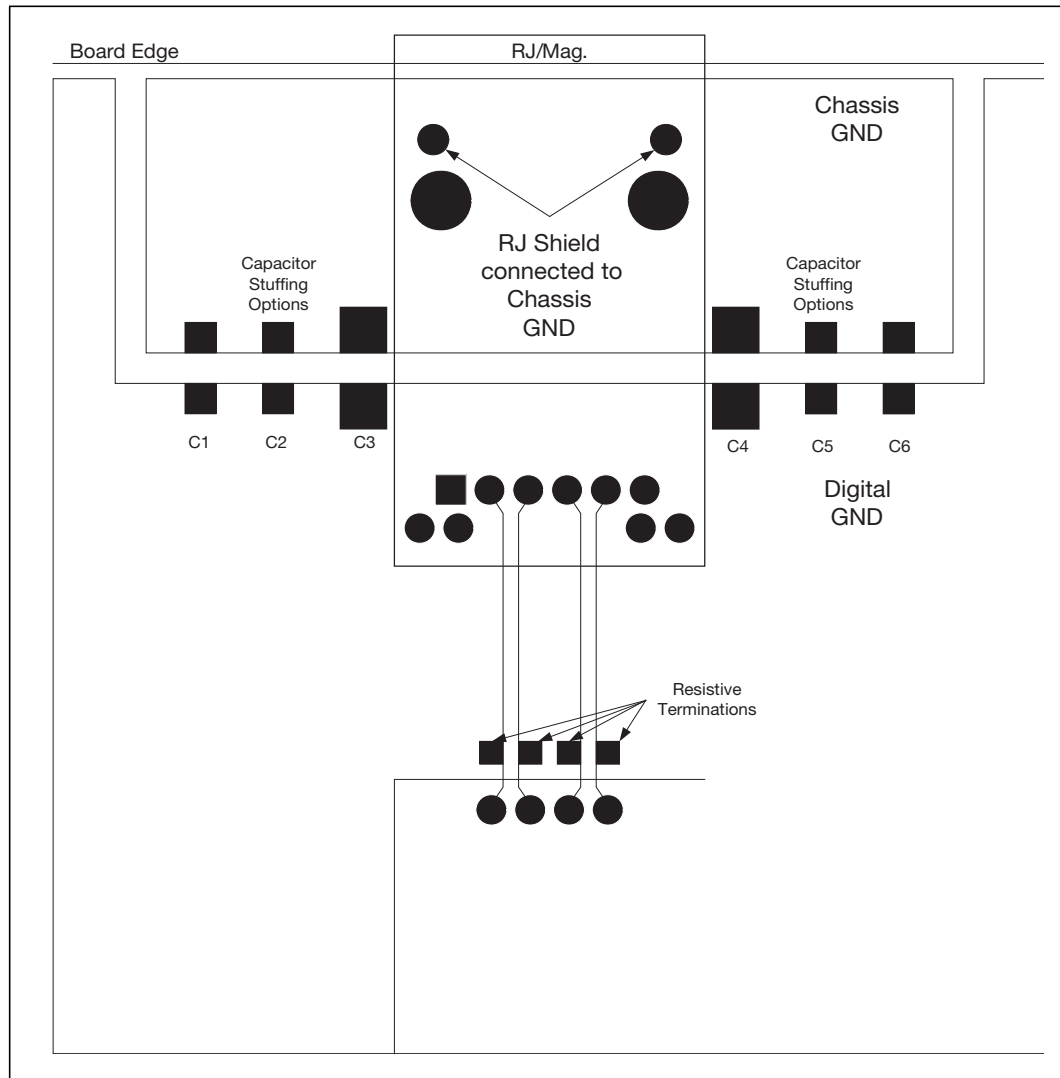


Table 14-9. Capacitor Stuffing Option Recommended Values


Capacitors	Value
C3, C4	4.7 μ F or 10 μ F
C1, C2, C5, C6	470 pF to 0.1 μ F

The placement of C1 through C6 may also differ for each board design (in other words, not all of the capacitors may need to be populated). Also, the capacitors may not be needed on both sides of the magnetics module.

Note: If using an integrated magnetics module without USB, provide a separate chassis ground “island” to ground around the RJ-45 connector. The split in the ground plane should be at least 20 mils wide.

Some integrated magnetics modules/RJ-45 connectors have recently incorporated the USB into the device. For this type of magnetics module, a chassis ground moat may not be feasible due to the digital ground required for the USB pins and their placement relative to the magnetics pins. Thus, a continuous digital ground without any moats or splits must be used. [Figure 14-19](#) provides an example of this.

