



Intel® 82578 GbE PHY

Datasheet

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
- 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
- Intel® VPro, Intel® Viiv and Virtualization support with appropriate Intel® chipset(s) components
- 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

■ Performance

- Supports up to 4 KB jumbo frames (full duplex)¹
- 802.1Q and 802.1p
- Receive side Scaling (RSS)
- Two queues (Tx and Rx)

■ Power

- Fully integrated linear regulation for 1.2 Vdc
- 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
- Low Power Link Up (LPLU)

■ Media Access Controller (MAC)² / Physical Layer (PHY) Interconnect

- PCI-based interface for active state operation (S0 state)
- SMBus 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
- Reduced Bill Of Material (BOM) cost by sharing SPI Flash with the Intel® 5 Series Express Chipset

1. Refer to the latest *Intel® 82578 Specification Update* for more details.

2. The MAC is incorporated into the Intel® 5 Series Express Chipset.



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.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

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

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>.

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 © 2012, 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 Codes	3
1.6	Product Matrix	4
2.0	Interconnects	5
2.1	Introduction	5
2.2	PCIe-Based	5
2.2.1	PCIe Interface Signals	5
2.2.2	PCIe Operation and Channel Behavior	6
2.3	SMBus	6
2.3.1	Overview	6
2.4	Transitions between SMBus and PCIe interfaces	8
2.4.1	Switching from SMBus to PCIe	8
2.4.2	Switching from PCIe to SMBus	8
2.5	Intel® 5 Series Express Chipset/82578 – SMBus/PCIe Interconnects	9
3.0	Pin Interface	11
3.1	Pin Assignment	11
3.1.1	Signal Type Definitions	11
3.1.2	PCIe Interface Pins (8)	12
3.1.3	SMBus Interface Pins (2)	12
3.1.4	Miscellaneous Pins (3)	12
3.1.5	PHY Pins (14)	13
3.1.6	Testability Pins (5)	14
3.1.7	Power Pins (13)	14
3.1.8	LVR Power and Control Pins (3)	14
4.0	Package	15
4.1	Package Type and Mechanical	15
4.2	Package Electrical and Thermal Characteristics	16
4.3	Power and Ground Requirements	17
4.4	Pinouts (Top View, Pins Down)	18
4.5	Ball Mapping	19
5.0	Initialization	21
5.1	Power Up	21
5.2	Reset Operation	23
5.3	Timing Parameters	24
5.3.1	Timing Requirements	24
5.3.2	Timing Guarantees	24
6.0	Power Management and Delivery	25
6.1	Power Targets	25
6.2	Power Delivery	26
6.2.1	2.5 Vdc and 1.8 Vdc Supply	26
6.2.2	1.2 Vdc Supply	26
6.3	Power Management	26
6.3.1	Global Power States	26
6.4	Power Saving Features	28
6.4.1	Intel® Auto Connect Battery Saver (ACBS)	28
6.4.2	Automatic Link Downshift	28



7.0	Device Functionality	33
7.1	Tx Flow	33
7.2	Rx Flow	33
7.3	Flow Control.....	33
	7.3.1 MAC Control Frames and Reception of Flow Control Packets.....	34
	7.3.2 Transmitting PAUSE Frames	35
7.4	Wake Up	35
	7.4.1 Host Wake Up	36
	7.4.2 Accessing The 82578's Wake Up Register Using MDIC.....	44
7.5	PHY Loopback	45
8.0	Programmer's Visible State	47
8.1	Terminology.....	47
8.2	MDIO Access.....	48
8.3	Addressing.....	48
8.4	Address Map	49
8.5	PHY Registers (Page 0).....	51
	8.5.1 Extended Debug Port Registers.....	65
8.6	Port Control Registers (Page 769)	68
8.7	Statistics Registers	69
8.8	PCIe Registers.....	71
8.9	General Registers	73
	8.9.1 Interrupts.....	75
8.10	Wake Up Registers.....	76
	8.10.1 Accessing Wake Up Registers Using MDIC	76
	8.10.2 Host Wake Up Control Status Register Description.....	77
9.0	Non-Volatile Memory (NVM)	87
9.1	Introduction	87
9.2	NVM Programming Procedure Overview	87
9.3	LAN NVM Format and Contents	89
	9.3.1 Hardware Accessed Words	90
	9.3.2 Software Accessed Words	99
	9.3.3 Basic Configuration Software Words	104
9.4	Intel® 5 Series Express Chipset/82578 NVM Contents.....	106
10.0	Intel® 5 Series Express Chipset MAC Programming Interface	107
10.1	Register Byte Ordering	107
10.2	Register Conventions	108
	10.2.1 PCI Configuration and Status Registers - CSR Space.....	109
11.0	Electrical and Timing Specifications	161
11.1	Introduction	161
11.2	Operating Conditions.....	161
	11.2.1 Absolute Maximum Ratings	161
	11.2.2 Recommended Operating Conditions	162
	11.2.3 ESD Specifications	162
11.3	Power Delivery	162
	11.3.1 Voltage Regulator Power Supply Specifications	163
	11.3.2 Power On/Off Sequence.....	165



11.4	I/O DC/AC Parameters	165
11.4.1	3.3 Vdc DC/IO	165
11.4.2	2.5 Vdc/IO	166
11.4.3	Input Buffer Only	167
11.4.4	SMBus AC I/O	168
11.4.5	PCIe DC/AC Specifications	168
11.5	Discrete/Integrated Magnetics Specifications.....	171
11.6	Oscillator/Crystal Specifications.....	172
12.0	Schematic and Board Layout Checklists	175
13.0	Reference Schematics	177
14.0	Models.....	179



Revision History

Date	Revision	Description
February 2012	2.5	<ul style="list-style-type: none"> Revised Table 67 (bit 5 description).
July 2011	2.4	<ul style="list-style-type: none"> Revised title page. Revised word 0x13 (new defaults for bits 7:6 and 15:9). Revised word 0x1B (new defaults for bits 1:0 and 4). Revised word 0x26 (new defaults for bits 10:9). Revised section 5.3.2 (Timing Guarantees, added T_{xtal} and T_{por} values). Revised section 11.4.2 (+3.3 Vdc level CMOS for signals JTAG_TCK, JTAG_TMS, and JTAG_TDI).
June 2010	2.3	<ul style="list-style-type: none"> Update title page. Updated section 10.3.1.10 (bit 13 description). Added new values to section 12.4.2 (VIH). Removed old section 8. Updated section 10.3.1.11 (bit 12 and 13 descriptions).
February 2010	2.2	<ul style="list-style-type: none"> Updated figure 1. Updated table 2. Updated section 7.4 and 10.3.1.2 (added Intel® 5 Series Express Chipset references). Added power sequencing note to section 5.3.2. Updated section 6.4.2.2 (added Windows* 7 reference). Updated sections 7.4.1.3.1.4 through 7.4.1.3.1.7 and 7.4.1.3.2.1 through 7.4.1.3.2.2 (swapped Possible VLAN Tag and Possible Len/LLC/SNAP Header in the tables). Added a note to section 10.3.1.15 (LED behaviour).
November 2009	2.1	<ul style="list-style-type: none"> Updated power consumption targets in section 6. Updated the NVM format and contents to match current NVM image. Added a PHY functionality section. Updated the recommended operating conditions in section 12.
June 2009	2.0	<ul style="list-style-type: none"> Initial Public Release.
May 2009	1.75	<ul style="list-style-type: none"> Major revision (all sections).
April 2009	1.0	<ul style="list-style-type: none"> Updated title page (advanced cable diagnostics). Added new Section "Intel® 5 Series Express Chipset/82578 – SMBus/PCIe Interconnects". Added SMBus specification reference to section 1.5. Updated pad size in section 4.1. Changed internal pin name from LAN_PWR_GOODn to LAN_DISABLE_N (all sections). Added new Section "Device Functionality". Added new Section "MAC Programming Interface". Added Appendix A, B, and C. Updated Section 11.6.
January 2009	0.98	<ul style="list-style-type: none"> Updated section 7.3.2 - removed power on/off sequence diagram.
December 2008	0.97	<ul style="list-style-type: none"> Updated tables 7 and 8 - added LAN disable estimated power numbers.
December 2008	0.96	<ul style="list-style-type: none"> Updated tables 7 and 8 (latest power numbers).
Sept 2008	0.95	<ul style="list-style-type: none"> Section 2.2.2 (Removed last paragraph and Table 2). Section 2.3 (changed SMBCLK to SMB_CLK and SMBDATA to SMB_DATA). Section 2.3.1 (updated paragraph). Section 2.3.1.6 (Removed). Removed old sections 2.3.1.6.1, 2.3.1.6.2, and 2.3.1.7). Section 2.3.2.2.1 (updated table). Section 4.1 (added new mechanical drawing and pad size values). Section 5.3.2 (changed T_{XTAL} parameter to 35 ms). Section 6.3.1.1 (updated paragraph). Section 6.3.1.2 (removed all mode 1 references and updated register references). Section 6.3.1.3 (added K1 Idle State information). Section 6.3.2 (changed KX to K0). Section 6.3.3 (updated register references). Section 6.4 (new section). Section 7.2.1 (added notes following table).



