

Intel® 82579 Gigabit Ethernet PHY

Datasheet v2.1

Product Features

■ General

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

■ Security & Manageability

- Intel® vPro support with appropriate Intel chipset components (82579LM SKU)
- MACSec support¹

■ Performance

- Jumbo Frames (up to 9 kB)²
- 802.1Q & 802.1p
- Receive Side Scaling (RSS)
- Two Queues (Tx & Rx)

■ Power

- Flexible power configuration: use either the PCH 1.05 Vdc or iSVR.
- Reduced power consumption during normal operation and power down modes
- Integrated Intel® Auto Connect Battery Saver (ACBS)
- Single-pin LAN Disable for easier BIOS implementation
- Fully integrated Switching Voltage Regulator (iSVR)
- Low Power LinkUp (LPLU)

■ MAC/PHY Interconnect

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

■ Package/Design

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

1. MACSec is not compatible with Intel® Active Management Technology

2. Jumbo Frames are not compatible with MACSec.



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

IMPORTANT - PLEASE READ BEFORE INSTALLING OR USING INTEL® PRE-RELEASE PRODUCTS.

Please review the terms at http://www.intel.com/netcomms/prerelease_terms.htm carefully before using any Intel® pre-release product, including any evaluation, development or reference hardware and/or software product (collectively, "Pre-Release Product"). By using the Pre-Release Product, you indicate your acceptance of these terms, which constitute the agreement (the "Agreement") between you and Intel Corporation ("Intel"). In the event that you do not agree with any of these terms and conditions, do not use or install the Pre-Release Product and promptly return it unused to Intel.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The 82579 GbE PHY may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

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

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

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

Intel and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2009-2011, Intel Corporation. All Rights Reserved.



Contents


1.0	Introduction	1
1.1	Scope	1
1.2	Overview	1
1.3	Main Flows	2
1.4	References	3
1.5	Product SKUs, Codes, and Device IDs	3
1.6	Supported Operating Systems	4
2.0	Interconnects	7
2.1	Introduction	7
2.2	PCIe-Based	7
2.3	SMBus	8
2.4	Transitions between SMBus and PCIe interfaces	9
2.5	Intel® 6 Series Express Chipset/82579 – SMBus/PCIe Interconnects	10
3.0	Pin Interface	13
3.1	Pin Assignment	13
4.0	Package	17
4.1	Package Type and Mechanical	17
4.2	Package Electrical and Thermal Characteristics	18
4.3	Power and Ground Requirements	18
4.4	Ball Mapping	19
5.0	Initialization	23
5.1	Power Up	23
5.2	Reset Operation	24
5.3	Timing Parameters	25
6.0	Power Management and Delivery	27
6.1	Power Delivery	28
6.2	Power Management	28
7.0	Device Functionality	35
7.1	Tx Flow	35
7.2	Rx Flow	35
7.3	Flow Control	36
7.4	Wake Up	38
7.5	Network Proxy Functionality	46
7.6	Loopback	52
8.0	Electrical and Timing Specifications	53
8.1	Introduction	53
8.2	Operating Conditions	53
8.3	Power Delivery	54
8.4	I/O DC Parameter	56
8.5	Discrete/Integrated Magnetics Specifications	60
8.6	Mechanical	60
8.7	Oscillator/Crystal Specifications	60
9.0	Programmer's Visible State	63
9.1	Terminology	63
9.2	MDIO Access	64
9.3	Addressing	64
9.4	Address Map	65
9.5	PHY Registers (Page 0)	67
9.6	Port Control Registers (Page 769)	88



9.7	Statistics Registers	89
9.8	PCIe Registers.....	91
9.9	General Registers	94
9.10	Wake Up Registers.....	97
9.11	LPI MMD PHY Registers.....	107
10.0	Non-Volatile Memory (NVM)	109
10.1	Introduction	109
10.2	NVM Programming Procedure Overview	109
10.3	LAN NVM Format and Contents	111
11.0	Time Synch (IEEE1588 and 802.1AS)	129
11.1	Overview	129
12.0	Intel® 6 Series Express Chipset MAC Programming Interface.....	137
13.0	Reference Schematics	225
14.0	Schematic and Board Layout Checklists	227
15.0	Models	229
16	Design Considerations and Guidelines (Non-Mobile Designs)	232
16.1	PHY Overview.....	233
16.2	Platform LAN Design Guidelines	236
16.3	PCH – SMBus/PCIe LOM Design Guidelines	244
16.4	SMBus Design Considerations	245
16.5	General Layout Guidelines.....	246
16.6	Layout Considerations	246
16.7	Guidelines for Component Placement.....	246
16.8	MDI Differential-Pair Trace Routing for LAN Design	249
16.9	Signal Trace Geometry	249
16.10	Trace Length and Symmetry	252
16.11	Impedance Discontinuities	253
16.12	Reducing Circuit Inductance	253
16.13	Signal Isolation	253
16.14	Power and Ground Planes	254
16.15	Traces for Decoupling Capacitors	256
16.16	Ground Planes under a Magnetics Module.....	256
16.17	Light Emitting Diodes	258
16.18	Frequency Control Device Design Considerations.....	258
16.19	Crystals and Oscillators	259
16.20	Quartz Crystal	259
16.21	Fixed Crystal Oscillator	259
16.22	Crystal Selection Parameters.....	260
16.23	Vibrational Mode.....	260
16.24	Nominal Frequency	260
16.25	Frequency Tolerance	260
16.26	Temperature Stability and Environmental Requirements.....	260
16.27	Calibration Mode.....	261
16.28	Load Capacitance.....	261
16.29	Shunt Capacitance	262
16.30	Equivalent Series Resistance	262
16.31	Drive Level	262
16.32	Aging	262
16.33	Reference Crystal	262
16.34	Oscillator Support.....	264
16.35	Oscillator Placement and Layout Recommendations	264




16.36	Troubleshooting Common Physical Layout Issues	265
16.37	Power Delivery	266
16.38	82579 Power Sequencing	267
17	Design Considerations and Guidelines (Mobile Designs)	268
17.1	PHY Overview	269
17.2	Platform LAN Design Guidelines	272
17.3	PCH – SMBus/PCIe LOM Design Guidelines	281
17.4	SMBus Design Considerations	282
17.5	General Layout Guidelines	283
17.6	Layout Considerations	283
17.7	Guidelines for Component Placement	283
17.8	MDI Differential-Pair Trace Routing for LAN Design	286
17.9	Signal Trace Geometry	286
17.10	Trace Length and Symmetry	288
17.11	Impedance Discontinuities	289
17.12	Reducing Circuit Inductance	289
17.13	Signal Isolation	289
17.14	Power and Ground Planes	290
17.15	Traces for Decoupling Capacitors	292
17.16	Ground Planes under a Magnetics Module	292
17.17	Light Emitting Diodes	294
17.18	Considerations for Layout	294
17.19	Frequency Control Device Design Considerations	295
17.20	Crystals and Oscillators	295
17.21	Quartz Crystal	295
17.22	Fixed Crystal Oscillator	296
17.23	Crystal Selection Parameters	296
17.24	Vibrational Mode	296
17.25	Nominal Frequency	296
17.26	Frequency Tolerance	297
17.27	Temperature Stability and Environmental Requirements	297
17.28	Calibration Mode	297
17.29	Load Capacitance	298
17.30	Shunt Capacitance	298
17.31	Equivalent Series Resistance	298
17.32	Drive Level	299
17.33	Aging	299
17.34	Reference Crystal	299
17.35	Oscillator Support	300
17.36	Oscillator Placement and Layout Recommendations	301
17.37	LAN Switch	302
17.38	Troubleshooting Common Physical Layout Issues	302
17.39	Power Delivery	303
17.40	Power Sequencing	304



Revision History

Date	Revision	Description
May 2009	0.5	Initial release (Intel Confidential).
January 2010	0.7	Major revision (all sections).
April 2010	1.0	Updated all 82579 references; added Product Codes and SKUs chart; updated 1.05 V power rail information throughout; updated voltage and capacitance values in Power Delivery diagram; updated Power Consumption Targets table; added information regarding how to disable WoL using BIOS; updated SIPS driver/OS information; added crystal requirements reference information.
November 2010	1.9	Updated electrical specification section; added solder flow information.
January 2011	2.0	Changed document to unclassified status. Updated Product Features listing. Updated the SKU matrix. Added Operating System support information. Updated maximum clock amplitude specification. Updated SMBus information.
May 2011	2.1	Updated mechanical package illustration for readability; no content changes to drawing. Added network proxy information. Removed information regarding pull-up resistor R15 on LAN Disable; resistor is not needed. Updated the SKU matrix. Updated register information (NVM and MAC). Added Time Sync (1588) chapter. Added Design Considerations chapters for desktop and mobile designs.





1.0 Introduction

1.1 Scope

This document describes the external architecture for the Intel® 82579 Gigabit Ethernet PHY. It's intended to be a reference for software developers of device drivers, board designers, test engineers, or anyone else who might need specific technical or programming information about the 82579.

1.2 Overview

The 82579 is a single port Gigabit Ethernet Physical Layer Transceiver (PHY). It connects to the Intel® 6 Series Express Chipset's integrated Media Access Controller (MAC) through a dedicated interconnect. The 82579 supports operation at 1000/100/10 Mb/s data rates. The PHY circuitry provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3, 802.3u, and 802.3ab). The 82579 also supports the Energy Efficient Ethernet (EEE) 802.3az specification.

The 82579 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 82579 interfaces with its integrated LAN controller 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 82579 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 82579's PCIe-based interface.





Figure 1. Intel® 82579 Block Diagram

1.3 Main Flows

The 82579 main interfaces are PCIe and SMBus on the host side and the MDI interface on the link side. Transmit traffic is received from the integrated LAN controller 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 82579 functionality through two mechanisms:

- The 82579 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 8.3](#). Power management is described in [Section 6.0](#).




1.4 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.
- Intel® 6 Series Express Chipset Chipset Family External Design Specification (Intel® 6 Series Express Chipset Chipset EDS), Intel Corporation
- Intel® 6 Series Express Chipset Chipset External Datasheet Specification, Intel Corporation
- Intel® 6 Series Express Chipset Chipset SPI Flash Programming Guide - Application Note, Intel Corporation
- Intel® 82579 Schematic and Layout Checklists, Intel Corporation
- Intel® 82579 MDI Differential Trace and Power Loss Calculators, Intel Corporation

1.5 Product SKUs, Codes, and Device IDs

The table below lists the product SKUs and ordering codes for the 82579 GbE controller. Refer to the *Intel® 82579 GbE PHY Specification Update* for further device ordering information.

					Performance				Extended Power Savings				Advanced Features								
Method of enabling/disabling features in SKUs =>					Driver	Driver	MAC	MAC	PHY	Driver	Software	PHY	HW	HW	Driver	Driver	Platform	Platform	BIOS	FW	FW
Link Speed	Platform	Segment	Description	Device ID	Jumbo Frames (up to 9k)	802.1Q & 802.1p	Receive Side Scaling (RSS)	2 Tx & 2 Rx Queues	Low 'No-Link' Power (ACBS)	Link Speed Battery Saver	Server Operating System Support	Low Power Linkup (LPLU)	802.1as	EEE - 802.3az	Ability to Initiate a Team	Intel® Virtualization	MACSec (802.3ae)	Intel® SIPP	iSCSI Boot	Standard Manageability	Intel® vPro*
Gigabit	Mobile Desktop Embedded	Corporate	82579LM with Intel® 6 Series Express Chipset	1502	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
		Consumer	82579V with Intel® 6 Series Express Chipset	1503	X	X	X	X	X	X		X	X	X	X	X					X
Gigabit	Entry Server Workstation	Corporate	82579LM with Server Chipset	1502	X	X	X	X	X		X	X	X	X	X		X	X	X	X	X



1 See Intel document *Intel® 82576 and 82579 Gigabit Ethernet Controllers – Intel Software Support for Cisco's MACsec Protocol Supplicant*; contact your Intel representative to obtain this document.

Note: MACSec supported only when Intel® AMT™ is disabled.


2 See Intel document *Intel® 82579 Gigabit Ethernet PHY (TA-192) – Jumbo Frames Technical Advisory*; contact your Intel representative to obtain this document.

1.6 Supported Operating Systems

1.6.1 Operating System Support Matrix

Table 1. Desktop / Workstation / Server / Embedded Operating System Support

Operating System	IA32	X64	Itanium
Windows* XP Professional x64 Edition SP2	n/a	Y	n/a
Windows* XP Professional SP3	Y	n/a	n/a
Windows* Vista* SP2	Y	Y	n/a
Windows* 7	Y	Y	n/a
Windows* Server 2003 SP2‡	Y	Y	N
Windows* Server 2008 SP2‡	Y	Y	N
Windows* Server 2008 SP2 Core ‡	Y	Y	n/a
Windows* Server 2008 R2 ‡	n/a	Y	N
Windows* Server 2008 R2 Core‡	n/a	Y	N
Linux* Stable Kernel version 2.6	Y	Y	n/a
Linux* RHEL 4.8	Y	Y	n/a
Linux* RHEL 5.4	Y	Y	N
Linux* SLES 10 SP3	Y	Y	n/a
Linux* SLES 11	Y	Y	N
FreeBSD* 8	Y	Y	n/a
DOS* NDIS 2	Y	n/a	n/a
DOS* ODI	Y	n/a	n/a
EFI 1.1	Y*	n/a	N
UEFI 2.1	N	Y	N
VMware* ESX 3.5	n/a	N	n/a
VMware* ESX 4.0	n/a	N	n/a




Notes: ‡ Supported on LM SKU only
 * Minimal validation
 + For more information, please see the Operating System Support Note section.

1.6.2 Mobile Operating System Support Matrix

Operating System	IA32	X64	Itanium
Windows* XP Professional x64 Edition SP2	n/a	Y	n/a
Windows* XP Professional SP3	Y	n/a	n/a
Windows* Vista* SP2	Y	Y	n/a
Windows* 7	Y	Y	n/a
Windows* Server 2003 SP2‡	Y	Y	N
Windows* Server 2008 SP2‡	Y	Y	N
Windows* Server 2008 SP2 Core ‡	Y	Y	n/a
Windows* Server 2008 R2 ‡	n/a	Y	N
Windows* Server 2008 R2 Core‡	n/a	Y	N
Linux* Stable Kernel version 2.6	Y	Y	n/a
Linux* RHEL 4.8	Y	Y	n/a
Linux* RHEL 5.4	Y	Y	N
Linux* SLES 10 SP3	Y	Y	n/a
Linux* SLES 11	Y	Y	N
DOS* NDIS 2	Y	n/a	n/a
DOS* ODI	Y	n/a	n/a
EFI 1.1	Y*	n/a	N
UEFI 2.1	N	Y	N

Notes: ‡ Supported on LM SKU only
 * Minimal validation
 + For more information, please see the Operating System Support Note section below.







2.0 Interconnects

2.1 Introduction

The 82579 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. 82579 Interconnect Modes

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

The 82579 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

2.2 PCIe-Based

Note: The 82579 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 82579 PCIe-based interface. Standard PCIe validation tools cannot be used to validate this interface. Refer to [Section 10.4.4](#) for PCIe-based specifications.

2.2.1 PCIe Interface Signals

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

- Serial differential pair running at 1.25 Gb/s for Rx
- Serial differential pair running at 1.25 Gb/s for Tx
- 100 MHz differential clock input to the PHY running at 100 MHz



- Power and clock good indication to the PHY PE_RST_N
- Clock control through CLK_REQ_N

2.2.2 PCIe Operation and Channel Behavior

The 82579 only runs at 1250 Mb/s speed, which is 1/2 of the PCIe Specification v1.1, 2.5Gb/s PCIe frequency. Each of the PCIe root ports in the integrated LAN controller have the ability to operate with the 82579. The port configuration is pre-loaded from NVM. The selected port adjusts the transmitter to run at the 1.25GHz 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.

2.3 SMBus

Note: The 82579 SMBus must only be connected to SMLink0 in the chipset. No other device (like an external BMC) can be connected to SMLink0 when the 82579 is connected to the chipset SMLink0.

2.3.1 Overview

SMBus is a low speed (100 kHz/400 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the 82579 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 82579 as well as passing in-band information between them.

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

The SMBus protocol includes various types of message protocols composed of individual bytes. The message protocols supported by the 82579 are described in the relevant sections.


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

2.3.1.1 SMBus Channel Behavior

The 82579 SMBus interface operates at a frequency of 100 kHz/400 kHz.

2.3.1.2 SMBus Addressing

The 82579's address is assigned using SMBus ARP protocol. The default SMBus address is 0xC8.



2.3.1.3 Bus Time Outs

The 82579 can detect (as a master or a slave) an SMB_CLK time out on the main SMBus. If the SMBus clock line is held low for less than 25 ms, 82579 does not abort the transaction. If the SMBus clock line is held low for 25 ms or longer, the 82579 aborts the transaction.

As a slave, the 82579 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 82579 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 82579 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 82579 has the capability of detecting a hung bus. If SMB_CLK or SMB_DATA are stuck low for more than 35 ms, the 82579 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 SMB_CLK is detected high for ~50 μ s

2.4 Transitions between SMBus and PCIe interfaces

2.4.1 Switching from SMBus to PCIe

Communication between the integrated LAN controller and the 82579 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 to enable host wake up from the 82579.

Possible states for activity over the SMBus:



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

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

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

2.4.2 Switching from PCIe to SMBus

The communication between the integrated LAN controller and the 82579 is done through PCIe each time the system is in active power state (S0); PE_RST_N signal is high. Switching the communication to SMBus is only needed to enable host wake up in low power states.


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

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

2.5 Intel® 6 Series Express Chipset / 82579 – SMBus / PCIe Interconnects

The 82579 can be connected to any x1 PCIe port in Intel® 6 Series Express Chipset. The PCIe port that connects to the 82579 is selected by PCHSTRP9, bits [11:8] in the SPI Flash descriptor region. For more information on this setting, please refer to the Intel® 6 Series Express Chipset External Datasheet Specification. The Intel® 6 Series Express Chipset-to-82579 PCIe port connection in the reference schematic must match the previously mentioned Intel® 6 Series Express Chipset SPI strap setting. Choosing another port can result in unexpected system behavior.


The SMBus/PCIe interface should be connected as shown [Figure 2](#).



Notes:

1. Any free PCIe ports (ports 1-8) can be used to connect to 82579 PCIe Interface.
2. Any CLKOUT_PCIE[7:0] and PCIECLKRQ[7:0] can be used to connect to PE_CLK for the 82579. Also, PCIECLKRQ[7:0] can be connected to CLK_REQ_N for the 82579. The 0 ohm resistor connected to the CLK_REQ_N output should not be populated for non-mobile designs.
3. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per the PCIe specification.
4. If connecting to PCIECLKRQ[1:2]#, the CLK_REQ_N pull-up resistor should be connected to +V3.3S. Refer to the CLK_REQ_N guidance section in the Intel® 6 Series Express Chipset PDG for more details.

Figure 2. Intel® 6 Series Express Chipset / 82579 Interconnects





3.0 Pin Interface

3.1 Pin Assignment

The 82579 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	Connect to PLT_RST# output of the PCH. When asserted, the 82579 interface is switched to SMBus mode.
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	Output	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 2.2 K Ω resistor.
SMB_DATA	31	O/d	BI-dir	SMBus data. Pull this signal up to 3.3 Vdc (auxiliary supply) through a 2.2 K Ω resistor.

3.1.4 Miscellaneous Pins (3)

Pin Name	Pin #	Type	Op Mode	Name and Function
RSVD1_VCC3P3	1	T/s		Pull this signal up to 3.3 Vdc (auxiliary supply) through a 4.7 K Ω resistor
RSVD2_VCC3P3	2	T/s		Pull this signal up to 3.3 Vdc (auxiliary supply) through a 4.7 K Ω resistor
LAN_DISABLE_N	3	I		When this pin is set to 0b, 82579 is disabled. Connect to the LAN_PHY_PWR_CTRL/GPIO12 pin in the Intel® 6 Series Express Chipset.



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] 1000BASE-T : In MDI configuration, MDI[0]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[0]+/- corresponds to BI_DB+/-. 100BASE-TX : In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair. 10BASE-T : In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair.
MDI_PLUS[1] MDI_MINUS[1]	17 18	Analog	Bi-dir	Media Dependent Interface[1] 1000BASE-T : In MDI configuration, MDI[1]+/- corresponds to BI_DB+/- and in MDI-X configuration MDI[1]+/- corresponds to BI_DA+/-. 100BASE-TX : In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair. 10BASE-T : In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair.
MDI_PLUS[2] MDI_MINUS[2] MDI_PLUS[3] MDI_MINUS[3]	20 21 23 24	Analog	Bi-dir	Media Dependent Interface[3:2] 1000BASE-T : In MDI configuration, MDI[3:2]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[3:2]+/- corresponds to BI_DB+/-. 100BASE-TX : Unused. 10BASE-T : Unused.
XTAL_OUT	9	O		Output crystal.
XTAL_IN	10	I		Input crystal.
RBIAS	12	Analog		Connect to ground through a 3.01 K Ω +/-1%.



3.1.6 Testability Pins (5)

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

Note: The 82579 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
VDD1P0	8, 11, 16, 22, 37, 40, 43, 46, 47	Power	1.05 Vdc supply.
VDD3P3	15, 19, 29	Power	3.3 Vdc supply.
VDD3P3_OUT	4	Power	3.3 Vdc out
VDD3P3_IN	5	Power	3.3 Vdc supply.
CTRL1P0	7	Analog	Internal SVR control pin. Connect to a 4.7 μ H inductor and to the 1.05 Vdc rail.
RSVD_NC	6	NC	Not Connected.

§ §

4.0 Package

4.1 Package Type and Mechanical

The 82579 package is a 6 mm x 6 mm, 48-pin QFN Halogen Free and Pb Free package with Expad size of 4.3 mm x 4.3 mm.




Figure 3. Package Dimensions



4.2 Package Electrical and Thermal Characteristics

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

Air Flow (m/s)	Maximum T_J	θ_{JA} ($^{\circ}\text{C}/\text{Watt}$)
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* Epad*. Following are the power supply options:

Option#1:

External Power Supply	3.3 Vdc, 1.05 Vdc
1.05 Vdc	Shared with Intel® 6 Series Express Chipset SVR output.

Option#2

Power Supply	3.3 Vdc, Pin 5
1.05 Vdc	1.05 Vdc power supply provided by internal SVR controlled through pin 7 using external inductor and capacitor (see Section 3.1.7 for more details).





Figure 4. 82579 Power Delivery Diagram

4.4 Ball Mapping

Pin Name	Side	Pin Number	Pin Name	Side	Pin Number
RSVD1_VCC3P3	Left	1	MDI_PLUS[0]	Bottom	13
RSVD2_VCC3P3	Left	2	MDI_MINUS[0]	Bottom	14
LAN_DISABLE_N	Left	3	VDD3P3	Bottom	15
VDD3P3_OUT	Left	4	VDD1P0	Bottom	16
VDD3P3_IN	Left	5	MDI_PLUS[1]	Bottom	17
RSVD_NC	Left	6	MDI_MINUS[1]	Bottom	18
CTRL1P0	Left	7	VDD3P3	Bottom	19
VDD1P0	Left	8	MDI_PLUS[2]	Bottom	20
XTAL_OUT	Left	9	MDI_MINUS[2]	Bottom	21



XTAL_IN	Left	10	VDD1P0	Bottom	22
VDD1P0	Left	11	MDI_PLUS[3]	Bottom	23
RBIAS	Left	12	MDI_MINUS[3]	Bottom	24
LED2	Right	25	VDD1P0	Top	37
LED0	Right	26	PETp	Top	38
LED1	Right	27	PETn	Top	39
SMB_CLK	Right	28	VDD1P0	Top	40
VDD3P3	Right	29	PERp	Top	41
TEST_EN	Right	30	PERn	Top	42
SMB_DATA	Right	31	VDD1P0	Top	43
JTAG_TDI	Right	32	PE_CLKP	Top	44
JTAG_TMS	Right	33	PE_CLKN	Top	45
JTAG_TDO	Right	34	VDD1P0	Top	46
JTAG_TCK	Right	35	VDD1P0	Top	47
PE_RST_N	Right	36	CLK_REQ_N	Top	48
GND	Epad (Center)	49			



82579 Pinouts

4.4.1 Soldering Profile (Pb-Free Reflow Profile)

A recommended reflow profile for Pb-free assembly of MLCC packages is shown in Figure 13. Again, this recommended profile should be used as a starting point and some degree of customization might be required based on the particular application and solder paste type. The solder paste supplier should be consulted for more specific reflow profile recommendations.





Figure 5. Recommended Solder Reflow Profile for Pb-Free Assembly

Important features to consider for proper joint attachment are:

1. Flux activation temperature and time that typically occurs above 150 °C for Pb-free solder pastes with a recommended activation time typically between 60 to 120 seconds. The goal is to consume most of the flux just as the solder alloy reaches its melting point.
2. A peak reflow temperature of ~30 °C above the melting point of the solder paste but not exceeding the maximum rated peak body temperature of the device per MSL requirements (for all MLCC packages the maximum rated peak reflow temperature is 260 °C).
3. A minimum time of 30 seconds above 235 °C is recommended to allow for proper wetting of the Pb-free solder.
4. The time above the melting point should be between 60 to 90 seconds. Shorter times might not allow for sufficient heat transfer on mixed technology boards, while longer times might result in damage to certain components due to over heating. In all cases care should be taken not to exceed 260 °C peak body temperature for the package during reflow optimization.



5.0 Initialization

5.1 Power Up

Initialization begins with power up.





Figure 6. Power-Up Sequence



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

Table 3. Notes

Note	
1	Platform power ramps up (3.3 Vdc/1.05 Vdc)
2	XTAL is stable after T _{XTAL} sec.
3	Internal Power On Reset triggers T _{POR} after XTAL is stable. Strapping options are latched.
4	PCIe training if PE_RST_N is de-asserted.
5	Wait for the integrated LAN Controller SMBus address valid.
6	Send Link Status message.
7	Integrated LAN controller configures 82579.
8	PHY goes through auto-negotiation to acquire link.

Power requirements during the power-up sequence are described in [Section 6.2.1.1](#).

5.2 Reset Operation

The reset sources for the 82579 are as follows:

- **Internal Power On Reset (POR)** – The 82579 has an internal mechanism for sensing the power pins. Until power is up and stable, the 82579 generates an internal active low reset. This reset acts as a master reset for the 82579. While the internal reset is 0b, all registers in the 82579 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 82579 registers are not.

Note: The integrated LAN controller configures the PHY registers. Other 82579 GbE PHY registers do not need to be configured.

- **PCIe Reset** from pin – After asserting a PCIe reset, the 82579 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 82579 except the wake up filters content.

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

Some of the bonding option registers are writable and can be loaded from the NVM or written by the integrated LAN Controller (SKU register). The effect of the various reset options on these and other registers is listed in [Table 4](#).



Table 4 lists the impact of each 82579 reset.

Table 4. 82579 Resets

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

1. Once the PHY completes its internal reset a reset complete indication is sent to the integrated LAN controller over the interconnect. The integrated LAN controller then configures the PHY.
2. Asserting a 3.3 Vdc power on reset should move the PHY out of power down mode.
3. PHY registers (page 0 in MDIO space and any aliases to page 0) are reset during a PHY soft reset. The rest of 82579's MDIO space is not reset.

5.3 Timing Parameters

5.3.1 Timing Requirements

The 82579 requires the following start-up and power-state transitions.

Table 5. Timing Requirements

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


5.3.2 Timing Guarantees

The 82579 guarantees the following start-up and power state transition related timing parameters.

Note: For platform power sequencing requirements for the Intel® 6 Series Express Chipset integrated LAN controller, refer to the Intel® 6 Series Express Chipset EDS.

Table 6. Timing Guarantees

Parameter	Description	Min.	Max.	Notes
T _{PHY_Reset}	Reset de-assertion to PHY reset complete		10 ms	PHY configuration should be delayed until PHY completes it's reset.
T _{c2an}	Cable connect at start of auto-negotiation	1.2 s	1.3 s	Per 802.3 specification.
T _{XTAL}	XTAL frequency stable after platform power ramp up		45 ms	
T _{POR}	Internal POR trigger after XTAL stable		40 ms	





6.0 Power Management and Delivery

This section describes power management in the 82579.

Table 7 lists the targets for device power for the 82579. Note that power is reduced according to link speed and link activity.

Table 7. 82579 Power Consumption Targets

System State		Link State	3.3 V Rail (mA) ¹	1.05 V Rail (mA) ¹ -external-	Device Power (mW)	Solution Power (mW) ²
S0 (Max)		1000 Mb/s active @ 90 °C [Ta]	94	305	630	711
S0 (Typ)		1000 Mbps Active	93	297	619	697
		1000 Mbps Idle	82	227	509	569
		1000 Mbps LPI (EEE link partner only)	10	15	49	53
		100 Mbps Active	52	137	315	351
		100 Mbps Idle	39	75	207	227
		100 Mbps LPI (EEE link partner only)	10	19	53	58
		10 Mbps Active	69	123	357	389
		10 Mbps Idle	18	57	119	134
		No Link (Cable Disconnected)	5	16	33	38
		LAN disabled using driver	4	10	24	26
Sx	WoL Enabled	100 Mbps WoL enabled	39	75	207	227
		100 Mbps LPI (EEE, Rx only)	39	45	176	188
		10 Mbps WoL enabled	18	57	119	134
	WoL Disabled	LAN disabled using driver	4	10	24	26
		LAN disabled using BIOS ³	0	0	0	0


1. Power was supplied by an external power supply.

2. Solution Power was calculated using this formula (assumes 80% efficiency): $3.3 \times I_{3.3} + (1.05 \times I_{1.05}) / 0.8$

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

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

The following sections describe requirements in specific power states.



6.1 Power Delivery

The 82579 operates from a 3.3 Vdc external power rail.

6.1.1 1.05 Vdc Supply

The 1.05 Vdc rail can be supplied in one of two ways:

- An external power supply not dependent on support from the 82579. For example, the PCH 1.05 Vdc SVR can be tied to the 1.05 Vdc PHY supply.
- Integrated SVR solution with external inductor and capacitor.

6.2 Power Management

6.2.1 Global Power States

The 82579 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 82579 and until the 82579 powers up its PHY. The 82579 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 82579 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 82579 is not actively transmitting or receiving data and might enter a lower power state (for example, an interface can be in electrical idle).
- **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_NLAN_PWR_GOODn pin to zero. The 82579 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 (see [section 9.4](#)).
- **LPI** – IEEE802.3az [Energy Efficient Ethernet (EEE)] defines an optional Low Power Idle (LPI) mode for 1000BASE-T, 100BASE-TX and other interfaces. LPI enables power saving by switching off part of the 82579 functionality when no data needs to be transmitted or/and received. When LPI support is enabled The 82579 will shut off RX circuitry and send an inband RX LPI Indication on detection that link the partner's TX moved into LPI state. The 82579 PHY will move TX into LPI state and power-down transmit circuitry when receiving an Inband TX LPI request from the integrated LAN controller. In SX states, LPI is supported only in 100 Mbps WoL-enabled mode while keeping the receive side active.

6.2.1.1 Power Up

Defined as the period from the time power is applied to the 82579 and until the 82579 powers up its PHY. The 82579 should consume less than ~40 mA during this period. Following the 82579 PHY entering reset, the power-up sequence is considered done and the requirement is removed. See [Section 5.1](#) for a description of the power-up sequence.

6.2.1.2 Cable Disconnect State

The 82579 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 82579 detects energy on the MDI link, and starts the following exit sequence:

- The 82579 signals the integrated LAN controller that link energy was detected by clearing the *Cable Disconnect* bit in the PCIe or SMBus interface.
- The PHY waits until the auto-negotiation break link timer expires (T_{c2an} time) and then starts to advertise data on the line.

6.2.1.3 Cable Disconnect State Active Power Down Mode

The 82579 sends a link down notification to the integrated LAN controller following detection of a cable disconnect condition on the Ethernet link. In the above case, the software device driver initiates a timer to check if the link is re-enabled. If there is no link after the timer expired, the software device driver configures the 82579 to Wake on LAN (WoL) on a link status change for the host and then moves the 82579 into an active power down mode.

When link is re-established, the 82579 sends a wake up message to the integrated LAN controller that the link was re-established. Afterwards, the Intel® 6 Series Express Chipset powers up the integrated LAN controller and restores the configuration space.

6.2.1.4 Power Down State

The 82579 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 82579 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 7 shows the power-down sequence in the two later cases.




Figure 7. Power-Down Sequence

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

Table 8. Figure 7 Notes

Note	Description
1, 2	Once the 82579 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 82579 through the MDIO interface.
6	PHY goes through auto-negotiation to acquire link.

6.2.1.5 EEE LPI State


EEE (Energy Efficient Ethernet) Low Power Idle (LPI) mode defined in IEEE 802.3az optionally allows power saving by switching off part of the integrated LAN controller and the 82579 functionality when no data needs to be transmitted or/and received. Decision on whether the 82579 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 82579 in order to allow power saving in the transmit circuitry. Information on whether Link Partner has entered Low Power Idle mode is detected by the 82579 and communicated to the integrated LAN controller to allow for power saving in the receive circuitry.

Figure 8 illustrates general principles of an EEE LPI operation on the Ethernet Link.

Table 9. LPI Parameters

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

Figure 8. EEE LPI Compliant Operation



In the transmit direction, entrance to Low Power Idle mode of operation is triggered by the reception of LPI TX Request from the integrated LAN controller. Following reception of the LPI TX in band Request, PHY transmits special Sleep symbols to communicate to the link partner that the local system is entering Low Power Idle mode. In 100BASE-TX LPI mode PHY enters low power operation in an asymmetric manner. After Sleep symbols transmission, the transmit function of the local PHY immediately enters a low power quiet mode. In 1000BASE-T LPI mode, PHY entry into low power mode is symmetric. Only after the local PHY transmits and receives sleep symbols from the remote PHY does the transmit function of the local PHY enter the quiet mode. Periodically the local PHY transmits Refresh symbols that are used by the link partner to update adaptive filters and timing circuits in order to maintain link integrity. This quiet-refresh cycle continues until the local integrated LAN controller sends an in band message with a clear (0) LPI TX Request, which signals to the 82579 that Low Power Idle mode should end. The 82579 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.

6.2.1.5.1 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 IEEE 802.3az clauses 28C.12 and 28C.13.

6.2.1.5.2 EEE LPI Unsupported Features

EEE LPI does not support:

- Half-duplex operation
- 10 Mb/s speed
- MDI-X mode

Note: In SX states, LPI is supported only in 100 Mbps WoL- enabled mode while keeping the receive side active.



6.2.2 Power Saving Features

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

6.2.2.1 Intel® Auto Connect Battery Saver (ACBS)

Intel® Auto Connect Battery Saver for the 82579 is a hardware-only feature that automatically reduces the PHY 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 82579 does not enter ACBS mode.

The 82579 ACBS works in both S0 and Sx states. Since the 82579 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 82579 is in ACBS mode.

6.2.2.2 Automatic Link Downshift

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

Note: Automatic Link Downshift is NOT enabled when in LPI mode.

6.2.2.3 Link Speed Battery Saver

Link speed battery saver is a power saving feature that negotiates to the lowest speed possible when the 82579 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 PHY 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 Windows driver.

Note: Link Speed Battery Saver is NOT enabled when in LPI mode.

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.

6.2.2.4 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 of the 82579. Refer to [Section 6.1](#) for the power dissipated in each link state.

Note: SIPS is NOT enabled when in LPI mode.

SIPS is designed to save power in the 82579 by negotiating to the lowest possible link speed when both the network is idle and the monitor is turned off due to inactivity. The SIPS feature is activated based on both of the following conditions:

- The Windows Vista/Windows 7 NDIS driver receives notification from the operating system when the monitor is turned off due to non-activity.
- The LAN driver monitors the current network activity and determines that the network is idle.

Then, with both the monitor off and the network idle, the LAN negotiates to the lowest possible link speed supported by both the PHY and the link partner (typically 10 Mb/s). If the link partner is hard-set to only advertise a certain speed, then the LAN negotiates to the advertised speed. This link speed is maintained until the LAN driver receives notification from the operating system that the monitor is turned on, thus exiting SIPS and re-negotiating to the highest possible link speed supported by both the PHY and the link partner. If SIPS is exited when transferring packets, any packet that was being transferred is re-transmitted by the protocol layer after re-negotiation to the higher link speed.


This feature is disabled by default through the driver INF. Please contact your Intel representative if support is needed to enable the feature.

6.2.2.5 Low Power Link Up (LPLU)

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

When speed negotiation starts, the PHY tries to negotiate for a 10 Mb/s link, independent of speed advertisement. If link establishment fails, the PHY tries to negotiate with different speeds. It enables all speeds up to the lowest speed supported by the partner. For example, if the 82579 advertises 10 Mb/s only and the link partner supports 1000/100 Mb/s only, a 100 Mb/s link is established.

LPLU is controlled through the LPLU bit in the PHY Power Management register. The integrated LAN controller sets and clears the bit according to hardware/software settings. The 82579 auto negotiates with the updated LPLU setting on the following auto-negotiation operation. The 82579 does not automatically auto-negotiate after a change in the LPLU value. LPLU is not dependent on whether the system is in Vac or Vdc mode. In S0 state, link speed battery saver overrides the LPLU functionality.



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

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

LPLU power consumption depends on what speed it negotiates at. [Section 6.1](#) includes all of the power numbers for the 82579 in the various speeds.

6.2.3 LAN Disable Requirements

LAN_DISABLE_N needs to be connected to the GPIO12/LAN_PHY_PWR_CTRL output of the integrated LAN controller. GPIO12 also needs to be configured using the integrated LAN controller soft straps as LAN_PHY_PWR_CTRL (bit [20] of PCHSTRP0 register - LAN_PHY_PWR_CTRL/GPIO12. Refer to the Intel® 6 Series Express Chipset External Design Specification (EDS).

§ §

7.0 Device Functionality

7.1 Tx Flow

When packets are ready for transmission in the integrated LAN controller it transfers them to the 82579 through the PCIe or the SMBus (depends on system state). The 82579 starts transmitting the arrived packet over the wire after it gathers eight 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 82579 design is based on the assumption that the integrated LAN controller has the full packet ready for transmission.

In several cases the 82579 has to stop transmission over the wire while still accepting data from the upper connection (PCIe or SMBus). For those cases, the 82579 maintains a 2 KB FIFO. The cases where the 82579 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 PHY Receive Control register should be set in addition to the configuration in the integrated LAN controller.
- In half-duplex mode while the 82579 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 82579 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 82579 might avoid sending the in-band message. An in-band message with the Tx *OFF* bit cleared is sent if when the collided packet was successfully transmitted or dropped after 16 retries (see [Section 7.3](#)).

In-band messages from the integrated LAN controller to the 82579 always come in between packets during the IPG. The 82579 does not accept in-band messages in the middle of a packet.

7.2 Rx Flow

The 82579 maintains a 2 KB FIFO on the receive side in order not to lose packets when PCIe is active but at K1 power save mode (similar to the L1 PCIe ASPM state). In this case the 82579 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 82579 starts the transmission after the entire packet is received. The 82579 assumes maximum recovery time (from the K1 state) of 10 μ s on both sides of the PCIe side. Higher recovery time causes a packet drop on the receive side.



The 82579 identifies PAUSE packets, stop transmission, and a send in-band message as described in the previous section.

In-band messages from the 82579 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. Some of the flow control functionality has moved to the 82579. The following registers are duplicated to the 82579 for the implementation of flow control:

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

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

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

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




Figure 9. 802.3x Integrated LAN Controller Control Frame Format

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

Note:

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

Once the 82579 has validated the reception of an XOFF, or PAUSE frame, it performs the following:

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

Resuming transmission might occur under the following conditions:



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

Once the 82579 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

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 82579 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).

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 82579 is operating in half-duplex mode.

7.4 Wake Up

The 82579 supports host wake up.

Host wake up uses in-band messages to wake the integrated LAN controller from a sleep state. The host can enable host wake up from the 82579 by setting the *Host_WU_Active* bit. When this bit is set, after the host transitions to a low power state, the SMBus interface is still active and the wake up indication from the 82579 to the integrated LAN controller would come in as an in-band message over the SMBus.

Setting the 82579'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 bit 2 (*MACPD_enable*) of the Port Control register (page 769, register 17) to enable the 82579 wake up capability and software accesses to page 800.
3. Set the *Slave Access Enable* bit (bit 2) in the Receive Control register (page 800, register 0) to enable access to the Flex Filter register, if setting those bits is needed in the next stage. The registers affected are:
 - a. Flexible Filter Value Table LSB– FFVT_L (filters 01)

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

- b. Flexible Filter Value Table MSBs – FFVT_H (filters 23)
 - c. Flexible Filter Value Table – FFVT_45 (filters 45)
 - d. Flexible Filter Value Table – FFVT_67 (filters 67)
 - e. Flexible TCO Filter Value/Mask Table LSBs – FTFT_L
 - f. Flexible TCO Filter Value/Mask Table MSBs – FTFT_H
4. Configure the 82579'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 82579's wake up functionality.

Note: Once wake up is enabled, the 82579 stops responding to SMBus commands.

Host wake up:

1. When a WoL packet/event is detected, the 82579 sends an in-band message to the integrated LAN controller indicating Host wake up.
2. In case of host wake up, the integrated LAN controller wakes the host.
3. Host should issue a LCD reset to the 82579 before clearing the *Host_WU_Active* bit.
4. Host reads the Wake Up Status register (WUS); wake up status from the 82579).

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).

While in wake up active mode new wake up packets received will not overwrite the packet in the FIFO. The 82579 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 82579 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 82579's wake up functionality is enabled, the *APM Enable* bits from the NVM are written to the 82579 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 82579 checks all incoming packets for Magic Packets. See [Section 7.4.1.3.1.4](#) for a definition of Magic Packets.



To enable APM wake up, programmers should write a 1b to bit 10 in register 26 on page 0 PHY address 01, and then the station address to registers 27, 28, 29 at page 0 PHY address 01. The order is mandatory since registers 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 82579 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 82579 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 82579 supports ACPI Power Management based wake ups and can generate system wake up events from three sources:

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

Activating ACPI Power Management wake up requires the following steps:

- Programming of the WUFC register to indicate the packets it needs to wake up and supplies the necessary data to the IPv4 Address Table (IP4AT) and the Flexible Filter Mask Table (FFMT), Flexible Filter Length Table (FFLT), and the Flexible Filter Value Table (FFVT). It can also set the *Link Status Change Wake up Enable* (LNKC) bit (bit 0) in the WUFC register to cause wake up when the link changes state.
- Setting bit 2 (*MACPD_enable*) of the Port Control register (page 769, register 17) to put the 82579 in wake up mode.



Once wake up is enabled, the 82579 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 82579:

- 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 Packets

The 82579 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 also depends on the *NoTCO* bit (11) in the WUFC register being inactive.

7.4.1.3.1 Pre-Defined Filters

The following packets are supported by the 82579'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.

7.4.1.3.1.1 Directed Exact Packet

The 82579 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



7.4.1.3.1.2 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 82579 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.

7.4.1.3.1.3 Broadcast

If the *Broadcast Wake Up Enable* bit (bit 4) in the WUFC register is set, the 82579 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	

7.4.1.3.1.4 Magic Packet

Magic packets are defined as follows:

- **Magic Packet Technology Details** - Once the 82579 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 integrated LAN controller 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 82579 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 82579 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 82579 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 82579 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 82579 starts up with the RAH0/RAL0 register 0 loaded from the NVM. This enables the 82579 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

7.4.1.3.1.5 ARP / IPv4 Request Packet

The 82579 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 four programmed IPv4 addresses. The 82579 also handles ARP Request 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	0x0806	Compare	ARP
14	2	Hardware Type	0x0001	Compare	
16	2	Protocol Type	0x0800	Compare	
18	1	Hardware Size	0x06	Compare	
19	1	Protocol Address Length	0x04	Compare	
20	2	Operation	0x0001	Compare	
22	6	Sender Hardware Address	-	Ignore	
28	4	Sender IP Address	-	Ignore	
32	6	Target Hardware Address	-	Ignore	
38	4	Target IP Address	IP4AT	Compare	Might match any of four values in IP4AT

7.4.1.3.2 Directed IPv4 Packet

The 82579 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 four programmed IPv4 addresses. The 82579 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	1	Version/ HDR length	0x4X	Compare	Check IPv4
15	1	Type of Service	-	Ignore	
16	2	Packet Length	-	Ignore	
18	2	Identification	-	Ignore	
20	2	Fragment Info	-	Ignore	
22	1	Time to live	-	Ignore	
23	1	Protocol	-	Ignore	
24	2	Header Checksum	-	Ignore	
26	4	Source IP Address	-	Ignore	
30	4	Destination IP Address	IP4AT	Compare	Might match any of four values in IP4AT

7.4.1.3.2.1 Directed IPv6 Packet

The 82579 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 82579 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	1	Version/ Priority	0x6X	Compare	Check IPv6
15	3	Flow Label	-	Ignore	
18	2	Payload Length	-	Ignore	
20	1	Next Header	-	Ignore	
21	1	Hop Limit	-	Ignore	
22	16	Source IP Address	-	Ignore	
38	16	Destination IP Address	IP6AT	Compare	Match value in IP6AT



7.4.1.3.3 Flexible Filter

The 82579 supports a total of eight 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.

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


7.4.1.3.3.1 IPX Diagnostic Responder Request Packet

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

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	
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	0x8137	Compare	IPX
14	16	Some IPX Information	-	Ignore	
30	2	IPX Diagnostic Socket	0x0456	Compare	

7.4.1.3.3.2 Directed IPX Packet

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



Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header – processed by main address filter
6	6	Source Address		Skip	
12	S = (0/4)	Possible VLAN Tag		Skip	
12 + S	D = (0/8)	Possible Len/LLC/SNAP Header		Skip	
12 + S + D	2	Type	0x8137	Compare	IPX
14	10	Some IPX Information	-	Ignore	
24	6	IPX Node Address	Receive Address 0	Compare	Must match Receive Address 0

7.4.1.3.3.3 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 addresses.

7.4.2 Accessing the 82579's Wake Up Register Using MDIC

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


7.5 Network Proxy Functionality

7.5.1 Introduction

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

Prior to network proxy, devices the wanted to maintain their network presence would have configured the ARP and neighbor discovery messages as wake up patterns to the system. Analysis show that many of these ARP wakeups are unnecessary as they are generated by automated processes whose sole purpose is to verify that the system is alive on the network.

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 Request For Comment (RFC) 826 as defined by Internet Engineering Task Force (IETF).



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

7.5.2 Network Proxy activation

As part of the system sleep flow and after receiving from the OS the network proxy and WoL patterns, the SW driver should go through the following steps to activate network proxy in the 82579:

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 IPAV[15:14]

Note: The 82579 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 82579 supports responding to ARP Request packets (proxy) if enabled through the driver. For more details refer to 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 82579 also handles ARP Request packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types



7.5.3.1 ARP Request Packet

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

7.5.3.2 ARP Response Packet

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



7.5.4 IPv6 Proxy - Neighbor Discovery

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

Neighbor Discovery is a set of five 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 10. Neighbor Discovery Flowchart



7.5.4.1 Ipv6 Neighbor Solicitation Packet

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

7.5.4.2 Ipv6 Neighbor Advertisement Packet

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



Offset	# of bytes	Field	Value (hex)	Action
21+D+S	1	Hop Limit	0xFF	Constant
22+D+S	16	Source Address		relevant IPv6AT entry (ND target address)
38+D+S	16	Destination Address		Copy from ND packet Source address If source address was undefined - send to All Nodes Multicast (FF02::1)
ICMPv6 header				
54+D+S	1	Type	0x88	Constant
55+D+S	1	Code	0x0	
56+D+S	2	Checksum		Calculate
58+D+S	4	Flags	0x60000000	Constant (Solicited, Override) if the source address was defined
			0x20000000	Constant (Override) if the source address was undefined
62+D+S	16	Target IP Address	IPv6AT	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

7.6 Loopback

PHY loopback is supported in the 82579. Software or Firmware should set the LAN Connected Device to the loopback mode via MDIC register writing to the PHY Control Register (Page 0 Register 00). The PHY supports a number of loopback modes configured through the Loopback Control Register (Page 0 Register 19).

For more information on the different loopback modes See [section 9.5.1](#).

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

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



8.0 Electrical and Timing Specifications

8.1 Introduction

This section describes the 82579's recommended operating conditions, power delivery, DC electrical characteristics, power sequencing and reset requirements, PCIe specifications, reference clock, and packaging information.

8.2 Operating Conditions

8.2.1 Absolute Maximum Ratings

Symbol	Parameter	Min	Max	Units
T_{case}	Case Temperature Under Bias	0	106	°C
$T_{storage}$	Storage Temperature Range	-40	125	°C
V_i/V_o	3.3 Vdc I/O Voltage	0.3	5.0	Vdc
VCC	3.3 Vdc Periphery DC Supply Voltage	0.3	5.0	Vdc
VCC1p0	1.05 Vdc Supply Voltage	0.3	1.8	Vdc

Notes:

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

8.2.2 Recommended Operating Conditions

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

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



8.3 Power Delivery

8.3.1 Voltage Regulator Power Supply Specifications

8.3.1.1 3.3 Vdc Rail

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

8.3.1.2 1.05 Vdc Rail (External/Shared)

Title	Description	Min	Max	Units
Rise Time	Time from 10% to 90% mark	0.05	40	mS
Monotonicity	Voltage dip allowed in ramp	N/A	0	mV
Slope	Ramp rate at any given time between 10% and 90% Min: $0.8 \cdot V(\text{min}) / \text{Rise time (max)}$ Max: $0.8 \cdot V(\text{max}) / \text{Rise time (min)}$	7.6	8400	V/S
Operational Range	Voltage range for normal operating conditions	0.990	1.103	Vdc
Ripple	Maximum voltage ripple (peak to peak)	N/A	50	mV
Overshoot	Maximum overshoot allowed	N/A	100	mV
Decoupling Capacitance	Capacitance range	20	30	μF
Capacitance ESR	Equivalent series resistance of output capacitance	5	50	m Ω

8.3.2 SVR Specification (Internal)

Parameter	Specifications			Units	Comments
	Min	Typ	Max		
Regulator Output Voltage	0.8		1.2	V dc	The 82579 default voltage is set to 1.05 Vdc
Output Voltage Accuracy	-3		+3	%	Not including line and load regulation errors.
Input Voltage Range	2.9	3.3	3.7	Vdc	Supply voltage range.
Load Current	0.01	0.3	0.5	A	Average value.
Output Voltage Under/Over Shoot	-10		+10	%	For min-to-max average load current change.



Transient Settling Time		100		μs	Duration of overshoot or undershoot.
Conversion Efficiency	80	85	90	%	
Switching Frequency		1.5625		MHz	
Output Filter Inductor	3.9	4.7		μH	
Output Filter Inductor DCR		0.100	0.318	Ω	+/-20%, values higher than the typical DCR value will lower the SVR conversion efficiency.
Output Filter Capacitor	20			μF	
Output Filter Capacitor ESR		5	50	mΩ	
Input Capacitor	22			μF	

8.3.3 Power On/Off Sequence

- There is no power sequencing requirement for the 82579.

Table 10. Power Detection Threshold

Symbol	Parameter	Specifications			Units
		Min	Typ	Max	
V1a	High-threshold for 3.3 Vdc supply	2.35	2.45	2.6	Vdc
V2a	Low-threshold for 3.3 Vdc supply	2.1	2.45	2.6	Vdc
V1b	High-threshold for 1.05 Vdc supply	0.6	0.75	0.85	Vdc
V2b	Low-threshold for 1.05 Vdc supply	0.45	0.65	0.75	Vdc



8.4 I/O DC Parameter

8.4.1 3.3 Vdc I/O (Open Drain)

Parameter	Minimum	Typical	Maximum	Unit
VIL	-0.4	0	0.8	Vdc
VIH	2	3.3	3.6	Vdc
VOL	-0.4	0	0.4	Vdc
VOH	2.4	3.3	3.6	Vdc
I _{pullup}	30	50	75	μA
I _{leakage}			10	μA
Ci		2	4	pF


Pin Name	Bus Size	Description
CLK_REQ_N	1	Open-drain I/O.
SMB_CLK	1	Open-drain I(H)/O with snap back NMOS ESD cell.
SMB_DATA	1	Open-drain I(H)/O with snap back NMOS ESD cell.