Figure 14-19. Ground Layout with USB



14.17 Light Emitting Diodes

The device has three high-current outputs to directly drive LEDs for link, activity and speed indication. Since LEDs are likely to be integral to a magnetics module, take care to route the LED traces away from potential sources of EMI noise. In some cases, it may be desirable to attach filter capacitors.

LAN LED traces should be placed at least 6x (side by side separation) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

LAN LED traces should be placed at least 7x (broadside coupling) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

14.18 Frequency Control Device Design Considerations

This section provides information regarding frequency control devices, including crystals and oscillators, for use with all Intel Ethernet controllers. Several suitable frequency control devices are available; none of which present any unusual challenges in selection. The concepts documented within this section are applicable to other data communication circuits, including the PHY.

The PHY contains amplifiers that form the basis for feedback oscillators when they are used with the specific external components. These oscillator circuits, which are both economical and reliable, are described in more detail in [Section 14.22](#).



The chosen frequency control device vendor should be consulted early in the design cycle. Crystal and oscillator manufacturers familiar with networking equipment clock requirements may provide assistance in selecting an optimum, low-cost solution.

Several types of third-party frequency reference components are currently available. Descriptions of each type follow in subsequent sections. They are also listed in order of preference.

14.19 Crystals and Oscillators

Clock sources should not be placed near I/O ports or board edges. Radiation from these devices may be coupled onto the I/O ports or out of the system chassis. Crystals should also be kept away from the Ethernet magnetics module to prevent interference.

Crystal and load capacitors should be placed on the printed circuit boards as close to the PHY as possible, which is within 1.0 inch. Traces from XTAL_IN (X1) and XTAL_OUT (X2) should be routed as symmetrically as possible. Do not route X1 and X2 as a differential trace. Doing so increases jitter and degrades LAN performance.

- The crystal trace lengths should be less than 1 inch.
- The crystal load capacitors should be placed less than 1" from the crystal.
- The clock lines must be at least 5 times the height of the thinnest adjacent dielectric layer away from other from other digital traces (especially reset signals), I/O ports, board edge, transformers and differential pairs.
- The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition.
- The clock lines should not cross or run in parallel (within 3x the dielectric thickness of the closest dielectric layer) with any trace (100Mhz signal or higher) on an adjacent layer.

14.20 Quartz Crystal

Quartz crystals are generally considered to be the mainstay of frequency control components due to their low cost and ease of implementation. They are available from numerous vendors in many package types and with various specification options.

14.21 Fixed Crystal Oscillator

A packaged fixed crystal oscillator comprises of an inverter, a quartz crystal, and passive components conveniently packaged together. The device renders a strong, consistent square wave output. Oscillators used with microprocessors are supplied in many configurations and tolerances.

Crystal oscillators should be restricted for use in special situations, such as shared clocking among devices or multiple controllers. Since clock routing can be difficult to accomplish, it is preferable to provide a separate crystal for each device.

Note: Contact your Intel Field Service Representative to obtain the most current device documentation prior to implementing this solution.



14.22 Crystal Selection Parameters

All crystals used with Intel Ethernet controllers are described as “AT-cut,” which refers to the angle at which the unit is sliced with respect to the long axis of the quartz stone.

The following table lists crystals which have been used successfully in past designs. (No particular product is recommended.)

Table 14-10. Crystal Manufacturers and Part Numbers

Manufacturer	Part No.	Note
TXC*	9C25000008	HC-49S SMD
NDK*	41CD25.0F1303018	HC-49S SMD
TXC*	7V25020004	Small SMD 3.2mmx2.5mmTop of Form ----- Bottom of Form
KDS*	DSX321G, 1B/C/N/H225000CC0M	Small SMD 3.2mmx2.5mm
River*	FCX-04-25MJ90141	Small SMD 3.2mmx2.5mm

The datasheet for the PHY lists the crystal electrical parameters and provides suggested values for typical designs. Designers should refer to criteria outlined in their respective PHY datasheet. The parameters are described in the following subsections.

14.23 Vibrational Mode

Crystals in the frequency range referenced above are available in both fundamental and third overtone. Unless there is a special need for third overtone, fundamental mode crystals should be used.

14.24 Nominal Frequency

Intel Ethernet controllers use a crystal frequency of 25.000 MHz. The 25 MHz input is used to generate a 125 MHz transmit clock for 100BASE-TX and 1000BASE-TX operation, and 10 MHz and 20-MHz transmit clocks, for 10BASE-T operation.

14.25 Frequency Tolerance

The frequency tolerance for an Ethernet Platform LAN Connect device is dictated by the IEEE 802.3 specification as ± 50 parts per million (ppm). This measurement is referenced to a standard temperature of 25 °C. Intel recommends a frequency tolerance of ± 30 ppm to ensure for any frequency variance contributed by the PCB.

