

# Intel® Ethernet Connection I219

## Datasheet

---

### Client Connectivity Division (CCD)

#### PRODUCT FEATURES

##### General

- 10 BASE-T IEEE 802.3 specification compliance
- 100 BASE-TX IEEE 802.3 specification compliance
- 1000 BASE-T IEEE 802.3 specification compliance
- 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)
- Network proxy/ARP Offload support
- Up to 32 programmable filters
- No support for Gb/s half-duplex operation

##### 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 Link-Up (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

Revision 2.02  
May 2015



## LEGAL

---

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting [www.intel.com/design/literature.htm](http://www.intel.com/design/literature.htm).

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

© 2015 Intel Corporation.



## Revision History

Revision	Date	Comments
2.02	May 2015	Updated title page.
2.01	April 2015	Revised sections: <ul style="list-style-type: none"> <li>• 10.3.1.9</li> <li>• 10.3.1.14</li> </ul> Revised table 12-12, 12-13, 13-14, and 13-15.
2.0	February 2015	<ul style="list-style-type: none"> <li>• Initial Release (Intel Public).<sup>1</sup></li> </ul>
0.9	January 2015	<ul style="list-style-type: none"> <li>• Updated Title Page (General Product Features).</li> <li>• Flexible Filters changed from 16 to 32 (Section 7.4.1.4.1.8).</li> <li>• Added Non-Volatile Memory (NVM) map (Section 10.0).</li> </ul>
0.7	July 2014	Revised sections: <ul style="list-style-type: none"> <li>• 11.2.1</li> <li>• 11.2.1.2</li> <li>• 11.3</li> <li>• 11.20.3</li> <li>• 11.20.15</li> <li>• 11.21</li> <li>• 11.23</li> <li>• 12.2</li> <li>• 12.4</li> <li>• 12.2.1.2</li> <li>• 12.21.3</li> <li>• 12.21.15</li> <li>• 12.23</li> </ul>
0.5	March 3, 2014	Initial release (Intel Confidential)

1. There have been no releases between revision 0.9 and revision 2.0.



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



## Contents

---

<b>1.0</b>	<b>Introduction</b>	11
1.1	Overview	11
1.2	Main Flows	12
1.3	References	12
<b>2.0</b>	<b>Interconnects</b>	13
2.1	Introduction	13
2.2	PCIe-Based	13
2.2.1	PCIe Interface Signals	14
2.2.2	PCIe Operation and Channel Behavior	14
2.3	SMBus	19
2.3.1	Overview	19
2.3.2	SMBus Pass Trough (PT)	23
2.3.3	Slave Transactions	25
2.3.4	Master Transactions	29
2.4	Transitions Between SMBus and PCIe interfaces	33
2.4.1	Switching from SMBus to PCIe	33
2.4.2	Switching from PCIe to SMBus	33
<b>3.0</b>	<b>Pin Interface</b>	35
3.1	Pin Assignment	35
3.1.1	Signal Type Definitions	35
3.1.2	PCIe Interface Pins (8)	36
3.1.3	SMBus Interface Pins (2)	36
3.1.4	Miscellaneous Pins (3)	36
3.1.5	PHY Pins (14)	37
3.1.6	Testability Pins (5)	38
3.1.7	Power and Power Control Pins (16)	38
3.2	Pinout	39
<b>4.0</b>	<b>Package</b>	41
4.1	Package Type and Mechanical	41
4.2	Package Electrical and Thermal Characteristics	42
4.3	Power and Ground Requirements	42
4.4	Ball Mapping	43
<b>5.0</b>	<b>Initialization</b>	45
5.1	Power Up	45
5.2	Reset Operation	46
5.3	Timing Parameters	47
5.3.1	Timing Requirements	47
5.3.2	Timing Guarantees	47
<b>6.0</b>	<b>Power Management and Delivery</b>	49
6.1	Power Information	49
6.2	Power Delivery	50
6.2.1	Core Vdc Supply	50
6.3	Power Management	50
6.3.1	Global Power States	50
6.4	I219 Ultra Low Power (ULP)	55
6.5	Off Board LAN Connected Device (OBLCD)	55



- 6.5.1 Indicating /Sensing OBLCD Disconnect Event..... 55
- 6.5.2 Indicating/Sensing OBLCD Connection ..... 56
- 7.0 Device Functionality ..... 57**
  - 7.1 Tx Flow ..... 57
  - 7.2 Rx Flow ..... 58
  - 7.3 Flow Control ..... 58
    - 7.3.1 MAC Control Frames and Reception of Flow Control Packets ..... 59
    - 7.3.2 Transmitting PAUSE Frames ..... 60
  - 7.4 Wake Up ..... 61
    - 7.4.1 Host Wake Up..... 62
    - 7.4.2 Management Engine (ME) Wake Up ..... 68
  - 7.5 Network Proxy Functionality ..... 72
    - 7.5.1 Introduction ..... 72
    - 7.5.2 Network Proxy Activation ..... 72
    - 7.5.3 IPv4 Proxy - ARP..... 73
    - 7.5.4 IPv6 Proxy - Neighbor Discovery ..... 74
    - 7.5.5 Multicast Listener Discovery Support..... 78
  - 7.6 Loopback ..... 81
- 8.0 PHY Functionality ..... 83**
  - 8.1 Reverse Auto-Negotiation ..... 83
- 9.0 Programmer’s Visible State ..... 87**
  - 9.1 Terminology ..... 87
  - 9.2 MDIO Access ..... 88
  - 9.3 Addressing ..... 88
  - 9.4 Registers Byte Ordering ..... 89
  - 9.5 Register Definitions ..... 90
    - 9.5.1 Address Map ..... 90
    - 9.5.2 PHY Registers ..... 95
    - 9.5.3 Port Control Registers..... 115
    - 9.5.4 Statistics Registers..... 119
    - 9.5.5 PCIe Registers ..... 122
    - 9.5.6 LPI Registers ..... 125
    - 9.5.7 ULP Registers ..... 128
    - 9.5.8 General Registers..... 131
    - 9.5.9 Wake Up Registers ..... 135
    - 9.5.10 Proxy Controller uCode ..... 166
    - 9.5.11 Host WoL Packet ..... 167
    - 9.5.12 LPI MMD PHY Registers ..... 168
- 10.0 Non-Volatile Memory (NVM) ..... 169**
  - 10.1 Introduction ..... 169
  - 10.2 NVM Programming Procedure Overview ..... 169
  - 10.3 LAN NVM Format and Contents ..... 171
    - 10.3.1 Hardware Accessed Words..... 173
    - 10.3.2 Software Accessed Words..... 183
    - 10.3.3 Basic Configuration Software Words ..... 189
- 11.0 Electrical and Timing Specifications ..... 191**
  - 11.1 Introduction ..... 191
  - 11.2 Operating Conditions ..... 191
    - 11.2.1 Absolute Maximum Ratings..... 191



11.2.2	Recommended Operating Conditions.....	191
11.3	Power Delivery .....	192
11.3.1	Voltage Regulator Power Supply Specifications.....	192
11.3.2	SVR Specification (Internal).....	193
11.3.3	Power On/Off Sequence .....	193
11.3.4	Power Delivery Schematic Drawing.....	194
11.4	I/O DC Parameter .....	195
11.4.1	3.3 Vdc I/O (Open Drain).....	195
11.4.2	3.3 Vdc I/O .....	195
11.4.3	Input Buffer Only .....	196
11.4.4	PCIe DC/AC Specifications.....	196
11.5	Discrete/Integrated Magnetics Specifications .....	199
11.6	Mechanical .....	199
11.7	Oscillator/Crystal Specifications .....	200
<b>12.0</b>	<b>Mobile Design Considerations and Guidelines .....</b>	<b>203</b>
12.1	I219 Overview .....	204
12.1.1	I219 Interconnects.....	204
12.1.2	PCIe-Based Interface.....	205
12.1.3	SMBus Interface .....	206
12.1.4	PCIe and SMBus Modes.....	206
12.1.5	Transitions Between PCIe and SMBus Interfaces.....	207
12.2	Platform LAN Design Guidelines .....	207
12.2.1	General Design Considerations for the Intel® Ethernet Connection I219 .....	208
12.2.2	NVM for PHY Implementations .....	210
12.2.3	LED.....	210
12.2.4	Exposed Pad* (e-Pad) Design and SMT Assembly Guide .....	211
12.3	PCH-SMBus/PCIe LOM Design Guidelines .....	216
12.4	SMBus Design Considerations .....	217
12.5	General Layout Guidelines .....	217
12.6	Layout Considerations .....	218
12.7	Guidelines for Component Placement .....	218
12.7.1	PHY Placement Recommendations .....	218
12.8	MDI Differential-Pair Trace Routing for LAN Design .....	220
12.9	Signal Trace Geometry .....	220
12.10	Trace Length and Symmetry .....	223
12.11	Impedance Discontinuities .....	224
12.12	Reducing Circuit Inductance .....	224
12.13	Signal Isolation .....	224
12.14	Power and Ground Planes .....	225
12.15	Traces for Decoupling Capacitors .....	227
12.16	Ground Planes Under a Magnetics Module .....	227
12.17	Light Emitting Diodes .....	229
12.18	Considerations for Layout .....	229
12.19	Frequency Control Device Design Considerations .....	230
12.20	Crystals and Oscillators .....	230
12.20.1	Quartz Crystal .....	230
12.20.2	Fixed Crystal Oscillator .....	231
12.20.3	Crystal Selection Parameters .....	231
12.20.4	Vibrational Mode .....	231
12.20.5	Nominal Frequency.....	232



12.20.6	Frequency Tolerance .....	232
12.20.7	Temperature Stability and Environmental Requirements .....	232
12.20.8	Calibration Mode .....	232
12.20.9	Load Capacitance .....	233
12.20.10	Shunt Capacitance .....	233
12.20.11	Equivalent Series Resistance .....	234
12.20.12	Drive Level.....	234
12.20.13	Aging .....	234
12.20.14	Reference Crystal.....	234
12.20.15	Oscillator Support .....	236
12.20.16	Oscillator Placement and Layout Recommendations .....	236
12.21	LAN Switch .....	237
12.22	Troubleshooting Common Physical Layout Issues .....	237
12.23	Power Delivery .....	238
12.24	I219 Power Sequencing .....	239
<b>13.0</b>	<b>Non-Mobile Design Considerations and Guidelines .....</b>	<b>241</b>
13.1	PHY Overview .....	242
13.1.1	PHY Interconnects.....	242
13.1.2	PCIe-Based Interface.....	243
13.1.3	SMBus Interface .....	244
13.1.4	PCIe and SMBus Modes.....	244
13.1.5	Transitions Between PCIe and SMBus Interfaces.....	245
13.2	Platform LAN Design Guidelines .....	246
13.2.1	General Design Considerations for PHYs .....	246
13.2.2	NVM for PHY Implementations .....	249
13.3	LAN Switch .....	249
13.3.1	LED.....	250
13.3.2	Exposed Pad* (e-Pad) Design and SMT Assembly Guide .....	251
13.4	PCH-SMBus/PCIe LOM Design Guidelines .....	256
13.5	SMBus Design Considerations .....	257
13.6	General Layout Guidelines .....	257
13.7	Layout Considerations .....	258
13.8	Guidelines for Component Placement .....	258
13.8.1	PHY Placement Recommendations .....	258
13.9	MDI Differential-Pair Trace Routing for LAN Design .....	260
13.10	Signal Trace Geometry .....	260
13.11	Trace Length and Symmetry .....	263
13.12	Impedance Discontinuities .....	264
13.13	Reducing Circuit Inductance .....	264
13.14	Signal Isolation .....	264
13.15	Power and Ground Planes .....	265
13.16	Traces for Decoupling Capacitors .....	267
13.17	Ground Planes Under a Magnetics Module .....	267
13.18	Light Emitting Diodes .....	269
13.19	Considerations for Layout .....	269
13.20	Frequency Control Device Design Considerations .....	270
13.21	Crystals and Oscillators .....	270
13.21.1	Quartz Crystal .....	270
13.21.2	Fixed Crystal Oscillator .....	271
13.21.3	Crystal Selection Parameters .....	271



13.21.4	Vibrational Mode .....	271
13.21.5	Nominal Frequency.....	272
13.21.6	Frequency Tolerance .....	272
13.21.7	Temperature Stability and Environmental Requirements .....	272
13.21.8	Calibration Mode .....	272
13.21.9	Load Capacitance .....	273
13.21.10	Shunt Capacitance .....	273
13.21.11	Equivalent Series Resistance.....	274
13.21.12	Drive Level.....	274
13.21.13	Aging .....	274
13.21.14	Reference Crystal.....	274
13.21.15	Oscillator Support .....	276
13.21.16	Oscillator Placement and Layout Recommendations .....	276
13.22	Troubleshooting Common Physical Layout Issues .....	277
13.23	Power Delivery .....	277
13.24	I219 Power Sequencing .....	278



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



## 1.0 Introduction

### 1.1 Overview

The Intel® Ethernet Connect I219 (I219) is a single-port Gigabit Ethernet Physical Layer Transceiver (PHY). It connects to an integrated Media Access Controller (MAC) through a dedicated interconnect. The I219 supports operation at 10/100/1000 Mb/s data rates. The PHY circuitry provides a standard IEEE 802.3 Ethernet interface for 10BASE-T, 100BASE-TX, and 1000BASE-T applications (802.3, 802.3u, and 802.3ab). The I219 also supports the Energy Efficient Ethernet (EEE) 802.az specification.

The I219 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 I219 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 I219 PCIe interface is not PCIe compliant. It operates at half of the PCI Express\* (PCIe\*) Specification v1.1 (2.5 GT/s) speed. In this datasheet the term PCIe-based is interchangeable with PCIe. There is no design layout differences between normal PCIe and the I219's PCIe-based interface.




Figure 1-1 I219 Block Diagram



## 1.2 Main Flows

The I219 main 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 I219 functionality through two mechanisms:

- The I219 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.
- *Intel® Ethernet Controllers Loopback Modes*, Intel Corporation
- Energy Efficient Ethernet (EEE) 802.az specification.
- SMBus specification revision 2.0.



## 2.0 Interconnects

### 2.1 Introduction

The I219 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.