Date	Revision	Description
July 2008	0.9	<ul style="list-style-type: none">• Added new section 8.0 "Non-Volatile Memory (NVM)".• Added section 6.4 "Power Saving Features".• Updated section 6.1 "Power Targets".• Updated section 3.1.7 "Power Pins".• Updated section 3.1.5.2 "Analog Pins".
May 2008	0.81	<ul style="list-style-type: none">• Updated pin description in Figure 3 and Section 4.5.
May 2008	0.8	<ul style="list-style-type: none">• Updated Tables 7 and 8 (added new power numbers).• Updated Sections 3.1.7 and 3.1.8 (clarified power, LVR, and control pins).• Updated Section 4.1 (Added Epad size specifications).• Updated Figure 1 (removed ferrite beads from diagram).• Updated Sections 1.2 (added note), 2.2 (added note), 2.3 (added note), and 7.2.2.
April 2008	0.7	Added a discrete/integrated magnetics specifications table to Section 7.0.
Mar 2008	0.6	Major revision (all sections).
Jan 2008	0.5	Initial release (Intel Confidential).



1.0 Introduction

1.1 Scope

This document describes the external architecture for the 82578. 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 82578.

1.2 Overview

The 82578 is a single port Gigabit Ethernet Physical Layer Transceiver (PHY). It connects to the Intel® 5 Series Express Chipset integrated Media Access Controller (MAC) through a dedicated interconnect. The 82578 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 82578 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 82578 interfaces with its MAC through two interfaces: PCIe-based and SMBus. The PCIe (main) interface is used for all link speeds when the system is in an active state (S0) while the SMBus is used only when the system is in a low power state (Sx). In SMBus mode, the link speed is reduced to 10 Mb/s (dependent on low power options). The PCIe interface incorporates two aspects: a PCIe SerDes (electrically) and a custom logic protocol.

Note: The 82578 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 82578's PCIe-based interface.

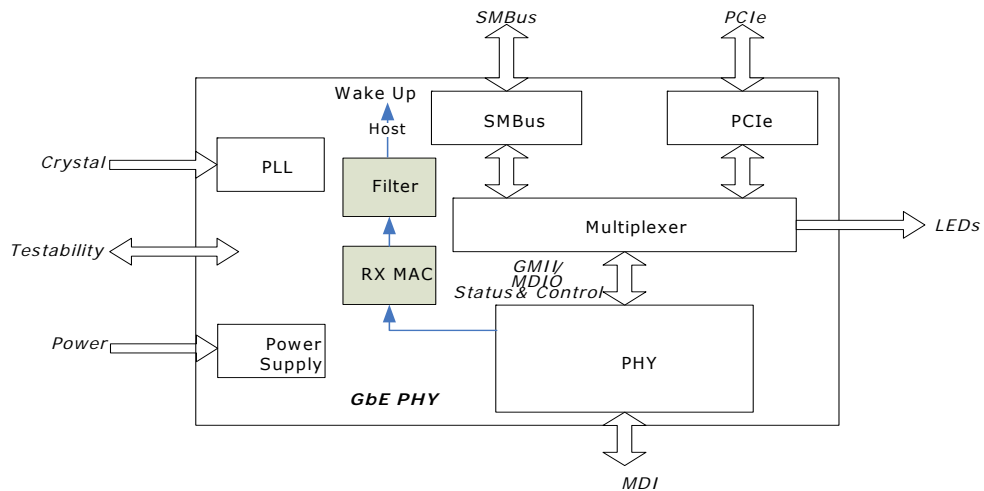


Figure 1. 82578 Block Diagram

1.3 Main Flows

The 82578 main interfaces are PCIe and SMBus on the host side and the MDI interface on the link side. Transmit traffic is received from the MAC device through either PCIe or SMBus on the host interconnect and then transmitted on the MDI link. Receive traffic arrives on the MDI link and transferred to the MAC through either the PCIe or SMBus interconnects.

The MAC and system software control the 82578 functionality through two mechanisms:

- The 82578 configuration registers are mapped into the MDIO space and can be accessed by the MAC 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:

- Power delivery options are described in [Section 4.3](#).
- Power management is described in [Section 6.3](#).



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.: 802.3-2002
- *Intel® Ethernet Controllers Loopback Modes*, Intel Corporation
- SMBus specification revision 2.0.
- *Intel® 5 Series Express Chipset Family External Design Specification (Intel® 5 Series Express Chipset EDS)*, Intel Corporation
- *Intel® 5 Series Express Chipset External Datasheet Specification*, Intel Corporation
- *Intel® 5 Series Express Chipset SPI Flash Programming Guide - Application Note*, Intel Corporation
- *Intel® 82578 Schematic and Layout Checklists*, Intel Corporation
- *Intel® 82578 MDI Differential Trace and Power Loss Calculators*, Intel Corporation

1.5 Product Codes

Table 1 lists the product ordering codes for the 82578 GbE controller. Refer to the *Intel® 82578 GbE PHY Specification Update* for device ordering information.

Table 1. Product Ordering Codes

Device	Market Segment	Product Code
82578DM	Corporate desktop and workstation	WG82578DM
82578DC	Consumer desktop	WG82578DC



1.6 Product Matrix

					Performance		Extended											
Method of enabling/disabling features in SKUs					Driver	Driver	MAC	MAC	PHY	Driver	PHY	Driver	Driver	Platform	Platform	BIOS	BIOS	FW
Link Speed	Platform	Segment	Description	Device ID	Jumbo Frames (up to 4k)	802.1Q & 802.1p	Receive Side Scaling (RSS)	2 Tx & 2 Rx Queues	Low 'No-Link' Power (ACBS)	Link Speed Battery Saver	Low Power Linkup (LPLU)	Ability to Initiate a Team	Intel® Virtualization	Viv (Home IT)	Intel® SIPP	PXE Boot	iSCSI Boot	Intel® vPro*
Gigabit	Desktop	Corporate	82578 for Corporate	10EF	X	X	X	X	X		X	X	X	X	X	X	X	X
	Desktop	Consumer	82578 for Consumer	10F0		X			X		X	X		X		X		



2.0 Interconnects

2.1 Introduction

The 82578 implements two interconnects to the MAC:

- 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. 82578 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 active	Not used	Active
Sx	Active	Power down
Sx and DMoff	Active	Power down

The 82578 automatically switches the in-band traffic between PCIe and SMBus based on the system power state.

2.2 PCIe-Based

Note:

The 82578 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 82578's PCIe-based interface. Standard PCIe validation tools cannot be used to validate this interface.

2.2.1 PCIe Interface Signals

The signals used to connect between the MAC 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 pin
- Clock control through CLK_REQ_N pin



2.2.2 PCIe Operation and Channel Behavior

The 82578 only runs at 1250 Mb/s speed, which is 1/2 of the PCIe Specification v1.1, 2.5 Gb/s PCIe frequency. Each of the PCIe root ports in the Intel® 5 Series Express Chipset-integrated MAC have the ability to operate with the 82578. The port configuration is pre-loaded from the NVM. The selected port adjusts the transmitter to run at the 1.25 GHz 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 82578 SMBus must only be connected to SMLink0 in the Intel® 5 Series Express Chipset. No other device (like an external BMC) can be connected to SMLink0 when the 82578 is connected to the Intel® 5 Series Express Chipset SMLink0.

