

Intel® Ethernet Connection I218

Datasheet v2.6

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
 - Up to 16 programmable filters
- 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
 - Ultra Low Power at cable disconnect (<1 mW) enables platform support for connected standby
 - 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



By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation
P.O. Box 5937
Denver, CO 80217-9808

or call in North America 1-800-548-4725, Europe 44-0-1793-431-155, France 44-0-1793-421-777, Germany 44-0-1793-421-333, other Countries 708-296-9333.

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 © 2009-2014, Intel Corporation. All Rights Reserved.



| Revision | Date | Description |
|----------|----------------|---|
| 0.95 | June 2012 | Initial release. |
| 0.96 | June 2012 | Updated Features table; updated Product Ordering Codes chart; removed external voltage regulator information; updated Transmit Flow information; updated Wake Up filter information; updated the number of flexible filters supported; updated Power table; updated Power Delivery diagram; added programming information regarding low power mode. |
| 0.97 | September 2012 | Corrected Table 1 Product Ordering Codes, Product Code Names |
| 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. Updated table 18. |
| 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). Revised section 6.4 (The Intel® Ethernet Connection I218 Ultra Low Power (ULP); LED tri-state mode bullet). |
| 2.3 | October 2013 | Revised section 7.5 (Network Proxy Functionality; added a note to section 7.5.1). |
| 2.4 | January 2013 | Revised Table 1 (Product Ordering Codes). |
| 2.5 | June 2014 | Updated section 3.1.7 (changed CTRL1P0 to CTRL0P9). |
| 2.6 | October 2014 | Changed the number of flexible filter references from 8 to 16. |

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

| | | |
|------------|---|----|
| 1.0 | Introduction | 1 |
| 1.1 | Overview | 1 |
| 1.2 | Main Flows | 2 |
| 1.3 | References | 2 |
| 1.4 | Product SKUs and Codes | 3 |
| 2.0 | 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 | 22 |
| 2.5 | SMBus/PCIe Interconnects | 22 |
| 3.0 | Pin Interface | 25 |
| 3.1 | Pin Assignment | 25 |
| 3.1.1 | Signal Type Definitions | 25 |
| 3.1.2 | PCIe Interface Pins (8) | 25 |
| 3.1.3 | SMBus Interface Pins (2) | 26 |
| 3.1.4 | Miscellaneous Pins (3) | 26 |
| 3.1.5 | PHY Pins (14) | 26 |
| 3.1.6 | Testability Pins (5) | 27 |
| 3.1.7 | Power and Power Control Pins (16) | 28 |
| 3.2 | Pinout | 28 |
| 4.0 | Package | 29 |
| 4.1 | Package Type and Mechanical | 29 |
| 4.2 | Package Electrical and Thermal Characteristics | 30 |
| 4.3 | Power and Ground Requirements | 30 |
| 4.4 | Ball Mapping | 31 |
| 5.0 | Initialization | 33 |
| 5.1 | Power Up | 33 |
| 5.2 | Reset Operation | 34 |
| 5.3 | Timing Parameters | 35 |
| 5.3.1 | Timing Requirements | 35 |
| 5.3.2 | Timing Guarantees | 35 |
| 6.0 | Power Management and Delivery | 37 |
| 6.1 | Power Information | 37 |
| 6.2 | Power Delivery | 37 |
| 6.2.1 | Core Vdc Supply | 38 |
| 6.3 | Power Management | 38 |
| 6.3.1 | Global Power States | 38 |
| 6.4 | The Intel® Ethernet Connection I218 Ultra Low Power (ULP) | 43 |
| 6.5 | Off Board LAN Connected Device (OBLCD) | 44 |
| 6.5.1 | Indicating/Sensing OBLCD disconnect event | 44 |
| 6.5.2 | Indicating/Sensing OBLCD connection | 44 |
| 7.0 | Device Functionality | 45 |



- 7.1 Tx Flow 45
- 7.2 Rx Flow 45
- 7.3 Flow Control..... 46
 - 7.3.1 MAC Control Frames and Reception of Flow Control Packets..... 46
 - 7.3.2 Transmitting PAUSE Frames 48
- 7.4 Wake Up 49
 - 7.4.1 Host Wake Up 50
 - 7.4.2 PCI-M Power Management (PCI-m) 57
- 7.5 Network Proxy Functionality 64
 - 7.5.1 Introduction 64
 - 7.5.2 Network Proxy activation 65
 - 7.5.3 IPv4 Proxy—ARP..... 65
 - 7.5.4 IPv6 Proxy—Neighbor Discovery 67
 - 7.5.5 Multicast Listener Discovery Support 70
- 7.6 Loopback..... 73
- 8.0 Programmers’ Visible State..... 75**
 - 8.1 Terminology..... 75
 - 8.2 MDIO Access..... 76
 - 8.3 Addressing..... 76
 - 8.4 Registers Byte Ordering 77
 - 8.5 Address Map 78
 - 8.6 PHY Registers (Page 0)..... 81
 - 8.6.1 Loopback Mode Settings 89
 - 8.7 Port Control Registers (Page 769) 102
 - 8.8 Statistics Registers..... 105
 - 8.9 PCIe Registers..... 106
 - 8.10 General Registers 108
 - 8.10.1 Interrupts 114
 - 8.11 Wake Up Registers..... 115
 - 8.11.1 Accessing Wake Up Registers Using MDIC 115
 - 8.11.2 Host Wake Up Control Status Register Description..... 116
 - 8.11.3 ME Wake Up Control Status Register Description 129
 - 8.11.4 Proxy Controller uCode 138
 - 8.12 LPI MMD PHY Registers..... 139
 - 8.12.1 LPI MMD PHY Registers..... 139
- 9.0 Non-Volatile Memory (NVM) 141**
 - 9.1 Introduction 141
 - 9.2 NVM Programming Procedure Overview 141
 - 9.3 LAN NVM Format and Contents 143
 - 9.3.1 Hardware Accessed Words 144
 - 9.3.2 Software Accessed Words 151
 - 9.3.3 Basic Configuration Software Words 155
- 10.0 Electrical and Timing Specifications 157**
 - 10.1 Introduction 157
 - 10.2 Operating Conditions..... 157
 - 10.2.1 Absolute Maximum Ratings 157
 - 10.2.2 Recommended Operating Conditions 157
 - 10.3 Power Delivery 158
 - 10.3.1 Voltage Regulator Power Supply Specifications..... 158
 - 10.3.2 SVR Specification (Internal) 158
 - 10.3.3 Power On/Off Sequence 158
 - 10.3.4 Power Delivery Schematic Drawing..... 159
 - 10.4 I/O DC Parameter 160



| | | |
|-------------|---|------------|
| 10.4.1 | 3.3 Vdc I/O (Open Drain) | 160 |
| 10.4.2 | 3.3 Vdc I/O | 160 |
| 10.4.3 | 3.3 Vdc I/O Pull-Up & Pull-Down Information | 161 |
| 10.4.4 | Input Buffer Only | 161 |
| 10.4.5 | Input Buffer Vdc I/O Pull-up & Pull-down | 161 |
| 10.4.6 | PCIe DC/AC Specifications | 161 |
| 10.5 | Discrete/Integrated Magnetics Specifications | 165 |
| 10.6 | Mechanical | 165 |
| 10.7 | Oscillator/Crystal Specifications | 165 |
| 11.0 | Time Synch (IEEE1588 and 802.1AS) | 167 |
| 11.1 | Overview | 167 |
| 11.1.1 | Flow and HW/SW responsibilities | 167 |
| 11.1.2 | HW Time sync elements | 169 |
| 11.1.3 | PTP packet structure | 171 |
| 12.0 | Mobile Design Considerations and Guidelines | 175 |
| 12.1 | Intel® Ethernet Connection I218 Overview | 176 |
| 12.1.1 | Intel® Ethernet Connection I218 Interconnects | 176 |
| 12.1.2 | PCIe-Based Interface | 177 |
| 12.1.3 | SMBus Interface | 177 |
| 12.1.4 | PCIe and SMBus Modes | 178 |
| 12.1.5 | Transitions between PCIe and SMBus Interfaces | 178 |
| 12.2 | Platform LAN Design Guidelines | 179 |
| 12.2.1 | General Design Considerations for the Intel® Ethernet Connection I218 | 179 |
| 12.2.2 | NVM for PHY Implementations | 181 |
| 12.2.3 | LED | 181 |
| 12.2.4 | Exposed Pad* (e-Pad) Design and SMT Assembly Guide | 182 |
| 12.3 | PCH—SMBus/PCIe LOM Design Guidelines | 187 |
| 12.4 | SMBus Design Considerations | 188 |
| 12.5 | General Layout Guidelines | 189 |
| 12.6 | Layout Considerations | 189 |
| 12.7 | Guidelines for Component Placement | 189 |
| 12.7.1 | PHY Placement Recommendations | 190 |
| 12.8 | MDI Differential-Pair Trace Routing for LAN Design | 191 |
| 12.9 | Signal Trace Geometry | 192 |
| 12.10 | Trace Length and Symmetry | 194 |
| 12.11 | Impedance Discontinuities | 195 |
| 12.12 | Reducing Circuit Inductance | 195 |
| 12.13 | Signal Isolation | 195 |
| 12.14 | Power and Ground Planes | 196 |
| 12.15 | Traces for Decoupling Capacitors | 198 |
| 12.16 | Ground Planes under a Magnetics Module | 198 |
| 12.17 | Light Emitting Diodes | 200 |
| 12.18 | Considerations for Layout | 200 |
| 12.19 | Frequency Control Device Design Considerations | 201 |
| 12.20 | Crystals and Oscillators | 201 |
| 12.21 | Quartz Crystal | 201 |
| 12.22 | Fixed Crystal Oscillator | 202 |
| 12.23 | Crystal Selection Parameters | 202 |
| 12.24 | Vibrational Mode | 202 |
| 12.25 | Nominal Frequency | 203 |
| 12.26 | Frequency Tolerance | 203 |
| 12.27 | Temperature Stability and Environmental Requirements | 203 |
| 12.28 | Calibration Mode | 203 |
| 12.29 | Load Capacitance | 204 |



- 12.30 Shunt Capacitance204
- 12.31 Equivalent Series Resistance204
- 12.32 Drive Level205
- 12.33 Aging205
- 12.34 Reference Crystal205
 - 12.34.1 Reference Crystal Selection205
 - 12.34.2 Circuit Board206
 - 12.34.3 Temperature Changes206
- 12.35 Oscillator Support206
- 12.36 Oscillator Placement and Layout Recommendations207
- 12.37 LAN Switch207
- 12.38 Troubleshooting Common Physical Layout Issues208
- 12.39 Power Delivery208
- 12.40 Intel® Ethernet Controller I218 Power Sequencing210
- 13.0 Non-Mobile Design Considerations and Guidelines211**
 - 13.1 PHY Overview.....212
 - 13.1.1 PHY Interconnects.....212
 - 13.1.2 PCIe-Based Interface212
 - 13.1.3 SMBus Interface213
 - 13.1.4 PCIe and SMBus Modes214
 - 13.1.5 Transitions between PCIe and SMBus Interfaces.....214
 - 13.2 Platform LAN Design Guidelines215
 - 13.2.1 General Design Considerations for PHYs215
 - 13.2.2 NVM for PHY Implementations217
 - 13.2.3 LAN Switch217
 - 13.2.4 LED218
 - 13.2.5 Exposed Pad* (e-Pad) Design and SMT Assembly Guide219
 - 13.3 PCH—SMBus/PCIe LOM Design Guidelines223
 - 13.4 SMBus Design Considerations224
 - 13.5 General Layout Guidelines.....225
 - 13.6 Layout Considerations225
 - 13.7 Guidelines for Component Placement.....225
 - 13.7.1 PHY Placement Recommendations226
 - 13.8 MDI Differential-Pair Trace Routing for LAN Design227
 - 13.9 Signal Trace Geometry228
 - 13.10 Trace Length and Symmetry230
 - 13.11 Impedance Discontinuities231
 - 13.12 Reducing Circuit Inductance231
 - 13.13 Signal Isolation231
 - 13.14 Power and Ground Planes232
 - 13.15 Traces for Decoupling Capacitors234
 - 13.16 Ground Planes under a Magnetics Module234
 - 13.17 Light Emitting Diodes236
 - 13.18 Considerations for Layout236
 - 13.19 Frequency Control Device Design Considerations.....237
 - 13.20 Crystals and Oscillators237
 - 13.21 Quartz Crystal237
 - 13.22 Fixed Crystal Oscillator238
 - 13.23 Crystal Selection Parameters238
 - 13.24 Vibrational Mode.....238
 - 13.25 Nominal Frequency238
 - 13.26 Frequency Tolerance239
 - 13.27 Temperature Stability and Environmental Requirements.....239
 - 13.28 Calibration Mode.....239



| | |
|--|------------|
| 13.29 Load Capacitance | 240 |
| 13.30 Shunt Capacitance..... | 240 |
| 13.31 Equivalent Series Resistance | 240 |
| 13.32 Drive Level | 240 |
| 13.33 Aging..... | 241 |
| 13.34 Reference Crystal | 241 |
| 13.34.1 Reference Crystal Selection | 241 |
| 13.34.2 Circuit Board | 242 |
| 13.34.3 Temperature Changes | 242 |
| 13.35 Oscillator Support..... | 242 |
| 13.36 Oscillator Placement and Layout Recommendations..... | 243 |
| 13.37 LAN Switch..... | 243 |
| 13.38 Troubleshooting Common Physical Layout Issues..... | 244 |
| 13.39 Power Delivery..... | 244 |
| 13.40 Power Sequencing | 246 |
| 14.0 Checklists | 247 |
| 15.0 Models..... | 249 |
| 16.0 Reference Schematics..... | 251 |



Note: This page intentionally left blank.



1.0 Introduction

1.1 Overview

The Intel® Ethernet Connection I218 is a single-port Gigabit Ethernet Physical Layer Transceiver (PHY). It connects to an integrated Media Access Controller (MAC) through a dedicated interconnect. The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 also supports the Energy Efficient Ethernet (EEE) 802.az specification.

- Note:*
- 1 Gb/s half-duplex operation is not supported.
 - Media Access Control Security (MACsec) encryption is not supported.

The Intel® Ethernet Connection I218 is packaged in a small footprint QFN package. Package size is 6 x 6 mm with a 0.4 mm lead pitch and a height of 0.85 mm, making it very attractive for small form-factor platforms.

The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 PCIe interface is not PCIe compliant. It operates at half the speed of the PCI Express* (PCIe*) Specification v1.1 (2.5 GT/s). In this datasheet, the term PCIe-based is interchangeable with PCIe. There are no design layout differences between normal PCIe and the Intel® Ethernet Connection I218's PCIe-based interface.

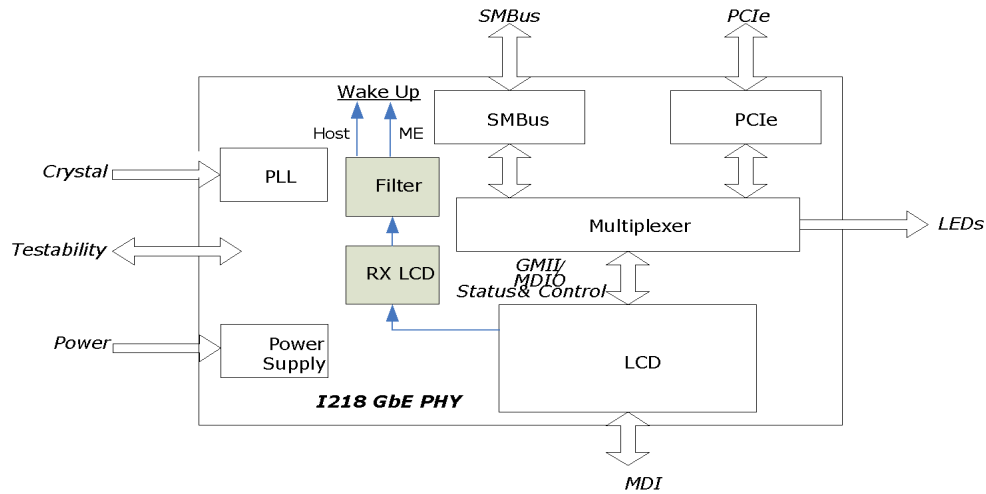


Figure 1. Intel® Ethernet Connection I218 Block Diagram

1.2 Main Flows

The main Intel® I218 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 Connection I218 functionality through two mechanisms:

- The Intel® Ethernet Connection I218 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 Codes

Table 1 lists the product SKUs and ordering codes for the Intel® Ethernet Connection I218. Refer to the *Intel® I218 GbE PHY Specification Update* for device ordering information.