**Table 2-1 I219 Interconnect Modes**

System	LAN Connected Device	
	SMBus	PCIe
S0 and PHY Power Down	Not Used	Idle
S0 and Idle or Link Discovery	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 I219 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

### 2.2 PCIe-Based

**Note:** The I219 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 LAN-connected device's (LCD) PCIe-based interface. Standard PCIe validation tools cannot be used to validate this interface. Refer to [Section 11.4.4](#) 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 I219 only runs at 1250 Mb/s KX (PCIe-based) speed, which is half 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 I219 and the integrated LAN Controller. The I219 only initiates PHY status in-bands and then waits for an acknowledgment. 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 acknowledgment 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	Bits 4:0 = Register address Bit 5: 0b = Read access 1b = Write access Bit 6: 0b = Command Bit 7: 0b = MDIO register
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 4:0 = PHY address Bits 7:5 = Reserved (0b)
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	Bits 4:0 = Register address Bit 5: 0b = Read access 1b = Write access Bit 6: 1b = Acknowledge Bit 7: 0b = MDIO register
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 I219 on every status change at the PHY side or when a timeout defined in the In-Band Control register has expired. The I219 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	Bits 5:0 = Reserved (must be 0x0) Bit 6 = 0b Bit 7: 1b = Status
1st Data Byte	Bit 0 = K1 entry request Bit 1 = TX Off (MAC back pressure required) Bit 2 = EI entry request Bit 3 = Reserved (must be 0b) Bit 4 = Reserved Bit 5 = Inband Host WoL indication Bits 7:6 = Reserved (must be 0x0)
2nd Data Byte	Bits 1:0 = Speed: 00b = 10 Mb/s 01b = 100 Mb/s 10b = 1000 Mb/s 11b = Reserved Bit 2 = Duplex mode: 0b = Half duplex 1b = Full duplex Bit 3 = PHY Link status: 0b = Link down 1b = Link up Bit 4 = PHY cable disconnected Bit 5 = GLCI link status: 0b = Link down 1b = Link up Bit 6 = Interrupt request (Not used) Bit 7 = Reset complete
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	Bits 5:0 = Reserved (must be 0x0) Bit 6: 1b = Acknowledge Bit 7: 1b = Status
1st Data Byte	Bit 0 = K1 entry request Bit 1 = TX Off (MAC back pressure required) Bit 2 = EI entry request Bit 3 = Reserved (must be 0b) Bit 4 = Reserved Bit 5 = Inband Host WoL indication Bits 7:6 = Reserved (must be 0x0)
2nd Data Byte	Bits 1:0 = Speed: 00b = 10 Mb/s 01b = 100 Mb/s 10b = 1000 Mb/s 11b = Reserved Bit 2 = Duplex mode: 0b = Half duplex 1b = Full duplex Bit 3 = PHY Link status: 0b = Link down 1b = Link up Bit 4 = PHY cable disconnected Bit 5 = GLCI link status: 0b = Link down 1b = Link up Bit 6 = Interrupt request (Not used) Bit 7 = Reset complete
Reserved Byte	Reserved byte (must be 0b).
CRC-8 Byte	CRC of in-band packet excluding the K28.1 special symbol.



### 2.2.2.1.5 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 and Controls	Bits 5:0 = Reserved (must be 0x0) Bit 6 = 0b Bit 7: 1b = Status
1st Data Byte	Bit 0 = K1 entry request Bit 1 = XOFF request Bit 2 = EI entry request Bit 3 = XON request Bits 5:4 = Reserved Bits 7:6 = Reserved (must be 0x0)
2nd Data Byte	Bits 3:0 = MAC LED Status (not used). Bits 5:4 = Power saving D-states. Bit 6 = Power down. Bit 7 = Port Reset.
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	Bits 5:0 = Reserved (must be 0x0) Bit 6: 1b = Acknowledge Bit 7: 1b = Status
1st Data Byte	Bit 7:0 = Reserved (must be 0x0)
2nd Data Byte	Bit 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 I219 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 I219 as well as passing in-band information between them.

The SMBus uses two primary signals to communicate: SMBCLK and SMBDAT. 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 I219 are described in the relevant sections.

For further details on SMBus behavior, refer to 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 I<sup>2</sup>C specification.

To change the I219'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. For details, refer to [Section 9.5.8.3](#).

#### 2.3.1.2 SMBus Addressing

The I219's address is assigned using SMBus ARP protocol. The default SMBus address is 0xC8.

#### 2.3.1.3 Bus Time Outs

The I219 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 I219 does not abort the transaction. If the SMBus clock line is held low for 25 ms or longer, the I219 aborts the transaction.

As a slave, the I219 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 I219 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 I219 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 I219 has the capability of detecting a hung bus. If SMB\_CLK or SMB\_DATA are stuck low for more than 35 ms, the I219 forces the bus to idle (both SMB\_CLK and SMB\_DATA set), if it is the cause of the bus hang.

### 2.3.1.5 Packet Error Code (PEC) Support

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

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

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

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

### 2.3.1.6 SMBus ARP Functionality

The I219 supports the SMBus ARP protocol as defined in the SMBus 2.0 specification. The I219 is a persistent slave address device, meaning its SMBus address is valid after power up (constant 0xC8). The I219 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 I219 has a valid SMBus address.
- **AR** (Address Resolved) — This flag is set when the I219 SMBus address is resolved (SMBus address was assigned by the SMBus ARP process).

**Note:** These flags are internal I219 flags and are not shown to external SMBus devices.

Since the I219 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 I219 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 I219 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 I219 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 I219 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 I219 an SMBus address by issuing an Assign Address command. The I219 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 I219 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 I219.




Figure 2-1 SMBus ARP Flow



### 2.3.1.6.2 SMBus ARP UDID Content

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

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

Where:

- **Vendor ID** — The device manufacturer's ID as assigned by the SBS Implementer's 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	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 <sup>1</sup>
MSB			LSB

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

### 2.3.1.7 SMBus ARP Transactions

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 I219. All traffic between the I219 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 I219 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 I219 receive bytes causes the packet to be retransmitted 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 I219 hardware to 1522 bytes. Packets larger than 1522 bytes are silently discarded.



## 2.3.2.2 Transmit Flow

The I219 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 I219, it is transmitted over the network link.

The I219 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 I219 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 I219.

If the network link is down when the I219 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 I219, 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 2-2 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:** 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 I219)

Table 2-3 lists the slave SMBus transactions supported by the I219.

**Table 2-3 SMBus Slave Transactions**

Command	Transaction	Command		Fragmentation	Section
Transmit Packet	Block Write	First Middle Last Single	0x84 0x04 0x44 0xC4	Multiple Multiple Multiple Single	<a href="#">Table 2.3.3.1.1</a>
MtP Control	Block Write	Single	0xC3	Single	<a href="#">Table 2.3.3.1.2</a>
MtP Status Acknowledge	Block Write	Single	0xC3	Single	<a href="#">Table 2.3.3.1.3</a>
MtP Configuration	Block Write	Single	0xC3	Single	<a href="#">Table 2.3.3.1.4</a>

#### 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 I219 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

Address and Control:

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

Command 1st Byte:

Field	Bit(s)	Description
Reserved	7:0	Reserved.

Command 2nd Byte:

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



### 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 I219'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

Address and Control:

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

PHY Status 1st Byte:

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

PHY Status 2nd Byte:

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



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

### 2.3.3.1.4 MtP Configuration Command

The Configuration command is a single fragment command enabling the integrated LAN Controller access to all I219 registers over the SMBus link.

MtP 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 <sup>1</sup>	PHY Address

1. Bits 7:5

Address and Control:

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



## 2.3.4 Master Transactions

### 2.3.4.1 SMBus Transactions (I219 to the integrated LAN Controller)

To avoid starvation on the SMBus in the opposite direction, the I219 adds a byte-time idle between any two fragments it sends.

Table 2-4 lists the master SMBus transactions supported by the I219.

**Table 2-4 SMBus Master Transactions**

Command	Transaction	Command		Fragmentation	Section
Receive Packet	Block Write	First Middle Last Single	0x90 0x10 0x50 0xD0	Multiple Multiple Multiple Single	<a href="#">Table 2.3.4.1.1</a>
PtM Status	Block Write	Single	0xD3	Single	<a href="#">Table 2.3.4.1.2</a>
PtM Configuration Acknowledge	Block Write	Single	0xD3	Single	<a href="#">Table 2.3.4.1.3</a>
PtM WU	Block Write	Single	0xD5	Single	<a href="#">Table 2.3.4.1.4</a>

#### 2.3.4.1.1 Receive Packet Transaction

When the I219 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

The receive packet fragments have the following format:

Function	Byte Count	Data 1 (Op-Code)	Data 12	...	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 I219 on every status change at the PHY side, or when a time out defined in the In-Band Control register has expired. The I219 re-transmits the Status command if no acknowledge arrived after a time out defined in the PCIe Diagnostic register has expired. If after 32 retries no acknowledge arrived, the I219 aborts this command and continues with the next operation.

PtM Control 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

Address and Control:

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

PHY Status 1st Byte:

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



PHY Status 2nd Byte:

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



### 2.3.4.1.3 PtM Configuration Acknowledge Command

The Configuration Acknowledge command is a single fragment command enabling the I219 to acknowledge the integrated LAN Controller access to I219 registers over the SMBus link.

MtP 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 <sup>1</sup>	PHY Address

1. Bits 7:5

Address and Control:

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

### 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

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 I219 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 I219.

Possible states for activity over the SMBus:

- After power on (G3 to S5).
- 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 I219.

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 I219 is done through PCIe each time the system is in active power state (S0); PE\_RSTn signal is high. Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME/PMC.

The switching from PCIe to SMBus is done when the PE\_RSTn signal is low.

- Any transmit/receive packet that is not completed when PE\_RSTn goes to 0b is discarded.
- Any in-band message that was sent over PCIe and was not acknowledged is retransmitted over SMBus.



**Notes:**

1. Not all PCH PCIe ports can be used for the I219. Refer to the SkyLake/Greenlow/Purley EDS documentation for the specific ports that can be used with the I219.
2. Any CLKOUT\_PCIE and PCIECLKRQ ports can be used to connect to the I219. These can be selected using the FITC tool.
3. PETp/n, PERp/n, PE\_CLKp/n should be routed as a differential pair as indicated in the PCIe specification.
4. Refer to the I219 reference schematics and design checklists for more details.

**Figure 2-2 PCIe/SMBus Interface**



## 3.0 Pin Interface

---

### 3.1 Pin Assignment

The I219 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 Vdc (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 Vdc (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		
LAN_DISABLE_N	3	I		When this pin is set to 0b, the I219 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_PLUS[0] MDI_MINUS[0]	13 14	Analog	Bi-dir	Media Dependent Interface[0]: <b>1000BASE-T:</b> In MDI configuration, MDI[0]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[0]+/- corresponds to BI_DB+/-. <b>100BASE-TX:</b> In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair. <b>10BASE-T:</b> 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_PLUS[1] MDI_MINUS[1]	17 18	Analog	Bi-dir	Media Dependent Interface[1]: <b>1000BASE-T:</b> In MDI configuration, MDI[1]+/- corresponds to BI_DB+/- and in MDI-X configuration MDI[1]+/- corresponds to BI_DA+/-. <b>100BASE-TX:</b> In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair. <b>10BASE-T:</b> 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_PLUS[2] MDI_MINUS[2] MDI_PLUS[3] MDI_MINUS[3]	20 21 23 24	Analog	Bi-dir	Media Dependent Interface[3:2]: <b>1000BASE-T:</b> In MDI configuration, MDI[3:2]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[3:2]+/- corresponds to BI_DB+/-. <b>100BASE-TX:</b> Unused. <b>10BASE-T:</b> Unused.
XTAL_OUT	9	O		Output crystal.
XTAL_IN	10	I		Input crystal.
RBIAS	12	Analog		Connect to ground through a 3,01 K $\Omega$ +/-1%.



### 3.1.6 Testability Pins (5)

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

**Note:** The I219 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
VDD3P3	4	Power	3.3 Vdc out.
VDD3P3_IN	5	Power	3.3 Vdc supply.
SVR_EN_N	6	Input	SVR Enable pin. Connect to GND for internal SVR use. Connect to VDD3P3_IN when using external power.
CTRL0P9	7	Analog	Internal SVR control pin. Connect to a 4.7 μH inductor and to the core rail.
VDD0P9	8 11 16 22 37 40 43 46 47	Power	Core Vdc supply.
VDD3P3	15 19 29	Power	3.3 Vdc supply.



## 3.2 Pinout




Figure 3-1 I219 Pinout



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



## 4.0 Package

### 4.1 Package Type and Mechanical

The I219 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 4-1 Package Illustration



## 4.2 Package Electrical and Thermal Characteristics

The thermal resistance from junction to case,  $\theta_{JC}$ , is 15.1 °C/Watt. The thermal resistance from junction to ambient,  $\theta_{JA}$ , is as follows, 4-layer PCB, 85 °C<sub>ambient</sub>.

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

No heat sink is required.

## 4.3 Power and Ground Requirements

All the grounds for the package is down-bonded to an Exposed Pad\* E-pad\*.



## 4.4 Ball Mapping

Pin Name	Pin Number	Side
RSVD1_VCC3P3	1	Left
LANWAKE_N	2	Left
LAN_DISABLE_N	3	Left
VDD3P3	4	Left
VDD3P3_IN	5	Left
SVR_EN_N	6	Left
CTRL0P9	7	Left
VDD0P9	8	Left
XTAL_OUT	9	Left
XTAL_IN	10	Left
VDD0P9	11	Left
RBIAS	12	Left
MDI_PLUS[0]	13	Bottom
MDI_MINUS[0]	14	Bottom
VDD3P3	15	Bottom
VDD0P9	16	Bottom
MDI_PLUS[1]	17	Bottom
MDI_MINUS[1]	18	Bottom
VDD3P3	19	Bottom
MDI_PLUS[2]	20	Bottom
MDI_MINUS[2]	21	Bottom
VDD0P9	22	Bottom
MDI_PLUS[3]	23	Bottom
MDI_MINUS[3]	24	Bottom

Pin Name	Pin Number	Side
LED2	25	Right
LED0	26	Right
LED1	27	Right
SMB_CLK	28	Right
VDD3P3	29	Right
TEST_EN	30	Right
SMB_DATA	31	Right
JTAG_TDI	32	Right
JTAG_TMS	33	Right
JTAG_TDO	34	Right
JTAG_TCK	35	Right
PE_RST_N	36	Right
VDD0P9	37	Top
PETp	38	Top
PETn	39	Top
VDD0P9	40	Top
PERp	41	Top
PERn	42	Top
VDD0P9	43	Top
PE_CLKP	44	Top
PE_CLKN	45	Top
VDD0P9	46	Top
VDD0P9	47	Top
CLK_REQ_N	48	Top
GND	49	Epad (Center)



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



## 5.0 Initialization

### 5.1 Power Up

Initialization begins with power up.