2.3.1 Overview

SMBus is used as an interface to pass traffic between the 82578 and the Intel® 5 Series Express Chipset when the system is in a low power state (Sx state). The interface is also used to enable the Intel® 5 Series Express Chipset to configure the 82578 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 float high with board-level pull-ups.

The SMBus specification has defined various types of message protocols composed of individual bytes. The message protocols supported by the 82578 are described in the relevant sections.

For more details about SMBus, see the SMBus specification.

2.3.1.1 SMBus Channel Behavior

The SMBus specification defines the maximum frequency of the SMBus as 100 KHz.

2.3.1.2 SMBus Addressing

The 82578's address is assigned using SMBus ARP protocol. The default SMBus address is 0xC8.



2.3.1.3 Bus Timeouts

The 82578 can detect (as a master or a slave) an SMB_CLK timeout on the main SMBus. If the SMBus clock line is held low for 25 ms, the 82578 aborts the transaction.

As a slave, the 82578 detects the timeout 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 timeout might have been processed depending on the transaction.

As a master, the 82578 detects a timeout 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 82578 detects a timeout on is aborted.

2.3.1.4 Bus Hangs

Although uncommon, SMBus bus hangs can happen in a system. The reason 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 timeouts 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 82578 has the capability of detecting a hung bus. If SMB_CLK or SMB_DATA are stuck low for more than 35 ms, the 82578 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



2.3.1.6 SMBus ARP Functionality

The 82578 doesn't support ARP protocol.

2.4 Transitions between SMBus and PCIe interfaces

2.4.1 Switching from SMBus to PCIe

Communication between the MAC and the 82578 is done through the SMBus each time the system is in a low power state (Sx); PE_RST_N signal is low. The MAC/PHY interface is needed to enable host wake up from the 82578.

Possible states for activity over the SMBus:

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

While in this state, the SMBus is used to transfer traffic, configuration, control and status between the MAC and the 82578.

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 MAC and the 82578 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 and is controlled by the Intel® 5 Series Express Chipset.

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

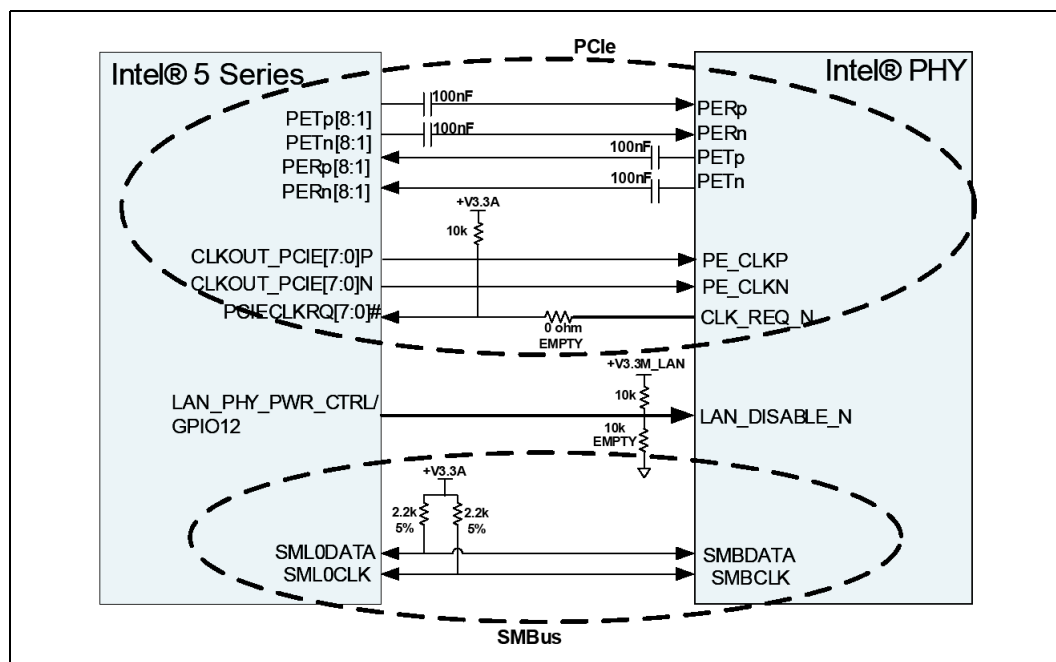
- Any transmit/receive packet that is not completed when PE_RST_N goes to 0b is discarded.
- Any in-band message that was sent over PCIe and was not acknowledged is re-transmitted over SMBus.

2.5 Intel® 5 Series Express Chipset/82578 – SMBus/PCIe Interconnects

The 82578 can be connected to any x1 PCIe port in Intel® 5 Series Express Chipset. The PCIe port that connects to the 82578 is selected by PCHSTRP9, bits [11:8] in the SPI Flash descriptor region. For more information on this setting, please refer to the *Intel® 5 Series Express Chipset External Datasheet Specification*. The Intel® 5 Series Express Chipset-to-82578 PCIe port connection in the reference schematic must match the previously mentioned Intel® 5 Series Express Chipset SPI strap setting. Choosing another port can result in unexpected system behavior.

The SMBus/PCIe interface can be configured in as shown [Figure 2](#).

Figure 2. Intel® 5 Series Express Chipset/82578 Interconnects



Notes:

1. Any free PCIe ports (ports 1-8) can be used to connect to the 82578 PCIe Interface.
2. Any CLKOUT_PCIE[7:0] and PCIECLKRQ[7:0] can be used to connect to PE_CLK for the 82578.
3. PETp/n, PERp/n, PE_CLKp/n should be routed as differential pair as per the PCIe specification.



Note: This page intentionally left blank.



3.0 Pin Interface

3.1 Pin Assignment

The 82578 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.

Note: Refer to the reference schematics for pin connection details. Contact your Intel representative for access.

3.1.1 Signal Type Definitions

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



3.1.2 PCIe Interface Pins (8)

Pin Name	Pin #	Type	Op Mode	Name and Function
PE_RST_N	36	I	Input	PCIe reset.
PETp PETn	38 39	A-out	Output	PCIe Tx.
PERp PERn	41 42	A-in	Input	PCIe Rx.
PE_CLKP PE_CLKN	44 45	A-in	Input	PCIe clock.
CLK_REQ_N	48	O/d	Output	Clock request. Connect to VCC3P3 through a 10 K Ω pull-up resistor.

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 ¹) through a 2.2 K Ω resistor (while in Sx mode).
SMB_DATA	31	O/d	BI-dir	SMBus data. Pull this signal up to 3.3 Vdc (auxiliary supply) through a 2.2 K Ω resistor (while in Sx mode).

1. AUX power means the power rail is available in all power states including G3 to S5 transitions and Sx states with Wake on LAN (WoL) enabled.

3.1.4 Miscellaneous Pins (3)

Pin Name	Pin #	Type	Op Mode	Name and Function
RSVD1_VCC3P3	1	T/s		Connect to VCC3P3 through a 5%, 3.01 K Ω resistor.
RSVD2_VCC3P3	2	T/s		Connect to VCC3P3 through a 5%, 3.01 K Ω resistor.
LAN_DISABLE_N	3	I		Note: When this pin is set to 0b, the 82578 is disabled.



3.1.5 PHY Pins (14)

3.1.5.1 LEDs (3)

This table lists the functionality of the LED output pins. Refer to the Intel® 5 Series Family Platform Design Guide (PDG) for LED connection details.

Pin Name	Pin #	Type	Op Mode	Name and Function
LED0	26	O	Output	This signal is used for the programmable LED (LINK_LINK/ACTIVITY).
LED1	27	O	Output	This signal is used for the programmable LED (LINK_1000).
LED2	25	O	Output	This signal is used for the programmable LED (LINK_100).