Note: SMBus leakage current when the 82579 is OFF is <180 uA.

8.4.2 3.3 Vdc I/O

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.3	0	0.4	Vdc
VIH		2	3.3	3.6	Vdc
VOL	I _{OL} = 9 mA VCC = Min	-0.4	0	0.4	Vdc
VOH	I _{OH} = -9 mA VCC = Min	2	2.6	2.8	Vdc
I _{pullup}		30	50	75	μA
I _{leakage}		15 (pull down)	25 (pull down)	35 (pull down)	μA
Ci			2	4	pF
PU			50		KΩ
PD			50		KΩ

Pin Name	Bus Size	Description
RSVD1_VCC3P3, RSVD2_VCC3P3	2	I/O, PU
LED0, LED1, LED2	3	I/O, PU
JTAG_TDI	1	I/O, PU
JTAG_TMS	1	I/O, PU
JTAG_TDO	1	I/O, PU
JTAG_TCK	1	I/O, PU



8.4.3 Input Buffer Only

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.3	0	0.8	Vdc
VIH		2	3.3	3.6	Vdc
Ipullup		30	50	75	μA
Ileakage				10	μA
Ci			2	4	pF

Pin Name	Bus Size	Description
LAN_DISABLE_N	1	I(H), PU
TEST_EN	1	I (no PU, no PD)
PE_RST_N	1	I(H), PU

8.4.4 PCIe DC/AC Specifications

8.4.4.1 PCIe Specifications (Transmitter)

Note: Refer to the Intel® 6 Series Express Chipset/82579 PCIe-Based Test Procedure for more details.

Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
UI	Unit interval	799.92	800.08	ps	Each UI is 800 pS +/- 100 ppm
V _{tx-diff-pp}	Differential peak-to-peak Tx voltage swing	0.8	1.2	V dc	
T _{tx-eye}	Transmitter eye including all jitter sources	0.75		UI	
T _{tx-eye-median-to-max-jitter}	Maximum time between the jitter median and maximum deviation from the median		0.125	UI	
RL _{tx-diff}	Tx package plus silicon differential return loss	7		db	
RL _{tx-cm}	Tx package plus silicon common mode return loss	6		db	
Z _{tx-diff-dc}	DC differential Tx impedance	75	120	W	
V _{tx-cm-ac-p}	Tx V ac common mode voltage (2.5 GT/s)		20	mV	
I _{tx-short}	Transmitter short-circuit current limit		90	mA	
V _{tx-dc-cm}	Transmitter DC common mode voltage	0	3.6	V dc	
V _{tx-cm-dc-active-idle-delta}	Absolute delta of DC common mode voltage during L0 and electrical idle	0	100	mV	



$V_{tx-cm-dc-line-delta}$	Absolute delta of DC common mode voltage between D+ and D-	0	25	mV	
$V_{tx-idle-diff-ac-p}$	Electrical idle differential peak output voltage	0	20	mV	
$T_{tx-idle-set-to-idle}$	Maximum time to transition to a valid electrical idle after sending an EIOS		35	ns	
$T_{tx-idle-to-diff-data}$	Maximum time to transition to valid differential signaling after leaving electrical idle		35	ns	

Note: Figure 11 is for informational purposes only. Do not use for actual eye comparisons.




Figure 11. Transmitter Eye Diagram

8.4.4.2 PCIe Specifications (Receiver)

Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
UI	Unit interval	799.92	800.08	ps	Each UI is 800 ps +/- 100 ppm
$V_{rx-diff-pp-cc}$	Differential peak-to-peak Rx voltage swing for common clock	0.175	1.2	V dc	
$V_{rx-diff-pp-dc}$	Differential peak-to-peak Rx voltage swing for data clock	0.175	1.2	V dc	
T_{rx-eye}	Receiver minimum eye time opening	0.4	N/A	UI	
$T_{rx-eye-median2maxjitter}$	Maximum time delta between median and deviation from median	N/A	0.3	UI	
$RL_{rx-diff}$	Rx differential return loss	6	N/A	dB	
RL_{rx-cm}	Rx CM return loss	5	N/A	dB	

$Z_{rx-diff-dc}$	Rx differential Vdc impedance	80	120	W	
$V_{rx-cm-ac-p}$	Rx Vac CM voltage	N/A	150	mVp	
$Z_{rx-high-imp-dc-pos}$	DC input CM impedance for $V > 0$	50 K	N/A	W	
$Z_{rx-high-imp-dc-neg}$	DC input CM impedance for $V < 0$	1 K	N/A	W	
$V_{rx-idle-det-diff-p-p}$	Electrical idle detect threshold	65	175	mV	

Note: The 82579 has integrated PCIe termination that results in attenuating the voltage swing of the PCIe clock supplied by the Intel® 6 Series Express Chipset. This is in compliance with the PCIe CEM 1.1 specification. More detail is available in the *Intel® 6 Series Express Chipset PDG*.

Note: [Figure 12](#) is intended to show the difference between the PCIe 1.0 and PCIe-based receiver sensitivity templates. It is for informational purposes only.




Figure 12. Receiver Eye Diagram

8.4.4.3 PCIe Clock Specifications

The PCIe clock specification can be found in the PCI Express Card Electromechanical Specification 1.1.



8.5 Discrete / Integrated Magnetics Specifications

Criteria	Condition	Values (Min / Max)
Voltage Isolation	At 50 to 60 Hz for 60 seconds	1500 Vrms (min)
	For 60 seconds	2250 V dc (min)
Open Circuit Inductance (OCL) or OCL (alternate)	With 8 mA DC bias at 25 °C	400 μ H (min)
	With 8 mA DC bias at 0 °C to 70 °C	350 μ H (min)
Insertion Loss	100 kHz through 999 kHz	1 dB (max)
	1.0 MHz through 60 MHz	0.6 dB (max)
	60.1 MHz through 80 MHz	0.8 dB (max)
	80.1 MHz through 100 MHz	1.0 dB (max)
	100.1 MHz through 125 MHz	2.4 dB (max)
Return Loss	1.0 MHz through 40 MHz 40.1 MHz through 100 MHz	18 dB (min) 12 to 20 * LOG (frequency in MHz / 80) dB (min)
	When reference impedance is 85 Ω , 100 Ω , and 115 Ω	
	Note that return loss values might vary with MDI trace lengths. The LAN magnetics might need to be measured in the platform where it is used.	
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	-50.3+(8.8*(freq in MHz / 30)) dB (max)
	30 MHz through 250 MHz	-26-(16.8*(LOG(freq in MHz / 250)))) dB (max)
	250.1 MHz through 375 MHz	-26 dB (max)
Crosstalk Isolation Integrated Modules	1.0 MHz through 10 MHz	-50.8+(8.8*(freq in MHz / 10)) dB (max)
	10.1 MHz through 100 MHz	-26-(16.8*(LOG(freq in MHz / 100)))) dB (max)
	100.1 MHz through 375 MHz	-26 dB (max)
Diff to CMR	1.0 MHz through 29.9 MHz	-40.2+(5.3*((freq in MHz / 30)) dB (max)
	30 MHz through 500 MHz	-22-(14*(LOG((freq in MHz / 250)))) dB (max)
CM to CMR	1.0 MHz through 270 MHz	-57+(38*((freq in MHz / 270)) dB (max)
	270.1 MHz through 300 MHz	-17-2*((300-(freq in MHz) / 30) dB (max)
	300.1 MHz through 500 MHz	-17 dB (max)

8.6 Mechanical

Body Size (mm)	Ball Count	Ball Pitch	Ball Matrix	Center Matrix	Substrate
6x6 mm	48	0.4 mm	N/A, Peripheral	N/A, Exposed Pad	N/A Lead frame-Based Package

8.7 Oscillator / Crystal Specifications

Figure 14 shows the external crystal design and Table 11 lists its parameters.




Table 11. External Crystal Specifications

Parameter Name	Symbol	Recommended Value	Max / Min Range	Conditions
Frequency	f_o	25 [MHz]		@25 [°C]
Vibration Mode		Fundamental		
Frequency Tolerance @25 °C	Df/f_o @25°C	±30 [ppm]		@25 [°C]
Temperature Tolerance	Df/f_o	±30 [ppm]		0 to +70 [°C]
Series Resistance (ESR)	R_s		50 [Ω] max	@25 [MHz]
Crystal Load Capacitance	C_{load}	18 [pF]		
Shunt Capacitance	C_o		6 [pF] max	
Drive Level	D_L		200 [μW] max	
Aging	Df/f_o	±5 ppm per year	±5 ppm per year max	
Calibration Mode		Parallel		
Insulation Resistance			500 [MΩ] min	@ 100 Vdc

Note: Crystal must meet or exceed the specified drive Level (D_L). Refer to the crystal design guidelines in the Intel® 5 Series Family PDG.

Table 12. Clock Oscillator Specifications

Parameter Name	Symbol / Parameter	Conditions	Min	Typ	Max	Unit
Frequency	f_o	@25 [°C]		25.0		MHz
Clock Amplitude	Vmax		0.8		2.2	Vdc
	Vmin				0	Vdc
Frequency Tolerance	f/f_o	20 to +70		±50		[ppm]
Operating Temperature	T_{opr}	-20 to +70				°C
Aging	f/f_o			±5 ppm per year		[ppm]
TH_XTAL_IN	XTAL_IN High Time		13	20		nS
TL_XTAL_IN	XTAL_IN Low Time		13	20		nS
TR_XTAL_IN	XTAL_IN Rise	10% to 90%			5	nS
TF_XTAL_IN	XTAL_IN Fall	10% to 90%			5	nS
TJ_XTAL_IN	XTAL_IN Total Jitter				200 ¹	pS

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






Figure 13. XTAL Timing Diagram



Note: The value of C1 above can be adjusted so that the waveform is within specified limits (see table above). To determine the final value of C1, measure the waveform at XTAL1 and compare to the waveform in Figure 12.

Figure 14. Clock Oscillator Schematic

§ §



9.0 Programmer's Visible State

9.1 Terminology

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

This document names registers as follows.

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



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

Register bits with default values marked with an asterisk * are loaded by the integrated LAN controller during the 82579 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 82579 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 82579 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 82579 contains registers controlling the custom interface as well as other the 82579 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 82579 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 $\text{page_num} \times 32$ to Register 31. This is because only the 11 MSB's of register 31 are used for defining the page. During write to the page register, the five LSB's are ignored.

In pages 800 and 801, the register address space is more than 32. See [section 9.10](#) 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. See [section 9.5](#) for details about these registers.



9.4 Address Map

Table 13. Address Map

PHY Address	Page	Register	Name	Page
02	Any	0	Control	67
02	Any	1	Status	68
02	Any	2	PHY Identifier 1	69
02	Any	3	PHY Identifier 2	69
02	Any	4	Auto-Negotiation Advertisement	70
02	Any	5	Auto-Negotiation Link Partner Ability	70
02	Any	6	Auto-Negotiation Expansion	71
02	Any	7	Auto-Negotiation Next Page Transmit	71
02	Any	8	Link Partner Next Page	72
02	Any	9	1000BASE-T Control	72
02	Any	10	1000BASE-T Status	73
02	Any	14:11	Reserved	
02	Any	15	Extended Status	73
02	0	17:16	Reserved	
02	0	18	PHY Control 2	74
02	0	19	Loopback Control	75
02	0	20	Rx Error Counter	76
02	0	21	Management Interface (MI) Control	76
02	0	22	PHY Configuration	77
02	0	23	PHY Control	78
02	0	24	Interrupt Mask	79
02	0	25	Interrupt Status	80
02	0	26	PHY Status	81
02	0	27	LED Control 1	82
02	0	28	LED Control 2	83
02	0	29	LED Control 3	84
02	0	30	Diagnostics Control (Linking Disabled)	84
02	0	31	Diagnostics Status	85
Page 769 – Port Control Registers				
01	769	16	Custom Mode Control	88
01	769	17	Port General Configuration	88
01	769	20	DFT Control	81
01	769	21	Power Management Control	88
01	769	25	Rate Adaptation Control	89
01	769	27	Flow Control Transmit Timer Value	89
Page 778 – Statistics Registers				
01	778	16 - 17	Single Collision Count	89
01	778	18 - 19	Excessive Collisions Count	90

**Table 13. Address Map**

PHY Address	Page	Register	Name	Page
01	778	20 - 21	Multiple Collisions Count	90
01	778	23 - 24	Late Collision Count	90
01	778	25 - 26	Collision Count	90
01	778	27 - 28	Defer Count	90
01	778	29 - 30	Transmit with No CRS - TNCRS	91
PCIe Registers				
01	770	16	PCIe FIFOs Control/Status	91
01	770	17	PCIe Power Management Control	91
01	770	18	In-Band Control	93
01	770	20	PCIe Diagnostics	93
01	770	21	Time Outs	93
01	770	23	PCIe K-State Minimum Duration Timeout	94
LPI Registers				
01	772	18	LowPower Idle GPIO Control	94
01	772	20	Low Power Idle Control	94
General Registers				
01	776	18	82579 Energy detect mode	88
01	776	19	82579 Capability Register	99
01	0	25	OEM Bits	95
01	0	26	SMBus Address	95
01	0	27-28	Shadow Register for RAL0[31:0].	95
01	0	29	Shadow Register for RAH0[15:0].	96
01	0	30	LED Configuration	96
Page 800 - Wake Up Registers				
01	800	0	Receive Control Register	98
01	800	1	Wake Up Control Register	99
01	800	2	Wake Up Filter Control Register	100
01	800	3	Wake Up Status Register	100
01	800	16	Receive Address Low 0	101
01	800	18	Receive Address High 0	101
01	800	44 - 45	Shared Receive Address Low 0	102
01	800	46 - 47	Shared Receive Address High 0	102
01	800	58 - 59	Shared Receive Address High 3	102
01	800	64	IP Address Valid – IPAV	102
01	800	82 - 83	IPv4 Address Table – IP4AT 0	103
01	800	88 - 89	IPv6 Address Table – IP6AT 0	103
01	800	128 - 191	Multicast Table Array – MTA[31:0]	103
01	800	256 + 2*n (n = 0 - 127)	Flexible Filter Value Table LSB– FFVT_01	104
01	800	257 + 2*n (n = 0 - 127)	Flexible Filter Value Table MSB – FFVT_23	104

**Table 13. Address Map**

PHY Address	Page	Register	Name	Page
01	800	$512 + 2*n$ ($n = 0 - 127$)	Flexible Filter Value Table – FFVT_45	105
01	800	$1024 + 2*n$ ($n = 0 - 127$)	Flexible Filter Value Table – FFVT_67	105
01	800	$768 + n$ ($n = 0 - 127$)	Flexible Filter Mask Table – FFMT	105
01	800	$896 + n$ ($n = 0 - 3$)	Flexible Filter Length Table – FFLT03	106
01	800	$904 + n$ ($n = 0 - 1$)	Flexible Filter Length Table – FFLT45	106
01	800	$908 + n$ ($n = 0 - 1$)	Flexible Filter Length Table – FFLT67	107

9.5 PHY Registers (Page 0)

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

Table 14. Control Register - Address 0

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

**Table 14. Control Register - Address 0**

Bits	Field	Type	Default	Description
7	Collision Test	R/W	0b	Enables IEEE 22.2.4.1.9 collision test. 1b = Enable collision test. 0b = Disable collision test.
6	Speed Selection (MSB)	R/W	1b	See description in bit 13.
5:0	Reserved	RO	Always 0x0	Reserved, always set to 0x0.

Table 15. Status Register - Address 1

Bits	Field	Type	Default	Description
15	100BASE-T4	RO	0b	100BASE-T4. This protocol is not supported. This register bit is always set to 0b. 0b = Not 100BASE-T4 capable.
14	100BASE-TX Full-Duplex	RO	1b	1b = 100BASE-TX full duplex capable. 0b = Not 100BASE-TX full duplex capable.
13	100BASE-TX Half-Duplex	RO	1b	1b = 100BASE-TX half duplex capable. 0b = Not 100BASE-TX half duplex capable.
12	10 Mb/s Full-Duplex	RO	1b	1b = 10BASE-T full duplex capable. 0b = Not 10BASE-T full duplex capable.
11	10 Mb/s Half-Duplex	RO	1b	1b = 10BASE-T half duplex capable. 0b = Not 10BASE-T half duplex capable.
10	100BASE-T2 Full-Duplex	RO	0b	Not able to perform 100BASE-T2.
9	100BASE-T2 Half-Duplex	RO	0b	Not able to perform 100BASE-T2.
8	Extended Status	RO	1b	Extended status information in the register Extended Status (0xF).
7	Reserved	RO	0b	Must always be set to 0b.
6	MF Preamble Suppression	RO	1b	1b = PHY accepts management frames with preamble suppressed.
5	Auto-Negotiation Complete	RO	0b	This bit is set after auto-negotiation completes. 1b = Auto-negotiation process complete. 0b = Auto-negotiation process not complete.
4	Remote Fault	RO,LH	0b	This bit indicates that a remote fault has been detected. Once set, it remains set until it is cleared by reading register 1 via the management interface or by PHY reset. 1b = Remote fault condition detected. 0b = Remote fault condition not detected.
3	Auto-Negotiation Ability	RO	1b	1b = PHY able to perform auto-negotiation. 0b = PHY not able to perform auto-negotiation.

**Table 15. Status Register - Address 1**

Bits	Field	Type	Default	Description
2	Link Status	RO,LL	0b	This bit indicates that a valid link has been established. Once cleared, due to link failure, this bit remains cleared until register 1 is read via the management interface. 1b = Link is up. 0b = Link is down.
1	Jabber Detect	RO,LH	0b	1b = Jabber condition detected. 0 = Jabber condition not detected.
0	Extended Capability	RO	1b	Indicates that the PHY provides an extended set of capabilities that might be accessed through the extended register set. For a PHY that incorporates a GMII/RGMII, the extended register set consists of all management registers except registers 0, 1, and 15. 1b = Extended register capabilities.

Table 16. PHY Identifier Register 1 - Address 2

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

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

Table 17. PHY Identifier Register 2 - Address 3

Bits	Field	Type	Default	Description
15:10	PHY ID Number ¹	RO	0x0	The PHY identifier composed of bits 19 through 24 of the OUI.
9:4	Model Number	RO	0x09	The value is part of the PHY identifier and represents the Device Model Number.
3:0	Revision Number	RO	0x0	The value is part of the PHY identifier and represents the Device Revision Number.

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



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

Table 18. Auto-Negotiation Advertisement Register -Address 4

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

Table 19. Auto-Negotiation Link Partner Ability Register - Address 5

Bits	Field	Type	Default	Description
15	Next Page	RO	0b	1b = Link partner has next page ability. 0b = Link partner does not have next page ability.
14	Acknowledge	RO	0b	1b = Link partner has received link code word. 0b = Link partner has not received link code word.
13	Remote Fault	RO	0b	1b = Link partner has detected remote fault. 0b = Link partner has not detected remote fault.
12	Reserved	RO	0b	Reserved.
11	Asymmetric Pause	RO	0b	1b = Link partner requests asymmetric pause. 0b = Link partner does not request asymmetric pause.
10	Pause Capable	RO	0b	1b = Link partner is capable of full duplex pause operation. 0b = Link partner is not capable of pause operation.
9	100BASE-T4 Capability	RO	0b	1b = Link partner is 100BASE-T4 capable. 0b = Link partner is not 100BASE-T4 capable.

**Table 19. Auto-Negotiation Link Partner Ability Register - Address 5**


Bits	Field	Type	Default	Description
8	100BASE-TX Full-Duplex Capability	RO	0b	1b = Link partner is 100BASE-TX full-duplex capable. 0b = Link partner is not 100BASE-TX full-duplex capable.
7	100BASE-TX Half-Duplex Capability	RO	0b	1b = Link partner is 100BASE-TX half-duplex capable. 0b = Link partner is not 100BASE-TX half-duplex capable.
6	10BASE-T Full-Duplex Capability	RO	0b	1b = Link partner is 10BASE-T full-duplex capable. 0b = Link partner is not 10BASE-T full-duplex capable.
5	10BASE-T Half-Duplex Capability	RO	0b	1b = Link partner is 10BASE-T half-duplex capable. 0b = Link partner is not 10BASE-T half-duplex capable.
4:0	Protocol Selector Field	RO	0x00	Link partner protocol selector field.

Table 20. Auto-Negotiation Expansion Register - Address 6

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

Table 21. Auto-Negotiation Next Page Transmit Register - Address 7

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

**Table 22. Link Partner Next Page Register - Address 8**

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

Table 23. 1000BASE-T Control PHY Register - Address 9

Bits	Field	Type	Default	Description
15:13	Test Mode	R/W	000b	000b = Normal mode. 001b = Test Mode 1 - Transmit waveform test. 010b = Test Mode 2 - Master transmit jitter test. 011b = Test Mode 3 - Slave transmit jitter test. 100b = Test Mode 4 - Transmit distortion test. 101b, 110b, 111b = Reserved.
12	Master/Slave Manual Configuration Enable	R/W	0b	1b = Enables master/slave configuration. 0b = Automatic master/slave configuration.
11	Master/Slave Configuration Value	R/W	0b	Setting this bit has no effect unless address 9, bit 12 is set. 1b = configures PHY as a master. 0b = Configures PHY as a slave.
10	Port Type	R/W	0b	1b = Multi-port device (prefer master). 0b = Single port device (prefer slave).
9	Advertise 1000BASE-T Full- Duplex Capability	R/W	0b	1b = Advertises 1000BASE-T full-duplex capability. 0b = Advertises no 1000BASE-T full-duplex capability.
8	Advertise 1000BASE-T Half-Duplex Capability	R/W	0b	1b = Advertises 1000BASE-T half-duplex capability. 0b = Advertises no 1000BASE-T half-duplex capability. Note: 1000BASE-T half-duplex not supported.
7:0	Reserved	RO	0x00	Set these bits to 0x00.

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




Table 24. 1000BASE-T Status Register - Address 10

Bits	Field	Type	Default	Description
15	Master/Slave Configuration Fault	RO,LH,SC	0b	Once set, this bit remains set until cleared by the following actions: - Read of Register 10 via the management interface. - Reset. - Auto-negotiation completed. - Auto-negotiation enabled. 1b = Master/slave configuration fault detected. 0b = No master/slave configuration fault detected.
14	Master/Slave Configuration Resolution	RO	0b	This bit is not valid when bit 15 is set. 1b = Local PHY resolved to master. 0b = Local PHY resolved to slave.
13	Local Receiver Status	RO	0b	1b = Local receiver is correct. 0b = Local receiver is incorrect.
12	Remote Receiver Status	RO	0b	1b = Remote receiver is correct. 0b = Remote receiver is incorrect.
11	Link Partner 1000BASE-T Full-Duplex Capability	RO	0b	1b = Link partner 1000BASE-T full-duplex capable. 0b = Link partner not 1000BASE-T full-duplex capable. Note: Logically, bits 11:10 might be regarded as an extension of the Technology Ability field in Register 5.
10	Link Partner 1000BASE-T Half-Duplex Capability	RO	0b	1b = Link partner 1000BASE-T half-duplex capable. 0b = Link partner not 1000BASE-T half-duplex capable. Note: Logically, bits 11:10 might be regarded as an extension of the Technology Ability field in Register 5.
9:8	Reserved	RO	00b	Reserved
7:0	Idle Error Count	RO	0x00	These bits contain a cumulative count of the errors detected when the receiver is receiving idles and both local and remote receiver status are operating correctly. The count is held at 255 in the event of overflow and is reset to zero by reading Register 10 via the management interface or by reset. MSB of idle error count.

Table 25. Extended Status Register - Address 15


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

**Table 26. PHY Control Register 2 - Address 18**

Bits	Field	Type	Default	Description
15	Resolve MDI/MDI-X Before Forced Speed	R/W	1b	1b = Resolves MDI/MDI-X configuration before forcing speed. 0b = Does not resolve MDI/MDI-X configuration before forcing speed.
14	Count False Carrier Events	R/W	0b	Count symbol errors (bit 13) and count false carrier events (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (refer to Table 27). The default is to count CRC errors. 1b = Rx error counter counts false carrier events. 0b = Rx error counter does not count false carrier events.
13	Count Symbol Errors	R/W	0b	Count symbol errors (bit 13) and count false carrier events (bit 14) control the type of errors that the Rx error counter (Register 20, bits 15:0) counts (refer to Table 27). The default is to count CRC errors. 1b = Rx error counter counts symbol errors. 0b = Rx error counter counts CRC errors.
12:11	Reserved			Reserved.
10	Automatic MDI/MDI-X	R/W	1b	1b = Enables automatic MDI/MDI-X configuration. 0b = Disables automatic MDI/MDI-X configuration.
9	MDI-MDI-X Configuration	R/W	0b	1b = Manual MDI-X configuration. 0b = Manual MDI configuration. Refer to Table 28 for further information.
8:3	Reserved		0x0	Reserved.
2	Enable Diagnostics	R/W	0b	This bit enables PHY diagnostics, which include IP phone detection and TDR cable diagnostics. It is not recommended to enable this bit in normal operation (when the link is active). This bit does not need to be set for link analysis cable diagnostics. 1b = Enables diagnostics. 0b = Disables diagnostics.
1:0	Reserved		0x0	Reserved.

Table 27. Rx Error Counter Characteristics

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



Bit 9 of the PHY Control register manually sets the MDI/MDI-X configuration if automatic MDI-X is disabled (refer to [Table 28](#)).

Table 28. MDI / MDI-X Configuration Parameters

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

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

Table 29. MDI / MDI-X Pin Mapping

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

Table 30. Loopback Control Register - Address 19

Bits	Field	Type	Default	Description
15	MII	R/W	0b	1b = MII loopback selected. 0b = MII loopback not selected.
14:13	Reserved			Reserved
12	All Digital	R/W	1b	1b = All digital loopback enabled. 0b = All digital loopback disabled.
11	Reserved			Reserved
10	Line Driver	R/W	0b	1b = Line driver loopback enabled. 0b = Line driver loopback disabled.
9	Remote	R/W	0b	1b = Remote loopback enabled. 0b = Remote loopback disabled.
8	Reserved	R/W		Reserved
7	External Cable	R/W	0b	1b = External cable loopback enabled. 0b = External cable loopback disabled.

**Table 30. Loopback Control Register - Address 19**

Bits	Field	Type	Default	Description
6	Tx Suppression	R/W	1b	1b = Suppress Tx during all digital loopback. 0b = Do not suppress Tx during all digital loopback.
5:1	Reserved			Reserved
0	Force Link Status	R/W	1b	This bit can be used to force link status operational during MII loopback. In MII loopback, the link status bit is not set unless force link status is used. In all other loopback mode, the link status bit is set when the link comes up. 1b = Forces link status operational in MII loopback. 0b = Forces link status not operational in MII loopback.

9.5.1 Loopback Mode Settings

Table 31 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 31. Loopback Bit (Register 0, Bit 14) Settings for Loopback Mode

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

Table 32. Rx Error Counter Register - Address 20

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

Table 33. Management Interface (MI) Register - Address 21

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

**Table 34. PHY Configuration Register - Address 22**

Bits	Field	Type	Default	Description
15	CRS Transmit Enable	R/W	0b	1b = Enables CRS on transmit in half-duplex mode. 0b = Disables CRS on transmit.
14	Ignore 10G Frames	R/W	1b	1b = Management frames with ST = <00> are ignored. 0b = Management frames with ST = <00> are treated as incorrect frames.
13:12	Transmit FIFO Depth (1000BASE-T)	R/W	00b	00b = +/-8. 01b = +/-16. 10b = +/-24. 11b = +/-32.
11:10	Automatic Speed Downshift Mode	R/W	11b	If automatic downshift is enabled and the PHY fails to auto-negotiate at 1000BASE-T, the PHY falls back to attempt connection at 100BASE-TX and, subsequently, 10BASE-T. This cycle repeats. If the link is broken at any speed, the PHY restarts this process by re-attempting connection at the highest possible speed (1000BASE-T). 00b = Automatic speed downshift disabled. 01b = 10BASE-T downshift enabled. 10b = 100BASE-TX downshift enabled. 11b = 100BASE-TX and 10BASE-T enabled.
9:8	Reserved			Reserved
7	Alternate Next Page	RO	0b	1b = Enables manual control of 1000BASE-T next pages only. 0b = Normal operation of 1000BASE-T next page exchange.
6	Group MDIO Mode Enable	R/W	0b	1b = Enables group MDIO mode. 0b = Disables group MDIO mode.
5	Transmit Clock Enable	R/W	0b	When this bit is set, the transmit test clock is available on pin TX_TCLK. 1b = Enables output of mixer clock (transmit clock in 1000BASE-T). 0b = Disables output.
4:0	Reserved		0x0	Reserved

**Table 35. PHY Control Register - Address 23**

Bits	Field	Type	Default	Description
15	IP Phone Detected	RO	0b	This bit is only valid when linking is disabled (bit 13 = 0b) and when IP phone detect is enabled (bit 14 = 1b). 1b = IP phone detected. 0b = IP phone not detected.
14	IP Phone Detect Enable	R/W	0b	When this bit is set, the PHY performs automatic IP phone detection each time linking is disabled. Linking is disabled when LNK_EN is cleared (bit 13 = 0b). If an IP phone is detected, it is indicated by bit 15. 1b = Enables automatic IP phone detect. 0b = Disables automatic IP phone detect.
13	LNK_EN (Link Enable)	R/W	1b	If LNK_EN is set, the PHY attempts to bring up a link with a remote partner and monitors the MDI for link pulses. If LNK_EN is cleared, the PHY takes down any active link, goes into standby, and does not respond to link pulses from a remote link partner. In standby, IP phone detect and TDR functions are available. 1b = Enables linking. 0b = Disables linking.
12:10	Automatic Speed Downshift Attempts Before Downshift	R/W	100b	000b = 1. 001b = 2. 010b = 3. 011b = 4. 100b = 5. 101b = 6. 110b = 7. 111b = 8.
9:8	Reserved		0x0	Reserved
7	Link Partner Detected	RO, LH	0b	When linking is disabled, the PHY automatically monitors for the appearance of a link partner and sets this bit if detected. Linking is disabled when LNK_EN is cleared (bit 13 = 0b). 1b = Link partner detected. 0b = Link partner not detected.
6	Jabber (10BASE-T)	R/W	1b	1b = Disables jabber. 0b = Normal operation.
5	SQE (10BASE-T)	R/W	0b	1b = Enables heartbeat. 0b = Disables heartbeat.
4	TP_LOOPBACK (10BASE-T)	R/W	0b	1b = Disables TP loopback during half duplex. 0b = Normal operation.
3:2	10BASE-T Preamble Length	R/W	10b	00b = 10BASE-T preamble length of zero bytes sent. 01b = 10BASE-T preamble length of one byte sent. 10b = 10BASE-T preamble length of twobytes sent. 11b = 10BASE-T preamble length of seven bytes sent.
1	Reserved		0b	Reserved
0	Force Interrupt	R/W	0b	1b = Asserts MDINT_N pin. 0b = Deasserts MDINT_N pin.




Table 36. Interrupt Mask Register - Address 24

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

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

**Table 37. Interrupt Status Register - Address 25**

Bits	Field	Type	Default	Description
15:11	Reserved		0x0	Reserved.
10	TDR/IP Phone	RO, LH	0b	1b = Event completed. 0b = Event has not completed.
9	MDIO Sync Lost	RO, LH	0b	If the management frame preamble is suppressed (MF preamble suppression, Register 0, bit 6), it is possible for the PHY to lose synchronization if there is a glitch at the interface. The PHY can recover if a single frame with a preamble is sent to the PHY. The MDIO sync lost interrupt can be used to detect loss of synchronization and, thus, enable recovery. 1b = Event has occurred. 0b = Event has not occurred.
8	Auto-Negotiation Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
7	CRC Errors	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
6	Next Page Received	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
5	Error Count Full	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
4	FIFO Overflow/ Underflow	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
3	Receive Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
2	Link Status Change	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
1	Automatic Speed Downshift	RO, LH	0b	1b = Event has occurred. 0b = Event has not occurred.
0	MII Interrupt Pending	RO, LH	0b	An event has occurred and the corresponding interrupt mask bit is enabled (set to 1b). 1b = Interrupt pending. 0b = No interrupt pending.



Table 38. PHY Status Register - Address 26

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

**Table 39. LED Control Register 1 - Address 27**

Bits	Field	Type	Default	Description
15	Two-Color Mode LED_100/LED_10	R/W	0b	If two-color mode is enabled for pair LED_LNK/ACT and LED_1000, the signal output for LED_LNK/ACT is equal to LED_LNK/ACT and LED_1000. When LED_LNK/ACT and LED_1000 are not mutually exclusive (such as duplex and collision), this mode can simplify the external circuitry because it ensures either LED_LNK/ACT and LED_1000 is on, and not both at the same time. The same rule applies to pair LED_100 and LED_10. 1b = Two-color mode for LED_100 and LED_10. 0b = Normal mode for LED_100 and LED_10.
14	Two-Color Mode LED_LNK/ACT/ LED_1000	R/W	0b	If two-color mode is enabled for pair LED_LNK/ACT and LED_1000, the signal output for LED_LNK/ACT is equal to LED_LNK/ACT and LED_1000. When LED_LNK/ACT and LED_1000 are not mutually exclusive (such as duplex and collision), this mode can simplify the external circuitry because it ensures either LED_LNK/ACT and LED_1000 is on, and not both at the same time. The same rule applies to pair LED_100 and LED_10. 1b = Two-color mode for LED_LNK/ACT and LED_1000. 0b = Normal mode for LED_LNK/ACT and LED_1000.
13	LED_10 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. The default value of this bit is set by the LED_CFG pin at reset. 1b = Extended modes for LED_10. 0b = Standard modes for LED_10.
12	LED_100 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_100. 0b = Standard modes for LED_100.
11	LED_1000 Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_1000. 0b = Standard modes for LED_1000.
10	LED_LNK/ACT Extended Modes	R/W	0b	The LED function is programmed using this bit and Register 28. 1b = Extended modes for LED_LNK/ACT. 0b = Standard modes for LED_LNK/ACT.
9:8	Reserved		0x0	Reserved
7:4	LED Blink Pattern Pause	R/W	0x0	LED blink pattern pause cycles.

**Table 39. LED Control Register 1 - Address 27**

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

Table 40. LED Control Register 2 - Address 28

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

**Table 41. LED Control Register 3 - Address 29**

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

Table 42. Diagnostics Control Register (Linking Disabled) - Address 30

Bits	Field	Type	Default	Description
15:14	TDR Request	R/W, SC	00b	Automatic TDR analysis is enabled by setting TDR request to 11b. All ten combinations of pairs are analyzed in sequence, and the results are available in Register 31. TDR analysis for a single pair combination can be enabled by setting TDR request to 10b. Linking must be disabled (Register 23, bit 13 = 0b) and IP phone detect must be disabled (Register 23, bit 14 = 0b.) to do TDR operations. Bit 15 self clears when the TDR operation completes. When TDR completes, bit 14 indicates if the results are valid. 11b = Automatic TDR analysis in progress. 10b = Single pair TDR analysis in progress. 01b = TDR analysis complete, results valid. 00b = TDR analysis complete, results valid.
13:12	TDR Tx Dim	R/W	00b	Transmit dimension for single-pair TDR analysis/ first dimension to be reported for automatic TDR analysis: 00b = TDR transmit on pair A. 01b = TDR transmit on pair B. 10b = TDR transmit on pair C. 11b = TDR transmit on pair D. The TDR transmit dimension is only valid for single pair TDR analysis. For automatic TDR analysis, these bits specify the first dimension that are reported in Register 31.
11:10	TDR Rx Dim	R/W	00b	Receive dimension for single pair TDR analysis: 00b = TDR receive on pair A. 01b = TDR receive on pair B. 10b = TDR receive on pair C. 11b = TDR receive on pair D. The TDR receive dimension is only valid for single pair TDR analysis. It is ignored for automatic TDR analysis when all 10 pair combinations are analyzed.
9:0	Reserved		0x0	Reserved.



Table 43. Diagnostics Status Register (Linking Disabled) - Address 31

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

**Table 43. Diagnostics Status Register (Linking Disabled) - Address 31**

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

**Table 44. Diagnostics Status Register (Linking Enabled) - Address 31**

Bits	Field	Type	Default	Description
15	Reserved			Reserved.
14	Pair Swap on Pairs C and D	RO	0b	If this bit is set, the PHY has detected that received pair 2 (RJ-45 pins 4 and 5) and pair 3 (RJ-45 pins 7 and 8) have crossed over. 1b = Pairs C and D are swapped (1000BASE-T only). 0b = Pairs C and D are not swapped (1000BASE-T only).
13	Polarity on Pair D	RO	0b	1b = Polarity on pair D is inverted (1000BASE-T only). 0b = Polarity on pair D is normal (1000BASE-T only).
12	Polarity on Pair C	RO	0b	1b = Polarity on pair C is inverted (1000BASE-T only). 0b = Polarity on pair C is normal (1000BASE-T only).
11	Polarity on Pair B	RO	0b	1b = Polarity on pair B is inverted (10BASE-T or 1000BASE-T only). 0b = Polarity on pair B is normal (10BASE-T or 1000BASE-T only).
10	Polarity on Pair A	RO	0b	1b = Polarity on pair A is inverted (10BASE-T or 1000BASE-T only). 0b = Polarity on pair A is normal (10BASE-T or 1000BASE-T only).
9:2	Cable Length	RO	0b	This 8-bit integer value is the cable length in meters when the link is active. The value 0xFF indicates an unknown result. Cable length when the link is active.
1	Reserved			Reserved.
0	Excessive Pair Skew	RO	0b	Excessive pair skew in 1000BASE-T is detected by detecting that the scrambler has not acquired a 1000BASE-T link and cannot be brought up. In this case, the PHY usually falls back to 100BASE-TX or 10BASE-T. It is possible for other scrambler acquisition errors to be mistaken for excessive pair skew. 1b = Excessive pair skew (1000BASE-T only). 0b = No excessive pair skew (1000BASE-T only).



9.6 Port Control Registers (Page 769)

Table 45. Custom Mode Control PHY Address 01, Page 769, Register 16

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

Table 46. Port General Configuration PHY Address 01, Page 769, Register 17

Bits	Name	Type	Default	Description
15:11	Tx Gate Wait IFS	RW	01110b	Determines the size (in nibbles) of non-deferring window from CRS de-assertion.
10:8	BP extension Wait	R/W	100b	Additional waiting byte times after TX Gate Wait IPG expires until the <i>Back Pressure In-band</i> bit is cleared.
7	Reserved	R/W	0b	Reserved
6	Active_PD_enable	R/W	0b	Active Power Down Enable (sD3 Enable) When set to 1b, The integrated LAN controller needs to enter integrated LAN controller power down mode.
5	Reserved		1b	Reserved
4	Host_WU_Active	R/W	0b	Enables host wake up from the 82579. This bit is reset by power on reset only.
3	Wakeup clocks stop	R/W	1b	Wake-up clocks are stopped while wake up is disabled.
2	MACPD_enable	R/W	1b	Written as 1b when the integrated LAN controller needs to globally enable the integrated LAN controller power down feature while the 82579 supports WoL. When set to 1b, pages 800 and 801 are enabled for configuration and <i>Host_WU_Active</i> is not blocked for writes.
1:0	Reserved	RO	00b	Reserved

Table 47. Power Management Control Register PHY Address 01, Page 769, Register 21

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

**Table 48. SMBus Control Register PHY Address 01, Page 769, Register 23**

Name	Default	Bits	Description	Type
Reserved	0x0000	15:2	Reserved	RO
dis_SMB_filtering	0b	1	When set, disables filtering of RX packets for the SMBus. In wake up mode, this configuration is ignored and the filters are enabled.	R/W
Reserved	0b	0	Reserved.	RO

Table 49. Rate Adaptation Control Register PHY Address 01, Page 769, Register 25

Name	Default	Bits	Description	Type
Reserved	0100010b	15:9	Reserved, write as read.	RWP
rx_en_rxdv_preamble	1	8	Enable generation of early preamble based on RX_DV in the receive path.	R/W
rx_en_crs_preamble	0	7	Enable generation of early preamble based on CRS in the receive path.	R/W
reserved	0	6	Reserved, write as read.	RWP
rx_flip_bad_sfd	1	5	Align the packet's start of frame delimiter to a byte boundary in the receive path.	R/W
read_delay_fd	10001b	4:0	Reserved, write as read.	RWP

Table 50. Flow Control Transmit Timer Value PHY Address 01, Page 769, Register 27

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

9.7 Statistics Registers

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

Table 51. Single Collision Count - SCC PHY Address 01, Page 778, Registers 16 - 17

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



This register counts the number of times that a successfully transmitted packet encountered a single collision. This register only increments if transmits are enabled and the 82579 is in half-duplex mode.

Table 52. Excessive Collisions Count - ECOL PHY Address 01, Page 778, Register 18 - 19

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

When 16 or more collisions have occurred on a packet, this register increments, regardless of the value of collision threshold. If collision threshold is set below 16, this counter won't increment. This register only increments if transmits are enabled and the 82579 is in half-duplex mode.

Table 53. Multiple Collision Count - MCC PHY Address 01, Page 778, Register 20 - 21

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

This register counts the number of times that a transmit encountered more than one collision but less than 16. This register only increments if transmits are enabled and the 82579 is in half-duplex mode.

Table 54. Late Collisions Count - LATECOL PHY Address 01, Page 778, Register 23 - 24

Bit	Type	Reset	Description
31:0	RO/V	0x00	LCC Number of packets with late collisions.

Late collisions are collisions that occur after one slot time. This register only increments if transmits are enabled and the 82579 is in half-duplex mode.

Table 55. Collision Count - COLC PHY Address 01, Page 778, Register 25 - 26


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

This register counts the total number of collisions seen by the transmitter. This register only increments if transmits are enabled and the 82579 is in half-duplex mode. This register applies to clear as well as secure traffic.

Table 56. Defer Count - DC PHY Address 01, Page 778, Register 27 - 28

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

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



enabled. The behavior of this counter is slightly different in the 82579 relative to the 82542. For the 82579, this counter does not increment for streaming transmits that are deferred due to TX IPG.

Table 57. Transmit with No CRS - TNCRS PHY Address 01, Page 778, Register 29 - 30

Bit	Type	Reset	Description
31:0	RO/V	0x00	TNCRS Number of transmissions without a CRS assertion from the 82579.

This register counts the number of successful packet transmission in which the CRS input from the 82579 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 82579.

The 82579 should assert CRS during every transmission. Failure to do so might indicate that the link has failed, or the 82579 has an incorrect link configuration. This register only increments if transmits are enabled. This register is only valid when the 82579 is operating at half duplex.

9.8 PCIe Registers

Table 58. PCIe FIFOs Control/Status PHY Address 01, Page 770, Register 16)

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

Table 59. PCIe Power Management Control PHY Address 01, Page 770, Register 17


Name	Default	Bits	Description	Type
Reserved	1b	15	Reserved.	RO
K1 enable ¹	0b	14	Enable K1 Power Save Mode 1b = Enable. 0b = Disable.	RW
Giga_K1_disable	0b	13	When set, the 82579 does not enter K1 while link speed is at 1000Mb/s.	RW
Reserved	100b	12:10	Reserved.	RO
Request a PCIe clock in K1	1b	9	Enables stopping the SERDES PLL in K1 state while link is in 1 Gb/s.	R/W
PLL stop in K1 giga ²	0b	8	Enables stopping the SERDES PLL in K1 state.	R/W
PLL stop in K1	1b	7	Enables stopping SerDes PLL in K1 state (in 10Mbps and 100Mbps). 1b = Disable. 0b = Enable.	R/W
Reserved	0b	6	Reserved	R/W



Table 59. PCIe Power Management Control PHY Address 01, Page 770, Register 17

Reserved	0b	5	Reserved	R/W
Reserved1y	0010b	4:1	Reserved 0x0 - Minimal delay is 816 ns (51 Idle code groups), provided by HW. Each incremental values have a weight of additional 4 Idle code group (4x2x8=64 ns). Added delay varies by setting (0 to 15) from 0ns to 960ns	R/W
Enable Electrical Idle in Cable Disconnect	0b	0	Consider reserved (no entry to Electrical Idle due to Cable Disconnect)	R/W

1. While in SMBus mode, this bit is cleared. To re-enable K1 after switching back to PCIe, this register needs to be re-configured.
2. Not requesting clock in GbE mode is not supported. For example, in GbE mode, the clock should always be present.

**Table 60. In-Band Control PHY Address 01, Page 770, Register 18¹**

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

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

Table 61. PCIe Diagnostic PHY Address 01, Page 770, Register 20¹

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

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

Table 62. Timeouts PHY Address 01, Page 770, Register 21¹

Name	Default	Bits	Description	Type
Reserved	0000b	15:12	Reserved, write as read.	RWP
K1 exit timeout	010100b	11:6	These bits define how much time IDLE symbols are sent on the TX pair after exiting from K1 state before The 82579 starts sending data to the integrated LAN controller (each bit represents 80ns).	R/W
Reserved	010100b	5:0	Reserved	R/W

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



9.9 General Registers

Table 63. PCIe Kstate Minimum Duration Timeout PHY Address 01, Page 770, Register 23¹

Name	Default	Bits	Description	Type
Reserved	0x00	15:5	Reserved, write as read.	R/W
EI_min_dur timeout	0x10	4:0	These bits define the minimum time the 82579 stays in electrical idle state once entered (each bit represents 80 ns).	R/W

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

Table 64. Low Power Idle GPIO Control PHY Address 01, Page 772, Register 18

Name	Default	Bits	Description	Type
Reserved	0x0	15:10	Reserved	R/W
Auto EN LPI	0x0	11	Auto Enable LPI after link up. When set to 0x1 772.20[14:13] will be automatically set by HW after link up.	
TX_LPI_GPIO0	0x0	10	Rouote Tx LPI indication to GPIO 0	R/W
Reserved	0x0	9:0	Reserved	R/W

Table 65. Low Power Idle Control PHY Address 01, Page 772, Register 20

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

**Table 66. OEM Bits PHY Address 01, Page 0, Register 25**

Bits	Field	Mode	HW Rst	Description
15:11	Reserved	R/W	00000b	
10	Aneg_now	R/W	0b	Restart auto-negotiation. This bit is self clearing.
9:7	Reserved	R/W	000b	
6	a1000_dis	R/W	0b ¹	When set to 1b, 1000 Mb/s speed is disabled.
5:3	Reserved	R/W	000b	
2	rev_aneg	R/W	0b	Low Power Link Up Mechanism. Enables a link to come up at the lowest possible speed in cases where power is more important than performance.
1:0	Reserved	R/W	00b	

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

Table 67. SMBus Address PHY Address 01, Page 0, Register 26¹

Name	Default	Bits	Description	Type
Reserved	0x00	15:12	Reserved	RO
SMB fragments size	0b	11	Select SMBus Fragments Size When set to 1b, the fragment size is 64 bytes, otherwise 32 bytes.	RW
APM Enable	0b	10	APM WoL enable.	RW
PEC Enable	1b	9	Defines if the 82579 supports PEC on the SMBus.	RW
SMBus Frequency	0b	8	0b = 100 KHz. 1b = 400 KHz.	RW
SMBus Address Valid	0b	7	0b = Address not valid. 1b = SMBus address valid. This bit is written by the integratedLAN controller when the SMBus Address field is updated. The 82579 cannot send SMBus transactions to the integratedLAN controller unless this bit is set.	RW
SMBus Address	0x00	6:0	This is the integratedLAN controller SMBus address. The 82579 uses it for master functionality.	RW

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

Table 68. Shadow Receive Address Low0 – SRAL0 PHY Address 01, Page 0, Registers 27-28

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

**Table 69. Shadow Receive Address High0 – RAH0 PHY Address 01, Page 0, Registers 29**

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

Table 70. LED Configuration PHY Address 01, Page 0, Register 30

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

NOTES:

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

LED Modes Table

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



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

9.9.1 Interrupts

The 82579 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.

9.10 Wake Up Registers

9.10.1 Accessing Wake Up Registers Using MDIC

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

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

For the 82579 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 82579'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 82579'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 82579's address from the MDI register
- d. REGADD = 0x12 (data cycle for read)
- e. DATA = YYYY (data is valid when the ready bit is set)

For the 82579, 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 82579'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 82579'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 82579's address from the MDI register
 - d. REGADD = 0x12 (data cycle for write)
 - e. DATA = YYYY (data to be written to the register)

9.10.2 Host Wake Up Control Status Register Description

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

Table 71. Receive Control – RCTL PHY Address 01, Page 800, Register 0

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	Unicast Promiscuous Enable (UPE) 0b = Disabled. 1b = Enabled.
RW	1	0b	Multicast Promiscuous Enable (MPE) 0b = Disabled. 1b = Enabled.
RW	2	1b	Slave Access Enable 0b = Access disabled, the filters are active. 1b = Access enabled, the filters are not active.




Table 71. Receive Control – RCTL PHY Address 01, Page 800, Register 0

RW	4:3	00b	Multicast Offset (MO) This determines which bits of the incoming multicast address are used in looking up the bit vector. 00b = [47:38]. 01b = [46:37]. 10b = [45:36]. 11b = [43:34].
RW	5	0b	Broadcast Accept Mode (BAM) 0b = Ignore broadcast (unless it matches through exact or imperfect filters) 1b = Accept broadcast packets.
RW	6	0b	Pass MAC Control Frames. 0b = Do not (specially) pass MAC control frames. 1b = Pass any integrated MAC frame (type field value of 0x8808).
RW	7	0b	Receive Flow Control Enable (RFCE) Indicates that the 82579 responds to the reception of flow control packets. If auto-negotiation is enabled, this bit is set to the negotiated duplex value.
RW	15:9	0x00	Reserved

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

Table 72. Wake Up Control – WUC PHY Address 01, Page 800, Register 1

Attribute	Bit(s)	Initial Value	Description
RW/SN	0	0b	Advance Power Management Enable (APME) If set to 1b, APM wake up is enabled.
RW/V	1	0b	PME_En If set to 1b, ACPI wake up is enabled.
RWC	2	0b	PME_Status This bit is set when the 82579 receives a wake up event.
RO	3	0b	Reserved
RW/SN	4	0b	Link Status Change Wake Enable (LSCWE) Enables wake on link status change as part of APM wake capabilities.
RW/SN	5	0b	Link Status Change Wake Override (LSCWO) If set to 1b, wake on link status change does not depend on the LNK_C bit in the WUFC register. Instead, it is determined by the APM settings in the WUC register.
RO	13:6	0x00	Reserved
RW	14	0b	FLX6 Flexible filter 6enable.
RW	15	0b	FLX7 Flexible filter 7enable.

**Table 73. Wake Up Filter Control – WUFC PHY Address 01, Page 800, Register 2**

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

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

Table 74. Wake Up Status – WUS PHY Address 01, Page 800, Register 3

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

**Table 74. Wake Up Status – WUS PHY Address 01, Page 800, Register 3**

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

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

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

Table 75. Receive Address Low – RAL PHY Address 01, Page 800, Registers 16-17 + $4 * n^1$ ($n=0...6$)

Attribute	Bit(s)	Initial Value	Description
RW	31:0	0	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.

1. While “ n ” is the exact unicast/multicast address entry and it is equals to 0,1,...6.

Table 76. Receive Address High – RAH PHY Address 01, Page 800, Registers 18-19 + $4 * n$ ($n=0...6$)

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address n ($n=0, 1...6$). RAH 0 is loaded from word 0x2 in the NVM.
RO	30:16	0x00	Reserved, reads as 0b and ignored on writes.
RW	31	0	Address valid (AV) When this bit is set, the relevant RAL and RAH are valid (compared against the incoming packet).




Table 77. Shared Receive Address Low – SHRAL PHY Address 01, Page 800, Registers 44-45 + 4*n (n=0...3)

Attribute	Bit(s)	Initial Value	Description
RW	31:0	X	Receive Address Low (RAL) The lower 32 bits of the 48-bit Ethernet address n (n=0...3).