Figure 5-1 Power-Up Sequence



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

Note	
1	Platform power ramps up (3.3 Vdc/Core Vdc).
2	XTAL is stable after $T_{XTAL}$ seconds.
3	Internal Power On Reset triggers $T_{POR}$ after XTAL is stable. Strapping options are latched.
4	PCIe training if PE reset is de-asserted.
5	SMBus ARP if initiated.
6	Wait for the integrated LAN Controller SMBus address valid.
7	Send Link Status message.
8	MAC configures the I219.
9	PHY goes through auto-negotiation to acquire link.

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

## 5.2 Reset Operation

The reset sources for the I219 are as follows:

- **Internal Power On Reset (POR)** — The I219 has an internal mechanism for sensing the power pins. Until power is up and stable, the I219 generates an internal active low reset. This reset acts as a master reset for the I219. While the internal reset is 0b, all registers in the I219 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 I219 registers are not.

**Note:** The integrated LAN controller configures the LCD registers. Other I219 GbE LCD registers do not need to be configured.

- **PCIe Reset** from pin — After asserting a PCIe reset, the I219 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 I219 except the wake up filters content.

**Note:** LAN\_DISABLE\_N is the only external signal that can reset the PHY. Refer to [Section 6.3.1](#) 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 6-1 on page 49](#).

[Table 5-1](#) lists the impact of each I219 reset.



**Table 5-1 I219 Resets**

Effects/Sources	PCIe-Based Interface	Non-PHY Registers and State	PHY Registers and State	Reset Complete Indication <sup>1</sup>	Strapping Options	Fuse Registers	Move Out of Power Down Mode	Wake Up Register
Internal POR <sup>2</sup>	X	X	X	X	X	X		X
PHY Soft Reset <sup>3</sup>			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 the I219's MDIO space is not reset.

## 5.3 Timing Parameters

### 5.3.1 Timing Requirements

The I219 requires the following start-up and power-state transitions.

**Table 5-2 Timing Requirements**

Parameter	Description	Min	Max	Notes
T <sub>r2init</sub>	Completing a PHY configuration following a reset complete indication.		0.5 s	

### 5.3.2 Timing Guarantees

The I219 guarantees the following start-up and power state transition related timing parameters.

**Note:** For platform power sequencing requirements for the Cougar Point/Patsburg integrated LAN controller, refer to the Cougar Point/Patsburg EDS.

**Table 5-3 Timing Requirements**

Parameter	Description	Min	Max	Notes
T <sub>PHY_Reset</sub>	Reset de-assertion to PHY reset complete.		10 ms	PHY configuration should be delayed until PHY completes its reset.
T <sub>c2an</sub>	Cable connect at start of auto-negotiation.	1.2 s	1.3 s	Per 802.3 specification.
T <sub>XTAL</sub>	XTAL frequency stable after platform power ramp up.		45 ms	
T <sub>POR</sub>	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 the I219.

### 6.1 Power Information

**Table 6-1 I219 Power Consumption Target**

System State		Link State	Device Power Using Internal SVR (mW)
S0 (Max)		Thermal Design Power (TDP)	542
S0 (Typ)		1000 Mb/s Active	542
		1000 Mb/s Idle	439
		1000 Mb/s LPI (EEE link partner only)	55
		100 Mb/s Active	264
		100 Mb/s Idle	177
		100 Mb/s LPI (EEE link partner only)	51
		10 Mb/s Active	306
		10 Mb/s Idle	84
Sx	WoL Enabled	100 Mb/s WoL enabled	169
		10 Mb/s 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 I219 operates from a 3.3 Vdc external power rail.

### 6.2.1 Core Vdc Supply

The Core Vdc rail can be supplied by integrated SVR solution with external inductor and capacitor.

## 6.3 Power Management

### 6.3.1 Global Power States

The I219 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 I219 and until the I219 powers up its PHY. The I219 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 I219 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 I219 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 I219 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).
- **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\_DISABLE\_N pin to zero. The I219 loses all functionality in this mode other than the ability to power up again.
- **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 (refer to [Section 9.5.2.18](#)).
- **LPI** — IEEE802.3az [Energy Efficient Ethernet (EEE)] defines an optional Low Power Idle (LPI) mode for 100BASE-T, 100BASE-TX and other interfaces. LPI enables power saving by switching off part of the I219 functionality when no data needs to be transmitted or/and received. When LPI support is enabled, the I219 will shut off RX circuitry and send an in-band RX LPI Indication on detection that link the partner's TX moved into LPI state. The I219 PHY 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 I219 and until the I219 powers up its PHY. The I219 should consume less than ~40 mA during this period. Following the I219 LCD entering reset, the power-up sequence is considered done and the requirement is removed. Refer to [Section 5.1](#) for a description of the power-up sequence.

### 6.3.1.2 Cable Disconnect State

The I219 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 I219 detects energy on the MDI link, and starts the following exit sequence:

- The I219 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 (Tc2an time) and then starts to advertise data on the line.

### 6.3.1.3 Power-Down State


The I219 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 I219 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 6-1](#) shows the power-down sequence in the two later cases.



**Figure 6-1 Power-Down Sequence**



**Notes:** In cases where the LAN\_DISABLE\_N pin cannot be used a power down in-band can be used. When used the power savings are lower since not all logic can be turned off at this mode.


**Table 6-2 Power-Down Notes**

Note	
1	The LAN_DISABLE_N pin is set to zero
2	Once the I219 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 I219 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 I219 functionality when no data needs to be transmitted or/and received. Decision on whether the I219 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 I219 to allow power saving in the transmit circuitry. Information on whether Link Partner has entered Low Power Idle mode is detected by the I219 and communicated to the integrated LAN controller to allow for power saving in the receive circuitry.

Figure 6-2 illustrates general principles of an EEE LPI operation on the Ethernet Link.



**Figure 6-2 EEE LPI Compliant Operation**

**Table 6-3 LPI Parameters**

Parameter	Description
Sleep Time ( $T_s$ )	Duration PHY sends Sleep symbols before going Quiet.
Quiet Duration ( $T_q$ )	Duration PHY remains Quiet before it must wake for Refresh period.
Refresh Duration ( $T_r$ )	Duration PHY sends Refresh symbols for timing recovery and coefficient synchronization.
PHY Wake Time ( $T_{w\_PHY}$ )	Minimum duration PHY takes to resume to Active state after decision to Wake.
Receive System Wake Time ( $T_{w\_System\_rx}$ )	Wait period where no data is expected to be received to give the local receiving system time to wake up.
Transmit System Wake Time ( $T_{w\_System\_tx}$ )	Wait period where no data is transmitted to give the remote receiving system time to wake up.



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 I219 that Low Power Idle mode should end. The I219 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

**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 I219 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 I219 does not enter ACBS mode.



The I219 ACBS works in both S0 and Sx states. Since the I219 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 I219 is in ACBS mode.

### 6.3.1.6 Automatic Link Down-Shift

Automatic link down-shift is a collection of power saving features that enable a link down-shift 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 I219 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.



## 6.4 I219 Ultra Low Power (ULP)

To reduce the power consumption of the I219 during link disconnect, the I219 introduces a new Ultra Low Power (ULP) mode. In ULP mode the I219 will power gate most of its functionality and only maintain a small portion of the die powered sensing for future cable connection (energy on the Ethernet link). Once energy is detected the I219 will exit ULP mode and either send an in-band message (link connect status) or assert LANWAKE# according to configuration.

The I219 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. The SW will configure I219 registers for proper entry and exit conditions.

The I219 can be configured to enter ULP on future link disconnect while the link is still connected so that during 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 the I219 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 (either located 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 would be 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# would de-assert (board controlled) not as part of the K1 entry flow. The LAN Controller would refer to such event as an indication of OBLCD disconnect.

During K1, CLKREQ# is de-asserted. To indicate presence, the LAN Connected Device will send an in-band beacon message once every BCN\_INTER for BCN\_DUR to indicate that the LAN Controller that the dongle is still connected. The LAN Controller would not exit K1 when sensing these squelch pulses. K1 exit will only start after sensing a beacon pulse longer than LC.OFFSET+BCN\_DUR. At this state the LAN Controller should enable the PLL and start the K1 exit flow.



If the LAN Controller misses these beacon pulse indications for LC.MISS\_BCN\*BCN\_INTER it should refer to such event as an indication of dongle disconnect. In addition the LAN Controller will implement 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 should be 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 should refer to such event as an indication of OBLCD disconnect.

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

- It will try to take the LAN Connected Device out of ULP without success.
- OBLCD will be reconnected (in S0) and the LAN Connected Device will assert CLKREQ#, indicating the LAN Controller to move to PCIe.

**Note:** When the LAN Connected Device is in WoL mode the LAN Controller is powered off, OBLCD disconnection will have no indication.

## 6.5.2 Indicating/Sensing OBLCD Connection

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

- In PCIe, the LAN Connected Device will cause K1 exit by asserting the CLKREQ# and then sending status in-bands over PCIe
- While the interface is defined as SMB, OBLCD connection will not be reflected 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 I219 through the PCIe or the SMBus (depends on system state). The I219 starts transmitting the arrived packet over the wire after it gathers 8 bytes of data if the PCIe interface is active or after all packet data is received if it was transferred over the SMBus; this behavior has no dependency on the link speed. The I219 design is based on the assumption that the integrated LAN controller has the full packet ready for transmission.

In several cases the I219 has to stop transmission over the wire while still accepting data from the upper connection (PCIe or SMBus). For those cases, the I219 maintains a 3 KB FIFO. The cases where the I219 needs to stop Tx are:

- PAUSE packet was received on the Rx side while flow control is enabled. For full support of flow control, the *Receive Flow Control Enable (RFCE)* bit (bit 7) in the LCD Receive Control register should be set in addition to the configuration in the integrated LAN controller.
- In half-duplex mode while the I219 is in the middle of a receive (DEFER).
- In half-duplex mode while a collision was detected on the wire.

In addition to stop transmission, the I219 sends an in-band message to the integrated LAN controller with the Tx *OFF* bit set. This in-band message must be sent at the first gap between received packets if (at the same time) the event caused the stop transmit is not valid and transmission over the wire is activated, the I219 might avoid sending the in-band message. An in-band message with the Tx *OFF* bit cleared is sent when the collided packet was successfully transmitted or dropped after 16 retries (refer to [Section 7.3.1](#)).

In-band messages from the integrated LAN controller to the I219 always come in between packets during the IPG. The I219 does not accept in-band messages in the middle of a packet.



## 7.2 Rx Flow

The I219 maintains a 3 KB FIFO on the receive side in order not to lose packets when PCIe is active but at K1 power save mode. In this case, the I219 initiates recovery of the PCIe when a reception has started. If the link is at 1 Gb/s, the transmission of the packet over the PCIe bus starts immediately after recovery. If the link speed is lower, the I219 starts the transmission after the entire packet is received. The I219 assumes maximum recovery time (from the K1 state) of 10  $\mu$ s on both sides of the PCIe side. Higher recovery time causes a packet drop on the receive side.

The I219 identifies PAUSE packets, stop transmission, and a send in-band message as described in [Section 7.1](#).

In-band messages from the I219 to the integrated LAN controller always come in between packets during the IPG.

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

## 7.3 Flow Control

Flow control as defined in 802.3x, as well as the specific operation of asymmetrical flow control defined by 802.3z, is supported in the integrated LAN controller during PCIe mode and in the LAN Connected Device during SMBus mode. Some of the flow control functionality has moved to the I219. The following registers are duplicated to the I219 for the implementation of flow control:

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

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




## 7.3.1 MAC Control Frames and Reception of Flow Control Packets

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

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

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

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



**Figure 7-1 802.3x MAC Control Frame Format**

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

**Note:** Flow control capability must be negotiated between link partners via the autonegotiation 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 I219 has validated the reception of an XOFF, or PAUSE frame, it performs the following:

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

Resuming transmission might occur under the following conditions:

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

Once the I219 has validated the reception of an XON frame, it performs the following:

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

## 7.3.2 Transmitting PAUSE Frames

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

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

When the in-band message from the integrated LAN controller is received, the I219 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 I219 receive and transmit FIFO status. If the sum of data in both FIFOs is greater than the configured *FCTH.HTHRSH*, the I219 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 I219 sends a PAUSE frame with a timer value of zero (XON). The I219 will send an XOFF refresh message with the PAUSE Timer set to *FCTTV* if after *FCRTV* time from the previous XOFF message the transmit and receive buffer fullness is still above the low threshold value.

**Note:** Transmitting flow control frames should only be enabled in full-duplex mode per the IEEE 802.3 standard. Software should make sure that the transmission of flow control packets is disabled when the I219 is operating in half-duplex mode.

---

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



## 7.4 Wake Up

The I219 supports host wake up. The host can enable host wake up from the I219 by setting the *Host\_WU\_Active* bit.

To set the I219'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 the configuration of the filters to wake up mode.
2. Set *MACPD\_enable* bit (bit 2) of the Port Control register (page 769, register 17) to enable the I219 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:
  - Flexible Filter Value Table LSB — FFVT\_L (filters 01)
  - Flexible Filter Value Table MSBs — FFVT\_H (filters 23)
  - Flexible Filter Value Table — FFVT\_45 (filters 45)
  - Flexible Filter Value Table — FFVT\_67 (filters 67)
  - Flexible TCO Filter Value/Mask Table LSBs — FTFT\_L
  - Flexible TCO Filter Value/Mask Table MSBs — FTFT\_H
  - Extended Flexible Filter Value Table — EFFVT (Filters 8-15)
4. Configure the I219'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 I219's wake up functionality.

On a detection of Host wake up event, the I219 will:

1. Assert the LANWAKE# pin indicating wake to PMC.
2. If in DeepSx, PMC will power up the SUS well.
3. PMC will send a PINSTOP message to the I219 through SMBus
4. On a reception of a PINSTOP message, the I219 will stop asserting the LANWAKE# pin and send a WU message over SMBus indicating the WU source.
  - The I219 will send a WU message on every PINSTOP message reception.
5. The I219 will not assert the LANWAKE# pin again until a WU message was sent and acknowledged. The In case of host wake up the integrated LAN controller wakes the host.
6. Host should issue a LCD reset to the I219 before clearing the *Host\_WU\_Active* bit.
7. Host reads the Wake Up Status register (WUS); wake up status from the I219.

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



While in wake up active mode new wake up packets received will not overwrite the packet in the FIFO. The I219 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 I219 supports two types of wake up mechanisms:

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

### 7.4.1.1 Advanced Power Management Wake Up

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

At power up, if the I219's wake up functionality is enabled, the *APM Enable* bits from the NVM are written to the I219 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 I219 checks all incoming packets for Magic Packets. Refer to [Section 7.4.1.4.1.4](#) for a definition of Magic Packets.