Table 1. Product Ordering Codes

| 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) | Ultra Low Power | Power Optimizer | 802.1as | EEE - 802.3az | Ability to Inhibit a Team | MACSec (802.1ae) | Intel® SFP | iSCSI Boot | Standard Manageability | Intel® vPro™ |
|------------|-----------------------------------|-----------|---------------------------------------|-----------|-------------------------|-----------------|----------------------------|--------------------|----------------------------|--------------------------|---------------------------------|-------------------------|-----------------|-----------------|---------|---------------|---------------------------|------------------|------------|------------|------------------------|--------------|
| Gigabit | Ultrabook Mobile Desktop Embedded | Corporate | G218-LM Lynx Point-LP +218LM | 155A | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| | | Consumer | G218-V Lynx Point-LP +218V | 1559 | X | X | X | X | X | X | | X | X | X | X | X | X | | | | X | |
| Gigabit | Ultrabook Mobile Desktop Embedded | Corporate | G218-LM Lynx Point-LP-TH +218LM | 15A0 | X | X | X | X | X | X | X | X | | X | X | X | X | X | X | X | X | X |
| | | Consumer | G218-V Lynx Point-LP-TH +218LM | 15A1 | X | X | X | X | X | X | | X | | X | X | X | X | | | | X | |
| Gigabit | Ultrabook Mobile Desktop Embedded | Corporate | G218-LM Wildcat Point +218LM | 15A2 | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| | | Consumer | G218-V Wildcat Point +218V | 15A3 | X | X | X | X | X | X | | X | X | X | X | X | X | | | | X | |
| Gigabit | Entry Server Workstation | Corporate | G218-LM Lynx Point-LP-TH +218LM | 15A0 | 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 is not inclusive of all features for the product. The listed features are only those enabled or disabled for the associated SKU.
- For the list of supported Operating Systems, contact your Intel Field Service Representative



Note: This page intentionally left blank.



2.0 Interconnects

2.1 Introduction

The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

2.2 PCIe-Based

Note: The Intel® Ethernet Connection I218 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 10.4.4](#) 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 Intel® Ethernet Connection I218 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 Connection I218 and the integrated LAN Controller. Intel® Ethernet Connection I218 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 Connection I218 on every status change at the PHY side or when a time out defined in the In-Band Control register has expired. The Intel® Ethernet Connection I218 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. |



| Byte | Description |
|---------------|--|
| 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). |
| 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 Connection I218 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 Connection I218 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 Connection I218 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 Connection I218'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 Intel® Ethernet Connection I218's address is assigned using SMBus ARP protocol. The default SMBus address is 0xC8.

2.3.1.3 Bus Time Outs

The Intel® Ethernet Connection I218 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 Connection I218 does not abort the transaction. If the SMBus clock line is held low for 25 ms or longer, the Intel® Ethernet Connection I218 aborts the transaction.

As a slave, the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 that the Intel® Ethernet Connection I218 detects a time out on, 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 Intel® Ethernet Connection I218 has the capability of detecting a hung bus. If SMB_CLK or SMB_DATA are stuck low for more than 35 ms, the Intel® Ethernet Connection I218 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 ~50 μ s

2.3.1.6 SMBus ARP Functionality

The Intel® Ethernet Connection I218 supports the SMBus ARP protocol as defined in the SMBus 2.0 specification. The Intel® Ethernet Connection I218 is a persistent slave address device, meaning its SMBus address is valid after power up (constant 0xC8). The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 has a valid SMBus address.
- AR—Address Resolved: This flag is set when the Intel® Ethernet Connection I218 SMBus address is resolved (SMBus address was assigned by the SMBus ARP process).

Note: These flags are internal Intel® Ethernet Connection I218 flags and are not shown to external SMBus devices.

Since the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 an SMBus address by issuing an Assign Address command. The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 shows the SMBus ARP behavior of the Intel® Ethernet Connection I218.

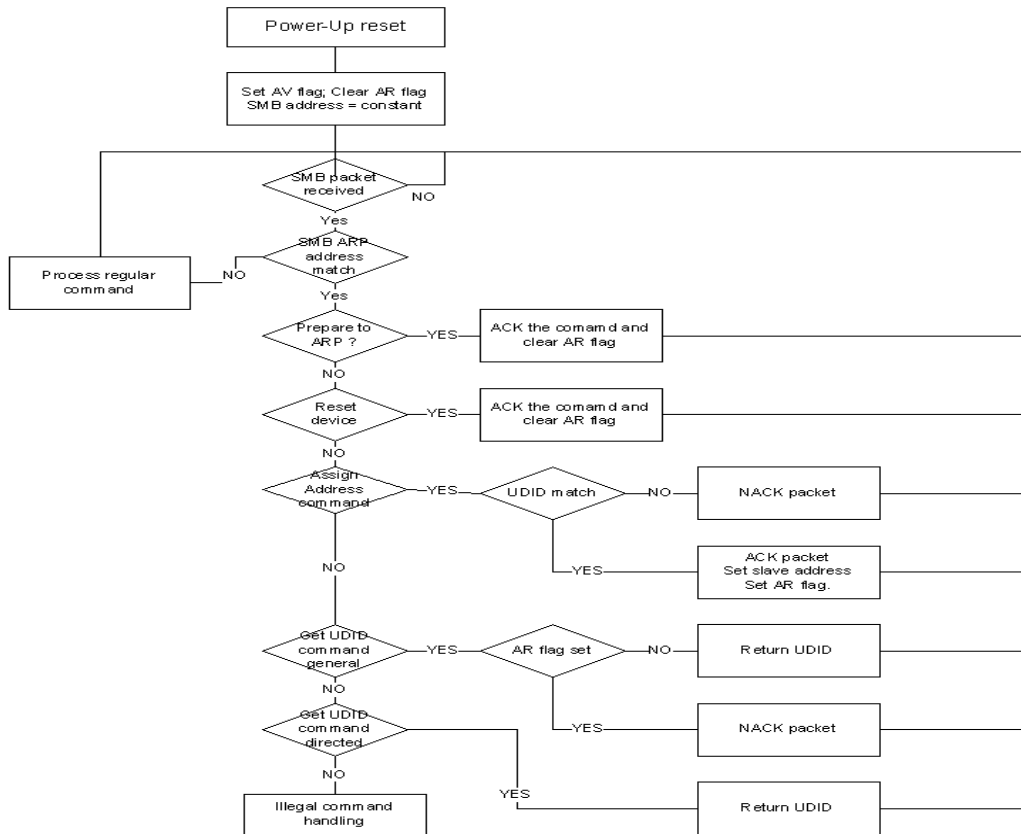


Figure 2. SMBus ARP Flow

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). Value for the LAN Connected device is constant: 0x10D4.
- Interface—Identifies 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 | |

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 Intel® Ethernet Connection I218. All traffic between the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 [Section 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 hardware to 1522 bytes. Packets larger than 1522 bytes are silently discarded.

2.3.2.2 Transmit Flow

The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218, it is transmitted over the network link.

The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218.

If the network link is down when the Intel® Ethernet Connection I218 is receiving SMBus fragments of the packet, it silently discards the packet.

The transmit SMBus transaction is described in [Section 2.3.3.1.1](#).

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 Intel® Ethernet Connection I218 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 2. 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 Intel® Ethernet Connection I218)

The following table lists the slave SMBus transactions supported by the Intel® Ethernet Connection I218.

Table 3. SMBus Slave Transactions

| Command | Transaction | Command | | Fragmentation | Section |
|------------------------|-------------|-----------------------------------|------------------------------|--|-----------|
| Transmit Packet | Block Write | First Middle Last Single | 0x84 0x04 0x44 0xC4 | Multiple Multiple Multiple Single | 2.3.3.1.1 |
| 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 Intel® Ethernet Connection I218 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 4. 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 5. Command 1st Byte

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

Table 6. 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 Intel® Ethernet Connection I218'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 7. 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 8. 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 9. 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 Intel® Ethernet Connection I218 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 10. 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 (Intel® Ethernet Connection I218 to the integrated LAN Controller)

In order to avoid starvation on the SMBus in the opposite direction, the Intel® Ethernet Connection I218 adds a byte-time idle between any two fragments it sends.

The following table lists the master SMBus transactions supported by Intel® Ethernet Connection I218.

Table 11. 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 Connection I218 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 Connection I218 on every status change at the PHY side or when a time out defined in the In-Band Control register has expired. The Intel® Ethernet Connection I218 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, Intel® Ethernet Connection I218 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 12. 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 13. 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 14. 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. |

2.3.4.1.3 PtM Configuration Acknowledge Command

The Configuration Acknowledge command is a single fragment command enabling the Intel® Ethernet Connection I218 to acknowledge the integrated LAN Controller access to Intel® Ethernet Connection I218 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 |

1. Bits 7:5.

**Table 15. 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 16. 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 Connection I218 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 Connection I218.

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 Connection I218.

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 Connection I218 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/PCH.

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 Intel® Ethernet Connection I218 can be connected to any x1 PCIe port in the chipset. The PCIe port that connects to the Intel® Ethernet Connection I218 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-Intel® Ethernet Connection I218 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 3](#).

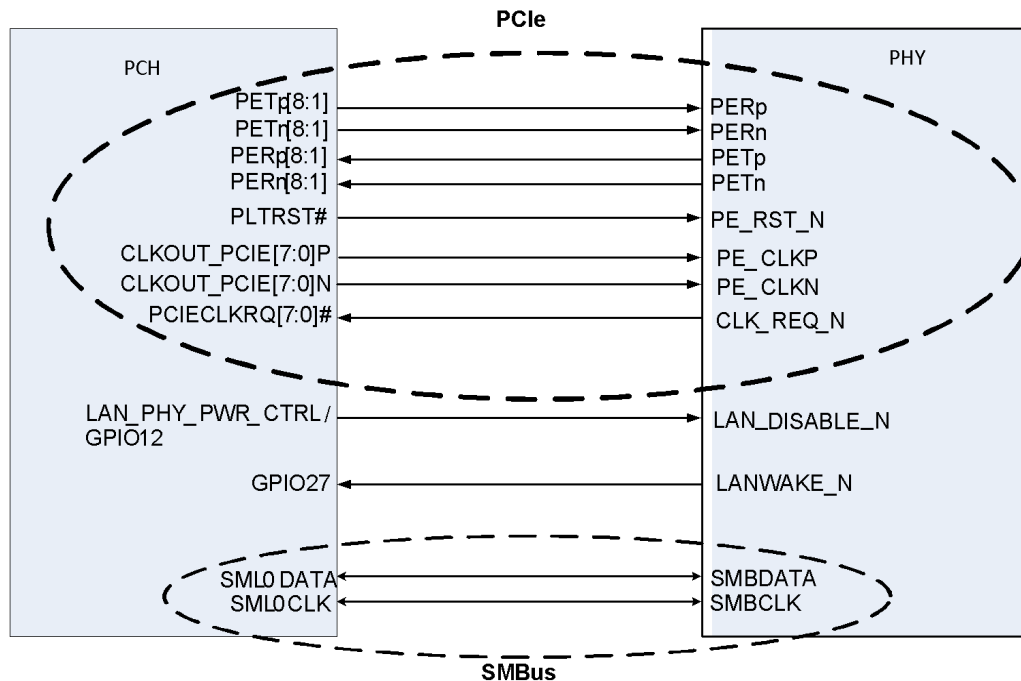




Figure 3. PCIe/SMBus Interface

Notes:

1. For latest PHY schematic connection recommendations, refer to the *Intel® I218 GbE PHY Reference Schematic*, and the *Intel® I218 GbE PHY Schematic and Layout Checklist*, available on CDI.
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.



Note: *Note:* This page intentionally left blank.



3.0 Pin Interface

3.1 Pin Assignment

The Intel® Ethernet Connection I218 is packaged in a 48-pin package, 6 x 6 mm with a 0.4 mm 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 499Ω 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 499Ω 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, Intel® Ethernet Connection I218 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] 1000BASE-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] 1000BASE-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] 1000BASE-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 k Ω resistor when connected to logic 1b and test mode is enabled. |


Note: The Intel® Ethernet Connection I218 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 VDD0P9_11 VDD0P9_16 VDD0P9_22 VDD0P9_37 VDD0P9_40 VDD0P9_43 VDD0P9_46 VDD0P9_47 | 8 11 16 22 37 40 43 46 47 | Power | Core Vdc supply |
| VDD3P3_4 VDD3P3_IN VDD3P3_15 VDD3P3_19 VDD3P3_29 | 4 5 15 19 29 | Power | 3.3 Vdc supply |
| CTRL0P9 | 7 | Analog | Internal SVR control pin. Connect to a 4.7 μ H inductor and to the core rail. |
| SVR_EN_N | 6 | Input | SVR Enable pin. Connect to GND for internal SVR |

3.2 Pinout





4.0 Package

4.1 Package Type and Mechanical

The Intel® Ethernet Connection I218 package is a 6 mm x 6 mm, 48-pin QFN Halogen Free and Pb Free package with Epad size of 3 mm x 3 mm.




Figure 4. 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 $^{\circ}C_{ambient}$.

| 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 are down-bonded to an Exposed Pad* e-PAD*.

| External Power Supply | 3.3 V dc, Pin 5 |
|-----------------------|--|
| Core V dc | 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 | Pin Name | Side | Pin Number |
|---------------|---------------|------------|------------|--------|------------|
| RSVD1_VCC3P3 | Left | 1 | MDI_PLUS0 | Bottom | 13 |
| LANWAKE_N | Left | 2 | MDI_MINUS0 | Bottom | 14 |
| LAN_DISABLE_N | Left | 3 | VDD3P3_15 | Bottom | 15 |
| VDD3P3_4 | Left | 4 | VDD0P9_16 | Bottom | 16 |
| VDD3P3_IN | Left | 5 | MDI_PLUS1 | Bottom | 17 |
| SVR_EN_N | Left | 6 | MDI_MINUS1 | Bottom | 18 |
| CTRL_0P9 | Left | 7 | VDD3P3_19 | Bottom | 19 |
| VDD0P9_8 | Left | 8 | MDI_PLUS2 | Bottom | 20 |
| XTAL_OUT | Left | 9 | MDI_MINUS2 | Bottom | 21 |
| XTAL_IN | Left | 10 | VDD0P9_22 | Bottom | 22 |
| VDD0P9_11 | Left | 11 | MDI_PLUS3 | Bottom | 23 |
| RBIAS | Left | 12 | MDI_MINUS3 | Bottom | 24 |
| LED2 | Right | 25 | VDD0P9_37 | Top | 37 |
| LED0 | Right | 26 | PETp | Top | 38 |
| LED1 | Right | 27 | PETn | Top | 39 |
| SMB_CLK | Right | 28 | VDD0P9_40 | Top | 40 |
| VDD3P3_29 | Right | 29 | PERp | Top | 41 |
| TEST_EN | Right | 30 | PERn | Top | 42 |
| SMB_DATA | Right | 31 | VDD0P9_43 | Top | 43 |
| JTAG_TDI | Right | 32 | PE_CLKP | Top | 44 |
| JTAG_TMS | Right | 33 | PE_CLKN | Top | 45 |
| JTAG_TDO | Right | 34 | VDD0P9_46 | Top | 46 |
| JTAG_TCK | Right | 35 | VDD0P9_47 | Top | 47 |
| PE_RST_N | Right | 36 | CLK_REQ_N | Top | 48 |
| GND_EPAD | Epad (Center) | 49 | | | |



Note: This page intentionally left blank.