Table 78. Shared Receive Address High – SHRAH PHY Address 01, Page 800, Registers 46-47 + 4*n (n=0...2)

Attribute	Bit(s)	Initial Value	Description
RW	15:0	X	Receive Address High (RAH) The upper 16 bits of the 48-bit Ethernet address n (n=0...3).
RO	30:16	0x00	Reserved, reads as 0b and is ignored on writes.
RW	31	0b	Address valid (AV) When this bit is set, the relevant RAL and RAH are valid (compared against the incoming packet).

Table 79. Shared Receive Address High 3 – SHRAH[3] PHY Address 01, Page 800, Registers 58-59

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

Table 80. IP Address Valid – IPAV¹ PHY Address 01, Page 800, Register 64

Attribute	Bit(s)	Initial Value	Description
RO	0	0b	Reserved
RW	1	0b	V41 IPv4 address 1 valid.
RW	2	0b	V42 IPv4 address 2 valid.
RW	3	0b	V43 IPv4 address 3 valid.
RO	4:14	0x00	Reserved
RW	15	0b	V60 IPv6 address valid.



1. The IP address valid indicates whether the IP addresses in the IP address table are valid.

Table 81. IPv4 Address Table – IP4AT¹ PHY Address 01, Page 800, Registers 82-83 + 2*n (n=0, 1, 2)

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

1. The IPv4 address table is used to store the three IPv4 addresses for ARP/IPv4 request packets and directed IPv4 packet wake ups. It is a 3-entry table with the following format:

Table 82. IPv6 Address Table – IP6AT PHY Address 01, Page 800, Registers 88-89 + 2*n (n=0...3)

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

The IPv6 address table is used to store the IPv6 addresses for directed IPv6 packet wake ups and manageability traffic filtering.

IP6AT might be used by the host.

Table 83. Multicast Table Array – MTA[31:0] PHY Address 01, Page 800, Registers 128-191

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

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

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

Figure 15 shows the multicast lookup algorithm. The destination address shown represents the internally stored ordering of the received destination address. Note that Byte 1 bit 0 shown in Figure 15 is the first on the wire. The bits that are directed to the multicast table array in this diagram match a multicast offset in the CTRL register equals 00b. The complete multicast offset options are:

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





Figure 15. Multicast Table Array Algorithm

Table 84. Flexible Filter Value Table LSB— FFVT_01 PHY Address 01, Page 800, Registers 256 + 2*n (n=0...127)

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

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

In the 82579 since each address contains 16 bits, only the least significant bytes are stored in those addresses.

Table 85. Flexible Filter Value Table MSBs – FFVT_23 PHY Address 01, Page 800, Registers 257 + 2*n (n=0...127)

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



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

In the 82579 since each address contains 16 bits, only the most significant bytes are stored in those addresses.

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

Table 86. Flexible Filter Value Table – FFVT_45 PHY Address 01, Page 800, Registers 512 + 2*n (n=0...127)

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

Table 87. Flexible Filter Value Table – FFVT_67 PHY Address 01, Page 800, Registers 1024 + 2*n (n=0...127)

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

Table 88. Flexible Filter Mask Table – FFMT PHY Address 01, Page 800, Registers 768 + n (n=0...127)

Attribute	Bit(s)	Initial Value	Description
RW	0	X	Mask 0 Mask for filter 0 byte n (n=0, 1... 127).
RW	1	X	Mask 1 Mask for filter 1 byte n (n=0, 1... 127).
RW	2	X	Mask 2 Mask for filter 2 byte n (n=0, 1... 127).
RW	3	X	Mask 3 Mask for filter 3 byte n (n=0, 1... 127).
RW	4	X	Mask 4 Mask for filter 4 byte n (n=0, 1... 127).
RW	5	X	Mask 5 Mask for filter 5 byte n (n=0, 1... 127).



Table 88. Flexible Filter Mask Table – FFMT PHY Address 01, Page 800, Registers 768 + n (n=0...127)

RW	6	X	Mask 6 Mask for filter 6 byte n (n=0, 1... 127).
RW	7	X	Mask 7 Mask for filter 7 byte n (n=0, 1... 127).
RO	15:8	X	Reserved.

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

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

Table 89. Flexible Filter Length Table – FFLT03 PHY Address 01, Page 800, Registers 896 + n (n=0...3)

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

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

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

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

Table 90. Flexible Filter Length Table – FFLT45 PHY Address 01, Page 800, Registers 904 + n (n=0...1)

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

**Table 91. Flexible Filter Length Table – FFLT67 PHY Address 01, Page 800, Registers 908 + n (n=0...1)**


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


9.11 LPI MMD PHY Registers

LPI MMD PHY registers are part of the 82579 EMI registers. These registers are accessed via MDIO by programming the EMI address to register MI16 and reading/writing the data from/to register MI17.

Table 92. 82579 EMI Registers PHY Address 02, Page 0, Registers 16 / 17

IEEE MMD	MMD Bits	EMI Address	EMI Bits	Description	Type
3.0	10	1829	0	Clock stoppable	
3.1	11	182D	3	Tx LP idle received	
3.1	10	182D	2	Rx LP idle received	
3.1	9	182D	1	Tx LP idle indication	
3.1	8	182D	0	RX LP idle indication	
3.20	15:0	0410	15:0	EEE capability register	
3.22	15:0	4C08	15:0	EEE wake error counter in 100BASE-TX mode	
3.22	15:0	4802	15:0	EEE wake error counter in 1000BASE-T mode	
7.60	15:0	040E	15:0	EEE advertisement	
7.61	15:0	040F	15:0	EEE LP advertisement	





10.0 Non-Volatile Memory (NVM)

10.1 Introduction

This section is intended for designs using a 10/100/1000 Mb/s Intel® 6 Series Express Chipset integrated LAN controller in conjunction with the 82579.

There are several LAN clients that might access the NVM such as hardware, LAN driver, and BIOS. Refer to the *Intel® 6 Series Express Chipset External Design Specification (Intel® 6 Series Express Chipset EDS)* and the *Intel® 6 Series Express Chipset SPI Programming Guide* for more details.

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

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

10.2 NVM Programming Procedure Overview

The LAN NVM shares space on an SPI Flash device (or devices) along with the BIOS, Manageability Firmware, and a Flash Descriptor Region. It is programmed through the Intel® 6 Series Express Chipset. This combined image is shown in [Figure 16](#). The Flash Descriptor Region is used to define vendor specific information and the location, allocated space, and read and write permissions for each region. The Manageability (ME) Region contains the code and configuration data for ME functions such as Intel® Active Management Technology. The system BIOS is contained in the BIOS Region. The ME Region and BIOS Region are beyond the scope of this document and a more detailed explanation of these areas can be found in the *Intel® 6 Series Express Chipset Family External Design Specification (Intel® 6 Series Express Chipset EDS)*. This document describes the LAN image contained in the Gigabit Ethernet (GbE) region.





Figure 16. LAN NVM Regions

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

1. A valid Flash Descriptor Region must be present. Details for the Flash Descriptor Region are contained in the *Intel® 6 Series Express Chipset EDS*. This process is described in detail in the *Intel® Active Management Technology OEM Bring-Up Guide*.

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

2. The GbE region must be part of the original image flashed onto the part.
3. For Intel LAN tools and drivers to work correctly, the BIOS must set the VSCC register(s) correctly. There are two sets of VSCC registers, the upper (UVSCC) and lower (LVSCC). Note that the LVSCC register is only used if the NVM attributes change. For example, the use of a second flash component, a change in erase size between segments, etc. Due to the architecture of the Intel® 6 Series Express Chipset, if these registers are not set correctly, the LAN tools might not report an error message even though the NVM contents remain unchanged. Refer to the *Intel® 6 Series Express Chipset EDS* for more information



4. The GbE region of the NVM must be accessible. To keep this region accessible, the Protected Range register of the GbE LAN Memory Mapped Configuration registers must be set to their default value of 0x0000 0000. (The GbE Protected Range registers are described in the *Intel® 6 Series Express Chipset EDS*).
5. The sector size of the NVM must equal 256 bytes, 4 KB, or 64 KB. When a Flash device that uses a 64 KB sector erase is used, the GbE region size must equal 128 KB. If the Flash part uses a 4 KB or 256-byte sector erase, then the GbE region size must be set to 8 KB.

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

10.3 LAN NVM Format and Contents

Table 17 lists the NVM maps for the LAN region. Each word listed is described in detail in the following sections.

Table 17. LAN NVM Address Map

LAN Word Offset	NVM Byte Offset	Used By	15	0	Image Value
0x00	0x00	HW-Shared	Ethernet Address Byte 2, 1		IA (2, 1)
0x01	0x02	HW-Shared	Ethernet Address Byte 4, 3		IA (4, 3)
0x02	0x04	HW-Shared	Ethernet Address Byte 6, 5		IA (6, 5)
0x03	0x06	SW	Reserved		0x0800
0x04	0x08	SW	Reserved		0xFFFF
0x05	0x0A	SW	Image Version Information 1		
0x06	0x0C	SW	Reserved		0xFFFF
0x07	0x0E	SW	Reserved		0xFFFF
0x08	0x10	SW	PBA Low		
0x09	0x12	SW	PBA High		
0x0A	0x14	HW-PCI	PCI Init Control Word		
0x0B	0x16	HW-PCI	Subsystem ID		
0x0C	0x18	HW-PCI	Subsystem Vendor ID		
0x0D	0x1A	HW-PCI	Device ID		0x10EF 0x10EA
0x0E	0x1C	HW-PCI	Reserved		
0x0F	0x1E	HW-PCI	Reserved		
0x10	0x20	HW-PCI	LAN Power Consumption		
0x11	0x22	HW	Reserved		
0x12	0x24		Reserved		
0x13	0x26	HW-Shared	Shared Init Control Word		
0x14	0x28	HW-Shared	Extended Configuration Word 1		
0x15	0x2A	HW-Shared	Extended Configuration Word 2		
0x16	0x2C	HW-Shared	Extended Configuration Word 3		



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

Table notes:

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

10.3.1 Hardware Accessed Words

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

10.3.1.1 Ethernet Address (Words 0x00-0x02)

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


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

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

10.3.1.2 PCI Init Control Word (Word 0x0A)

This word contains initialization values that:

- Sets defaults for some internal registers
- Enables/disables specific features
- Determines which PCI configuration space values are loaded from the NVM



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

10.3.1.3 Subsystem ID (Word 0x0B)

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

10.3.1.4 Subsystem Vendor ID (Word 0x0C)

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

10.3.1.5 Device ID (Word 0x0D)

If the Load Device ID in word 0x0A is set, this word is read in to initialize the Device ID of the 82579 PHY. Default value is 0x1502.

Note: When the 82579V SKU is used in combination with certain chipset SKUs, the default value for this word is 0x1503.

10.3.1.6 Words 0x0E and 0x0F Are Reserved

Default value is 0x0.

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

10.3.1.7 LAN Power Consumption (Word 0x10)

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



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


10.3.1.8 Word 0x12 and Word 0x11 Are Reserved

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

10.3.1.9 Shared Init Control Word (Word 0x13)

This word controls general initialization values.

Bits	Name	Default	Description
15:14	Sign	10b	Valid Indication A 2-bit valid indication field indicates to the device that there is a valid NVM present. If the valid field does not equal 10b the integrated LAN controller does not read the rest of the NVM data and default values are used for the device configuration.
13	MACsec Disable	1b	This bit enables the MACsec logic when set to 0b. It is loaded to the <i>LSecCK</i> bit in the <i>CTRL_EXT</i> register. When set, the MACsec logic is disabled and its clocks are gated. When cleared, the MACsec logic is enabled. Note: Refer to MACsec doc.
12:10	Reserved	001b	Reserved, set to 001b.
9	PHY PD Ena	0b	Enable PHY Power Down When set, enables PHY power down at DMoff/D3 or Dr and no WoL. This bit is loaded to the <i>PHY Power Down Enable</i> bit in the Extended Device Control (<i>CTRL_EXT</i>) register. 1b = Enable PHY power down. 0b = PHY always powered up.
8	Reserved	1b	Reserved, should be set to 1b.
7:6	PHYT	00b	PHY Device Type Indicates that the PHY is connected to the integrated LAN controller and resulted mode of operation of the integrated LAN controller/PHY link buses. 00b = 82579. 01b = Reserved. 10b = Reserved. 11b = Reserved.
5	Reserved	0b	Reserved, should be set to 0b.
4	FRCS PD	0b	Default setting for the <i>Force Speed</i> bit in the Device Control register (<i>CTRL[11]</i>).



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



10.3.1.10 Extended Configuration Word 1 (Word 0x14)

Bits	Name	Default	Description
15:14	Reserved	00b	Reserved, set to 00b.
13	PHY Write Enable	1b	When set, enables loading of the extended PHY configuration area in the 82579. When disabled, the extended PHY configuration area is ignored. Loaded to the EXTCNF_CTRL register.
12	OEM Write Enable	1b	When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[3]). 1b = OEM bits written to the 82579. 0b = No OEM bits configuration.
11:0	Extended Configuration Pointer	0x0028	Defines the base address (in Dwords) of the Extended Configuration area in the NVM. The base address defines an offset value relative to the beginning of the LAN space in the NVM. A value of 0x00 is not supported when operating with the 82579. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[27:16]).


10.3.1.11 Extended Configuration Word 2 (Word 0x15)

Bits	Name	Default	Description
15:8	Extended PHY Length	0x12	Size (in Dwords) of the Extended PHY configuration area loaded to the Extended Configuration Size register (EXTCNF_SIZE[23:16]). If an extended configuration area is disabled by bit 13 in word 0x14, its length must be set to zero.
7:0	Reserved	0x00	Reserved, must be set to 0x00.

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

10.3.1.12 Extended Configuration Word 3 (Word 0x16)

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



10.3.1.13 OEM Configuration Defaults (Word 0x17)

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

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



10.3.1.14 LED 0 - 2 Configuration Defaults (Word 0x18)

This NVM word specifies the hardware defaults for the LED Control (LEDCTL) register fields controlling the LED1 (LINK_1000), LED0 (LINK/ACTIVITY) and LED2 (LINK_100) output behaviors. Refer to the *Intel® 6 Series Express Chipset Family PDG* and the *82579 Reference Schematics* for LED connection details. Also, [Table 18](#) lists mode encodings for LED outputs.

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

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

**Table 18. Mode Encodings for LED Outputs**

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

10.3.1.15 Reserved (Word 0x19)

Bits	Name	Default	Description
15:0	Reserved	0x2B00	Reserved, set to 0x2B00.


Note: When software calculates the checksum, bit 6 of this word is set to 1b to indicate that the checksum is valid after the image is successfully programmed.

10.3.1.16 Reserved (Word 0x1A)

Bits	Name	Default	Description
15:1	Reserved	0x0421	Reserved, set to 0x0421.
0	APM Enable	1b	APM Enable Initial value of Advanced Power Management Wake Up Enable in the Wake Up Control (WUC.APME) register. 1b = Advanced power management enabled. 0b = Advanced power management disabled.

10.3.1.17 Reserved (Word 0x1B)

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



10.3.1.18 Reserved (Word 0x1C)

Bits	Name	Default	Description
15:0	Reserved	0x1502	Reserved

10.3.1.19 Reserved (Word 0x1D)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

10.3.1.20 Reserved (Word 0x1E)

Bits	Name	Default	Description
15:0	Reserved	0x1502	Reserved

10.3.1.21 Reserved (Word 0x1F)

Bits	Name	Default	Description
15:0	Reserved	0x1503	Reserved

10.3.1.22 Reserved (Word 0x20)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

10.3.1.23 Reserved (Word 0x21)

Bits	Name	Default		Description
15:0	Reserved	0xBAAD		Reserved

10.3.1.24 Reserved (Word 0x22)

Bits	Name	Default		Description
15:0	Reserved	0xBAAD		Reserved

10.3.1.25 Reserved (Word 0x23)

Bits	Name	Default		Description
15:0	Reserved	0x1502		Reserved



10.3.1.26 Reserved (Word 0x24)

Bits	Name	Default	Description
15	Reserved	1b	Reserved, set to 1b.
14	Reserved	0b	Reserved, set to 0b.
13:0	Reserved	0x0000	Reserved, set to 0x0000.

10.3.1.27 Reserved (Word 0x25)

Bits	Name	Default	Description
15	Reserved	1b	Reserved, set to 1b.
14:8	Reserved	0x00	Reserved, set to 0x00.
7	Reserved	1b	Reserved, set to 1b.
6:5	Reserved	00b	Reserved, set to 00b.
4	Reserved	1b	Reserved, set to 1b.
3:0	Reserved	0000b	Reserved, set to 0000b.

10.3.1.28 Reserved (Word 0x26)

Bits	Name	Default	Description
15	Reserved	0b	Reserved
14	Reserved	1b	Reserved
13:12	Reserved	00b	Reserved
11	Reserved	1b	Reserved
10	Reserved	1b	Reserved
9	Reserved	1b	Reserved
8:0	Reserved	0x00	Reserved

10.3.1.29 Reserved (Word 0x27)

Bits	Name	Default	Description
15:0	Reserved	0x80	Reserved

10.3.2 Software Accessed Words

10.3.2.1 PXE Words (Words 0x30 Through 0x3E)

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



10.3.2.1.1 Boot Agent Main Setup Options (Word 0x30)

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

Table 19. Boot Agent Main Setup Options

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




10.3.2.1.2 Boot Agent Configuration Customization Options (Word 0x31)

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

Table 20. Boot Agent Configuration Customization Options (Word 0x31)

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



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



10.3.2.1.3 Boot Agent Configuration Customization Options (Word 0x32)

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

Table 21. Boot Agent Configuration Customization Options (Word 0x32)

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

10.3.2.1.4 IBA Capabilities (Word 0x33)

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

**Table 22. IBA Capabilities**

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

10.3.2.2 Checksum Word Calculation (Word 0x3F)

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

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

10.3.3 Basic Configuration Software Words

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

10.3.3.1 Reserved (Word 0x3)

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



10.3.3.2 Reserved (Word 0x04)

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved

10.3.3.3 Image Version Information (Word 0x05)

0x00D3 refers to mobile LAN switch version. 0x00D4 is the image version for all other platform configurations.

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


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

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

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

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

§ §



11.0 Time Synch (IEEE1588 and 802.1AS)

11.1 Overview

IEEE 1588 addresses the clock synchronization requirements of measurement and control systems. The protocol supports system-wide synchronization accuracy in the sub-microsecond range with minimal network and local clock computing resources. The protocol is spatially localized and allows simple systems to be installed and operated without requiring the administrative attention of users.

The IEEE802.1AS standard specifies the protocol used to ensure that synchronization requirements are met for time sensitive applications, such as audio and video, across Bridged and Virtual Bridged Local Area Networks consisting of LAN media where the transmission delays are fixed and symmetrical. For example, IEEE 802.3 full duplex links. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration. It specifies the use of IEEE 1588 specifications where applicable.


Activation of the LAN Controller Time Sync mechanism is possible in full duplex mode only. No limitations on wire speed exist, although wire speed might affect the accuracy. Time Sync protocol is tolerant of dropping packets as well as missing timestamps.

11.1.1 Flow and HW /SW responsibilities

The operation of a PTP (precision time protocol) enabled network is divided into two stages, Initialization and time synchronization.

At the initialization stage every master enabled node starts by sending Sync packets that include the clock parameters of its clock. Upon receipt of a Sync packet a node compares the received clock parameters to its own and if the received parameters are better, then this node moves to Slave state and stops sending Sync packets. When in slave state the node continuously compares the incoming Sync packets to its currently chosen master and if the new clock parameters are better then the master selection is transferred to this better master clock. Eventually the best master clock is chosen. Every node has a defined time-out interval after which if no Sync packet is received from its chosen master clock it moves back to master state and starts sending Sync packets until a new best master clock (BMC) is chosen.

The time synchronization stage is different between master and slave nodes. If a node is at a master state it should periodically send a Sync packet which is time stamped by hardware on the TX path (as close as possible to the PHY). After the Sync packet a Follow_Up packet is sent which includes the value of the timestamp kept from the Sync packet. In addition the master should timestamp Delay_Req packets on its RX path and return, to the slave that sent it, the timestamp value using a Delay_Response packet. A node in Slave state should timestamp every incoming Sync packet and, if it came from its selected master, software will use this value for time offset calculation. In addition it should periodically send Delay_Req packets in order to calculate the path delay from its master. Every sent Delay_Req packet sent by the slave is time stamped and kept. With the timestamp value received from the master with the Delay_Response packet,



the slave can now calculate the path delay from the master to the slave. The synchronization protocol flow and the offset calculation are described in the following figure.




Figure 11-1. Sync flow and offset calculation

The HW responsibilities are:

- Identify the packets that require time stamping.
- Time stamp the packets on both RX and TX paths.
- Store the time stamp value for SW.
- Keep the system time in HW and give a time adjustment service to the SW.

The SW is responsible for:

- Best Master Clock (BMC) protocol execution which determines which clock is the highest quality clock within the network. As a result of the protocol the SW will set the node state (master or slave) and selection of the master clock if the node is in slave state.
- Generate PTP packets, consume PTP packets.
- Calculate the time offset and adjust the system time using HW mechanism.

Table 11-1. Chronological order of events for Sync and path delay

Action	Responsibility	Node Role
Generate a Sync packet with timestamp notification in descriptor.	SW	Master
Timestamp the packet and store the value in registers (T1).	HW	Master
Timestamp incoming Sync packet, store the value in register and store the sourceID and sequenceID in registers (T2).	HW	Slave
Read the timestamp from register T1 and put it in a Follow_Up packet and send.	SW	Master
Once the Follow_Up packet arrives, store T2 from registers and T1 from the Follow_up packet.	SW	Slave
Generate a Delay_Req packet with timestamp notification in descriptor.	SW	Slave



Action	Responsibility	Node Role
Timestamp the packet and store the value in registers (T3).	HW	Slave
Timestamp incoming Delay_Req packet, store the value in register and store the sourceID and sequenceID in registers (T4).	HW	Master
Read the timestamp from the register and send back to the Slave using a Delay_Response packet.	SW	Master
Once the Delay_Response packet arrives, calculate the offset using T1, T2, T3 and T4 values.	SW	Slave

11.1.1.1 TimeSync indications in RX and TX packet descriptors

Some indications need to be transferred between SW and HW regarding PTP packets.

On the transmit path the SW should set the TST bit in the ExtCMD field in the TX advanced descriptor to indicate that the transmit packet time stamp should be captured.

On the RX path the HW has two indications to indicate to the SW in the receive descriptor:

1. An indication that this packet is a PTP packet (no matter if a timestamp was taken or not), this is also for other types of PTP packets needed for management of the protocol, this bit is set only for the L2 type of packets (the PTP packet is identified according to its Ethertype). PTP packets have the PACKETTYPE field set to 0xE to indicate that the Etype matches the filter number set by the SW to filter PTP packets. The UDP type of PTP packets don't need such an indication since the port number (319 for event and 320 all the rest PTP packets) will direct the packets toward the time sync application.
2. An indication using the TST bit in the Extended Status field of the Rx descriptor to indicate to the SW that a time stamp was taken for this packet. SW needs to access the time stamp registers to get the time stamp values.

11.1.2 HW Time sync elements


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

11.1.2.1 System time structure and mode of operation

The time sync logic contains an up counter to maintain the system time value. This is a 64 bit counter that is built using the **SYSTIML** and **SYSTIMH** registers. When in Master state the **SYSTIMH** and **SYSTIML** registers should be set once by the software according to the general system, when in slave state software should update the system time on every sync event as described in [Section 11.1.2.3](#). Setting the system time is done by a direct write to the **SYSTIMH** register and fine tune setting of the **SYSTIM** register using the adjustment mechanism described in [Section 11.1.2.3](#).

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

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



The SYSTIM timer value in **SYSTMIL** and **SYSTMH** registers is updated periodically every **TIMINCA.incperiod** (If **TIMINCA.incperiod** is one, then an increment event should occur on every clock cycle).

Upon increment event the system time value should increment its value by the value stored in **TIMINCA.incvalue**. The *incvalue* defines the granularity in which the time is represented by the **SYSTMH/L** registers. For example if the cycle time is 16ns and the *incperiod* is 1 then if the *incvalue* is 16 then the time is represented in nanoseconds if the *incvalue* is 160 then the time is represented in 0.1ns units and so on. The *incperiod* helps to avoid inaccuracy in cases where T value can not be represented as a simple integer and should be multiplied to get to an integer representation. The *incperiod* value should be as small as possible to achieve best accuracy possible. For more details please refer to [Section 1.10.3.8.13](#)

The clock used has stable 96MHz or 25Mhz frequency that is always active at S0 state, independent of the link speed.

Note: System time registers should be implemented on a free running clock to make sure the system time is kept valid on traffic idle times (dynamic clock gating).

11.1.2.2 Time stamping mechanism

The time stamping logic is located on TX and RX paths at a location as close as possible to the PHY. This is to reduce delay uncertainties originating from implementation differences. The operation of this logic is slightly different on TX and on RX.

The TX part decides to timestamp a packet if the TX timestamp is enabled and the time stamp bit in the packet descriptor is set. On the TX side only the time is captured.

On the RX side the logic parses the traversing frame and if it is matching the message type defined in register described in [Section 1.10.3.8.6](#) the time, sourceId and sequenceId are latched in the timestamp registers. In addition two bits in the RX descriptor are set, one to identify that this is a PTP packet (this bit is set only for L2 packets since on the UDP packets the port number will direct the packet to the application) and the second (TS) to identify that a time stamp was taken for this packet. If this PTP packet is not Sync or Delay_Req or for some reason time stamp was not taken only the first bit is set.

For more details please refer to the timestamp registers section. The following figure defines the exact point where the time value should be captured.

On both sides the timestamp values is locked in the registers until SW access. This means that if a new PTP packet that requires time stamp has arrived before SW access is not time stamped. In some cases on the RX path a packet that was timestamped might be lost and not get to the host, to avoid lock condition the SW should keep a watch dog timer to clear locking of the time stamp register. The value of such timer should be at least higher then the expected interval between two Sync or Delay_Req packets depends the state (Master or Slave).




A node in the Time Sync network can be in one of two states Master or Slave. When a Time Sync entity is at Master state it synchronizes other entities to its System Clock through the sending out of TimeSync, Follow-up and delay response packets. Master nodes require no time adjustments. Slave nodes adjust their system clocks by using the data arrived with the Follow_Up and Delay_Response packets and to the time stamp values of Sync and Delay_Req packets. When having all the values the SW on the slave node can calculate its offset in the following manner:

1. Write the lower portion of the offset to **TIMADJL**.
2. Write the high portion of the offset to **TIMADJH** to the lower 31 bits and the sign to the most significant bit.

11.1.3 PTP packet structure

The time sync implementation supports both the 1588 V1 and V2 PTP frame formats. The V1 structure can come only as UDP payload over IPv4 while the V2 can come over L2 with its Ethertype or as a UDP payload over IPv4 or IPv6. The 802.1AS uses only the layer2 V2 format.

Offset in bytes	V1 fields	V2 fields	
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	
0	versionPTP	transportSpecific ¹	messageId
1		Reserved	versionPTP
2	versionNetwork	messageLength	
3			



Offset in bytes	V1 fields	V2 fields
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0
4	Subdomain	SubdomainNumber
5		Reserved
6		flags
7		
8		
9		
10		correctionNs
11		
12		
13		
14		correctionSubNs
15		
16		reserved
17		
18		
19		
20	messageType	Reserved
21	Source communication technology	Source communication technology
22	Sourceuuid	Sourceuuid
23		
24		
25		
26		
27		
28	sourceportid	sourceportid
29		
30	sequenceId	sequenceId
31		
32	control	control
33	reserved	logMessagePeriod
34	flags	N/A
35		

1. Should be all zero.

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

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

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

Table 11-4. PTP message over layer 4

Ethernet (L2)	IP (L3)	UDP	PTP message
---------------	---------	-----	-------------

When a PTP packet is recognized (by Ethertype or UDP port address) on the RX side, the version should be checked. If it is V1, then the control field at offset 32 should be compared to the control field in the register described at [Section 1.10.3.8.6](#). Otherwise the byte at offset 0 (messageId) should be used for comparison to the messageId field.

The rest of the needed fields are at the same location and the same size for both V1 and V2 versions.


Table 11-5. Message decoding for V1 (the control field at offset 32)

Enumeration	Value
PTP_SYNC_MESSAGE	0
PTP_DELAY_REQ_MESSAGE	1
PTP_FOLLOWUP_MESSAGE	2
PTP_DELAY_RESP_MESSAGE	3
PTP_MANAGEMENT_MESSAGE	4
reserved	5–255

Table 11-6. Message decoding for V2 (messageId field at offset 0)

MessageId	Message Type	Value (hex)
PTP_SYNC_MESSAGE	Event	0
PTP_DELAY_REQ_MESSAGE	Event	1
PTP_PATH_DELAY_REQ_MESSAGE	Event	2
PTP_PATH_DELAY_RESP_MESSAGE	Event	3
Unused		4–7
PTP_FOLLOWUP_MESSAGE	General	8
PTP_DELAY_RESP_MESSAGE	General	9
PTP_PATH_DELAY_FOLLOWUP_MESSAGE	General	A
PTP_ANNOUNCE_MESSAGE	General	B
PTP_SIGNALLING_MESSAGE	General	C
PTP_MANAGEMENT_MESSAGE	General	D
Unused		E–F

If V2 mode is configured in [Section 1.10.3.8.1](#), “RX Time Sync Control register - [TSYNCRXCTL \(0xB620; RW\)](#)” on page 322 then a time stamp should be taken on PTP_PATH_DELAY_REQ_MESSAGE and PTP_PATH_DELAY_RESP_MESSAGE for any value in the message field in register described at [Section 1.10.3.8.6](#).



§ §



12.0 Intel® 6 Series Express Chipset MAC Programming Interface

12.0.1 Registers Byte Ordering

This section defines the structure of registers that contain fields carried over the network. Some examples are L2, L3, L4 fields and MACsec fields.

The following example is used to describe byte ordering over the wire (hex notation):

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

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

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

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

	Byte 3	Byte 2	Byte 1	Byte0
DW address (N)	0x03	0x02	0x01	0x00
DW address (N+4)	...	...	0x05	0x04

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

	Byte 3	Byte 2	Byte 1	Byte0
(DW aligned)	...	...	0x00	0x01

or


(Word aligned)	0x00	0x01	...	...
----------------	------	------	-----	-----

The following exceptions use network ordering:

- All ETHERType fields

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

Byte 3	Byte 2	Byte 1	Byte0
--------	--------	--------	-------



DW address (N)	00	C9	A0	00
DW address (N+4)	...	...	00	00

12.0.2 Register Conventions

All registers in the LAN Controller are defined to be 32 bits, so write cycles should be accessed as 32 bit double-words, There are some exceptions to this rule:

- Register pairs where two 32 bit registers make up a larger logical size

Reserved bit positions: Some registers contain certain bits that are marked as "reserved". These bits should never be set to a value of "one" by software. Reads from registers containing reserved bits may return indeterminate values in the reserved bit-positions unless read values are explicitly stated. When read, these reserved bits should be ignored by software.

Reserved and/or undefined addresses: any register address not explicitly declared in this specification should be considered to be reserved, and should not be written to. Writing to reserved or undefined register addresses may cause indeterminate behavior. Reads from reserved or undefined configuration register addresses may return indeterminate values unless read values are explicitly stated for specific addresses. Reserved fields within defined registers are defined as Read-Only (RO). When writing to these registers the RO fields should be set to their init value. Reading from reserved fields may return indeterminate values.

Initial values: most registers define the initial hardware values prior to being programmed. In some cases, hardware initial values are undefined and will be listed as such via the text "undefined", "unknown", or "X". Some such configuration values may need to be set via NVM configuration or via software in order for proper operation to occur; this need is dependent on the function of the bit. Other registers may cite a hardware default which is overridden by a higher-precedence operation. Operations which may supersede hardware defaults may include a valid NVM load, completion of a hardware operation (such as hardware auto-negotiation), or writing of a different register whose value is then reflected in another bit.

For registers that should be accessed as 32 bit double words, partial writes (less than a 32 bit double word) will not take effect (i.e. the write is ignored). Partial reads will return all 32 bits of data regardless of the byte enables.

Note: Partial reads to read-on-clear registers (e.g. ICR) can have unexpected results since all 32 bits are actually read regardless of the byte enables. Partial reads should not be done.

Note: All statistics registers are implemented as 32 bit registers. Though some logical statistics registers represent counters in excess of 32-bits in width, registers must be accessed using 32-bit operations (e.g. independent access to each 32-bit field).

Note: The LAN Controller supports a single memory access at a time.

See special notes for Multicast Table Arrays in their specific register definitions.

12.0.3 PCI Configuration and Status Registers - CSR Space

12.0.3.1 PCI Register Map

All configuration registers are listed in the table below. These registers are ordered by grouping and are not necessarily listed in order that they appear in the address space.



Register Based Legend:

RW – Read Write register.

RO – Read Only Register.

RO/CR – Read Only Register, Clear on Read.

RO/V – Read Only Register, Read status is not constant.

RW/RO – Read write by FW, Read only by SW.

R/WC – Read Write Clear registers. Writing '0's has no affect. Writing '1's clears the appropriate fields (see detailed description of the specific registers).

RW/V – Read Write register This bit self-clears immediately.

RW/SN – Read Write register initial value loaded from NVM.

RC/WC – Read Clear/ Write Clear registers. Writing '0's has no affect. Writing '1's clears the appropriate fields. Read may also clear the register depending on enablement (see specific registers).

RWC/CR/V – Read Write register clear on read, clear on write.


WO – Write only registers. Reading from these registers does not reflect any meaningful data. Mostly it would be all zero's (see detailed description of the specific registers).

Table 12-7. Register Summary

Offset	Abbreviation	Name	RW
General Register Descriptions			
0x00000	CTRL	Device Control Register	RW
0x00008	STATUS	Device Status Register	RO
0x0000C	STRAP	Strapping Option Register	RO
0x00018	CTRL_EXT	Extended Device Control Register	RW
0x00020	MDIC	MDI Control Register	RW
0x00024	FEXTNVM4	Future Extended4 NVM Register	RW
0x00028	FEXTNVM	Future Extended NVM Register	RW
0x0002C	FEXT	Future Extended Register	RW
0x00030	FEXTNVM2	Future Extended2 NVM Register	RW
0x00034	KUMCTRLSTA	Kumeran control and status registers	RW
0x00038	BUSNUM	Device and Bus Number	RO
0x000F8	LTRV	Latency Tolerance Reporting Value	RW
0x000FC	LPIC	Low Power Idle Control	RW
0x00170	FCTTV	Flow Control Transmit Timer Value	RW
0x05F40	FCRTV	Flow Control Refresh Threshold Value	RW
0x00F00	EXTCNF_CTRL	Extended Configuration Control	RW



Offset	Abbreviation	Name	RW
0x00F08	EXTCNF_SIZE	Extended Configuration Size	RW
0x00F10	PHY_CTRL	PHY Control Register	RW
0x00F18	PCIEANACFG	PCIE Analog Configuration	RW
0x01000	PBA	Packet Buffer Allocation	RW
0x01008	PBS	Packet Buffer Size	RW
0x05B00	DCR	DMA Control Register	RW
Interrupt Register Descriptions			
0x000C0	ICR	Interrupt Cause Read Register	RC/ WC
0x000C4	ITR	Interrupt Throttling Register	RW
0x000C8	ICS	Interrupt Cause Set Register	WO
0x000D0	IMS	Interrupt Mask Set/Read Register	RW
0x000D8	IMC	Interrupt Mask Clear Register	WO
0x000E0	Mask - IAM	Interrupt Acknowledge Auto	RW
Receive Register Descriptions			
0x00100	RCTL	Receive Control Register	RW
0x00104	RCTL1	Receive Control Register 1	RW
0x02008	ERT	Early Receive Threshold	RW
0x02170 + n*0x4 [n=0..1]	PSRCTL	Packet Split Receive Control Register	RW
0x02160	FCRTL	Flow Control Receive Threshold Low	RW
0x02168	FCRTH	Flow Control Receive Threshold High	RW
0x02800 + n*0x100[n=0..1]	RDBAL	Receive Descriptor Base Address Low queue	RW
0x02804 + n*0x100[n=0..1]	RDBAH	Receive Descriptor Base Address High queue	RW
0x02808 + n*0x100[n=0..1]	RDLEN	Receive Descriptor Length queue	RW
0x02810 + n*0x100[n=0..1]	RDH	Receive Descriptor Head queue	RW
0x02818 + n*0x100[n=0..1]	RDT	Receive Descriptor Tail queue	RW
0x02820 + n*0x100[n=0..1]	RDTR	Interrupt Delay Timer (Packet Timer)	RW
0x02828 + n*0x100[n=0..1]	RXDCTL	Receive Descriptor Control	RW
0x0282C	RADV	Receive Interrupt Absolute Delay Timer	RW
0x02C00	RSRPD	Receive Small Packet Detect Interrupt	RW
0x02C08	RAID	Receive ACK Interrupt Delay Register	RW
0x02C10	CPUVEC	CPU Vector Register	RW
0x05000	RXCSUM	Receive Checksum Control	RW
0x05008	RFCTL	Receive Filter Control Register	RW



Offset	Abbreviation	Name	RW
0x05200-0x0527C	MTA[31:0]	Multicast Table Array	RW
0x05400 + 8*n (n=0...6)	RAL	Receive Address Low	RW
0x05404 + 8*n (n=0...6)	RAH	Receive Address High	RW
0x05438 + 8*n (n=0...3)	SRAL	Shared Receive Address Low	RW
0x0543C + 8*n (n=0...2)	SRAH	Shared Receive Address High 0...2	RW
0x05454	SHRAH[3]	Shared Receive Address High 3	RW
0x05818	MRQC	Multiple Receive Queues Command Register	RW
0x05864	RSSIM	RSS Interrupt Mask Register	RW
0x05868	RSSIR	RSS Interrupt Request Register	RW
0x05C00 + 4*n (n=0...31)	RETA	Redirection Table	RW
0x05C80 + 4*n (n=0...9)	RSSRK	Random Key Register	RW
Transmit Register Descriptions			
0x00400	TCTL	Transmit Control Register	RW
0x00410	TIPG	Transmit IPG Register	RW
0x00458	AIT	Adaptive IFS Throttle	RW
0x03800 + n*0x100[n=0..1]	TDBAL	Transmit Descriptor Base Address Low	RW
0x03804 + n*0x100[n=0..1]	TDBAH	Transmit Descriptor Base Address High	RW
0x03808 + n*0x100[n=0..1]	TDLEN	Transmit Descriptor Length	RW
0x03810 + n*0x100[n=0..1]	TDH	Transmit Descriptor Head	RW
0x03818 + n*0x100[n=0..1]	TDT	Transmit Descriptor Tail	RW
0x03840 + n*0x100[n=0..1]	TARC	Transmit Arbitration Count	RW
0x03820	TIDV	Transmit Interrupt Delay Value	RW
0x03828 + n*0x100[n=0..1]	TXDCTL	Transmit Descriptor Control	RW
0x0382C	TADV	Transmit Absolute Interrupt Delay Value	RW
Statistic Register Descriptions			
0x04000	CRCERRS	CRC Error Count	RO
0x04004	ALGNERRC	Alignment Error Count	RO
0x0400C	RXERRC	RX Error Count	RO
0x04010	MPC	Missed Packets Count	RO



Offset	Abbreviation	Name	RW
0x0403C	CEXERR	Carrier Extension Error Count	RO
0x04040	RLEC	Receive Length Error Count	RO
0x04048	XONRXC	XON Received Count	RO
0x0404C	XONTXC	XON Transmitted Count	RO
0x04050	XOFFRXC	XOFF Received Count	RO
0x04054	XOFFTXC	XOFF Transmitted Count	RO
0x04058	FCRUC	FC Received Unsupported Count	RO
0x04074	GPRC	Good Packets Received Count	RO
0x04078	BPRC	Broadcast Packets Received Count	RO
0x0407C	MPRC	Multicast Packets Received Count	RO
0x04080	GPTC	Good Packets Transmitted Count	RO
0x04088	GORCL	Good Octets Received Count Low	RO
0x0408C	GORCH	Good Octets Received Count High	RO
0x04090	GOTCL	Good Octets Transmitted Count Low	RO
0x04094	GOTCH	Good Octets Transmitted Count High	RO
0x040A0	RNBC	Receive No Buffers Count	RO
0x040A4	RUC	Receive Undersize Count	RO
0x040A8	RFC	Receive Fragment Count	RO
0x040AC	ROC	Receive Oversize Count	RO
0x040B0	RJC	Receive Jabber Count	RO
0x040B4	MNGPRC	Management Packets Received Count	RO
0x040B8	MNGPDC	Management Packets Dropped Count	RO
0x040BC	MNGPTC	Management Packets Transmitted Count	RO
0x040D8	TCBPD	Tx Circuit Breaker Packets Dropped	RO
0x040C0	TORL	Total Octets Received Low	RO
0x040C4	TORH	Total Octets Received High	RO
0x040C8	TOTL	Total Octets Transmitted	RO
0x040CC	TOTH	Total Octets Transmitted	RO
0x040D0	TPR	Total Packets Received	RO
0x040D4	TPT	Total Packets Transmitted	RO
0x040F0	MPTC	Multicast Packets Transmitted Count	RO
0x040F4	BPTC	Broadcast Packets Transmitted Count	RO
0x040F8	TSCTC	TCP Segmentation Context Transmitted Count	RO
0x04100	IAC	Interrupt Assertion Count	RO
Management Register Descriptions			
0x05800	WUC	Wake Up Control Register	RW
0x05808	WUFC	Wake Up Filter Control Register	RW
0x05810	WUS	Wake Up Status Register	RW



Offset	Abbreviation	Name	RW
0x5838	IPAV	IP Address Valid	RW
0x05840 + 8*n (n=1...3)	IP4AT	IPv4 Address Table	RW
0x05880 + 4*n (n=0...3)	IP6AT	IPv6 Address Table	RW
0x05F00 + 8*n (n=0...7)	FFLT	Flexible Filter Length Table	RW
0x09000 + 8*n (n=0...127)	FFMT	Flexible Filter Mask Table	RW
0x09800 + 8*n (n=0...127)	FFVT	Flexible Filter Value Table	RW
0x09804 + 8*n (n=0...127)	FFVT2	Flexible Filter Value Table	RW
0x0B620	TSYNCRXCTL	RX Time Sync Control register	RW
0x0B624	RXSTMPL	RX timestamp Low	RO
0x0B628	RXSTMPH	RX timestamp High	RO
0x0B62C	RXSATRL	RX timestamp attributes low	RO
0x0B630	RXSATRH	RX timestamp attributes high	RO
0x0B634	RXMTRL	RX message type register low	RW
0x0B638	RXUDP	RX UDP port	RW
0x0B614	TSYNCTXCTL	TX Time Sync Control register	RW
0x0B618	TXSTMPL	TX timestamp value Low	RO
0x0B61C	TXSTMPH	TX timestamp value High	RO
0x0B600	SYSTIML	System time register Low	RO
0x0B604	SYSTIMH	System time register High	RO
0x0B608	TIMINCA	Increment attributes register	RW
0x0B60C	TIMADJL	Time adjustment offset register low	RW
0x0B610	TIMADJH	Time adjustment offset register high	RW
Diagnostic Register Descriptions			
0x02410	RDFH	Receive Data FIFO Head Register	RW
0x02418	RDFT	Receive Data FIFO Tail Register	RW
0x02420	RDFHS	Receive Data FIFO Head Saved Register	RW
0x02428	RDFTS	Receive Data FIFO Tail Saved Register	RW
0x02430	RDFPC	Receive Data FIFO Packet Count	RW
0x03410	TDFH	Transmit Data FIFO Head Register	RW
0x03418	TDFT	Transmit Data FIFO Tail Register	RW
0x03420	TDFHS	Transmit Data FIFO Head Saved Register	RW
0x03428	TDFTS	Transmit Data FIFO Tail Saved Register	RW
0x03430	TDFPC	Transmit Data FIFO Packet Count	RW
0x10000 - 0x15FFC	PBM	Packet Buffer Memory	RW



Offset	Abbreviation	Name	RW
0x0C000 - 0x0C3FC	GMD	Ghost Memory Data	RW
0x0C400 - 0x0C5FC	DMD	Descriptor Memory Data	RW
Hidden Testability Register Descriptions			
0x05B60	MNGCCR	MNG CSR Control register	RW
0x05B64	MNGCAR	MNG CSR Address Register	RW
0x05B68	MNGCDO	MNG CSR Data Out Register	RO
0x05B6C	MNGCDI	MNG CSR Data In Register	RW
0x05F04	CHIKN	Debug Chicken Register	RW
MACsec Register Descriptions			
0x0B000	LSECTXCAP	MACsec TX Capabilities register	RW
0x0B300	LSECRXCAP	MACsec RX Capabilities register	RW
0x0B004	LSECTXCTRL	MACsec TX Control register	RW
0x0B304	LSECRXCTRL	MACsec RX Control register	RW
0x0B008	LSECTXSCL	MACsec TX SCI Low	RW
0x0B00C	LSECTXSCH	MACsec TX SCI High	RW
0x0B010	LSECTXSA	MACsec TX SA	RW
0x0B018	LSECTXPN0	MACsec TX SA PN 0	RW
0x0B01C	LSECTXPN1	MACsec TX SA PN 1	RW
0x0B01C	LSECTXKEY0 [n]	MACsec TX Key 0 0 + 4*n (n=0...3) 0x0B02	WO
0x0B01C	LSECTXKEY1 [n]	MACsec TX Key 1 0 + 4*n (n=0...3) 0x0B03	WO
0x0B3D0 + 4*n (n=0...3)	LSECRXSCL[n]	MACsec RX SCI Low	RW
0x0B3E0 + 4*n (n=0...3)	LSECRXSCH[n]	MACsec RX SCI High	RW
0x0B310 + 4*n (n=0...7)	LSECRXSA[n]	MACsec RX SA	RW
0x0B330 + 4*n (n=0...7)	LSECRXSAPN	MACsec RX SA PN	RW
0x0B350 + 0x10*n (n=0...7) + 4*m (m=0..3)	LSECRXKEY[n,m]	MACsec RX Key	WO
0x04300	LSECTXUT	Tx Untagged Packet Counter	RC
0x04304	LSECTXPKTE	Encrypted Tx Packets	RC
0x04308	LSECTXPKTP	Protected Tx Packets	RC
0x0430C	LSECTXOCTE	Encrypted Tx Octets	RC
0x04310	LSECTXOCTP	Protected Tx Octets	RC
0x04314	LSECRXUT	MACsec Untagged RX Packet	RC
0x0431C	LSECRXOCTE	MACsec RX Octets Decrypted	RC




Offset	Abbreviation	Name	RW
0x04320	LSECRXOCTP	MACsec RX Octets Validated	RC
0x04324	LSECRXBAD	MACsec RX Packet with Bad Tag	RC
0x04328	LSECRXNOSCI	MACsec RX Packet No SCI	RC
0x0432C	LSECRXUNSCI	MACsec RX Packet Unknown SCI count	RC
0x04330	LSECRXUNCH	MACsec RX Unchecked Packets	RC
0x04340 + 4*n (n=0...3)	LSECRXDELAY[n]	MACsec RX Delayed Packets	RC
0x04350 + 4*n (n=0...3)	LSECRXLATE[n]	MACsec RX Late Packets	RC
0x04360 + 4*n (n=0...7)	LSECRXOK[n]	MACsec RX Packet OK	RC
0x043A0 + 4*n (n=0...7)	LSECRXINV[n]	MACsec Check RX Invalid	RC
0x04380 + 4*n [n=0...7]	LSECRXNV[n]	MACsec RX Not valid count	RC
0x043C0 + 4*n (n=0...3)	LSECRXNUSA[n]	MACsec RX Not Using SA	RC
0x043D0 + 4*n (n=0...3)	LSECRXUNSA[n]	MACsec RX Unused SA	RC

12.0.3.2 General Register Descriptions

12.0.3.2.1 Device Control Register - CTRL (0x00000; RW)

Bit	Type	Default	Description
0	RW/SN	1	Full Duplex (FD). 0 – half duplex 1 – full duplex. Controls the MAC duplex setting when explicitly setting by software. Loaded from the NVM word 13h.
1	RO	0	Reserved. Write as 0 for future compatibility
2	RW	0	Master Disable. When set, the LAN Controller blocks new master requests on the PCI device. Once no master requests are pending by this function, the <i>Master Enable Status</i> bit is cleared.
5:3	RO	0	Reserved. Write as 0 for future compatibility.
6	RO	1	Reserved.
7	RO	0	Reserved. Must be set to '0'.
9:8	RW	10b	Speed selection (SPEED). These bits may determine the speed configuration and are written by software after reading the PHY configuration through the MDIO interface. These signals are ignored when Auto-Speed Detection is enabled. (00)b – 10Mb/s (01)b – 100Mb/s (10)b – 1000Mb/s (11)b – not used



10	RO	0	Reserved. Write as 0 for future compatibility.
11	RW/SN	0	Force Speed (FRCSPD). This bit is set when software wants to manually configure the MAC speed settings according to the SPEED bits above. When using a PHY device, note that the PHY device must resolve to the same speed configuration or software must manually set it to the same speed as the MAC. The value is loaded from word 13h in the NVM. Note that this bit is superseded by the CTRL_EXT.SPD_BYPS bit which has a similar function.
12	RW	0	Force Duplex (FRCPLX). When set to 1, software may override the duplex indication from the PHY that is indicated in the FDX to the MAC. Otherwise, the duplex setting is sampled from the PHY FDX indication into the MAC on the asserting edge of the PHY LINK signal. When asserted, the CTRL.FD bit sets duplex.
13	RO	0	Reserved.
15:14	RO	0	Reserved. Reads as 0.
16	RW	0	LANPHYPC override When set to 1 this bit provides the SW driver the ability to control the LANPHYPC pin value.
17	RW	0	LANPHYPC Value When LANPHYPC override is set to 1 this bit will define the value of the LANPHYPC pin.
23:18	RO	0x0	Reserved.
24	RW	0	LCD Power Down (LCDPD). When the bit is cleared to '0', the LCD power down setting is controlled by the internal logic of the LAN controller. When set to '1' and the CTRL_EXT.PHYPDEN is set as well, the LAN controller sets the external LCD to power down mode using the LANPHYPC.
26	RW/V	0	Host Software Reset (SWRST). This bit performs a reset to the PCI data path and the relevant shared logic. Writing 1 initiates the reset. This bit is self-clearing.
27	RW	0	Receive Flow Control Enable (RFCE). Indicates the device will respond to the reception of flow control packets. If Auto-Negotiation is enabled, this bit is set to the negotiated duplex value.
28	RW	0	Transmit Flow Control Enable (TFCE). Indicates the device will transmit flow control packets (XON & XOFF frames) based on receiver fullness. If Auto-Negotiation is enabled, this bit is set to the negotiated duplex value.
29	RO	0	Reserved.
30	RW	0	VLAN Mode Enable (VME). When set to 1, all packets transmitted from LAN Controller that have VLE set will be sent with an 802.1Q header added to the packet. The contents of the header come from the transmit descriptor and from the VLAN type register. On receive, VLAN information is stripped from 802.1Q packets.



31	RW/V	0	LAN Connected Device Reset (LCD_RST). Controls a 0 – normal (operational) 1 – reset to PHY is asserted. The LCD_RST functionality is gated by the FWSM.RSPCIPHY bit. If the FWSM.RSPCIPHY bit is not set to '1', then setting the LCD_RST has no impact. For proper operation Software or Firmware must also set the SWRST bit in the register at the same time. This bit is self-clearing.
----	------	---	---

Note: Fields loaded from the NVM are set by the NVM only if the signature bits of the NVM's Initialization Control Word match (01)b.

This register, as well as the Extended Device Control register (CTRL_EXT), controls the major operational modes for the device. While software writes to this register to control device settings, several bits (such as FD and SPEED) may be overridden depending on other bit settings and the resultant link configuration is determined by the PHY's Auto-Negotiation resolution.