To enable APM wake up, programmers should write a 1b to bit 10 in register 26 on page 0 PHY address 01, and then the station address to registers 27, 28, 29 at page 0 PHY address 01. The order is mandatory since registers 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 I219 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 I219 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 I219 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:

1. 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.
2. Setting bit 2 (*MACPD\_enable*) of the Port Control register (page 769, register 17) to put the I219 in wake up mode.

Once wake up is enabled, the I219 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 I219:

- 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 I219 by setting 776.19[0] prior to WoL entry.

Post wakeup, the once the driver is ready to get the wake up packet (post LAN Controller configurations) the driver should set 776.19[1] and the wake up packet would be sent through the LC-LCD interface as an incoming packet.



## 7.4.1.4 Wake Up Packets

The I219 supports various wake up packets using two types of filters:

- Pre-defined filters
- Flexible filters

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

### 7.4.1.4.1 Pre-Defined Filters

The following packets are supported by the I219'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 I219 generates a wake up event after receiving any packet whose destination address matches one of the seven 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 I219 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 Packet

If the *Broadcast Wake Up Enable* bit (bit 4) in the WUFC register is set, the I219 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 Packet Technology Details:

Once the I219 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 I219 expects the destination address to either:

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

If the packet destination address met one of the three criteria previously listed, the I219 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 I219 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 I219 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 I219 starts up with the RAH0/RAL0 register 0 loaded from the NVM. This enables the I219 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 I219 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 I219 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.



### 7.4.1.4.1.6 Directed IPv4 Packet

The I219 supports receiving Directed IPv4 packets for wake up if the *IPV4* bit (bit 6) is set in the WUFC register. Three IPv4 addresses are supported, which are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must contain the station’s MAC address, a Protocol Type of 0x0800, and one of the three host programmed Ipv4 addresses. The I219 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.

### 7.4.1.4.1.7 Directed IPv6 Packet

The I219 supports receiving Directed IPv6 packets for wake up if the *IPV6* bit (bit 7) is set in the WUFC register. One IPv6 address is supported, which is programmed in the IPv6 Address Table (IP6AT). A successfully matched packet must contain the station’s MAC address, a protocol type of 0x0800, and the programmed IPv6 address. The I219 also handles Directed IPv6 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

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



Offset	# of Bytes	Field	Value	Action	Comment
21 + S + D	1	Hop Limit	-	Ignore	
22 + S + D	16	Source IP Address	-	Ignore	
38 + S + D	16	Destination IP Address	IP6AT	Compare	Match value in IP6AT.

#### 7.4.1.4.1.8 Flexible Filter

The I219 supports a total of 32 flexible filters. Each filter can be configured to recognize any arbitrary pattern within the first 128 bytes of the packet. To configure the flexible filter, software programs the mask values into the Flexible Filter Mask Table (FFMT) and the required values into the Flexible Filter Value Table (FFVT), and the minimum packet length into the Flexible Filter Length Table (FFLT). These contain separate values for each filter. Software must also enable the filter in the WUFC register, and enable the overall wake up functionality must be enabled by setting PME\_En in the Power Management Control Status Register (PMCSR) or the WUC register.

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

#### 7.4.1.4.1.9 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 Management Engine (ME) Wake Up

The I219 can also wake up the ME. Any packet that should be routed to the ME during normal operation should wake up the ME when it is in the Moff state. The ME firmware should configure the MANC register and the relevant manageability packet filters before setting the integrated LAN controller to the DMoff state.

The integrated LAN controller can also wake up the ME by a link status change in the same manner as host wake up. Wake up the ME on link status change is enabled by the *WoLS* bit (bit 12) in the MANC register.

### 7.4.2.1 Manageability Wake Up Receive Filter

The I219 has the ability to wake up the ME. Any packet that can cause ME wake up must first match the MAC address filtering this includes Exact Unicast/Multicast filtering; hash based Multicast filtering any unicast MAC address if Promiscuous is enabled and Broadcast packets if enabled. Other filters can be enabled based on the MANC register setting.

Figure 7-2 shows a top level diagram of the ME wake up filters. The following sections describe these filters.




Figure 7-2 ME Receive Wake Up Filters Top Level



### 7.4.2.1.1 Manageability Registers

The Manageability filtering configuration is done by the manageability firmware. The following is a list of registers in the I219 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.

### 7.4.2.1.2 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.

### 7.4.2.1.3 ARP Packets Filtering

The I219 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.



## 7.4.2.2 Flexible TCP UDP Port and IP Address Filtering

The I219 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 dependent 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.

## 7.4.2.3 Dedicated TCP UDP Port and IP Address Filtering

The I219 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 dependent 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 dependent on a specific IP address and to pass filtering the packet must pass IP filtering as well as port and protocol filtering.

### 7.4.2.3.1 L2 Filtering

The I219 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, the I219 wakes up the ME (the packet must pass any of the MAC filters as well).



### 7.4.2.3.2 Flexible TCO Filtering

The I219 includes two flexible filters as defined by the Flexible TCO Filter Table (FTFT). These filters compare an exact pattern match within the 1st 128 bytes of the packet. Enabling these filters is done by setting *FLEX\_TCO1/0\_FILTER\_EN* bits (bits 6 and 7) in the MANC register.

### 7.4.2.4 Accessing the I219'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, refer to [Section 9.5.9.1](#).

## 7.5 Network Proxy Functionality

### 7.5.1 Introduction

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

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

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

Ethernet devices that implement IPv6 NS offload must implement it as defined in Power Management specification on the NDIS Program, Connect site. Specifically, the offload must respond to a Neighbor Solicitation (operation = 135) by responding with an NS Advertisement (operation = 136) as defined in RFC 2461. Devices must support at least 2 NS offloads, each with up to 2 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 I219:

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 I219 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 I219 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 I219 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	



Offset	# of Bytes	Field	Value	Action	Comment
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.




Figure 7-3 Neighbor Discovery



### 7.5.4.1 Ipv6 Neighbor Solicitation Packet

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



## 7.5.4.2 Ipv6 Neighbor Advertisement Packet

Offset	# of bytes	Field	Value (hex)	Action
0	6	Destination Address		Copy from ND packet
6	6	Source Address		Station Address
12	S=(0/4)	Possible VLAN Tag		Copy from ND packet
12+S	D=(0/8)	Possible LLC/SNAP Header		Copy from ND packet
IPv6 header				
12+D+S	2	Type	0x86DD	Constant (Copy from ND packet)
14+D+S	1	Version/ Traffic Class	0x6	Constant (Copy from ND packet)
15+D+S	3	Traffic Class/Flow Label		Constant (Copy from ND packet)
18+D+S	2	Payload Length		
20+D+S	1	Next Header	0x3A	Constant
21+D+S	1	Hop Limit	0xFF	Constant
22+D+S	16	Source Address		relevant IPv6AT entry (ND target address)
38+D+S	16	Destination Address		Copy from ND packet Source address If source address was undefined - send to All Nodes Multicast (FF02::1)
ICMPv6 header				
54+D+S	1	Type	0x88	Constant
55+D+S	1	Code	0x0	
56+D+S	2	Checksum		Calculate
58+D+S	4	Flags	0x60000000	Constant (Solicited, Override) if the source address was defined
			0x20000000	Constant (Override) if the source address was undefined
62+D+S	16	Target IP Address	IP6AT	Same as source address
78+D+S	1	Type	0x2	Target Link Layer Address option
79+D+S	1	Length	0x1	
80+D+S	6	Link Layer Address	From ND	
86+D+S	4	CRC		Calculate



## 7.5.5 Multicast Listener Discovery Support

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

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

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

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

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

MLDv2 is designed to be inter-operable with MLDv1.

MLD messages supported by the I219

### As a Receiver:

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

There are two subtypes of Multicast Listener Query messages supported by the I219:

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

### As a Transmitter:

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



### 7.5.5.1 MLD Query Packet

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



## 7.5.5.2 MLD Report Packet

Offset	# of bytes	Field	Value (hex)	Action
0	6	Destination Address		Copy from MLD query packet
6	6	Source Address		Station Address
12	S=(0/4)	Possible VLAN Tag		Copy from MLD packet
12+S	D=(0/8)	Possible LLC/SNAP Header		Copy from MLD packet
IPv6 header				
12+D+S	2	Type	0x86DD	Constant (Copy from MLD packet)
14+D+S	1	Version/ Traffic Class	0x6	Constant (Copy from MLD packet)
15+D+S	3	Traffic Class/Flow Label		Constant (Copy from MLD packet)
18+D+S	2	Payload Length		
20+D+S	1	Next Header	0x00	Constant (Copy from MLD packet)
21+D+S	1	Hop Limit	0x01	Constant (Copy from MLD packet)
22+D+S	16	Source Address		Link Local Address
38+D+S	16	Destination Address		Multicast Address being reported
Hop by Hop header				
54+D+S	1	Next Header	0x3A	Constant (Copy from MLD packet)
55+D+S	1	Header Extended Length	0x00	Constant (Copy from MLD packet)
56+D+S	1	Type	0x05	Constant (Copy from MLD packet)
57+D+S	1	Length	0x02	Constant (Copy from MLD packet)
58+D+S	2	MLD	0x0000	Constant (Copy from MLD packet)
60+D+S	1	PadN Option	0x01	Constant (Copy from MLD packet)
61+D+S	1	PadN Length	0x00	Constant (Copy from MLD packet)
ICMPv6 header				
62+D+S	1	Type	0x83	Constant
63+D+S	1	Code	0x00	Constant
64+D+S	2	Checksum		Calculate
66+D+S	2	Maximum Respond Delay	0x0000	
68+D+S	2	Reserved	0x0000	
70+D+S	16	Multicast Address		Copy from MLD query packet, if undefined send a single packet for each of Product Name 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, refer to [Section 9.5.2.14.1](#).

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 PHY Functionality

### 8.1 Reverse Auto-Negotiation

When reverse auto-negotiation (LPLU) is enabled, the auto-negotiation process proceeds as usual except the priority resolution is resolved in the order shown in [Table 8-1](#). The IEEE registers 4.8:5 and 9.9:8 are used to advertise the capability just like in normal auto-negotiation. Fullchip register 25\_0.6 (PHY address 1, page 0, register 25, it 6) controls whether GbE mode should be advertised regardless of the values in registers 9.9:8.

**Table 8-1 Reverse Auto-Negotiation Priority Resolution**

Priority	Speed/Duplex
1 (highest)	10 full-duplex
2	10 half-duplex
3	100 full-duplex
4	100 half-duplex
5 (lowest)	1000 full-duplex

Full-chip registers 25\_0.2 (PHY address 1, page 0, register 25, bit 1) and 25\_0.6 (PHY address 1, page 0, register 25, bit 6) controls the auto-negotiation behavior as listed in [Table 8-1](#). Note that the LPLU and 1000dis signals can toggle respectively. A 1b to 0b transition sets the register to zero, and a 0b to 1b transition set the register to one.

Since these registers can be changed via signal toggling or via MDIO write access, the latest occurring event (signal toggling or register write) will determine the values in the registers.

Registers 25\_0.2 (PHY address 1, page 0, register 25, bit 2) and 25\_0.6 (PHY address 1, page 0, register 25, bit 6) will not take effect unless one of the following occurs.

- Software reset is asserted (Register 0.15)
- Restart Auto-Negotiation-Now is asserted (Register 25\_0.10)

The enabling of reverse auto-negotiation and the disabling of 1000BASE-T creates some conflicts. [Table 8-2](#) clarifies the behavior.



**Table 8-2 Reverse Auto-Negotiation, Disable 1000BASE-T, and Down-Shift Behavior**

Reverse/ Normal Auto-Negotiation	Disable 1000BASE-T	Action
Reverse	Disable	Negotiates to the lowest of 10/100/1000.
Reverse	Enable	Negotiates to the lowest of 10/100.
Normal	Disable	Negotiates to the highest of 10/100/1000.
Normal	Enable	Negotiates to the highest of 10/100.

The reverse auto-negotiation mechanism works by temporarily stalling the base page exchange until the link partners 10/100 capabilities is learned. Once the PHY learns the link partner's capability it advertises only the desired capability, in this case the lowest speed with the highest duplex. [Figure 8-1](#) illustrates the process in more detail.

Each time auto-negotiation is restarted all advertised capabilities are advertised to the link partner. If reverse auto negotiation is enabled then the pause bits for every other FLP burst is inverted. Since the link partner never sees 3 consecutive FLP burst having the same bit pattern its auto-negotiation is stalled. At the same time the link partner advertises its capabilities. Once the PHY learns the link partner's capabilities it determines the lowest common speed/duplex. It then no longer advertises any higher capabilities in the FLP burst. This new set of advertisement is used and the remainder of the auto-negotiation process continues. (for example, the FLP burst pause bits are no longer toggling so the link partner's auto-negotiation process is no longer stalled.)

**Note:** If two PHYs with reverse auto-negotiation enabled are connected to each other, the process described above will still work. Since all advertised capabilities are initially sent to the link partner and the speed/duplex bits are not toggling it is possible for the PHY to determine each other's capabilities.




Figure 8-1 Reverse Auto-Negotiation Add-On State Machine



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



## 9.0 Programmer's Visible State

### 9.1 Terminology

Shorthand	Description
RW	Read/Write. A register with this attribute can be read and written. If written since reset, the value read reflects the value written.
RWS	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, I f a register is read only, writes to this register have no effect.
WO	Write Only. Reading this register might not return a meaningful value.
RWC	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.
RW/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 I219 power up and following reset. Other fields in the same 16-bit register must be loaded with their default values.

## 9.2 MDIO Access

After LCD reset to the I219 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.

## 9.3 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 I219 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 I219 contains registers controlling the custom interface as well as other I219 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 I219 registers in PHY address 01, are divided into pages. Each page has 32 registers. Registers 0-15 are identical in all the pages and are the IEEE defined registers. Register 31 is the page register in all pages of PHY address 01. All other registers are page-specific.

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



In pages 800 and 801, the register address space is more than 32. Refer to [Section 9.5.9](#) 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 address 16 to 31 are used as Address Offset port and Data port for extended set of registers. Refer to [Section 9.5.2](#) for details about these registers.