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] 100BASE-T : In MDI configuration, MDI[0]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[0]+/- corresponds to BI_DB+/-. 100BASE-TX : In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair. 10BASE-T : In MDI configuration, MDI[0]+/- is used for the transmit pair and in MDI-X configuration MDI[0]+/- is used for the receive pair.
MDI_PLUS[1] MDI_MINUS[1]	17 18	Analog	Bi-dir	Media Dependent Interface[1] 100BASE-T : In MDI configuration, MDI[1]+/- corresponds to BI_DB+/- and in MDI-X configuration MDI[1]+/- corresponds to BI_DA+/-. 100BASE-TX : In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair. 10BASE-T : In MDI configuration, MDI[1]+/- is used for the receive pair and in MDI-X configuration MDI[1]+/- is used for the transmit pair.
MDI_PLUS[2] MDI_MINUS[2] MDI_PLUS[3] MDI_MINUS[3]	20 21 23 24	Analog	Bi-dir	Media Dependent Interface[3:2] 100BASE-T : In MDI configuration, MDI[3:2]+/- corresponds to BI_DA+/- and in MDI-X configuration MDI[3:2]+/- corresponds to BI_DB+/-. 100BASE-TX : Unused. 10BASE-T : Unused.
XTAL_OUT	9	O		Output crystal.
XTAL_IN	10	I		Input crystal.
RBIA5	12	Analog		Connect to ground through a 2.37 KΩ +/-1%.



3.1.6 Testability Pins (5)

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

Note: The 82578 uses the JTAG interface to support XOR files for manufacturing test. BSDL is not supported.

3.1.7 Power Pins (13)

Pin Name	Pin #	Type	Name and Function
AVDD1P2	8, 11, 16, 22, 40, 43	Power	1.2 Vdc supply.
AVDD1P2	8	Power	1.2 Vdc sense feedback.
AVDD2P5	15, 19	Power	2.5 Vdc supply.
DVDD1P2	37, 46, 47	Power	1.2 Vdc supply; connected to DVDD using 0 Ω 0805 resistors.
DVDD2P5	29	Power	2.5 Vdc supply to I/O.
VDD3P3_IN	5	Power	3.3 Vdc supply.
VDD2P5_OUT	4	Power	2.5 Vdc out.

3.1.8 LVR Power and Control Pins (3)

Pin Name	Pin #	Type	Name and Function
CTRL1P2	7	Analog	Connect to the base of the PNP.
VCT	6	Analog	Regulator output; connect to 1.8 Vdc supply and a center tap, 1 μ F capacitor.

4.0 Package

4.1 Package Type and Mechanical

The 82578 is a 6 mm x 6 mm, 48-pin QFN Halogen Free, Pb Free package with a pad size of 3.80 mm x 3.80 mm.




Figure 3. Package Dimensions



4.2 Package Electrical and Thermal Characteristics

The thermal resistance from junction to case, θ_{JC} , is 15.1 $\times C/Watt$.

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} ($\times C/Watt$)
0	119	34
1	118	33
2	116	31

No heat sink is required.

4.3 Power and Ground Requirements

The 82578 requires three power supplies plus one internal power rail that is brought out for decoupling. Figure 4 shows a typical power delivery configuration that can be implemented.

Note: Power delivery can be customized based on a specific OEM platform configuration.