The FD (duplex) and SPEED configurations of the device are normally determined from the link configuration process. Software may specifically override/set these MAC settings via certain bits in a forced-link scenario; if so, the values used to configure the MAC must be consistent with the PHY settings.

Manual link configuration is controlled through the PHY's MII management interface.


Host Software Reset (bit 26), may be used to globally reset the entire host data path . This register is provided primarily as a last-ditch software mechanism to recover from an indeterminate or suspected hung hardware state. Most registers (receive, transmit, interrupt, statistics, etc.), and state machines will be set to their power-on reset values, approximating the state following a power-on or PCI reset. One internal configuration register, the Packet Buffer Allocation (PBA) register, retains its value through a Software reset.

Note: To ensure that global device reset has fully completed and that the controller will respond to subsequent accesses, one must wait approximately 1 microsecond after setting before attempting to check to see if the bit has cleared or to access (read or write) any other device register.

Note: This register's address is reflected also at address 0x00004 for legacy reasons. Neither the SW driver nor the FW should use it since it may be unsupported in next generations.

12.0.3.2.2 Device Status Register - STATUS (0x00008; RO)

Bits	Attribute	Default	Description
0	RO/V	X	Full Duplex (FD). 0 – half duplex 1 – full duplex Reflects duplex setting of the MAC and/or link.
1	RO/V	X	Link up (LU). 0 – no link established 1 – link established For this to be valid, the <i>Set Link Up</i> bit of the <i>Device Control Register</i> (CTRL.SU) must be set.



3,2	RO/V	00b	PHY Type Indication (PHYTYPE). Indicates the LAN Connected Device attached to LAN controller and the resulting mode of operation of the MAC/LAN Connected Device Link buses. <table><tr><td>PHYTYPE</td><td>PHY-Device</td></tr><tr><td>00</td><td>82579</td></tr><tr><td>01</td><td>Reserved</td></tr><tr><td>10</td><td>Reserved.</td></tr><tr><td>11</td><td>Reserved</td></tr></table> This field is loaded from the Shared Init control word in the NVM.	PHYTYPE	PHY-Device	00	82579	01	Reserved	10	Reserved.	11	Reserved
PHYTYPE	PHY-Device												
00	82579												
01	Reserved												
10	Reserved.												
11	Reserved												
4	RO/V	X	Transmission Paused (TXOFF). Indication of pause state of the transmit function when symmetrical flow control is enabled.										
5	RO/V	1	PHY Power Up not (PHYPWR). RO bit that indicates the power state of the PHY. 0 – The PHY is powered on in the active state. 1 – The PHY is in the power down state The PHYPWR bit is valid only after PHY reset is asserted. Note: The PHY Power Up indication reflects the status of the LANPHYPC signaling to the LCD.										
7:6	RO/V	X	Link speed setting (SPEED). This bit reflects the speed setting of the MAC and/or link. (00)b – 10Mb/s (01)b – 100Mb/s (10)b – 1000Mb/s (11)b – 1000Mb/s										
8	RO/V	X	Master Read Completions Blocked. This bit is set when the device receives a completion with an error (EP = 1 or status unsuccessful); It is cleared on PCI reset.										
9	RW/V/C	0	LAN Init Done. This bit is Asserted following completion of the LAN initialization from the FLASH. See “LAN Init Done Event” section for a complete description. Software is expected to clear this field to make it usable for the next Init done event.										
10	RW/V/C	1	PHY Reset Asserted (PHYRA). This bit is R/W. The HW sets this bit following the assertion of LCD reset (either HW or in-band). The bit is cleared on writing '0' to it.										
18:11	RO	0	Reserved.										
19	RO/V	1	Master Enable Status. Cleared by LAN Controller when the <i>Master Disable</i> bit is set and no master requests are pending by this function, otherwise this bit is set. This bit indicates that no master requests will be issued by this function as long as the <i>Master Disable</i> bit is set.										
29:20	RO	0	Reserved. Reads as 0.										
30	RO	0	Reserved.										
31	RO/SN	1	Clock Control ¼ (CLK_CNT_1_4). This bit is loaded from the NVM word 13h and indicates the device supports lowering its DMA clock to ¼ of its value.										



FD reflects the actual MAC duplex configuration. This normally reflects the duplex setting for the entire link, as it normally reflects the duplex configuration negotiated between the PHY and link partner (copper link) or MAC and link partner (fiber link).

Link up provides a useful indication of whether something is attached to the port. Successful negotiation of features/link parameters results in link activity. The link startup process (and consequently the duration for this activity after reset) may be several 100's of ms. It reflects whether the PHY's LINK indication is present. Refer to section for more details.

TXOFF indicates the state of the transmit function when symmetrical flow control has been enabled and negotiated with the link partner. This bit is set to 1 when transmission is paused due to the reception of an XOFF frame. It is cleared upon expiration of the pause timer or the receipt of an XON frame.

SPEED indicates the actual MAC speed configuration. These bits normally reflect the speed of the actual link, negotiated by the PHY and link partner, and reflected internally from the PHY to the MAC. These bits may represent the speed configuration of the MAC only, if the MAC speed setting has been forced via software (CTRL.SPEED). Speed indications are mapped as shown below:

(00)b – 10Mb/s

(01)b – 100Mb/s

(10)b – 1000Mb/s

(11)b – 1000Mb/s

12.0.3.2.3 Strapping Option Register - STRAP (0x0000C; RO)

This register reflects the values of the soft strapping options fetched from the NVM descriptor in the CH space.

Attribute	Bit(s)	Initial Value	Description
RO	0	1	Reserved.
RO	5:1	0	LAN NVM Size (NVMS). LAN NVM Space size is indicated in multiples of 4KB. LAN NVM size may vary from 4KB to 128KB while 0 value means 4KB.
RO	16	0	LC SMBus address enable (LCSMBADDEN)
RO	23:17	0	LC SMBus address (LCSMBADD)
RO	24	0	LCD SMBus address enable (LCDSMBADDEN)
RO	31:25	0	LCD SMBus address (LCDSMBADD)

12.0.3.2.4 Extended Device Control Register - CTRL_EXT (0x00018; RW)

Bits	Type	Default	Description
11:0	RO	0	Reserved.
12	RW/V	1	MACsec Clock Gate (LSecCK). When cleared, the MACsec logic gets its clocks. When the LSecCK is set the MACsec logic (including all its CSR registers) do not get any clocks. This bit is loaded from NVM word 13h.
14:13	RO	0	Reserved.



15	RW	0	Speed Select Bypass (SPD_BYPS) . When set to 1, all speed detection mechanisms are bypassed, and the device is immediately set to the speed indicated by CTRL.SPEED. This provides a method for software to have full control of the speed settings of the device when the change takes place by overriding the hardware clock switching circuitry.
18:16	RO	0	Reserved.
19	RW/SN	0	Dynamic Clock Gating (DynCK) . When set, this bit enables dynamic clock gating of the DMA and MAC units. Also see the description of the DynWakeCK in this register. The bit is loaded from NVM word 13h.
20	RW/SN	1	PHY Power Down Enable (PHYPDEN) . When set, this bit enables the LCD to enter a low-power state when the LAN controller is at the DMoff / D3 or Dr with no WoL. This bit is loaded from word 13h in the NVM.
21	RO	0	Reserved.
22	RO	1	Reserved.
23	RO	0	Reserved.
24	RO	1	Reserved.
25	RW	0	DMA Clock Control (DMACKCTL) . Controls the DMA clock source in none GbE mode (10/100 and no Link). In GbE mode the DMA clock source is always Kumeran PLL divided by 2. In nominal operation this bit should be in the default state in which the DMA clock source in none GbE is mosc_clk. In test mode the DMACKCTL and PLLGateDis should be set to 1 and CLK_CNT_1_4 in the NVM should not be set. In this mode the DMA clock source is Kumeran PLL divided by 2.
26	RW	0	Disable Static Kumeran PLL Gating (PLLGateDis) . By default the PLL is functional only when Kumeran link is required, and inactive when it is not required (at non GbE mode if Jordan is available). When set to 1 the Kumeran PLL is always active.
27	RW	0	IAME . When the IAM (interrupt acknowledge auto-mask enable) bit is set, a read or write to the ICR register will have the side effect of writing the value in the IAM register to the IMC register. When this bit is 0, the feature is disabled.
28	RW	0	Driver loaded (DRV_LOAD) . This bit should be set by the driver after it was loaded, Cleared when the driver unloads or soft reset. The MNG controller loads this bit to indicate to the manageability controller that the driver has loaded.
29	RW	0	INT_TIMERS_CLEAR_ENA . When set this bit enables the clear of the interrupt timers following an IMS clear. In this state, successive interrupts will occur only after the timers will expire again. When clear, successive interrupts following IMS clear may happen immediately.
30	RO	0	Reserved.
31	RO	0	Reserved. Reads as 0.

This register provides extended control of device functionality beyond that provided by the Device Control register (CTRL).

Note:

If software uses the EE_RST function and desires to retain current configuration information, the contents of the control registers should be read and stored by



software. Control register values are changed by a read of the NVM which occurs upon assertion of the EE_RST bit.

Note: The EEPROM reset function may read configuration information out of the NVM which affects the configuration of PCI configuration space BAR settings. The changes to the BAR's are not visible unless the system is rebooted and the BIOS is allowed to re-map them.

Note: The SPD_BYPS bit performs a similar function as the CTRL.FRCSPD bit in that the device's speed settings are determined by the value software writes to the CTRL.SPEED bits. However, with the SPD_BYPS bit asserted, the settings in CTRL.SPEED take effect immediately rather than waiting until after the device's clock switching circuitry performs the change.

12.0.3.2.5 MDI Control Register - MDIC (0x00020; RW)


Bits	Type	Default	Description
15:0	RW/V	X	Data (DATA). In a Write command, software places the data bits and the MAC shifts them out to the LAN Connected Device. In a Read command, the MAC reads these bits serially from the LAN Connected Device and software can read them from this location.
20:16	RW/V	0	LAN Connected Device Register address (REGADD). i.e., Reg 0, 1, 2, ... 31.
25:21	RW/V	0	LAN Connected Device Address (PHYADD).
27:26	RW/V	0	Op-code (OP). 01 for MDI Write 10 for MDI Read. Other values are reserved.
28	RW/V	1	Ready bit (R). Set to 1 by LAN Controller at the end of the MDI transaction (i.e., indicates a Read or Write has been completed). It should be reset to 0 by software at the same time the command is written.
29	RW/V	0	Interrupt Enable (I). When set to 1 by software, it will cause an Interrupt to be asserted to indicate the end of an MDI cycle.
30	RW/V	0	Error (E). This bit set is to 1 by the HW when it fails to complete an MDI read. Software should make sure this bit is clear (0) before making a MDI read or write command.
31	RO	0	Reserved. Write as 0 for future compatibility.

This register is used by software to read or write MDI (Management Data Interface) registers in a GMII/MII LAN Connected Device.

Note: Internal logic uses the MDIC to communicate w the LCD. All fields in these registers are indicated as "/V" since the internal logic may use them to access the LCD. Since the HW uses this register, all HW, SW and FW must use semaphore logic (the Ownership flags) before accessing the MDIC.

For an MDI Read cycle the sequence of events is as follows:

1. the CPU performs a write cycle to the MII register with:
 - Ready = 0
 - Interrupt Enable bit set to 1 or 0.



- Op-Code = 10b (read)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = the register address of the specific register to be accessed (0 through 31)
2. the MAC applies the following sequence on the MDIO signal to the LAN Connected Device:
 - <PREAMBLE><01><10><PHYADD><REGADD><Z>

where the Z stands for the MAC tri-stating the MDIO signal.

3. the LAN Connected Device returns the following sequence on the MDIO signal:
 - <0><DATA><IDLE>
4. the MAC discards the leading bit and places the following 16 data bits in the MII register.
5. LAN Controller asserts an Interrupt indicating MDI "Done", if the Interrupt Enable bit was set.
6. LAN Controller sets the Ready bit in the MII register indicating the Read is complete.
7. the CPU may read the data from the MII register and issue a new MDI command.

For an MDI Write cycle the sequence of events is as follows:

1. the CPU performs a write cycle to the MII register with:
 - Ready = 0
 - Interrupt Enable bit set to 1 or 0.
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = the register address of the specific register to be accessed (0 through 31)
 - Data = specific Data for Desired Control of LAN Connected Device
2. the MAC applies the following sequence on the MDIO signal to the LAN Connected Device:
 - <PREAMBLE><01><01><PHYADD><REGADD><10><DATA><IDLE>
3. the LAN Controller asserts an Interrupt indicating MDI "Done" if the Interrupt Enable bit was set.
4. the LAN Controller sets the Ready bit in the MII register to indicate Step 2. has been completed.
5. the CPU may issue a new MDI command.

Note: An MDI Read or Write may take as long as 64 microseconds from the CPU Write to the Ready bit assertion.

If an invalid opcode is written by software, the MAC will not execute any accesses to the LAN Connected Device registers.

If the LAN Connected Device does not generate a zero as the second bit of the turnaround cycle for reads, the MAC will abort the access, set the E (error) bit, write 0xFFFF to the data field to indicate an error condition, and set the ready bit.



Accessing LCD Wakeup register using MDIC

A new page is defined in the LAN Connected Device to hold the wakeup register space.

When SW wants to configure the wakeup state (either read or write to these registers) the MDIO page should be set to 800 (for host accesses). While the page remains the same Wakeup register access is enabled.

While the page is set to the Wakeup page the address field is no longer translated as reg_addr (register address) but as an instruction. If the given address is in [0..15] range meaning PHY registers, the functionality remains unchanged.

There are 2 valid instructions:

Address Set – 0x11 – Wakeup space address is set for either reading or writing.

Data cycle – 0x12 – Wakeup space accesses read or write cycle.

For the LAN Connected Device, in the wake area Read cycle, the sequence of events is as follows:

Setting page 0x0800

1. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = page setting
 - DATA = 0x0800 (Wakeup page)

Address setting

2. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 01b (write)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x11 (Address set)
 - DATA = XXXX (Address of the register to be read)


Reading a register

1. the Driver performs a write cycle to the MDI register with:
 - Ready = 0
 - Op-Code = 10b (read)
 - PHYADD = the LAN Connected Device address from the MDI register
 - REGADD = 0x12 (data cycle for read)
 - DATA = YYYY (Data will be valid when the ready bit is set)

For the LAN Connected Device, in the wake area Write cycle, the sequence of events is as follows:

Setting page 0x0800

1. the Driver performs a write cycle to the MDI register with:



- Ready = 0
- Op-Code = 01b (write)
- PHYADD = the LAN Connected Device address from the MDI register
- REGADD = page setting
- DATA = 0x0800 (Wakeup page)

Address setting

2. the Driver performs a write cycle to the MDI register with:

- Ready = 0
- Op-Code = 01b (write)
- PHYADD = the LAN Connected Device address from the MDI register
- REGADD = 0x11 (Address set)
- DATA = XXXX (Address of the register to be read)

Writing a register


3. the Driver performs a write cycle to the MDI register with:

- Ready = 0
- Op-Code = 01b (write)
- PHYADD = the LAN Connected Device address from the MDI register
- REGADD = 0x12 (data cycle for write)
- DATA = YYYY (Data to be written to the register)


12.0.3.2.6 Future Extended NVM Register - FEXTNVM (0x00028; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 19h and bits 31:16 are loaded from the NVM word 1Ah.

Bits	Type	Default	Description
0	RW/SN	0	Reserved
1	RW/SN	0	dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
2	RW/SN	0	wake_dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
3	RW/SN	0	gpt_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
4	RW/SN	0	mac_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).
5	RW/SN	0	m2k_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1, clk is always ticking, default value is '0' (HW and NVM).



Bits	Type	Default	Description
6	RW/SN	0	Invalid image CSUM. When cleared this bit indicates to the LAD NVM programming tools (eeupdate..) that the Image CSUM needs to be corrected. When set the CSUM is assumed to be correct.
7	RW/SN	0	Reserved.
8	RW/SN	0	Reserved.
9	RW/SN	0	Reserved.
10	RW/SN	0	Enable MDIO Watchdog Timer (MDIOWatchEna). When set to '0', the 100msec MDIO watchdog timer is enabled.
11	RW/SN	0	Reserved.
12	RW/SN	0	Reserved.
13	RW/SN	0	Reserved.
14	RW/SN	0	Reserved.
15	RW/SN	0	Reserved.
16	RW/SN	0	Reserved.
17	RW/SN	0	Reserved.
18	RW/SN	0	Reserved.
19	RW/SN	0	Reserved.
20	RW/SN	0	Disable CLK gate Enable Due to D3hot. When set it disables assertion of bb_clkgaten due to D3hot. Default NVM setting is '0'.
21	RW/SN	0	LAN Disable Mode. When set to '1', legacy flow managed by BIOS routine should be performed to disable the Lan. Otherwise, the whole flow will be managed by hardware when Lan-Disable RTC well bit is set to '1'in ICH9. Default NVM setting is '0'.
22	RW/SN	0	Reserved.
23	RW/SN	0	Reserved.
24	RW/SN	0	Reserved.
25	RW/SN	0	Reserved.
26	RW/SN	0	Reserved.
27	RW/SN	0	SW LCD Config Enable. This bit has no impact on the hardware but rather influences the software flow. The software should initialize the LCD using the "Extended Configuration" image in the NVM only when both the "SW LCD Config Enable" bit is set and the "LCD Write Enable" bit in the EXTCNF_CTRL register is cleared.
28	RW/SN	0	Reserved.
29	RW/SN	0	Reserved.
30	RW/SN	0	Enables assertion of "gbe_pmcmsus_powerdown_rdy_mosc" in WoL enabled configuration.
31	RW/SN	0	Enables assertion of "gbe_pmcmsus_powerdown_rdy_mosc" in D3/DMoff/no_wake configuration.



12.0.3.2.7 Future Extended Register - FEXT (0x0002C; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values.

Bits	Type	Default	Description
0	RO	0	Reserved.
1	RO	0	Reserved.
2	RO/V	0	Reserved.
3	RO/V	0	LAN PHY Power Control (PHYPC). Set to 1 indicates external power to the LAN Controller is On. 0 – external power is off.
4	RW	0	Reserved.
5	RW	0	Reserved.
6	RW	0	Reserved.
7	RW	0	Reserved.
8	RW	0	HW SW CRC mismatch trigger – when set to 1 the LC generates a trigger signal whenever there is a mismatch between the software calculated CRC and the hardware calculated CRC. This feature is ignored when CRC calculation is off-loaded to HW.
9	RW	0	Write disable Ghost and DMA RAMs on CRC mismatch – when set to 1: disable any writes to the following rams in the event of CRC mismatch until reset: Ghost read pci descriptor Ghost read pci data the four rams in the descriptor engine the packet buffer
10	RW	0	When set to 1: Enables the data visibility of the Ghost read PCI descriptor and PCI data RAMs to the NOA.
11	RW	0	Visibility in/out read data select (1 = in) FEXT.10 must be set to 1.
12	RW	0	Visibility data/desc read ram select (1 = data) FEXT.10 must be set to 1
13	RW	0	When set to 1: The Ghost read RAMs are readable by the slave bus.
14	RW	0	Reserved.
15	RW	0	Reserved.
16	RW	0	Reserved.
17	RW	0	Reserved.
31:18	RW	0x00	Future Extended. Reserved for future setting.

12.0.3.2.8 Future Extended NVM 2 - FEXTNVM2 (0x00030; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 24h and bits 31:16 are loaded from the NVM word 25h.



Bits	Type	Default	Description
31:0	RW/SN	0	Reserved.

12.0.3.2.9 Future Extended NVM 3 - FEXTNVM3 (0x0003C; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 26h and bits 31:16 are loaded from the NVM word 27h.

Bits	Type	Default	Description
31:0	RW/SN	0	Reserved.


12.0.3.2.10 Future Extended NVM 4- FEXTNVM4 (0x00024; RW)

This register is initialized to the HW default only at LAN PWR Good reset. Software should not modify these fields to values other than their recommended values. Bits 15:0 of this register are loaded from the NVM word 2Ah and bits 31:16 are loaded from the NVM word 2Bh.

Bits	Type	Default	Description
2:0	RW/SN	000	Reserved.
3	RW/SN	0	Reserved.
5:4	RW/SN	00	Reserved.
6	RW/SN	0	Reserved.
7	RW/SN	0	Reserved.
14:8	RW/SN	0x0	Reserved.
15	RW/SN	0	Reserved.
23:16	RW/SN	0x0	Reserved.
25:24	RW/SN	00	Reserved.
26	RW/SN	0	Reserved.
27	RW/SN	0	Reserved.
28	RW/SN	0	Reserved.
29	RW/SN	0	Enable PLL shut on 1000Mbps link up. When set to 0 and 1000Mbps link is up the LAN Controller will not approve PLL shut in K1, when set to 1 PLL shut on K1 will not be gated in 1000Mbps speed. This bit is loaded from NVM word 2Bh bit 13
31:30	RW/SN	000	Reserved.

12.0.3.2.11 Device and Bus Number - BUSNUM (0x00038; RO)

Bit	Type	Default	Description
28:0	RO	0x00	Reserved.



10:8	RO	000b	Function Number. The LAN controller is a single PCI function being function 0.
15:11	RO	0x19	Device Number. During nominal operation the LAN controller has a predefined Device number equal to 25 (0x19).
23:16	RO	0x00	Bus Number. The LAN controller captures its bus number during host configuration write cycles type 0 aimed at the device. This field is initialized by LAN power good reset, PCI reset and D3 to D0 transition.
31:24	RO	0x00	Reserved.

12.0.3.2.12 Flow Control Transmit Timer Value - FCTTV (0x00170; RW)

Bit	Type	Default	Description
15:0	RW	X	Transmit Timer Value (TTV). to be included in XOFF frame.
31:16	RO	0	Reserved. Read as 0. Should be written to 0 for future compatibility.

The 16-bit value in the TTV field is inserted into a transmitted frame (either XOFF frames or any PAUSE frame value in any software transmitted packets). It counts in units of slot time. If software wishes to send an XON frame, it must set TTV to 0 prior to initiating the PAUSE frame.

Note: The LAN Controller uses a fixed slot time value of 64 byte times.

12.0.3.2.13 Flow Control Refresh Threshold Value - FCRTV (0x05F40; RW)

Bit	Type	Default	Description
15:0	RW	X	Flow Control Refresh Threshold (FCRT). This value indicates the threshold value of the flow control shadow counter. When the counter reaches this value, and the conditions for a pause state are still valid (buffer fullness above low threshold value), a pause (XOFF) frame is sent to the link partner. The FCRTV timer count interval is the same as other flow control timers and counts at slot times of 64 byte times. If this field contains a zero value, the Flow Control Refresh is disabled.
31:16	RO	0	Reserved.

12.0.3.2.14 Extended Configuration Control - EXTCNF_CTRL (0x00F00; RW)

Bit	Type	Default	Description
0	RW/SN	0	LCD Write Enable. When set, enables the Extended LAN Connected Device Configuration area in the LAN Controller. When disabled, the Extended LAN Connected Device Configuration area is ignored. Loaded from NVM word 14h.
1	RW/SN	0	Reserved.
2	RW/SN	0	Reserved.




3	RW/SN	1	OEM Write Enable. When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. Loaded from NVM word 14h.
4	RO	0	Reserved.
5	RW/V	0	SW Semaphore FLAG (SWFLAG). This bit is set by the device driver to gain access permission to shared CSR registers with the firmware and hardware The bit is initialized on power-up PCI reset and software reset.
6	RO/V	0	MDIO HW Ownership. HW requests access to MDIO. Part of the arbitration scheme for MDIO access. This is a RO bit.
7	RW/V	0	Gate Phy Configuration (PPW, SKU read, OEM configuration).
15:8	RO	0x0	Reserved.
27:16	RW/SN	0x001	Extended Configuration Pointer. Defines the base address (in DW) of the Extended Configuration area in the NVM.
31:28	RW	0	Reserved.

12.0.3.2.15 Extended Configuration Size - EXTCNF_SIZE (0x00F08; RW)

Bit	Type	Default	Description
31:24	RO	0x0	Reserved.
23:16	RW/SN	0x0	Extended LCD Length. Size (in DW) of the Extended LAN Connected Device Configuration area loaded from Extended Configuration word 2 in the NVM. If an extended configuration area is disabled by "LCD Write Enable" field in word 14h in the NVM, this length must be set to zero.
15:0	RW/SN	0x0	Reserved.

12.0.3.2.16 +-PHY Control Register - PHY_CTRL (0x00F10; RW) This register is initialized to the HW default at LAN PWR Good reset.

Bit	Type	Default	Description
31:29	RO	0x0	Reserved
28:25	RO	0x0	SKU read data. These four bits contain the SKU value read from the LANConnected Device SKU register. Using these bits, the SKU mechanism determines the Device ID.
24	RO	0x0	Reserved.
23	RO	0x0	SKU done. This bit indicates the termination of SKU read.
22	RW	0x0	Reserved.
21	RW	0	Reserved
20	RW	0	Reserved.
19:17	RW	0x2	Reserved.
16	RW	0	Reserved.
15:8	RO	0x0	Reserved
7	RW/SN	0	B2B Ena. Enables SPD in Back To Back link setup. Bit is initialized by word 17h bit 15 in the NVM.



Bit	Type	Default	Description
6	RW/SN	0	Global GbE Disable. Prevents PHY auto negotiating 1000 Mb/s link in all power states (including D0a). Bit is initialized by word 17h bit 14 in the NVM.
5:4	RO	00b	Reserved.
3	RW/SN	1	GbE Disable at non D0a. Prevents PHY from auto negotiating 1000Mb/s link in all power states except D0a (DR, D0u and D3). Bit is initialized by word 17h bit 11 in the NVM. This bit must be set since GbE is not supported in Sx by the platform.
2	RW/SN	1	LPLU in non D0a. Enables PHY to negotiate for slowest possible link (Reverse AN) in all power states except D0a (DR, D0u and D3). Bit is initialized by word 17h bit 10 in the NVM.
1	RW/SN	0	LPLU in D0a. Enables PHY to negotiate for slowest possible link (Reverse AN) in all power states (including D0a). This bit overrides the LPLU in non D0a bit. Bit is initialized by word 17h bit 9 in the NVM.
0	RW/SN	0	SPD Ena. Enables PHY Smart Power Down mode. Bit is initialized by word 17h bit 8 in the NVM.

12.0.3.2.17 PCIE Analog Configuration - PCIEANACFG (0x00F18; RW)

Bit	Type	Default	Description
31:6	RO	0	Reserved.
5:0	RW/SN	0x0	Reserved.

12.0.3.2.18 Packet Buffer Allocation - PBA (0x01000; RW)

Bit	Type	Reset	Description
4:0	RW	0x12	Receive packet buffer allocation (RXA). Defines the size of the Rx buffer in K byte units. Default is 18Kbytes.
15:5	RO	X	Reserved.
20:16	RO	0x	Transmit packet buffer allocation (TXA). Defines the size of the Tx buffer in K byte units. This field is read only and equals to the Packet Buffer Size (PBS) minus RXA (the default value of the PBS is KB).
31:21	RO	X	Reserved.

This register sets the on-chip receive and transmit storage allocation ratio.

Programming this register does not automatically re-load or initialize internal packet-buffer RAM pointers. The software must reset both transmit and receive operation (using the global device reset CTRL.SWRST bit) after changing this register in order for it to take effect. The PBA register itself will not be reset by assertion of the Software reset, but will only be reset upon initial hardware power-on.

If Early Receive functionality is not enabled (indicate field/register), the Receive packet buffer should be larger than the max expected received packet + 32B.



For best performance the transmit buffer allocation should be set to accept two full sized packets.

Transmit Packet Buffer size should be configured to be more than 4 KB.

12.0.3.2.19 Packet Buffer Size - PBS (0x01008; RW)

Bit	Type	Reset	Description
4:0	RW	0x12	Receive packet buffer allocation (RXA) . Defines the size of the Rx buffer in K byte units. Default is 18Kbytes.
15:5	RO	X	Reserved.
20:16	RO	0x	Transmit packet buffer allocation (TXA) . Defines the size of the Tx buffer in K byte units. This field is read only and equals to the Packet Buffer Size (PBS) minus RXA (the default value of the PBS is KB).
31:21	RO	X	Reserved.


This register sets the on-chip receive and transmit storage allocation size, The allocation value is read/write for the lower 6 bits. The division between transmit and receive is done according to the PBA register.

Note: Programming this register does not automatically re-load or initialize internal packet-buffer RAM pointers. The software must reset both transmit and receive operation (using the global device reset CTRL.SWRST bit) after changing this register in order for it to take effect. The PBS register itself will not be reset by assertion of the Software reset, but will only be reset upon initial hardware power-on.

Note: Programming this register should be aligned with programming the PBA register HW operation, if PBA and PBS are not coordinated is not determined.

12.0.3.2.20 Packet Buffer ECC Status - PBECCSTS (0x0100C; RW)

Bit	Type	Reset	Description
7:0	RC	0x0	Corr_err_cnt - Correctable Error Count: This counter is increment every time a correctable error is detected, the counter stops counting when reaching 0xff. Cleared by read.
15:8	RC	0x0	uncorr_err_cnt - Uncorrectable Error Count: This counter is increment every time an uncorrectable error is detected, the counter stops counting when reaching 0xff. Cleared by read.
16	RW	0	ECC enable
17	RW	0	Stop on first Error (SOFE). When set the ECC test will capture the failing address into LFA
19:18	RO	0x0	Reserved. read as zero.
31:20	RO	0x0	Last failure address (LFA). When Stop on first Error (SOFE) bit is set to 1, when there is ECC failure, the LFA register will capture the failing address of the failure




12.0.3.2.21 Packet Buffer ECC Error Inject - PBEEI (0x01004; RW)

Bit	Type	Reset	Description
0	RW	0	Inject an error on TX Buffer on header line When this bit is set an error will be injected in the next write cycle to a header line of the TX buffer. Auto cleared by HW when error is injected if PBECCINJ.ENECCADD is clear (0).
1	RW	0	Inject an error on TX Buffer on data line When this bit is set an error will be injected in the next write cycle to a data line of the TX buffer. Auto cleared by HW when error is injected if PBECCINJ.ENECCADD is clear (0).
2	RW	0	Inject an error on RX Buffer on header line When this bit is set an error will be injected in the next write cycle to a header line of the RX buffer. Auto cleared by HW when error is injected if PBECCINJ.ENECCADD is clear (0).
3	RW	0	Inject an error on RX Buffer on data line When this bit is set an error will be injected in the next write cycle to a data line of the RX buffer. Auto cleared by HW when error is injected if PBECCINJ.ENECCADD is clear (0).
15:4	RO	0x0	Reserved
23:16	RW	0	Error 1 bit location (value of 0xFF - No Error injection on this bit)
31:24	RW	0x0	Error 2 bit location (value of 0xFF - No Error injection on this bit)

12.0.3.2.22 Packet Buffer ECC Injection - PBECCINJ (0x01010; RW)

Bit	Type	Reset	Description
11:0	RW	0	Address 0 Injection - Error injection first address in packet buffer.
23:12	RW	0	Address 1 Injection - Error injection second address in packet buffer.
24	RW	0	Enable ECC injection to address (ENACCADD). When set to 0 the addresses for ECC injection from this register will be ignored
31:25	RO	0x0	Reserved




12.0.3.3 Interrupt Register Descriptions

12.0.3.3.1 Interrupt Cause Read Register - ICR (0x000C0; RC / WC)

This register is Read-Clear or Write-Clear (see details after the tables). If enabled, read access also clears the ICR content after it is posted to the SW. Otherwise a Write cycle is required to clear the relevant bit fields. Write a '1' clears the written bit while writing '0' has no affect (with the exception of the INT_ASSERTED bit as detailed below).

Bit	Type	Default	Description
0	RWC/CR/V	0	Transmit Descriptor Written Back (TXDW). Set when hardware processes a descriptor with either RS set. If using delayed interrupts (IDE set), the interrupt is delayed until after one of the delayed-timers (TIDV or TADV) expires.
1	RWC/CR/V	0	Transmit Queue Empty (TXQE). Set when, the last descriptor block for a transmit queue has been used. When configured to use more than one transmit queue this interrupt indication will be issued if one of the queues is empty and will not be cleared until all the queues have valid descriptors.
2	RWC/CR/V	0	Link Status Change (LSC). This bit is set whenever the link status changes (either from up to down, or from down to up). This bit is affected by the LINK indication from the PHY.
3	RO	0	Reserved.
4	RWC/CR/V	0	Receive Descriptor Minimum Threshold hit (RXDMT0). Indicates that the minimum number of receive descriptors RCTL.RDMTS are available and software should load more receive descriptors.
5	RWC/CR/V	0	Disable SW Write Access (DSW). The DSW bit indicates that the FW changed the status of the DISSW or the DISSWLNK bits in the FWSM register.
6	RWC/CR/V	0	Receiver Overrun (RXO). Set on receive data FIFO overrun. Could be caused either because there are no available buffers or because receive bandwidth is inadequate.
7	RWC/CR/V	0	Receiver Timer Interrupt (RXT0). Set when the timer expires.
8	RWC/CR/V	0	Reserved.
9	RWC/CR/V	0	MDIO Access Complete (MDAC). Set when the MDIO access is completed.
11:10	RO	0	Reserved.
12	RWC/CR/V	0	PHY Interrupt (PHYINT). Set when the LAN Connected Device generates an interrupt.
13	RO	0	Reserved.
14	RWC/CR/V	0	MACsec Packet Number (LSECPN). The Tx Packet Number hit the "PN exhaustion threshold" as defined in the LSECTXCTRL register and the host is the KaY.
15	RWC/CR/V	0	Transmit Descriptor Low Threshold hit (TXD_LOW). Indicates that the descriptor ring has reached the threshold specified in the Transmit Descriptor Control register.
16	RWC/CR/V	0	Small Receive Packet Detected (SRPD). Indicates that a packet of size < RSRPD.SIZE register has been detected and transferred to host memory. The interrupt is only asserted if RSRPD.SIZE register has a non-zero value.



Bit	Type	Default	Description
17	RWC/CR/V	0	Receive ACK Frame Detected (ACK). Indicates that an ACK frame has been received and the timer in RAID.ACK_DELAY has expired.
18	RWC/CR/V	0	Reserved.
19	RWC/CR/V	0	Reserved.
20	RWC/CR/V	0	Reserved.
21	RWC/CR/V	0	Reserved.
22	RWC/CR/V	0	ECC Error (ECCER). Indicates an uncorrectable EEC error had occurred
30:23	RO	0	Reserved. Reads as 0.
31	RWC/CR/V	0	Interrupt Asserted (INT_ASSERTED). This bit is set when the LAN port has a pending interrupt. If the Interrupt is enabled in the PCI configuration space an Interrupt is asserted.

This register contains all interrupt conditions for LAN Controller. Whenever an interrupt causing event occurs, the corresponding interrupt bit is set in this register. An interrupt is generated whenever one of the bits in this register is set, and the corresponding interrupt is enabled via the Interrupt Mask Set/Read Register (see [Section 12.0.3.4.5](#)).

Whenever an interrupt causing event occurs, all timers of delayed interrupts are cleared and their cause event is set in the ICR.

- Read ICR register is affected differently in the following cases:
 - Case 1 - Interrupt Mask Register equals 0x0000 (mask all) - ICR content will be cleared.
 - Case 2 - Interrupt was asserted (ICR.INT_ASSERTED=1) - ICR content will be cleared and Auto Mask is active, meaning, the IAM register is written to the IMC register.
 - Case 3 - Interrupt was not asserted (ICR.INT_ASSERTED=0) - Read has no side affect.

Writing a 1 to any bit in the register will also clear that bit. Writing a 0 to any bit will have no effect on that bit. The INT_ASSERTED bit is a special case. Writing a 1 or 0 to this bit has no affect. It is cleared only when all interrupt sources are cleared.

12.0.3.3.2 Interrupt Throttling Register - ITR (0x000C4; RW)

Bit	Type	Default	Description
15:0	RW	0	INTERVAL. Minimum inter-interrupt interval. The interval is specified in 256ns units. Zero disables interrupt throttling logic.
31:16	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Software can use this register to pace (or even out) the delivery of interrupts to the host CPU. This register provides a guaranteed inter-interrupt delay between interrupts asserted by the network controller, regardless of network traffic conditions. To independently validate configuration settings, software can use the following algorithm to convert the inter-interrupt interval value to the common 'interrupts/sec' performance metric:

$$\text{Interrupts/sec} = (256 \times 10^{-9} \text{sec} \times \text{interval})^{-1}$$



For example, if the interval is programmed to 500d, the network controller guarantees the CPU will not be interrupted by the network controller for 128 microseconds from the last interrupt.

Inversely, inter-interrupt interval value can be calculated as:

$$\text{inter-interrupt interval} = (256 \times 10^{-9} \text{sec} \times \text{interrupts/sec})^{-1}$$


The optimal performance setting for this register is very system and configuration specific. An initial suggested range for the interval value is 65--5580 (28B - 15CC).

Note: When working at 10/100Mbps and running at ¼ clock the interval time is multiplied by four.

12.0.3.3.3 Interrupt Cause Set Register - ICS (0x000C8; WO)

Bit	Type	Default	Description
0	WO	X	TXDW. Sets Transmit Descriptor Written Back.
1	WO	X	TXQE. Sets Transmit Queue Empty.
2	WO	X	LSC. Sets Link Status Change.
3	RO	X	Reserved.
4	WO	X	RXDMT. Sets Receive Descriptor Minimum Threshold hit.
5	WO	X	DSW. Sets Block SW Write accesses.
6	WO	X	RXO. Sets Receiver Overrun. Set on receive data FIFO overrun.
7	WO	X	RXT. Sets Receiver Timer Interrupt.
8	WO	X	Reserved.
9	WO	X	MDAC. Sets MDIO Access Complete Interrupt.
11:10	RO	X	Reserved.
12	WO	X	PHYINT. Sets PHY Interrupt.
13	RO	X	Reserved.
14	WO	X	LSECPN. Sets MACsec Packet Number Interrupt.
15	WO	X	TXD_LOW. Transmit Descriptor Low Threshold Hit.
16	WO	X	SRPD. Small Receive Packet Detected and Transferred.
17	WO	X	ACK. Set Receive ACK frame detected.
18	WO	X	MNG. Set the Manageability Event Interrupt.
19	WO	X	Reserved.
20	WO	X	Reserved.
21	RO	X	Reserved.
22	WO	X	ECCER Set uncorrectable EEC error.
31:23	RO	X	Reserved. Should be written with 0 to ensure future compatibility.

Software uses this register to set an interrupt condition. Any bit written with a “1” sets the corresponding interrupt. This results in the corresponding bit being set in the Interrupt Cause Read Register (see [Section 12.0.3.4](#)), and an interrupt is generated if one of the bits in this register is set, and the corresponding interrupt is enabled via the Interrupt Mask Set/Read Register (see [Section 12.0.3.4.5](#)).



Bits written with “0” are unchanged.

12.0.3.3.4 Interrupt Mask Set/Read Register - IMS (0x000D0; RW)

Bit	Type	Default	Description
0	RWS	0	TXDW. Sets mask for Transmit Descriptor Written Back.
1	RWS	0	TXQE. Sets mask for Transmit Queue Empty.
2	RWS	0	LSC. Sets mask for Link Status Change.
3	RO	0	Reserved.
4	RWS	0	RXDMT0. Sets mask for Receive Descriptor Minimum Threshold hit.
5	RWS	0	DSW. Sets mask for Block SW Write accesses.
6	RWS	0	RXO. Sets mask for Receiver Overrun. Set on receive data FIFO overrun.
7	RWS	0	RXT0. Sets mask for Receiver Timer Interrupt.
8	RWS	0	Reserved.
9	RWS	0	MDAC. Sets mask for MDIO Access Complete Interrupt.
11:10	RO	0	Reserved.
12	RWS	0	PHYINT. Sets mask for PHY Interrupt.
13	RO	0	Reserved.
14	RWS	0	LSECPN. Sets the mask for MACsec Packet Number Int.
15	RWS	0	TXD_LOW. Sets the mask for Transmit Descriptor Low Threshold hit.
16	RWS	0	SRPD. Sets mask for Small Receive Packet Detection.
17	RWS	0	ACK. Sets the mask for Receive ACK frame detection.
18	RWS	0	MNG. Sets mask for Manageability Event Interrupt.
19	RWS	0	Reserved.
20	RWS	0	Reserved.
21	RO	0	Reserved.
22	RWS	0	ECCER Sets mask for uncorrectable EEC error
31:23	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Reading this register returns which bits have an interrupt mask set. An interrupt is enabled if its corresponding mask bit is set to 1, and disabled if its corresponding mask bit is set to 0. An interrupt is generated whenever one of the bits in this register is set, and the corresponding interrupt condition occurs. The occurrence of an interrupt condition is reflected by having a bit set in the Interrupt Cause Read Register (see [Section 12.0.3.4](#)).

A particular interrupt may be enabled by writing a 1 to the corresponding mask bit in this register. Any bits written with a 0 are unchanged.

Note: If software desires to disable a particular interrupt condition that had been previously enabled, it must write to the Interrupt Mask Clear Register (see [Section 12.0.3.4.6](#)), rather than writing a 0 to a bit in this register.



When the CTRL_EXT.INT_TIMERS_CLEAR_ENA bit is set, then following writing all 1's to the IMS register (enable all Interrupts) all interrupt timers are cleared to their initial value. This auto clear provides the required latency before the next INT event.


12.0.3.3.5 Interrupt Mask Clear Register - IMC (0x000D8; WO)

Bit	Type	Default	Description
0	WO	0	TXDW. Clears mask for Transmit Descriptor Written Back.
1	WO	0	TXQE. Clears mask for Transmit Queue Empty.
2	WO	0	LSC. Clears mask for Link Status Change.
3	RO	0	Reserved.
4	WO	0	RXDMT0. Clears mask for Receive Descriptor Minimum Threshold hit.
5	WO	0	DSW. Clears mask for Block SW Write accesses.
6	WO	0	RXO. Clears mask for Receiver Overrun.
7	WO	0	RXT0. Clears mask for Receiver Timer Interrupt.
8	WO	0	Reserved.
9	WO	0	MDAC. Clears mask for MDIO Access Complete Interrupt.
11:10	RO	0	Reserved. Reads as 0.
12	WO	0	PHYINT. Clears PHY Interrupt.
13	RO	0	Reserved.
14	WO	0	LSECPN. Clears the MACsec Packet Number Interrupt.
15	WO	0	TXD_LOW. Clears the mask for Transmit Descriptor Low Threshold hit.
16	WO	0	SRPD. Clears mask for Small Receive Packet Detect Interrupt.
17	WO	0	ACK. Clears the mask for Receive ACK frame detect Interrupt.
18	WO	0	MNG. Clears mask for the Manageability Event Interrupt.
19	WO	0	Reserved.
20	WO	0	Reserved.
21	RO	0	Reserved.
22	WO	0	ECCER Clears the mask for uncorrectable EEC error
31:23	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Software uses this register to disable an interrupt. Interrupts are presented to the bus interface only when the mask bit is a one and the cause bit is a one. The status of the mask bit is reflected in the Interrupt Mask Set/Read Register, and the status of the cause bit is reflected in the Interrupt Cause Read register (see [Section 12.0.3.4](#)).

Software blocks interrupts by clearing the corresponding mask bit. This is accomplished by writing a 1 to the corresponding bit in this register. Bits written with 0 are unchanged (i.e. their mask status does not change).

In summary, the sole purpose of this register is to allow software a way to disable certain, or all, interrupts. Software disables a given interrupt by writing a 1 to the corresponding bit in this register.




12.0.3.3.6 Interrupt Acknowledge Auto-Mask - IAM (0x000E0; RW)

Bit	Type	Default	Description
0-31	RW	0	IAM_VALUE. When the CTRL_EXT.IAME bit is set and the ICR.INT_ASSERTED=1, an ICR read or write will have the side effect of writing the contents of this register to the IMC register.


12.0.3.4 Receive Register Descriptions

12.0.3.4.1 Receive Control Register - RCTL (0x00100; RW)

Bit	Type	Default	Description
0	RO	0	Reserved. This bit represented a hardware reset of the receive-related portion of the device in previous controllers, but is no longer applicable. Only a full device reset CTRL.SWRST is supported. Write as 0 for future compatibility.
1	RW	0	Enable (EN). The receiver is enabled when this bit is 1. Writing this bit to 0, stops reception after receipt of any in progress packets. All subsequent packets are then immediately dropped until this bit is set to 1. Note that this bit controls only DMA functionality to the host. Packets are counted by the statistics even when this bit is cleared.
2	RW	0	Store bad packets (SBP). 0 – do not store bad packets 1 – store bad packets Note that CRC errors before the SFD are ignored. Any packet must have a valid SFD in order to be recognized by the device (even bad packets). Note: Packet errors will not be routed to the MNG even if this bit is set.
3	RW	0	Unicast promiscuous enable (UPE). 0 – disabled 1 – enabled
4	RW	0	Multicast promiscuous enable (MPE). 0 – disabled 1 – enabled
5	RW	0	Long packet enable (LPE). 0 – disabled 1 – enabled
7:6	RW	00b	Reserved.
9:8	RW	0	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set whenever the fractional number of free descriptors becomes equal to RDMTS. Table 12-8 below lists which fractional values correspond to RDMTS values. See Section 12.0.3.5.8 for details regarding RDLEN.
11:10	RW	0	Descriptor Type (DTYP). 00 – Legacy or Extended descriptor type 01 – Packet Split descriptor type 10 and 11 – Reserved



Bit	Type	Default	Description
13:12	RW	0	Multicast Offset (MO). This determines which bits of the incoming multicast address are used in looking up the bit vector. 00 – [47:38] 01 – [46:37] 10 – [45:36] 11 – [43:34]
14	RW	0	Reserved.
15	RW	0	Broadcast Accept Mode (BAM). 0 – ignore broadcast (unless it matches through exact or imperfect filters) 1 – accept broadcast packets.
17:16	RW	0	Receive Buffer Size (BSIZE). RCTL.BSEX – 0 00 – 2048 Bytes 01 – 1024 Bytes 10 – 512 Bytes 11 – 256 Bytes RCTL.BSEX – 1 00 – reserved 01 – 16384 Bytes 10 – 8192 Bytes 11 – 4096 Bytes BSIZE is only used when DTYP – 00. When DTYP – 01, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than 0, in that case, FLXBUF determines the buffer size.
21:18	RO	0	Reserved. Should be written with 0.
22	RW	0	Reserved.
23	RW	0	Pass MAC Control Frames (PMCF). 0 – do not (specially) pass MAC control frames. 1 – pass any MAC control frame (type field value of 0x8808) that does not contain the pause opcode of 0x0001.
24	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
25	RW	0	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE above). 0 – Buffer size is as defined in BSIZE 1 – Original BSIZE values are multiplied by 16.
26	RW	0	Strip Ethernet CRC from incoming packet (SECRC). 0 – does not strip CRC 1 – Strips CRC The stripped CRC is not DMA'd to host memory and is not included in the length reported in the descriptor
30:27	RW	0	FLXBUF. Determines a flexible buffer size. When this field is "0000", the buffer size is determined by BSIZE. If this field is different from "0000", the receive buffer size is the number represented in Kbytes: i.e. "0001" = 1KB (1024 Bytes).



Bit	Type	Default	Description
31	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

LPE controls whether long packet reception is permitted. Hardware discards long packets if LPE is 0. A long packet is one longer than 1522 bytes. If LPE is 1, the maximum packet size that the device can receive is 9018 bytes.

RDMTS{1,0} determines the threshold value for free receive descriptors according to the following table:

Table 12-8. RDMTS Values

RDMTS	Free Buffer Threshold
00	1/2
01	1/4
10	1/8
11	Reserved


BSIZE controls the size of the receive buffers and permits software to trade-off descriptor performance versus required storage space. Buffers that are 2048 bytes require only one descriptor per receive packet maximizing descriptor efficiency. Buffers that are 256 bytes maximize memory efficiency at a cost of multiple descriptors for packets longer than 256 bytes.

PMCF controls the DMA function of MAC control frames (other than flow control). A MAC control frame in this context must be addressed to either the MAC control frame multicast address or the station address, match the type field and NOT match the PAUSE opcode of 0x0001. If PMCF = 1 then frames meeting this criteria will be DMA'd to host memory.

The SECRC bit controls whether the hardware strips the Ethernet CRC from the received packet. This stripping occurs prior to any checksum calculations. The stripped CRC is not DMA'd to host memory and is not included in the length reported in the descriptor.

12.0.3.4.2 Receive Control Register 1 - RCTL1 (0x00104; RW)

Bit	Type	Default	Description
7:0	RO	0	Reserved. This bit represented a hardware reset of the receive-related portion of the device in previous controllers, but is no longer applicable. Only a full device reset CTRL.SWRST is supported. Write as 0 for future compatibility.
9:8	RW	0	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set whenever the fractional number of free descriptors becomes equal to RDMTS. Table 12-8 lists which fractional values correspond to RDMTS values. See Section 12.0.3.5.8 for details regarding RDLEN.
11:10	RW	0	Descriptor Type (DTYP). 00 – Legacy or Extended descriptor type 01 – Packet Split descriptor type 10 and 11 – Reserved The value of RCTL1.DTYP should be the same as RCTL.DTYP (same descriptor types used in both descriptor queues).




15:12	RO	0	Reserved.
17:16	RW	0	Receive Buffer Size (BSIZE). RCTL.BSEX – 0 00 – 2048 Bytes 01 – 1024 Bytes 10 – 512 Bytes 11 – 256 Bytes RCTL.BSEX – 1 00 – reserved 01 – 16384 Bytes 10 – 8192 Bytes 11 – 4096 Bytes BSIZE is only used when DTYP – 00. When DTYP – 01, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than 0, in that case, FLXBUF determines the buffer size.
24:18	RO	0	Reserved. Should be written with 0.
25	RW	0	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE above). 0 – Buffer size is as defined in BSIZE 1 – Original BSIZE values are multiplied by 16.
26	RW	0	Reserved. Should be written with 0.
30:27	RW	0	FLXBUF. Determine a flexible buffer size. When this field is "0000", the buffer size is determined by BSIZE. If this field is different from "0000", the receive buffer size is the number represented in Kbytes: i.e. "0001" – 1KB (1024 Bytes).
31	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

This register is used to configure queue1 registers when working in VMDq mode.

12.0.3.4.3 Early Receive Threshold - ERT (0x02008; RW)

Bit	Type	Default	Description
12:0	RW	0x0000	Receive Threshold Value (RxThreshold). This threshold is in units of 8 bytes.
13	RW	0	Reserved. When SW activates the "Early Receive" mechanism (by setting the RxThreshold field to a non-zero value) it must set this bit as well.
14	RO	0	Reserved.
21:15	RW	0	Reserved.
31:22	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the RxThreshold value. This threshold determines how many bytes of a given packet should be in the LAN Controller's on-chip receive packet buffer before it attempts to begin transmission of the frame on the host bus. This register allows software to configure the "early receive" mode.



This field has a granularity of 8 bytes. So, if this field is written to 0x20, which corresponds to a threshold of 256 (decimal) bytes. If the size of a given packet is smaller than the threshold value, or if this register is set to 0, then LAN Controller will start the PCI transfer only after the entire packet is contained in LAN Controller's receive packet buffer. LAN Controller examines this register on a cycle by cycle basis to determine if there is enough data to start a transfer for the given frame over the PCI bus.

Once the device acquires the bus, it will attempt to DMA all of the data collected in the internal receive packet buffer so far.

The only negative affect of setting this value too low is that it will cause additional PCI bursts for the packet. In other words, this register allows software to trade-off latency versus bus utilization. Too high a value will effectively eliminate the early receive benefits (at least for short packets) and too low a value will deteriorate PCI bus performance due to a large number of small bursts for each packet. The RUTEC statistic counts certain cases where the ERT has been set too low, and thus provides software a feedback mechanism to better tune the value of the ERT.

It should also be noted that this register will have an effect only when the receive packet buffer is nearly empty (the only data in the packet buffer is from the packet that is currently on the wire).

Note: When Early receive is used in parallel to the Packet split feature the minimum value of the ERT register should be bigger than the header size to enable the actual packet split.