## 9.4 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 LSbit first. That is, the bit order over the wire for this example is

Last	First
..., 0000 0011, 0000 0010, 0000 0001, 0000 0000	

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

	Byte 1	Byte0
DW address (N)	0x01	0x00
DW address (N+1)	0x03	0x02
DW address (N+2)	0x05	0x04

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

	Byte 1	Byte0
(Word aligned)	0x00	0x01

The following exceptions use network ordering:

- All EtherType fields

**Note:** The “normal” notation as it appears in text books, etc. is to use network ordering. Example: Suppose a MAC address of 00-A0-C9-00-00-00. The order on the network is 00, then A0, then C9, etc. However, the host ordering presentation would be

	Byte 1	Byte0
DW address (N)	A0	00
DW address (N+1)	00	C9
DW address (N+2)	00	00



## 9.5 Register Definitions

### 9.5.1 Address Map

Table 9-1 Address Map

PHY Address	Page	Register	Name	Link
<b>PHY Registers</b>				
02	Any	0	Control	<a href="#">95</a>
02	Any	1	Status	<a href="#">96</a>
02	Any	2	PHY Identifier 1	<a href="#">97</a>
02	Any	3	PHY Identifier 2	<a href="#">97</a>
02	Any	4	Auto-Negotiation Advertisement	<a href="#">98</a>
02	Any	5	Auto-Negotiation Link Partner Ability	<a href="#">99</a>
02	Any	6	Auto-Negotiation Expansion	<a href="#">100</a>
02	Any	7	Auto-Negotiation Next Page Transmit	<a href="#">100</a>
02	Any	8	Link Partner Next Page	<a href="#">101</a>
02	Any	9	1000BASE-T Control	<a href="#">101</a>
02	Any	10	1000BASE-T Status	<a href="#">102</a>
02	Any	11 - 14	Reserved	
02	Any	15	Extended Status	<a href="#">103</a>
02	0	16 - 17	Reserved	
02	0	18	PHY Control 2	<a href="#">103</a>
02	0	19	Loopback Control	<a href="#">105</a>
02	0	20	Rx Error Counter	<a href="#">106</a>
02	0	21	Management Interface (MI)	<a href="#">106</a>
02	0	22	PHY Configuration	<a href="#">106</a>
02	0	23	PHY Control	<a href="#">107</a>
02	0	24	Interrupt Mask	<a href="#">108</a>
02	0	25	Interrupt Status	<a href="#">109</a>
02	0	26	PHY Status	<a href="#">110</a>
02	0	27	LED Control 1	<a href="#">111</a>
02	0	28	LED Control 2	<a href="#">112</a>
02	0	29	LED Control 3	<a href="#">113</a>
02	0	30	Late Collision Counter	<a href="#">113</a>
02	0	31	Link Polarity and Length Status	<a href="#">114</a>



**Table 9-1 Address Map (Continued)**

PHY Address	Page	Register	Name	Link
<b>Port Control Registers</b>				
01	769	16	Custom Mode Control	<a href="#">115</a>
01	769	17	Port General Configuration	<a href="#">115</a>
01	769	21	Power Management Control	<a href="#">116</a>
01	769	23	SMBus Control	<a href="#">116</a>
01	769	25	Rate Adaptation Control	<a href="#">116</a>
01	769	27	Flow Control Transmit Timer Value	<a href="#">117</a>
01	769	28	System Low Power Control – SxCTRL	<a href="#">117</a>
01	764	23 - 24	SERDES MDI Control Register – SMDIC	<a href="#">117</a>
<b>Statistics Registers</b>				
01	778	16 - 17	Single Collision Count – SCC	<a href="#">119</a>
01	778	18 - 19	Excessive Collisions Count – ECOL	<a href="#">119</a>
01	778	20 - 21	Multiple Collisions Count – MCC	<a href="#">119</a>
01	778	23 - 24	Late Collision Count – LATECOL	<a href="#">120</a>
01	778	25 - 26	Collision Count – COLC	<a href="#">120</a>
01	778	27 - 28	Defer Count – DC	<a href="#">120</a>
01	778	29 - 30	Transmit with No CRS – TNCRS	<a href="#">121</a>
<b>PCIe Registers</b>				
01	770	16	PCIe FIFOs Control/Status	<a href="#">122</a>
01	770	17	PCIe Power Management Control	<a href="#">122</a>
01	770	18	In-Band Control	<a href="#">123</a>
01	770	20	PCIe Diagnostics	<a href="#">123</a>
01	770	21	Timeouts	<a href="#">123</a>
01	770	23	PCIe K-State Minimum Duration Timeout	<a href="#">124</a>
<b>LPI Registers</b>				
01	772	18	Low Power Idle GPIO Control	<a href="#">125</a>
01	772	20	Low Power Idle Control	<a href="#">125</a>
01	772	23	Flow Control Refresh Threshold Value – FCRTV	<a href="#">126</a>
01	772	24	Flow Control Thresholds – FCTH	<a href="#">126</a>
01	772	25	LANWAKE# Control – LANWAKEC	<a href="#">126</a>
01	772	26	Memories Power	<a href="#">127</a>
01	772	29	Configuration	<a href="#">127</a>



Table 9-1 Address Map (Continued)

PHY Address	Page	Register	Name	Link
<b>ULP Registers</b>				
01	779	16	ULP Configuration 1	<a href="#">128</a>
01	779	17	ULP Configuration 2	<a href="#">128</a>
01	779	18	ULP SW Control	<a href="#">129</a>
01	779	19	SW Control	<a href="#">129</a>
01	779	20	Off Board LAN Connected Device Control	<a href="#">130</a>
<b>General Registers</b>				
01	776	19	I219 Capability	<a href="#">131</a>
01	0	25	OEM Bits	<a href="#">131</a>
01	0	26	SMBus Address	<a href="#">132</a>
01	0	27 - 28	Shadow Receive Address Low0 – SRAL0	<a href="#">133</a>
01	0	29	Shadow Receive Address High0 – SRAH0	<a href="#">133</a>
01	0	30	LED Configuration	<a href="#">133</a>
<b>Wake Up Registers</b>				
01	800	0	Receive Control – RCTL	<a href="#">137</a>
01	800	1	Wake Up Control – WUC	<a href="#">138</a>
01	800	2	Wake Up Filter Control – WUFC	<a href="#">138</a>
01	800	3	Wake Up Status – WC	<a href="#">139</a>
01	800	16 - 17	Receive Address Low – RAL	<a href="#">140</a>
01	800	18 - 19	Receive Address High – RAH	<a href="#">140</a>
01	800	20-21 + 4*n (n=0...10)	Shared Receive Address Low – SHRAL	<a href="#">140</a>
01	800	22-23 + 4*n (n=0...8,10)	Shared Receive Address High – SHRAH	<a href="#">141</a>
01	800	58 - 59	Shared Receive Address High 9 – SHRAH[9]	<a href="#">141</a>
01	800	64	IP Address Valid – IPAV	<a href="#">141</a>
01	800	70	Proxy Control – PRXC	<a href="#">142</a>
01	800	71	Proxy Code Checksum – PRCC	<a href="#">143</a>
01	800	72	Proxy Control 2 – PRXC2	<a href="#">143</a>
01	800	75	Flex Filters Proxy Control – FFPRXC	<a href="#">143</a>
01	800	76	Wake Up Filter Control 2 – WUFC2	<a href="#">144</a>
01	800	77	Wake Up Filter Status 2 – WUS2	<a href="#">144</a>
01	800	78	Wake Up Filter Control 3 – WUFC3	<a href="#">144</a>
01	800	79	Wake Up Filter Status 3 – WUS3	<a href="#">145</a>
01	800	80	Wake Up Filter Control 4 – WUFC4	<a href="#">145</a>
01	800	81	Wake Up Filter Status 4 – WUS4	<a href="#">145</a>
01	800	82-83 + 2*n (n=0, 1, 2)	IPv4 Address Table – IP4AT	<a href="#">145</a>
01	800	88-89 + 2*n (n=0...3)	IPv6 Address Table – IP6AT[3:0]	<a href="#">145</a>



**Table 9-1 Address Map (Continued)**

PHY Address	Page	Register	Name	Link
01	800	128 - 191	Multicast Table Array – MTA[31:0]	<a href="#">146</a>
01	800	256 + 2*n (n=0...127)	Flexible Filter Value Table LSB – FFVT_01	<a href="#">147</a>
01	800	257 + 2*n (n=0...127)	Flexible Filter Value Table MSBs – FFVT_23	<a href="#">147</a>
01	800	512 + 2*n (n=0...127)	Flexible Filter Value Table – FFVT_45	<a href="#">147</a>
01	800	1024 + 2*n (n=0...127)	Flexible Filter Value Table – FFVT_67	<a href="#">148</a>
01	800	768 + n (n=0...127)	Flexible Filter Mask Table – FFMT	<a href="#">148</a>
01	800	896 + n (n=0...3)	Flexible Filter Length Table – FFLT03	<a href="#">149</a>
01	800	904 + n (n=0...1)	Flexible Filter Length Table – FFLT45	<a href="#">149</a>
01	800	908 + n (n=0...1)	Flexible Filter Length Table – FFLT67	<a href="#">149</a>
01	800	2304 + 2*n (n=0...127)	Flexible Filter Value Table 89 – FFVT_89	<a href="#">150</a>
01	800	2305 + 2*n (n=0...127)	Flexible Filter Value Table 1011 – EFFVT_1011	<a href="#">150</a>
01	800	2560 + 2*n (n=0...127)	Flexible Filter Value Table 1213 – FFVT_1213	<a href="#">150</a>
01	800	2816 + n (n=0...127)	Flexible Filter Mask Table 2 – FFMT2	<a href="#">151</a>
01	800	2944 + n (n=0...3)	Flexible Filter Length Table 891011 – FFLT891011	<a href="#">151</a>
01	800	2952 + n (n=0...1)	Flexible Filter Length Table 1213 – FFLT1213	<a href="#">151</a>
01	800	2956 + n (n=0...1)	Flexible Filter Length Table 1415 – FFLT1415	<a href="#">152</a>
01	800	3072 + 2*n (n=0...127)	Flexible Filter Value Table 1415 – FFVT_1415	<a href="#">152</a>
01	800	4352 + 2*n (n=0...127)	Flexible Filter Value Table 1617 – FFVT_1617	<a href="#">152</a>
01	800	4353 + 2*n (n=0...127)	Flexible Filter Value Table 1819 – FFVT_1819	<a href="#">152</a>
01	800	4608 + 2*n (n=0...127)	Flexible Filter Value Table 2021 – FFVT_2021	<a href="#">152</a>
01	800	4609 + 2*n (n=0...127)	Flexible Filter Value Table 2223 – FFVT_2223	<a href="#">152</a>
01	800	4864 + n (n=0...127)	Flexible Filter Mask Table 3 – FFMT3	<a href="#">153</a>
01	800	4992 + n (n=0...3)	Flexible Filter Length Table 1619 – FFLT1619	<a href="#">153</a>
01	800	5000 + n (n=0...1)	Flexible Filter Length Table 2021 – FFLT2021	<a href="#">153</a>
01	800	5004 + n (n=0...1)	Flexible Filter Length Table 2223 – FFLT2223	<a href="#">153</a>
01	800	6400 + 2*n (n=0...127)	Flexible Filter Value Table 2425 – FFVT_2425	<a href="#">153</a>
01	800	6401 + 2*n (n=0...127)	Flexible Filter Value Table 2627 – FFVT_2627	<a href="#">153</a>
01	800	6656 + 2*n (n=0...127)	Flexible Filter Value Table 2829 – FFVT_2829	<a href="#">153</a>
01	800	6657 + 2*n (n=0...127)	Flexible Filter Value Table 3031 – FFVT_3031	<a href="#">153</a>
01	800	6912 + n (n=0...127)	Flexible Filter Mask Table 4 – FFMT4	<a href="#">154</a>
01	800	7040 + n (n=0...3)	Flexible Filter Length Table 2427 – FFLT2427	<a href="#">154</a>
01	800	7048 + n (n=0...1)	Flexible Filter Length Table 2829 – FFLT2829	<a href="#">154</a>
01	800	7052 + n (n=0...1)	Flexible Filter Length Table 3031 – FFLT3031	<a href="#">154</a>
01	801	30 - 31	Management 2 Host Control Register – MANC2H	<a href="#">154</a>
01	801	32 - 33	Management 2 Host Control Register 2 – MANC2H2	<a href="#">157</a>
01	801	40 - 41	IPv4 ME Binding Control – IPV4MBC	<a href="#">158</a>
01	801	42 - 43	IPv4 Host Binding Control – IPV4HBC	<a href="#">161</a>



Table 9-1 Address Map (Continued)

PHY Address	Page	Register	Name	Link
01	801	50 - 51	IPv6 Binding Control - IPV6BC	<a href="#">163</a>
01	801	52	SHRA Filter Enable Register - SHRAFER	<a href="#">165</a>
<b>Proxy Controller uCode</b>				
01	802	0 - 1536	Proxy Micro Code	<a href="#">166</a>
<b>Host WoL Packet</b>				
01	803	0 - 63	Host WoL Packet Data	<a href="#">167</a>
01	803	64	Host WoL Packet Length	<a href="#">167</a>
01	803	66	Host WoL Packet Indication Clear	<a href="#">167</a>
<b>LPI MMD PHY Registers</b>				
02	0	16-17	I219 EMI Registers	<a href="#">168</a>



## 9.5.2 PHY Registers

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

### 9.5.2.1 Control Register - Address 0

Field Name	Bit(s)	Type	Default	Description
Reserved	5:0	RO	0x0	Reserved. Always set to 0x0.
Speed Selection (MSB)	6	RW	1b	The speed selection address 0 (bits 13 and 6) are used to configure the link manually. Setting these bits has no effect unless bit 12 (AN En) is cleared. 00b = 10 Mb/s 01b = 100 Mb/s 10b = 1000 Mb/s 11b = Reserved
Collision Test	7	RW	0b	Enables IEEE 22.2.4.1.9 collision test. 0b = Disable collision test. 1b = Enable collision test.
Duplex Mode	8	RW	1b	This bit is used to configure the link manually. Setting this bit has no effect unless bit 12 (AN En) is cleared. 0b = Half-duplex 1b = Full-duplex
Restart Auto-Negotiation	9	RW/SC	0b	0b = Normal operation, 1b = Restarts auto-negotiation process.
Isolate	10	RW	0b	Setting this bit isolates the PHY from the MII or GMII interfaces. 0b = Normal operation, 1b = Isolates the PHY from MII.
Power Down	11	RW	0b	0b = Normal operation, 1b = Power down.
Auto-Negotiation Enable	12	RW	1b	When this bit is cleared, the link configuration is determined manually. 0b = Disables auto-negotiation process. 1b = Enables auto-negotiation process.
Speed Select (LSB)	13	RW	0b	The speed selection address 0 (bits 13 and 6) are used to configure the link manually. Setting these bits has no effect unless bit 12 (AN En) is cleared. 00b = 10 Mb/s 01b = 100 Mb/s 10b = 1000 Mb/s 11b = Reserved
Loopback	14	RW	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). 0b = Disables loopback. 1b = Enables loopback.



Field Name	Bit(s)	Type	Default	Description
Reset	15	RW/SC	0b	Writing a 1b to this bit causes immediate PHY reset. Once the operation completes, this bit clears to 0b automatically. 0b = Normal operation. 1b = PHY reset.