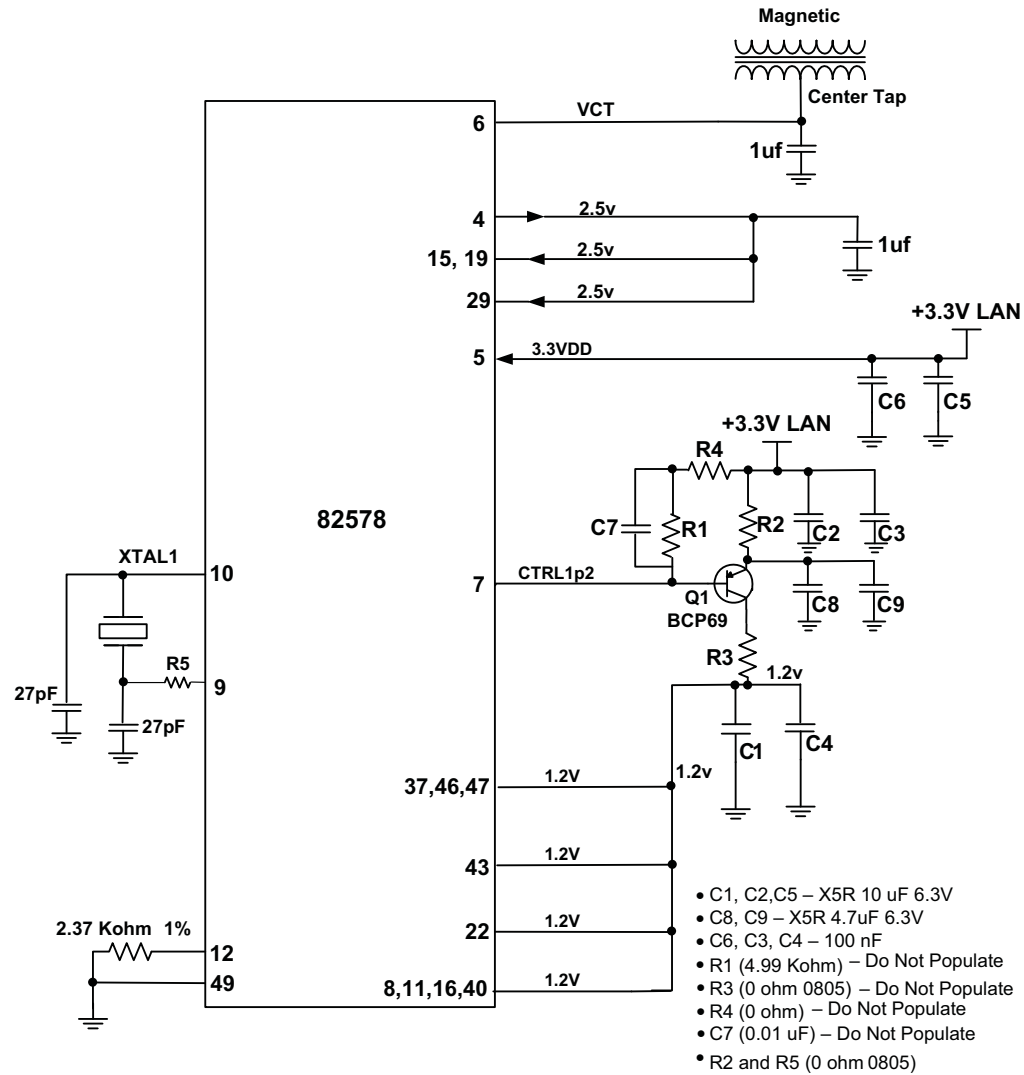


Figure 4. 82578 Power Delivery Diagram



4.4 Pinouts (Top View, Pins Down)

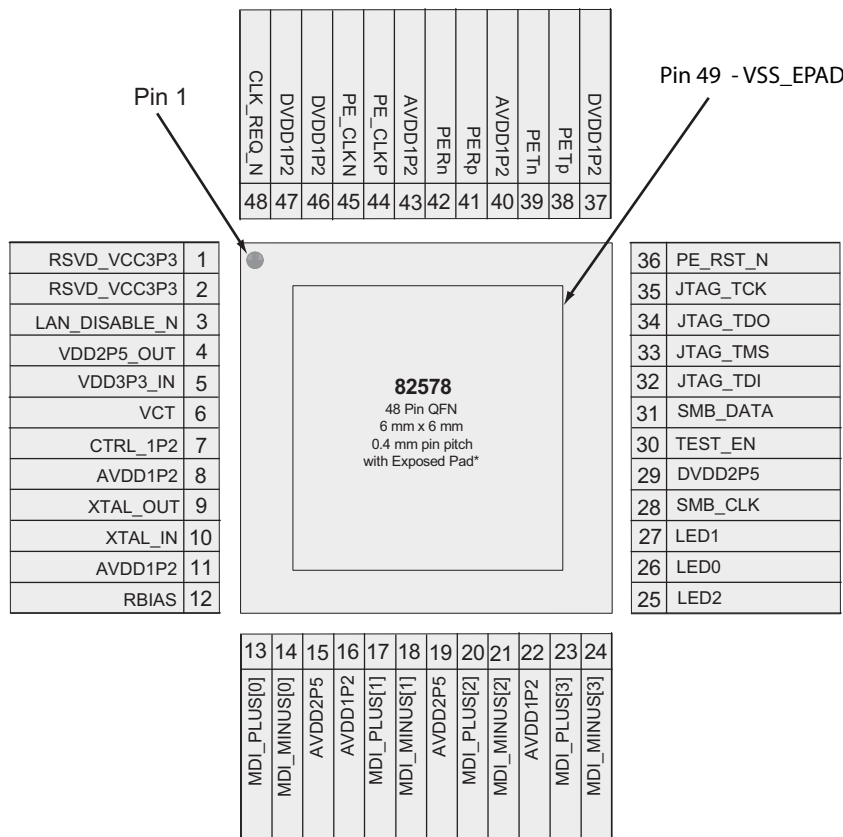


Figure 5. 82578 Pinouts



4.5 Ball Mapping

Pin Name	Side	Pin Number	Pin Name	Side	Pin Number
RSVD_VCC3P3	Left	1	MDI_PLUS[0]	Bottom	13
RSVD_VCC3P3	Left	2	MDI_MINUS[0]	Bottom	14
LAN_DISABLE_N	Left	3	AVDD2P5	Bottom	15
VDD2P5_OUT	Left	4	AVDD1P2	Bottom	16
VDD3P3_IN	Left	5	MDI_PLUS[1]	Bottom	17
VCT	Left	6	MDI_MINUS[1]	Bottom	18
CTRL1P2	Left	7	AVDD2P5	Bottom	19
AVDD1P2	Left	8	MDI_PLUS[2]	Bottom	20
XTAL_OUT	Left	9	MDI_MINUS[2]	Bottom	21
XTAL_IN	Left	10	AVDD1P2	Bottom	22
AVDD1P2	Left	11	MDI_PLUS[3]	Bottom	23
RBIAS	Left	12	MDI_MINUS[3]	Bottom	24
LED2	Right	25	DVDD1P2	Top	37
LED0	Right	26	PETp	Top	38
LED1	Right	27	PETn	Top	39
SMB_CLK	Right	28	AVDD1P2	Top	40
DVDD2P5	Right	29	PERp	Top	41
TEST_EN	Right	30	PERn	Top	42
SMB_DATA	Right	31	AVDD1P2	Top	43
JTAG_TDI	Right	32	PE_CLKP	Top	44
JTAG_TMS	Right	33	PE_CLKN	Top	45
JTAG_TDO	Right	34	DVDD1P2	Top	46
JTAG_TCK	Right	35	DVDD1P2	Top	47
PE_RST_N	Right	36	CLK_REQ_N	Top	48
VSS_EPAD	EPAD	49			



Note: This page intentionally left blank.

5.0 Initialization

5.1 Power Up

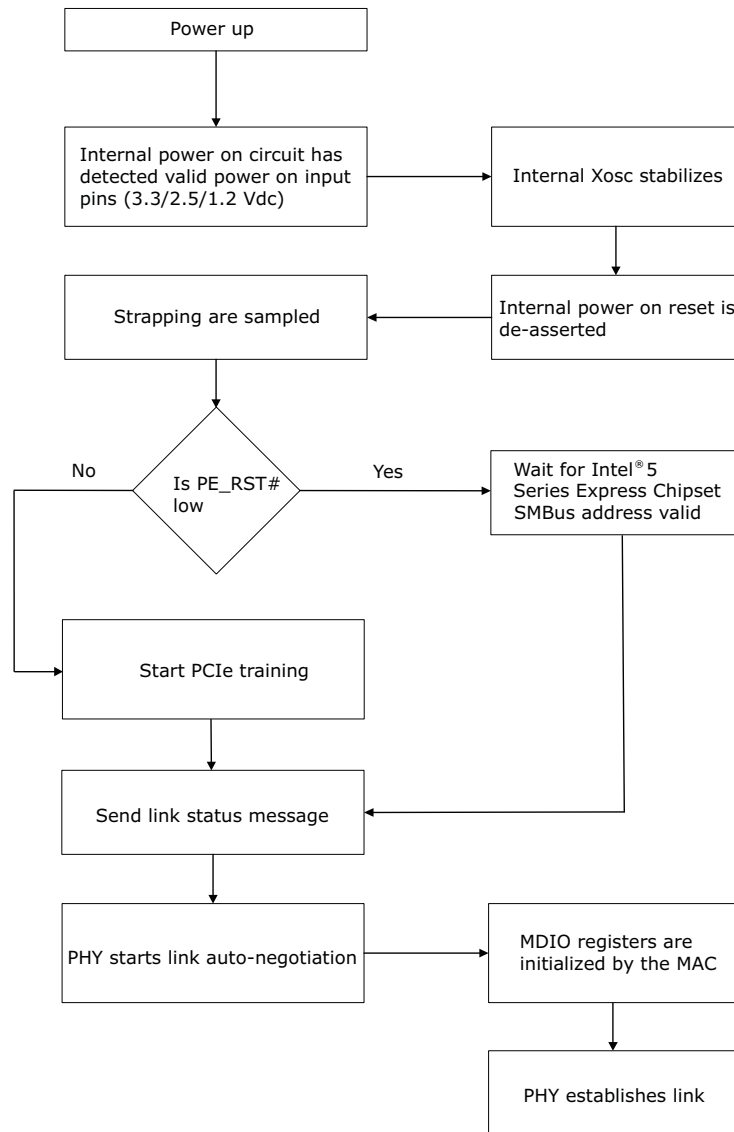


Figure 6. Power Up Sequence



Table 3. Figure Notes

Note	
1	Platform power ramps up (3.3 V dc/2.5/1.2 Vdc)
2	XTAL is stable after T_{XTAL} sec.
3	Internal Power On Reset triggers T_{POR} after XTAL is stable. Strapping options are latched.
4	PCIe training if PE reset is de-asserted.
5	Wait for Intel® 5 Series Express Chipset SMBus address valid.
6	Send Link Status message.
7	MAC configures the 82578.
8	PHY goes through auto-negotiation to acquire link.



5.2 Reset Operation

The reset sources for the 82578 are as follows:

- **Internal Power On Reset (POR)** – The 82578 has an internal mechanism for sensing the power pins. Until power is up and stable, the 82578 generates an internal active low reset. This reset acts as a master reset for the 82578. While the internal reset is 0b, all registers in the 82578 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. Soft reset is used mainly to program the PHY to a different work point without affecting functionality of the rest of the device. Once the PHY completes its internal reset, a reset complete indication is sent to the MAC over the interconnect. The MAC then configures the PHY.

Note:

The MAC configures the PHY registers. Other the 82578 registers do not need to be configured.

- **PCIe Reset** - After asserting a PCIe reset, the 82578 stops the PCIe interface and if in the middle of transmitting a packet it will 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 82578 except the wake up filters content.

The effect of the various reset options on these and other registers is listed in [Table 4](#).

[Table 4](#) lists the impact of each the 82578 reset.

Table 4. 82578 Resets

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

1. Asserting a 3.3 Vdc power on reset should move the PHY out of power down mode.
2. PHY registers (page 0 in MDIO space and any aliases to page 0) are reset during a PHY soft reset. The rest of the 82578's MDIO space is not reset.



5.3 Timing Parameters

5.3.1 Timing Requirements

The 82578 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 82578 guarantees the following start-up and power state transition related timing parameters.

Note: For platform power sequencing requirements for the Intel® 5 Series Express Chipset MAC, refer to the Intel® 5 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}	X_{tal} frequency stable after platform power ramp up		45 ms	
T_{por}	Internal POR trigger after X_{tal} stable		40 ms	



6.0 Power Management and Delivery

This section describes how power management is implemented in the 82578.

6.1 Power Targets

Table 7 lists the targets for device power for the 82578. Note that power is reduced according to link speed and link activity.

Note: Device power is the power dissipated by the 82578.