Note: Early Receive should be enabled only when working in a Jumbo frames enabled environment and the ERT.RxThreshold should be set to 0xFA so that only packets bigger than 2K bytes would trigger the early receive mechanism.

12.0.3.4.4 Packet Split Receive Control Register - PSRCTL (0x02170 + n*0x4 [n=0..1]; RW)

Bit	Type	Default	Description
6:0	RW	2	Receive Buffer Size for Buffer 0 (BSIZE0). The value is in 128 byte resolution. Value can be from 128 bytes to 16256 bytes (15.875 Kbytes). Default buffer size is 256B. SW should not program this field to a zero value.
7	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
13:8	RW	4	Receive Buffer Size for Buffer 1 (BSIZE1). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 4KB. SW should not program this field to a zero value.
15:14	RO	0	Reserved. Should be written with 0 to ensure future compatibility.
21:16	RW	4	Receive Buffer Size for Buffer 2 (BSIZE2). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 4KB. SW may program this field to any value.
23:22	RO	0	Reserved. Should be written with 0 to ensure future compatibility.



Bit	Type	Default	Description
29:24	RW	0	Receive Buffer Size for Buffer 3 (BSIZE3). The value is in 1K resolution. Value can be from 1K byte to 63K bytes. Default buffer size is 0KB. SW may program this field to any value.
31:30	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

Note: If SW sets a buffer size to zero, all buffers following that one must be set to zero as well. Pointers in the Receive descriptors to buffers with a zero size should be set to anything but NULL pointers.

12.0.3.4.5 Flow Control Receive Threshold Low - FCRTL (0x02160; RW)

Bit	Type	Default	Description
2:0	RO	0	Reserved. The underlying bits might not be implemented in all versions of the chip. Must be written with 0.
15:3	RW	0	Receive Threshold Low (RTL). FIFO low water mark for flow control transmission.
30:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
31	RW	0	XON Enable (XONE). 0 – disabled 1 – enabled.

This register contains the receive threshold used to determine when to send an XON packet. It counts in units of bytes. The lower 3 bits must be programmed to 0 (8 byte granularity). Software must set XONE to enable the transmission of XON frames. Whenever hardware crosses the receive high threshold (becoming more full), and then crosses the receive low threshold and XONE is enabled (= 1), hardware transmits an XON frame.

Note that flow control reception/transmission are negotiated capabilities by the Auto-Negotiation process. When the device is manually configured, flow control operation is determined by the RFCE and TFCE bits of the device control register.

This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x00168.

12.0.3.4.6 Flow Control Receive Threshold High - FCRTTH (0x02168; RW)

Bit	Type	Default	Description
2:0	RO	0	Reserved. The underlying bits might not be implemented in all versions of the chip. Must be written with 0.
15:3	RW	0	Receive Threshold High (RTH). FIFO high water mark for flow control transmission.
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the receive threshold used to determine when to send an XOFF packet. It counts in units of bytes. This value must be at least 8 bytes less than the maximum number of bytes allocated to the Receive Packet Buffer (PBA, RXA), and the



lower 3 bits must be programmed to 0 (8 byte granularity). Whenever the receive FIFO reaches the fullness indicated by RTH, hardware transmits a PAUSE frame if the transmission of flow control frames is enabled.

Note that flow control reception/transmission are negotiated capabilities by the Auto-Negotiation process. When the device is manually configured, flow control operation is determined by the RFCE and TFCE bits of the device control register.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x00160.

12.0.3.4.7 Receive Descriptor Base Address Low queue - RDBAL (0x02800 + n*0x100[n=0..1]; RW)

Bit	Type	Default	Description
3:0	RO	0	Reserved. Ignored on writes. Returns 0 on reads.
31:4	RW	X	Receive Descriptor Base Address Low (RDBAL).

This register contains the lower bits of the 64 bit descriptor base address. The lower 4 bits are always ignored. The Receive Descriptor Base Address must point to a 16B aligned block of data.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, RDBAL[0] may also be accessed at its alias offset of 0x00110.

12.0.3.4.8 Receive Descriptor Base Address High queue - RDBAH (0x02804 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
31:0	RW	X	Receive Descriptor Base Address [63:32] (RDBAH).

This register contains the upper 32 bits of the 64 bit Descriptor base address.


Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, RDBAH[0] may also be accessed at its alias offset of 0x00114.

12.0.3.4.9 Receive Descriptor Length queue- RDLEN (0x02808 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
6:0	RO	0	Reserved. Ignore on write. Reads back as 0.
19:7	RW	0	Descriptor Length (LEN)
31:20	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register sets the number of bytes allocated for descriptors in the circular descriptor buffer. It must be 128B aligned.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, RDLEN[0] may also be accessed at its alias offset of 0x00118.



Note: The descriptor ring must be equal to or larger than 8 descriptors.

12.0.3.4.10 Receive Descriptor Head queue - RDH (0x02810 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
15:0	RW/V	0	Receive Descriptor Head (RDH).
31:16	RO	0	Reserved. Should be written with 0

This register contains the head pointer for the receive descriptor buffer. The register points to a 16B datum. Hardware controls the pointer. The only time that software should write to this register is after a reset (hardware reset or CTRL.SWRST) and before enabling the receive function (RCTL.EN). If software were to write to this register while the receive function was enabled, the on-chip descriptor buffers may be invalidated and the hardware could become confused.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, RDH[0] may also be accessed at its alias offset of 0x00120.

12.0.3.4.11 Receive Descriptor Tail queue - RDT (0x02818 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
15:0	RW	0	Receive Descriptor Tail (RDT).
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register contains the tail pointer for the receive descriptor buffer. The register points to a 16B datum. Software writes the tail register to add receive descriptors for the hardware to process.


Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, RDT[0] may also be accessed at its alias offset of 0x00128.

12.0.3.4.12 Interrupt Delay Timer (Packet Timer) - RDTR (0x02820 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
15:0	RW	0	Receive Delay Timer. Receive Packet delay timer measured in increments of 1.024 microseconds.
30:16	RO	0	Reserved. Reads as 0.
31	WO	0	Flush Partial Descriptor Block (FPD). when set to 1; ignored otherwise. Reads 0.

This register is used to delay interrupt notification for the receive descriptor ring by coalescing interrupts for multiple received packets. Delaying interrupt notification helps maximize the number of receive packets serviced by a single interrupt.

This feature operates by initiating a countdown timer upon successfully receiving each packet to system memory. If a subsequent packet is received *before* the timer expires, the timer is re-initialized to the programmed value and re-starts its countdown. If the timer expires due to *not* having received a subsequent packet within the programmed interval, pending receive descriptor write backs are flushed and a receive timer interrupt is generated.



Setting the value to 0 represents no delay from a receive packet to the interrupt notification, and results in immediate interrupt notification for each received packet.

Writing this register with FPD set initiates an immediate expiration of the timer, causing a write back of any consumed receive descriptors pending write back, and results in a receive timer interrupt in the ICR.

Receive interrupts due to a Receive Absolute Timer (RADV) expiration will cancel a pending RDTR interrupt. The RDTR countdown timer is reloaded but halted, so as to avoid generation of a spurious second interrupt after the RADV has been noted, but may be restarted by a subsequent received packet.

Note: FPD is self-clearing.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x00108.

12.0.3.4.13 Receive Descriptor Control - RXDCTL (0x02828 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
14	RW	0	Reserved.
15	RW	0	Reserved.
21:16	RW	0x01	Write-back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0x0	Granularity (GRAN). Units for the thresholds in this register. 0 – cache lines 1 – descriptors
31:25	RO	0x00	Reserved.

Note: This register was not fully validated. SW should set it to 0x0000 during nominal operation.

This register controls the fetching and write-back of receive descriptors. The three threshold values are used to determine when descriptors will be read from and written to host memory. The values may be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag. If GRAN=0 (specifications are in cache-line granularity), the thresholds specified (based on the cache line size specified in the PCI configuration space CLS field) must not represent greater than 31 descriptors.

Note: When (WTHRESH = 0) or (WTHRESH = 1 and GRAN = 1) only descriptors with the 'RS' bit set will be written back.

PTHRESH is used to control when a prefetch of descriptors will be considered. This threshold refers to the number of valid, unprocessed receive descriptors the chip has in its on-chip buffer. If this number drops below PTHRESH, the algorithm will consider pre-fetching descriptors from host memory. This fetch will not happen however unless there are at least HTHRESH valid descriptors in host memory to fetch.

Note: HTHRESH should be given a non zero value when ever PTHRESH is used.



WTHRESH controls the write-back of processed receive descriptors. This threshold refers to the number of receive descriptors in the on-chip buffer which are ready to be written back to host memory. In the absence of external events (explicit flushes), the write-back will occur only after at least WTHRESH descriptors are available for write-back.

Note: Possible values:

GRAN = 1 (descriptor granularity):

PTHRESH = 0...31

WTHRESH = 0...31

HTHRESH = 0...31

GRAN = 0 (cache line granularity):

PTHRESH = 0...3 (for 16 descriptors cache line - 256 bytes)

WTHRESH = 0...3

HTHRESH = 0...4

Note: For any WTHRESH value other than 0 - The packet and absolute timers must get a non zero value for WTHRESH feature to take affect.

Note: Since the default value for write-back threshold is 1, the descriptors are normally written back as soon as one cache line is available. WTHRESH must contain a non-zero value to take advantage of the write-back bursting capabilities of LAN Controller.

Note: RXDCTL1 is only accessible when VMDq is enabled (MRQC.MRxQueue = 10).

12.0.3.4.14 Receive Interrupt Absolute Delay Timer- RADV (0x0282C; RW)


Bits	Type	Default	Description
15:0	RW	0	Receive Absolute Delay Timer. Receive Absolute delay timer measured in increments of 1.024 microseconds (0 = disabled)
31:16	RO	0	Reserved. Reads as 0

If the packet delay timer is used to coalesce receive interrupts, it will ensure that when receive traffic abates, an interrupt will be generated within a specified interval of no receives. During times when receive traffic is continuous, it may be necessary to ensure that no receive remains unnoticed for too long an interval. This register may be used to ensure that a receive interrupt occurs at some predefined interval after the first packet is received.

When this timer is enabled, a separate absolute countdown timer is initiated upon successfully receiving each packet to system memory. When this absolute timer expires, pending receive descriptor write backs are flushed and a receive timer interrupt is generated.

Setting this register to 0 disables the absolute timer mechanism (the RDTR register should be used with a value of 0 to cause immediate interrupts for all receive packets).

Receive interrupts due to a Receive Packet Timer (RDTR) expiration will cancel a pending RADV interrupt. If enabled, the RADV countdown timer is reloaded but halted, so as to avoid generation of a spurious second interrupt after the RDTR has been noted.



12.0.3.4.15 Receive Small Packet Detect Interrupt- RSRPD (0x02C00; RW)

Bits	Type	Default	Description
11:0	RW	0	SIZE. If the interrupt is enabled any receive packet of size <= SIZE will assert an Interrupt. SIZE is specified in bytes and includes the headers and the CRC. It does not include the VLAN header in size calculation if it is stripped.
31:12	RO	X	Reserved.

12.0.3.4.16 Receive ACK Interrupt Delay Register - RAID (0x02C08; RW)

Bits	Type	Default	Description
15:0	RW	0	ACK_DELAY. ACK delay timer measured in increments of 1.024 microseconds. When the Receive ACK frame detect Interrupt is enabled in the IMS register, ACK packets being received will use a unique delay timer to generate an interrupt. When an ACK is received, an absolute timer will load to the value of ACK_DELAY. The interrupt signal will be set only when the timer expires. If another ACK packet is received while the timer is counting down, the timer will not be reloaded to ACK_DELAY.
31:16	RO		Reserved.

If an immediate (non-scheduled) interrupt is desired for any received Ack frame, the ACK_DELAY should be set to 0.

12.0.3.4.17 CPU Vector Register - CPUVEC (0x02C10; RW)

Bits	Type	Default	Description
31:0	Res.	0x00	Reserved.

12.0.3.4.18 Receive Checksum Control - RXCSUM (0x05000; RW)

Bits	Type	Default	Description
7:0	RW	0x00	Packet Checksum Start (PCSS).
8	RW	1	IP Checksum Offload Enable (IPOFL).
9	RW	1	TCP/UDP Checksum Offload Enable (TUOFL).
11:10	RO	0	Reserved.
12	RW	0	IP Payload Checksum Enable (IPPCSE).
13	RW	0	Packet Checksum Disable (PCSD).
14	RW	0	Reserved.
15	RW	0	Reserved.
31:16	RO	0	Reserved.

The Receive Checksum Control register controls the receive checksum off loading features of LAN Controller. LAN Controller supports the off loading of three receive checksum calculations: the Packet Checksum, the IP Header Checksum, and the TCP/UDP Checksum.



PCSD: The Packet Checksum and IP Identification fields are mutually exclusive with the RSS hash. Only one of the two options is reported in the Rx descriptor. The RXCSUM.PCSD affect is shown in the table below:

RXCSUM.PCSD	0 (Checksum Enable)	1 (Checksum Disable)
Legacy Rx Descriptor (RCTL.DTYP = 00b)	Packet Checksum is reported in the Rx Descriptor	Forbidden Configuration
Extended or Header Split Rx Descriptor (RCTL.DTYP = 01b)	Packet Checksum and IP Identification are reported in the Rx Descriptor	RSS Hash value is reported in the Rx Descriptor

PCSS IPPCSE: The PCSS and the IPPCSE control the Packet Checksum calculation. As noted above, the packet checksum shares the same location as the RSS field. The Packet checksum is reported in the Receive descriptor when the RXCSUM.PCSD bit is cleared.

If RXCSUM.IPPCSE cleared (the default value), the checksum calculation that is reported in the Rx Packet checksum field is the unadjusted “16 bit ones complement” of the packet. The Packet Checksum starts from the byte indicated by RXCSUM.PCSS (0 corresponds to the first byte of the packet), after VLAN stripping if enabled (by CTRL.VME). For example, for an Ethernet II frame encapsulated as an 802.3ac VLAN packet and with RXCSUM.PCSS set to 14, the packet checksum would include the entire encapsulated frame, excluding the 14-byte Ethernet header (DA, SA, Type/Length) and the 4-byte VLAN tag. The Packet Checksum will not include the Ethernet CRC if the RCTL.SECRC bit is set. Software must make the required offsetting computation (to back out the bytes that should not have been included and to include the pseudo-header) prior to comparing the Packet Checksum against the TCP checksum stored in the packet.

If the RXCSUM.IPPCSE is set, the Packet checksum is aimed to accelerate checksum calculation of fragmented UDP packets.

Note:

The PCSS value should not exceed a pointer to IP header start or else it will erroneously calculate IP header checksum or TCP/UDP checksum.

RXCSUM.IPOFLD is used to enable the IP Checksum off-loading feature. If RXCSUM.IPOFLD is set to one, LAN Controller will calculate the IP checksum and indicate a pass/fail indication to software via the IP Checksum Error bit (IPE) in the ERROR field of the receive descriptor. Similarly, if RXCSUM.TUOFLD is set to one, LAN Controller will calculate the TCP or UDP checksum and indicate a pass/fail indication to software via the TCP/UDP Checksum Error bit (TCPE). Similarly, if RFCTL.IPv6_DIS and RFCTL.IP6Xsum_DIS are cleared to zero and RXCSUM.TUOFLD is set to one, LAN Controller will calculate the TCP or UDP checksum for IPv6 packets. It then indicates a pass/fail condition in the TCP/UDP Checksum Error bit (RDESC.TCPE).

This applies to checksum off loading only. Supported Frame Types:

- Ethernet II
- Ethernet SNAP

This register should only be initialized (written) when the receiver is not enabled (e.g. only write this register when RCTL.EN = 0).



12.0.3.4.19 Receive Filter Control Register - RFCTL (0x05008; RW)

Bits	Type	Default	Description
0	RW	0	iSCSI Disable (ISCSI_DIS) . Disable the iSCSI filtering for header split functionality.
1:5	RW	0	iSCSI DWord count (ISCSI_DWC) . This field indicated the Dword count of the iSCSI header, which is used for packet split mechanism.
6	RW	0	NFS Write disable (NFSW_DIS) . Disable filtering of NFS write request headers for header split functionality.
7	RW	0	NFS Read disable (NFSR_DIS) . Disable filtering of NFS read reply headers for header split functionality.
9:8	RW	00	NFS Version (NFS_VER) . 00 – NFS version 2 01 – NFS version 3 10 – NFS version 4 11 – Reserved for future use
10	RW	0	Reserved.
11	RW	0	Reserved.
12	RW	0	ACK accelerate disable (ACKDIS) . When this bit is set LAN Controller will not accelerate interrupt on TCP ACK packets.
13	RW	0	ACK data Disable (ACKD_DIS) . 1 – LAN Controller will recognize ACK packets according to the ACK bit in the TCP header + No –CP data 0 – LAN Controller will recognize ACK packets according to the ACK bit only. This bit is relevant only if the ACKDIS bit is not set.
14	RW	0	IP Fragment Split Disable (IPFRSP_DIS) . When this bit is set the header of IP fragmented packets will not be set.
15	RW	0	Extended status Enable (EXSTEN) . When the EXSTEN bit is set or when the Packet Split receive descriptor is used, LAN Controller writes the extended status to the Rx descriptor.
16	RO	0	Reserved.
17	RO	0	Reserved.
31:18	RO	0	Reserved. Should be written with 0 to ensure future compatibility.

12.0.3.4.20 Multicast Table Array - MTA[31:0] (0x05200-0x0527C; RW)

Bits	Type	Default	Description
31:0	RW	X	Bit Vector. Word wide bit vector specifying 32 bits in the multicast address filter table.

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



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

Note: These registers' addresses have been moved from where they were located in the 82542. However, for 82542 compatibility, these registers may also be accessed at their alias offsets of 0x00200-0x0027C.

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

Multicast Offset	Bits directed to the Multicast Table Array
00b	DA[47:38] = Byte 6 bits 7:0, Byte 5 bits 7:6
01b	DA[46:37] = Byte 6 bits 6:0, Byte 5 bits 7:5
10b	DA[45:36] = Byte 6 bits 5:0, Byte 5 bits 7:4
11b	DA[43:34] = Byte 6 bits 3:0, Byte 5 bits 7:2





Figure 12-1. Multicast Table Array Algorithm

12.0.3.4.21 Receive Address Low - RAL (0x05400 + 8*n (n=0...6); RW)

While "n" is the exact unicast/Multicast address entry and it is equals to 0,1,...6



Bits	Type	Default	Description
31:0	RW	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 0 and 1 in the NVM.

Note: These registers' addresses have been moved from where they were located in the 82542. However, for 82542 compatibility, these registers may also be accessed at their alias offsets of 0x0040-0x000BC.

12.0.3.4.22 Receive Address High - RAH (0x05404 + 8*n (n=0...6); RW)

While "n" is the exact unicast/Multicast address entry and it is equals to 0,1,...6

Bits	Type	Default	Description
15:0	RW	X	Receive Address High (RAH). The upper 16 bits of the 48 bit Ethernet address n (n=0, 1...6). RAH 0 is loaded from word 2 in the NVM.
17:16	RW	X	Address Select (ASEL). Selects how the address is to be used. Decoded as follows: 00 – Destination address (must be set to this in normal mode) 01 – Source address 10 – Reserved 11 – Reserved
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
30:19	RO	0	Reserved. Reads as 0. Ignored on write.
31	RW	see description	Address valid (AV). Cleared after master reset. If the NVM is present, the <i>Address Valid</i> field of <i>Receive Address Register 0</i> will be set to 1 after a software or PCI reset or NVM read. This bit is cleared by master (software) reset.

AV determines whether this address is compared against the incoming packet. AV is cleared by a master (software) reset.

ASEL enables the device to perform special filtering on receive packets.

Note: The first receive address register (RAR0) is also used for exact match pause frame checking (DA matches the first register). Therefore RAR0 should always be used to store the individual Ethernet MAC address of the adapter.

Note: These registers' addresses have been moved from where they were located in the 82542. However, for 82542 compatibility, these registers may also be accessed at their alias offsets of 0x0040-0x000BC.

After reset, if the NVM is present, the first register (Receive Address Register 0) will be loaded from the IA field in the NVM, its *Address Select* field will be 00, and its *Address Valid* field will be 1. If no NVM is present the *Address Valid* field will be 0. The *Address Valid* field for all of the other registers will be 0.

**12.0.3.4.23 Shared Receive Address Low - SHRAL[n] (0x05438 + 8*n (n=0...3); RW)**

Bits	Type	Default	Description
31:0	RW	X	Receive Address Low (RAL) . The lower 32 bits of the 48 bit Ethernet address n (n=0...3).

These registers may be WR locked by the LockMAC field in the FWSM register.

12.0.3.4.24 Shared Receive Address High 0...2 - SHRAH[n] (0x0543C + 8*n (n=0...2); RW)

Bits	Type	Default	Description
15:0	RW	X	Receive Address High (RAH) . The upper 16 bits of the 48 bit Ethernet address n (n=0...3).
17:16	RO	00	Address Select (ASEL) . Selects how the address is to be used. 00b means that it is used to decode the Destination MAC address.
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
30:19	RO	0	Reserved . Reads as 0. Ignored on write.
31	RW	0	Address valid (AV) . When this bit is set the relevant RAL,RAH are valid (compared against the incoming packet). Init trigger of this field depends on the state of the LockMAC state (see below).


These registers may be WR locked by the LockMAC field in the FWSM register.

The description of SHRAH[9] is different and is described in [Section 12.0.3.4.25](#).

12.0.3.4.25 Shared Receive Address High 3 - SHRAH[3] (0x05454; RW)

Bits	Type	Default	Description
15:0	RW	X	Receive Address High (RAH) . The upper 16 bits of the 48 bit Ethernet address n (n=0...3).
17:16	RO	00	Address Select (ASEL) . Selects how the address is to be used. 00b means that it is used to decode the Destination MAC address.
18	RW	0	VMDq output index (VIND) – defines the VMDq output index associated with a receive packet that matches this MAC address (RAH and RAL).
29:19	RO	0	Reserved . Reads as 0. Ignored on write.
30	RW	0	All Nodes Multicast Address valid (MAV) . The All Nodes Multicast address (33:33:00:00:00:01) is valid when this bit is set. Init trigger of this bit depends on the state of the LockMAC state (see below). Note that 0x33 is the first byte on the wire.
31	RW	0	Address valid (AV) . When this bit is set the relevant Address 3 is valid (compared against the incoming packet). Init trigger of this field depends on the state of the LockMAC state (see below).

These registers may be WR locked by the LockMAC field in the FWSM register.



12.0.3.4.26 Receive Address Initial CRC Calculation - RAICC[n] (0x05F50 + 4*n (n=0...10); RW)

Bits	Type	Reset	Description
31:0	RW	0x0	CRC32 calculation of the relevant configured RA address

12.0.3.4.27 Multiple Receive Queues Command register - MRQC (0x05818; RW)

Bits	Type	Default	Description
1:0	RW	0x00	Multiple Receive Queues Enable (MRxQueue) . Enables support for Multiple Receive Queues and defines the mechanism that controls queue allocation. This field can be modified only when receive to host is not enabled (RCTL.EN = 0). 00b – Multiple Receive Queues are disabled 01b – Multiple Receive Queues as defined by MSFT RSS. The RSS Field Enable bits define the header fields used by the hash function. 10b – VMDq enable, enables VMDq operation as defined in section Receive Queuing for Virtual Machine Devices 11b – Reserved
2	Reserved.	0x0	Reserved.
15:3	RO	0x0	Reserved.
21:16	RW	0x0	Reserved.
31:22	RO	0x0	Reserved.

12.0.3.4.28 RSS Interrupt Mask Register - RSSIM (0x05864; RW)

Bits	Type	Default	Description
31:0	Reserved.	0x00	Reserved.

12.0.3.4.29 RSS Interrupt Request Register - RSSIR (0x05868; RW)

Bits	Type	Default	Description
31:0	Reserved.	0x00	Reserved.

12.0.3.4.30 Redirection Table - RETA (0x05C00 + 4*n (n=0...31); RW)

The redirection table is a 32 entry table. Each entry is composed of 4 Tags each 8-bits wide. Only the first or last 6 bits of each Tag are used (5 bits for the CPU index and 1 bit for Queue index).

Offset	31:24	23:16	15:8	7:0
0x05C00 + n*4	Tag 4*n+3	Tag 4*n+2	Tag 4*n+1	Tag 4*n

Bits	Type	Default	Description
4:0	RW	X	CPU INDX 0 . CPU index for Tag 4*n (n=0,1,...31)
6:5	RO	X	Reserved.
7	RW	X	QUE INDX 0 . Queue Index for Tag 4*n (n=0,1,...31)



Bits	Type	Default	Description
12:8	RW	X	CPU INDX 1. CPU index for Tag $4*n+1$ ($n=0,1,...31$)
14:13	RO	X	Reserved.
15	RW	X	QUE INDX 1. Queue Index for Tag $4*n+1$ ($n=0,1,...31$)
20:16	RW	X	CPU INDX 2. CPU index for Tag $4*n+2$ ($n=0,1,...31$)
22:21	RO	X	Reserved.
23	RW	X	QUE INDX 2. Queue Index for Tag $4*n+2$ ($n=0,1,...31$)
28:24	RW	X	CPU INDX 3. CPU index for Tag $4*n+3$ ($n=0,1,...31$)
30:29	RO	X	Reserved.
31	RW	X	QUE INDX 3. Queue Index for Tag $4*n+3$ ($n=0,1,...31$)

Note: RETA cannot be read when RSS is enabled.

12.0.3.4.31 Random Key Register - RSSRK ($0x05C80 + 4*n$ ($n=0...9$); RW)


The RSS Random Key Register stores a 40 byte key (10 Dword entry table) used by the RSS hash function.

Bits	Type	Default	Description
7:0	RW	0x00	K0. Byte $n*4$ of the RSS random key ($n=0,1,...9$)
15:8	RW	0x00	K1. Byte $n*4+1$ of the RSS random key ($n=0,1,...9$)
23:16	RW	0x00	K2. Byte $n*4+2$ of the RSS random key ($n=0,1,...9$)
31:24	RW	0x00	K3. Byte $n*4+3$ of the RSS random key ($n=0,1,...9$)

12.0.3.5 Transmit Register Descriptions

12.0.3.5.1 Transmit Control Register - TCTL ($0x00400$; RW)

Bits	Type	Default	Description
0	RW	0	IP Identification 15 bit (IPID15). When '1', the IP Identification field will be incremented and wrapped around on 15-bit base. For example, if IP ID is equal to 0x7FFF then the next value will be 0x0000; if IP ID is equal to 0xFFFF then the next value will be 0x8000. When '0', the IP Identification field will be incremented and wrapped around on 16-bit base. In this case, the value following 0x7FFF is 0x8000, and the value following 0xFFFF is 0x0000. The purpose of this feature is to enable the software to manage 2 subgroups of connections.
1	RW	0	Enable (EN). The transmitter is enabled when this bit is 1. Writing this bit to 0 will stop transmission after any in progress packets are sent. Data remains in the transmit FIFO until the device is re-enabled. Software should combine this with reset if the packets in the FIFO should be flushed.
2	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.



Bits	Type	Default	Description
3	RW	1	Pad short packets (PSP). (with valid data, NOT padding symbols). 0 – do not pad; 1 – pad. Padding makes the packet 64B. This is not the same as the minimum collision distance. If Padding of short packets is allowed, the value in TX descriptor length field should be not less than 17 bytes.
11:4	RW	0x0F	Collision Threshold (CT). This determines the number of attempts at retransmission prior to giving up on the packet (not including the first transmission attempt). While this can be varied, it should be set to a value of 15 in order to comply with the IEEE specification requiring a total of 16 attempts. The Ethernet back-off algorithm is implemented and clamps to the maximum number of slot-times after 10 retries. This field only has meaning when in half-duplex operation.
21:12	RW	0x3F	Collision Distance (COLD). Specifies the minimum number of byte times which must elapse for proper CSMA/CD operation. Packets are padded with special symbols, not valid data bytes. Hardware checks and pads to this value plus one byte even in full-duplex operation. Default value is 64B – 512B times.
22	RW/V	0	Software XOFF Transmission (SWXOFF). When set to a 1 the device will schedule the transmission of an XOFF (PAUSE) frame using the current value of the PAUSE timer. This bit self clears upon transmission of the XOFF frame.
23	RW	0	Reserved.
24	RW	0	Re-transmit on Late Collision (RTLCL). Enables the device to retransmit on a late collision event.
27:25	RW	0x0	Reserved. Used to be UNORTX and TXDSCMT in predecessors.
28	RO	1	Reserved.
30:29	RW	01	Read Request Threshold (RRTHRESH). These bits will define the threshold size for the intermediate buffer to determine when to send the read command to the Packet buffer. Threshold is defined as follow: RRTHRESH – 00b Threshold – 2 lines of 16 bytes RRTHRESH – 01b Threshold – 4 lines of 16 bytes RRTHRESH – 10b Threshold – 8 lines of 16 bytes RRTHRESH – 11b Threshold – No threshold (transfer data after all of the request is in the RFIFO)
31	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Two fields deserve special mention: CT and COLD. Software may choose to abort packet transmission in less than the Ethernet mandated 16 collisions. For this reason, hardware provides CT.

Wire speeds of 1000 Mbps result in a very short collision radius with traditional minimum packet sizes. COLD specifies the minimum number of bytes in the packet to satisfy the desired collision distance. It is important to note that the resulting packet has special characters appended to the end. These are NOT regular data characters. Hardware strips special characters for packets that go from 1000 Mbps environments to 100 Mbps environments. Note that the hardware evaluates this field against the packet size in Full Duplex as well.

Note: While 802.3x flow control is only defined during full duplex operation, the sending of PAUSE frames via the SWXOFF bit is not gated by the duplex settings within the device.



Software should not write a 1 to this bit while the device is configured for half duplex operation.

RTLIC configures the LAN Controller to perform retransmission of packets when a late collision is detected. Note that the collision window is speed dependent: 64 bytes for 10/100 Mb/s and 512 bytes for 1000Mb/s operation. If a late collision is detected when this bit is disabled, the transmit function assumes the packet is successfully transmitted. This bit is ignored in full-duplex mode.

12.0.3.5.2 Transmit IPG Register - TIPG (0x00410; RW)

Bits	Type	Default	Description
9:0	RW	0x8	IPG Transmit Time (IPGT) . Specifies the IPG length for back-to-back transmissions equal to $[(IPGT+4) \times 8]$ bit time.
19:10	RW	0x8	IPG Receive Time 1 (IPGR1) . Specifies the defer IPG part 1 (during which carrier sense is monitored). Equal to $(IPGR1 \times 8)$ when DJHDX=0 and equals to $(IPGR1+2) \times 8$ when DJHDX=1.
29:20	RW	0x9	IPG Receive Time 2 (IPGR2) . Specifies the defer IPG. Equal to $(IPGR2+3) \times 8$ when DJHDX=0 and equal to $(IPGR2+5) \times 8$ when DJHDX=1.
31:30	RO	0	Reserved . Reads as 0. Should be written to 0 for future compatibility.

This register controls the IPG (Inter Packet Gap) timer. IPGT specifies the IPG length for back-to-back transmissions in both full and half duplex. Note that an offset of 4 byte times is added to the programmed value to determine the total IPG. Therefore, a value of 8 is recommended to achieve a 12 byte time IPG.

IPGR1 specifies the portion of the IPG in which the transmitter will defer to receive events. This should be set to 2/3 of the total effective IPG, or 8.

IPGR specifies the total IPG time for non back-to-back transmissions (transmission following deferral) in half duplex.


An offset of 5 byte times is added to the programmed value to determine the total IPG after a defer event. Therefore, a value of 7 is recommended to achieve a 12 byte time effective IPG for this case. Note the IPGR should never be set to a value greater than IPGT. If IPGR is set to a value equal to or larger than IPGT, it will override the IPGT IPG setting in half duplex, resulting in inter packet gaps that are larger than intended by IPGT in that case. Full Duplex will be unaffected by this, and will always rely on IPGT only.

In summary, the recommended TIPG value to achieve 802.3 compliant minimum transmit IPG values in full and half duplex is 0x00702008.

12.0.3.5.3 Adaptive IFS Throttle - AIT (0x00458; RW)

Bits	Type	Default	Description
15:0	RW	0x0000	Adaptive IFS value (AIFS) . This value is in units of 8 nanoseconds.
31:16	RO	0x0000	Reserved . This field should be written with 0.

Adaptive IFS throttles back-to-back transmissions in the transmit packet buffer and delays their transfer to the CSMA/CD transmit function, and thus can be used to delay the transmission of back-to-back packets on the wire. Normally, this register should be set to 0. However, if additional delay is desired between back-to-back transmits, then this register may be set with a value greater than zero.



The Adaptive IFS field provides a similar function to the IPGT field in the TIPG register (see [Section 12.0.3.6.2](#)). However, it only affects the initial transmission timing, not re-transmission timing.

Note: If the value of the AdaptiveIFS field is less than the IPGTransmitTime field in the Transmit IPG registers then it will have no effect, as the chip will select the maximum of the two values.

12.0.3.5.4 Transmit Descriptor Base Address Low - TDBAL (0x03800 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
3:0	RO	0	Reserved. Ignored on writes. Returns 0 on reads
31:4	RW	X	Transmit Descriptor Base Address Low (TDBAL)

This register contains the lower bits of the 64 bit descriptor base address. The lower 4 bits are ignored. The Transmit Descriptor Base Address must point to a 16B aligned block of data.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, TDBAL[0] may also be accessed at its alias offset of 0x00420.

12.0.3.5.5 Transmit Descriptor Base Address High - TDBAH (0x03804 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
31:0	RW	X	Transmit Descriptor Base Address [63:32] (TDBAH).

This register contains the upper 32 bits of the 64 bit Descriptor base address.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, TDBAH[0] may also be accessed at its alias offset of 0x00424.

12.0.3.5.6 Transmit Descriptor Length - TDLEN (0x03808 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
6:0	RO	0	Reserved. Ignore on write. Reads back as 0.
19:7	RW	0	Descriptor Length (LEN).
31:20	RO	0	Reserved. Reads as 0. Should be written to 0.

This register contains the descriptor length and must be 128B aligned.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, TDLEN[0] may also be accessed at its alias offset of 0x00428.

Note: The descriptor ring must be equal to or larger than 8 descriptors.

**12.0.3.5.7 Transmit Descriptor Head - TDH (0x03810 + n*0x100[n=0..1]; RW)**

Bits	Type	Default	Description
15:0	RW/V	0	Transmit Descriptor Head (TDH).
31:16	RO	0	Reserved. Should be written with 0.

This register contains the head pointer for the transmit descriptor ring. It points to a 16B datum. Hardware controls this pointer. The only time that software should write to this register is after a reset (hardware reset or CTRL.SWRST) and before enabling the transmit function (TCTL.EN). If software were to write to this register while the transmit function was enabled, the on-chip descriptor buffers may be invalidated and the hardware could become confused.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, TDH[0] may also be accessed at its alias offset of 0x00430.

12.0.3.5.8 Transmit Descriptor Tail - TDT (0x03818 + n*0x100[n=0..1]; RW)


Bits	Type	Default	Description
15:0	RW	0	Transmit Descriptor Tail (TDT).
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

Note: This register contains the tail pointer for the transmit descriptor ring. It points to a 16B datum. Software writes the tail pointer to add more descriptors to the transmit ready queue. Hardware attempts to transmit all packets referenced by descriptors between head and tail.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, TDT[0] may also be accessed at its alias offset of 0x00438.

12.0.3.5.9 Transmit Arbitration Count - TARC (0x03840 + n*0x100[n=0..1]; RW)

Bits	Type	Default	Description
6:0	RW	3	Transmit Arbitration Count (COUNT). number of packets that can be sent from queue 0 to make the N over M arbitration between the queues. Writing 0 to this register is forbidden.
7	RW	0	Compensation mode (COMP). when set to 1 the HW will compensate this queue according to the compensation ratio, if the number of packets in a TCP segmentation in queue 1 caused the counter in queue 1 to go below zero
9:8	RW	00	Compensation Ratio (RATIO). this value will determine the ratio between the number of packets transmitted on queue1 in a TCP segmentation offload to the number of compensated packets transmitted from queue 0 00 – 1/1 compensation ratio 01 – 1/2 compensation ratio 10 – 1/4 compensation ratio 11 – 1/8 compensation ratio



Bits	Type	Default	Description
10	RW	1	Descriptor enable (ENABLE). The ENABLE bit of the Transmit queue 0 should always be set.
26:11	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
27	RW	0	Reserved. Reserved for Multiple Tx request disable. This bit should not be modified by SW.
31:28	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

The default hardware value for TARC0.COUNT is **3** (this value is also reflected after reset).

The counter is subtracted as a part of the transmit arbitration.

It is reloaded to its high (last written) value when it decreases below zero.

- Upon read, the hardware returns the current counter value.
- Upon write, the counter will update the high value in the **next** counter-reload.
- The counter may be decreased in chunks (when transmitting TCP segmentation packets). It should never roll because of that.
The size of chunks is determined according to the TCP segmentation (number of packets sent).

When the counter reaches zero, other TX queues should be selected for transmission as soon as possible (usually after current transmission).

COMP is the enable bit to compensate between the two queues, when enabled (set to 1) the HW will compensate between the two queues if one of the queues is transmitting TCP segmentation packets and its counter went below zero, the HW will compensate the other queue according to the ratio in the TARC1.RATIO.

For example if the TARC0.COUNT reached (-5) after sending TCP segmentation packets and both TARC0.COMP and TARC1.COMP are enabled (set to 1) and TARC1.RATIO is 01 (1/2 compensation) TARC1.COUNT will be adjusted by adding $5/2=2$ to the current count.

RATIO is the multiplier to compensate between the two queues. The compensation method is described in the explanation above.

For DHG 802.3p using qWAVE API the following configuration will be used:

TARC0: COUNT = 1, COMP = 0, RATIO = 00.

TARC1: COUNT = 4, COMP = 1, RATIO = 00.

12.0.3.5.10 Transmit Interrupt Delay Value - TIDV (0x03820; RW)

Bits	Type	Default	Description
15:0	RW	0	Interrupt Delay Value (IDV). Counts in units of 1.024 microseconds. A value of 0 is not allowed.
30:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.
31	WO	0	Flush Partial Descriptor Block (FPD). when set to 1; ignored otherwise. Reads 0.



This register is used to delay interrupt notification for transmit operations by coalescing interrupts for multiple transmitted buffers. Delaying interrupt notification helps maximize the amount of transmit buffers reclaimed by a single interrupt. This feature only applies to transmit descriptor operations where (a) interrupt-based reporting is requested (RS set) and (b) the use of the timer function is requested (IDE is set).

This feature operates by initiating a countdown timer upon successfully transmitting the buffer. If a subsequent transmit delayed-interrupt is scheduled before the timer expires, the timer is re-initialized to the programmed value and re-starts its countdown. When the timer expires, a transmit-complete interrupt (ICR.TXDW) is generated.

Setting the value to 0 is not allowed. If an immediate (non-scheduled) interrupt is desired for any transmit descriptor, the descriptor IDE should be set to 0.

The occurrence of either an immediate (non-scheduled) or absolute transmit timer interrupt will halt the TIDV timer and eliminate any spurious second interrupts.

Transmit interrupts due to a Transmit Absolute Timer (TADV) expiration or an immediate interrupt (RS/RSP=1, IDE=0) will cancel a pending TIDV interrupt. The TIDV countdown timer is reloaded but halted, though it may be restarted by a processing a subsequent transmit descriptor.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x00440.


Writing this register with FPD set initiates an immediate expiration of the timer, causing a write back of any consumed transmit descriptors pending write back, and results in a transmit timer interrupt in the ICR.

Note: FPD is self-clearing.

12.0.3.5.11 Transmit Descriptor Control - TXDCTL (0x03828 + n*0x100[n=0..1]; RW)

Note: This register was not fully validated. SW should set it to 0x0000 during nominal operation.

Bits	Type	Default	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
15:14	RO	0x00	Reserved.
21:16	RW	0x00	Write-back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0x0	Granularity (GRAN). Units for the thresholds in this register. 0 – cache lines 1 – descriptors
31:25	RW	0x0	Transmit descriptor Low Threshold (LWTHRESH). Interrupt asserted when the number of descriptors pending service in the transmit descriptor queue (processing distance from the TDT) drops below this threshold.



This register controls the fetching and write-back of transmit descriptors. The three threshold values are used to determine when descriptors will be read from and written to host memory. The values may be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag.

Note: When GRAN = 1 all descriptors will be written back (even if not requested).

PTHRESH is used to control when a prefetch of descriptors will be considered. This threshold refers to the number of valid, unprocessed transmit descriptors the chip has in its on-chip buffer. If this number drops below PTHRESH, the algorithm will consider pre-fetching descriptors from host memory. This fetch will not happen however, unless there are at least HTHRESH valid descriptors in host memory to fetch.

Note: HTHRESH should be given a non zero value when ever PTHRESH is used.

WTHRESH controls the write-back of processed transmit descriptors. This threshold refers to the number of transmit descriptors in the on-chip buffer which are ready to be written back to host memory. In the absence of external events (explicit flushes), the write-back will occur only after at least WTHRESH descriptors are available for write-back.

Possible values:

GRAN = 1 (descriptor granularity):

PTHRESH = 0..31

WTHRESH = 0..31

HTHRESH = 0..31

GRAN = 0 (cacheline granularity):

PTHRESH = 0..3 (for 16 descriptors cacheline - 256 bytes)

WTHRESH = 0..3

HTHRESH = 0..4

Note: For any WTHRESH value other than 0 - The packet and absolute timers must get a non zero value for the WTHRESH feature to take affect.

Note: Since the default value for write-back threshold is 0, descriptors are normally written back as soon as they are processed. WTHRESH must be written to a non-zero value to take advantage of the write-back bursting capabilities of the LAN Controller. If the WTHRESH is written to a non-zero value then all of the descriptors are written back consecutively no matter the setting of the RS bit.

Since write-back of transmit descriptors is optional (under the control of RS bit in the descriptor), not all processed descriptors are counted with respect to WTHRESH. Descriptors start accumulating after a descriptor with RS is set. Furthermore, with transmit descriptor bursting enabled, all of the descriptors are written back consecutively no matter the setting of the RS bit.

Note: Leaving this value at its default will cause descriptor processing to be similar to the 82542.

LWTHRESH controls the number of pre-fetched transmit descriptors at which a transmit descriptor-low interrupt (ICR.TXD_LOW) is reported. This may allow software to operate more efficiently by maintaining a continuous addition of transmit work, interrupting only when the hardware nears completion of all submitted work.



LWTHRESH specifies a multiple of 8 descriptors. An interrupt is asserted when the number of descriptors available transitions from (threshold level=8*LWTHRESH)+1 to (threshold level=8*LWTHRESH). Setting this value to 0 will disable this feature.

12.0.3.5.12 Transmit Absolute Interrupt Delay Value-TADV (0x0382C; RW)

Bits	Type	Default	Description
15:0	RW	0	Interrupt Delay Value (IDV). Counts in units of 1.024 microseconds. (0 – disabled)
31:16	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

The transmit interrupt delay timer (TIDV) may be used to coalesce transmit interrupts. However, it may be necessary to ensure that no completed transmit remains unnoticed for too long an interval in order ensure timely release of transmit buffers. This register may be used to *ensure* that a transmit interrupt occurs at some predefined interval after a transmit is completed. Like the delayed-transmit timer, the absolute transmit timer *only* applies to transmit descriptor operations where (a) interrupt-based reporting is requested (RS set) and (b) the use of the timer function is requested (IDE is set).

This feature operates by initiating a countdown timer upon successfully transmitting the buffer. When the timer expires, a transmit-complete interrupt (ICR.TXDW) is generated. The occurrence of either an immediate (non-scheduled) or delayed transmit timer (TIDV) expiration interrupt will halt the TADV timer and eliminate any spurious second interrupts.

Setting the value to 0 disables the transmit absolute delay function. If an immediate (non-scheduled) interrupt is desired for any transmit descriptor, the descriptor IDE should be set to 0.

12.0.3.6 Statistic Register Descriptions


Note: All Statistics registers reset when read. In addition, they stick at 0xFFFF_FFFF when the maximum value is reached.

Note: For the receive statistics it should be noted that a packet is indicated as “received” if it passes the device’s filters and is placed into the packet buffer memory. A packet does not have to be DMA’d to host memory in order to be counted as “received”.

Note: Due to divergent paths between interrupt-generation and logging of relevant statistics counts, it may be possible to generate an interrupt to the system for a noteworthy event prior to the associated statistics count actually being incremented. This is extremely unlikely due to expected delays associated with the system interrupt-collection and ISR delay, but might be observed as an interrupt for which statistics values do not quite make sense. Hardware guarantees that any event noteworthy of inclusion in a statistics count will be reflected in the appropriate count within 1 microsecond; a small time-delay prior to read of statistics may be necessary to avoid the potential for receiving an interrupt and observing an inconsistent statistics count as part of the ISR.

12.0.3.6.1 CRC Error Count - CRCERRS (0x04000; RO)

Bits	Type	Default	Description
31:0	RO/V	0	CEC. CRC error count.



Counts the number of receive packets with CRC errors. In order for a packet to be counted in this register, it must pass address filtering and must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. If receives are not enabled, then this register will not increment.

12.0.3.6.2 RX Error Count - RXERRC (0x0400C; RO)

Bits	Type	Default	Description
31:0	RO/V	0	RXEC. RX error count.

Counts the number of packets received in which RX_ER was asserted by the PHY. In order for a packet to be counted in this register, it must pass address filtering and must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. If receives are not enabled, then this register will not increment.

12.0.3.6.3 Missed Packets Count - MPC (0x04010; RO)

Bits	Type	Default	Description
31:0	RO/V	0	MPC. Missed Packets Count.

Counts the number of missed packets. Packets are missed when the receive FIFO has insufficient space to store the incoming packet. This could be caused by too few buffers allocated, or because there is insufficient bandwidth on the IO bus. Events setting this counter cause RXO, the receiver overrun interrupt, to be set. This register does not increment if receives are not enabled.

Note: Note that these packets will also be counted in the Total Packets Received register as well as in Total Octets Received.

12.0.3.6.4 Carrier Extension Error Count - CEXTERR (0x0403C; RO)

Bits	Type	Default	Description
31:0	RO/V	0	CEXTERR. Number of packets received with a carrier extension error.

This register counts the number of packets received in which the carrier extension error was signaled across the GMII interface. The PHY propagates carrier extension errors to the MAC when an error is detected during the carrier extended time of a packet reception. An extension error is signaled by the PHY by the encoding of 0x1F on the receive data inputs while RX_ER is asserted to the MAC. This register will only increment if receives are enabled and the device is operating at 1000Mb/s.

12.0.3.6.5 XON Received Count - XONRXC (0x04048; RO)

Bits	Type	Default	Description
31:0	RO/V	0	XONRXC. Number of XON packets received.

This register counts the number of XON packets received. XON packets can use the global address, or the station address. This register will only increment if receives are enabled.

**12.0.3.6.6 XON Transmitted Count - XONTXC (0x0404C; RO)**

Bits	Type	Default	Description
31:0	RO/V	0	XONTXC . Number of XON packets transmitted.

This register counts the number of XON packets transmitted. These can be either due to queue fullness, or due to software initiated action (using SWXOFF). This register will only increment if transmits are enabled.

12.0.3.6.7 XOFF Received Count - XOFFRXC (0x04050; RO)

Bits	Type	Default	Description
31:0	RO/V	0	XOFFRXC . Number of XOFF packets received.

This register counts the number of XOFF packets received. XOFF packets can use the global address, or the station address. This register will only increment if receives are enabled.

12.0.3.6.8 XOFF Transmitted Count - XOFFTXC (0x04054; RO)

Bits	Type	Default	Description
31:0	RO/V	0	XOFFTXC . Number of XOFF packets transmitted.

This register counts the number of XOFF packets transmitted. These can be either due to queue fullness, or due to software initiated action (using SWXOFF). This register will only increment if transmits are enabled.

12.0.3.6.9 FC Received Unsupported Count - FCRUC (0x04058; RO)

Bits	Type	Default	Description
31:0	RO/V	0	FCRUC . Number of unsupported flow control frames received.


This register counts the number of unsupported flow control frames that are received.

The FCRUC counter is incremented when a flow control packet is received which matches either the reserved flow control multicast address (in FCAH/L) or the MAC station address, and has a matching flow control type field match (to the value in FCT), but has an incorrect opcode field. This register will only increment if receives are enabled.

12.0.3.6.10 Good Packets Received Count - GPRC (0x04074; RO)

Bits	Type	Default	Description
31:0	RO/V	0	GPRC . Number of good packets received (of any legal length).

This register counts the number of good (non-erred) packets received of any legal length. The legal length for the received packet is defined by the value of LongPacketEnable (see 10.2.7.13 Receive Length Error Count). This register does not include received flow control packets and only counts packets that pass filtering. This register will only increment if receives are enabled. This register does not count packets counted by the *Missed Packet Count (MPC)* register.



12.0.3.6.11 Broadcast Packets Received Count - BPRC (0x04078; RO)

Bits	Type	Default	Description
31:0	RO/V	0	BPRC. Number of broadcast packets received.

This register counts the number of good (non-erred) broadcast packets received. This register counts every broadcast packets received that passed filtering. This register does not count packets counted by the *Missed Packet Count (MPC)* register.

12.0.3.6.12 Multicast Packets Received Count - MPRC (0x0407C; RO)

Bits	Type	Default	Description
31:0	RO/V	0	MPRC. Number of multicast packets received.

This register counts the number of good (non-erred) multicast packets received. This register does not count multicast packets received that fail to pass address filtering nor does it count received flow control packets. This register does not count packets counted by the *Missed Packet Count (MPC)* register.

12.0.3.6.13 Good Packets Transmitted Count - GPTC (0x04080; RO)

Bits	Type	Default	Description
31:0	RO/V	0	GPTC. Number of good packets transmitted.

This register counts the number of good (non-erred) packets transmitted. A good transmit packet is considered one that is 64 or more bytes in length (from <Destination Address> through <CRC>, inclusively) in length. This does not include transmitted flow control packets. This register will only increment if transmits are enabled. This register does not count packets counted by the *Missed Packet Count (MPC)* register. The register counts clear as well as secure packets.

12.0.3.6.14 Good Octets Received Count - GORCL (0x04088; RO)

12.0.3.6.15 Good Octets Received Count - GORCH (0x0408C; RO)

Bits	Type	Default	Description
31:0	RO/V	0	GORCL. Number of good octets received – lower 4 bytes.
31:0	RO/V	0	GORCH. Number of good octets received – upper 4 bytes.

These registers make up a logical 64-bit register which counts the number of good (non-erred) octets received. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (GORCL) at this point the HW will snapshot the upper 32 bits to be read (GORCH).

In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached. Only packets that pass address filtering are counted in this register. This register will only increment if receives are enabled.

These octets do not include octets in received flow control packets.

**12.0.3.6.16 Good Octets Transmitted Count - GOTCL (0x04090; RO);****12.0.3.6.17 Good Octets Transmitted Count - GOTCH (0x04094; RO);**

Bits	Type	Default	Description
31:0	RO/V	0	GOTCL . Number of good octets transmitted – lower 4 bytes.
31:0	RO/V	0	GOTCH . Number of good octets transmitted – upper 4 bytes.

These registers make up a logical 64-bit register which counts the number of good (non-erred) packets transmitted. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (GOTCL) at this point the HW will snapshot the upper 32 bits to be read (GOTCH).

In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached. This register includes bytes transmitted in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register counts octets in successfully transmitted packets which are 64 or more bytes in length. This register will only increment if transmits are enabled. The register counts clear as well as secure octets.

These octets do not include octets in transmitted flow control packets.

12.0.3.6.18 Receive No Buffers Count - RNBC (0x040A0; RO)

Bits	Type	Default	Description
31:0	RO/V	0	DefaultRNBC . Number of receive no buffer conditions.

This register counts the number of times that frames were received when there were no available buffers in host memory to store those frames (receive descriptor head and tail pointers were equal). The packet will still be received if there is space in the FIFO. This register will only increment if receives are enabled.