5.0 Initialization

5.1 Power Up

Initialization begins with power up.




Figure 5. Power-Up Sequence



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

| Note | |
|------|---|
| 1 | Platform power ramps up (3.3 Vdc/Core Vdc) |
| 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 Intel® Ethernet Connection I218. |
| 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 Intel® Ethernet Connection I218 are as follows:

- **Internal Power On Reset (POR)**—the Intel® Ethernet Connection I218 has an internal mechanism for sensing the power pins. Until power is up and stable, Intel® Ethernet Connection I218 generates an internal active low reset. This reset acts as a master reset for Intel® Ethernet Connection I218. While the internal reset is 0b, all registers in Intel® Ethernet Connection I218 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 Connection I218 registers are not.

Note: The integrated LAN controller configures the LCD registers. Other Intel® Ethernet Connection I218 LCD registers do not need to be configured.

- **PCIe Reset** from pin—After asserting a PCIe reset, Intel® Ethernet Connection I218 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 the Intel® Ethernet Connection I218 except the wake up filters content.

Note: LAN_DISABLE_N is the only external signal that can reset the PHY. Refer to [Section 6.2.1](#) 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 [Table 18](#).

[Table 17](#) lists the impact of each Intel® Ethernet Connection I218 reset.



Table 17. Reset Impact

| 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 | |

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 Connection I218’s MDIO space is not reset.

5.3 Timing Parameters

5.3.1 Timing Requirements

The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 it’s 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.



6.0 Power Management and Delivery

This section describes power management in Intel® Ethernet Connection I218.

6.1 Power Information

The following table lists the targets for device power for the Intel® Ethernet Connection I218. Note that power is reduced according to link speed and link activity.

Table 18. Power Information

| System State | | Link State | Device Power using Internal SVR (mW) |
|-----------------|---------------------|---------------------------------------|--------------------------------------|
| S0 (Max) | | Thermal Design Power (TDP) | 542 |
| S0 (Typ) | | 1000 Mbps Active | 542 |
| | | 1000 Mbps Idle | 439 |
| | | 1000 Mbps LPI (EEE Link Partner Only) | 55 |
| | | 100 Mbps Active | 264 |
| | | 100 Mbps Idle | 177 |
| | | 100 Mbps LPI (EEE Link Partner Only) | 51 |
| | | 10 Mbps Active | 306 |
| | | 10 Mbps Idle | 84 |
| | | | |
| Sx | WoL Enabled | 100 Mbps WOL enabled | 169 |
| | | 10 Mbps WOL enabled | 76 |
| | | Cable Disconnect ULP Mode | 0.13 |
| | WoL Disabled | LAN disabled using BIOS | 0 |

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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 the Intel® Ethernet Connection I218 and until the Intel® Ethernet Connection I218 powers up its PHY. The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 is not actively transmitting or receiving data and might enter a lower power state (for example, an interface can be in electrical idle).
- **Cable Disconnect**—The PHY identified that a cable is not connected. The Intel® Ethernet Connection I218 signals the integrated LAN controller that the link is down. The PHY might enter energy detect mode or the integrated LAN controller might initiate a move into active power down mode (sD3).
- **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](#)).
- **Power Down (LAN Disable)**—Entry into power down is initiated by the integrated LAN controller through an in-band message or by setting the LAN_PWR_GOODn pin to zero. The Intel® Ethernet Connection I218 loses all functionality in this mode other than the ability to power up again.
- **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 Intel® Ethernet Connection I218 functionality when no data needs to be transmitted or/and received. When LPI support is enabled the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 and until the Intel® Ethernet Connection I218 powers up its PHY. The Intel® Ethernet Connection I218 should consume less than ~40 mA during this period. Following the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 detects energy on the MDI link, and starts the following exit sequence:

- The Intel® Ethernet Connection I218 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 about two to five seconds. This is expected behavior for a link-down.

6.3.1.3 Power Down State

The Intel® Ethernet Connection I218 enters a power-down state when the LAN_DISABLE_N 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 Intel® Ethernet Connection I218 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. Power-Down Sequence

Note: In cases where the LAN_DISABLE_N pin cannot be used a power down inband can be used. When used the power savings are lower since not all logic can be turned off at



this mode.

| Note | Description |
|------|--|
| 2 | Once the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 functionality when no data needs to be transmitted or/and received. Decision on whether the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 and communicated to the integrated LAN controller to allow for power saving in the receive circuitry.

Figure 7 illustrates general principles of an EEE LPI operation on the Ethernet Link.

Table 19. 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 7. 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 Intel® Ethernet Connection I218 that Low Power Idle mode should end. The Intel® Ethernet Connection I218 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 $T_{w_System_tx}$ 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 I218 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 Intel® Ethernet Connection I218 does not enter ACBS mode.



The Intel® Ethernet Connection I218 ACBS works in both S0 and Sx states. Since the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 I218 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 I218 auto negotiates with the updated LPLU setting on the following auto-negotiation operation. The I218 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 I218 in the various speeds.

6.4 The Intel® Ethernet Connection I218 Ultra Low Power (ULP)

To reduce power consumption during link disconnect, the Intel® Ethernet Connection I218 uses an Ultra Low Power (ULP) mode. In ULP mode, the Intel® Ethernet Connection I218 power gates most of its functionality and only maintains a small portion of the die-powered sensing for future cable connection (energy on the Ethernet link). Once energy is detected, the Intel® Ethernet Connection I218 exits ULP mode and either send an in-band message (link connect status) or assert LANWAKE# according to its configuration.

The Intel® Ethernet Connection I218 functionality in ULP mode (In-band/WoL exit, SMBus/PCIe exit, etc.) is controlled by the host driver (on non-ME systems) or the ME FW. Software configures the Intel® Ethernet Connection I218 registers for proper entry and exit conditions.

The Intel® Ethernet Connection I218 can be configured to enter ULP on future link disconnects while the link is still connected so that ensuing entry to Sx future link disconnects will benefit from the ULP mode.

For proper functionality and power consumption during ULP mode, the board design of Intel® Ethernet Connection I218 should be such that:

- LANWAKE_N, SMB_CLK, SMB_DATA, CLK_REQ_N are pulled up on board.
- TEST_EN is driven low or pulled down on board.
- LAN_DISABLE_N, PE_RST_N can be driven either low or high.
- PE_CLKP/N are both either driven low or tri-stated in ULP mode.
- LED0, LED1, LED2 are in tri-state mode during ULP.



6.5 Off Board LAN Connected Device (OBLCD)

To enable OBLCD for platforms where the LAN Connected Device is not located on the motherboard (located either on a docking station or on a dongle), additional communication is required between the LAN Controller and the LAN Connected Device so that the LAN Controller is aware of OBLCD connect and disconnect events.

6.5.1 Indicating/Sensing OBLCD disconnect event

During K0, CLKREQ# is asserted by the LAN Connected Device. On OBLCD disconnect the CLKREQ# de-asserts (board controlled) not as part of the K1 entry flow. The LAN Controller refers to such an event as an indication of OBLCD disconnect

During K1, CLKREQ# is de-asserted. To indicate presence, the LAN Connected Device sends an in-band beacon message once every BCN_INTER for BCN_DUR to indicate to the LAN Controller that the dongle is still connected. The LAN Controller does not exit K1 when sensing these squelch pulses. K1 exit starts only after sensing a beacon pulse longer than LC.OFFSET+BCN_DUR. At this state the LAN Controller enables the PLL and starts the K1 exit flow. If the LAN Controller misses these beacon pulse indications for LC.MISS_BCN*BCN_INTER, it perceives such an event as an indication of dongle disconnect. In addition, the LAN Controller implements a mode to exit K1 on CLKREQ# assertion to avoid the additional latency in the K1 exit sequence.

While in SMBus, the LAN Connected Device is configured to send periodic status in-bands and RX packets. If the LAN Controller doesn't sense RX activity (packets or statuses) for NO_RX_ON_SMBUS time, it refers to such an event as an indication of OBLCD disconnect.

While in ULP, the LAN Connected Device is power gated so there is no notification on disconnect and the LAN Controller is not notified of such an event until:

- It tries to take the LAN Connected Device out of ULP without success
- OBLCD is reconnected (in S0) and the LAN Connected Device asserts CLKREQ# indicating that the LAN Controller move to PCIe.

Note: When the LAN Connected Device is in WoL mode and the LAN Controller is powered off, there is no indication of OBLCD disconnect.

6.5.2 Indicating/Sensing OBLCD connection

OBLCD connection is detected by using CLKREQ# assertion on connection. The default interface of the OBLCD is PCIe. Connection is made on the following flows:

In PCIe, the LAN Connected Device causes K1 exit by asserting the CLKREQ# and then sending status in-bands over PCIe.

While the interface is defined as SMB, OBLCD connection is not reflected as up while the interface is SMB (and PERST# is asserted).



7.0 Device Functionality

7.1 Tx Flow

When packets are ready for transmission in the integrated LAN controller, it transfers them to the Intel® Ethernet Connection I218 through either the PCIe or the SMBus depending on system state. The Intel® Ethernet Connection I218 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 link speed. The Intel® Ethernet Connection I218 design is based on the assumption that the integrated LAN controller has the full packet ready for transmission.

In several cases, the Intel® Ethernet Connection I218 stops transmission over the wire while still accepting data from the upper connection (PCIe or SMBus). In those cases, it maintains a 3 KB FIFO. The cases in which the Intel® Ethernet Connection I218 stops 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 if the Intel® Ethernet Connection I218 is in the middle of a receive (DEFER).
- In half-duplex mode if a collision was detected on the wire.

In addition to stopping transmission, the Intel® Ethernet Connection I218 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 that caused the stop transmit is not valid and transmission over the wire is activated, the Intel® Ethernet Connection I218 might avoid sending the in-band message. An in-band message with the Tx *OFF* bit cleared is sent when the collided packet is successfully transmitted or dropped after 16 retries (see [Section 7.3.1](#)).

In-band messages from the integrated LAN controller to the Intel® Ethernet Connection I218 always come between packets during the IPG. The Intel® Ethernet Connection I218 does not accept in-band messages in the middle of a packet.

7.2 Rx Flow

The Intel® Ethernet Connection I218 maintains a 3 KB FIFO on the receive side in order to not lose packets when PCIe is active but in K1 power save mode. In this case, the Intel® Ethernet Connection I218 initiates recovery of the PCIe when reception starts. 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 Intel® Ethernet Connection



I218 starts the transmission after the entire packet is received. The Intel® Ethernet Connection I218 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.

The Intel® Ethernet Connection I218 identifies PAUSE packets, stop transmission, and a send in-band message as described in the previous section.

In-band messages from the Intel® Ethernet Connection I218 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 substantial 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 Intel® Ethernet Connection I218. The following registers are duplicated to the Intel® Ethernet Connection I218 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 included in the transmitted PAUSE frame
- Flow Control Refresh Threshold Value (FCRTV): a 16 bit PAUSE refresh threshold value

Flow control is implemented as a means of reducing the possibility of receive buffer overflows resulting in dropped received packets, as well as for local control of network congestion levels. This is 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, by not having 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 with 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 8. 802.3x MAC Control Frame Format

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 Intel® Ethernet Connection I218 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.
- Forwards 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 Intel® Ethernet Connection I218 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 are 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 receive and transmit FIFO status. If the sum of data in both FIFOs is greater than the configured *FCTH.HTHRSH*, the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 sends a PAUSE frame with a timer value of zero (XON). The Intel® Ethernet Connection I218 sends 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 Intel® Ethernet Connection I218 is operating in half-duplex mode.

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



7.4 Wake Up

The Intel® Ethernet Connection I218 supports host wake up.

The host can enable host wake up from the Intel® Ethernet Connection I218 by setting the *Host_WU_Active* bit.

To set the Intel® Ethernet Connection I218's wake up:

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 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218's wake up functionality.

On a detection of a Host wakeup event, the Intel® Ethernet Connection I218:

1. Asserts the LANWAKE# pin indicating wake to PCH.
2. If in DeepSx, PCH powers up the SUS well.
3. PCH sends a PINSTOP message to the Intel® Ethernet Connection I218 through SMBus

On reception of a PINSTOP message, the Intel® Ethernet Connection I218 stops asserting the LANWAKE# pin and sends a WU message over SMBus indicating the WU source.

The Intel® Ethernet Connection I218 sends a WU message on every PINSTOP message reception.

4. The Intel® Ethernet Connection I218 does not assert the LANWAKE# pin again until a WU message is sent and power-up PCH acknowledged. In case of host wake up, the integrated LAN controller wakes the host.
5. The Host should issue an LCD reset to the Intel® Ethernet Connection I218 before clearing the *Host_WU_Active* bit.



6. Host reads the Wake Up Status register (WUS); wake up status from the Intel® Ethernet Connection I218.

When a wake up packet is identified, the wake up in-band message is sent and the host clears 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 won't overwrite the packet in the FIFO. The Intel® Ethernet Connection I218 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

The Intel® Ethernet Connection I218 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 do so).

At power up, if the Intel® Ethernet Connection I218's wake up functionality is enabled, the *APM Enable* bits from the NVM are written to the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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, 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 *RAL0[31:0]* and *RAH0[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 *RAH0*) 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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

The Intel® Ethernet Connection I218 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 supplies 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 Intel® Ethernet Connection I218 in wake up mode.

Once wake up is enabled, the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218:

- 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 Intel® Ethernet Connection I218 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 clears the wake-up packet indication in preparation for a subsequent low power state entry.



7.4.1.4 Wake-Up Packets

The Intel® Ethernet Connection I218 supports various wake-up packets using two types of filters:

- Pre-defined filters
- Flexible filters

Each of these filters is 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.NoHostWoLonME traffic bit that enables 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 Intel® Ethernet Connection I218’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.

7.4.1.4.1.1 Directed Exact Packet

The Intel® Ethernet Connection I218 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 |

7.4.1.4.1.2 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 Intel® Ethernet Connection I218 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. |



7.4.1.4.1.3 Broadcast

If the *Broadcast Wake Up Enable* bit (bit 4) in the WUFC register is set, the Intel® Ethernet Connection I218 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 | |

7.4.1.4.1.4 Magic Packet

Magic packets are defined as follows:

- Magic Packet Technology Details**—Once the Intel® Ethernet Connection I218 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.

The Intel® Ethernet Connection I218 expects the destination address to either:

- Be the broadcast address (FF.FF.FF.FF.FF.FF)
- 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.
- 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 starts up with the RAH0/RAL0 register 0 loaded from the NVM. This enables the Intel® Ethernet Connection I218 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 |

7.4.1.4.1.5 ARP/IPv4 Request Packet

The Intel® Ethernet Connection I218 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. The Intel® Ethernet Connection I218 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

The Intel® Ethernet Connection I218 supports reception of Directed IPv4 packets for wake up if the IPV4 bit (bit 6) is set in the WUFC register. Three IPv4 addresses are supported and 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. The Intel® Ethernet Connection I218 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 |

7.4.1.4.2.1 Directed IPv6 Packet

The Intel® Ethernet Connection I218 supports reception of Directed IPv6 packets for wake up if the *IPV6* bit (bit 7) is set in the WUFC register. One IPv6 address is supported and it 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 Intel® Ethernet Connection I218 also handles Directed IPv6 packets having 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 | 4 | Possible VLAN Tag | | Skip | |
| 12 | 8 | Possible LLC/SNAP Header | | Skip | |
| 12 | 2 | Type | 0x86DD | Compare | IP |
| 14 | 1 | Version/ Priority | 0x6X | Compare | Check Ipv6 |
| 15 | 3 | Flow Label | - | Ignore | |
| 18 | 2 | Payload Length | - | Ignore | |
| 20 | 1 | Next Header | - | Ignore | |
| 21 | 1 | Hop Limit | - | Ignore | |
| 22 | 16 | Source IP Address | - | Ignore | |
| 38 | 16 | Destination IP Address | IP6AT | Compare | Match value in IP6AT |



7.4.1.4.3 Flexible Filter

The Intel® Ethernet Connection I218 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, the 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. The software must also enable the filter in the Wake Up Filter Control Register (WUFC), and enable the overall wake up functionality must be enabled by setting PME_En in the Power Management Control Status Register or the Wake Up Control 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), 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.