Table 7. 82578 Power Consumption Targets – External 1.8 Vdc¹

System State		Link State	3.3 Vdc Current [mA]	1.8 Vdc Current [mA]	1.2 Vdc Current [mA]	Device Power [mW]	Solution Power (mW) BCP69 Solution
S0 Maximum		1000Mb – active @ 90 °C [Ta]	22	193	284	761	1647
S0 Typical		1000 Mb/s active	22	193	282	758	1640
		100 Mb/s active	14	43	121	259	587
		10 Mb/s active	15	114	72	341	663
		1000 Mb/s idle	22	194	237	706	1495
		100 Mb/s idle	13	42	71	204	416
		10 Mb/s idle	11	1	25	68	122
		Cable disconnect	9	0	10	42	63
		LAN disable	9	0	14	47	76
Sx	Wake on LAN (WoL) enabled	100 Mb/s - WoL	13	42	68	200	406
		10 Mb/s - WoL	11	1	23	66	116
	WoL disabled	Disabled in BIOS ²	0	0	0	0	0
		Disabled in driver	9	0	14	47	76

1. Measured power could be higher or lower based on measurement setup and PHY power delivery configuration.
 2. Assumes the system is in the Moff state and SLP_LAN# is used to gate PHY power.



The following sections describe requirements in specific power states.

6.2 Power Delivery

The 82578 operates from a 3.3 Vdc external power rail (see [Figure 7](#)).

6.2.1 2.5 Vdc and 1.8 Vdc Supply

The 2.5 Vdc and 1.8 Vdc are supplied by internal LVRs as shown in the schematics that follow.

6.2.2 1.2 Vdc Supply

The 1.2 Vdc rail can be supplied in one of two ways (see [Figure 4](#)):

- An external power supply not dependent on support from the 82578. For example, the platform designer might choose to route a platform-available 1.2 Vdc supply to the 82578. Intel® 5 Series Express Chipset
- A discrete LVR solution, where the base current of PNP power transistor is driven by the 82578, while the power transistor is placed externally.

6.3 Power Management

6.3.1 Global Power States

The 82578 transitions between power states based on a status packet received over the interconnect and based on the Ethernet link state. The following power states are defined:

- **Power Up** – Defined as the period from the time power is applied to the 82578 and until the 82578 powers up its PHY. The 82578 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 82578 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 82578 is not actively transmitting or receiving data and might enter a lower power state (for example, the custom interface can be in electrical idle).
- **Cable Disconnect** – The PHY identified that a cable is not connected. The 82578 signals the MAC that the link is down. The PHY might enter energy detect mode or the MAC might initiate a move into active power down mode (sD3).
- **Power Down (LAN Disable)** – Entry into power down is initiated by the MAC by setting the LAN_DISABLE_N pin to zero. The 82578 loses all functionality in this mode other than the ability to power up again.
-

6.3.1.1 Power Up

Defined as the period from the time power is applied to the 82578 and until the 82578 powers up its PHY. the 82578 should consume less than ~40 mA during this period. Following the 82578 PHY entering reset, the power-up sequence is considered done and the requirement is removed.

6.3.1.2 Cable Disconnect State

The 82578 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 82578 detects energy on the MDI link, and starts the following exit sequence:

- The 82578 signals the MAC 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.

6.3.1.3 Power Down State

The 82578 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.

Figure 7 shows the power-down sequence.




Figure 7. Power-Down Sequence

Note: If the LAN_DISABLE_N pin cannot be used, a power-down in-band can be used. When used, the power savings are lower since all logic cannot be turned off in this mode.



Table 8. Figure 7 Notes

Note	Description
1	MAC sends an in-band power-down message through SMBus or PCIe or LAN_DISABLE_N pin set to zero.
2	Once the 82578 detects the power-down message or 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	MAC configures the 82578 through the MDIO interface.
6	PHY goes through auto-negotiation to acquire link.

6.4 Power Saving Features

This section provides information about the low power configurations for the 82578.

6.4.1 Intel® Auto Connect Battery Saver (ACBS)

Intel® Auto Connect Battery Saver for the 82578 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.

Note: ACBS is only supported during auto negotiation. If link is forced, the 82578 does not enter ACBS mode.

82578 ACBS works in both S0 and Sx states. Since 82578 ACBS has no driver control, the feature is always enabled, allowing power savings by default.

Note: The crystal clock drivers are intermittently disabled when the network cable is unplugged and the 82578 is in ACBS mode.

6.4.2 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.



6.4.2.1 Link Speed Battery Saver

Link speed battery saver is a power saving feature that negotiates to the lowest speed possible when the 82578 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 Window's driver.

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

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

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

6.4.2.2 System Idle Power Saver (SIPS)

SIPS is a software-based power saving feature that is enabled only with Microsoft* Windows* Vista* and 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 82578. Refer to [Section 6.1](#) for the power dissipated in each link state.

SIPS is designed to save power in the 82578 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.



6.4.2.3 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 82578 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 MAC sets and clears the bit according to hardware/software settings. The 82578 auto negotiates with the updated LPLU setting on the following auto-negotiation operation. The 82578 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 82578 in the various speeds.



6.4.2.4 LAN Disable Recommendations

LAN_DISABLE_N needs to be connected to the GPIO12/LAN_PHY_PWR_CTRL output of the Intel® 5 Series Express Chipset. GPIO12 also needs to be configured using Intel® 5 Series Express Chipset soft straps as LAN_PHY_PWR_CTRL (bit [20] of PCHSTRP0 register - LAN_PHY_PWR_CTRL/GPIO12. Refer to the *Intel® 5 Series Express Chipset Family External Design Specification (Intel® 5 Series Express Chipset EDS)*.



Note: This page intentionally left blank.



7.0 Device Functionality

7.1 Tx Flow

When packets are ready for transmission in the MAC it transfers them to the 82578 through the PCIe or the SMBus (depending on system state). The 82578 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; however, this behavior has no dependency on link speed. The 82578 design is based on the assumption that the MAC has the full packet ready for transmission.

7.2 Rx Flow

The 82578 maintains a FIFO on the receive side in order not to lose packets when PCIe is active. In this case, the 82578 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 82578 starts the transmission after the entire packet is received.

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 MAC. Some of the flow control functionality has moved to the 82578. The following registers are duplicated to the 82578 for the implementation of flow control:

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

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



7.3.1 MAC Control Frames and Reception of Flow Control Packets

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

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

The 802.3x standard defines the MAC control frame multicast address as 01-80-C2-00-00-01. The flow control packet's *Type* field is checked to determine if it is a valid flow control packet: XON or XOFF. 802.3x reserves this as 0x8808. The final check for a valid PAUSE frame is the *MAC Control Opcode* field. At this time only the PAUSE control frame opcode is defined and 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 hardwired to 64 byte times.

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

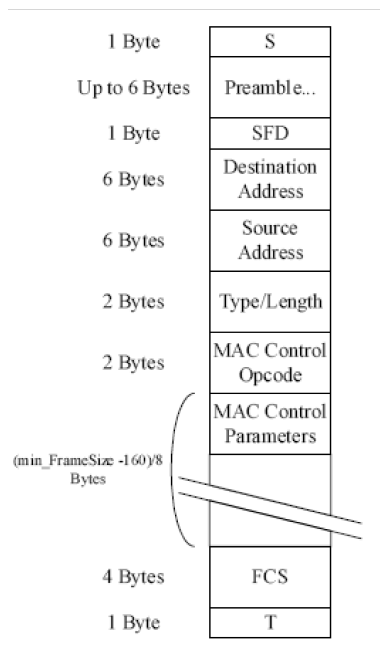


Figure 8. 802.3x MAC Control Frame Format

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 82578 has validated the reception of an XOFF, or PAUSE frame, it does the following:

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

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 82578 has validated the reception of an XON frame, it does the following:

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

7.3.2 Transmitting PAUSE Frames

Transmitting PAUSE frames is done as a result of an In-Band Control command from the MAC. The MAC 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 MAC is received, the 82578 sends a PAUSE frame with its *Pause Timer* field equal to FCTTV. Once the receive buffer fullness reaches the low water mark, the MAC 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 82578 is operating in half-duplex mode.

7.4 Wake Up

The 82578 supports host wake up.

This mechanism uses in-band messages to wake the Intel® 5 Series Express Chipset from a sleep state. The host can enable host wake up from the 82578 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 82578 to the Intel® 5 Series Express Chipset would come in as an in-band message over the SMBus.