This register does not increment when flow control packets are received.

12.0.3.6.19 Receive Undersize Count - RUC (0x040A4; RO)


Bits	Type	Default	Description
31:0	RO/V	0	RUC . Number of receive undersize errors.

This register counts the number of received frames that passed address filtering, and were less than the minimum size (64 bytes from <Destination Address> through <CRC>, inclusively), and had a valid CRC. This register will only increment if receives are enabled.

12.0.3.6.20 Receive Fragment Count - RFC (0x040A8; RO)

Bits	Type	Default	Description
31:0	RO/V	0	RFC . Number of receive fragment errors.

This register counts the number of received frames that passed address filtering, and were less than the minimum size (64 bytes from <Destination Address> through <CRC>, inclusively), but had a bad CRC (this is slightly different from the Receive Undersize Count register). This register will only increment if receives are enabled.



12.0.3.6.21 Receive Oversize Count - ROC (0x040AC; RO)

Bits	Type	Default	Description
31:0	RO/V	0	ROC. Number of receive oversize errors.

This register counts the number of received frames that passed address filtering, and were greater than maximum size. Packets over 1522 bytes are oversized if LongPacketEnable is 0. If LongPacketEnable (LPE) is 1, then an incoming, packet is considered oversized if it exceeds 16384 bytes.

If receives are not enabled, this register will not increment. These lengths are based on bytes in the received packet from <Destination Address> through <CRC>, inclusively.

12.0.3.6.22 Receive Jabber Count - RJC (0x040B0; RO)

Bits	Type	Default	Description
31:0	RO/V	0	RJC. Number of receive jabber errors.

This register counts the number of received frames that passed address filtering, and were greater than maximum size and had a bad CRC (this is slightly different from the Receive Oversize Count register).

Packets over 1522 bytes are oversized if LongPacketEnable is 0. If LongPacketEnable (LPE) is 1, then an incoming packet is considered oversized if it exceeds 16384 bytes.

If receives are not enabled, this register will not increment. These lengths are based on bytes in the received packet from <Destination Address> through <CRC>, inclusively.

12.0.3.6.23 Management Packets Received Count - MNGPRC (0x040B4; RO)

Bits	Type	Default	Description
15:0	RO/V	0	Reserved.
31:16	RO	0	Reserved.

This register counts the total number of packets received that pass the management filters or receive circuit breaker redirection. Any packets with errors are not counted, except for packets dropped because the management receive FIFO is full.

12.0.3.6.24 Management Packets Dropped Count - MNGPDC (0x040B8; RO)

Bits	Type	Default	Description
15:0	RO/V	0	MPDC. Number of management packets dropped.
31:16	RO	0	Reserved.

This register counts the total number of packets received that pass the management filters and then are dropped because the management receive FIFO is full.

**12.0.3.6.25 Management Packets Transmitted Count - MNGPTC (0x040BC; RO)**

Bits	Type	Default	Description
15:0	RO/V	0	MPTC. Number of management packets transmitted.
31:16	RO	0	Reserved.

This register counts the total number of packets that are transmitted that are received over the LAN Controller PCI-M interface.

12.0.3.6.26 Tx Circuit Breaker Packets Dropped -TCBPD (0x040D8; RO)

Bits	Type	Default	Description
15:0	RO/V	0	TCBD. Dropped redirected transmit CB packets. Packets are dropped due to lack of room in the redirection MNG FIFO. It could be either due to M-Link BW or no memory resources.
31:16	RO	0	Reserved.

12.0.3.6.27 Total Octets Received - TORL (0x040C0; RO);**12.0.3.6.28 Total Octets Received - TORH (0x040C4; RO);**

Bits	Type	Default	Description
31:0	RO/V	0	TORL. Number of total octets received – lower 4 bytes.
31:0	RO/V	0	TORH. Number of total octets received – upper 4 bytes.

These registers make up a logical 64-bit register which count the total number of octets received. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (TORL) at this point the HW will snapshot the upper 32 bits to be read (TORH). In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached.


All packets received that pass address filtering will have their octets summed into this register, regardless of their length, whether they are erred, or whether they are flow control packets. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register will only increment if receives are enabled.

Note: Broadcast rejected packets will be counted in this counter (in contradiction to all other rejected packets that are not counted).

12.0.3.6.29 Total Octets Transmitted - TOTL (0x040C8; RO)

Bits	Type	Default	Description
31:0	RO/V	0	TOTL. Number of total octets transmitted – lower 4 bytes.

The TOTL and TOTH registers make up a logical 64-bit register which count the total number of octets transmitted. This register must be accessed using two independent 32-bit accesses. SW should first read the low register (TORL) at this point the HW will snapshot the upper 32 bits to be read (TOTH). In addition, it sticks at 0xFFFF_FFFF_FFFF_FFFF when the maximum value is reached.



All transmitted packets will have their octets summed into this register, regardless of their length or whether they are flow control packets. This register includes bytes transmitted in a packet from the <Destination Address> field through the <CRC> field, inclusively.

Octets transmitted as part of partial packet transmissions (e.g. collisions in half-duplex mode) are not included in this register. This register will only increment if transmits are enabled.

12.0.3.6.30 Total Octets Transmitted - TOTH (0x040CC; RO)

Bits	Type	Default	Description
31:0	RO/V	0	TOTH. Number of total octets transmitted – upper 4 bytes

See explanation of the TOTL above.

12.0.3.6.31 Total Packets Received - TPR (0x040D0; RO)

Bits	Type	Default	Description
31:0	RO/V	0	TPR. Number of all packets received.

This register counts the total number of all packets received. All packets received will be counted in this register, regardless of their length, whether they are erred, or whether they are flow control packets. This register will only increment if receives are enabled.

Note: Broadcast rejected packets will be counted in this counter (in contradiction to all other rejected packets that are not counted).

12.0.3.6.32 Total Packets Transmitted - TPT (0x040D4; RO)

Bits	Type	Default	Description
31:0	RO/V	0	TPT. Number of all packets transmitted.


This register counts the total number of all packets transmitted. All packets transmitted will be counted in this register, regardless of their length, or whether they are flow control packets.

Partial packet transmissions (e.g. collisions in half-duplex mode) are not included in this register. This register will only increment if transmits are enabled. This register counts all packets, including standard packets, secure packets, packets received over the SMB, and packets generated by the ASF function.

12.0.3.6.33 Multicast Packets Transmitted Count - MPTC (0x040F0; RO)

Bits	Type	Default	Description
31:0	RO/V	0	MPTC. Number of multicast packets transmitted.

This register counts the number of multicast packets transmitted. This register does not include flow control packets and increments only if transmits are enabled. Counts clear as well as secure traffic.

**12.0.3.6.34 Broadcast Packets Transmitted Count - BPTC (0x040F4; RO)**

Bits	Type	Default	Description
31:0	RO/V	0	BPTC. Number of broadcast packets transmitted count.

This register counts the number of broadcast packets transmitted. This register will only increment if transmits are enabled. This register counts all packets, including standard and secure packets. (Management packets will never be more than 200 bytes).

12.0.3.6.35 TCP Segmentation Context Transmitted Count - TSCTC (0x040F8; RO)

Bits	Type	Default	Description
31:0	RO/V	0	TSCTC. Number of TCP Segmentation contexts transmitted count.


This register counts the number of TCP segmentation offload transmissions and increments once the last portion of the TCP segmentation context payload is segmented and loaded as a packet into the on-chip transmit buffer. Note that it is not a measurement of the number of packets sent out (covered by other registers). This register will only increment if transmits and TCP Segmentation offload are enabled.

12.0.3.6.36 Interrupt Assertion Count - IAC (0x04100; RO)

Bits	Type	Default	Description
0-31	RO/V	0	IAC. This is a count of the interrupt assertions that have occurred. It counts the total number of interrupts generated in the system.

12.0.3.7 Management Register Descriptions**12.0.3.7.1 Wake Up Control Register - WUC (0x05800; RW)**

Bits	Type	Default	Description
0	RW/ SN	0	Advance Power Management Enable (APME). 1 – APM Wakeup is enabled 0 – APM Wakeup is disabled Loaded from the NVM word 1Ah.
1	RW/V	0	PME_En. This read/write bit is used by the driver to access the PME_En bit of the <i>Power Management Control / Status Register</i> (PMCSR) without writing to PCI configuration space.
2	RWC	0	PME_Status. This bit is set when LAN Controller receives a wakeup event. It is the same as the PME_Status bit in the <i>Power Management Control / Status Register</i> (PMCSR). Writing a "1" to this bit will clear it, and also clear the PME_Status bit in the PMCSR.
3	RW	1	Assert PME On APM Wakeup (APMPME). If it is 1, LAN Controller will set the PME_Status bit in the <i>Power Management Control / Status Register</i> (PMCSR) and assert Host_Wake when APM Wakeup is enabled and LAN Controller receives a matching magic packet.
4	RW/ SN	0	Link Status Change Wake Enable (LSCWE). Enables wake on link status change as part of APM wake capabilities.



Bits	Type	Default	Description
5	RW/ SN	0	Link Status Change Wake Override (LSCWO). If "1", wake on Link Status Change does not depend on the LNK_C bit in the Wake Up Filter Control Register (WUFC). Instead, it is determined by the APM settings in the WUC register.
6	RO	0	Reserved. Was "APM Flexible Filter Allocation (APMFFA)".
7	RO	0	Reserved. Was "Flexible APM filter Enable (FLEX_APM_FILTER_EN)".
8	RW/ SN	0	Phy_Wake. This bit indicates if the Phy connected to the LAN controller supports wakeup. This bit is loaded from NVM word 13h bit 8.
29:9	RO	0	Reserved. Reads as 0.
31:30	RO	0	Reserved.

The PME_Status bits are cleared in the following conditions:

- If there is VAUX then the PME Status bits should be cleared by:
 - PWR Good
 - Explicit Software Clear
- If there is NO VAUX then the PME Status bits should be cleared by:
 - PWR Good
 - PCI Reset de-assertion
 - Explicit Software Clear

12.0.3.7.2 Wake Up Filter Control Register - WUFC (0x05808; RW)

Bits	Type	Default	Description
0	RW	0	LNKC. Link Status Change Wake Up Enable
1	RW	0	MAG. Magic Packet Wake Up Enable
2	RW	0	EX. Directed Exact Wake Up Enable
3	RW	0	MC. Directed Multicast Wake Up Enable
4	RW	0	BC. Broadcast Wake Up Enable
5	RW	0	ARP. ARP/IPv4 Request Packet Wake Up Enable
6	RW	0	IPV4. Directed IPv4 Packet Wake Up Enable
7	RW	0	IPV6. Directed IPv6 Packet Wake Up Enable
8	RO	0	Reserved.
9:14	RO	0	Reserved.
15	RW	0	NoTCO. Ignore TCO Packets for TCO. If the NoTCO bit is set, then any packet that passes the manageability packet filtering will not cause a Wake Up event even if it passes one of the Wake Up Filters.
16	RW	0	FLX0. Flexible Filter 0 Enable
17	RW	0	FLX1. Flexible Filter 1 Enable
18	RW	0	FLX2. Flexible Filter 2 Enable
19	RW	0	FLX3. Flexible Filter 3 Enable
22	RW	0	FLX6. Flexible Filter 6 Enable



Bits	Type	Default	Description
23	RW	0	FLX7. Flexible Filter 7 Enable
31:24	RO	0	Reserved.

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

12.0.3.7.3 Wake Up Status Register - WUS (0x05810; RW)

Bits	Type	Default	Description
0	RW	0	LNKC. Link Status Changed
1	RW	0	MAG. Magic Packet Received
2	RW	0	EX. Directed Exact Packet Received. The packet's address matched one of the 7 pre-programmed exact values in the <i>Receive Address</i> registers.
3	RW	0	MC. Directed Multicast Packet Received. The packet was a multicast packet that hashed to a value corresponding to a 1 bit in the <i>Multicast Table Array</i> .
4	RW	0	BC. Broadcast Packet Received
5	RW	0	ARP. ARP/IPv4 Request Packet Received
6	RW	0	IPV4. Directed IPv4 Packet Received
7	RW	0	IPV6. Directed IPv6 Packet Received
15:8	RO	0	Reserved. Read as 0
16	RW	0	FLX0. Flexible Filter 0 Match
17	RW	0	FLX1. Flexible Filter 1 Match
18	RW	0	FLX2. Flexible Filter 2 Match
19	RW	0	FLX3. Flexible Filter 3 Match
20	RW	0	FLX4. Flexible Filter 4 Match
21	RW	0	FLX5. Flexible Filter 5 Match
22	RW	0	FLX6. Flexible Filter 6 Match
23	RW	0	FLX7. Flexible Filter 7 Match
31:24	RO	0	Reserved.


This register is used to record statistics about all Wake Up packets received. A packet that matches multiple criteria may set multiple bits. Writing a 1 to any bit will clear that bit.

This register will not be cleared when PCI_RST_N is asserted. It will only be cleared when LAN PWR Good is de-asserted or when cleared by the driver.

12.0.3.7.4 IP Address Valid - IPAV (0x5838; RW)

The IP Address Valid indicates whether the IP addresses in the IP Address Table are valid:

Bits	Type	Default	Description
0	RO	0	Reserved.



Bits	Type	Default	Description
1	RW	0	V41. IPv4 Address 1 Valid
2	RW	0	V42. IPv4 Address 2 Valid
3	RW	0	V43. IPv4 Address 3 Valid
4:15	RO	0x00	Reserved.
16	RW	0	V60. IPv6 Address Valid
31:17	RO	0x00	Reserved.

12.0.3.7.5 IPv4 Address Table - IP4AT (0x05840 + 8*n (n=1...3); RW)

The IPv4 Address Table is used to store the three IPv4 addresses for ARP/IPv4 Request packet and Directed IPv4 packet wake up. It is a 4 entry table with the following format:

Bits	Type	Default	Description
31:0	RW	X	IPADD. IP Address n (n=1, 2, 3)

The register at address 0x5840 (n=0) was used in predecessors and reserved in the LAN Controller.

12.0.3.7.6 IPv6 Address Table - IP6AT (0x05880 + 4*n (n=0...3); RW)

The IPv6 Address Table is used to store the IPv6 address for Directed IPv6 packet wake up and Manageability traffic filtering. The IP6AT has the following format:

Bits	Type	Default	Description
31:0	RW	X	IPV6 Address. IPv6 Address bytes n*4...n*4+3 (n=0, 1, 2, 3) while byte 0 is first on the wire and byte 15 is last.

The IP6AT may be used by both host and manageability engine. An interrupt mechanism is added to inform the manageability engine on any change of these registers by the host software.

12.0.3.7.7 Flexible Filter Length Table - FFLT (0x05F00 + 8*n (n=0...7); RW)

There are 8 flexible filters Lengths. The Flexible Filter Length Table stores the minimum packet lengths required to pass each of the Flexible Filters. Any packets that are shorter than the programmed length will not pass that filter. Each Flexible Filter will consider a packet that does not have any mismatches up to that point to have passed the Flexible Filter when it reaches the required length. It will not check any bytes past that point.

Bits	Type	Default	Description
10:0	RW	X	LEN. Minimum Length for Flexible Filter n.
31:11	RO	X	Reserved.

All reserved fields read as 0's and ignore writes.

Note: Before writing to the Flexible Filter Length Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).



12.0.3.7.8 Flexible Filter Mask Table - FFMT (0x09000 + 8*n (n=0...127); RW)

There are 128 mask entries. The Flexible Filter Mask and Table is used to store the four 1-bit masks for each of the first 128 data bytes in a packet, one for each Flexible Filter. If the mask bit is 1, the corresponding Flexible Filter will compare the incoming data byte at the index of the mask bit to the data byte stored in the Flexible Filter Value Table.

Bits	Type	Default	Description
0	RW	X	Mask 0. Mask for filter 0 byte n (n=0, 1... 127)
1	RW	X	Mask 1. Mask for filter 1 byte n (n=0, 1... 127)
2	RW	X	Mask 2. Mask for filter 2 byte n (n=0, 1... 127)
3	RW	X	Mask 3. Mask for filter 3 byte n (n=0, 1... 127)
6	RW	X	Mask 6. Mask for filter 6 byte n (n=0, 1... 127)
7	RW	X	Mask 7. Mask for filter 7 byte n (n=0, 1... 127)
31:8	RO	X	Reserved.

Note: The table is organized to permit expansion to 8 (or more) filters and 256 bytes in a future product without changing the address map.

Note: Before writing to the Flexible Filter Mask Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

12.0.3.7.9 Flexible Filter Value Table - FFFT (0x09800 + 8*n (n=0...127); RW)

There are 128 filter values. The Flexible Filter Value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is 1, the Flexible Filter will compare the incoming data byte to the values stored in this table.


Bits	Type	Default	Description
7:0	RW	X	Value 0. Value of filter 0 byte n (n=0, 1... 127)
15:8	RW	X	Value 1. Value of filter 1 byte n (n=0, 1... 127)
23:16	RW	X	Value 2. Value of filter 2 byte n (n=0, 1... 127)
31:24	RW	X	Value 3. Value of filter 3 byte n (n=0, 1... 127)

Before writing to the Flexible Filter Value Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

12.0.3.7.10 Flexible Filter Value Table - FFFT (0x09800 + 8*n (n=0...127); RW)

There are 128 filter values. The Flexible Filter Value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is 1, the Flexible Filter will compare the incoming data byte to the values stored in this table.

Bits	Type	Reset	Description
7:0	RW	X	Value 0. Value of filter 0 byte n (n=0, 1... 127)
15:8	RW	X	Value 1. Value of filter 1 byte n (n=0, 1... 127)
23:16	RW	X	Value 2. Value of filter 2 byte n (n=0, 1... 127)



31:24	RW	X	Value 3. Value of filter 3 byte n (n=0, 1... 127)
-------	----	---	--

Before writing to the Flexible Filter Value Table the driver must first disable the flexible filters by writing 0's to the Flexible Filter Enable bits of the Wake Up Filter Control Register (WUFC.FLXn).

12.0.3.8 Time Sync Register Description

Note: All the registers bellow are for logical needs only, upon implementation their offset or structure can be changed according to the project specific needs.

12.0.3.8.1 RX Time Sync Control register - TSYNCRXCTL (0xB620; RW)


Bits	Type	Default	Description
0	(RO/ V)	0	RXTT. Rx timestamp valid equals '1' when a valid value for Rx timestamp is captured in the Rx timestamp register, clear by read of Rx timestamp register RXSTMPH.
3:1	RW	0x0	Type. Type of packets to timestamp - 000b – Time stamp L2 (V2) packets only (Sync or Delay_req depends on message type in Section 12.0.3.8.6 and packets with message ID 2 and 3) 001b – Time stamp L4 (V1) packets only (Sync or Delay_req depends on message type in Section 12.0.3.8.6) 010b – Time stamp V2 (L2 and L4) packets (Sync or Delay_req depends on message type in Section 12.0.3.8.6 and packets with message ID 2 and 3) 100b – Time stamp all packets (in this mode no locking is done to the value in the timestamp registers and no indications in receive descriptors will be transferred) 101b - Time stamp all packets which message id bit 3 is zero, which means timestamp all event packets. This is applicable for V2 packets only. 011b, 110b and 111b – reserved
4	RW	0x0	En. Enable RX timestamp 0x0 – time stamping disabled. 0x1 – time stamping enabled.
31:5	RO	0x0	Reserved.

12.0.3.8.2 RX timestamp Low - RXSTMPL (0x0B624; RO)

Bits	Type	Default	Description
31:0	RO	0x0	RXSTMPL. Rx timestamp LSB value

12.0.3.8.3 RX timestamp High - RXSTMPH (0x0B628; RO)

Bits	Type	Default	Description
31:0	RO	0x0	RXSTMPH. Rx timestamp MSB value

**12.0.3.8.4 RX timestamp attributes low - RXSATRL (0x0B62C; RO)**

Bits	Type	Default	Description
31:0	RO	0x0	SourceIDL. Sourceuuid low The value of this register is in host order.

12.0.3.8.5 RX timestamp attributes high- RXSATRH (0x0B630; RO)

Bits	Type	Default	Description
15:0	RO	0x0	SourceIDH. Sourceuuid high The value of this register is in host order.
31:16	RO	0x0	SequenceID. SequenceI The value of this register is in host order.

12.0.3.8.6 RX message type register low - RXMTRL (0x0B634; RW)

Bits	Type	Default	Description
15:0	RW	0x88F7	PTP L2 EtherType to timestamp The value of this register is programmed/read in network order.
23:16	RW	0x0	V1 control to timestamp
31:24	RW	0x0	V2 messageId to timestamp

12.0.3.8.7 RX UDP port - RXUDP (0x0B638; RW)

Bits	Type	Default	Description
15:0	RW	0x0319	UPORT. UDP port number to time stamp The value of this register is programmed/read in network order.
31:16	RO	0x0	Reserved.

12.0.3.8.8 TX Time Sync Control register - TSYNCTXCTL (0x0B614; RW)

Bits	Type	Default	Description
0	RO/V	0	TXTT. Tx timestamp valid equals '1' when a valid value for Rx timestamp is captured in the Rx timestamp register, clear by read of Tx timestamp register TXSTMPH.
3:1	RO	0	Reserved.
4	RW	0	EN. Enable TX timestamp 0x0 – time stamping disabled. 0x1 – time stamping enabled.
5	RW	0	Packet mode. In this mode (deterministic SYSTIM) the SYSTIM register is incremented by TIMINCA.IV (increment value) every time a packet is time stamped. There is no significance whatsoever to the TIMINCA.IP (increment period) value. TIMADJ (time adjustment) registers work as they do in normal mode: every time they're set to a certain value - that value is immediately (next clock - not packet) added (/subtracted) to the SYSTIM value. 0x0 - Packet mode disabled 0x1 - Packet mode enabled



Bits	Type	Default	Description
31:6	RO	0	Reserved.

12.0.3.8.9 TX timestamp value Low - TXSTMPL (0x0B618; RO)

Bits	Type	Default	Description
31:0	RO	0x0	TXSTMPL. Tx timestamp LSB value

12.0.3.8.10 TX timestamp value High - TXSTMPH (0x0B61C; RO)

Bits	Type	Default	Description
31:0	RO	0x0	TXSTMPH. Tx timestamp MSB value

12.0.3.8.11 System time register Low - SYSTIML (0x0B600; RO)

Bits	Type	Default	Description
31:0	RW	0x0	STL. System time LSB register

12.0.3.8.12 System time register High - SYSTIMH (0x0B604; RO)

Bits	Type	Default	Description
31:0	RW	0x0	STH. System time MSB register

12.0.3.8.13 Increment attributes register - TIMINCA (0x0B608; RW)

Bits	Type	Default	Description
23:0	RW	0x0	IV. Increment value – <i>incvalue</i>
31:24	RW	0x0	IP. Increment period – <i>incperiod</i>

12.0.3.8.14 Time adjustment offset register low - TIMADJL (0x0B60C; RW)

Bits	Type	Default	Description
31:0	RW	0x00	TADJL. Time adjustment value – Low

12.0.3.8.15 Time adjustment offset register high - TIMADJH (0x0B610;RW)

Bits	Type	Default	Description
30:0	RW	0x00	TADJH. Time adjustment value - High
31	RW	0x0	Sign. Sign ("0"="+", "1"="-")

12.0.3.9 Diagnostic Register Descriptions

LAN Controller contains several diagnostic registers. These registers allow software to directly access the contents of the LAN Controller's internal Packet Buffer Memory (PBM), also referred to as FIFO space. These registers also give software visibility into what locations in the PBM that the HW currently considers to be the "head" and "tail" for both transmit and receive operations.



12.0.3.9.1 Receive Data FIFO Head Register - RDFH (0x02410; RW)

Bits	Type	Default	Description
12:0	RW	0	Receive FIFO Head pointer (FIFO Head).
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores the head pointer of the on-chip receive data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Receive FIFO Head. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Receive FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x08000. In addition, with the LAN Controller, the value in this register contains the offset of the Receive FIFO head, relative to the beginning of the entire PBM space. Alternatively, with the 82542, the value in this register contains the relative offset from the beginning of the Receive FIFO space (within the PBM space).

12.0.3.9.2 Receive Data FIFO Tail Register - RDFT (0x02418; RW)

Bits	Type	Default	Description
12:0	RW	0	Receive FIFO Tail pointer (FIFO Tail).
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.


This register stores the tail pointer of the on-chip receive data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Receive FIFO Tail. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Receive FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x08008. In addition, with the LAN Controller, the value in this register contains the offset of the Receive FIFO tail, relative to the beginning of the entire PBM space. Alternatively, with the 82542, the value in this register contains the relative offset from the beginning of the Receive FIFO space (within the PBM space).

12.0.3.9.3 Receive Data FIFO Head Saved Register - RDFHS (0x02420; RW)

Bits	Type	Default	Description
12:0	RW	0	FIFO Head. A "saved" value of the Receive FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores a copy of the Receive Data FIFO Head register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.



12.0.3.9.4 Receive Data FIFO Tail Saved Register - RDFTS (0x02428; RW)

Bits	Type	Default	Description
12:0	RW	0	FIFO Tail. A “saved” value of the Receive FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register stores a copy of the Receive Data FIFO Tail register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

12.0.3.9.5 Receive Data FIFO Packet Count - RDFPC (0x02430; RW)

Bits	Type	Default	Description
12:0	RW	0	RX FIFO Packet Count. The number of received packets currently in the RX FIFO.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register reflects the number of receive packets that are currently in the Receive FIFO. This register is available for diagnostic purposes only, and should not be written during normal operation.

12.0.3.9.6 Transmit Data FIFO Head Register - TDFH (0x03410; RW)

Bits	Type	Default	Description
12:0	RW/V	0x0900 ¹	FIFO Head. Transmit FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

1. The initial value equals PBA. RXA times 128.

This register stores the head pointer of the on-chip transmit data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Transmit FIFO Head. So a value of “0x8” in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Transmit FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

Note:

This register’s address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x08010. In addition, with the LAN Controller, the value in this register contains the offset of the Transmit FIFO head relative to the beginning of the entire PBM space. Alternatively, with the 82542, the value in this register contains the relative offset from the beginning of the Transmit FIFO space (within the PBM space).

12.0.3.9.7 Transmit Data FIFO Tail Register - TDFT (0x03418; RW)

Bits	Type	Default	Description
12:0	RW/V	0x000 ¹	FIFO Tail. Transmit FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

1. The initial value equals PBA. RXA times 128.



This register stores the head pointer of the on-chip transmit data FIFO. Since the internal FIFO is organized in units of 64 bit words, this field contains the 64 bit offset of the current Transmit FIFO Tail. So a value of "0x8" in this register corresponds to an offset of 8 QWORDS or 64 bytes into the Transmit FIFO space. This register is available for diagnostic purposes only, and should not be written during normal operation.

Note: This register's address has been moved from where it was located in the 82542. However, for 82542 compatibility, this register may also be accessed at its alias offset of 0x08018. In addition, with the LAN Controller, the value in this register contains the offset of the Transmit FIFO tail relative to the beginning of the entire PBM space. Alternatively, with the 82542, the value in this register contains the relative offset from the beginning of the Transmit FIFO space (within the PBM space).

12.0.3.9.8 Transmit Data FIFO Head Saved Register - TDFHS (0x03420; RW)

Bits	Type	Default	Description
12:0	RW/V	0x000 ¹	FIFO Head. A "saved" value of the Transmit FIFO Head pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

1. The initial value equals PBA. RXA times 128.

This register stores a copy of the Transmit Data FIFO Head register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

12.0.3.9.9 Transmit Data FIFO Tail Saved Register - TDFTS (0x03428; RW)

Bits	Type	Default	Description
12:0	RW/V	0x000 ¹	FIFO Tail. A "saved" value of the Transmit FIFO Tail pointer.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.


1. The initial value equals PBA. RXA times 128.

This register stores a copy of the Transmit Data FIFO Tail register in the case that the internal register needs to be restored. This register is available for diagnostic purposes only, and should not be written during normal operation.

12.0.3.9.10 Transmit Data FIFO Packet Count - TDFPC (0x03430; RW)

Bits	Type	Default	Description
12:0	RW	0	TX FIFO Packet Count. The number of packets to be transmitted that are currently in the TX FIFO.
31:13	RO	0	Reserved. Reads as 0. Should be written to 0 for future compatibility.

This register reflects the number of packets to be transmitted that are currently in the Transmit FIFO. This register is available for diagnostic purposes only, and should not be written during normal operation.



12.0.3.9.11 Ghost Memory Data - GMD (0x0C000 - 0x0C3FC; RW)

Bits	Type	Default	Description
31:0	RW	X	Data. Ghost Memory Data.

Ghost memory data is available to diagnostics. Locations can be accessed as 32 bit words.

12.0.3.9.12 Descriptor Memory Data - DMD (0x0C400 - 0x0C5FC; RW)

Bits	Type	Default	Description
31:0	RW	X	Data. Descriptor Memory Data.

Descriptor memory data is available to diagnostics. Locations can be accessed as 32 bit words.

12.0.3.10 MACsec Register Descriptions

12.0.3.10.1 MACsec TX Capabilities register - LSECTXCAP (0x0B000; RW)

Bits	Type	Default	Description
2:0	RO	1b	TX CA-supported Number of CA's supported by the device.
6:3	RO	1b	TX SC Capable. Number of SC's supported by the device on the transmit data path. The LAN controller supports twice the number of SA's as the TX SC for seamless re-keying, i.e., 2 SA's.
15:7	RO	0x0	Reserved.
23:16	RO/V	0x0	Tx LSEC Key SUM (LSECTXSUM). A bit wise XOR of the LSECTXKEY 0 bytes and LSECTXKEY 1 bytes. This register may be used by KaY (the programming entity) to validate key programming.
31:24	RO/V	0x0	Reserved.

12.0.3.10.2 MACsec RX Capabilities register - LSECRXCAP (0x0B300; RW)

Bits	Type	Default	Description
2:0	RO	1b	RX CA-supported Number of CA's supported by the device.
6:3	RO	4b1b	RX SC Capable. Number of SC's supported by the device on the receive data path. The LAN controller supports twice SA's as the RX SC for seamless re-keying, i.e., 8 2 SA's.
15:7	RO	0x0	Reserved.
23:16	RO/V	0x0	Rx LSEC Key SUM (LSECRXSUM). A bit wise XOR of the Rx MACsec keys 0...7 as defined in registers LSECRXKEY [n, m]. Each byte is XORed with the respective byte of the other keys. This register may be used by KaY (the programming entity) to validate key programming.
31:24	RO	0x0	Reserved.



12.0.3.10.3 MACsec TX Control register - LSECTXCTRL (0x0B004; RW)

Bits	Type	Default	Description
1:0	RW	00b	Enable Tx MACsec. Enable Tx MACsec off loading. 00b – Disable Tx MACsec (Tx all packets w/o MACsec offload). 01b – Add integrity signature. 10b – Encrypt and add integrity signature. 11b – Reserved. When this field equals 00b (MACsec offload is disabled). The "Tx Untagged Packet" register is not incremented for transmitted packets when the "Enable Tx MACsec" equals 00b.
2	RW	0	PNID PN Increase Disable 0 - Normal operation 1 - PN is not incremented, used for testability mode only.
3	RW	0	(Testability feature) when set to 1 the HW ignores the ILSEC bit in the TX descriptor and transmits the packet as if the ILSEC bit was set.
4	RW	0b	Reserved.
5	RW	1b	Always Include SCI. This field controls whether SCI is explicitly included in the transmitted SecTag. 0b – False 1b – True, SCI is explicitly included
6	RW	0b	Reserved.
7	RW	1	Reserved.
31:8	RW	11..1b	PN exhaustion threshold. MSB of the threshold over which HW needs to interrupt the KaY to warn of TX SA PN exhaustion and will trigger a new SA renegotiation. Bits 7:0 of the threshold are all 1's.

12.0.3.10.4 MACsec RX Control register - LSECRXCTRL (0x0B304; RW)

Bits	Type	Default	Description
1:0	RW	00b	Reserved.
3:2	RW	00b	Reserved.
4	RO	1b	Reserved.
5	RO	1b	Reserved.
6	RW	0b	Reserved.
7	RW	1	Replay Protect. Enable replay protection.
28:8	RO	0x0	Reserved
31:29	RW	0	Reserved.



12.0.3.10.5 MACsec TX SCI Low - LSECTXSCL (0x0B008; RW)

Bits	Type	Default	Description
31:0	RW	0b	MAC Address SecY Low. The 4 LS bytes of the MAC address copied to the SCI field in the MACsec header. The value of this register is programmed/read in host order.

12.0.3.10.6 MACsec TX SCI High - LSECTXSCH (0x0B00C; RW)

Bits	Type	Default	Description
15:0	RW	0b	MAC Address SecY High. The 2 MS bytes of the MAC address copied to the SCI field in the MACsec header. The value of this register is programmed/read in host order.
31:16	RW	0b	Port Identifier. Always zero for transmitted packets.

12.0.3.10.7 MACsec TX SA - LSECTXSA (0x0B010; RW)

Bits	Type	Default	Description
1:0	RW	0b	AN0 – Association Number 0. This 2 bit field is posted to the AN field in the transmitted MACsec header when SA 0 is active.
3:2	RW	0b	AN1 – Association Number 1. This 2 bit field is posted to the AN field in the transmitted MACsec header when SA 1 is active.
4	RW	0b	SA Select (SelSA). This bit selects between SA 0 or SA 1 smoothly, i.e., on a packet boundary. A value of '0' selects SA 0 and a value of '1' selects SA 1.
5	RO/V	0b	Active SA (ActSA). This bit indicates the active SA. The ActSA follows the value of the SelSA on a packet boundary. The KaY (the programming entity) may use this indication to retire the old SA.
6	RW	0	Reserved.
7	RW	0	Reserved.
31:8	RW	0x0	Reserved.

**12.0.3.10.8 MACsec TX SA PN 0 - LSECTXPNO (0x0B018; RW)**


Bits	Type	Default	Description
31:0	RW	0b	PN – Packet number. This field is posted to the PN field in the transmitted MACsec header when SA 0 is active. It is initialized by the KaY at SA creation and then increments by 1 for each transmitted packet using this SA. Packets should never be transmitted if the PN repeats itself. In order to protect against such an event, the HW generates an LSECPN interrupt to KaY when the PN reaches the exhaustion threshold as defined in the LSECTXCTRL register. There is additional level of defense against repeating the PN. The HW will never transmit packets after the PN reaches a value of 0xFF...FF. In order to guarantee this, the HW clears the "Enable Tx MACsec" field in the LSECTXCTRL register to 00b once a packet is transmitted with a PN that equals to 0xFF...F0. The value of this register is programmed/read in host order.

12.0.3.10.9 MACsec TX SA PN 1 - LSECTXPNI (0x0B01C; RW)

Bits	Type	Default	Description
31:0	RW	0b	PN – Packet number. This field is posted to the PN field in the transmitted MACsec header when SA 1 is active. It is initialized by the KaY at SA creation and then increments by 1 for each transmitted packet using this SA. Packets should never be transmitted if the PN repeats itself. In order to protect against such an event the HW generates an LSECPN interrupt to KaY when the PN reaches the exhaustion threshold as defined in the LSECTXCTRL register. There is additional level of defense against repeating the PN. The HW will never transmit packets after the PN reaches a value of 0xFF...FF. In order to guarantee this, the HW clears the "Enable Tx MACsec" field in the LSECTXCTRL register to 00b once a packet is transmitted with a PN that equals to 0xFF...F0. The value of this register is programmed/read in host order.

12.0.3.10.10 MACsec TX Key 0 - LSECTXKEY0 [n] (0x0B020 + 4*n (n=0...3); WO)

Bits	Type	Default	Description
31:0	WO	0x0	LSEC Key 0. Transmit MACsec key of SA 0. n – 0 LSEC Key defines bits 31:0 of the Tx MACsec Key n – 1 LSEC Key defines bits 63:32 of the Tx MACsec Key n – 2 LSEC Key defines bits 95:64 of the Tx MACsec Key n – 3 LSEC Key defines bits 127:96 of the Tx MACsec Key This field is WO for confidentiality protection. For data integrity check, the hash value may read the LSECTXSUM field in the LSECCAP register. If for some reason a read request is aimed to this register a value of all zeros will be returned. The value of this register is programmed/read in host order.



12.0.3.10.11 MACsec TX Key 1 - LSECTXKEY1 [n] (0x0B030 + 4*n (n=0...3); WO)

Bits	Type	Default	Description
31:0	WO	0x0	LSEC Key 1. Transmit MACsec key of SA 1. n – 0 LSEC Key defines bits 31:0 of the Tx MACsec Key n – 1 LSEC Key defines bits 63:32 of the Tx MACsec Key n – 2 LSEC Key defines bits 95:64 of the Tx MACsec Key n – 3 LSEC Key defines bits 127:96 of the Tx MACsec Key This field is WO for confidentiality protection. For data integrity check, the hash value may read the LSECTXSUM field in the LSECCAP register. If for some reason a read request is aimed to this register a value of all zeros will be returned. The value of this register is programmed/read in host order.

12.0.3.11 MACsec Rx SC Registers

The registers in this section relate to MACsec Receive SC context. There are 4 SC(s) in the receive data path defined as SC0, SC1, SC2 and SC3. The registers below with index n relates to the SC index, while n=0, 1, 2, 3.

12.0.3.11.1 MACsec RX SCI Low - LSECRXSCL (0x0B3D0 + 4*n (n=0...3); RW)

Bits	Type	Default	Description
31:0	RW	0b	MAC Address SecY low. The 4 LS bytes of the MAC address in the SCI field in the incoming packet that are compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set. Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively. The value of this register is programmed/read in host order.

12.0.3.11.2 MACsec RX SCI High - LSECRXSCH (0x0B3E0 + 4*n (n=0...3); RW)

Bits	Type	Default	Description
15:0	RW	0b	MAC Address SecY High. The 2 MS bytes of the MAC address in the SCI field in the incoming packet that are compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set. Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively. The value of this register is programmed/read in host order.
31:16	RW	0b	Port Identifier. The Port Number in the SCI field in the incoming packet that is compared with this field for SCI matching. Comparison result is meaningful only if the SC bit in the TCI header is set. Index n=0, 1, 2, 3 for SCI0, SCI1, SCI2 and SCI3 respectively. The value of this register is programmed/read in host order.



12.0.3.12 MACsec Rx SA Registers


The registers in this section relates to MACsec Receive SA context. There are 8 SA(s) in the receive data path defined as SA0 and SA1... SA7. The registers below with index n relates to the SA index, while n=0, 1... 7. SA0 and SA1 are assigned to SCI0, SA2 and SA3 are assigned to SCI1, SA4 and SA5 are assigned to SCI2 and SA6 and SA7 are assigned to SCI3.

12.0.3.12.1 MACsec RX SA - LSECRXSA (0x0B310 + 4*n (n=0...7); RW)

Bits	Type	Default	Description
1:0	RW	0b	AN – Association Number. This field is compared with the AN field in the TCI field in the incoming packet for match.
2	RW	0b	SA Valid. This bit is set or cleared by the KaY to validate or invalidate the SA.
3	RO/V	0b	Frame received. This bit is cleared when the SA Valid (bit 2) transitions from 0->1, and is set when a frame is received with this SA. When the Frame received bit is set the Retired bit of the other SA of the same SC is set. Note that a single frame reception with the new SA is sufficient to retire the old SA since we assume the Replay Window is 0.
4	RO/V	0b	Retired. When this bit is set the SA is invalid (retired). This bit is cleared when a new SA is configured by the KaY (SA Valid transition to 1). It is set to '1' when a packet is received with the other SA of the same SC. Note that a single frame reception with the new SA is sufficient to retire the old SA since we assume the Replay Window is 0.
31:5	RO	0	Reserved.

12.0.3.12.2 MACsec RX SA PN - LSECRXSAPN (0x0B330 + 4*n (n=0...7); RW)

Bits	Type	Default	Description
31:0	RW	0b	PN – Packet number. This register holds the PN field of the next incoming packet that uses this SA. The PN field in the incoming packet must be greater or equal to the PN register. The PN register is set by KaY at SA creation. It is updated by the hardware for each received packet using this SA to be Received PN + 1. The value of this register is programmed/read in host order.



12.0.3.12.3 MACsec RX Key - LSECRXKEY[n,m] (0x0B350 + 0x10*n (n=0...7) + 4*m (m=0...3); WO)

Bits	Type	Default	Description
31:0	WO	0x0	LSEC Key. Receive MACsec key of SA n, while n=0,1,2...7. m – 0 LSEC Key defines bits 31:0 of the Rx MACsec Key m – 1 LSEC Key defines bits 63:32 of the Rx MACsec Key m – 2 LSEC Key defines bits 95:64 of the Rx MACsec Key m – 3 LSEC Key defines bits 127:96 of the Rx MACsec Key This field is WO for confidentiality protection. For data integrity check, the KaY hash value may read the LSECRXSUM field in the LSECCAP registers. If for some reason a read request is aimed to this register a value of all zeros will be returned. The value of this register is programmed/read in host order.

12.0.3.13 MACsec Tx Port Statistics

These counters are defined by spec as 64bit while implementing only 32 bit in the hardware. The KaY must implement the 64 bit counter in SW by regularly polling the hardware statistic counters. The HW section of the statistics counter is cleared upon read action.

12.0.3.13.1 Tx Untagged Packet Counter - LSECTXUT (0x04300; RC)

This statistic implements the SecyStatsRxUntaggedPkts statistic of the 802.1ae MIB while in non strict mode. In strict mode, this implements the secyStatsRxNoTagPkts statistic of the 802.1ae MIB.

Bits	Type	Default	Description
31:0	RC	0x0	Untagged Packet CNT. Increments for each transmitted packet that is transmitted with the ILSec bit cleared in the packet descriptor while "Enable Tx MACsec" field in the LSECTXCTRL register is either 01b or 10b. The KaY must implement a 64 bit counter. It can do that by reading the LSECTXUT register regularly.

12.0.3.13.2 Encrypted Tx Packets - LSECTXPKTE (0x04304; RC)

Bits	Type	Default	Description
31:0	RC	0x0	Encrypted Packet CNT. Increments for each transmitted packet through the controlled port with E bit set (i.e. confidentiality was prescribed for this packet by SW/FW).

**12.0.3.13.3 Protected Tx Packets - LSECTXPOTP (0x04308; RC)**

Bits	Type	Default	Description
31:0	RC	0x0	Protected Packet CNT. Increments for each transmitted packet through the controlled port with E bit cleared (i.e. integrity only was prescribed for this packet by SW/FW).

12.0.3.13.4 Encrypted Tx Octets - LSECTXOCTE (0x0430C; RC)

Bits	Type	Default	Description
31:0	RC	0x0	Encrypted Octet CNT. Increments for each byte of user data through the controlled port with E bit set (i.e. confidentiality was prescribed for this packet by SW/FW).

12.0.3.13.5 Protected Tx Octets - LSECTXOCTP (0x04310; RC)

Bits	Type	Default	Description
31:0	RC	0x0	Protected Octet CNT. Increments for each byte of user data through the controlled port with E bit (i.e. integrity only was prescribed for this packet by SW/FW).

12.0.3.14 MACsec Rx Port Statistic Counters


These counters are defined by spec as 64bit while implementing only 32 bit in the hardware. The KaY must implement the 64 bit counter in SW by regularly polling the hardware statistic counters.

12.0.3.14.1 MACsec Untagged RX Packet - LSECRXUT (0x04314; RC)

Bits	Type	Default	Description
31:0	RC	0b	Untagged Packet CNT. Increments for each packet received having no tag. Increments only when "Enable Rx MACsec" field in the LSECRXCTRL register is either 01b or 10b.

12.0.3.14.2 MACsec RX Octets Decrypted - LSECRXOCTD (0x0431C; RC)

Bits	Type	Default	Description
31:0	RC	0b	Decrypted Rx Octet CNT. The number of octets of User Data recovered from received frames that were both integrity protected and encrypted. This includes the octets from SecTag to ICV not inclusive. These counts are incremented even if the User Data recovered failed the integrity check or could not be recovered.



12.0.3.14.3 MACsec RX Octets Validated - LSECRXOCTV (0x04320; RC)

Bits	Type	Default	Description
31:0	RC	0b	Validated Rx Octet CNT. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted. This includes the octets from SecTag to ICV not inclusive. These counts are incremented even if the User Data recovered failed the integrity check or could not be recovered.

12.0.3.14.4 MACsec RX Packet with Bad Tag - LSECRXBAD (0x04324; RC)

Bits	Type	Default	Description
31:0	RC	0b	Bad Rx Packet CNT. Number of packets received having an invalid tag.

12.0.3.14.5 MACsec RX Packet No SCI - LSECRXNOSCI (0x04328; RC)

Bits	Type	Default	Description
31:0	RC	0b	No SCI Rx Packet CNT. Number of packets received having unrecognized SCI and dropped due to that condition.

12.0.3.14.6 MACsec RX Packet Unknown SCI count - LSECRXUNSCI (0x432C; RC)

Bits	Type	Default	Description
31:0	RC	0b	Unknown SCI Rx Packet CNT. Number of packets received with an unrecognized SCI but still forwarded to the host.

12.0.3.15 MACsec Rx SC Statistic Counters


12.0.3.15.1 MACsec RX Unchecked Packets - LSECRXUNCH (0x04330; RC)

SW/FW needs to maintain the full sized register.

Bits	Type	Default	Description
31:0	RC	0b	Unchecked Rx Packet CNT. Rx Packet CNT. Number of packets received with MACsec encapsulation (SecTag) while ValidateFrames is disabled (LSECRXCTRL bits 3:2 equal 00b)."

12.0.3.15.2 MACsec RX Delayed Packets - LSECRXDELAY (0x04340 + 4*n (n=0...3); RC)

SW/FW needs to maintain the full sized register.



Bits	Type	Default	Description
31:0	RC	0b	Delayed Rx Packet CNT. Number of packets received and accepted for validation having failed replay-protection and ReplayProtect is false (LSECRXCTRL bit 7 is zero).

12.0.3.15.3 MACsec RX Late Packets - LSECRXLATE (0x04350 + 4*n (n=0...3); RC)

SW/FW needs to maintain the full sized register.

Bits	Type	Default	Description
31:0	RC	0b	Late Rx Packet CNT. Number of packets received and accepted for validation having failed replay-protection and ReplayProtect is true (LSECRXCTRL bit 7 is '1').

12.0.3.16 MACsec Rx SA Statistic Counters

12.0.3.16.1 MACsec RX Packet OK - LSECRXOK[n] (0x04360 + 4*n (n=0...7); RC)

Bits	Type	Default	Description
31:0	RC	0b	OK Rx Packet CNT. Number of packets received that were valid (authenticated) and passed replay protection.

12.0.3.16.2 MACsec Check RX Invalid - LSECRXINV[n] (0x43A0 + 4*n (n=0...7); RC)

Bits	Type	Default	Description
31:0	RC	0b	Invalid Rx Packet CNT. Number of packets received that were not valid (authentication failed) and were forwarded to host.


12.0.3.16.3 MACsec RX Not valid count - LSECRXNV[n] (0x04380 + 4*n [n=0...7]; RC)

Bits	Type	Default	Description
31:0	RC	0b	Invalid Rx Packet CNT. Number of packets received that were not valid (authentication failed) and were dropped.

12.0.3.16.4 MACsec RX Not using SA - LSECRXNUSA[n] (0x043C0 + 4*n (n=0...3); RC)

Bits	Type	Default	Description
31:0	RC	0b	Invalid SA Rx Packet CNT. Number of packets received that were associated with an SA that is not "inUse" (No match on AN or not valid or retired) and were dropped. ¹

1. The implementation maintains one such counter per SC.


**12.0.3.16.5 MACsec RX Unused SA - LSECRXUNSA[n] (0x043D0 + 4*n (n=0...3); RC)**


Bits	Type	Default	Description
31:0	RC	0b	Invalid SA Rx Packet CNT. Number of packets received that were associated with an SA that is not "inUse" (No match on AN or not valid or retired) and where forwarded to host. ¹

1. The implementation maintains one such counter per SC.



Note: This page intentionally left blank.







13.0 Reference Schematics

A reference schematic is available as a separate document; please check Developer.intel.com or contact your Intel representative.

§ §






14.0 Schematic and Board Layout Checklists

Design review checklists for the 82579 are available at:

<http://www.developer.intel.com>

§ §






15.0 Models

Contact your Intel Representative for access.

§ §





16 Design Considerations and Guidelines (Non-Mobile Designs)

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel® 82579 Physical Layer Device (PHY) shown in [Figure 16-1](#). Its bus master capabilities enable the component to process high-level commands and perform multiple operations, which lowers processor use by off loading communication tasks from the processor.

The PCH, which hereinafter refers to the integrated MAC within the PCH, supports the SMBus interface for manageability while in an Sx state and PCI Express* (PCIe*) for 10/100/1000 Mbps traffic in an S0 state.

Note: Design guidance is available for non-mobile designs (this chapter) and mobile designs (following chapter). Be sure you are using the proper information for your design type.

Note: The 82579 PCIe interface is not PCIe compliant. It operates at half of the PCI Express* (PCIe*) Specification v1.0 (2.5 GT/s) speed. In this chapter, the term “PCIe-based” is interchangeable with “PCIe.” There are no design layout differences between normal PCIe and the 82579 PCIe-based interface.

The 82579 PHY interfaces with the integrated MAC through two interfaces: PCIe and SMBus. In SMBus mode, the link speed is reduced to 10 Mbps. The PCIe interface incorporates two aspects: a PCIe-based SerDes (electrically) and a custom logic protocol for messaging between the integrated MAC and the PHY.

Note: Gigabit Ethernet requires an SPI Flash to host firmware and does not work without an SPI Flash on board.

The integrated MAC supports multi-speed operation (10/100/1000 Mbps). The integrated MAC also operates in full-duplex at all supported speeds or half-duplex at 10/100 Mbps as well as adhering to the IEEE 802.3x Flow Control Specification.

Note: References to the AUX power rail means the power rail is available in all power states including G3 to S5 transitions and Sx states with Wake on LAN (WoL) enabled. For example, V3P3_AUX in this chapter refers to a rail that is powered under the conditions previously mentioned.

Figure 16-1. PCH/PHY Interface Connections






Table 16-1. SMBus Data Signals on the PCH

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

Table 16-2. PCIe Data Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Data	PETp PETn	PETp PETn	PCIe transmit pair
Data	PERp PERn	PERp PERn	PCIe receive pair

Notes:

1. Can be connected to any PCIe port on the integrated MAC. The appropriate NVM descriptor soft strap (PCHSTRP9) should define which PCIe port is configured as GbE LAN.

Table 16-3. Clock and Reset Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Clock	SMB_CLK	SML0_CLK	SMBus clock
Clock	PE_CLKP PE_CLKN	CLKOUT_PCIE[7:0]_P ¹ CLKOUT_PCIE[7:0]_N ¹	PCIe* clock
Clock ²	CLK_REQ_N	Not Connected	PCIe clock request
Reset	PE_RST_N	PLTRST# ³	PCIe reset

Notes:

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

16.1 PHY Overview


The PHY is a single port compact component designed for 10/100/1000 Mbps operation. It enables a single port Gigabit Ethernet (GbE) implementation in a very small area, easing routing constraints from the PCH chipset to the PHY.

The PHY provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3ab, 802.3u, and 802.3i, respectively).

16.1.1 PHY Interconnects

The main interfaces for either PHY are PCIe and SMBus on the host side and Media Dependent Interface (MDI) on the link side. Transmit traffic is received from the PCH as either PCIe or SMBus packets on the host interconnect and transmitted as Ethernet packets on the MDI link. Receive traffic arrives as Ethernet packets on the MDI link and transferred to the PCH through either the PCIe or SMBus interconnects.

The PHY switches the in-band traffic automatically between PCIe and SMBus based on platform reset. The transition protocol is done through SMBus. The PCIe interface is powered down when the Ethernet link is running in an Sx state.



16.1.2 PCIe-Based Interface

A high-speed SerDes interface that uses PCIe electrical signaling at half speed while utilizing a custom logical protocol for active state operation mode.