The following packets are listed for reference purposes only. The flexible filter could be used to filter these packets.

IPX Diagnostic Responder Request Packet

An IPX Diagnostic Responder Request Packet must contain a valid MAC address, a Protocol Type of 0x8137, and an IPX Diagnostic Socket of 0x0456. It may include LLC/SNAP Headers and VLAN Tags. Since filtering this packet relies on the flexible filters, which uses offsets specified by the operating system directly, the operating system must account for the extra offset LLC/SNAP Headers and VLAN tags.

| Offset | # of bytes | Field | Value | Action | Comment |
|--------|------------|--------------------------|--------|---------|---------|
| 0 | 6 | Destination Address | | Compare | |
| 6 | 6 | Source Address | | Skip | |
| 12 | 4 | Possible VLAN Tag | | Skip | |
| 12 | 8 | Possible LLC/SNAP Header | | Skip | |
| 12 | 2 | Type | 0x8137 | Compare | IPX |
| 14 | 16 | Some IPX Stuff | - | Ignore | |
| 30 | 2 | IPX Diagnostic Socket | 0x0456 | Compare | |

Directed IPX Packet

A valid Directed IPX Packet contains the station's MAC address, a Protocol Type of 0x8137, and an IPX Node Address that equals to the station's MAC address. It may include LLC/SNAP Headers and VLAN Tags. Since filtering this packet relies on the flexible filters, which use offsets specified by the operating system directly, the operating system must account for the extra offset LLC/SNAP Headers and VLAN tags.



| 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 | 4 | Possible VLAN Tag | | Skip | |
| 12 | 8 | Possible LLC/SNAP Header | | Skip | |
| 12 | 2 | Type | 0x8137 | Compare | IPX |
| 14 | 10 | Some IPX Stuff | - | Ignore | |
| 24 | 6 | IPX Node Address | Receive Address 0 | Compare | Must match Receive Address 0 |

7.4.1.4.3.1 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.

7.4.2 PCI-M Power Management (PCI-m)

The LAN Controller supports a similar concept to ACPI for the PCI-M space accessed by the ME.

7.4.2.1 Introduction to PCIm power states

The PCI-M device supports the power states defined by the PCIm and M-Link spec: DMoff, DM1, DM2 and DM3.

At DMoff the LAN controller may provide on the PCIm interface only WoL functionality to ME while in all other states the LAN controller may provide its complete functionality. It is the ME firmware responsibility to enable the memory BAR and bus master accesses for complete functionality. It is also the responsibility of the ME firmware to limit the LAN controller from ME host memory accesses while it is in M2 and M3 power states.

The diagram below shows the power states and transitions between them.

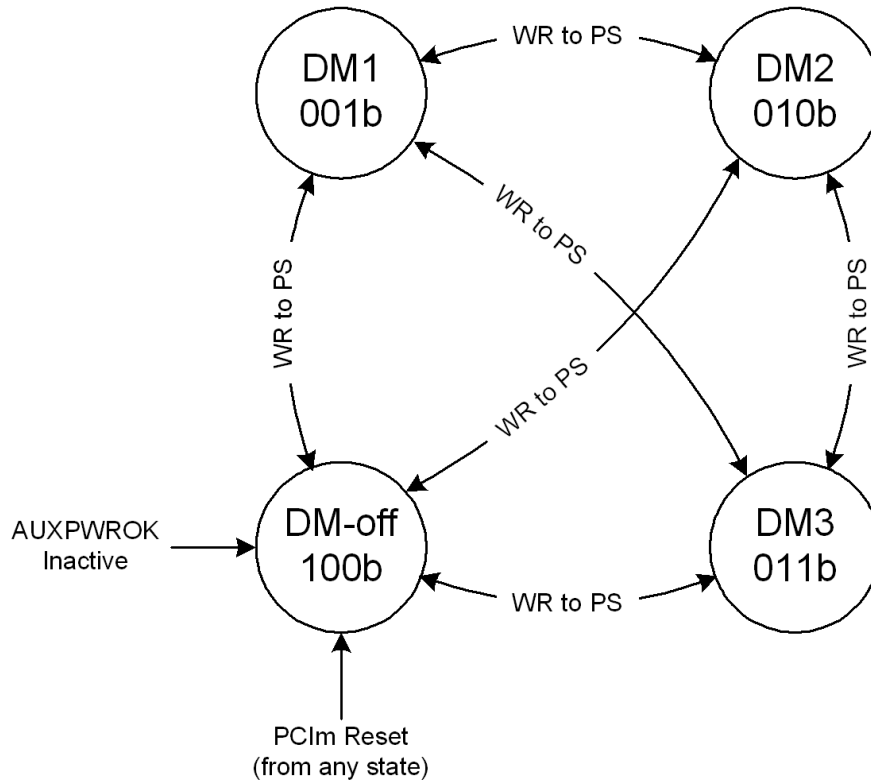


Figure 7-1. PCIe Power Management State Diagram

7.4.2.2 Unit power reduction measures

This section describes the power reduction techniques employed by the LAN Controller. The internal ME DMA engine (GPT block) will have a running clock as long as the PCIe function is in functional power states: DM1, DM2 or DM3.

When the CLK_CNT_1_4 bit in the Status CSR (loaded from NVM) is "0," the DMA clock runs at full frequency regardless of the link speed / state. The HW default value is "1".

- 1000Mb/s link: The DMA gets the MAC clock of 62.5MHz
- At 10/100Mbps or link negotiation or link down the DMA gets the mosc_clk (62.5MHz)

When the CLK_CNT_1_4 bit is set (the HW default) it enables dynamic reduction of clock frequency in the functional power states.

- 1000Mb/s link: it runs at full speed of 62.5MHz from the MAC
- Else (link at 10/100Mbps, link negotiation, or link down), it operates at ¼ of the mosc_clk frequency



7.4.2.3 Power States

7.4.2.3.1 DMoff State

The DMoff state is a low-power state following the assertion of PCI_m reset or following programming of the DM state in the PCI_m configuration space to 100b. The LAN controller exits the DMoff state by ME setting of the power state. While in this state the LAN controller may provide WoL to ME functionality as described on [page 60](#). The PCI_m logic is clock gated if WoL is not enabled for power reduction. A system without an ME will have a PCI_m function that will remain in DMoff low power state forever. ME should set the LAN controller to any other state as required only after the ready bit in the PCI_m status register indicates that the device completed its basic initialization from the NVM.

7.4.2.3.2 Switch to DMoff State

Before the ME sets the device to the DMoff state it should follow the steps below:

1. Set the RSPCIPHY bit in the FWSM register to enable the host to have complete control on the LCD.
2. If WoL to ME is not required then turn off the RCV_TCO_EN bit in the MANC register to enable setting the LCD to power down. Else, turn on the RCV_TCO_EN bit to avoid setting the PHY to power down and enable the wake-up filtering.
3. Stop gracefully all master activities on the MBB by disabling transmit, receive and mask all interrupts.
4. Clear the Bus Master Enable in the PCI_m Command Register and poll the Transaction Pending bit in the PCI_m Status Register.
5. Clear the MSI Enable bit in the PCI_m Message Control register and poll the pending interrupt indication until there are no pending MSI.
6. Clear the Memory and IO Access Enable in the PCI_m Command Register.
7. Write to the device's PCI-M Power Management PM State register to put the device in DMOff.

7.4.2.3.3 DM1, DM2 and DM3 States

The LAN controller may enter the DM1/DM2/DM3 states following programming of the DM state in the PCI_m configuration space to 01b/10b/11b respectively. The LAN controller may provide its complete functionality over PCI_m in these states. The device may change its state to any other state following programming of the DM state in the PCI_m configuration space. The device returns to the DMoff state following assertion of PCI_m reset or de-assertion of AUXPWROK.

Prior to transition from active states to DMoff state, the ME should do the following:

- Disable scheduling of further tasks to the LAN Controller.
- Mask all interrupts.
- Wait till the transmit ring buffer is empty (all packets were fetched by the LAN controller).
- Turn off the transmit and receive enable flags in the MTBCR and the MRBCR registers.



7.4.2.4 ME Wake-Up

The Intel® Ethernet Connection I218 has the ability to wake-up the ME. Any packet routed to the ME during normal operation should wake-up the ME when it is in the Mox state. Any packet causing ME wake-up must first match the MAC address filtering including Exact Unicast/Multicast filtering, hash-based Multicast filtering of 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 9. ME Receive Wake Up Filters Top Level



7.4.2.4.1 Manageability Registers

The Manageability filtering configuration is done by the manageability firmware. The following is a list of registers in the Intel® Ethernet Connection I218 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.

7.4.2.4.2 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.

7.4.2.4.3 ARP Packets Filtering

The Intel® Ethernet Connection I218 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.



7.4.2.5 Flexible TCP UDP port and IP address filtering

The Intel® Ethernet Connection I218 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.

7.4.2.6 Dedicated TCP UDP port and IP address filtering

The Intel® Ethernet Connection I218 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, 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.

7.4.2.6.1 L2 Filtering

The Intel® Ethernet Connection I218 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 Connection I218 wakes up the ME (the packet must pass any of the MAC filters as well).

7.4.2.6.2 Flexible TCO Filtering

The Intel® Ethernet Connection I218 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/O_FILTER_EN* bits (bits 6 and 7) in the MANC register.

7.4.2.7 Accessing the Intel® Ethernet Connection I218'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 the wanted to maintain their network presence would have configured the ARP and neighbor discovery messages as wake up patterns to the system. Analysis show 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 an 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 Intel® Ethernet Connection I218:

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 PRXC [6:5]

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

Note: The Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 |



| Offset | # of bytes | Field | Value (hex) | Action |
|----------------------|------------|---------------------|-------------|--|
| 21+D+S | 1 | Hop Limit | 0xFF | Constant |
| 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 | | |

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.



7.5.5.1 MLD messages supported by the Intel® Ethernet Connection I218

As a receiver:

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

There are two subtypes of Multicast Listener Query messages supported by the Intel® Ethernet Connection I218:

- 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).

7.5.5.2 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 | |



| Offset | # of bytes | Field | Value (hex) | Action | Comment |
|--------|------------|-----------------------|-------------|---------|--|
| 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 |

7.5.5.3 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 Intel® Ethernet Connection I218 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 [section 9.5.1](#).

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.



8.0 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 contains registers controlling the custom interface as well as other the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 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 $page_num \times 32$ 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](#) for a description of registers addressing in these pages.

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



8.4 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 LS bit 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 |



8.5 Address Map

| PHY Address | Page | Register | Name | Page |
|--|------|----------|---------------------------------------|------|
| 02 | Any | 0 | Control | 81 |
| 02 | Any | 1 | Status | 82 |
| 02 | Any | 2 | PHY Identifier 1 | 82 |
| 02 | Any | 3 | PHY Identifier 2 | 83 |
| 02 | Any | 4 | Auto-Negotiation Advertisement | 83 |
| 02 | Any | 5 | Auto-Negotiation Link Partner Ability | 83 |
| 02 | Any | 6 | Auto-Negotiation Expansion | 84 |
| 02 | Any | 7 | Auto-Negotiation Next Page Transmit | 85 |
| 02 | Any | 8 | Link Partner Next Page | 85 |
| 02 | Any | 9 | 1000BASE-T Control | 85 |
| 02 | Any | 10 | 1000BASE-T Status | 86 |
| 02 | Any | 14:11 | Reserved | |
| 02 | Any | 15 | Extended Status | 86 |
| 02 | 0 | 17:16 | Reserved | |
| 02 | 0 | 18 | PHY Control 2 | 87 |
| 02 | 0 | 19 | Loopback Control | 88 |
| 02 | 0 | 20 | Rx Error Counter | 89 |
| 02 | 0 | 21 | Management Interface (MI) Control | 90 |
| 02 | 0 | 22 | PHY Configuration | 90 |
| 02 | 0 | 23 | PHY Control | 91 |
| 02 | 0 | 24 | Interrupt Mask | 91 |
| 02 | 0 | 25 | Interrupt Status | 93 |
| 02 | 0 | 26 | PHY Status | 93- |
| 02 | 0 | 27 | LED Control 1 | 95 |
| 02 | 0 | 28 | LED Control 2 | 96 |
| 02 | 0 | 29 | LED Control 3 | 97 |
| 02 | 0 | 30 | Late Collision Counter | 97 |
| 02 | 0 | 31 | Link Polarity and Length Status | 97 |
| Page 769—Port Control Registers | | | | |
| 01 | 769 | 16 | Custom Mode Control | 102 |
| 01 | 769 | 17 | Port General Configuration | 102 |
| 01 | 769 | 21 | Power Management Control | 102 |
| 01 | 769 | 23 | SMBus Control | 103 |
| 01 | 769 | 25 | Rate Adaptation Control | 103 |
| 01 | 769 | 27 | Flow Control Transmit Timer Value | 103 |
| 01 | 769 | 28 | System Low Power Control | 104 |
| 01 | 774 | 23-24 | SERDES MDI Control Register | 104 |
| Page 778—Statistics Registers | | | | |
| 01 | 778 | 16 - 17 | Single Collision Count | 105 |



| PHY Address | Page | Register | Name | Page |
|-----------------------------------|------|----------|---|------|
| 01 | 778 | 18 - 19 | Excessive Collisions Count | 105 |
| 01 | 778 | 20 - 21 | Multiple Collisions Count | 105 |
| 01 | 778 | 23 - 24 | Late Collision Count | 105 |
| 01 | 778 | 25 - 26 | Collision Count | 105 |
| 01 | 778 | 27 - 28 | Defer Count | 106 |
| 01 | 778 | 29 - 30 | Transmit with No CRS—TNCRS | 106 |
| PCIe Registers | | | | |
| 01 | 770 | 16 | PCIe FIFOs Control/Status | 106 |
| 01 | 770 | 17 | PCIe Power Management Control | 106 |
| 01 | 770 | 18 | In-Band Control | 107 |
| 01 | 770 | 20 | PCIe Diagnostics | 107 |
| 01 | 770 | 21 | Time Outs | 107 |
| 01 | 770 | 23 | PCIe K-State Minimum Duration Timeout | 108 |
| LPI Registers | | | | |
| 01 | 772 | 18 | Low Power Idle GPIO Control | 108 |
| 01 | 772 | 20 | Low Power Idle Control | 108 |
| 01 | 772 | 23 | Flow Control Refresh Threshold Value | 109 |
| 01 | 772 | 24 | Flow Control Thresholds | 109 |
| 01 | 772 | 25 | LANWAKE# Control | 109 |
| 01 | 772 | 26 | Memories Power | 110 |
| 01 | 772 | 29 | Config Register | 110 |
| ULP Registers | | | | |
| 01 | 779 | 16 | ULP Configuration 1 | 110 |
| 01 | 779 | 17 | ULP Configuration 2 | 111 |
| 01 | 779 | 18 | ULP SW Control | 111 |
| General Registers | | | | |
| 01 | 779 | 20 | Off Board LAN Connected Device Control | 111 |
| 01 | 776 | 18 | Intel I218 Energy Detect Mode | 111 |
| 01 | 776 | 19 | Intel® Ethernet Connection I218 Capability Register | 112 |
| 01 | 0 | 25 | OEM Bits | 112 |
| 01 | 0 | 26 | SMBus Address | 113 |
| 01 | 0 | 27-28 | Shadow Register for RAL0[31:0]. | 113 |
| 01 | 0 | 29 | Shadow Register for RAH0[15:0]. | 113 |
| 01 | 0 | 30 | LED Configuration | 114 |
| Page 800—Wake Up Registers | | | | |
| 01 | 800 | 0 | Receive Control Register | 116 |
| 01 | 800 | 1 | Wake Up Control Register | 117 |
| 01 | 800 | 2 | Wake Up Filter Control Register | 117 |
| 01 | 800 | 3 | Wake Up Status Register | 118 |