### 9.5.2.2 Status Register - Address 1

Field Name	Bit(s)	Type	Default	Description
Extended Capability	0	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.
Jabber Detect	1	RO/LH	0b	0b = Jabber condition not detected. 1b = Jabber condition detected.
Link Status	2	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. 0b = Link is down. 1b = Link is up.
Auto-Negotiation Ability	3	RO	1b	0b = PHY not able to perform auto-negotiation. 1b = PHY able to perform auto-negotiation.
Remote Fault	4	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. 0b = Remote fault condition not detected. 1b = Remote fault condition detected.
Auto-Negotiation Complete	5	RO	0b	This bit is set after auto-negotiation completes. 0b = Auto-negotiation process not complete. 1b = Auto-negotiation process complete.
MF Preamble Suppression	6	RO	1b	1b = PHY accepts management frames with preamble suppressed.
Reserved	7	RO	0b	Reserved. Must always be set to 0b.
Extended Status	8	RO	1b	Extended status information in the register Extended Status (0xF).
100BASE-T2 Half-Duplex	9	RO	0b	Not able to perform 100BASE-T2.
100BASE-T2 Full-Duplex	10	RO	0b	Not able to perform 100BASE-T2.
10 Mb/s Half-Duplex	11	RO	1b	0b = Not 10BASE-T half duplex capable. 1b = 10BASE-T half duplex capable.
10 Mb/s Full-Duplex	12	RO	1b	0b = Not 10BASE-T full duplex capable. 1b = 10BASE-T full duplex capable.



Field Name	Bit(s)	Type	Default	Description
100BASE-TX Half-Duplex	13	RO	1b	0b = Not 100BASE-TX half duplex capable. 1b = 100BASE-TX half duplex capable.
100BASE-TX Full-Duplex	14	RO	1b	0b = Not 100BASE-TX full duplex capable. 1b = 100BASE-TX full duplex capable.
100BASE-T4	15	RO	0b	100BASE-T4. This protocol is not supported. This register bit is always set to 0b. 0b = Not 100BASE-T4 capable.

### 9.5.2.3 PHY Identifier Register 1 - Address 2

Field Name	Bit(s)	Type	Default	Description
PHY ID Number <sup>1</sup>	15:0	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.

### 9.5.2.4 PHY Identifier Register 2 - Address 3

Field Name	Bit(s)	Type	Default	Description
Revision Number	3:0	RO	0x1	The value is part of the PHY identifier and represents the Device Revision Number.
Model Number	9:4	RO	0xA	The value is part of the PHY identifier and represents the Device Model Number.
PHY ID Number <sup>1</sup>	15:10	RO	0x0	The PHY identifier composed of bits 19 through 24 of the OUI.

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



## 9.5.2.5 Auto-Negotiation Advertisement Register - Address 4

**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.

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



## 9.5.2.6 Auto-Negotiation Link Partner Ability Register - Address 5

Field Name	Bit(s)	Type	Default	Description
Protocol Selector Field	4:0	RO	0x00	Link partner protocol selector field.
10BASE-T Half-Duplex Capability	5	RO	0b	0b = Link partner is not 10BASE-T half-duplex capable. 1b = Link partner is 10BASE-T half-duplex capable.
10BASE-T Full-Duplex Capability	6	RO	0b	0b = Link partner is not 10BASE-T full-duplex capable. 1b = Link partner is 10BASE-T full-duplex capable.
100BASE-TX Half-Duplex Capability	7	RO	0b	0b = Link partner is not 100BASE-TX half-duplex capable. 1b = Link partner is 100BASE-TX half-duplex capable.
100BASE-TX Full-Duplex Capability	8	RO	0b	0b = Link partner is not 100BASE-TX full-duplex capable. 1b = Link partner is 100BASE-TX full-duplex capable.
100BASE-T4 Capability	9	RO	0b	0b = Link partner is not 100BASE-T4 capable. 1b = Link partner is 100BASE-T4 capable.
Pause Capable	10	RO	0b	0b = Link partner is not capable of pause operation. 1b = Link partner is capable of full duplex pause operation.
Asymmetric Pause	11	RO	0b	0b = Link partner does not request asymmetric pause. 1b = Link partner requests asymmetric pause.
Reserved	12	RO	0b	Reserved.
Remote Fault	13	RO	0b	0b = Link partner has not detected remote fault. 1b = Link partner has detected remote fault.
Acknowledge	14	RO	0b	0b = Link partner has not received link code word. 1b = Link partner has received link code word.
Next Page	15	RO	0b	0b = Link partner does not have next page ability. 1b = Link partner has next page ability.



## 9.5.2.7 Auto-Negotiation Expansion Register - Address 6

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

## 9.5.2.8 Auto-Negotiation Next Page Transmit Register - Address 7

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



### 9.5.2.9 Link Partner Next Page Register - Address 8

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

### 9.5.2.10 1000BASE-T Control PHY Register - Address 9

Field Name	Bit(s)	Type	Default	Description
Reserved	7:0	RO	0x00	Reserved. Set these bits to 0x00.
Advertise 1000BASE-T Half-Duplex Capability	8	RW	0b	0b = Advertises no 1000BASE-T half-duplex capability. 1b = Advertises 1000BASE-T half-duplex capability. <b>Note:</b> 1000BASE-T half-duplex not supported.
Advertise 1000BASE-T Full-Duplex Capability	9	RW	0b	0b = Advertises no 1000BASE-T full-duplex capability. 1b = Advertises 1000BASE-T full-duplex capability.
Port Type	10	RW	0b	0b = Single port device (prefer slave). 1b = Multi-port device (prefer master).
Master/Slave Configuration Value	11	RW	0b	Setting this bit has no effect unless address 9, bit 12 is set. 0b = Configures PHY as a slave. 1b = Configures PHY as a master.
Master/Slave Manual Configuration Enable	12	RW	0b	0b = Automatic master/slave configuration. 1b = Enables master/slave configuration.
Test Mode	15:13	RW	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.

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



### 9.5.2.11 1000BASE-T Status Register - Address 10

Field Name	Bit(s)	Type	Default	Description
Idle Error Count	7:0	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.
Reserved	9:8	RO	00b	Reserved.
Link Partner 1000BASE-T Half-Duplex Capability	10	RO	0b	0b = Link partner not 1000BASE-T half-duplex capable. 1b = Link partner 1000BASE-T half-duplex capable. <b>Note:</b> Logically, bits 11:10 might be regarded as an extension of the <i>Technology Ability</i> field in Register 5.
Link Partner 1000BASE-T Full-Duplex Capability	11	RO	0b	0b = Link partner not 1000BASE-T full-duplex capable. 1b = Link partner 1000BASE-T full-duplex capable. <b>Note:</b> Logically, bits 11:10 might be regarded as an extension of the <i>Technology Ability</i> field in Register 5.
Remote Receiver Status	12	RO	0b	0b = Remote receiver is incorrect. 1b = Remote receiver is correct.
Local Receiver Status	13	RO	0b	0b = Local receiver is incorrect. 1b = Local receiver is correct.
Master/Slave Configuration Resolution	14	RO	0b	Master/Slave Configuration Resolution: 0b = Local PHY resolved to slave. 1b = Local PHY resolved to master. This bit is not valid when bit 15 is set.
Master/Slave Configuration Fault	15	RO/LH/SC	0b	Master/Slave Configuration Fault: 0b = No master/slave configuration fault detected. 1b = Master/slave configuration fault detected. Once set, this bit remains set until cleared by the following actions: <ul style="list-style-type: none"> <li>• Read of Register 10 via the management interface.</li> <li>• Reset.</li> <li>• Auto-negotiation completed.</li> <li>• Auto-negotiation enabled.</li> </ul>



### 9.5.2.12 Extended Status Register - Address 15

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

### 9.5.2.13 PHY Control Register 2 - Address 18

Field Name	Bit(s)	Type	Default	Description
Reserved	1:0		0x0	Reserved.
Enable Diagnostics	2	RW	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. 0b = Disables diagnostics. 1b = Enables diagnostics.
Reserved	8:3		0x0	Reserved.
MDI/MDI-X Configuration	9	RW	0b	MDI/MDI-X Configuration: 0b = Manual MDI configuration. 1b = Manual MDI-X configuration. See <a href="#">Table 9-2</a> for further information.
Automatic MDI/MDI-X	10	RW	1b	0b = Disables automatic MDI/MDI-X configuration. 1b = Enables automatic MDI/MDI-X configuration.
Reserved	12:11			Reserved.
Count Symbol Errors	13	RW	0b	<i>Count Symbol Errors</i> (bit 13) and <i>Count False Carrier Events</i> (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (see <a href="#">Table 9-3</a> ). The default is to count CRC errors. 0b = Rx error counter counts CRC errors. 1b = Rx error counter counts symbol errors.
Count False Carrier Events	14	RW	0b	<i>Count Symbol Errors</i> (bit 13) and <i>Count False Carrier Events</i> (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (see <a href="#">Table 9-3</a> ). The default is to count CRC errors. 0b = Rx error counter does not count false carrier events. 1b = Rx error counter counts false carrier events.



Field Name	Bit(s)	Type	Default	Description
Resolve MDI/MDIX Before Forced Speed	15	RW	1b	0b = Does not resolve MDI/MDI-X configuration before forcing speed. 1b = Resolves MDI/MDI-X configuration before forcing speed.

Bit 9 of the PHY Control register manually sets the MDI/MDI-X configuration if automatic MDI-X is disabled (see Table 9-2).

**Table 9-2 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).

**Table 9-3 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.

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 Table 9-4. Note that even in manual MDI/MDI-X configuration, the PHY automatically detects and corrects for C and D pair swaps.

**Table 9-4 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+/-



## 9.5.2.14 Loopback Control Register - Address 19

Field Name	Bit(s)	Type	Default	Description
Force Link Status	0	RW	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. 0b = Forces link status not operational in MII loopback. 1b = Forces link status operational in MII loopback.
Reserved	5:1			Reserved.
Tx Suppression	6	RW	1b	0b = Do not suppress Tx during all digital loopback. 1b = Suppress Tx during all digital loopback.
External Cable	7	RW	0b	0b = External cable loopback disabled. 1b = External cable loopback enabled.
Reserved	8	RW		Reserved.
Remote	9	RW	0b	0b = Remote loopback disabled. 1b = Remote loopback enabled.
Line Driver	10	RW	0b	0b = Line driver loopback disabled. 1b = Line driver loopback enabled.
Reserved	11			Reserved.
All Digital	12	RW	1b	0b = All digital loopback disabled. 1b = All digital loopback enabled.
Reserved	14:13			Reserved.
MII	15	RW	0b	0b = MII loopback not selected. 1b = MII loopback selected.

### 9.5.2.14.1 Loopback Mode Settings

Table 9-5 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 9-5 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



### 9.5.2.15 Rx Error Counter Register - Address 20

Field Name	Bit(s)	Type	Default	Description
Rx Error Counter	15:0	RO/SC	0x00	16-bit Rx error counter. <b>Note:</b> For error type descriptions, see Register 18, bits 13 and 14 (Section 9.5.2.13).

### 9.5.2.16 Management Interface (MI) Register - Address 21

Field Name	Bit(s)	Type	Default	Description
Reserved	1:0		0x0	Reserved.
Energy-Detect Power-Down Mode Transmit Enable	2	RW	1b	0b = Disables NLP transmission during energy-detect power-down. 1b = Enables NLP transmission during energy-detect power-down.
Energy Detect Power-Down Enable	3	RW	1b	0b = Disables energy detect power-down. 1b = Enables energy detect power-down.
Reserved			0x0	Reserved.

### 9.5.2.17 PHY Configuration Register - Address 22

Field Name	Bit(s)	Type	Default	Description
Reserved	4:0		0x0	Reserved.
Transmit Clock Enable	5	RW	0b	When this bit is set, the transmit test clock is available on pin TX_TCLK. 0b = Disables output. 1b = Enables output of mixer clock (transmit clock in 1000BASE-T).
Group MDIO Mode Enable	6	RW	0b	0b = Disables group MDIO mode. 1b = Enables group MDIO mode.
Alternate Next Page	7	RO	0b	0b = Normal operation of 1000BASE-T next page exchange. 1b = Enables manual control of 1000BASE-T next pages only.
Reserved	9:8			Reserved.



Field Name	Bit(s)	Type	Default	Description
Automatic Speed Down-Shift Mode	11:10	RW	11b	If automatic down-shift 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 reattempting connection at the highest possible speed (1000BASE-T). 00b = Automatic speed down-shift disabled. 01b = 10BASE-T down-shift enabled. 10b = 100BASE-TX down-shift enabled. 11b = 100BASE-TX and 10BASE-T enabled.
Reserved	14:12	RO	0x0	Reserved.
CRS Transmit Enable	15	RW	0b	0b = Disables CRS on transmit. 1b = Enables CRS on transmit in half-duplex mode.

### 9.5.2.18 PHY Control Register - Address 23

Field Name	Bit(s)	Type	Default	Description
Force Interrupt	0	RW	0b	0b = De-asserts MDINT_N pin. 1b = Asserts MDINT_N pin.
Reserved	1		0b	Reserved.
10BASE-T Preamble Length	3:2	RW	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.
TP_LOOPBACK (10BASE-T)	4	RW	0b	0b = Normal operation. 1b = Disables TP loopback during half duplex.
SQE (10BASE-T)	5	RW	0b	0b = Disables heartbeat. 1b = Enables heartbeat.
Jabber (10BASE-T)	6	RW	1b	0b = Normal operation. 1b = Disables jabber.
Link Partner Detected	7	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). 0b = Link partner not detected. 1b = Link partner detected.
Reserved	9:8		0x0	Reserved.



Field Name	Bit(s)	Type	Default	Description
Automatic Speed Down-Shift Attempts Before Down-Shift	12:10	RW	100b	000b = 1 001b = 2 010b = 3 011b = 4 100b = 5 101b = 6 110b = 7 111b = 8
LNK_EN (Link Enable)	13	RW	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 stand-by, and does not respond to link pulses from a remote link partner. In standby, IP phone detect and TDR functions are available. 0b = Disables linking. 1b = Enables linking.
Reserved	15:14	RO	0x0	Reserved.

### 9.5.2.19 Interrupt Mask Register - Address 24

Field Name	Bit(s)	Type	Default	Description
MDINT_N Enable	0	RW	0b	0b = MDINT_N disabled. 1b = MDINT_N enabled. <sup>1</sup>
Automatic Speed Down-Shift	1	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Link Status Change	2	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Receive Status Change	3	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
FIFO Overflow/Underflow	4	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Error Count Full	5	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Next Page Received	6	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
CRC Errors	7	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Auto-Negotiation Status Change	8	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
MDIO Sync Lost	9	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.