Note: PCIe validation tools cannot be used for electrical validation of this interface; however, PCIe layout rules apply for on-board routing.

16.1.2.0.1 PCIe Interface Signals

The signals used to connect between the PCH and the PHY in this mode are:

- Serial differential pair running at 1.25 Gb/s for Rx.
- Serial differential pair running at 1.25 Gb/s for Tx.
- 100-MHz differential clock input to the PHY is generated by the PCH.
- Power and clock good indication to the PHY PE_RST_N.
- Clock control through CLK_REQ_N (refer to [Table 16-3](#)). This PHY output should be pulled up with a 10-k Ω resistor connected to 3.3-V DC AUX power (present in G3 to S5).

16.1.2.0.2 PCIe Operation and Channel Behavior

The PHY only runs at 1250 Mbps speed, which is 1/2 of the Gen 1 2.5 Gb/s PCIe frequency. Each of the PCIe root ports in the PCH has the ability to run at 1250 Mbps. Configuring a PCH PCIe port that is attached to a PCIe Intel PHY only device is pre-loaded from the GbE region of the NVM. The selected port adjusts the transmitter to run at 1/2 the Gen 1 PCIe speed and does not need to be PCIe compliant.

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

16.1.2.0.3 PCIe Connectivity

The PHY transmit/receive pins are output/input signals and are connected to the PCH as listed in [Table 16-1](#) through [Table 16-3](#).

16.1.2.0.4 PCIe Reference Clock

The PCIe Interface uses a 100-MHz differential reference clock, denoted PE_CLKP and PE_CLKN. This signal is typically generated on the platform and routed to the PCIe port.

The frequency tolerance for the PCIe reference clock is ± 300 ppm.

16.1.3 SMBus Interface

SMBus is a low speed (100 kHz/400 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the PHY and the PCH when the platform is in a low power state (Sx). The interface is also used to enable the PCH to configure the PHY as well as passing in-band information between them.

The SMBus uses two primary signals: SMBCLK and SMBDATA, to communicate. Both of these signals float high with board-level 2.2 k $\Omega \pm 5\%$ pull-up resistors.



The SMBus specification has defined various types of message protocols composed of individual bytes. For more details about SMBus, see the SMBus specification.

16.1.3.0.1 SMBus Connectivity

Table 16-1 through Table 16-3 list the relationship between PHY SMBus pins to the PCH LAN SMBus pins.

Note: The 82579 SMBus signals (SMB_DATA and SMB_CLK) cannot be connected to any other devices other than the integrated MAC. Connect the SMB_DATA and SMB_CLK pins to the integrated MAC SML0DATA and SML0CLK pins, respectively.

16.1.4 PCIe and SMBus Modes

In GbE operation, PCIe is used to transmit and receive data and for MDIO status and control. The PHY automatically switches the in-band traffic between PCIe and SMBus based on the platform power state. The table below lists the operating modes of PCIe and SMBus.

The 82579 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

System / Intel Management Engine State	PHY	
	SMBus	PCIe
S0 and PHY Power Down	Not used	Electrical Idle (EI)
S0 and Idle or Link Disconnect	Not used	EI
S0 and Link in Low Power Idle (LPI)	Not used	EI
S0 and active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

16.1.5 Transitions between PCIe and SMBus Interfaces

16.1.5.0.1 Switching from SMBus to PCIe


Communication between the integrated MAC and the PHY is done through the SMBus each time the system is in a low power state (Sx). The integrated MAC/PHY interface is needed while the Manageability Engine (ME) is still active to transfer traffic, configuration, control and status or to enable host wake up from the PHY.

Possible states for activity over the SMBus:

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

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

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



16.1.5.0.2 Switching from PCIe to SMBus

The communication between the integrated MAC and the PHY is done through PCIe each time the platform is in active power state (S0). Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME.

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

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

16.2 Platform LAN Design Guidelines

These sections provide recommendations for selecting components and connecting special pins. For GbE designs, the main elements are the PCH chipset, 82579 PHY, a magnetics module and RJ-45 connector, a GbE region NVM (Non Volatile Memory) image, and a clock source.

16.2.1 General Design Considerations for PHYs

Sound engineering practices must be followed with respect to unused inputs by terminating them with pull-up or pull-down resistors, unless otherwise specified in a datasheet, design guide or reference schematic. Pull-up or pull-down resistors must not be attached to any balls identified as “No Connect.” These devices might have special test modes that could be entered unintentionally.

16.2.1.0.1 Clock Source

All designs require a 25-MHz clock source. The PHY uses the 25-MHz source to generate clocks up to 125 MHz and 1.25 GHz for both the PHY circuits and the PCIe interface. For optimum results with lowest cost, a 25-MHz parallel resonant crystal can be used along with the appropriate load capacitors at the XTAL_OUT (X2) and XTAL_IN (X1) leads. The frequency tolerance of the timing device should equal 30 ppm or better. Further detail is found in [Section 16.18](#) and [Section 16.34](#).

Note: XTAL_OUT and XTAL_IN are the signal names for the PHY.

There are three steps to crystal qualification:

1. Verify that the vendor’s published specifications in the component datasheet meet the required conditions for frequency, frequency tolerance, temperature, oscillation mode and load capacitance as specified in the respective datasheet.
2. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems.
3. Independently measure the component’s electrical parameters in real systems. Measure frequency at a test output to avoid test probe loading effects at the PHY. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications. For crystals, it is also important to examine startup behavior while varying system voltage and temperature.



16.2.1.0.2 Magnetics Module

The magnetics module has a critical effect on overall IEEE and emissions conformance. The device should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Carefully qualifying new magnetics modules prevents problems that might arise because of interactions with other components or the printed circuit board itself.

The steps involved in magnetics module qualification are similar to those for crystal qualification:

1. Verify that the vendor's published specifications in the component datasheet meet or exceed the required IEEE specifications.
2. Independently measure the component's electrical parameters on the test bench, checking samples from multiple lots. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications.
3. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems. Vary temperature and voltage while performing system level tests.

Magnetics modules for 1000BASE-T Ethernet as used by the PHY only are similar to those designed solely for 10/100 Mbps, except that there are four differential signal pairs instead of two. Refer to the 82579 datasheet for specific electrical requirements that the magnetics need to meet.

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

Table 16-4. Magnetic Modules and Manufacturers

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

16.2.1.0.3 Criteria for Integrated Magnetics Electrical Qualification

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

**Table 16-5. Integrated Magnetics Recommended Qualification Criteria**

Open Circuit Inductance (OCL)	w/8 mA DC bias; at 25C	400uH Min
	w/8 mA DC bias; at 0C to 70C	350uH Min
Insertion Loss	100 kHz through 999 kHz	1dB Max
	1.0 MHz through 60.0 MHz	0.6dB Max
	60.1 MHz through 80.0 MHz	0.8dB Max
	80.1 MHz through 100.0 MHz	1.0dB Max
	100.1 MHz through 125.0 MHz	2.4dB Max
Return Loss	1.0 MHz through 40.0 MHz	18.0 dB Min
	40.1 MHz through 100.0 MHz When reference impedance is 85 Ohms, 100 Ohms, and 115 Ohms. Note that R.L. values may vary with MDI trace lengths. The LAN magnetics may need to be measured in the platform where it will be used.	$12 - 20 * \text{LOG} (\text{Freq in MHz} / 80) \text{ dB Min}$
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	$-50.3 + (8.8 * (\text{Freq in MHz} / 30)) \text{ dB Max}$
	30.0 MHz through 250.0 MHz	$-(26 - (16.8 * (\text{LOG}(\text{Freq in MHz} / 250 \text{ MHz})))) \text{ dB Max}$
Crosstalk Isolation Integrated Modules (Proposed)	250.1 MHz through 375.0 MHz	-26.0 dB Max
	1.0 MHz through 10 MHz	$-50.8 + (8.8 * (\text{Freq in MHz} / 10)) \text{ dB Max}$
Diff to CMR	10.0 MHz through 100.0 MHz	$-(26 - (16.8 * (\text{LOG}(\text{Freq in MHz} / 100 \text{ MHz})))) \text{ dB Max}$
	100 MHz through 375.0 MHz	-26.0 dB Max
CM to CMR	1 MHz through 29.9 MHz	$-40.2 + (5.3 * ((\text{Freq in MHz} / 30))) \text{ dB Max}$
	30.0 MHz through 500 MHz	$-(22 - (14 * (\text{LOG}((\text{Freq in MHz} / 250)))) \text{ dB Max}$
Hi-Voltage Isolation	1 MHz through 270 MHz	$-57 + (38 * ((\text{Freq in MHz} / 270))) \text{ dB Max}$
	270.1 MHz through 300 MHz	$-17 - 2 * ((300 - (\text{Freq in MHz} / 30)) \text{ dB Max}$
	300.1 MHz through 500 MHz	-17 dB Max
	1500 Vrms at 50 or 60 Hz for 60 sec. or: 2250 Vdc for 60 seconds	Minimum

16.2.2 NVM for PHY Implementations

The LAN only supports an SPI Flash, which is connected to the PCH. Several words of the NVM are accessed automatically by the device after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the NVM space is available to software for storing the MAC address, serial numbers, and additional information. More details may be obtained from the 82579 *Datasheet*.

Intel has an MS-DOS* software utility called EEUupdate that is used to program the SPI Flash images in development or production line environments. A copy of this program can be obtained through your Intel representative.

16.2.3 LED

The PHY has three LED outputs that can be configured via the NVM. The hardware configuration is shown in [Figure 16-2](#).


Refer to the 82579 Reference Schematic for default LED color based on reference design.

Refer to the 82579 *Intel® 82579 Datasheet* for details regarding the programming of the LED's and the various modes. The default values for the PHY (based on the LED NVM setting--word 0x18 of the LAN region) are listed in the table below:

Table 16-6. LED Default Values

LED	Mode	Color	Blink	Polarity
LED0	Link Up/Activity	Green	200 ms on/200 ms off	Active low
LED1	Link 1000	Yellow	No	Active low
LED2	Link 100	Green	No	Active low

Figure 16-2. LED Hardware Configuration



16.2.3.0.1 RBIAS

RBIAS requires external resistor connection to bias the internal analog section of the device. The input is sensitive to the resistor value. Resistors of 1% tolerance must be used. Connect RBIAS through a 3.01 k Ω 1% pull-down resistor to ground and then place it no more than one inch away from the PHY.

16.2.3.0.2 LAN Disable

The PHY enters a power-down state when the LAN_DISABLE_N pin is asserted low. Exiting this mode requires setting the LAN_DISABLE_N pin to a logic one. Connect LAN_DISABLE_N to LAN_PHY_PWR_CTRL in the PCH.

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

16.2.4.0.1 Overview

This section provides general information about ePAD and SMT assemblies. Chip packages have exposed die pads on the bottom of each package to provide electrical interconnections with the printed circuit board. These ePADs also provide excellent thermal performance through efficient heat paths to the PCB.

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

to implement the ePAD package design. Each fab and assembly house might need to tune the land pattern/stencil and create a solution that best suits their methodology and process.

16.2.4.0.2 PCB Design Requirements

In order to maximize both heat removal and electrical performance, a land pattern must be incorporated on the PCB within the footprint of the package corresponding to the exposed metal pad or exposed heat slug of the package as shown in the following figures. Refer to the specific product datasheet for actual dimensions.

Note: Due to the 82579 package size, a via-in-pad configuration must be used [Figure 16-3](#) and [Figure 16-4](#) are general guidelines see [Figure 16-5](#) for 82579-specific via-in-pad thermal pattern recommendations.

Figure 16-3. Typical ePAD Land Pattern






Figure 16-4. Typical Thermal Pad and Via Recommendations



Note: Encroached and uncapped via configurations have voids less than the maximum allowable void percentage. Uncapped via provides a path for trapped air to escape during the reflow soldering process.

Note: Secondary side solder bumps might be seen in an uncapped via design. This needs to be considered when placing components on the opposite side of the PHY.

Figure 16-5. Recommended Thermal Via Patterns for the 82579



16.2.4.0.3 Board Mounting Guidelines

The following are general recommendations for mounting a QFN-48 device on the PCB. This should serve as the starting point in assembly process development and it is recommended that the process should be developed based on past experience in mounting standard, non-thermally/electrically enhanced packages.

16.2.4.0.4 Stencil Design


For maximum thermal/electrical performance, it is required that the exposed pad/slug on the package be soldered to the land pattern on the PCB. This can be achieved by applying solder paste on both the pattern for lead attachment as well as on the land pattern for the exposed pad. While for standard (non-thermally/ -electrically enhanced) lead-frame based packages the stencil thickness depends on the lead pitch and package co-planarity, the package standoff must also be considered for the thermally/ electrically enhanced packages to determine the stencil thickness. In this case, a stencil foil thickness in the range of 5 - 6 mils (or 0.127—0.152 mm) is recommended; likely or practically, a choice of either 5 mils or 6 mils. Tolerance wise, it should not be worse than ± 0.5 mil.

Note: Industry specialists typically use ± 0.1 -mil tolerance on stencil for its feasible precision.

The aperture openings should be the same as the solder mask openings on the land pattern. Since a large stencil opening may result in poor release, the aperture opening should be subdivided into an array of smaller openings, similar to the thermal land pattern shown in the figure below.

Note: Refer to the specific product datasheet for actual dimensions.


Figure 16-6. Stencil Design Recommendation



16.2.4.0.5 Assembly Process Flow

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


Figure 16-7. Assembly Flow



16.2.4.0.6 Reflow Guidelines

The typical reflow profile consists of four sections. In the preheat section, the PCB assembly should be preheated at the rate of 1 to 2 °C/sec to start the solvent evaporation and to avoid thermal shock. The assembly should then be thermally soaked for 60 to 120 seconds to remove solder paste volatiles and for activation of flux. The reflow section of the profile, the time above liquidus should be between 45 to 60 seconds with a peak temperature in the range of 245 to 250 °C, and the duration at the peak should not exceed 30 seconds. Finally, the assembly should undergo cool down in the fourth section of the profile. A typical profile band is provided in the figure below, in which 220 °C is referred to as an approximation of the liquidus point. The actual profile parameters depend upon the solder paste used and specific recommendations from the solder paste manufacturers should be followed.

Figure 16-8. Typical Profile Band



1. Preheat: 125 °C - 220 °C, 150 - 210 s at 0.4 K/s to 1.0 K/s
2. Time at T > 220 °C: 60 - 90 s
3. Peak Temperature: 245-250 °C
4. Peak time: 10 - 30 s
5. Cooling rate: ≤ 6 K/s
6. Time from 25 °C to Peak: 240 - 360 s
7. Intel recommends a maximum solder void of 50% after reflow.

Note: Contact your Intel representative for any designs unable to meet the recommended guidance for E-pad implementation.


16.3 PCH – SMBus/PCIe LOM Design Guidelines

This section contains guidelines on how to implement a PCH/PHY single solution on a system motherboard. It should not be treated as a specification, and the system designer must ensure through simulations or other techniques that the system meets the specified timings. The following are guidelines for both PCH SMBus and PCIe interfaces. Note that PCIe is only applicable to the PHY.

The SMBus/PCIe Interface can be configured in as shown [Figure 16-9](#).

Refer to [Section 16.6](#) for PCI Express Routing Guidelines.

Figure 16-9. Single Solution Interconnect


Notes:


1. Any free PCIe ports (Ports 1-8) can be used to hook up to the 82579 PCIe Interface.
2. Any CLKOUT_SRC[7:0] and SRC[7:0]CLKRQ# can be used to connect to PE_CLK and CLK_REQ_N on the 82579.
3. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.
4. For latest PHY schematic connection recommendations, refer to the 82579 reference schematic, which is available on CDI.

16.4 SMBus Design Considerations

No single SMBus design solution works for all platforms. Designers must consider the total bus capacitance and device capabilities when designing SMBus segments. Routing SMBus to the PCI slots makes the design process even more challenging since they add so much capacitance to the bus. This extra capacitance has a large affect on the bus time constant which in turn affects the bus rise and fall times.

Primary considerations in the design process are:

- Amount of $V_{CC_SUS3_3}$ current available, that is, minimizing load of $V_{CC_SUS3_3}$.
- The pull-up resistor size for the SMBus data and clock signals is dependent on the bus load (this includes all device leakage currents). Generally the SMBus device that can sink the least amount of current is the limiting agent on how small the resistor can be. The pull-up resistor cannot be made so large that the bus time constant (Resistance X Capacitance) does not meet the SMBus rise and time specification.
- The maximum bus capacitance that a physical segment can reach is 400 pF.
- SMBus devices that can operate in STR must be powered by the $V_{CC_SUS3_3}$ supply.
- It is recommended that I²C (Inter-Integrated Circuit) devices be powered by the V_{CC_core} supply. During an SMBus transaction in which the device is sending



information to the integrated MAC, the device may not release the SMBus if the integrated MAC receives an asynchronous reset. V_{CC_core} is used to enable the BIOS to reset the device if necessary. SMBus 2.0- compliant devices have a timeout capability that makes them in-susceptible to this I²C issue, enabling flexibility in choosing a voltage supply.

- No other devices (except the integrated MAC and pull-up resistors) should be connected to the SMBus that connects to the PHY.
- **For system LAN on motherboard (LOM) designs:** The traces should be less than 70 inches for stripline and less than 100 inches for Microstrip. These numbers depend on the stackup, dielectric layer thickness, and trace width. The total capacitance on the trace and input buffers should be under 400 pF.
- **For system LAN on daughterboard designs:** Being conservative, the traces should be less than 7 inches for stripline designs and less than 10 inches for Microstrip designs. The lengths depend on the stackup, dielectric layer thickness, and trace width. Longer traces can be used as long as the total capacitance on the trace and input buffers is under 30 pF.

Note: Refer to [Section 16.1.3](#) for additional SMBus design considerations.

16.5 General Layout Guidelines

PHY interface signals must be carefully routed on the motherboard to meet the timing and signal quality requirements of their respective interface specifications. The following are some general guidelines that should be followed in designing a LAN solution. It is recommended that the board designer simulate the board routing to verify that the specifications are met for flight times and skews due to trace mismatch and crosstalk.

16.6 Layout Considerations

Critical signal traces should be kept as short as possible to decrease the likelihood of effects by high frequency noise of other signals, including noise carried on power and ground planes. This can also reduce capacitive loading.

Since the transmission line medium extends onto the printed circuit board, layout and routing of differential signal pairs must be done carefully.

Designing for GbE (1000BASE-T) operation is very similar to designing for 10/100 Mbps. For the PHY, system level tests should be performed at all three speeds.

16.7 Guidelines for Component Placement

Component placement can affect signal quality, emissions, and component operating temperature. Careful component placement can:

Decrease potential problems directly related to electromagnetic interference (EMI), which could cause failure to meet applicable government test specifications. In this case, place the PHY more than one inch from the edge of the board.

Simplify the task of routing traces. To some extent, component orientation affects the complexity of trace routing. The overall objective is to minimize turns and crossovers between traces.

16.7.1 PHY Placement Recommendations

Minimizing the amount of space needed for the PHY is important because other interfaces compete for physical space on a motherboard near the connector. The PHY circuits need to be as close as possible to the connector.

The figure below illustrates some basic placement distance guidelines. To simplify the diagram, it shows only two differential pairs, but the layout can be generalized for a GbE system with four analog pairs. The ideal placement for the PHY (LAN silicon) is approximately one inch behind the magnetics module.

While it is generally a good idea to minimize lengths and distances, this figure also illustrates the need to keep the PHY away from the edge of the board and the magnetics module for best EMI performance.





Figure 16-10. LAN Device Placement: At Least One Inch from Chassis Openings or Unshielded Connectors--Non-Mobile

Figure 16-11.

**Note:**

* this distance is variable and follows the general guidelines.

The PHY, referred to as "LAN Device" in the above figure, must be at least one inch from the I/O back panel. To help reduce EMI, the following recommendations should be followed:

- Minimize the length of the MDI interface. See detail in table below: MDI Routing Summary
- Place the MDI traces no closer than 0.5 inch (1.3 cm) from the board edge.
- The 82579 PHY must be placed greater than 1" away from any hole to the outside of the chassis larger than 0.125 inches (125 mils) The larger the hole the higher the probability the EMI and ESD immunity will be negatively affected.
- The 82579 PHY should be placed greater than 250mils from the board edge.
- If the connector or integrated magnetics module is not shielded, the 82579 should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the 82579 closer than one inch to Unshielded magnetics or connectors will increase the probability of failed EMI and common mode noise. If the LAN switch is too far away it will negatively affect IEEE return loss performance.
- The RBIAS trace length must be less than 1"
- Place the crystal less than 0.75 inch (1.9 cm) from the PHY.

Figure 16-12. PLC Placement: At Least One Inch from I/O Backplane





Figure 16-13. Effect of LAN Device Placed Too Close To a RJ-45 Connector or Chassis Opening



16.8 MDI Differential-Pair Trace Routing for LAN Design


Trace routing considerations are important to minimize the effects of crosstalk and propagation delays on sections of the board where high-speed signals exist. Signal traces should be kept as short as possible to decrease interference from other signals, including those propagated through power and ground planes.

16.9 Signal Trace Geometry

One of the key factors in controlling trace EMI radiation are the trace length and the ratio of trace-width to trace-height above the reference plane. To minimize trace inductance, high-speed signals and signal layers that are close to a reference or power plane should be as short and wide as practical. Ideally, the trace-width to trace-height above the ground plane ratio is between 1:1 and 3:1. To maintain trace impedance, the width of the trace should be modified when changing from one board layer to another if the two layers are not equidistant from the neighboring planes.

Each pair of signals should have a differential impedance of $100\ \Omega \pm 15\%$.

A set of trace length calculation tools are available from Intel (via the Intel Business Link (IBL)) to aid with MDI topology design.



When performing a board layout, the automatic router feature of the CAD tool must not route the differential pairs without intervention. In most cases, the differential pairs will require manual routing.

Note: Measuring trace impedance for layout designs targeting 100 Ω often results in lower actual impedance due to over-etching. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of 105 Ω to 110 Ω should compensate for over-etching.

It is necessary to compensate for trace-to-trace edge coupling, which can lower the differential impedance by up to 10 Ω , when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 16-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Signal group	MDI_PLUS[0:3] MDI_MINUS[0:3]		
Microstrip/stripline uncoupled single-ended impedance specification	50 $\Omega \pm 10\%$		
Microstrip/stripline uncoupled differential impedance specification	100 $\Omega \pm 15\%$		2,3
Microstrip nominal trace width	Design dependent	Design dependent	4
Microstrip nominal trace space	Design dependent	Design dependent	3,5
Microstrip/stripline trace length	4 in (102 mm)		6,7
Microstrip pair-to-pair space (edge-to-edge)	≥ 7 times the thickness of the thinnest adjacent dielectric layer		Figure 16-14
Stripline pair-to-pair space (edge-to-edge)	≥ 6 times the thickness of the thinnest adjacent dielectric layer		
Microstrip bus-to-bus spacing	≥ 7 times the thickness of the thinnest adjacent dielectric layer		
Stripline bus-to-bus spacing	≥ 6 times the thickness of the thinnest adjacent dielectric layer		

Notes:

1. Pair-to-pair spacing ≥ 3 times the dielectric thickness for a maximum distance of 500 mils from the pin.
2. Board designers should ideally target 100 $\Omega \pm 15\%$. If it's not feasible (due to board stack-up) it is recommended that board designers use a 95 $\Omega \pm 10\%$ target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95 Ω . The $\pm 10\%$ tolerance is provided to allow for board manufacturing process variations and not lower target impedances. The minimum value of impedance cannot be lower than 85 Ω .
3. Simulation shows 80 Ω differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90 Ω .
4. Stripline is NOT recommended due to thinner more resistive signal layers.
5. Use a minimum of 21 mil (0.533 mm) pair-to-pair spacing for board designs that use the CRB design stack-up. Using dielectrics that are thicker than the CRB stack-up might require larger pair-to-pair spacing.




Table 16-8. Maximum Trace Lengths Based on Trace Geometry and Board Stack-Up

Dielectric Thickness (mils)	Dielectric Constant (DK) at 1 MHz	Width / Space / Width (mils)	Pair-to-Pair Space (mils)	Nominal Impedance (Ohms)	Impedance Tolerance ($\pm$ %)	Maximum Trace Length (inches) ¹
2.7	4.05	4/10/4	19	95 ²	17 ²	3.5
2.7	4.05	4/10/4	19	95 ²	15 ²	4
2.7	4.05	4/10/4	19	95	10	5
3.3	4.1	4.2/9/4.2	23	100 ²	17 ²	4
3.3	4.1	4.2/9/4.2	23	100	15	4.6
3.3	4.1	4.2/9/4.2	23	100	10	6
4	4.2	5/9/5	28	100 ²	17 ²	4.5
4	4.2	5/9/5	28	100	15	5.3
4	4.2	5/9/5	28	100	10	7

Notes:

1. Longer MDI trace lengths may be achievable, but may make it more difficult to achieve IEEE conformance. Simulations have shown deviations are possible if traces are kept short. Longer traces are possible; use cost considerations and stack-up tolerance for differential pairs to determine length requirements.
2. Deviations from 100 Ω nominal and/or tolerances greater than 15% decrease the maximum length for IEEE conformance.

Note:


Use the MDI Differential Trace Calculator to determine the maximum MDI trace length for your trace geometry and board stack-up. Contact your Intel representative for access.

The following factors can limit the maximum MDI differential trace lengths for IEEE conformance:

- Dielectric thickness
- Dielectric constant
- Nominal differential trace impedance
- Trace impedance tolerance
- Copper trace losses
- Additional devices, such as switches, in the MDI path may impact IEEE conformance.

Board geometry should also be factored in when setting trace length.

Figure 16-14. MDI Trace Geometry



16.10 Trace Length and Symmetry

The differential traces should be equal in total length to within 10 mils (0.254 mm) per segment within each pair and as symmetrical as possible. Asymmetrical and unequal length traces in the differential pairs contribute to common mode noise. If a choice has to be made between matching lengths and fixing symmetry, more emphasis should be placed on fixing symmetry. Common mode noise can degrade the receive circuit's performance and contribute to radiated emissions.


The intra-pair length matching on the pairs must be within 10 mils on a segment by segment basis. An MDI segment is defined as any trace within the same layer. For example, transitioning from one layer to another through a via is considered as two separate MDI segments.

The end to end total trace lengths within each differential pair must match as shown in the figure titled MDI Trace Geometry. The end to end trace length is defined as the total MDI length from one component to another regardless of layer transitions.

The pair to pair length matching is not as critical as the intra-pair length matching but it should be within 2 inches.

When using Microstrip, the MDI traces should be at least 7x the thinnest adjacent dielectric away from the edge of an adjacent reference plane. When using stripline, the MDI traces should be at least 6x the thinnest adjacent dielectric away from the edge of an adjacent reference plane.

Figure 16-15. MDI Differential Trace Geometry



16.11 Impedance Discontinuities

Impedance discontinuities cause unwanted signal reflections. Vias (signal through holes) and other transmission line irregularities should be minimized. If vias must be used, a reasonable budget is four or less per differential trace. Unused pads and stub traces should also be avoided.


16.12 Reducing Circuit Inductance

Traces should be routed over a continuous reference plane with no interruptions. If there are vacant areas on a reference or power plane, the signal conductors should not cross the vacant area. This causes impedance mismatches and associated radiated noise levels.

16.13 Signal Isolation

To maintain best signal integrity, keep digital signals far away from the analog traces. Also, keep the MDI traces away from the edge of an adjacent reference plane by a distance that is at least 7x the thickness of the thinnest adjacent dielectric layer (7x when using Microstrip; 6x when using stripline). If digital signals on other board layers cannot be separated by a ground plane, they should be routed perpendicular to the differential pairs. If there is another LAN controller on the board, the differential pairs from that circuit must be kept away.

Other rules to follow for signal isolation include:



- Separate and group signals by function on separate layers if possible. If possible, maintain at least a gap of 30 mils between all differential pairs (Ethernet) and other nets, but group associated differential pairs together.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.
- Isolate I/O signals from high-speed signals to minimize crosstalk, which can increase EMI emission and susceptibility to EMI from other signals.
- Avoid routing high-speed LAN traces near other high-frequency signals associated with a video controller, cache controller, processor, switching power supplies, or other similar devices.

16.14 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and power inputs bypassed to signal return. This will significantly reduce EMI radiation.

The following guidelines help reduce circuit inductance in both backplanes and motherboards:

- Route traces over a continuous plane with no interruptions. Do not route over a split power or ground plane. If there are vacant areas on a ground or power plane, avoid routing signals over the vacant area. This will increase inductance and EMI radiation levels.
- All ground vias should be connected to every ground plane; and every power via, to all power planes at equal potential. This helps reduce circuit inductance.
- Physically locate grounds between a signal path and its return. This will minimize the loop area.
- Split the ground plane beneath a magnetics module. The RJ-45 connector side of the transformer module should have chassis ground beneath it.

Caution: DO NOT do this, if the RJ-45 connector has integrated USB.

Note: All impedance-controlled signals should be routed in reference to a solid plane. If there are plane splits on a reference layer and the signal traces cross those splits then stitching capacitors should be used within 40 mils of where the crossing occurs. See Figure 1-13.

If signals transition from one reference layer to another reference layer then stitching capacitors or connecting vias should be used based on the following:

If the transition is from power-referenced layer to a ground-referenced layer or from one voltage-power referenced layer to a different voltage-power referenced layer, then stitching capacitors should be used within 40 mils of the transition.

If the transition is from one ground-referenced layer to another ground-referenced layer or is from a power-referenced layer to the same net power-referenced layer, then connecting vias should be used within 40 mils of the transition.

Figure 16-16.Trace Transitioning Layers and crossing Plane Splits




Figure 16-17.Via Connecting GND to GND





Figure 16-18. Stitching Capacitor between Vias Connecting GND to GND



16.15 Traces for Decoupling Capacitors

Traces between decoupling and I/O filter capacitors should be as short and wide as practical. Long and thin traces are more inductive and reduce the intended effect of decoupling capacitors. Also, for similar reasons, traces to I/O signals and signal terminations should be as short as possible. Vias to the decoupling capacitors should be sufficiently large in diameter to decrease series inductance. Refer to the Power Delivery section for the PHY in regards to actual placement requirements of the capacitors.

16.16 Ground Planes under a Magnetics Module

The magnetics module chassis or output ground (secondary side of transformer) should be separated from the digital or input ground (primary side) by a physical separation of 100 mils minimum. Splitting the ground planes beneath the transformer minimizes noise coupling between the primary and secondary sides of the transformer and between the adjacent coils in the magnetics. This arrangement also improves the common mode choke functionality of magnetics module.

Caution: DO NOT do this if the RJ-45 connector has integrated USB.

The figure below illustrates the split plane layout for a discrete magnetics module. Capacitors are used to interconnect chassis ground and signal ground.

Figure 16-19 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.

Figure 16-19. Ideal Ground Split Implementation





Table 16-9. Capacitor Stuffing Option Recommended Values

Capacitors	Value
C3, C4	4.7 μ F or 10 μ F
C1, C2, C5, C6	470 pF to 0.1 μ F


The placement of C1 through C6 may also differ for each board design (in other words, not all of the capacitors may need to be populated). Also, the capacitors may not be needed on both sides of the magnetics module.

Note: If using an integrated magnetics module without USB, provide a separate chassis ground “island” to ground around the RJ-45 connector. The split in the ground plane should be at least 20 mils wide.



Some integrated magnetics modules/RJ-45 connectors have recently incorporated the USB into the device. For this type of magnetics module, a chassis ground moat may not be feasible due to the digital ground required for the USB pins and their placement relative to the magnetics pins. Thus, a continuous digital ground without any moats or splits must be used. [Figure 16-20](#) provides an example of this.

Figure 16-20. Ground Layout with USB



16.17 Light Emitting Diodes

The device has three high-current outputs to directly drive LEDs for link, activity and speed indication. Since LEDs are likely to be integral to a magnetics module, take care to route the LED traces away from potential sources of EMI noise. In some cases, it may be desirable to attach filter capacitors.


LAN LED traces should be placed at least 6x (side by side separation) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

LAN LED traces should be placed at least 7x (broadside coupling) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

16.18 Frequency Control Device Design Considerations

This section provides information regarding frequency control devices, including crystals and oscillators, for use with all Intel Ethernet controllers. Several suitable frequency control devices are available; none of which present any unusual challenges in selection. The concepts documented within this section are applicable to other data communication circuits, including the PHY.

The PHY contains amplifiers that form the basis for feedback oscillators when they are used with the specific external components. These oscillator circuits, which are both economical and reliable, are described in more detail in [Section 16.22](#).



The chosen frequency control device vendor should be consulted early in the design cycle. Crystal and oscillator manufacturers familiar with networking equipment clock requirements may provide assistance in selecting an optimum, low-cost solution.

Several types of third-party frequency reference components are currently available. Descriptions of each type follow in subsequent sections. They are also listed in order of preference.

16.19 Crystals and Oscillators

Clock sources should not be placed near I/O ports or board edges. Radiation from these devices may be coupled onto the I/O ports or out of the system chassis. Crystals should also be kept away from the Ethernet magnetics module to prevent interference.

Crystal and load capacitors should be placed on the printed circuit boards as close to the PHY as possible, which is within 1.0 inch. Traces from XTAL_IN (X1) and XTAL_OUT (X2) should be routed as symmetrically as possible. Do not route X1 and X2 as a differential trace. Doing so increases jitter and degrades LAN performance.

- The crystal trace lengths should be less than 1 inch.
- The crystal load capacitors should be placed less than 1" from the crystal.
- The clock lines must be at least 5 times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, board edge, transformers and differential pairs.
- The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition.
- The clock lines should not cross or run in parallel (within 3x the dielectric thickness of the closest dielectric layer) with any trace (100Mhz signal or higher) on an adjacent layer.

16.20 Quartz Crystal

Quartz crystals are generally considered to be the mainstay of frequency control components due to their low cost and ease of implementation. They are available from numerous vendors in many package types and with various specification options.

16.21 Fixed Crystal Oscillator

A packaged fixed crystal oscillator comprises of an inverter, a quartz crystal, and passive components conveniently packaged together. The device renders a strong, consistent square wave output. Oscillators used with microprocessors are supplied in many configurations and tolerances.

Crystal oscillators should be restricted for use in special situations, such as shared clocking among devices or multiple controllers. Since clock routing can be difficult to accomplish, it is preferable to provide a separate crystal for each device.

Note: Contact your Intel customer representative to obtain the most current device documentation prior to implementing this solution.



16.22 Crystal Selection Parameters

All crystals used with Intel Ethernet controllers are described as "AT-cut," which refers to the angle at which the unit is sliced with respect to the long axis of the quartz stone.

The following table lists crystals which have been used successfully in past designs. (No particular product is recommended.)

Table 16-10. Crystal Manufacturers and Part Numbers

Manufacturer	Part No.	Note
Raltron*	AS-25.000-20-SMD-TR-NS6	HC-49S package
TXC Corporation - USA*	9C25000355	HC-49S package
KDS America*	DSX321G, 1B/C/N/H225000CC0M	Small package
River*	FCX-04-25MJ90141	Small package

The datasheet for the PHY lists the crystal electrical parameters and provides suggested values for typical designs. Designers should refer to criteria outlined in their respective PHY datasheet. The parameters are described in the following subsections.

16.23 Vibrational Mode

Crystals in the frequency range referenced above are available in both fundamental and third overtone. Unless there is a special need for third overtone, fundamental mode crystals should be used.

16.24 Nominal Frequency

Intel Ethernet controllers use a crystal frequency of 25.000 MHz. The 25 MHz input is used to generate a 125-MHz transmit clock for 100BASE-TX and 1000BASE-TX operation, and 10-MHz and 20-MHz transmit clocks, for 10BASE-T operation.

16.25 Frequency Tolerance

The frequency tolerance for an Ethernet Platform LAN Connect device is dictated by the IEEE 802.3 specification as ± 50 parts per million (ppm). This measurement is referenced to a standard temperature of 25 °C. Intel recommends a frequency tolerance of ± 30 ppm to ensure for any frequency variance contributed by the PCB.

16.26 Temperature Stability and Environmental Requirements

Temperature stability is a standard measure of how the oscillation frequency varies over the full operational temperature range (and beyond). Several optional temperature ranges are currently available, including -40 °C to +85 °C for industrial environments. Some vendors separate operating temperatures from temperature stability. Manufacturers may also list temperature stability as 50 ppm in their data sheets.

Note: Crystals also carry other specifications for storage temperature, shock resistance, and reflow solder conditions. Crystal vendors should be consulted early in the design cycle to discuss its application and environmental requirements.


16.27 Calibration Mode

The terms “series-resonant” and “parallel-resonant” are often used to describe crystal oscillator circuits. Specifying parallel mode is critical to determining how the crystal frequency is calibrated at the factory.

A crystal specified and tested as series resonant oscillates without problem in a parallel-resonant circuit, but the frequency is higher than nominal by several hundred parts per million. The purpose of adding load capacitors to a crystal oscillator circuit is to establish resonance at a frequency higher than the crystal’s inherent series resonant frequency.

Figure 16-21 illustrates a simplified schematic of the internal oscillator circuit. Pin X1 and X2 refers to XTAL_IN and XTAL_OUT in the Ethernet device, respectively. The crystal and the capacitors form a feedback element for the internal inverting amplifier. This combination is called parallel-resonant, because it has positive reactance at the selected frequency. In other words, the crystal behaves like an inductor in a parallel LC circuit. Oscillators with piezoelectric feedback elements are also known as “Pierce” oscillators.


Figure 16-21. Thermal Oscillator Circuit



16.28 Load Capacitance

The formula for crystal load capacitance is as follows:

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



where $C_1 = C_2 = 33 \text{ pF}$ (as suggested in most Intel reference designs)
and C_{stray} = allowance for additional capacitance in pads, traces and the chip carrier within the Ethernet device package and C_{damp} .

16.29 Shunt Capacitance

The shunt capacitance parameter is relatively unimportant compared to load capacitance. Shunt capacitance represents the effect of the crystal's mechanical holder and contacts. The shunt capacitance should equal a maximum of 6 pF (7 pF is also acceptable).

16.30 Equivalent Series Resistance

Equivalent Series Resistance (ESR) is the real component of the crystal's impedance at the calibration frequency, which the inverting amplifier's loop gain must overcome. ESR varies inversely with frequency for a given crystal family. The lower the ESR, the faster the crystal starts up. Crystals with an ESR value of 50 Ω or better should be used.

16.31 Drive Level

Drive level refers to power dissipation in use. The allowable drive level for a Surface Mounted Technology (SMT) crystal is less than its through-hole counterpart. This is due to the fact that surface mount crystals are typically made from narrow, rectangular AT strips, rather than circular AT quartz blanks.

When selecting a crystal, board designers must ensure that the crystal specification meets at least the drive level specified. For example, if the crystal drive level specification states that the drive level is 200 μW maximum, then the crystal drive level must be at least 200 μW . So, a 500 μW crystal is sufficient, but a 100 μW crystal is not.

16.32 Aging

Aging is a permanent change in frequency (and resistance) occurring over time. This parameter is most important in its first year because new crystals age faster than old crystals. Crystals with a maximum value of $\pm 5 \text{ ppm}$ per year aging should be used.

16.33 Reference Crystal

The normal tolerances of the discrete crystal components can contribute to small frequency offsets with respect to the target center frequency. To minimize the risk of tolerance-caused frequency offsets causing a small percentage of production line units to be outside of the acceptable frequency range, it is important to account for those shifts while empirically determining the proper values for the discrete loading capacitors, C_1 and C_2 .

Even with a perfect support circuit, most crystals will oscillate slightly higher or lower than the exact center of the target frequency. Therefore, frequency measurements, which determine the correct value for C_1 and C_2 , should be performed with an ideal reference crystal. When the capacitive load is exactly equal to the crystal's load rating, an ideal reference crystal will be perfectly centered at the desired target frequency.

16.33.1 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal:

- If a Saunders and Associates (S&A) crystal network analyzer is available, then discrete crystal components can be tested until one is found with zero or nearly zero ppm deviation (with the appropriate capacitive load). A crystal with zero or near zero ppm deviation will be a good reference crystal to use in subsequent frequency tests to determine the best values for C1 and C2.
- If a crystal analyzer is not available, then the selection of a reference crystal can be done by measuring a statistically valid sample population of crystals, which has units from multiple lots and approved vendors. The crystal, which has an oscillation frequency closest to the center of the distribution, should be the reference crystal used during testing to determine the best values for C1 and C2.
- It may also be possible to ask the approved crystal vendors or manufacturers to provide a reference crystal with zero or nearly zero deviation from the specified frequency when it has the specified C_{Load} capacitance.

When choosing a crystal, customers must keep in mind that to comply with IEEE specifications for 10/100 Mbps operation and 10/100/1000 Mbps operation if applicable, the transmitter reference frequency must be precise within ± 50 ppm. Intel recommends customers use a transmitter reference frequency that is accurate to within ± 30 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance.

16.33.2 Circuit Board

Since the dielectric layers of the circuit board are allowed some reasonable variation in thickness, the stray capacitance from the printed board (to the crystal circuit) will also vary. If the thickness tolerance for the outer layers of dielectric are controlled within $\pm 15\%$ of nominal, then the circuit board should not cause more than ± 2 pF variation to the stray capacitance at the crystal. When tuning crystal frequency, it is recommended that at least three circuit boards are tested for frequency. These boards should be from different production lots of bare circuit boards.

Alternatively, a larger sample population of circuit boards can be used. A larger population will increase the probability of obtaining the full range of possible variations in dielectric thickness and the full range of variation in stray capacitance.

Next, the exact same crystal and discrete load capacitors (C1 and C2) must be soldered onto each board, and the LAN reference frequency should be measured on each circuit board.

The circuit board, which has a LAN reference frequency closest to the center of the frequency distribution, should be used while performing the frequency measurements to select the appropriate value for C1 and C2.

16.33.3 Temperature Changes

Temperature changes can cause the crystal frequency to shift. Therefore, frequency measurements should be done in the final system chassis across the system's rated operating temperature range.



16.34 Oscillator Support

The PHY clock input circuit is optimized for use with an external crystal. However, an oscillator can also be used in place of the crystal with the proper design considerations (refer to the PHY Datasheet for detailed clock oscillator specifications):

- The clock oscillator has an internal voltage regulator to isolate it from the external noise of other circuits to minimize jitter. If an external clock is used, this imposes a maximum input clock amplitude. For example, if a 3.3 V DC oscillator is used, its output signal should be attenuated to a maximum value with a resistive divider circuit.
- The input capacitance introduced by the PHY (approximately 20 pF) is greater than the capacitance specified by a typical oscillator (approximately 15 pF).
- The input clock jitter from the oscillator can impact the PHY clock and its performance.

Note: The power consumption of additional circuitry equals about 1.5 mW.

Table 1-10 lists oscillators that can be used with the PHY. Note that no particular oscillator is recommended):

Table 16-11. Oscillator Manufacturers and Part Numbers

Manufacturer	Part No.
Kyocera*	K30-3C0-SE-25.0000M
MtronPTI*	M214TCN25.0000MHz
TXC*	7C25000230 7X25080001

16.35 Oscillator Placement and Layout Recommendations

Oscillator clock sources should not be placed near I/O ports or board edges. Radiation from these devices can be coupled into the I/O ports and radiate beyond the system chassis. Oscillators should also be kept away from the Ethernet magnetics module to prevent interference.

The oscillator must have its own decoupling capacitors and they must be placed within 0.25 inches. If a power trace is used (not power plane), the trace from the capacitor to the oscillator must not exceed 0.25 inches in length. The decoupling capacitors help to improve the oscillator stability. The oscillator clock trace should be less than two inches from the PHY. If it is greater than 2 inches, then verify the signal quality, jitter, and clock frequency measurements at the PHY.

The clock lines should also target 50 Ω +/- 15% and should have 33 Ω series back termination placed close to the series oscillator. To help reduce EMI, the clock lines must be a distance of at least five times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, the board edge, transformers and differential pairs.


The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition. The clock lines should not cross or run in parallel with any trace (100 MHz signal or higher) on an adjacent layer.

There should be a ferrite bead within 250 mils of the oscillator power pin and there must be a 1 uF or greater capacitor within 250 mils of the oscillator, connected to the power trace between the oscillator input and ferrite bead. With a ferrite bead on the power trace for the oscillator, there should be a power pour (or fat trace) to supply power to the oscillator.

Note: Figure 13-19 shows a connection between CLK Oscillator Out and the 82579 XTAL_IN. The oscillator must meet the requirements described in the Intel® 82579 GbE PHY Datasheet.

When placing the oscillator, ensure that the 6.8 pF capacitor is included as a series component, as shown in the following diagram.

Figure 16-22. Oscillator Solution



16.36 Troubleshooting Common Physical Layout Issues

The following is a list of common physical layer design and layout mistakes in LAN on Motherboard (LOM) designs.

1. Lack of symmetry between the two traces within a differential pair. Asymmetry can create common-mode noise and distort the waveforms. For each component and via that one trace encounters, the other trace should encounter the same component or a via at the same distance from the Ethernet silicon.
2. Unequal length of the two traces within a differential pair. Inequalities create common-mode noise and will distort the transmit or receive waveforms.
3. Excessive distance between the Ethernet silicon and the magnetics. Long traces on FR4 fiberglass epoxy substrate will attenuate the analog signals. In addition, any impedance mismatch in the traces will be aggravated if they are longer than the four-inch guideline.



4. Routing any other trace parallel to and close to one of the differential traces. Crosstalk getting onto the receive channel will cause degraded long cable BER. Crosstalk getting onto the transmit channel can cause excessive EMI emissions and can cause poor transmit BER on long cables. At a minimum, for stripline other signals should be kept at least 6x the height of the thinnest adjacent dielectric layer. For microstrip it is 7x. The only possible exceptions are in the vicinities where the traces enter or exit the magnetics, the RJ-45 connector, and the Ethernet silicon.
5. Using a low-quality magnetics module.
6. Reusing an out-of-date physical layer schematic in a Ethernet silicon design. The terminations and decoupling can be different from one PHY to another.
7. Incorrect differential trace impedances. It is important to have about a 100- Ω impedance between the two traces within a differential pair. This becomes even more important as the differential traces become longer. To calculate differential impedance, many impedance calculators only multiply the single-ended impedance by two. This does not take into account edge-to-edge capacitive coupling between the two traces. When the two traces within a differential pair are kept close to each other, the edge coupling can lower the effective differential impedance by 5 Ω to 20 Ω . Short traces will have fewer problems if the differential impedance is slightly off target.

16.37 Power Delivery

The 82579 requires a 3.3 V power rail and a 1.05V power rail. The internal 3.3 V power rail is brought out for decoupling. [Figure 16-23](#) shows a typical power delivery configuration that can be implemented. However, power delivery can be customized based on a specific OEM. In general planes should be used to deliver 3.3 Vdc and 1.0 Vdc. Not using planes can cause resistive voltage drop and/or inductive voltage drop (due to transient or static currents). Some of the symptoms of these voltage drops can include higher EMI, radiated immunity, radiated emissions, IEEE conformance issues, and register corruption.

Decoupling capacitors (0.1 μ F and smaller) should be placed within 250 mils of the LAN device. They also should be distributed around the PHY and some should be in close proximity to the power pins.

The bulk capacitors (1.0 μ F or greater) should be placed within 1 inch if using a trace (50 mils wide or wider) or within 1.5 inches if using a plane.


The 1.05 V dc power rail for the 82579 can be generated using either the integrated SVR (iSVR) or can be generated by sharing the PCH SVR output. If using the internal SVR to generate the 1.05 V power, the inductor must be placed within 0.5" of the input pin to the PHY and connected with a trace wider than or equal to 20 mil wide. (Please see the reference schematic for further details regarding the 1.05 V power rail.) If using the shared SVR configuration, the power rail should be delivered through a power plane. Care should be taken to minimize any voltage drops to within 50 mV. For calculating the voltage drop through copper traces, refer to the Power Delivery Loss Calculator. Contact your Intel representative for access.

The following list shows inductors that have been used successfully with Intel designs:

Table 16-12. Inductors and Manufacturers

Manufacturer	Part Number
muRata*	LQH32PN4R7NN0
muRata*	LQH32CN4R7M53
TDK*	FLF3215T-4R7M

Figure 16-23. Intel 82579 Power Delivery Diagram




Note: For latest PHY schematic connection recommendations, refer to the 82579 reference schematic which is available on CDI.

Note: This is the default connection for sharing the 1.05 V rail with the PCH. If using the iSVR, stuff the 4.7 uH inductor and unstuff the 0 ohm resistor to PCH 1.05 V.

16.38 82579 Power Sequencing

if 1.05 Vdc is generated from the 82579's internal SVR or from the PCH's shared power rail. See the datasheet and reference schematic for details. For platform power sequencing information, refer to the PCH datasheet.



17 Design Considerations and Guidelines (Mobile Designs)

The PCH incorporates an integrated 10/100/1000 Mbps MAC controller that can be used with an external Intel® 82579 Physical Layer Device (PHY) shown in [Figure 17-1](#). Its bus master capabilities enable the component to process high-level commands and perform multiple operations, which lowers processor use by off loading communication tasks from the processor.

The PCH, which hereinafter refers to the integrated MAC within the PCH, supports the SMBus interface for manageability while in an Sx state and PCI Express* (PCIe*) for 10/100/1000 Mbps traffic in an S0 state.

Note: Design guidance is available for mobile designs (this chapter) and non-mobile designs (previous chapter). Be sure you are using the proper information for your design type.

Note: The PCIe interface is not PCIe compliant. It operates at half of the PCI Express* (PCIe*) Specification v1.0 (2.5 GT/s) speed. In this chapter, the term “PCIe-based” is interchangeable with “PCIe.” There are no design layout differences between normal PCIe and the PCIe-based interface.


The PHY interfaces with the integrated MAC through two interfaces: PCIe and SMBus. In SMBus mode, the link speed is reduced to 10 Mbps. The PCIe interface incorporates two aspects: a PCIe-based SerDes (electrically) and a custom logic protocol for messaging between the integrated MAC and the PHY.


Note: Gigabit Ethernet requires an SPI Flash to host firmware and does not work without an SPI Flash on board.

The integrated MAC supports multi-speed operation (10/100/1000 Mbps). The integrated MAC also operates in full-duplex at all supported speeds or half-duplex at 10/100 Mbps as well as adhering to the IEEE 802.3x Flow Control Specification.

Note: References to the AUX power rail means the power rail is available in all power states including G3 to S5 transitions and Sx states with Wake on LAN (WoL) enabled. For example, V3P3_AUX in this chapter refers to a rail that is powered under the conditions previously mentioned.

Figure 17-1. PCH/PHY Interface Connections



**Table 17-1. SMBus Data Signals on the PCH**

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

Table 17-2. PCIe Data Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Data	PETp PETn	PETp PETn	PCIe transmit pair
Data	PERp PERn	PERp PERn	PCIe receive pair

Notes:

1. Can be connected to any PCIe port on the integrated MAC. The appropriate NVM descriptor soft strap (PCHSTRP9) should define which PCIe port is configured as GbE LAN.

Table 17-3. Clock and Reset Signals on the PCH

Group	PHY Signal Name	PCH Signal Name	Description
Clock	SMB_CLK	SML0_CLK	SMBus clock
Clock	PE_CLKP PE_CLKN	CLKOUT_PCIE[7:0]_P ¹ CLKOUT_PCIE[7:0]_N ¹	PCIe* clock
Clock ²	CLK_REQ_N	PCIECLKRQ[7:0]#	PCIe clock request
Reset	PE_RST_N	PLTRST# ³	PCIe reset

Notes:

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

17.1 PHY Overview


The PHY is a single port compact component designed for 10/100/1000 Mbps operation. It enables a single port Gigabit Ethernet (GbE) implementation in a very small area, easing routing constraints from the PCH chipset to the PHY.

The PHY provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3ab, 802.3u, and 802.3i, respectively).

17.1.1 PHY Interconnects

The main interfaces for either PHY are PCIe and SMBus on the host side and Media Dependent Interface (MDI) on the link side. Transmit traffic is received from the PCH as either PCIe or SMBus packets on the host interconnect and transmitted as Ethernet packets on the MDI link. Receive traffic arrives as Ethernet packets on the MDI link and transferred to the PCH through either the PCIe or SMBus interconnects.