| PHY Address | Page | Register | Name | Page |
|---------------------------------|------|--------------------------|---|------|
| 01 | 800 | 16 | Receive Address Low 0 | 119 |
| 01 | 800 | 18 | Receive Address High 0 | 119 |
| 01 | 800 | 20-21 + 4*n (n=0...10) | Shared Receive Address Low | 119 |
| 01 | 800 | 22-23 + 4*n (n=0...8,10) | Shared Receive Address High | 119 |
| 01 | 800 | 58 - 59 | Shared Receive Address High 9 | 120 |
| 01 | 800 | 64 | IP Address Valid—IPAV | 120 |
| 01 | 800 | 70 | Proxy Control | 121 |
| 01 | 800 | 71 | Proxy Code Checksum | 121 |
| 01 | 800 | 72 | Proxy Control 2 | 121 |
| 01 | 800 | 75 | Flex Filters Proxy Control | 122 |
| 01 | 800 | 76 | Wake Up Filter Control 2 | 122 |
| 01 | 800 | 77 | Wake Up Filter Status 2 | 123 |
| 01 | 800 | 82 - 83 | IPv4 Address Table—IP4AT 0 | 122 |
| 01 | 800 | 88 - 89 | IPv6 Address Table—IP6AT 0 | 123 |
| 01 | 800 | 128 - 191 | Multicast Table Array—MTA[31:0] | 123 |
| 01 | 800 | 256 + 2*n (n = 0 - 127) | Flexible Filter Value Table LSB—FFVT_01 | 124 |
| 01 | 800 | 257 + 2*n (n = 0 - 127) | Flexible Filter Value Table MSB—FFVT_23 | 125 |
| 01 | 800 | 512 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_45 | 125 |
| 01 | 800 | 1024 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_67 | 125 |
| 01 | 800 | 768 + n (n = 0 - 127) | Flexible Filter Mask Table—FFMT | 126 |
| 01 | 800 | 896 + n (n = 0 - 3) | Flexible Filter Length Table—FFLT03 | 126 |
| 01 | 800 | 904 + n (n = 0 - 1) | Flexible Filter Length Table—FFLT45 | 127 |
| 01 | 800 | 908 + n (n = 0 - 1) | Flexible Filter Length Table—FFLT67 | 127 |
| 01 | 800 | 2304 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_89 | 127 |
| 01 | 800 | 2305 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_1011 | 127 |
| 01 | 800 | 2560 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_1213 | 128 |
| 01 | 800 | 3072 + 2*n (n = 0 - 127) | Flexible Filter Value Table—FFVT_1415 | 128 |
| 01 | 800 | 2816 + n (n = 0 - 127) | Flexible Filter Mask Table 2—FFMT2 | 128 |
| 01 | 800 | 2944 + n (n = 0 - 3) | Flexible Filter Length Table—FFLT891011 | 128 |
| 01 | 800 | 2952 + n (n = 0 - 1) | Flexible Filter Length Table—FFLT1213 | 128 |
| 01 | 800 | 2956 + n (n = 0 - 1) | Flexible Filter Length Table—FFLT1415 | 129 |
| 01 | 801 | 30-31 | Management 2 Host Control Register | 129 |
| 01 | 801 | 32-33 | Management 2 Host Control Register 2 | 130 |
| 01 | 801 | 40-41 | IPv4 ME Binding Control | 132 |
| 01 | 801 | 42-43 | IPv4 Host Binding Control | 133 |
| 01 | 801 | 44-45 | IPv6 Binding Control | 134 |
| 01 | 801 | 46 | SHRA Filter Enable Register | 136 |
| Page 803—Host WoL packet | | | | |



| PHY Address | Page | Register | Name | Page |
|-------------|------|----------|----------------------------------|------|
| 01 | 803 | 0-63 | Host WoL packet | 109 |
| 01 | 803 | 64 | Host WoL packet length | 109 |
| 01 | 803 | 66 | Host WoL packet indication clear | 109 |

8.6 PHY Registers (Page 0)

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

Table 20. 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 21. 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 22. PHY Identifier Register 1—Address 284

| 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) |

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

**Table 23. 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. |

1. 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 24. 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 25. Auto-Negotiation Link Partner Ability Register—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. |
| 13 | Remote Fault | RO | 0b | 1b = Link partner has detected remote fault. 0b = Link partner has not detected remote fault. |



Table 25. Auto-Negotiation Link Partner Ability Register—Address 5

| Bits | Field | Type | Default | Description |
|------|-----------------------------------|------|---------|--|
| 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 26. Auto-Negotiation Expansion Register—Address 6

| 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 27. Auto-Negotiation Next Page Transmit Register—Address 7

| 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. |

**Table 27. Auto-Negotiation Next Page Transmit Register—Address 7**

| Bits | Field | Type | Default | Description |
|------|-------------------------------|------|---------|---|
| 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 28. Link Partner Next Page Register—Address 8

| 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 29. 1000BASE-T Control PHY Register—Address 9

| 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. |
| 10 | Port Type | R/W | 0b | 1b = Multi-port device (prefer master). 0b = Single port device (prefer slave). |



Table 29. 1000BASE-T Control PHY Register—Address 9

| Bits | Field | Type | Default | Description |
|------|---|------|---------|---|
| 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 30. 1000BASE-T Status Register—Address 10

| 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 <i>Technology Ability</i> 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 <i>Technology Ability</i> 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 31. Extended Status Register—Address 15**

| 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 32. PHY Control Register 2—Address 18

| 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 33. 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 34. 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 35. 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 36. 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 |

**Table 36. Loopback Control Register—Address 19**

| Bits | Field | Type | Default | Description |
|------|-------------------|------|---------|---|
| 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. |

8.6.1 Loopback Mode Settings

The table below 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 37. 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 38. 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 39. 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 Power-down Mode Transmit Enable | R/W | 1b | 1: Enables NLP transmission during energy-detect power-down. 0: Disables NLP transmission during energy-detect power-down. |
| 1:0 | Reserved | | 0x0 | Reserved |

Table 40. PHY Configuration 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 41. PHY Control Register—Address 23**

| Bits | Field | Type | Default | Description |
|-------|---|--------|---------|--|
| 15:14 | Reserved | R/O | 0x0 | Reserved |
| 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 = De-asserts MDINT_N pin. |

Table 42. 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. |



Table 42. Interrupt Mask Register—Address 24

| Bits | Field | Type | Default | Description |
|------|--------------------------------|------|---------|--|
| 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. |

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

**Table 43. 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. |
| 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 44. 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. |



Table 44. PHY Status Register—Address 26

| Bits | Field | Type | Default | Description |
|------|--|------|---------|---|
| 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. |
| 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 45. 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. |
| 9:8 | Reserved | | 0x0 | Reserved |
| 7:4 | LED Blink Pattern Pause | R/W | 0x0 | LED blink pattern pause cycles. |



Table 45. LED Control Register 1—Address 27

| Bits | Field | Type | Default | Description |
|------|--------------------|------|---------|--|
| 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 46. 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 | <p>Standard modes:</p> <ul style="list-style-type: none"> 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. <p>Extended modes:</p> <ul style="list-style-type: none"> 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 47. 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 48. 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 49. 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. |
| 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. |



Table 49. Link Polarity and length Status—Address 31

| Bits | Field | Type | Default | Description |
|------|-----------------|------|---------|--|
| 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 50. 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> |
| 11 | Short Between Pairs X and B | 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 B. 0b = No short between pairs X and B.</p> |



Table 50. Diagnostics Status Register (Linking Disabled)—Address 31

| Bits | Field | Type | Default | Description |
|------|-----------------------------|------|---------|--|
| 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 51. 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). |
| 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). |



8.7 Port Control Registers (Page 769)

Table 52. 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 53. Port General Configuration 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. |
| 5 | Reserved | R/W | 1b | Reserved |
| 4 | Host_WU_Active | R/W | 0b | Enables host wake up from the Intel® Ethernet Connection I218. 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 Intel® Ethernet Connection I218 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 |

Note: Register resets on Power Good only.

Table 54. DFT Control Register PHY Address 01, Page 769, Register 20

| Name | Default | Bits | Description | Type |
|----------|---------|------|-------------|------|
| Reserved | 0x140 | 15:7 | Reserved | R/W |
| Reserved | 0x0 | 6 | Reserved | R/W |
| Reserved | 0x0 | 5:0 | Reserved | R/W |

Table 55. 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 56. 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 57. 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 58. 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 <i>TTV</i> 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 <i>TTV</i> to 0x0000 prior to initiating the pause frame. | RW |

Note: Register resets on Power Good only.



Table 59. 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 indication 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 |

Note: Register resets on Power Good only.

Table 60. 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. |
| 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 |

This register is used by software to access the SERDES registers in the LAN Connected Device.



8.8 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 61. 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 Intel® Ethernet Connection I218 is in half-duplex mode.

Table 62. 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 Intel® Ethernet Connection I218 is in half-duplex mode.

Table 63. 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 Intel® Ethernet Connection I218 is in half-duplex mode.

Table 64. Late Collisions Count—LATECOL PHY Address 01, Page 778, Register 23-24

| Bit | Type | Reset | Description |
|------|------|-------|--|
| 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 Intel® Ethernet Connection I218 is in half-duplex mode.

Table 65. 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 Intel® Ethernet Connection I218 is in half-duplex mode. This register applies to clear as well as secure traffic.

Table 66. 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. For the Intel® Ethernet Connection I218, this counter does not increment for streaming transmits that are deferred due to TX IPG.

Table 67. 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 Intel® Ethernet Connection I218. |

This register counts the number of successful packet transmission in which the CRS input from the Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218.

The Intel® Ethernet Connection I218 should assert CRS during every transmission. Failure to do so might indicate that the link has failed, or the Intel® Ethernet Connection I218 has an incorrect link configuration. This register only increments if transmits are enabled. This register is only valid when the Intel® Ethernet Connection I218 is operating at half duplex.

8.9 PCIe Registers

Table 68. 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 69. 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 |

**Table 69. PCIe Power Management Control—PHY Address 01, Page 770, Register 17**

| | | | | |
|--|-------|-------|---|-----|
| Giga_K1_disable | 0b | 13 | When set, the Intel® Ethernet Connection I218 does not enter K1 while link speed is at 1000Mb/s. | R/W |
| 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 |

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 70. In-Band Control—PHY Address 01, Page 770, Register 18.

| 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 |

Table 71. 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 time-outs are multiplied by 1000 while in SMBus mode.

Table 72. 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 Intel® Ethernet Connection I218 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.



8.10 General Registers

Table 73. 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. | R/WP |
| EI_min_dur timeout | 0x10 | 4:0 | These bits define the minimum time the Intel® Ethernet Connection I218 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 74. 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 |
| 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 75. 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 μs 001b = 16 μs 010b = 24 μs 011b = 32 μs 100b = 40 μs 101b = 48 μs 110b = 56 μs 111b = 64 μs Before transmitting XOFF or a valid packet. | R/W |
| PLLLockCnt | 0x2 | 6:4 | PLL Lock Counter—when LPI 100Enable or 1000Enable is asserted, this counter forces the PLL Lock count to be in the range of 10.02 μs to 81.88 μs in steps of 10.02 μs. | |
| Reserved | 0x0 | 3:0 | Reserved | RO |



Table 76. 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 77. Flow Control Thresholds FCTH PHY—Address 01, Page 772, Register 24

| Name | Default | Bits | Description | Type |
|-------|---------|------|--|------|
| HTRSH | 0x00 | 7:0 | High threshold for sending XOFF (in units of 32 bytes) | RW |
| LTRSH | 0x00 | 15:8 | Low threshold for sending XON (in units of 32 bytes) | RW |

Note: Register resets on Power Good only.

Table 78. 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 82579 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 79. 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 80. Configuration Register PHY—Address 01, Page 772, Register 29

| 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 81. ULP Configuration 1 PHY Address 01, Page 779, Register 16

| Name | Default | Bits | Description | Type |
|-----------------|---------|-------|--|------|
| RESET_ULP_IND | 0b | 15 | Reset the ULP indication | R/W |
| FORCE_ULP | 0b | 14 | Internal Force PHY to energy power down | R/W |
| RESERVED | 0x0 | 13:11 | Reserved | R/W |
| EN_ULP_LANPHYPC | 0b | 10 | Enable ULP on LAN disable (LANPHYPC) | R/W |
| EN_1G_SMBUS | 0b | 9 | Enable 1 GbE in SMB mode | R/W |
| RESER_TO_SMBUS | 0b | 8 | Reset to SMBus by default (on power on or on ULP exit, functional only if power is supplied to the device) | R/W |
| WOL_ME | 0b | 7 | WoL ME on ULP exit | R/W |
| WOL_HOST | 0b | 6 | WoL host on ULP exit | R/W |
| INBAND_EXIT | 0b | 5 | In-band on ULP exit | R/W |
| STICKY_ULP | 0b | 4 | Sticky ULP Enter ULP on Link disconnect and wake on ULP exit if WoL ME/Host on ULP is set | R/W |
| RESERVED | 0b | 3 | Reserved | R/W |
| ULP_IND | 0b | 2 | Power Up from ULP indication | R/W |
| SW_ACCESS | 0b | 1 | Internal If this bit is set the software can access direct to the ULP 3P3 block | R/W |
| START | 0b | 0 | Start. When set the HW will start the auto ULP configuration Bit is auto cleared once configuration is done. | R/W |


Table 82. ULP Configuration 2 PHY Address 01, Page 779, Register 17

| Name | Default | Bits | Description | Type |
|----------|---------|------|------------------------------|------|
| Reserved | 0x00 | 15:5 | Reserved | R/W |
| MESHADOW | 0x0 | 4:0 | ME 3.3v Shadow configuration | |

Table 83. ULP SW Control PHY Address 01, Page 779, Register 18

| Name | Default | Bits | Description | Type |
|----------------------------|---------|------|--|------|
| ULP MAILBOX | 0x00 | 15:8 | This register is writeable and readable—can be used by the software by read modify write | R/W |
| Reserved | 0b | 7 | Reserved | R/W |
| ULP SW READ TO REGISTER | 0b | 6 | Set this bit for write the data on the selected ulp_bus on bits [15:8] on this register. Cleared by HW after implement the current read. | R/W |
| ULP SW WRITE FROM REGISTER | 0b | 5 | Select if the write will be from default ulp_bus or from bits [15:8] on this register. If set this bit will choose bit [15:8] on this register. Cleared by HW after implement the current write by: software_write_en_to_ulp signal. | R/W |
| SW GEN WRITE ENABLE | 0b | 4 | Cleared by HW after implement the current write enable by: software_write_en_to_ulp signal | R/W |
| SW BUS SELECTOR' | 0x0 | 3:0 | Select the bus address from the ulp bus table | R/W |

Table 84. 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 |
| 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 |

Table 85. Energy Detect Mode PHY Address 01, Page 776, Register 18

| Name | Default | Bits | Description | Type |
|---------------------------------------|---------|------|---|------|
| Reserved | 0x08 | 15:1 | Reserved | Rsvd |
| Disable Energy Detect Mode in the PHY | 0b | 0 | Overrides the energy detection disable from the bonding option, only if it was bonded to 1b (for example, energy detection is enabled). Note: Used for testing only. When set to 1b, overwrites the PHY power saving value, unless bond_opt[1:0] = 10b. This causes the PHY not to operate in power saving mode. | RW |



Table 86. 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 87. 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 ¹ | When set to 1b, 1000 Mb/s speed is disabled. |
| 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 | |

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

**Table 88. 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 Intel® Ethernet Connection I218 supports PEC on the SMBus. | RW |
| SMBus Frequency Low | 0b | 8 | Together with SMBus Furnace 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 Intel® Ethernet Connection I218 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 Intel® Ethernet Connection I218 uses it for master functionality. | RW |

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

Table 89. 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 90. 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 91. 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:

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

Table 92. 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 |
| 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 |

8.10.1 Interrupts

The Intel® Ethernet Connection I218 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.