1. The XON frame is also forwarded to the MAC.



Setting the 82578's wake up:

1. Clear the *Host_WU_Active* bit (bit 4) in the Port General Configuration register (page 769, register 17) to enable wake up mode.
2. Set bit 2 (*MACPD_enable*) of the Port Control register (page 769, register 17) to enable the 82578 wake up capability and software accesses to page 800.
3. Set the *Slave Access Enable* bit (bit 2) in the Receive Control register (page 800, register 0) to enable access to the Flex Filter register, if setting those bits is needed in the next stage. The registers affected are:
 - a. Flexible Filter Value Table LSB- FFVT_L (filters 01)
 - b. Flexible Filter Value Table MSBs - FFVT_H (filters 23)
 - c. Flexible Filter Value Table - FFVT_45 (filters 45)
 - d. Flexible TCO Filter Value/Mask Table LSBs - FTFT_L
 - e. Flexible TCO Filter Value/Mask Table MSBs - FTFT_H
4. Configure the 82578'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 82578's wake up functionality.

Note: Once wake up is enabled, the 82578 stops responding to SMBus commands.

Host wake up:

1. When a WoL packet/event is detected, the 82578 sends an in-band message to the Intel® 5 Series Express Chipset indicating a host wake up.
2. The Intel® 5 Series Express Chipset wakes the host.
3. The host should issue an PHY reset to the 82578 before clearing the *Host_WU_Active* bit.
4. Host reads the Wake Up Status (WUS) register; wake up status from the 82578).

The 82578 keeps and forwards the wake up packet. 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). As a result, the 82578 resumes transmitting the packet. Each time this bit is set and if a wake up in-band message has already sent, any new packets received does not overwrite the packet in the FIFO. The 82578 re-transmits the wake up in-band message after 50 ms if no change in the *Host_WU_Active* bit occurred.

7.4.1 Host Wake Up

The 82578 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 82578's wake up functionality is enabled, the *APM Enable* bits from the NVM are written to the 82578 by the Intel® 5 Series Express Chipset 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 82578 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 82578 receives a matching magic packet, it:

- Sets the *Magic Packet Received* bit in the WUS register.
- Initiates the Intel® 5 Series Express Chipset 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 82578 detects a link status change it:

- Sets the *Link Status Changed* (LNKC) bit (bit 0) in the WUS register.
- Initiates the Intel® 5 Series Express Chipset 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 82578 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 82578 in wake up mode.

Once wake up is enabled, the 82578 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 82578:

- Initiates a the Intel® 5 Series Express Chipset 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 82578 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 82578's pre-defined filters:

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

Each of these 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 82578 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 82578 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 previous paragraph.

7.4.1.3.1.3 Broadcast

If the *Broadcast Wake Up Enable* bit (bit 4) in the WUFC register is set, the 82578 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 82578 has been put into Magic Packet mode, it scans all incoming frames addressed to the node for a specific data sequence, which indicates to the MAC that this is a Magic Packet frame. A Magic Packet frame must also meet the basic requirements for the LAN technology chosen, such as Source address, Destination Address (which might be the receiving station's IEEE address or a Multicast address that includes the Broadcast address) and CRC. The specific data sequence consists of 16 duplications of the IEEE address of this node with no breaks or interruptions. This sequence can be located anywhere within the packet, but must be preceded by a synchronization stream. The synchronization stream enables the scanning state machine to be much simpler. The synchronization stream is defined as 6 bytes of 0xFF. The device also accepts a Broadcast frame, as long as the 16 duplications of the IEEE address match the address of the system that needs to wake up.

The 82578 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 82578 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 82578 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 82578 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 82578 starts up with the RAH0/RAL0 register 0 loaded from the NVM. This enables the 82578 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	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	4	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 IPv4 Request Packet

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, and one of the four programmed IPv4 addresses.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header – processed by main address filter
6	6	Source Address		Skip	
12	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	2	Type	0x0806	Compare	
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



7.4.1.3.1.6 Directed IPv4 Packet

The 82578 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 82578 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	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	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

7.4.1.3.1.7 Directed IPv6 Packet

The 82578 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 82578 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	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	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	



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

7.4.1.3.2 Flexible Filter

The 82578 supports a total of six 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.

7.4.1.3.2.1 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	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	2	Type	0x8137	Compare	IPX
14	16	Some IPX Information	-	Ignore	
30	2	IPX Diagnostic Socket	0x0456	Compare	

7.4.1.3.2.2 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	4	Possible VLAN Tag		Skip	
12	8	Possible Len/LLC/SNAP Header		Skip	
12	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

7.4.1.3.2.3 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 82578'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) until the page is not changed to a different value wake up register access is enabled. Refer to [Section 8.10.1](#) for more details.

After the page is 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 [0..15] range meaning PHY registers, the functionality remains unchanged.

There are two valid instructions:

Instruction	Address	Description
Address set	0x11	Wake up space address is set for either reading or writing.
Data cycle	0x12	Wake up space accesses read or write cycle.




7.5 PHY Loopback

PHY loopback is supported in the 82578. Software or firmware should set the 82578 to the loopback mode (via the MDIC register) writing to the PHY Loopback Control register (address 19). The MAC must be in forced link and in full duplex mode for PHY loopback to operate. The following bits must be configured to enable PHY loopback:

CTRL.FRCDPLX = 1b: // force duplex mode by the MAC

CTRL.FD = 1b: // Set full-duplex mode



Note: This page intentionally left blank.



8.0 Programmer's Visible State

8.1 Terminology

Shorthand	Description
R/W	Read/Write. A register with this attribute can be read and written. If written since reset, the value read reflects the value written.
R/W S	Read/Write Status. A register with this attribute can be read and written. This bit represents status of some sort, so the value read might not reflect the value written.
RO	Read Only. If a register is read only, writes to this register have no effect.
WO	Write Only. Reading this register might not return a meaningful value.
R/WC	Read/Write Clear. A register bit with this attribute can be read and written. However, a write of 1b clears (sets to 0b) the corresponding bit and a write of 0b has no effect.
R/W SC	Read/Write Self Clearing. When written to 1b the bit causes an action to be initiated. Once the action is complete the bit return to 0b.
RO/LH	Read Only, Latch High. The bit records an event or the occurrence of a condition to be recorded. When the event occurs the bit is set to 1b. After the bit is read, it returns to 0b unless the event is still occurring.
RO/LL	Read Only, Latch Low. The bit records an event. When the event occurs the bit is set to 0b. After the bit is read, it reflects the current status.
RO/SC	Read Only, Self Clear. Writes to this register have no effect. Reading the register clears (set to 0b) the corresponding bits.
RW0	Ignore Read, Write Zero. The bit is a reserved bit. Any values read should be ignored. When writing to this bit always write as 0b.
RWP	Ignore Read, Write Preserving. This bit is a reserved bit. Any values read should be ignored. However, they must be saved. When writing the register the value read out must be written back. (There are currently no bits that have this definition.)
Update	Value written to the register field does not take effect until a software reset is executed. The value can still be read after it is written.
Retain	Value written to a register field does take effect without a software reset.

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 MAC during the 82578 power up and following reset. Other fields in the same 16-bit register must be loaded with their default values.

8.2 MDIO Access

After PHY reset, a delay of 10 ms is required before any register access using MDIO.

8.3 Addressing

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

The 82578 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 82578 contains registers controlling the custom interface as well as other the 82578 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 82578 registers in PHY address 01, are divided into pages. Each page has 32 registers. Registers 0-15 are identical in all the pages and are the IEEE defined registers. Register 31 is the page register in all pages of PHY address 01. All other registers are page specific.

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

In pages 800 and 801, the register address space is more than 32. See [section 8.9](#) for a description of registers addressing in these pages.