The PHY switches the in-band traffic automatically between PCIe and SMBus based on platform reset. The transition protocol is done through SMBus. The PCIe interface is powered down when the Ethernet link is running in an Sx state.



17.1.2 PCIe-Based Interface

A high-speed SerDes interface that uses PCIe electrical signaling at half speed while utilizing a custom logical protocol for active state operation mode.

Note: PCIe validation tools cannot be used for electrical validation of this interface; however, PCIe layout rules apply for on-board routing.

17.1.2.0.1 PCIe Interface Signals

The signals used to connect between the PCH and the PHY in this mode are:

- Serial differential pair running at 1.25 Gb/s for Rx.
- Serial differential pair running at 1.25 Gb/s for Tx.
- 100-MHz differential clock input to the PHY is generated by the PCH.
- Power and clock good indication to the PHY PE_RST_N.
- Clock control through CLK_REQ_N (refer to [Table 17-3](#)). This PHY output should be tied to the PCH input and pulled up with a 10-k Ω resistor connected to 3.3-V DC AUX power (present in G3 to S5).

17.1.2.0.2 PCIe Operation and Channel Behavior

The PHY only runs at 1250 Mbps speed, which is 1/2 of the Gen 1 2.5 Gb/s PCIe frequency. Each of the PCIe root ports in the PCH has the ability to run at 1250 Mbps. Configuring a PCH PCIe port that is attached to a PCIe Intel PHY only device is pre-loaded from the GbE region of the NVM. The selected port adjusts the transmitter to run at 1/2 the Gen 1 PCIe speed and does not need to be PCIe compliant.

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

17.1.2.0.3 PCIe Connectivity

The PHY transmit/receive pins are output/input signals and are connected to the PCH as listed in [Table 17-1](#) through [Table 17-3](#).

17.1.2.0.4 PCIe Reference Clock


The PCIe Interface uses a 100-MHz differential reference clock, denoted PE_CLKP and PE_CLKN. This signal is typically generated on the platform and routed to the PCIe port.

The frequency tolerance for the PCIe reference clock is ± 300 ppm.

17.1.3 SMBus Interface

SMBus is a low speed (100 kHz/400 kHz) serial bus used to connect various components in a system. SMBus is used as an interface to pass traffic between the PHY and the PCH when the platform is in a low power state (Sx). The interface is also used to enable the PCH to configure the PHY as well as passing in-band information between them.

The SMBus uses two primary signals: SMBCLK and SMBDATA, to communicate. Both of these signals float high with board-level 2.2 k $\Omega \pm 5\%$ pull-up resistors.



The SMBus specification has defined various types of message protocols composed of individual bytes. For more details about SMBus, see the SMBus specification.

17.1.3.0.1 SMBus Connectivity

Table 17-1 through Table 17-3 list the relationship between PHY SMBus pins to the PCH LAN SMBus pins.

Note: The SMBus signals (SMB_DATA and SMB_CLK) cannot be connected to any other devices other than the integrated MACCougar Point-M. Connect the SMB_DATA and SMB_CLK pins to the integrated MACCougar Point-M SML0DATA and SML0CLK pins, respectively.

17.1.4 PCIe and SMBus Modes

In GbE operation, PCIe is used to transmit and receive data and for MDIO status and control. The PHY automatically switches the in-band traffic between PCIe and SMBus based on the platform power state. The table below lists the operating modes of PCIe and SMBus.

The 82579 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

System / Intel Management Engine State	PHY	
	SMBus	PCIe
S0 and PHY Power Down	Not used	Electrical Idle (EI)
S0 and Idle or Link Disconnect	Not used	EI
S0 and Link in Low Power Idle (LPI)	Not used	EI
S0 and active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

17.1.5 Transitions between PCIe and SMBus Interfaces

17.1.5.0.1 Switching from SMBus to PCIe


Communication between the integrated MACCougar Point-M and the PHY is done through the SMBus each time the system is in a low power state (Sx). The integrated MACCougar Point-M/PHY interface is needed while the Manageability Engine (ME) is still active to transfer traffic, configuration, control and status or to enable host wake up from the PHY.

Possible states for activity over the SMBus:

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

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

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



17.1.5.0.2 Switching from PCIe to SMBus

The communication between the integrated MACCougar Point-M and the PHY is done through PCIe each time the platform is in active power state (S0). Switching the communication to SMBus is only needed for ME activity or to enable host wake up in low power states and is controlled by the ME.

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

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

17.2 Platform LAN Design Guidelines

These sections provide recommendations for selecting components and connecting special pins. For GbE designs, the main elements are the PCH chipset, 82579 PHY, a magnetics module and RJ-45 connector, a GbE region NVM (Non Volatile Memory) image, and a clock source.

17.2.1 General Design Considerations for PHYs

Sound engineering practices must be followed with respect to unused inputs by terminating them with pull-up or pull-down resistors, unless otherwise specified in a datasheet, design guide or reference schematic. Pull-up or pull-down resistors must not be attached to any balls identified as "No Connect." These devices might have special test modes that could be entered unintentionally.

17.2.1.0.1 Clock Source

All designs require a 25-MHz clock source. The PHY uses the 25-MHz source to generate clocks up to 125 MHz and 1.25 GHz for both the PHY circuits and the PCIe interface. For optimum results with lowest cost, a 25-MHz parallel resonant crystal can be used along with the appropriate load capacitors at the XTAL_OUT (X2) and XTAL_IN (X1) leads. The frequency tolerance of the timing device should equal 30 ppm or better. Further detail is found in [Section 17.19](#) and [Section 17.35](#).

Note: XTAL_OUT and XTAL_IN are the signal names for the PHY.

There are three steps to crystal qualification:

1. Verify that the vendor's published specifications in the component datasheet meet the required conditions for frequency, frequency tolerance, temperature, oscillation mode and load capacitance as specified in the respective datasheet.
2. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems.
3. Independently measure the component's electrical parameters in real systems. Measure frequency at a test output to avoid test probe loading effects at the PHY. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications. For crystals, it is also important to examine startup behavior while varying system voltage and temperature.



17.2.1.0.2 Magnetics Module

The magnetics module has a critical effect on overall IEEE and emissions conformance. The device should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Carefully qualifying new magnetics modules prevents problems that might arise because of interactions with other components or the printed circuit board itself.

The steps involved in magnetics module qualification are similar to those for crystal qualification:

1. Verify that the vendor's published specifications in the component datasheet meet or exceed the required IEEE specifications.
2. Independently measure the component's electrical parameters on the test bench, checking samples from multiple lots. Check that the measured behavior is consistent from sample to sample and that measurements meet the published specifications.
3. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems. Vary temperature and voltage while performing system level tests.

Magnetics modules for 1000BASE-T Ethernet as used by the PHY only are similar to those designed solely for 10/100 Mbps, except that there are four differential signal pairs instead of two. Refer to the 82579 datasheet for specific electrical requirements that the magnetics need to meet.

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

Table 17-4. Magnetic Modules and Manufacturers

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

17.2.1.0.3 Criteria for Integrated Magnetics Electrical Qualification

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

**Table 17-5. Integrated Magnetics Recommended Qualification Criteria**

Open Circuit Inductance (OCL)	w/8 mA DC bias; at 25C	400uH Min
	w/8 mA DC bias; at 0C to 70C	350uH Min
Insertion Loss	100 kHz through 999 kHz	1dB Max
	1.0 MHz through 60.0 MHz	0.6dB Max
	60.1 MHz through 80.0 MHz	0.8dB Max
	80.1 MHz through 100.0 MHz	1.0dB Max
	100.1 MHz through 125.0 MHz	2.4dB Max
Return Loss	1.0 MHz through 40.0 MHz	18.0 dB Min
	40.1 MHz through 100.0 MHz When reference impedance is 85 Ohms, 100 Ohms, and 115 Ohms. Note that R.L. values may vary with MDI trace lengths. The LAN magnetics may need to be measured in the platform where it will be used.	$12 - 20 * \text{LOG} (\text{Freq in MHz} / 80) \text{ dB Min}$
Crosstalk Isolation Discrete Modules	1.0 MHz through 29.9 MHz	$-50.3 + (8.8 * (\text{Freq in MHz} / 30)) \text{ dB Max}$
	30.0 MHz through 250.0 MHz	$-(26 - (16.8 * (\text{LOG}(\text{Freq in MHz} / 250 \text{ MHz})))) \text{ dB Max}$
Crosstalk Isolation Integrated Modules (Proposed)	250.1 MHz through 375.0 MHz	-26.0 dB Max
	1.0 MHz through 10 MHz	$-50.8 + (8.8 * (\text{Freq in MHz} / 10)) \text{ dB Max}$
Diff to CMR	10.0 MHz through 100.0 MHz	$-(26 - (16.8 * (\text{LOG}(\text{Freq in MHz} / 100 \text{ MHz})))) \text{ dB Max}$
	100 MHz through 375.0 MHz	-26.0 dB Max
CM to CMR	1 MHz through 29.9 MHz	$-40.2 + (5.3 * ((\text{Freq in MHz} / 30))) \text{ dB Max}$
	30.0 MHz through 500 MHz	$-(22 - (14 * (\text{LOG}((\text{Freq in MHz} / 250)))) \text{ dB Max}$
Hi-Voltage Isolation	1 MHz through 270 MHz	$-57 + (38 * ((\text{Freq in MHz} / 270))) \text{ dB Max}$
	270.1 MHz through 300 MHz	$-17 - 2 * ((300 - (\text{Freq in MHz} / 30)) \text{ dB Max}$
	300.1 MHz through 500 MHz	-17 dB Max
	1500 Vrms at 50 or 60 Hz for 60 sec. or: 2250 Vdc for 60 seconds	Minimum

17.2.2 NVM for PHY Implementations

The LAN only supports an SPI Flash, which is connected to the PCH. Several words of the NVM are accessed automatically by the device after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the NVM space is available to software for storing the MAC address, serial numbers, and additional information. More details may be obtained from the *Datasheet*.

Intel has an MS-DOS* software utility called Eeupdate that is used to program the SPI Flash images in development or production line environments. A copy of this program can be obtained through your Intel representative.

17.2.3 LAN Switch

In order to achieve IEEE conformance for applications that must operate both docked and undocked, a LAN switch is recommended. Note that Intel does not recommend specific switches, but those in the following list have been used successfully in previous designs.

Manufacturer	Part Number
Pericom*	PI3L500-AZ
Texas Instrument*	S3L500AE

17.2.4 LED

The PHY has three LED outputs that can be configured via the NVM. The hardware configuration is shown in [Figure 17-2](#).


Refer to the 82579 Reference Schematic for default LED color based on reference design.

Refer to the *Intel® 82579 Datasheet* for details regarding the programming of the LED's and the various modes. The default values for the PHY (based on the LED NVM setting--word 0x18 of the LAN region) are listed in the table below:

Table 17-6. LED Default Values


LED	Mode	Color	Blink	Polarity
LED0	Link Up/Activity	Green	200 ms on/200 ms off	Active low
LED1	Link 1000	Yellow	No	Active low
LED2	Link 100	Green	No	Active low

Figure 17-2. LED Hardware Configuration



17.2.4.0.1 RBIAS

RBIAS requires external resistor connection to bias the internal analog section of the device. The input is sensitive to the resistor value. Resistors of 1% tolerance must be used. Connect RBIAS through a 3.01 k Ω 1% pull-down resistor to ground and then place it no more than one inch away from the PHY.



17.2.4.0.2 LAN Disable

The PHY enters a power-down state when the LAN_DISABLE_N pin is asserted low. Exiting this mode requires setting the LAN_DISABLE_N pin to a logic one. Connect LAN_DISABLE_N to LAN_PHY_PWR_CTRL.

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

17.2.5.0.1 Overview

This section provides general information about ePAD and SMT assemblies. Chip packages have exposed die pads on the bottom of each package to provide electrical interconnections with the printed circuit board. These ePADs also provide excellent thermal performance through efficient heat paths to the PCB.

Packages with ePADs are very popular due to their low cost. Note that this section only provides basic information and references in regards to the ePAD. It is recommended that each customer consult their fab and assembly house to obtain more details on how to implement the ePAD package design. Each fab and assembly house might need to tune the land pattern/stencil and create a solution that best suits their methodology and process.

17.2.5.0.2 PCB Design Requirements

In order to maximize both heat removal and electrical performance, a land pattern must be incorporated on the PCB within the footprint of the package corresponding to the exposed metal pad or exposed heat slug of the package as shown in the following figures. Refer to the specific product datasheet for actual dimensions.

Note: Due to the package size, a via-in-pad configuration must be used [Figure 17-3](#) and [Figure 17-4](#) are general guidelines see [Figure 17-5](#) for specific via-in-pad thermal pattern recommendations.

Figure 17-3. Typical ePAD Land Pattern






Figure 17-4. Typical Thermal Pad and Via Recommendations



Note: Encroached and uncapped via configurations have voids less than the maximum allowable void percentage. Uncapped via provides a path for trapped air to escape during the reflow soldering process.

Note: Secondary side solder bumps might be seen in an uncapped via design. This needs to be considered when placing components on the opposite side of the PHY.

Figure 17-5. Recommended Thermal Via Patterns for**17.2.5.0.3 Board Mounting Guidelines**

The following are general recommendations for mounting a QFN-48 device on the PCB. This should serve as the starting point in assembly process development and it is recommended that the process should be developed based on past experience in mounting standard, non-thermally/electrically enhanced packages.

17.2.5.0.4 Stencil Design


For maximum thermal/electrical performance, it is required that the exposed pad/slug on the package be soldered to the land pattern on the PCB. This can be achieved by applying solder paste on both the pattern for lead attachment as well as on the land pattern for the exposed pad. While for standard (non-thermally/ -electrically enhanced) lead-frame based packages the stencil thickness depends on the lead pitch and package co-planarity, the package standoff must also be considered for the thermally/ electrically enhanced packages to determine the stencil thickness. In this case, a stencil foil thickness in the range of 5 - 6 mils (or 0.127—0.152 mm) is recommended; likely or practically, a choice of either 5 mils or 6 mils. Tolerance wise, it should not be worse than ± 0.5 mil.

Note: Industry specialists typically use ± 0.1 -mil tolerance on stencil for its feasible precision.

The aperture openings should be the same as the solder mask openings on the land pattern. Since a large stencil opening may result in poor release, the aperture opening should be subdivided into an array of smaller openings, similar to the thermal land pattern shown in the figure below.

Note: Refer to the specific product datasheet for actual dimensions.


Figure 17-6. Stencil Design Recommendation



17.2.5.0.5 Assembly Process Flow

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


Figure 17-7. Assembly Flow



17.2.5.0.6 Reflow Guidelines

The typical reflow profile consists of four sections. In the preheat section, the PCB assembly should be preheated at the rate of 1 to 2 °C/sec to start the solvent evaporation and to avoid thermal shock. The assembly should then be thermally soaked for 60 to 120 seconds to remove solder paste volatiles and for activation of flux. The reflow section of the profile, the time above liquidus should be between 45 to 60 seconds with a peak temperature in the range of 245 to 250 °C, and the duration at the peak should not exceed 30 seconds. Finally, the assembly should undergo cool down in the fourth section of the profile. A typical profile band is provided in the figure below, in which 220 °C is referred to as an approximation of the liquidus point. The actual profile parameters depend upon the solder paste used and specific recommendations from the solder paste manufacturers should be followed.

Figure 17-8. Typical Profile Band



1. Preheat: 125 °C - 220 °C, 150 - 210 s at 0.4 k/s to 1.0 k/s
2. Time at T > 220 °C: 60 - 90 s
3. Peak Temperature: 245-250 °C
4. Peak time: 10 - 30 s
5. Cooling rate: ≤ 6 k/s
6. Time from 25 °C to Peak: 240 - 360 s
7. Intel recommends a maximum solder void of 50% after reflow.

Note: Contact your Intel representative for any designs unable to meet the recommended guidance for E-pad implementation.


17.3 PCH – SMBus/PCIe LOM Design Guidelines

This section contains guidelines on how to implement a PCH/PHY single solution on a system motherboard. It should not be treated as a specification, and the system designer must ensure through simulations or other techniques that the system meets the specified timings. The following are guidelines for both PCH SMBus and PCIe interfaces. Note that PCIe is only applicable to the PHY.

The SMBus/PCIe Interface can be configured in as shown [Figure 17-9](#).

Refer to [Section 17.6](#) for PCI Express Routing Guidelines.

Figure 17-9. Single Solution Interconnect

**Notes:**

1. Any free PCIe ports (Ports 1-8) can be used to hook up to the 82579 PCIe Interface.
2. Any CLKOUT_SRC[7:0] and SRC[7:0]CLKRQ# can be used to connect to PE_CLK and CLK_REQ_N on the 82579.
3. Stuff empty resistor pad with respective resistor in mobile design.
4. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per PCIe specification.
5. For latest PHY schematic connection recommendations, refer to the 82579 reference schematic, which is available on CDI.

17.4 SMBus Design Considerations

No single SMBus design solution works for all platforms. Designers must consider the total bus capacitance and device capabilities when designing SMBus segments. Routing SMBus to the PCI slots makes the design process even more challenging since they add so much capacitance to the bus. This extra capacitance has a large effect on the bus time constant which in turn affects the bus rise and fall times.

Primary considerations in the design process are:

- Amount of $V_{CC_SUS3_3}$ current available, that is, minimizing load of $V_{CC_SUS3_3}$.
- The pull-up resistor size for the SMBus data and clock signals is dependent on the bus load (this includes all device leakage currents). Generally the SMBus device that can sink the least amount of current is the limiting agent on how small the resistor can be. The pull-up resistor cannot be made so large that the bus time constant (Resistance X Capacitance) does not meet the SMBus rise and time specification.
- The maximum bus capacitance that a physical segment can reach is 400 pF.
- SMBus devices that can operate in STR must be powered by the $V_{CC_SUS3_3}$ supply.

- It is recommended that I²C (Inter-Integrated Circuit) devices be powered by the V_{CC_core} supply. During an SMBus transaction in which the device is sending information to the integrated MAC, the device may not release the SMBus if the integrated MAC receives an asynchronous reset. V_{CC_core} is used to enable the BIOS to reset the device if necessary. SMBus 2.0- compliant devices have a timeout capability that makes them in-susceptible to this I²C issue, enabling flexibility in choosing a voltage supply.
- No other devices (except the integrated MAC and pull-up resistors) should be connected to the SMBus that connects to the PHY.
- **For system LAN on motherboard (LOM) designs:** The traces should be less than 70 inches for stripline and less than 100 inches for Microstrip. These numbers depend on the stackup, dielectric layer thickness, and trace width. The total capacitance on the trace and input buffers should be under 400 pF.
- **For system LAN on daughterboard designs:** Being conservative, the traces should be less than 7 inches for stripline designs and less than 10 inches for Microstrip designs. The lengths depend on the stackup, dielectric layer thickness, and trace width. Longer traces can be used as long as the total capacitance on the trace and input buffers is under 30 pF.

Note: Refer to [Section 17.1.3](#) for additional SMBus design considerations.

17.5 General Layout Guidelines

PHY interface signals must be carefully routed on the motherboard to meet the timing and signal quality requirements of their respective interface specifications. The following are some general guidelines that should be followed in designing a LAN solution. It is recommended that the board designer simulate the board routing to verify that the specifications are met for flight times and skews due to trace mismatch and crosstalk.

17.6 Layout Considerations

Critical signal traces should be kept as short as possible to decrease the likelihood of effects by high frequency noise of other signals, including noise carried on power and ground planes. This can also reduce capacitive loading.

Since the transmission line medium extends onto the printed circuit board, layout and routing of differential signal pairs must be done carefully.

Designing for GbE (1000BASE-T) operation is very similar to designing for 10/100 Mbps. For the PHY, system level tests should be performed at all three speeds.

17.7 Guidelines for Component Placement

Component placement can affect signal quality, emissions, and component operating temperature. Careful component placement can:

Decrease potential problems directly related to electromagnetic interference (EMI), which could cause failure to meet applicable government test specifications. In this case, place the PHY more than one inch from the edge of the board.

Simplify the task of routing traces. To some extent, component orientation affects the complexity of trace routing. The overall objective is to minimize turns and crossovers between traces.


17.7.1 PHY Placement Recommendations

Minimizing the amount of space needed for the PHY is important because other interfaces compete for physical space on a motherboard near the connector. The PHY circuits need to be as close as possible to the connector.

Figure 1-10 illustrates some basic placement distance guidelines. To simplify the diagram, it shows only two differential pairs, but the layout can be generalized for a GbE system with four analog pairs. The ideal placement for the PHY (LAN silicon) is approximately one inch behind the magnetics module.

While it is generally a good idea to minimize lengths and distances, this figure also illustrates the need to keep the PHY away from the edge of the board and the magnetics module for best EMI performance.

Figure 17-10. LAN Device Placement: At Least One Inch from Chassis Openings or Unshielded Connectors--Mobile



Note: * this distance is variable and follows the general guidelines.

The PHY, referred to as "LAN Device" in the above figure, must be at least one inch from the I/O back panel. To help reduce EMI, the following recommendations should be followed:

- Minimize the length of the MDI interface. See detail in table below: MDI Routing Summary
- Place the MDI traces no closer than 0.5 inch (1.3 cm) from the board edge.
- The 82579 PHY must be placed greater than 1" away from any hole to the outside of the chassis larger than 0.125 inches (125 mils) The larger the hole the higher the probability the EMI and ESD immunity will be negatively affected.
- The 82579 PHY should be placed greater than 250mils from the board edge.
- If the connector or integrated magnetics module is not shielded, the 82579 should be placed at least one inch from the magnetics (if a LAN switch is not used).
- Placing the 82579 closer than one inch to Unshielded magnetics or connectors will increase the probability of failed EMI and common mode noise. If the LAN switch is too far away it will negatively affect IEEE return loss performance.
- The RBIAS trace length must be less than 1"
- Place the crystal less than 0.75 inch (1.9 cm) from the PHY.

Figure 17-11. PLC Placement: At Least One Inch from I/O Backplane





Figure 17-12. Effect of LAN Device Placed Too Close To a RJ-45 Connector or Chassis Opening





17.8 MDI Differential-Pair Trace Routing for LAN Design

Trace routing considerations are important to minimize the effects of crosstalk and propagation delays on sections of the board where high-speed signals exist. Signal traces should be kept as short as possible to decrease interference from other signals, including those propagated through power and ground planes.

17.9 Signal Trace Geometry

One of the key factors in controlling trace EMI radiation are the trace length and the ratio of trace-width to trace-height above the reference plane. To minimize trace inductance, high-speed signals and signal layers that are close to a reference or power plane should be as short and wide as practical. Ideally, the trace-width to trace-height above the ground plane ratio is between 1:1 and 3:1. To maintain trace impedance, the width of the trace should be modified when changing from one board layer to another if the two layers are not equidistant from the neighboring planes.

Each pair of signals should have a differential impedance of $100\ \Omega \pm 15\%$.

A set of trace length calculation tools are available from Intel (via the Intel Business Link (IBL)) to aid with MDI topology design.

When performing a board layout, the automatic router feature of the CAD tool must not route the differential pairs without intervention. In most cases, the differential pairs will require manual routing.

Note: Measuring trace impedance for layout designs targeting $100\ \Omega$ often results in lower actual impedance due to over-etching. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of $105\ \Omega$ to $110\ \Omega$ should compensate for over-etching.

It is necessary to compensate for trace-to-trace edge coupling, which can lower the differential impedance by up to $10\ \Omega$, when the traces within a pair are closer than 30 mils (edge-to-edge).

Table 17-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Signal group	MDI_PLUS[0:3] MDI_MINUS[0:3]		
Microstrip/stripline uncoupled single-ended impedance specification	$50\ \Omega \pm 10\%$		
Microstrip/stripline uncoupled differential impedance specification	$100\ \Omega \pm 15\%$		2,3
Microstrip nominal trace width	Design dependent	Design dependent	4
Microstrip nominal trace space	Design dependent	Design dependent	3,5
Microstrip/stripline trace length	8 in (203 mm) maximum		6,7
Microstrip pair-to-pair space (edge-to-edge)	≥ 7 times the thickness of the thinnest adjacent dielectric layer		Figure 17-13




Table 17-7. MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines ¹	Notes
Stripline pair-to-pair space (edge-to-edge)	≥ 6 times the thickness of the thinnest adjacent dielectric layer		
Microstrip bus-to-bus spacing	≥ 7 times the thickness of the thinnest adjacent dielectric layer		
Stripline bus-to-bus spacing	≥ 6 times the thickness of the thinnest adjacent dielectric layer		

Notes:

1. Pair-to-pair spacing ≥ 3 times the dielectric thickness for a maximum distance of 500 mils from the pin.
2. Board designers should ideally target 100 Ω ±15%. If it's not feasible (due to board stack-up) it is recommended that board designers use a 95 Ω ±10% target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95 Ω. The ±10% tolerance is provided to allow for board manufacturing process variations and not lower target impedances. The minimum value of impedance cannot be lower than 85 Ω.
3. Simulation shows 80 Ω differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90 Ω.
4. Stripline is NOT recommended due to thinner more resistive signal layers.
5. Use a minimum of 21 mil (0.533 mm) pair-to-pair spacing for board designs that use the CRB design stack-up. Using dielectrics that are thicker than the CRB stack-up might require larger pair-to-pair spacing.
6. For applications that require a longer MDI trace length of more than 8 inches (20.32 mm), it is recommended that thicker dielectric or lower Er materials be used. This permits higher differential trace impedance and wider, lower loss traces. Refer to Table 17-8 for examples of microstrip trace geometries for common circuit board materials.
7. If a LAN switch is not used, then the maximum trace length is 4 inches (102 mm). Mobile designs without LAN switch can range up to ~8 inches. Refer to Table 1-8 for trace length information.

Table 17-8. Maximum Trace Lengths Based on Trace Geometry and Board Stack-Up

Dielectric Thickness (mils)	Dielectric Constant (DK) at 1 MHz	Width / Space / Width (mils)	Pair-to-Pair Space (mils)	Nominal Impedance (Ohms)	Impedance Tolerance (± %)	Maximum Trace Length (inches) ¹
2.7	4.05	4/10/4	19	95 ²	17 ²	3.5
2.7	4.05	4/10/4	19	95 ²	15 ²	4
2.7	4.05	4/10/4	19	95	10	5
3.3	4.1	4.2/9/4.2	23	100 ²	17 ²	4
3.3	4.1	4.2/9/4.2	23	100	15	4.6
3.3	4.1	4.2/9/4.2	23	100	10	6
4	4.2	5/9/5	28	100 ²	17 ²	4.5
4	4.2	5/9/5	28	100	15	5.3
4	4.2	5/9/5	28	100	10	7

Notes:

1. Longer MDI trace lengths may be achievable, but may make it more difficult to achieve IEEE conformance. Simulations have shown deviations are possible if traces are kept short. Longer traces are possible; use cost considerations and stack-up tolerance for differential pairs to determine length requirements.
2. Deviations from 100 Ω nominal and/or tolerances greater than 15% decrease the maximum length for IEEE conformance.

Note:


Use the MDI Differential Trace Calculator to determine the maximum MDI trace length for your trace geometry and board stack-up. Contact your Intel representative for access.

The following factors can limit the maximum MDI differential trace lengths for IEEE conformance:

- Dielectric thickness
- Dielectric constant
- Nominal differential trace impedance
- Trace impedance tolerance
- Copper trace losses
- Additional devices, such as switches, in the MDI path may impact IEEE conformance.

Board geometry should also be factored in when setting trace length.

Figure 17-13. MDI Trace Geometry



17.10 Trace Length and Symmetry

The differential traces should be equal in total length to within 10 mils (0.254 mm) per segment within each pair and as symmetrical as possible. Asymmetrical and unequal length traces in the differential pairs contribute to common mode noise. If a choice has to be made between matching lengths and fixing symmetry, more emphasis should be placed on fixing symmetry. Common mode noise can degrade the receive circuit's performance and contribute to radiated emissions.


The intra-pair length matching on the pairs must be within 10 mils on a segment by segment basis. An MDI segment is defined as any trace within the same layer. For example, transitioning from one layer to another through a via is considered as two separate MDI segments.

The end to end total trace lengths within each differential pair must match as shown in the figure titled MDI Trace Geometry. The end to end trace length is defined as the total MDI length from one component to another regardless of layer transitions.

The pair to pair length matching is not as critical as the intra-pair length matching but it should be within 2 inches.

When using Microstrip, the MDI traces should be at least 7x the thinnest adjacent dielectric away from the edge of an adjacent reference plane. When using stripline, the MDI traces should be at least 6x the thinnest adjacent dielectric away from the edge of an adjacent reference plane.

Figure 17-14. MDI Differential Trace Geometry



Note: Similar topology applies to MDI routing from the 82579 to the dock RJ45 connector.

17.11 Impedance Discontinuities


Impedance discontinuities cause unwanted signal reflections. Vias (signal through holes) and other transmission line irregularities should be minimized. If vias must be used, a reasonable budget is four or less per differential trace. Unused pads and stub traces should also be avoided.

17.12 Reducing Circuit Inductance

Traces should be routed over a continuous reference plane with no interruptions. If there are vacant areas on a reference or power plane, the signal conductors should not cross the vacant area. This causes impedance mismatches and associated radiated noise levels.

17.13 Signal Isolation

To maintain best signal integrity, keep digital signals far away from the analog traces. Also, keep the MDI traces away from the edge of an adjacent reference plane by a distance that is at least 7x the thickness of the thinnest adjacent dielectric layer (7x when using Microstrip; 6x when using stripline). If digital signals on other board layers cannot be separated by a ground plane, they should be routed perpendicular to the differential pairs. If there is another LAN controller on the board, the differential pairs from that circuit must be kept away.



Other rules to follow for signal isolation include:

- Separate and group signals by function on separate layers if possible. If possible, maintain at least a gap of 30 mils between all differential pairs (Ethernet) and other nets, but group associated differential pairs together.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.
- Isolate I/O signals from high-speed signals to minimize crosstalk, which can increase EMI emission and susceptibility to EMI from other signals.
- Avoid routing high-speed LAN traces near other high-frequency signals associated with a video controller, cache controller, processor, switching power supplies, or other similar devices.

17.14 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and power inputs bypassed to signal return. This will significantly reduce EMI radiation.

The following guidelines help reduce circuit inductance in both backplanes and motherboards:

- Route traces over a continuous plane with no interruptions. Do not route over a split power or ground plane. If there are vacant areas on a ground or power plane, avoid routing signals over the vacant area. This will increase inductance and EMI radiation levels.
- All ground vias should be connected to every ground plane; and every power via, to all power planes at equal potential. This helps reduce circuit inductance.
- Physically locate grounds between a signal path and its return. This will minimize the loop area.
- Split the ground plane beneath a magnetics module. The RJ-45 connector side of the transformer module should have chassis ground beneath it.

Caution: DO NOT do this, if the RJ-45 connector has integrated USB.

Note: All impedance-controlled signals should be routed in reference to a solid plane. If there are plane splits on a reference layer and the signal traces cross those splits then stitching capacitors should be used within 40 mils of where the crossing occurs. See Figure 1-13.

If signals transition from one reference layer to another reference layer then stitching capacitors or connecting vias should be used based on the following:

If the transition is from power-referenced layer to a ground-referenced layer or from one voltage-power referenced layer to a different voltage-power referenced layer, then stitching capacitors should be used within 40 mils of the transition.

If the transition is from one ground-referenced layer to another ground-referenced layer or is from a power-referenced layer to the same net power-referenced layer, then connecting vias should be used within 40 mils of the transition.

Figure 17-15. Trace Transitioning Layers and crossing Plane Splits




Figure 17-16. Via Connecting GND to GND





Figure 17-17. Stitching Capacitor between Vias Connecting GND to GND



17.15 Traces for Decoupling Capacitors

Traces between decoupling and I/O filter capacitors should be as short and wide as practical. Long and thin traces are more inductive and reduce the intended effect of decoupling capacitors. Also, for similar reasons, traces to I/O signals and signal terminations should be as short as possible. Vias to the decoupling capacitors should be sufficiently large in diameter to decrease series inductance. Refer to the Power Delivery section for the PHY in regards to actual placement requirements of the capacitors.

17.16 Ground Planes under a Magnetics Module

The magnetics module chassis or output ground (secondary side of transformer) should be separated from the digital or input ground (primary side) by a physical separation of 100 mils minimum. Splitting the ground planes beneath the transformer minimizes noise coupling between the primary and secondary sides of the transformer and between the adjacent coils in the magnetics. This arrangement also improves the common mode choke functionality of magnetics module.

Caution: DO NOT do this if the RJ-45 connector has integrated USB.

Figure 1-19 illustrates the split plane layout for a discrete magnetics module. Capacitors are used to interconnect chassis ground and signal ground.

Figure 17-18 shows the preferred method for implementing a ground split under an integrated magnetics module/RJ-45 connector.

Figure 17-18. Ideal Ground Split Implementation





Table 17-9. Capacitor Stuffing Option Recommended Values

Capacitors	Value
C3, C4	4.7 μ F or 10 μ F
C1, C2, C5, C6	470 pF to 0.1 μ F


The placement of C1 through C6 may also differ for each board design (in other words, not all of the capacitors may need to be populated). Also, the capacitors may not be needed on both sides of the magnetics module.

Note: If using an integrated magnetics module without USB, provide a separate chassis ground "island" to ground around the RJ-45 connector. The split in the ground plane should be at least 20 mils wide.



Some integrated magnetics modules/RJ-45 connectors have recently incorporated the USB into the device. For this type of magnetics module, a chassis ground moat may not be feasible due to the digital ground required for the USB pins and their placement relative to the magnetics pins. Thus, a continuous digital ground without any moats or splits must be used. [Figure 17-19](#) provides an example of this.

Figure 17-19. Ground Layout with USB



17.17 Light Emitting Diodes

The device has three high-current outputs to directly drive LEDs for link, activity and speed indication. Since LEDs are likely to be integral to a magnetics module, take care to route the LED traces away from potential sources of EMI noise. In some cases, it may be desirable to attach filter capacitors.

LAN LED traces should be placed at least 6x (side by side separation) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

LAN LED traces should be placed at least 7x (broadside coupling) the dielectric height from sources of noise (ex: signaling traces) and susceptible signal traces (ex: reset signals) on the same or adjacent layers.

17.18 Considerations for Layout

The PHY MDI routing using microstrip requires a differential impedance of $100\ \Omega \pm 15\%$. A 35 mils (0.889 mm) separation is required between pairs. The 35-mil separation can be reduced for 24 mils (0.61 mm) in breakout routing. All MDI traces must be referenced to ground.

17.19 Frequency Control Device Design Considerations

This section provides information regarding frequency control devices, including crystals and oscillators, for use with all Intel Ethernet controllers. Several suitable frequency control devices are available; none of which present any unusual challenges in selection. The concepts documented within this section are applicable to other data communication circuits, including the PHY.

The PHY contains amplifiers that form the basis for feedback oscillators when they are used with the specific external components. These oscillator circuits, which are both economical and reliable, are described in more detail in [Section 17.23](#).

The chosen frequency control device vendor should be consulted early in the design cycle. Crystal and oscillator manufacturers familiar with networking equipment clock requirements may provide assistance in selecting an optimum, low-cost solution.

Several types of third-party frequency reference components are currently available. Descriptions of each type follow in subsequent sections. They are also listed in order of preference.

17.20 Crystals and Oscillators

Clock sources should not be placed near I/O ports or board edges. Radiation from these devices may be coupled onto the I/O ports or out of the system chassis. Crystals should also be kept away from the Ethernet magnetics module to prevent interference.

Crystal and load capacitors should be placed on the printed circuit boards as close to the PHY as possible, which is within 1.0 inch. Traces from XTAL_IN (X1) and XTAL_OUT (X2) should be routed as symmetrically as possible. Do not route X1 and X2 as a differential trace. Doing so increases jitter and degrades LAN performance.

- The crystal trace lengths should be less than 1 inch.
- The crystal load capacitors should be placed less than 1" from the crystal.
- The clock lines must be at least 5 times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, board edge, transformers and differential pairs.
- The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition.
- The clock lines should not cross or run in parallel (within 3x the dielectric thickness of the closest dielectric layer) with any trace (100Mhz signal or higher) on an adjacent layer.

17.21 Quartz Crystal

Quartz crystals are generally considered to be the mainstay of frequency control components due to their low cost and ease of implementation. They are available from numerous vendors in many package types and with various specification options.



17.22 Fixed Crystal Oscillator

A packaged fixed crystal oscillator comprises of an inverter, a quartz crystal, and passive components conveniently packaged together. The device renders a strong, consistent square wave output. Oscillators used with microprocessors are supplied in many configurations and tolerances.

Crystal oscillators should be restricted for use in special situations, such as shared clocking among devices or multiple controllers. Since clock routing can be difficult to accomplish, it is preferable to provide a separate crystal for each device.

Note: Contact your Intel customer representative to obtain the most current device documentation prior to implementing this solution.

17.23 Crystal Selection Parameters

All crystals used with Intel Ethernet controllers are described as “AT-cut,” which refers to the angle at which the unit is sliced with respect to the long axis of the quartz stone.

The following table lists crystals which have been used successfully in past designs. (No particular product is recommended.)

Table 17-10. Crystal Manufacturers and Part Numbers

Manufacturer	Part No.	Note
Raltron*	AS-25.000-20-SMD-TR-NS6	HC-49S package
TXC Corporation - USA*	9C25000355	HC-49S package
KDS America*	DSX321G, 1B/C/N/H225000CC0M	Small package
River*	FCX-04-25MJ90141	Small package

The datasheet for the PHY lists the crystal electrical parameters and provides suggested values for typical designs. Designers should refer to criteria outlined in their respective PHY datasheet. The parameters are described in the following subsections.

17.24 Vibrational Mode

Crystals in the frequency range referenced above are available in both fundamental and third overtone. Unless there is a special need for third overtone, fundamental mode crystals should be used.

17.25 Nominal Frequency

Intel Ethernet controllers use a crystal frequency of 25.000 MHz. The 25 MHz input is used to generate a 125-MHz transmit clock for 100BASE-TX and 1000BASE-TX operation, and 10-MHz and 20-MHz transmit clocks, for 10BASE-T operation.

17.26 Frequency Tolerance

The frequency tolerance for an Ethernet Platform LAN Connect device is dictated by the IEEE 802.3 specification as ± 50 parts per million (ppm). This measurement is referenced to a standard temperature of 25 °C. Intel recommends a frequency tolerance of ± 30 ppm to ensure for any frequency variance contributed by the PCB.

17.27 Temperature Stability and Environmental Requirements

Temperature stability is a standard measure of how the oscillation frequency varies over the full operational temperature range (and beyond). Several optional temperature ranges are currently available, including -40 °C to +85 °C for industrial environments. Some vendors separate operating temperatures from temperature stability. Manufacturers may also list temperature stability as 50 ppm in their data sheets.

Note: Crystals also carry other specifications for storage temperature, shock resistance, and reflow solder conditions. Crystal vendors should be consulted early in the design cycle to discuss its application and environmental requirements.


17.28 Calibration Mode

The terms “series-resonant” and “parallel-resonant” are often used to describe crystal oscillator circuits. Specifying parallel mode is critical to determining how the crystal frequency is calibrated at the factory.

A crystal specified and tested as series resonant oscillates without problem in a parallel-resonant circuit, but the frequency is higher than nominal by several hundred parts per million. The purpose of adding load capacitors to a crystal oscillator circuit is to establish resonance at a frequency higher than the crystal’s inherent series resonant frequency.

Figure 17-20 illustrates a simplified schematic of the internal oscillator circuit. Pin X1 and X2 refers to XTAL_IN and XTAL_OUT in the Ethernet device, respectively. The crystal and the capacitors form a feedback element for the internal inverting amplifier. This combination is called parallel-resonant, because it has positive reactance at the selected frequency. In other words, the crystal behaves like an inductor in a parallel LC circuit. Oscillators with piezoelectric feedback elements are also known as “Pierce” oscillators.

Figure 17-20. Thermal Oscillator Circuit



17.29 Load Capacitance

The formula for crystal load capacitance is as follows:

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

where $C1 = C2 = 33 \text{ pF}$ (as suggested in most Intel reference designs) and C_{stray} = allowance for additional capacitance in pads, traces and the chip carrier within the Ethernet device package and C_{damp} .

17.30 Shunt Capacitance

The shunt capacitance parameter is relatively unimportant compared to load capacitance. Shunt capacitance represents the effect of the crystal's mechanical holder and contacts. The shunt capacitance should equal a maximum of 6 pF (7 pF is also acceptable).

17.31 Equivalent Series Resistance

Equivalent Series Resistance (ESR) is the real component of the crystal's impedance at the calibration frequency, which the inverting amplifier's loop gain must overcome. ESR varies inversely with frequency for a given crystal family. The lower the ESR, the faster the crystal starts up. Crystals with an ESR value of 50 Ω or better should be used.

17.32 Drive Level

Drive level refers to power dissipation in use. The allowable drive level for a Surface Mounted Technology (SMT) crystal is less than its through-hole counterpart. This is due to the fact that surface mount crystals are typically made from narrow, rectangular AT strips, rather than circular AT quartz blanks.

When selecting a crystal, board designers must ensure that the crystal specification meets at least the drive level specified. For example, if the crystal drive level specification states that the drive level is 200 μW maximum, then the crystal drive level must be at least 200 μW . So, a 500 μW crystal is sufficient, but a 100 μW crystal is not.

17.33 Aging

Aging is a permanent change in frequency (and resistance) occurring over time. This parameter is most important in its first year because new crystals age faster than old crystals. Crystals with a maximum value of ± 5 ppm per year aging should be used.

17.34 Reference Crystal


The normal tolerances of the discrete crystal components can contribute to small frequency offsets with respect to the target center frequency. To minimize the risk of tolerance-caused frequency offsets causing a small percentage of production line units to be outside of the acceptable frequency range, it is important to account for those shifts while empirically determining the proper values for the discrete loading capacitors, C1 and C2.

Even with a perfect support circuit, most crystals will oscillate slightly higher or lower than the exact center of the target frequency. Therefore, frequency measurements, which determine the correct value for C1 and C2, should be performed with an ideal reference crystal. When the capacitive load is exactly equal to the crystal's load rating, an ideal reference crystal will be perfectly centered at the desired target frequency.

17.34.1 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal:

- If a Saunders and Associates (S&A) crystal network analyzer is available, then discrete crystal components can be tested until one is found with zero or nearly zero ppm deviation (with the appropriate capacitive load). A crystal with zero or near zero ppm deviation will be a good reference crystal to use in subsequent frequency tests to determine the best values for C1 and C2.
- If a crystal analyzer is not available, then the selection of a reference crystal can be done by measuring a statistically valid sample population of crystals, which has units from multiple lots and approved vendors. The crystal, which has an oscillation frequency closest to the center of the distribution, should be the reference crystal used during testing to determine the best values for C1 and C2.
- It may also be possible to ask the approved crystal vendors or manufacturers to provide a reference crystal with zero or nearly zero deviation from the specified frequency when it has the specified C_{Load} capacitance.



When choosing a crystal, customers must keep in mind that to comply with IEEE specifications for 10/100 Mbps operation and 10/100/1000 Mbps operation if applicable, the transmitter reference frequency must be precise within ± 50 ppm. Intel recommends customers use a transmitter reference frequency that is accurate to within ± 30 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance.

17.34.2 Circuit Board

Since the dielectric layers of the circuit board are allowed some reasonable variation in thickness, the stray capacitance from the printed board (to the crystal circuit) will also vary. If the thickness tolerance for the outer layers of dielectric are controlled within $\pm 15\%$ of nominal, then the circuit board should not cause more than ± 2 pF variation to the stray capacitance at the crystal. When tuning crystal frequency, it is recommended that at least three circuit boards are tested for frequency. These boards should be from different production lots of bare circuit boards.

Alternatively, a larger sample population of circuit boards can be used. A larger population will increase the probability of obtaining the full range of possible variations in dielectric thickness and the full range of variation in stray capacitance.

Next, the exact same crystal and discrete load capacitors (C1 and C2) must be soldered onto each board, and the LAN reference frequency should be measured on each circuit board.

The circuit board, which has a LAN reference frequency closest to the center of the frequency distribution, should be used while performing the frequency measurements to select the appropriate value for C1 and C2.

17.34.3 Temperature Changes

Temperature changes can cause the crystal frequency to shift. Therefore, frequency measurements should be done in the final system chassis across the system's rated operating temperature range.

17.35 Oscillator Support

The PHY clock input circuit is optimized for use with an external crystal. However, an oscillator can also be used in place of the crystal with the proper design considerations (refer to the PHY Datasheet for detailed clock oscillator specifications):

- The clock oscillator has an internal voltage regulator to isolate it from the external noise of other circuits to minimize jitter. If an external clock is used, this imposes a maximum input clock amplitude. For example, if a 3.3 V DC oscillator is used, its output signal should be attenuated to a maximum value with a resistive divider circuit.
- The input capacitance introduced by the PHY (approximately 20 pF) is greater than the capacitance specified by a typical oscillator (approximately 15 pF).
- The input clock jitter from the oscillator can impact the PHY clock and its performance.

Note: The power consumption of additional circuitry equals about 1.5 mW.



Table 1-10 lists oscillators that can be used with the PHY. Note that no particular oscillator is recommended):

Table 17-11.Oscillator Manufacturers and Part Numbers

Manufacturer	Part No.
Kyocera*	K30-3C0-SE-25.0000M
MtronPTI*	M214TCN25.0000MHz
TXC*	7C25000230 7X25080001

17.36 Oscillator Placement and Layout Recommendations

Oscillator clock sources should not be placed near I/O ports or board edges. Radiation from these devices can be coupled into the I/O ports and radiate beyond the system chassis. Oscillators should also be kept away from the Ethernet magnetics module to prevent interference.

The oscillator must have its own decoupling capacitors and they must be placed within 0.25 inches. If a power trace is used (not power plane), the trace from the capacitor to the oscillator must not exceed 0.25 inches in length. The decoupling capacitors help to improve the oscillator stability. The oscillator clock trace should be less than two inches from the PHY. If it is greater than 2 inches, then verify the signal quality, jitter, and clock frequency measurements at the PHY.

The clock lines should also target $50\ \Omega \pm 15\%$ and should have $33\ \Omega$ series back termination placed close to the series oscillator. To help reduce EMI, the clock lines must be a distance of at least five times the height of the thinnest adjacent dielectric layer away from other digital traces (especially reset signals), I/O ports, the board edge, transformers and differential pairs.


The clock lines must not cross any plane cuts on adjacent power or ground reference layers unless there are decoupling capacitors or connecting vias near the transition. The clock lines should not cross or run in parallel with any trace (100 MHz signal or higher) on an adjacent layer.

There should be a ferrite bead within 250 mils of the oscillator power pin and there must be a 1 μF or greater capacitor within 250 mils of the oscillator, connected to the power trace between the oscillator input and ferrite bead. With a ferrite bead on the power trace for the oscillator, there should be a power pour (or fat trace) to supply power to the oscillator.

Note: Figure 13-19 shows a connection between CLK Oscillator Out and the 82579 XTAL_IN. The oscillator must meet the requirements described in the Intel® 82579 GbE PHY Datasheet.

When placing the oscillator, ensure that the 6.8 pF capacitor is included as a series component, as shown in the following diagram.

Figure 17-21. Oscillator Solution



17.37 LAN Switch

The following table lists LAN switches that can be used with the 82579. Note that no particular LAN switch is recommended:

Table 17-12. LAN Switch Manufacturers and Part Numbers

Manufacturer	Part No.
Pericom*	PI3L500-AZ
Texas Instruments*	TS3L500AE

17.38 Troubleshooting Common Physical Layout Issues

The following is a list of common physical layer design and layout mistakes in LAN on Motherboard (LOM) designs.

1. Lack of symmetry between the two traces within a differential pair. Asymmetry can create common-mode noise and distort the waveforms. For each component and via that one trace encounters, the other trace should encounter the same component or a via at the same distance from the Ethernet silicon.
2. Unequal length of the two traces within a differential pair. Inequalities create common-mode noise and will distort the transmit or receive waveforms.
3. Excessive distance between the Ethernet silicon and the magnetics. Long traces on FR4 fiberglass epoxy substrate will attenuate the analog signals. In addition, any impedance mismatch in the traces will be aggravated if they are longer than the four-inch guideline.
4. Routing any other trace parallel to and close to one of the differential traces. Crosstalk getting onto the receive channel will cause degraded long cable BER. Crosstalk getting onto the transmit channel can cause excessive EMI emissions and

can cause poor transmit BER on long cables. At a minimum, for stripline other signals should be kept at least 6x the height of the thinnest adjacent dielectric layer. For microstrip it is 7x. The only possible exceptions are in the vicinities where the traces enter or exit the magnetics, the RJ-45 connector, and the Ethernet silicon.

5. Using a low-quality magnetics module.
6. Reusing an out-of-date physical layer schematic in a Ethernet silicon design. The terminations and decoupling can be different from one PHY to another.
7. Incorrect differential trace impedances. It is important to have about a 100- Ω impedance between the two traces within a differential pair. This becomes even more important as the differential traces become longer. To calculate differential impedance, many impedance calculators only multiply the single-ended impedance by two. This does not take into account edge-to-edge capacitive coupling between the two traces. When the two traces within a differential pair are kept close to each other, the edge coupling can lower the effective differential impedance by 5 Ω to 20 Ω . Short traces will have fewer problems if the differential impedance is slightly off target.

17.39 Power Delivery

The 82579 requires a 3.3 V power rail and a 1.05V power rail. The internal 3.3 V power rail is brought out for decoupling. [Figure 17-22](#) shows a typical power delivery configuration that can be implemented. However, power delivery can be customized based on a specific OEM. In general planes should be used to deliver 3.3 V_D and 1.0 V_{dc}. Not using planes can cause resistive voltage drop and/or inductive voltage drop (due to transient or static currents). Some of the symptoms of these voltage drops can include higher EMI, radiated immunity, radiated emissions, IEEE conformance issues, and register corruption.

Decoupling capacitors (0.1 μ F and smaller) should be placed within 250 mils of the LAN device. They also should be distributed around the PHY and some should be in close proximity to the power pins.


The bulk capacitors (1.0 μ F or greater) should be placed within 1 inch if using a trace (50 mils wide or wider) or within 1.5 inches if using a plane.

The 1.05 V dc power rail for the 82579 can be generated using either the integrated SVR (iSVR) or can be generated by sharing the PCH SVR output. If using the internal SVR to generate the 1.05 V power, the inductor must be placed within 0.5" of the input pin to the PHY and connected with a trace wider than or equal to 20 mil wide. (Please see the reference schematic for further details regarding the 1.05 V power rail.) If using the shared SVR configuration, the power rail should be delivered through a power plane. Care should be taken to minimize any voltage drops to within 50 mV. For calculating the voltage drop through copper traces, refer to the Power Delivery Loss Calculator. Contact your Intel representative for access.

The following list shows inductors that have been used successfully with Intel designs:

Table 17-13. Inductors and Manufacturers

Manufacturer	Part Number
muRata*	LQH32PN4R7NN0
muRata*	LQH32CN4R7M53
TDK*	FLF3215T-4R7M

Figure 17-22. Intel 82579 Power Delivery Diagram

Note: For latest PHY schematic connection recommendations, refer to the 82579 reference schematic which is available on CDI.

Note: This is the default connection for sharing the 1.05 V rail with the PCH. If using the iSVR, stuff the 4.7 uH inductor and unstuff the 0 ohm resistor to PCH 1.05 V.

17.40 Power Sequencing

The 82579 PHY does not require any power sequencing between the 3.3-V DC and 1.05 Vdc power rails if 1.05 Vdc is generated from the 82579's internal SVR or from the PCH's shared power rail. See the datasheet and reference schematic for details. For platform power sequencing information, refer to the PCH datasheet.



Mouser Electronics

Authorized Distributor

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

[Intel:](#)

[WG82579V S LHA7](#) [WG82579V S LHA8](#) [WG82579LM S LHA6](#) [WG82579LM S LHA5](#)