8.11 Wake Up Registers

8.11.1 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 Intel® Ethernet Connection I218, 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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218, 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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218'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 Intel® Ethernet Connection I218's address from the MDI register
 - d. REGADD = 0x12 (data cycle for write)
 - e. DATA = YYYY (data to be written to the register)

8.11.2 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 cleared only when internal power on reset is de-asserted or when cleared by the software device driver.

Table 93. 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. |
| 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). |

**Table 93. Receive Control—RCTL PHY Address 01, Page 800, Register 0**

| | | | |
|-----------|------|------|---|
| RW | 7 | 0b | Receive Flow Control Enable (RFCE) Indicates that the Intel® Ethernet Connection I218 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 94. 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 Intel® Ethernet Connection I218 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 6 enable. |
| RW | 15 | 0b | FLX7 Flexible filter 7 enable. |

Table 95. 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 |



Table 95. Wake Up Filter Control—WUFC PHY Address 01, Page 800, Register 2

| | | | |
|----|----|----|---|
| 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, and a value of 0b means the filter is turned off.

Table 96. 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. |
| 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 97. 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 98. 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 99. 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 100. 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). |
| 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 101. 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 102. 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. |

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 103. 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). |



Table 104. 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 PProxy 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 I218 is ready for the code load. |

Table 105. 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 106. 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 107. 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 108. 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 109. 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 110. 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 111. 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 11 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 11 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 10. Multicast Table Array Algorithm

Table 112. 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, then the flexible filter compares the incoming data byte to the values stored in this table.

In the Intel® Ethernet Connection I218 since each address contains 16 bits, only the least significant bytes are stored in those addresses.

Table 113. 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, then the flexible filter compares the incoming data byte to the values stored in this table.

In the Intel® Ethernet Connection I218 since each address contains 16 bits, only the most significant bytes are stored in those addresses.

Note: 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 114. 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 115. 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 116. 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 117. 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 118. 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 119. 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 120. 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 121. 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 122. 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 123. 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 124. 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 125. Flexible Filter Length Table 891011—FFLT891011 PHY Address 01, Page 800, Registers 2944 + 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) for filters (8, 9, 10, 11). |
| RO | 15:11 | X | Reserved. |



Table 126. 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 127. 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. |

8.11.3 ME Wake Up Control Status Register Description

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. |
| 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 128. 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. |
| 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 129. 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 UDPSPORT0 (EN_IPUDSPORT0). When set, only packets that match the ME IPv4 filter may pass UDPSPORT0 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 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 130. 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_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/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 131. 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_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. |

Table 132. 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. |
| 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 133. 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 134. 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 135. 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 136. 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 137. 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 138. 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 139. 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 |

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

8.11.4 Proxy Controller uCode

Table 140. Proxy Micro Code—PMC PHY Address 01, Page 802, Register 0-1536Host WoL

| Attribute | Bit(s) | Initial Value | Description |
|-----------|--------|---------------|-------------|
| RW | 15:0 | 0x00 | Micro Code |

Packet Data—HWPDP PHY Address 01, Page 803, Register 0-63

| Attribute | Bit(s) | Initial Value | Description |
|-----------|--------|---------------|-------------------------------|
| RO | 15:0 | 0x00 | Host WoL captured packet data |

Table 141. Host WoL Packet Length—HWPL PHY Address 01, Page 803, Register 0-64

| Attribute | Bit(s) | Initial Value | Description |
|-----------|--------|---------------|--|
| RO | 15:0 | 0x00 | Host WoL captured packet Length, a value of 0x00 means no packet was captured. |

Table 142. Host WoL Packet Clear—HWPC PHY Address 01, Page 803, Register 0-66



8.12 LPI MMD PHY Registers

LPI MMD PHY registers are part of the Intel® Ethernet Connection I218’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 143. Intel® Ethernet Connection I218 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 | |

8.12.1 LPI MMD PHY Registers

LPI MMD PHY registers are part of the Intel® Ethernet Connection I218’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.

Intel® Ethernet Connection I218 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 | |



Note: This page intentionally left blank.



9.0 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 Intel® Ethernet Connection I218.

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 | | 0x155A 0x1559 |
| 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 PCH 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 PCH register of the PCI function. 0b = Disable. 1b = Enable. |
| 5: | Reserved | 0x0 | These bits are reserved and must be set to 0x0. |
| 3 | ENABLE_SSI D_UP | 0 | Enable SSID write once by the host |
| 2 | IO BAR Disable | 0b | IO BAR Disable. When set to 1 the LAN Controller will not advertise the request for IO space during enumeration time and BAR 2 will be reserved not used. |
| 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 I218LM PHY. Default value is 0x155A.

Note: When the I218V SKU is used in combination with certain chipset SKUs, the default value for this word is 0x1559.

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 CTRL_EXT register. When set, the Linksec logic is disabled and its clocks are gated. When cleared, the Linksec logic is enabled. |
| 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 (CTRL_EXT) 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 = I218 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 (CTRL[11]). |



| 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 *I218 Reference Schematics* for LED connection details. Mode encodings for LED outputs follow that.

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 0x03)

| Bits | Name | Default | Description |
|-------|----------|---------|------------------------|
| 15:12 | Reserved | 0x03 | 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.

§ §



10.0 Electrical and Timing Specifications

10.1 Introduction

This section describes the Intel® Ethernet Connection I218'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 |

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 Intel I218 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 Intel® Ethernet Connection I218.



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 I218 GbE PHY reference schematic. For access to documentation, contact your Intel Field Service Representative.




Figure 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 Intel® Ethernet Connection I218 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 | $I_{OL} = 9 \text{ mA}$ $V_{CC} = \text{Min}$ | -0.4 | 0 | 0.4 | Vdc |
| VOH | $I_{OH} = -9 \text{ mA}$ $V_{CC} = \text{Min}$ | 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: Figure 13 is for informational purposes only. Do not use for actual eye comparisons.




Figure 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 Intel® Ethernet Connection I218 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.




Figure 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) 12 to 20 * LOG (frequency in MHz / 80) dB (min) |
| | 40.1 MHz through 100 MHz | |
| | 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 122 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°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 [μW] max | |
| Aging | Df/f_o | ±5 ppm per year | ±5 ppm per year max | |
| Calibration Mode | | Parallel | | |
| Insulation Resistance | | | 500 [MΩ] 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 | Vmax | | 0.8 | | 1.8 | Vdc |
| Clock Amplitude | Vmin | | | | 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 |

1. Broadband peak-to-peak = 200 pS, Broadband rms = 3 pS, 12 KHz to 20 MHz rms = 1 ps




Figure 4. XTAL Timing Diagram



11.0 Time Synch (IEEE1588 and 802.1AS)

11.1 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.

11.1.1 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 11-1. 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 11-1. 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 |
| Generate a Delay_Req packet with timestamp notification in descriptor. | SW | Slave |



| Action | Responsibility | Node Role |
|---|----------------|-----------|
| 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 |

11.1.1.1 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 indicated 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.

11.1.2 HW Time sync elements

All time sync HW elements are reset to their initial values as defined in the registers section upon MAC reset.

11.1.2.1 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 **SYSTIML** and **SYSTIMH** registers. When in Master state the **SYSTIMH** and **SYSTIML** 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 11.1.2.3](#). Setting the system time is done by a direct write to the **SYSTIMH** register and fine tune setting of the **SYSTIM** register using the adjustment mechanism described in [Section 11.1.2.3](#).

Read access to the **SYSTIMH** and **SYSTIML** registers should executed in the following manner:

1. SW read register **SYSTIML**, at this stage the HW should latch the value of **SYSTIMH**.
2. SW read register **SYSTIMH**, the latched (from last read from **SYSTIML**) value should be returned by HW.



The SYSTIM timer value in **SYSTIML** and **SYSTIMH** 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 **SYSTEMH/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 cannot 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.

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).

11.1.2.2 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 11.1.2.3](#) 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 11-1. Time stamp point

11.1.2.3 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.

11.1.3 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 11-2. 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 | | |
| 14 | | correctionSubNs |
| 15 | | reserved |
| 16 | | |
| 17 | | |
| 18 | | |
| 19 | Reserved | |
| 20 | messageType | Reserved |
| 21 | Source communication technology | Source communication technology |
| 22 | Sourceuuid | Sourceuuid |
| 23 | | |
| 24 | | |
| 25 | | |
| 26 | | |
| 27 | | |
| 28 | sourceportid | sourceportid |
| 29 | sequenceId | sequenceId |
| 30 | | |
| 31 | control | control |
| 32 | reserved | logMessagePeriod |
| 33 | flags | N/A |
| 34 | | |
| 35 | | |

1. Should be all zero.

Note: Only the **BOLD** highlighted fields are of interest to the hardware.

**Table 11-3. PTP message over layer 2**

| | | | |
|---------------|-----------------|---------------|-------------|
| Ethernet (L2) | VLAN (Optional) | PTP Ethertype | PTP message |
|---------------|-----------------|---------------|-------------|

Table 11-4. 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 11.1.2.3](#). 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 11-5. 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 11-6. 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 11.1.2.3](#) 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 11.1.2.3](#).



Note: This page intentionally left blank.



12.0 Mobile Design Considerations and Guidelines

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel® Ethernet Connection I218 (PHY) shown in [Figure 5](#). 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 5. PCH/PHY Interface Connections

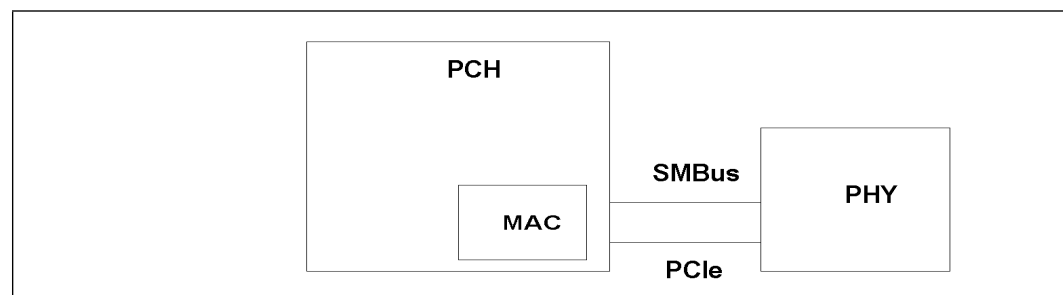




Table 4. SMBus Data Signals on the PCH

| Group | PHY Signal Name | PCH Signal Name | Description |
|-------|-----------------|-----------------|-------------|
| Data | SMB_DATA | SMLINKO_DATA | SMBus data |

Table 5. 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 6. 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# | PCIe reset |

Notes:

1. These signals come from the PCH and drive the PHY.
2. See Fig. 1-9 for connection information.

12.1 Intel® Ethernet Connection I218 Overview

The Intel® Ethernet Connection I218 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).

12.1.1 Intel® Ethernet Connection I218 Interconnects

The main interfaces for the Intel® Ethernet Connection I218 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 is 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.



12.1.2 PCIe-Based Interface

A high-speed SerDes interface 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.

12.1.2.0.1 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 6](#)). 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).

12.1.2.0.2 PCIe Operation and Channel Behavior

The Intel® Ethernet Connection I218 runs only at 1250 Mbps speed; 1/2 the Gen 1, 2.5 Gb/s PCIe frequency. Each PCIe root port in the PCH has the ability to run at 1250 Mbps. The configuration for a PCH PCIe port 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.

12.1.2.0.3 PCIe Connectivity

The PHY transmit/receive pins are output/input signals and are connected to the PCH as listed in [Table 4](#) through [Table 6](#).

12.1.2.0.4 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.

12.1.3 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.

12.1.3.0.1 SMBus Connectivity

Table 4 through Table 6 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.

12.1.4 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 Connection I218 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 |

12.1.5 Transitions between PCIe and SMBus Interfaces

12.1.5.0.1 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.



12.1.5.0.2 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.

12.2 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 Connection I218, a magnetics module and RJ-45 connector, a GbE region NVM (Non Volatile Memory) image, and a clock source.

12.2.1 General Design Considerations for the Intel® Ethernet Connection I218

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.

12.2.1.0.1 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 12.19](#) and [Section 12.35](#).

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.



12.2.1.0.2 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. The following magnetics modules are not recommended, however, they have been used successfully in previous designs:

Table 7. Magnetic Modules and Manufacturers

| Manufacturer | Part Number | Note |
|--------------|------------------|----------------------|
| Foxconn* | JFM38U1A-21C7-4F | USB stack, 8core |
| Foxconn* | JFM38U1A-7110-4F | USB2.0 stack, 8core |
| SpeedTech* | P25BPB4MFRT9 | USB2.0 stack, 8core |
| SpeedTech* | P25BFB4-RDW9 | USB2.0 stack, 12core |
| Tyco* | 1840023-1 | USB stack, 8core |

Table 8. Discrete Magnetics Modules: Manufacturers and Part Numbers

| Manufacturer | Part Number | Note |
|--------------|--------------|---------------|
| Bothhand* | GST5009LF | In Evaluation |
| Delta* | LFE9249-R-IN | In Evaluation |

12.2.1.0.3 Criteria for Integrated Magnetics Electrical Qualification

The following table gives the criteria used to qualify integrated magnetics.

**Table 9. 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 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. | 18.0 dB Min $12 - 20 * \text{LOG}(\text{Freq in MHz} / 80)$ dB Min |
| | 1.0 MHz through 29.9 MHz 30.0 MHz through 250.0 MHz 250.1 MHz through 375.0 MHz | -50.3+(8.8*(Freq in MHz / 30)) dB Max -(26 -(16.8*(LOG(Freq in MHz / 250 MHz)))) dB Max -26.0 dB Max |
| Crosstalk Isolation Discrete Modules | 1.0 MHz through 10 MHz 10.0 MHz through 100.0 MHz 100 MHz through 375.0 MHz | -50.8+(8.8*(Freq in MHz / 10)) dB Max -(26 -(16.8*(LOG(Freq in MHz / 100 MHz)))) dB Max -26.0 dB Max |
| | 1.0 MHz through 29.9 MHz 30.0 MHz through 500 MHz | -40.2+(5.3*((Freq in MHz / 30)) dB Max -(22-(14*(LOG((Freq in MHz / 250)))))) dB Max |
| Diff to CMR | 1 MHz through 270 MHz 270.1 MHz through 300 MHz 300.1 MHz through 500 MHz | -57+(38*((Freq in MHz / 270)) dB Max -17-2*((300-(Freq in MHz) / 30)) dB Max -17 dB Max |
| CM to CMR | 1500 Vrms at 50 or 60 Hz for 60 sec. or:2250 Vdc for 60 seconds | Minimum |

12.2.2 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 Field Service representative.

12.2.3 LED

The PHY has three LED outputs that can be configured via the NVM. The hardware configuration is shown in [Figure 6](#).


Refer to the Intel® Ethernet Connection I218 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 10. 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 6. LED Hardware Configuration



12.2.3.0.1 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 K.Ω 1% pull-down resistor to ground and then place it no more than one half inch (0.5”) away from the PHY.

12.2.3.0.2 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.

12.2.4 Exposed Pad* (e-Pad) Design and SMT Assembly Guide

12.2.4.0.1 Overview

This section provides general information about ePAD 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 ePADs are very popular due to their low cost. Note that this section only provides basic information and references in regards to the ePAD. It is recommended that each customer consult their fab and assembly house to obtain more details on how



to implement the ePAD 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.

12.2.4.0.2 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 package size, a via-in-pad configuration must be used [Figure 7](#) and [Figure 8](#) are general guidelines see [Figure 9](#) for -specific via-in-pad thermal pattern recommendations.

Figure 7. Typical ePAD Land Pattern






Figure 8. 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 9. Recommended Thermal Via Patterns



12.2.4.0.3 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.