Field Name	Bit(s)	Type	Default	Description
TDR/IP Phone	10	RW	0b	0b = Interrupt disabled. 1b = Interrupt enabled.
Reserved	15:11		0x0	Reserved.

1. MDINT\_N is asserted (active low) if MII interrupt pending = 1b.

## 9.5.2.20 Interrupt Status Register - Address 25

Field Name	Bit(s)	Type	Default	Description
MII Interrupt Pending	0	RO/LH	0b	An event has occurred and the corresponding interrupt mask bit is enabled (set to 1b). 0b = No interrupt pending. 1b = Interrupt pending.
Automatic Speed Down-Shift	1	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
Link Status Change	2	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
Receive Status Change	3	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
FIFO Overflow/Underflow	4	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
Error Count Full	5	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
Next Page Received	6	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
CRC Errors	7	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
Auto-Negotiation Status Change	8	RO/LH	0b	0b = Event has not occurred. 1b = Event has occurred.
MDIO Sync Lost	9	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. 0b = Event has not occurred. 1b = Event has occurred.
TDR/IP Phone	10	RO/LH	0b	0b = Event has not completed. 1b = Event completed.
Reserved	15:11		0x0	Reserved.



## 9.5.2.21 PHY Status Register - Address 26

Field Name	Bit(s)	Type	Default	Description
Link Partner Advertised Asymmetric PAUSE	0	RO	0b	0b = Link partner did not advertise asymmetric PAUSE. 1b = Link partner advertised asymmetric PAUSE.
Link Partner Advertised PAUSE	1	RO	0b	0b = Link partner did not advertise PAUSE. 1b = Link partner advertised PAUSE.
Auto-Negotiation Enabled	2	RO	0b	0b = Both partners do not have autonegotiation enabled. 1b = Both partners have auto-negotiation enabled.
Collision Status	3	RO	0b	0b = Collision not occurring. 1b = Collision occurring.
Receive Status	4	RO	0b	0b = PHY not receiving a packet. 1b = PHY receiving a packet.
Transmit Status	5	RO	0b	0b = PHY not transmitting a packet. 1b = PHY transmitting a packet.
Link Status	6	RO	0b	0b = Link down. 1b = Link up.
Duplex Status	7	RO	0b	0b = Half duplex. 1b = Full duplex.
Speed Status	9:8	RO	11b	00b = 10BASE-T 01b = 100BASE-TX. 10b = 1000BASE-T. 11b = Undetermined.
Polarity Status	10	RO	1b	0b = Polarity normal (10BASE-T only). 1b = Polarity inverted (10BASE-T only).
Pair Swap on Pairs A and B	11	RO	0b	0b = Pairs A and B not swapped. 1b = Pairs A and B swapped.
Auto-Negotiation Status	12	RO	0b	0b = Auto-negotiation not complete. 1b = Auto-negotiation complete.
Auto-Negotiation Fault Status	14:13	RO	00b	00b = No auto-negotiation fault. 01b = Parallel detect auto-negotiation fault. 10b = Master/slave auto-negotiation fault. 11b = Reserved.
PHY in Standby	15	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). 0b = PHY not in standby mode. 1b = PHY in standby mode.



## 9.5.2.22 LED Control Register 1 - Address 27

Field Name	Bit(s)	Type	Default	Description
Pulse Stretch 0	0	RW	1b	0b = Disables pulse stretching of LED functions: transmit activity, receive activity, and collision. 1b = Enables pulse stretching of LED functions: transmit activity, receive activity, and collision.
LED Output Disable	1	RW	0b	The default value of this bit is set by the LED_CFG pin at reset. 0b = Enables LED outputs. 1b = Disables LED outputs.
LED Pause Duration	3:2	RW	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.
LED Blink Pattern Pause	7:4	RW	0x0	LED blink pattern pause cycles.
Reserved	9:8		0x0	Reserved.
LED_LNK/ACT Extended Modes	10	RW	0b	The LED function is programmed using this bit and Register 28. 0b = Standard modes for LED_LNK/ACT. 1b = Extended modes for LED_LNK/ACT.
LED_1000 Extended Modes	11	RW	0b	The LED function is programmed using this bit and Register 28. 0b = Standard modes for LED_1000. 1b = Extended modes for LED_1000.
LED_100 Extended Modes	12	RW	0b	The LED function is programmed using this bit and Register 28. 0b = Standard modes for LED_100. 1b = Extended modes for LED_100.
LED_10 Extended Modes	13	RW	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. 0b = Standard modes for LED_10. 1b = Extended modes for LED_10.
Two-Color Mode LED_LNK/ACT/LED_1000	14	RW	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. 0b = Normal mode for LED_LNK/ACT and LED_1000. 1b = Two-color mode for LED_LNK/ACT and LED_1000.



Field Name	Bit(s)	Type	Default	Description
Two-Color Mode LED_100/LED_10	15	RW	0b	<p>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.</p> <p>0b = Normal mode for LED_100 and LED_10. 1b = Two-color mode for LED_100 and LED_10.</p>

### 9.5.2.23 LED Control Register 2 - Address 28

Field Name	Bit(s)	Type	Default	Description
LED_LNK/ACT	3:0	R/W	LED_CFG	<p><b>Standard Modes:</b></p> <p>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> <p><b>Extended modes:</b></p> <p>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). All other values are reserved.</p>
LED_100	7:4	R/W	LED_CFG	See description for bits 3:0.
LED_100	11:8	R/W	LED_CFG	See description for bits 3:0
LED_10	15:12	R/W	LED_CFG	See description for bits 3:0



### 9.5.2.24 LED Control Register 3 - Address 29

Field Name	Bit(s)	Type	Default	Description
LED Blink Pattern	7:0	RW	0x55	LED Blink Pattern The default pattern is 512 ms blink.
LED Blink Pattern Frequency	13:8	RW	0x1F	LED Blink Pattern clock frequency divide ratio. The default pattern is 512 ms blink.
LED Blink Pattern Address	15:14	RW	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.

### 9.5.2.25 Late Collision Counter - Address 30

Field Name	Bit(s)	Type	Default	Description
LateColCnt02	7:0	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.
LateColCnt13	15:8	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.



## 9.5.2.26 Link Polarity and Length Status - Address 31

Field Name	Bit(s)	Type	Default	Description
B1000DscrAcqErr	0	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
Reserved	1	RO	0b	Reserved.
DspCblLen	9:2	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.
Pair0PolInv	10	RO	0b	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.
Pair1PolInv	11	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.
Pair2PolInv	12	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.
Pair3PolInv	13	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.
Pair23Swap	14	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.
LateColCnt23Sel	15	RO	0b	This bit indicates that late collision counters 2 and 3 are selected for read. See <i>LateColCnt02</i> and <i>LateColCnt13</i> for more information.



## 9.5.3 Port Control Registers

### 9.5.3.1 Custom Mode Control PHY Address 01, Page 769, Register 16

Field Name	Bit(s)	Type	Default	Description
Reserved	9:0	RW	0x180	Reserved.
MDIO frequency access	10	RW	0b	0b = normal MDIO frequency access 1b = reduced MDIO frequency access (required for read during cable disconnect)
Reserved	15:11	RW	0x04	Reserved.

### 9.5.3.2 Port General Configuration PHY Address 01, Page 769, Register 17

Field Name	Bit(s)	Type	Default	Description
Reserved	1:0	RO	00b	Reserved.
MACPD_enable	2			Written as 1b when the integrated LAN controller needs to globally enable the MAC power down feature while the I219 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.
Wakeup clocks stop	3			Wake-up clocks are stopped while wake up is disabled.
Host_WU_Active	4			Enables host wake up from the I219. This bit is reset by power on reset only.
Reserved	5	RW	1b	Reserved.
Active_PD_enable	6			Active Power Down Enable (sD3 Enable) When set to 1b, The MAC needs to enter integrated LAN controller power down mode.
Reserved	7	RW	0b	Reserved.
BP extension Wait	10:8	RW	000b	Additional waiting byte times after TX Gate Wait IPG expires until the <i>Back Pressure In-band</i> bit is cleared.
Tx Gate Wait IFS	15:11	RW	00111b	Determines the size (in nibbles) of non-deferring window from CRS de-assertion.

**Note:** Register resets on Power Good only.



### 9.5.3.3 Power Management Control Register PHY Address 01, Page 769, Register 21

Field Name	Bit(s)	Type	Default	Description
Retry late collision	0	RW	0b	Retry late collision.
Collision threshold	8:1	RW	0x0F	Number of retries for a collided packet.
Reserved	15:9	RO	0x00	Reserved. Write to 0x00

### 9.5.3.4 SMBus Control Register PHY Address 01, Page 769, Register 23

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

### 9.5.3.5 Rate Adaptation Control Register PHY Address 01, Page 769, Register 25

Field Name	Bit(s)	Type	Default	Description
read_delay_fd	4:0	RWP	10001b	Reserved. Write as read.
rx_flip_bad_sfd	5	RW	1b	Align the packet's start of frame delimiter to a byte boundary in the receive path.
Reserved	6	RW	0b	Reserved. Write as read.
rx_en_crs_preamble	7	RW	0b	Enable generation of early preamble based on CRS in the receive path.
rx_en_rxdv_preamble	8	RW	1b	Enable generation of early preamble based on RX_DV in the receive path.
Reserved	15:9			Reserved. Write as read.



### 9.5.3.6 Flow Control Transmit Timer Value PHY Address 01, Page 769, Register 27

Field Name	Bit(s)	Type	Default	Description
Flow Control Transmit Timer Value	15:0	RW	0x0000	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.

**Note:** Register resets on Power Good only.

### 9.5.3.7 System Low Power Control – SxCTRL PHY Address 01, Page 769, Register 28

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

**Note:** Register resets on Power Good only.

### 9.5.3.8 SERDES MDI Control Register – SMDIC PHY Address 01, Page 774, Registers 23-24

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

Field Name	Bit(s)	Type	Default	Description
DATA	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.
REGADD	20:16	RW	00000b	LAN Connected Device Register Address (REGADD) Reg 0, 1, 2, ... 31.
PHYADD	25:21	RW	00000b	LAN Connected Device Address (PHYADD)



Field Name	Bit(s)	Type	Default	Description
OP	27:26	RW	00b	Op-Code (OP) 00b = Reserved. 01b = MDI Write. 10b = MDI Read. 11b = Reserved.
R	28	RW	1b	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.
ST	30:29	RW	00b	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 = Reserved. 11b = Reserved.
Reserved	31	RO	0b	Reserved.



## 9.5.4 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.

### 9.5.4.1 Single Collision Count – SCC PHY Address 01, Page 778, Registers 16-17

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 I219 is in half-duplex mode.

Field Name	Bit(s)	Type	Default	Description
SCC	31:0	RO/V	0x00	Number of times a transmit encountered a single collision.

### 9.5.4.2 Excessive Collisions Count – ECOL PHY Address 01, Page 778, Register 18-19

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 I219 is in half-duplex mode.

Field Name	Bit(s)	Type	Default	Description
ECC	31:0	RO/V	0x00	Number of packets with more than 16 collisions.

### 9.5.4.3 Multiple Collision Count – MCC PHY Address 01, Page 778, Register 20-21

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 I219 is in half-duplex mode.

Field Name	Bit(s)	Type	Default	Description
MCC	31:0	RO/V	0x00	Number of times a successful transmit encountered multiple collisions.



### 9.5.4.4 Late Collisions Count – LATECOL PHY Address 01, Page 778, Register 23-24

Late collisions are collisions that occur after one slot time. This register only increments if transmits are enabled and the I219 is in half-duplex mode.

Field Name	Bit(s)	Type	Default	Description
LCC	31:0	RO/V	0x00	Number of packets with late collisions.

### 9.5.4.5 Collision Count – COLC PHY Address 01, Page 778, Register 25-26

This register counts the total number of collisions seen by the transmitter. This register only increments if transmits are enabled and the I219 is in half-duplex mode. This register applies to clear as well as secure traffic.

Field Name	Bit(s)	Type	Default	Description
COLC	31:0	RO/V	0x00	Total number of collisions experienced by the transmitter.

### 9.5.4.6 Defer Count – DC PHY Address 01, Page 778, Register 27-28

This register counts defer events. A defer event occurs when the transmitter cannot immediately send a packet due to the medium busy either because another device is transmitting, the IPG timer has not expired, half-duplex deferral events, reception of XOFF frames, or the link is not up. This register only increment if transmits are enabled. The behavior of this counter is slightly different in the I219 relative to the 82542. For the I219, this counter does not increment for streaming transmits that are deferred due to TX IPG.

Field Name	Bit(s)	Type	Default	Description
CDC	31:0	RO/V	0x00	Number of defer events.



### 9.5.4.7 Transmit with No CRS – TNCRS PHY Address 01, Page 778, Register 29-30

This register counts the number of successful packet transmission in which the CRS input from the I219 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 I219.

The I219 should assert CRS during every transmission. Failure to do so might indicate that the link has failed, or the I219 has an incorrect link configuration. This register only increments if transmits are enabled. This register is only valid when the I219 is operating at half duplex.

Field Name	Bit(s)	Type	Default	Description
TNCRS	31:0	RO/V	0x00	Number of transmissions without a CRS assertion from the I219.



## 9.5.5 PCIe Registers

### 9.5.5.1 PCIe FIFOs Control/Status PHY Address 01, Page 770, Register 16

Field Name	Bit(s)	Type	Default	Description
Reserved	5:0	RO	000000b	Reserved.
Tx FIFO Overflow	6	RO/SC	0b	Tx FIFO overflow occurred.
Reserved	7	RO	0b	Reserved.
Rx FIFO Overflow	8	RO/SC	0b	Rx FIFO overflow occurred.
Reserved	9:15	RO	0000001b	Reserved.

### 9.5.5.2 PCIe Power Management Control PHY Address 01, Page 770, Register 17

Field Name	Bit(s)	Type	Default	Description
Enable Electrical Idle in Cable Disconnect	0	RW	0b	Consider reserved (no entry to Electrical Idle due to Cable Disconnect).
Reserved	4:1	RW	0010b	Reserved.
Reserved	6:5	RW	00b	Reserved.
PLL stop in K1	7	RW	1b	Enables stopping SerDes PLL in K1 state (in 10 Mb/s and 100 Mb/s). 0b = Enable 1b = Disable
PLL stop in K1 giga	8	RW	0b	Enables stopping the SERDES PLL in K1 state (in 1 Gbps) 0b = Disable 1b = Enable
Request a PCIe clock in K1	9	RW	1b	Use CLK_REQ to request PCIe clock in K1.
Reserved	12:10	RO	100b	Reserved.
Giga_K1_disable	13	RW	0b	When set, the I219 does not enter K1 while link speed at 1000 Mb/s.
K1 enable <sup>1</sup>	14	RW	0b	Enable K1 Power Save Mode: 0b = Disable 1b = Enable
Reserved	15	RO	1b	Reserved.