14.26 Temperature Stability and Environmental Requirements

Temperature stability is a standard measure of how the oscillation frequency varies over the full operational temperature range (and beyond). Several optional temperature ranges are currently available, including -40 °C to +85 °C for industrial



environments. Some vendors separate operating temperatures from temperature stability. Manufacturers may also list temperature stability as 50 ppm in their data sheets.

Note: Crystals also carry other specifications for storage temperature, shock resistance, and reflow solder conditions. Crystal vendors should be consulted early in the design cycle to discuss its application and environmental requirements.


14.27 Calibration Mode

The terms “series-resonant” and “parallel-resonant” are often used to describe crystal oscillator circuits. Specifying parallel mode is critical to determining how the crystal frequency is calibrated at the factory.

A crystal specified and tested as series resonant oscillates without problem in a parallel-resonant circuit, but the frequency is higher than nominal by several hundred parts per million. The purpose of adding load capacitors to a crystal oscillator circuit is to establish resonance at a frequency higher than the crystal’s inherent series resonant frequency.

Figure 14-20 illustrates a simplified schematic of the internal oscillator circuit. Pin X1 and X2 refers to XTAL_IN and XTAL_OUT in the Ethernet device, respectively. The crystal and the capacitors form a feedback element for the internal inverting amplifier. This combination is called parallel-resonant, because it has positive reactance at the selected frequency. In other words, the crystal behaves like an inductor in a parallel LC circuit. Oscillators with piezoelectric feedback elements are also known as “Pierce” oscillators.

Figure 14-20. Thermal Oscillator Circuit



14.28 Load Capacitance

The formula for crystal load capacitance is as follows:



$$C_L = \frac{(C1 \cdot C2)}{(C1 + C2)} + C_{stray}$$

where $C1 = C2 = 10$ pF (as suggested in most Intel reference designs)
and C_{stray} = allowance for additional capacitance in pads, traces and the chip carrier within the Ethernet device package and C_{damp} .

14.29 Shunt Capacitance

The shunt capacitance parameter is relatively unimportant compared to load capacitance. Shunt capacitance represents the effect of the crystal's mechanical holder and contacts. The shunt capacitance should be a maximum of 6 pF.

14.30 Equivalent Series Resistance

Equivalent Series Resistance (ESR) is the real component of the crystal's impedance at the calibration frequency, which the inverting amplifier's loop gain must overcome. ESR varies inversely with frequency for a given crystal family. The lower the ESR, the faster the crystal starts up. Crystals with an ESR value of 50 Ω or better should be used.

14.31 Drive Level

Drive level refers to power dissipation in use. The allowable drive level for a Surface Mounted Technology (SMT) crystal is less than its through-hole counterpart. This is due to the fact that surface mount crystals are typically made from narrow, rectangular AT strips, rather than circular AT quartz blanks.

When selecting a crystal, board designers must ensure that the crystal specification meets at least the drive level specified. For example, if the crystal drive level specification states that the drive level is 200 μ W maximum, then the crystal drive level must be at least 200 μ W. So, a 500 μ W crystal is sufficient, but a 100 μ W crystal is not.

14.32 Aging

Aging is a permanent change in frequency (and resistance) occurring over time. This parameter is most important in its first year because new crystals age faster than old crystals. Crystals with a maximum value of ± 5 ppm per year aging should be used.

14.33 Reference Crystal

The normal tolerances of the discrete crystal components can contribute to small frequency offsets with respect to the target center frequency. To minimize the risk of tolerance-caused frequency offsets causing a small percentage of production line units to be outside of the acceptable frequency range, it is important to account for those shifts while empirically determining the proper values for the discrete loading capacitors, C1 and C2.



Even with a perfect support circuit, most crystals will oscillate slightly higher or lower than the exact center of the target frequency. Therefore, frequency measurements, which determine the correct value for C1 and C2, should be performed with an ideal reference crystal. When the capacitive load is exactly equal to the crystal's load rating, an ideal reference crystal will be perfectly centered at the desired target frequency.

14.33.1 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal:

- If a Saunders and Associates (S&A) crystal network analyzer is available, then discrete crystal components can be tested until one is found with zero or nearly zero ppm deviation (with the appropriate capacitive load). A crystal with zero or near zero ppm deviation will be a good reference crystal to use in subsequent frequency tests to determine the best values for C1 and C2.
- If a crystal analyzer is not available, then the selection of a reference crystal can be done by measuring a statistically valid sample population of crystals, which has units from multiple lots and approved vendors. The crystal, which has an oscillation frequency closest to the center of the distribution, should be the reference crystal used during testing to determine the best values for C1 and C2.
- It may also be possible to ask the approved crystal vendors or manufacturers to provide a reference crystal with zero or nearly zero deviation from the specified frequency when it has the specified C_{Load} capacitance.