12.2.4.0.4 Stencil Design

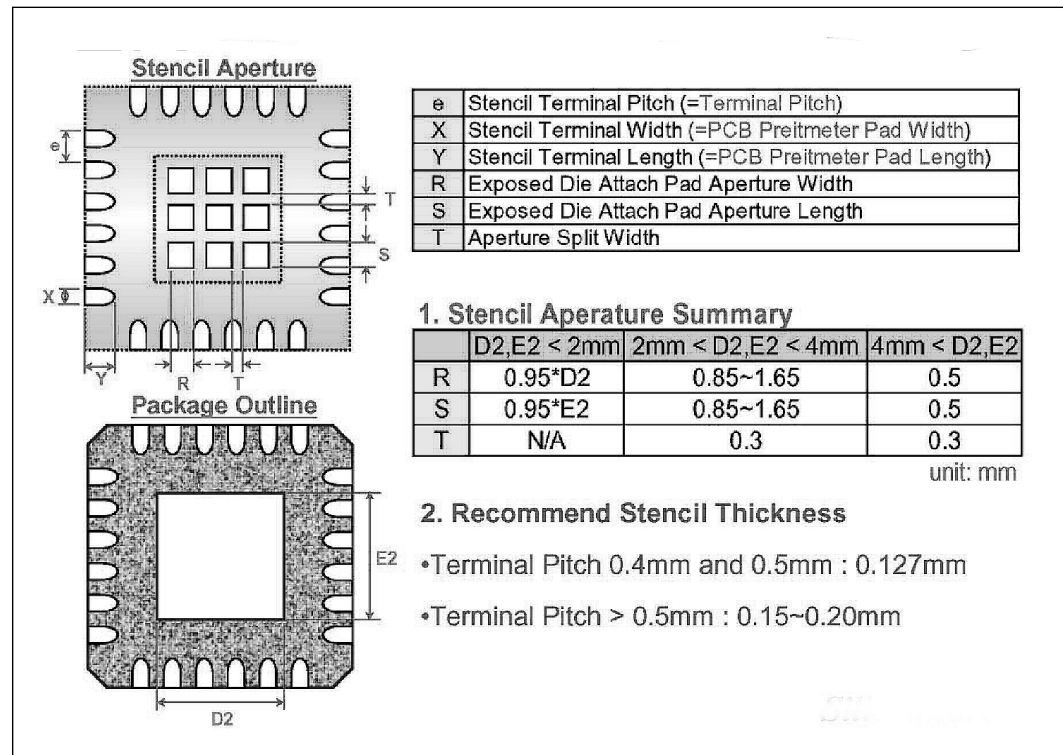
For maximum thermal/electrical performance, it is required that the exposed pad/slug on the package be soldered to the land pattern on the PCB. This can be achieved by applying solder paste on both the pattern for lead attachment as well as on the land pattern for the exposed pad. While for standard (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 10. Stencil Design Recommendation



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 Controller I218 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 Aperture's/e-PAD area is 81% of the exposed landing area (solder mask opening).

Note: This information is intended only as general guidance. Please consult with the manufacturer to confirm the final design meets requirements.

12.2.4.0.5 Assembly Process Flow

The following figure below shows the typical process flow for mounting packages to the PCB.

Figure 11. Assembly Flow

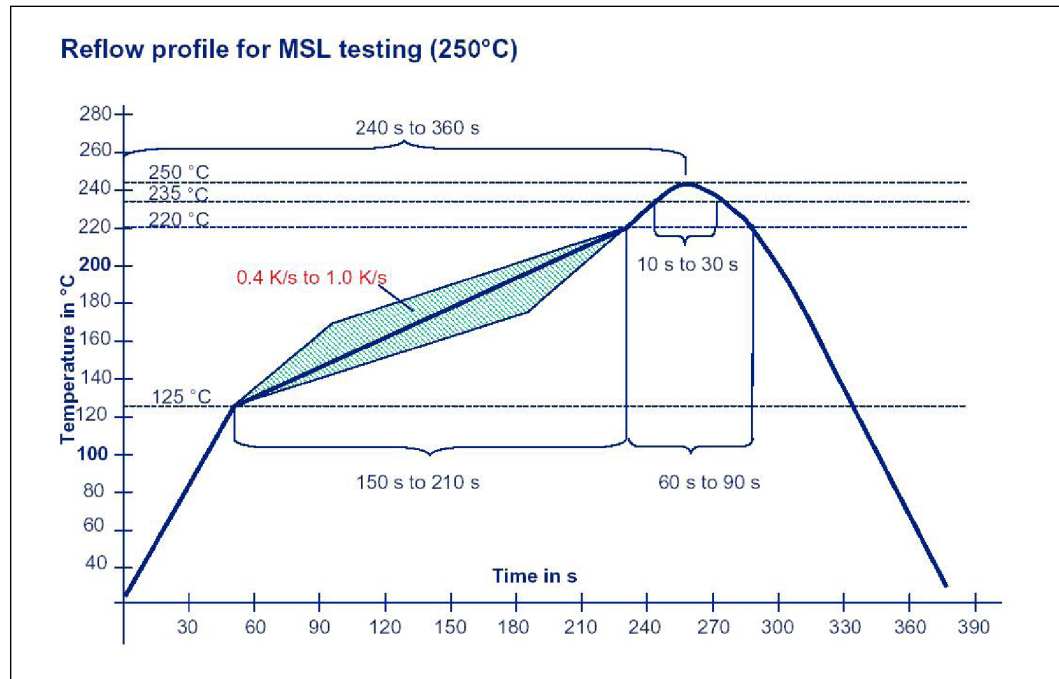


12.2.4.0.6 Reflow Guidelines

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 12. Typical Profile Band

**Notes:**

1. Preheat: 125 °C -220 °C, 150 - 210 s at 0.4 k/s to 1.0 k/s
2. Time at T > 220 °C: 60 - 90 s
3. Peak Temperature: 245-250 °C
4. Peak time: 10 - 30 s
5. Cooling rate: <= 6 k/s
6. Time from 25 °C to Peak: 240 - 360 s
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.


12.3 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 Connection Reference Schematic*.

Refer to [Section 12.6](#) for PCI Express Routing Guidelines.

Figure 13. Single Solution Interconnect



Notes:

1. For latest PHY schematic connection recommendations, refer to the *Intel® Ethernet Connection I218 Reference Schematic*, and the *Intel® Ethernet Connection I218 Schematic and Layout Checklist*, available through 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 Connection I218.
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 Connection I218 respectively.
4. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.

12.4 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 12.1.3](#) for additional SMBus design considerations.

12.5 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.

12.6 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.

12.7 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.


12.7.1 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. 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 Connection I218 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 Connection I218 should be placed greater than 250 mils from the board edge.
- If the connector or integrated magnetics module is not shielded, Intel® Ethernet Connection I218 should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the Intel® Ethernet Connection I218 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 15. PLC Placement: At Least One Inch from I/O Backplane





Figure 16. Effect of LAN Device Placed Too Close to an Rj-45 Connector or Chassis Opening



12.8 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.

12.9 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 to aid with MDI topology design. For access to documentation contact your Intel 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 100Ω 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 105Ω to 110Ω 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 10Ω , when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 11. 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 17 |
| 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 12 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 1-8 for trace length information.

Table 12. 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 17. MDI Trace Geometry



12.10 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 18. MDI Differential Trace Geometry



Note: Similar topology applies to MDI routing from the Intel® Ethernet Connection I218 to the dock RJ45 connector.

12.11 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.

12.12 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.

12.13 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.

12.14 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 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 19. Trace Transitioning Layers and crossing Plane Splits




Figure 20. Via Connecting GND to GND





Figure 21. **Stitching Capacitor between Vias Connecting GND to GND**



12.15 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.

12.16 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 15-19 illustrates the split plane layout for a discrete magnetics module. Capacitors are used to interconnect chassis ground and signal ground.

Figure 22 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.



Figure 22. Ideal Ground Split Implementation

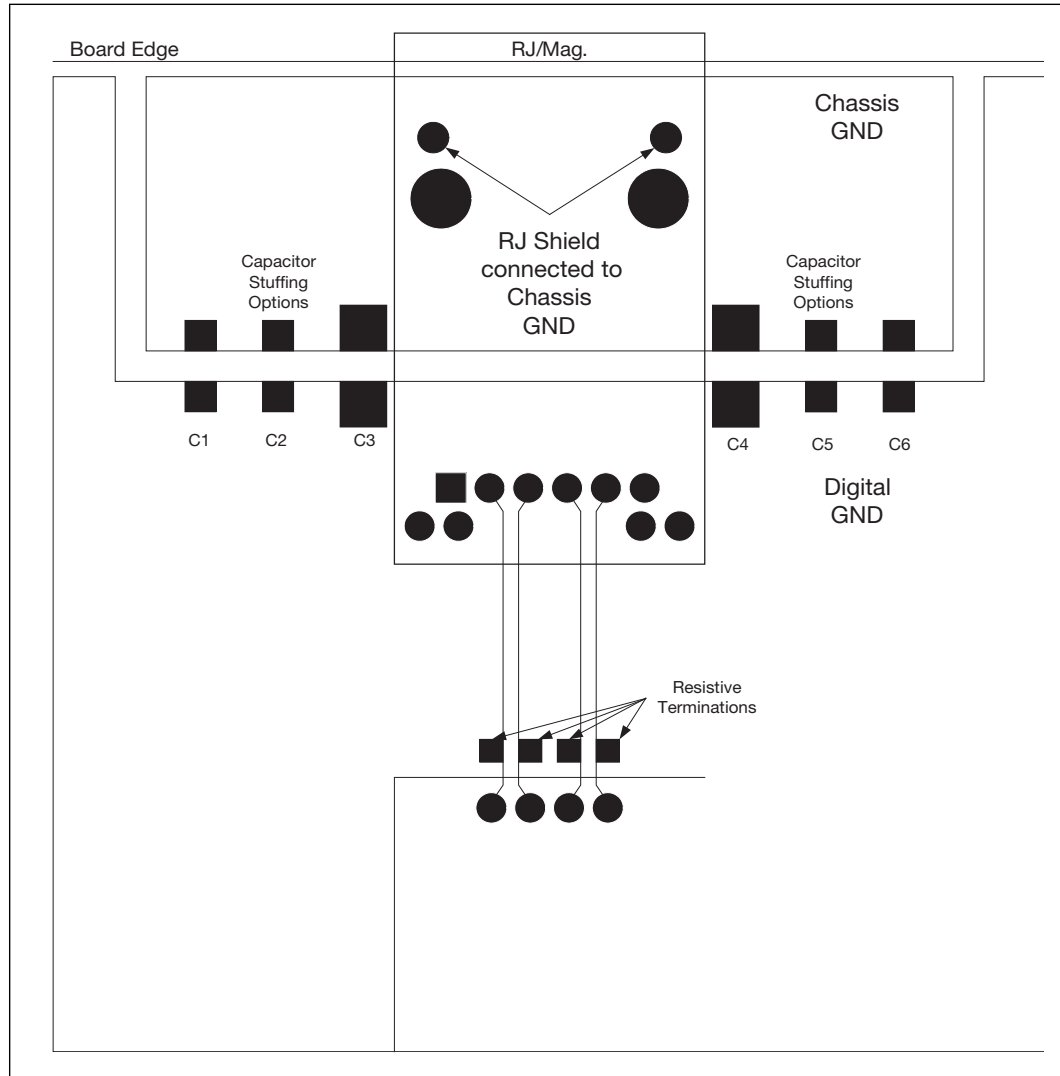


Table 13. Capacitor Stuffing Option Recommended Values


| 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 23 provides an example of this.

Figure 23. Ground Layout with USB



12.17 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.

12.18 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.



12.19 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 12.23](#).

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.

12.20 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 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.

12.21 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.



12.22 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.

12.23 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 14. Crystal Manufacturers and Part Numbers

| Manufacturer | Part Number | Note |
|-----------------------|-----------------------------|------------------------------|
| KDS America* | DSX321G, 1B/C/N/H225000CC0M | Small size, SMD, 3.2mmx2.5mm |
| NDK America Inc.* | 41CD25.0F1303018 | HC-49S type, SMD |
| River* | FCX-04-25MJ90141 | Small size, SMD, 3.2mmx2.5mm |
| TXC Corporation, USA* | 7V25020001 | Small size, SMD, 3.2mmx2.5mm |
| TXC Corporation, USA* | 7V25020004 | Small size, 3.2mmx2.5mm |
| TXC Corporation, USA* | 9C25000008 | HC-49S type, SMD |

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.

12.24 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.



12.25 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.

12.26 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.

12.27 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.


12.28 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 24 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 24. Thermal Oscillator Circuit



12.29 Load Capacitance

The formula for crystal load capacitance is as follows:

$$C_L = \frac{(C1 \cdot C2)}{(C1 + C2)} + C_{stray}$$

where C1 = C2 = 22pF (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}.

12.30 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.

12.31 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 50 Ω or better should be used.



12.32 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 μW maximum, then the crystal drive level must be at least 200 μW . So, a 500 μW crystal is sufficient, but a 100 μW crystal is not.

12.33 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.

12.34 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.

12.34.1 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, customers 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.

12.34.2 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.

12.34.3 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.

12.35 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 (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 3.3 V 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 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):

| Manufacturer | Part Number | Note |
|--------------|---------------------|-------------|
| Kyocera* | K30-3C0-SE-25.0000M | 5mmx3.2mm |
| MtronPTI* | M214TCN25.0000MHz | 7mmx5mm |
| TXC* | 7C25000230 | 5mmx3.2mm |
| TXC* | 7X25080001 | 3.2mmx2.5mm |

12.36 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 μF 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 the latest PHY schematic connection recommendations, refer to the Intel® Ethernet Connection I218 Reference Schematic and the Intel® Ethernet Connection I218 Schematic and Layout Checklist, available through your Intel representative.

12.37 LAN Switch

The following table lists LAN switches that can be used with the Intel® Ethernet Connection I218. Note that no particular LAN switch is recommended.

Table 15. LAN Switch: Manufacturers and Part Numbers

| Manufacturer | Part No. | Note |
|--------------------|-------------|--------------|
| Pericom* | PI3L500ZFEX | Limited ESD |
| Pericom* | PI3L500-AZ | Enhanced ESD |
| Texas Instruments* | TS3L500AE | Enhanced ESD |

12.38 Troubleshooting Common Physical Layout Issues

The following lists 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-Ω 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.

12.39 Power Delivery

The Intel® Ethernet Connection I218 requires a 3.3V power rail and a 0.93V (Core) power rail. The internal 3.3V power rail is brought out for decoupling. [Figure 25](#) 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 3.3Vdc 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 Connection I218 uses the integrated SVR (iSVR). When laying out the circuit 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.


While Intel does not endorse vendors or specific components, design compatibility has been verified for the connectors on the following list.

Table 16. 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 25. Intel® I218 Gbe PHY Power Delivery Diagram



12.40 Intel® Ethernet Controller I218 Power Sequencing

The Intel® Ethernet Controller I218 does not require any power sequencing between the 3.3V and Core power rails. See the reference schematic for details.

For latest PHY schematic connection recommendations, refer to the Intel® Ethernet Controller I218 reference schematic or contact your Intel Field Service Representative.



13.0 Non-Mobile Design Considerations and Guidelines

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel®Ethernet Connection I218 (PHY) shown in [Figure 26](#). 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 26. PCH/PHY Interface Connections




Table 17. SMBus Data Signals on the PCH

| Group | PHY Signal Name | PCH Signal Name | Description |
|-------|-----------------|-----------------|-------------|
| Data | SMB_DATA | SMLINK0_DATA | SMBus data |

**Table 18. 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 19. 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 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 19](#)). This PHY output should be tied to the PCH input and pulled up with a 10-k 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 runs only at 1250 Mbps speed; 1/2 the Gen 1, 2.5 Gb/s PCIe frequency. Each PCIe root port in the PCH has the ability to run at 1250 Mbps. The configuration for a PCH PCIe port 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 17](#) through [Table 19](#).