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.



### 9.5.5.3 In-Band Control PHY Address 01, Page 770, Register 18

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

**Note:** All in-band timeouts are multiplied by 1000 while in SMBus mode.

### 9.5.5.4 PCIe Diagnostic PHY Address 01, Page 770, Register 20

Field Name	Bit(s)	Type	Default	Description
In-band status acknowledge timeout	5:0	RW	0x04	Timeout in microseconds for receiving an acknowledge for an in-band status message.
Reserved	7:6	RW	0x0	Reserved.
Inband MDIO acknowledge timeout	15:8	RW	0x55	Timeout in microseconds for receiving acknowledge for an in-band MDIO message.

**Note:** All in-band timeouts are multiplied by 1000 while in SMBus mode.

### 9.5.5.5 Timeouts PHY Address 01, Page 770, Register 21

Field Name	Bit(s)	Type	Default	Description
Reserved	5:0	RW	0101000b	Reserved.
K1 exit timeout	11:6	RW	0101000b	These bits define how much time IDLE symbols are sent on the TX pair after exiting from K1 state before the I219 starts sending data to the integrated LAN controller (each bit represents 80 ns).
Reserved	15:12	RWP	0000b	Reserved. Write as read.

**Note:** All in-band timeouts are multiplied by 1000 while in SMBus mode.



## 9.5.5.6 PCIe Kstate Minimum Duration Timeout PHY Address 01, Page 770, Register 23

Field Name	Bit(s)	Type	Default	Description
EL_min_dur timeout	4:0	RW	0x10	These bits define the minimum time the I219 stays in electrical idle state once entered (each bit represents 80 ns).
Reserved	15:5	RWP	0x00	Reserved. Write as read

**Note:** All in-band timeouts are multiplied by 1000 while in SMBus mode.



## 9.5.6 LPI Registers

### 9.5.6.1 Low Power Idle GPIO Control PHY Address 01, Page 772, Register 18

Field Name	Bit(s)	Type	Default	Description
Reserved	9:0	RW	0xC9	Reserved.
TX_LPI_GPIO0	10	RW	0x0	Route Tx LPI indication to GPIO 0.
Auto EN LPI	11	RW	0x0	Auto Enable LPI after link up. When set to 0x1 772.20[14:13] will be automatically set by HW after link up.
Reserved	15:12	RW	0x40	Reserved.

**Note:** Register resets on Power Good only.

### 9.5.6.2 Low Power Idle Control PHY Address 01, Page 772, Register 20

Field Name	Bit(s)	Type	Default	Description
Reserved	3:0	RO	0x0	Reserved.
PLLLockCnt	6:4	RW	0x2	PLL Lock Counter iV 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	8:7	RW	0x0	Reserved.
PostLPICount	11:9	RW	0x1	Post LPI Counter. When in LPI active and an Ethernet packet of an in-band 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.
ForceLPI	12	RW	0x0	Force LPI Entry When set to 1b by software the PHY enters LPI mode even when not in K1.
100Enable	13	RW	0x0	100Enable Enable EEE on 100 Mb/s link speed. This bit auto clears on link down.



Field Name	Bit(s)	Type	Default	Description
1000Enable	14	RW	0x0	1000Enable Enable EEE on 1 Gb/s link speed. This bit auto clears on link down.
Reserved	15	RW	0x0	Reserved.

### 9.5.6.3 Flow Control Refresh Threshold Value – FCRTV PHY Address 01, Page 772, Register 23

Field Name	Bit(s)	Type	Default	Description
FCRTV	15:0	RW	0x0000	Flow Control Refresh Threshold Value (FCRTV) 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.

**Note:** Register resets on Power Good only.

### 9.5.6.4 Flow Control Thresholds – FCTH PHY Address 01, Page 772, Register 24

Field Name	Bit(s)	Type	Default	Description
HTHRSH	7:0	RW	0x00	High threshold for sending XOFF (in units of 32 bytes).
LTHRSH	15:8	RW	0x00	Low threshold for sending XON (in units of 32 bytes).

**Note:** Register resets on Power Good only.

### 9.5.6.5 LANWAKE# Control – LANWAKEC PHY Address 01, Page 772, Register 25

Field Name	Bit(s)	Type	Default	Description
Rx FIFO empty threshold	3:0	RW	0xA	For flow control in SMB.
Store Host WoL packet	4	RW	0x1	Store the Host WoL packet in the Proxy RAM.
Use LANWAKE pin	5	RW	0x1	Always use LANWAKE pin to indicate WoL.
Clear LANWAKE pin	6	RW	0x0	1b = Clear the LANWAKE pin. This bit is auto cleared to 0b.



Field Name	Bit(s)	Type	Default	Description
Use Legacy wake	7	RW	0x0	1b = Use 82679 WoL flows.
Reserved	8	RW	0x0	Reserved.
Filter on 1Gbps	9	RW	0x0	Filter packets in 1G for jumbo fix.
Crop ANM DA Jumbo	10	RW	0x0	Crop ANM DA in Jumbo fix.
Crop BCST DA in Jumbo	11	RW	0x0	Crop BCST DA in Jumbo fix.
Enable WoL if no packet capture	12	RW	0x1	0b = No capture means no WoL. 1b = WoL is enabled, even if the WoL packet is not captured.
Reserved	15:13	RW	0x6	Reserved.

**Note:** Register resets on Power Good only.

### 9.5.6.6 Memories Power PHY Address 01, Page 772, Register 26

Field Name	Bit(s)	Type	Default	Description
Reserved	3:0	RW	0xF	Reserved.
DIS_SMB_REL_ON_LCD	4	RW	1b	Disable the SMB release on LCD reset.
Reserved	11:15	RW	0x7F	Reserved.
MOEM	12	RW	0b	Mask OEM bits/Gig Disable/restart AN bits impact.
Reserved	15:13	RW	0x1	Reserved.

### 9.5.6.7 Configuration Register PHY Address 01, Page 772, Register 29

Field Name	Bit(s)	Type	Default	Description
Reserved	0	RW	0b	Reserved.
ENMTAONPWRGD	1	RW	0b	Enable MTA to reset only on power good.
Reserved	15:2	RW	0x3FC0	Reserved.



## 9.5.7 ULP Registers

### 9.5.7.1 ULP Configuration 1 PHY Address 01, Page 779, Register 16

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

### 9.5.7.2 ULP Configuration 2 PHY Address 01, Page 779, Register 17

Field Name	Bit(s)	Type	Default	Description
MESHADOW	4:0	RW	0x0	ME 3.3v Shadow configuration.
Reserved	15:5	RW	0x00	Reserved.



### 9.5.7.3 ULP SW Control PHY Address 01, Page 779, Register 18

Field Name	Bit(s)	Type	Default	Description
SW BUS SELECTOR	3:0	RW	0x0	Select the bus address from the ULP bus table.
SW GEN WRITE ENABLE	4	RW	0b	Cleared by HW after implement the current write enable by software_write_en_to_ulp signal.
ULP SW WRITE FROM REGISTER	5	RW	0b	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.
ULP SW READ TO REGISTER	6	RW	0x0	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.
Reserved	7	RW	0b	Reserved.
ULP MAILBOX	15:8	RW	0x00	This register is write-able and readable. Can be used by the software by read modify write.

### 9.5.7.4 SW Control PHY Address 01, Page 779, Register 19

Field Name	Bit(s)	Type	Default	Description
str_wol_pkt_feautre_en	0	RW	0b	Store WoL packet - Feature Enable Cleared on smb_sel (PCIe Reset). <b>Note:</b> RX gating would occur from receiving the WoL packet until Driver reads the WoL packet (unless gating timeout expired or stop gate bit is set).
sw_rd_pkt	1	RW	0b	Write indication from SW: Driver is ready to read the WoL packet. This bit is cleared upon one of following: 1. Reading of WoL packet from FIFO is done. 2. Gating timeout is expired. 3. stop_gate_rx_due2flex is set.
str_wol_pkt_no_ind_from_filtr_timeout	5:2	RW	0x3	How Many cycles after write to FIFO ends, last EOP address would be sampled as candidate for the WoL packet start address. Granularity of 8 cycles. <b>Note:</b> If the packet arrives without SFD, there is no indication from the filters. This timeout is needed to determine whether to sample the EOP address. <b>Note:</b> Cycles in wr-dk.



Field Name	Bit(s)	Type	Default	Description
str_wol_pkt_gating_timeout	9:6	RW	0x2	Timeout for RX gating. Granularity of 300 ms.
Reserved	13:10	RO	0x0	Reserved.
stop_gate_rx_due2flex	14	W1S	0b	When set, gating RX due2flex WoL packet stops. This bit is cleared by HW after gating ends. <b>Note:</b> Driver can only set this bit.
gating_status	14	RO	0b	Asserted while RX gating-due-to-flex-WoL-pkt is ongoing.

### 9.5.7.5 Off Board LAN Connected Device Control PHY Address 01, Page 779, Register 20

Field Name	Bit(s)	Type	Default	Description
BCN_DUR	3:0	RW	0x2	Beacon Duration (BCN_DUR) Defines the time of a single beacon cycle. Granularity in 0.5 $\mu$ s. 0x0 is not a valid setting.
BCN_INTER	7:4	RW	0x4	Beacon Interval (BCN_INTER) Defines the time interval between beacons. Granularity in 50 ms. 0x0 is not a valid setting.
REF_DLY	11:8	RW	0x1	Reference Clock Delay (REF_DLY) Defines the time K1 exit is delayed, waiting for reference clock. Granularity in 1 $\mu$ s.
Reserved	14:12	RW	0x0	Reserved.
OBLCD_En	15	RW	0b	OBLDC Enable (OBLCD_En) Enables the beacon feature for OBLCD connection indication.



## 9.5.8 General Registers

### 9.5.8.1 I219 Capability PHY Address 01, Page 776, Register 19

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

### 9.5.8.2 OEM Bits PHY Address 01, Page 0, Register 25

Field Name	Bit(s)	Type	Default	Description
Reserved	1:0	RW	00b	Reserved.
rev_aneg	2	RW	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.



Field Name	Bit(s)	Type	Default	Description
Reserved	5:3	RW		Reserved.
a1000_dis	6	RW	0b <sup>1</sup>	When set to 1b, 1000 Mb/s speed is disabled.
Reserved	9:7	RW	000b	Reserved.
Aneg_now	10	RW	0b	Restart auto-negotiation. This bit is self clearing.
Reserved	15:11	RW	00000b	Reserved.

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

### 9.5.8.3 SMBus Address PHY Address 01, Page 0, Register 26

Field Name	Bit(s)	Type	Default	Description
SMBus Address	6:0	RW	0x00	This is the integrated LAN controller SMBus address. The I219 uses it for master functionality.
SMBus Address Valid	7	RW	0b	This bit is written by the integrated LAN controller when the SMBus Address field is updated. The I219 cannot send SMBus transactions to the integrated LAN controller unless this bit is set. 0b = Address not valid. 1b = SMBus address valid.
SMBus Frequency Low	8	RW	0b	SMBus Frequency Low Together with <i>SMBus Frequency High</i> (bit 12) defines the SMBus frequency: <b>High Low Frequency</b> 0b 0b 100 KHz 0b 1b 400 KHz 1b 0b 1000 KHz 1b 1b Reserved
PEC Enable	9	RW	1b	Defines if the I219 supports PEC on the SMBus.
APM Enable	10	RW	0b	APM WoL enable.
SMB fragments size	11	RW	0b	Select SMBus Fragments Size: 0b = Fragment size is 32 bytes. 1b = Fragment size is 64 bytes.
SMBus Frequency High	12	RW	0b	SMBus Frequency High See description for <i>SMBus Frequency Low</i> (bit 8)
Reserved	15:13	RO	000b	Reserved.

**Note:** This register is reset only on internal power on reset.



### 9.5.8.4 Shadow Receive Address Low0 – SRAL0 PHY Address 01, Page 0, Registers 27-28

Field Name	Bit(s)	Type	Default	Description
RAL	31:0	R/W	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.

### 9.5.8.5 Shadow Receive Address High0 – SRAH0 PHY Address 01, Page 0, Registers 29

Field Name	Bit(s)	Type	Default	Description
RAH	15:0	R/W	X	Shadow 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.

### 9.5.8.6 LED Configuration PHY Address 01, Page 0, Register 30

Field Name	Bit(s)	Type	Default	Description
LED0 Mode	2:0	RW	100b	Mode specifying what event/state/pattern is displayed on LED0.
LED0 Invert	3	RW	0b	LED0_IVRT Field: 0b = Active low output. 1b = Active high output.
LED0 Blink	4	RW	1b	LED0_BLINK Field: 0b = No blinking. 1b = Blinking.
LED1 Mode	7:5	RW	111b	Mode specifying what event/state/pattern is displayed on LED1.
LED1 Invert	8	RW	0b	LED1_IVRT Field: 0b = Active low output. 1b = Active high output.
LED1 Blink	9	RW	0b	LED1_BLINK Field: 0b = No blinking. 1b = Blinking.
LED2 Mode	12:10	RW	110b	Mode specifying what event/state/pattern is displayed on LED2.
LED2 Invert	13	RW	0b	LED2_IVRT Field: 0b = Active low output. 1b = Active high output.



Field Name	Bit(s)	Type	Default	Description
LED2 Blink	14	RW	0b	LED2_BLINK Field: 0b = No blinking. 1b = Blinking.
Blink rate	15	RW	0b	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.

**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 9-6 LED Modes**

Mode	Selected Mode	Source Indication
000	Link 10/1000	Asserted when either 10 or 1000 Mb/s link is established and maintained.
001	Link 100/1000	Asserted when either 100 or 1000 Mb/s 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 10 Mb/s link is established and maintained.
110	Link 100	Asserted when a 100 Mb/s link is established and maintained.
111	Link 1000	Asserted when a 1000 Mb/s link is established and maintained.

### 9.5.8.7 Interrupts

The I219 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.



## 9.5.9 Wake Up Registers

### 9.5.9.1 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.

#### 9.5.9.1.1 Wake Area Read Cycle

For the I219 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 I219'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 I219'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 I219'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)



### 9.5.9.1.2 Wake Area Write Cycle

For the I219, 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 I219'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 I219'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 I219's address from the MDI register
  - d. REGADD = 0x12 (data cycle for write)
  - e. DATA = YYYY (data to be written to the register)