When choosing a crystal, designers must keep in mind that to comply with IEEE specifications for 10/100 Mbps operation and 10/100/1000 Mbps operation if applicable, the transmitter reference frequency must be precise within ± 50 ppm. Intel recommends customers use a transmitter reference frequency that is accurate to within ± 30 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance.

14.33.2 Circuit Board

Since the dielectric layers of the circuit board are allowed some reasonable variation in thickness, the stray capacitance from the printed board (to the crystal circuit) will also vary. If the thickness tolerance for the outer layers of dielectric are controlled within $\pm 15\%$ of nominal, then the circuit board should not cause more than ± 2 pF variation to the stray capacitance at the crystal. When tuning crystal frequency, it is recommended that at least three circuit boards are tested for frequency. These boards should be from different production lots of bare circuit boards.

Alternatively, a larger sample population of circuit boards can be used. A larger population will increase the probability of obtaining the full range of possible variations in dielectric thickness and the full range of variation in stray capacitance.

Next, the exact same crystal and discrete load capacitors (C1 and C2) must be soldered onto each board, and the LAN reference frequency should be measured on each circuit board.

The circuit board, which has a LAN reference frequency closest to the center of the frequency distribution, should be used while performing the frequency measurements to select the appropriate value for C1 and C2.

14.33.3 Temperature Changes



Temperature changes can cause the crystal frequency to shift. Therefore, frequency measurements should be done in the final system chassis across the system's rated operating temperature range.

14.34 Oscillator Support

The PHY clock input circuit is optimized for use with an external crystal. However, an oscillator can also be used in place of the crystal with the proper design considerations (refer to the PHY Datasheet for detailed clock oscillator specifications):

- The clock oscillator has an internal voltage regulator to isolate it from the external noise of other circuits to minimize jitter. If an external clock is used, this imposes a maximum input clock amplitude. For example, if a 3.3V DC oscillator is used, it's output signal should be attenuated to a maximum value with a resistive divider circuit.
- The input capacitance introduced by the PHY (approximately 11 to 13 pF) is greater than the capacitance specified by a typical oscillator (approximately 15 pF).
- The input clock jitter from the oscillator can impact the PHY clock and its performance.

Note: The power consumption of additional circuitry equals about 1.5 mW.

The following table lists oscillators that can be used with the PHY. Note that no particular oscillator is recommended):

Table 14-11.Oscillator Manufacturers and Part Numbers

Manufacturer	Part No.
Kyocera*	K30-3C0-SE-25.0000M
MtronPTI*	M214TCN25.0000MHz
TXC*	7C25000230 7X25080001

14.35 Oscillator Placement and Layout Recommendations

Oscillator clock sources should not be placed near I/O ports or board edges. Radiation from these devices can be coupled into the I/O ports and radiate beyond the system chassis. Oscillators should also be kept away from the Ethernet magnetics module to prevent interference.

The oscillator must have it's own decoupling capacitors and they must be placed within 0.25 inches. If a power trace is used (not power plane), the trace from the capacitor to the oscillator must not exceed 0.25 inches in length. The decoupling capacitors help to improve the oscillator stability. The oscillator clock trace should be less than two inches from the PHY. If it is greater than 2 inches, then verify the signal quality, jitter, and clock frequency measurements at the PHY.



The clock lines should also target $50\Omega \pm 15\%$ and should have 10Ω series back termination placed close to the series oscillator. To help reduce EMI, the clock lines must be a distance of at least five times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, the board edge, transformers and differential pairs.

The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition. The clock lines should not cross or run in parallel with any trace (100 MHz signal or higher) on an adjacent layer.

There should be a ferrite bead within 250 mils of the oscillator power pin and there must be a 1 uF or greater capacitor within 250 mils of the oscillator, connected to the power trace between the oscillator input and ferrite bead. With a ferrite bead on the power trace for the oscillator, there should be a power pour (or fat trace) to supply power to the oscillator.

Note: For latest PHY schematic connection recommendations, refer to the *Intel® Ethernet Controller I217 Reference Schematic* and the *Intel® Ethernet Controller I217 Schematic and Layout Checklist*. For access to documentation, contact your Intel Field Service Representative.

14.36 Troubleshooting Common Physical Layout Issues

The following is a list of common physical layer design and layout mistakes in LAN on Motherboard (LOM) designs.