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 17 through Table 19 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 I218 GbE PHY 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 I218 GbE PHY, 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 specifications section of this datasheet for specific electrical requirements that the magnetics need to meet.

While Intel does not endorse vendors or specific components, design compatibility has been verified for the connectors on the following list.

Table 20. 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 21. 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 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. | 18.0 dB Min $12 - 20 * \text{LOG}(\text{Freq in MHz} / 80)$ dB Min |
| Crosstalk Isolation Discrete Modules | 1.0 MHz through 29.9 MHz 30.0 MHz through 250.0 MHz 250.1 MHz through 375.0 MHz | -50.3+(8.8*(Freq in MHz / 30)) dB Max -(26 -(16.8*(LOG(Freq in MHz / 250 MHz)))) dB Max -26.0 dB Max |
| Crosstalk Isolation Integrated Modules (Proposed) | 1.0 MHz through 10 MHz 10.0 MHz through 100.0 MHz 100 MHz through 375.0 MHz | -50.8+(8.8*(Freq in MHz / 10)) dB Max -(26 -(16.8*(LOG(Freq in MHz / 100 MHz)))) dB Max -26.0 dB Max |
| Diff to CMR | 1 MHz through 29.9 MHz 30.0 MHz through 500 MHz | -40.2+(5.3*((Freq in MHz / 30)) dB Max -(22-(14*(LOG((Freq in MHz / 250)))) dB Max |
| CM to CMR | 1 MHz through 270 MHz 270.1 MHz through 300 MHz 300.1 MHz through 500 MHz | -57+(38*((Freq in MHz / 270)) dB Max -17-2*((300-(Freq in MHz) / 30) dB Max -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.

Intel has an MS-DOS* software utility called Eeupdate 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 27.


Refer to the Intel I218 GbE PHY Reference Schematic for default LED color based on reference design.

Refer to Section 8 of this 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 22. 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 27. 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 3.01 kΩ 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 PCHCougar Point-M.



13.2.5 Exposed Pad* (e-Pad) Design and SMT Assembly Guide

13.2.5.0.1 Overview

This section provides general information about ePAD 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 ePADs are very popular due to their low cost. Note that this section only provides basic information and references in regards to the ePAD. It is recommended that each customer consult their fab and assembly house to obtain more details on how to implement the ePAD 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 28](#) and [Figure 29](#) are general guidelines see [Figure 30](#) for -specific via-in-pad thermal pattern recommendations.

Figure 28. Typical ePAD Land Pattern






Figure 29. 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 30. Recommended Thermal Via Pattern



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. 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 31. Stencil Design Recommendation



Important General Guidelines:

- The Stencil Aperture Openings of the e-PAD must not go outside 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 I218 e-PAD has $D2=E2=3\text{mm}$. Therefore, the stencil design can have only four aperture openings for the e-PAD. Achieve this by setting $R=S=1.35\text{mm}$ and $T=0.3$. With this arrangement, the Aperture/e-PAD area is 81% of the exposed landing area (solder mask opening).

This information is only intended as general guidance. Please consult the manufacturer to confirm that the final design meets their 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 32. 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 33. Typical Profile Band



1. Preheat: 125 °C -220 °C, 150 - 210 s at 0.4 k/s to 1.0 k/s
2. Time at T > 220 °C: 60 - 90 s
3. Peak Temperature: 245-250 °C
4. Peak time: 10 - 30 s
5. Cooling rate: <= 6 k/s
6. Time from 25 °C to Peak: 240 – 360 s
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® I218 GbE PHY Reference Schematic.

Refer to [Section 13.6](#) for PCI Express Routing Guidelines.

Figure 34. Single Solution Interconnect



Notes:

1. For latest PHY schematic connection recommendations, refer to the Intel® I218 GbE PHY Reference Schematic, and the Intel® I218 GbE PHY Schematic and Layout Checklist, both available through your Intel representative.
2. Any free PCIe lanes [8:1] on the PCH can be used to connect the PCIe interface on the Intel I218 GbE 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 Intel I218 GbE PHY 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 35. 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 I218 GbE PHY 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 PHY should be placed greater than 250mils from the board edge.
- If the connector or integrated magnetics module is not shielded, the I218 GbE PHY should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the PHY 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 36. PLC Placement: At Least One Inch from I/O Backplane





Figure 37. Effect of LAN Device Placed Too Close to an 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.

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 100Ω 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 105Ω to 110Ω 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 10Ω , when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 23. 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 38 |
| 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 Ω ±15%. If it's not feasible (due to board stack-up) it is recommended that board designers use a 95 Ω ±10% target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95 Ω. 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 85 Ω.
3. Simulation shows 80 Ω differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90 Ω.
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 16-8 for trace length information
7. 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 24 for examples of microstrip trace geometries for common circuit board materials.
8. If a LAN switch is not used, then the maximum trace length is 4 inches (102 mm).

Table 24. 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 100 Ω 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 38. 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 39. MDI Differential Trace Geometry



Note: Similar topology applies to MDI routing from the I218 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 40. Trace Transitioning Layers and crossing Plane Splits

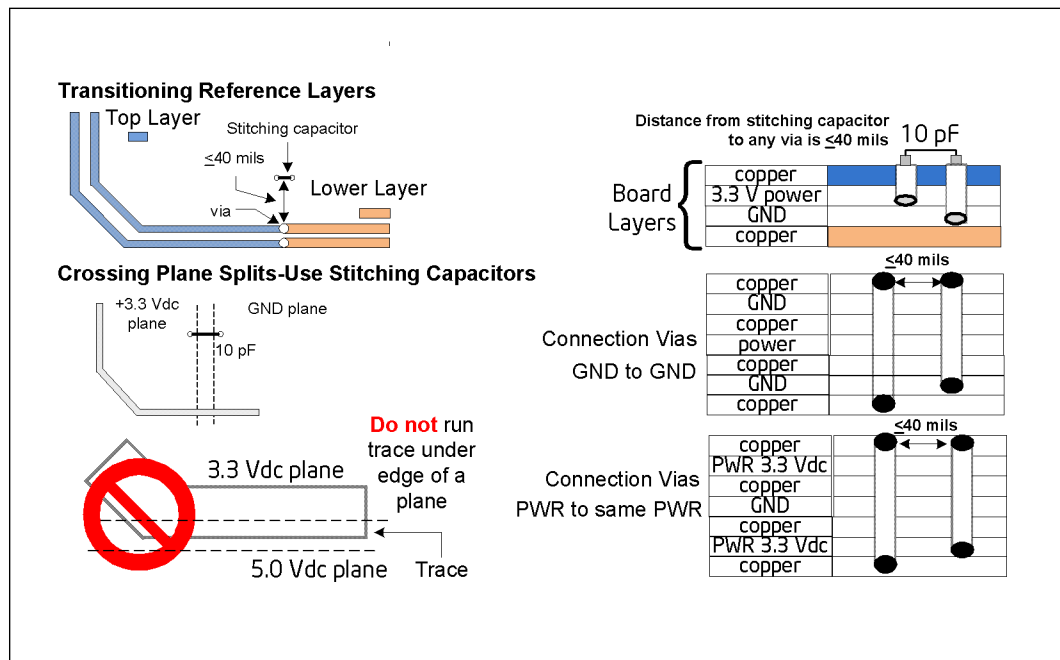


Figure 41. Via Connecting GND to GND





Figure 42. **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.

Figure 16-18 illustrates the split plane layout for a discrete magnetics module. Capacitors are used to interconnect chassis ground and signal ground.

Figure 43 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.



Figure 43. Ideal Ground Split Implementation

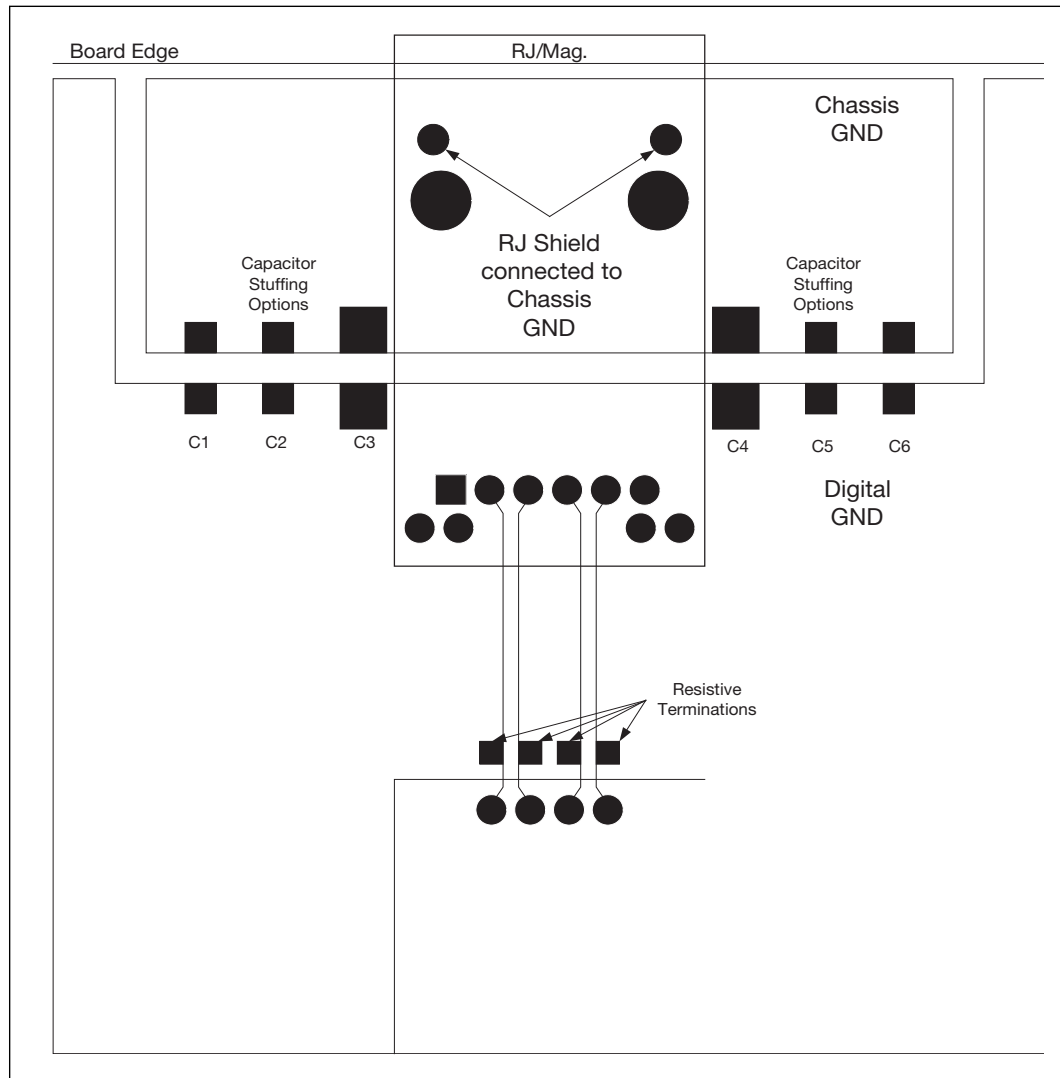


Table 25. Capacitor Stuffing Option Recommended Values


| 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 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 44 provides an example of this.

Figure 44. 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 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 26. Crystal Manufacturers and Part Numbers

| Manufacturer | Part No. | Note |
|--------------|-----------------------------|-------------------------|
| TXC* | 9C25000008 | HC-49S SMD |
| NDK* | 41CD25.0F1303018 | HC-49S SMD |
| TXC* | 7V25020004 | Small SMD 3.2mm x 2.5mm |
| 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 45 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 45. 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 = 22pF (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 50 Ω 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, customers 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 (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, 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 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 27. 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 μF 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® I218 GbE PHY Reference Schematic* and the *Intel® I218 GbE PHY Schematic and Layout Checklist*, available through your Intel representative.

13.37 LAN Switch

The following table lists LAN switches that can be used with the I218. Note that no particular LAN switch is recommended:

Table 28. 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 100- Ω 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.

13.39 Power Delivery

The Intel I218 GbE PHY requires a 3.3V power rail and a 0.9V (Core) power rail. The internal 3.3V power rail is brought out for decoupling. [Figure 46](#) 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 μ F 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 μ F 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 I218 GbE PHY uses the integrated SVR (iSVR). When laying out the circuit 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.)


While Intel does not endorse vendors or specific components, design compatibility has been verified for the connectors on the following list.

Table 29. 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 46. Intel® I218 GbE PHY Power Delivery Diagram





Note: For latest PHY schematic connection recommendations, contact your Intel Field Service Representative for the Intel I218 GbE PHY reference schematic,.

13.40 Power Sequencing

The PHY does not require any power sequencing between the 3.3V and Core power rails. See the datasheet and reference schematic for details.



14.0 Checklists

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I218 Design and Board Layout Checklists.



Note: This page intentionally left blank.

§ §



15.0 Models

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I218 model.



Note: This page intentionally left blank.

§ §



16.0 Reference Schematics

Contact your Intel Field Service Representative for access to the Intel® Ethernet Connection I218 reference schematics.



Note: This page intentionally left blank.

§ §

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Intel:](#)

[WGI218LM S LK3B](#) [WGI218V S LK3D](#) [WGI218LM SLJK3A](#) [WGI218V S LK3C](#)

Компания «Океан Электроники» предлагает заключение долгосрочных отношений при поставках импортных электронных компонентов на взаимовыгодных условиях!

Наши преимущества:

- Поставка оригинальных импортных электронных компонентов напрямую с производств Америки, Европы и Азии, а так же с крупнейших складов мира;
- Широкая линейка поставок активных и пассивных импортных электронных компонентов (более 30 млн. наименований);
- Поставка сложных, дефицитных, либо снятых с производства позиций;
- Оперативные сроки поставки под заказ (от 5 рабочих дней);
- Экспресс доставка в любую точку России;
- Помощь Конструкторского Отдела и консультации квалифицированных инженеров;
- Техническая поддержка проекта, помощь в подборе аналогов, поставка прототипов;
- Поставка электронных компонентов под контролем ВП;
- Система менеджмента качества сертифицирована по Международному стандарту ISO 9001;
- При необходимости вся продукция военного и аэрокосмического назначения проходит испытания и сертификацию в лаборатории (по согласованию с заказчиком);
- Поставка специализированных компонентов военного и аэрокосмического уровня качества (Xilinx, Altera, Analog Devices, Intersil, Interpoint, Microsemi, Actel, Aeroflex, Peregrine, VPT, Syfer, Eurofarad, Texas Instruments, MS Kennedy, Miteq, Cobham, E2V, MA-COM, Hittite, Mini-Circuits, General Dynamics и др.);

Компания «Океан Электроники» является официальным дистрибьютором и эксклюзивным представителем в России одного из крупнейших производителей разъемов военного и аэрокосмического назначения «JONHON», а так же официальным дистрибьютором и эксклюзивным представителем в России производителя высокотехнологичных и надежных решений для передачи СВЧ сигналов «FORSTAR».



JONHON

«JONHON» (основан в 1970 г.)

Разъемы специального, военного и аэрокосмического назначения:

(Применяются в военной, авиационной, аэрокосмической, морской, железнодорожной, горно- и нефтедобывающей отраслях промышленности)

«FORSTAR» (основан в 1998 г.)

ВЧ соединители, коаксиальные кабели,
кабельные сборки и микроволновые компоненты:

(Применяются в телекоммуникациях гражданского и специального назначения, в средствах связи, РЛС, а так же военной, авиационной и аэрокосмической отраслях промышленности).



Телефон: 8 (812) 309-75-97 (многоканальный)

Факс: 8 (812) 320-03-32

Электронная почта: ocean@oceanchips.ru

Web: <http://oceanchips.ru/>

Адрес: 198099, г. Санкт-Петербург, ул. Калинина, д. 2, корп. 4, лит. А