1. Lack of symmetry between the two traces within a differential pair. Asymmetry can create common-mode noise and distort the waveforms. For each component and via that one trace encounters, the other trace should encounter the same component or a via at the same distance from the Ethernet silicon.
2. Unequal length of the two traces within a differential pair. Inequalities create common-mode noise and will distort the transmit or receive waveforms.
3. Excessive distance between the Ethernet silicon and the magnetics. Long traces on FR4 fiberglass epoxy substrate will attenuate the analog signals. In addition, any impedance mismatch in the traces will be aggravated if they are longer than the four-inch guideline.
4. Routing any other trace parallel to and close to one of the differential traces. Crosstalk getting onto the receive channel will cause degraded long cable BER. Crosstalk getting onto the transmit channel can cause excessive EMI emissions and can cause poor transmit BER on long cables. At a minimum, for stripline other signals should be kept at least 6x the height of the thinnest adjacent dielectric layer. For microstrip it is 7x. The only possible exceptions are in the vicinities where the traces enter or exit the magnetics, the RJ-45 connector, and the Ethernet silicon.
5. Using a low-quality magnetics module.
6. Reusing an out-of-date physical layer schematic in a Ethernet silicon design. The terminations and decoupling can be different from one PHY to another.
7. Incorrect differential trace impedances. It is important to have about a $100\text{-}\Omega$ impedance between the two traces within a differential pair. This becomes even more important as the differential traces become longer. To calculate differential impedance, many impedance calculators only multiply the single-ended impedance by two. This does not take into account edge-to-edge capacitive coupling between



the two traces. When the two traces within a differential pair are kept close to each other, the edge coupling can lower the effective differential impedance by 5 Ω to 20 Ω. Short traces will have fewer problems if the differential impedance is slightly off target.

14.37 Power Delivery

The Intel® Ethernet Controller I217 requires a 3.3V power rail and a Core power rail. The internal 3.3V power rail is brought out for decoupling. Figure 14-21 shows a typical power delivery configuration that can be implemented. However, power delivery can be customized based on a specific OEM. In general, planes should be used to deliver 3.3 Vdc and the Core voltage. Not using planes can cause resistive voltage drop and/or inductive voltage drop (due to transient or static currents). Some of the symptoms of these voltage drops can include higher EMI, radiated immunity, radiated emissions, IEEE conformance issues, and register corruption.

Decoupling capacitors (0.1 uF and smaller) should be placed within 250 mils of the LAN device. They also should be distributed around the PHY and some should be in close proximity to the power pins.

The bulk capacitors (1.0 uF or greater) should be placed within 1 inch if using a trace (50 mils wide or wider) or within 1.5 inches if using a plane.

The Core power rail for the Intel® Ethernet Controller I217 is generated using the integrated SVR (iSVR). Using the internal SVR to generate the Core power, the inductor must be placed within 0.5" of the input pin to the PHY and connected with a trace wider than or equal to 20 mil wide. (Please see the reference schematic for further details regarding the Core power rail.) Care should be taken to minimize any voltage drops to within 50 mV. For calculating the voltage drop through copper traces, refer to the Power Delivery Loss Calculator. Contact your Intel Field Service Representative for access.


The following list shows inductors that have been used successfully with Intel designs:

Table 14-12. Inductors and Manufacturers

Manufacturer	Part Number
muRata*	LQH32PN4R7NN0
muRata*	LQH32CN4R7M53
TDK*	FLF3215T-4R7M
TDK*	VLS252015ET‡
Taiyo Yuden*	NRS2012T4R7MGJ
TDK*	VLS2012MT-4R7M

‡ The height of this part may not be appropriate for some designs.

Figure 14-21. Intel® Ethernet Controller I217 Power Delivery Diagram




Note: For latest PHY schematic connection recommendations, refer to the Intel® Ethernet Controller I217 reference schematic. For access to documentation, contact your Intel Field Service Representative.

14.38 Intel® Ethernet Controller I217 Power Sequencing

The Intel® Ethernet Controller I217 does not require any power sequencing between the 3.3V and Core power rails if the Core voltage is generated from the Intel® Ethernet Controller I217's internal SVR. See the datasheet and reference schematic for details.



NOTE: *This page intentionally left blank.*



15 Checklists

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I217 Design and Board Layout Checklists.



NOTE: *This page intentionally left blank.*

§ §




16 Models

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I217 models.



NOTE: *This page intentionally left blank.*

§ §



17 Reference Schematics

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I217 reference schematics.



NOTE: *This page intentionally left blank.*

§ §

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Intel:](#)

[WGI217V S LJWH](#) [WGI217V S LJWG](#)