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



8.4 Address Map

Table 9. Address Map

PHY Address	Page	Register	Name	Table #
02	Any	0	Control Register	Table 10
02	Any	1	Status Register	Table 11
02	Any	2	PHY Identifier [18:3]	Table 12
02	Any	3	PHY Identifier [19:24]	Table 13
02	Any	4	Auto-Negotiation Advertisement	Table 14
02	Any	5	Link Partner Ability (Base Page) Base	Table 15
02	Any	6	Auto-Negotiation Expansion	Table 16
02	Any	7	Next Page Transmit	Table 17
02	Any	8	Link Partner Next Page	Table 18
02	Any	9	1000BASE-T Control	Table 19
02	Any	10	1000BASE-T Status	Table 20
02	Any	15	Extended Status	Table 21
02	0	16	Function Control	Table 22
02	0	17	PHY-Specific Status	Table 23
02	0	18	Interrupt Enable	Table 24
02	0	19	Interrupt Status	Table 25
02	0	20	Extended PHY-Specific Control	Table 26
02	0	21	Receive Error Counter	Table 27
02	0	22	Cable Defect Tester Control	Table 28
02	0	24	LED Control	Table 29
02	0	25	Manual LED Override	Table 30
02	0	28	Cable Defect Tester Status	Table 31
02	0	29	Debug Port Address Offset	Table 32
02	0	30	Debug Port Data	Table 33
Page 769 – Port Control Registers				
01	769	17	Port General Configuration	Table 41
01	769	21	Power Management Control	Table 42
01	769	25	Rate Adaptation Control	Table 44
01	769	27	Flow Control Transmit Timer Value	Table 45
Page 778 – Statistics Registers				
01	778	16 - 17	Single Collision Count	Table 46
01	778	18 - 19	Excessive Collisions Count	Table 47
01	778	20 - 21	Multiple Collisions Count	Table 48
01	778	23 - 24	Late Collision Count	Table 49
01	778	25 - 26	Collision Count	Table 50
01	778	27 - 28	Defer Count	Table 51
01	778	29 - 30	Transmit with No CRS - TNCRS	Table 52
PCIe Registers				



Table 9. Address Map

01	770	16	PCIe FIFOs Control/Status	Table 53
01	770	17	PCIe Power Management Control	Table 54
01	770	18	In-Band Control	Table 55
01	770	20	PCIe Diagnostics	Table 56
01	770	21	Timeouts	Table 57
01	770	23	PCIe K-State Minimum Duration Timeout	Table 58
General Registers				
01	776	19	82578 Capability Register	Table 59
01	0	25	OEM Bits	Table 60
01	0	26	SMBus Address	Table 61
01	0	27-28	Shadow Register for RAL0[31:0].	Table 62
01	0	29	Shadow Register for RAH0[15:0].	Table 63
01	0	30	LED Configuration	Table 64
Page 800 - Wake Up Registers				
01	800	0	Receive Control Register	Table 65
01	800	1	Wake Up Control Register	Table 66
01	800	2	Wake Up Filter Control Register	Table 67
01	800	3	Wake Up Status Register	Table 68
01	800	16	Receive Address Low 0	Table 69
01	800	18	Receive Address High 0	Table 70
01	800	44 - 45	Shared Receive Address Low 0	Table 71
01	800	46 - 47	Shared Receive Address High 0	Table 72
01	800	58 - 59	Shared Receive Address High 3	Table 73
01	800	64	IP Address Valid - IPAV	Table 74
01	800	82 - 83	IPv4 Address Table - IP4AT 0	Table 75
01	800	88 - 89	IPv6 Address Table - IP6AT 0	Table 76
01	800	128 - 191	Multicast Table Array - MTA[31:0]	Table 77
01	800	256 + 2*n (n = 0 - 127)	Flexible Filter Value Table LSB- FTVT_01	Table 78
01	800	257 + 2*n (n = 0 - 127)	Flexible Filter Value Table MSB - FTVT_23	Table 79
01	800	512 + 2*n (n = 0 - 127)	Flexible Filter Value Table - FTVT_45	Table 80
01	800	768 + n (n = 0 - 127)	Flexible Filter Mask Table - FFMT	Table 81
01	800	896 + n (n = 0 - 3)	Flexible Filter Length Table - FFLT03	Table 82
01	800	904 + n (n=0...1)	Flexible Filter Length Table - FFLT45	Table 83



8.5 PHY Registers (Page 0)

Table 10. Control Register PHY Address 02, Page Any, Register 0

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reset	R/W, SC	0b	SC	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	0b	When loopback is active, the transmitter data on TXD loops back to RXD internally. The link breaks when loopback is enabled. 1b = Enable loopback. 0b = Disable loopback.
13	Speed Select (LSB)	R/W	0b	Retain	Bit 6, 13 11b = Reserved. 10b = 1000 Mb/s. 01b = 100 Mb/s. 00b = 10 Mb/s.
12	Auto-Negotiation Enable	R/W	0b	Retain	1b = Enable auto-negotiation process. 0b = Disable auto-negotiation process.
11	Power Down	R/W	0b	0b	1b = Power down. 0b = Normal operation.
10	Reserved	RO	0b	0b	Reserved
9	Restart Auto-Negotiation	R/W,SC	0b	SC	Auto-negotiation automatically restarts after a hardware or software reset regardless of whether or not this bit is set. 1b = Restart auto-negotiation process. 0b = Normal operation.
8	Duplex Mode	R/W	0b	Retain	1b = Full-duplex. 0b = Half-duplex.
7	Collision Test	RO	0b	0b	Setting this bit to 1b causes the COL pin to assert each time the TX_EN pin is asserted. This bit takes effect only while in 10 Mb/s loopback mode.
6	Speed Selection (MSB)	R/W	0b	Update	See description in bit 13.
5:0	Reserved	RO	Always 0x0	Always 0x0	Reserved, always set to 0x0.

Table 11. Status Register PHY Address 02, Page Any, Register 1

Bits	Field	Mode	HW Rst	SW Rst	Description
15	100BASE-T4	RO	0b	0b	100BASE-T4. This protocol is not available. 0b = PHY not able to perform 100BASE-T4.
14	100BASE-X Full-Duplex	RO	1b	1b	1b = PHY able to perform full-duplex 100BASE-X.
13	100BASE-X Half-Duplex	RO	1b	1b	1b = PHY able to perform half-duplex 100BASE-X.
12	10 Mbps Full-Duplex	RO	1b	1b	1b = PHY able to perform full-duplex 10BASE-T.
11	10 Mbps Half-Duplex	RO	1b	1b	1b = PHY able to perform half-duplex 10BASE-T.



Table 11. Status Register PHY Address 02, Page Any, Register 1

Bits	Field	Mode	HW Rst	SW Rst	Description
10	100BASE-T2 Full-Duplex	RO	0b	0b	Not able to perform 100 Base-T2.
9	100BASE-T2 Half-Duplex	RO	0b	0b	Not able to perform 100 Base-T2.
8	Extended Status	RO	1b	1b	Extended status information in the register Extended Status.
7	Reserved	RO	0b	0b	Must always be set to 0b.
6	MF Preamble Suppression	RO	1b	1b	1b = PHY accepts management frames with preamble suppressed.
5	Auto-Negotiation Complete	RO	0b	0b	1b = Auto-negotiation process complete. 0b = Auto-negotiation process not complete.
4	Remote Fault	RO,LH	0b	0b	1b = Remote fault condition detected. 0b = Remote fault condition not detected.
3	Auto-Negotiation Ability	RO	1b	1b	1b = PHY able to perform auto-negotiation.
2	Link Status	RO,LL	0b	0b	Indicates whether the link was lost since the last read. For the current link status, read LINK_REAL_TIME (bit [10]) of the register "PHY-Specific Status". Latching low function. 1b = Link is up. 0b = Link is down.
1	Jabber Detect	RO,LH	0b	0b	1b = Jabber condition detected 0 = Jabber condition not detected.
0	Extended Capability	RO	1b	1b	1b = Extended register capabilities.

Table 12. PHY Identifier [18:3] PHY Address 02, Page Any, Register 2

Bits	Field	Type	Default	Description
15:0	Unique Identifier Bits 18:3	RO	0x004D	Organizationally Unique Identifier (OUI), bits [18:3].

Table 13. PHY Identifier [19:24] PHY Address 02, Page Any, Register 3

Bits	Field	Type	Default	Description
15:10	PHY Identifier Bits 24:19	RO	110100b	OUI, bits [24:19].
9:4	Model Number	RO	000100b	The value is part of the PHY identifier and represents the Device Model Number.
3:0	Revision Number	RO	0x2	The value is part of the PHY identifier and represents the Device Revision Number.



Table 14. Auto-Negotiation Advertisement PHY Address 02, Page Any, Register 4

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Next Page	R/W	0b	Update	If 1000BASE-T is advertised then the required next pages are automatically transmitted. This bit should be set to 0b if no additional next pages are needed. 1b = Advertise. 0b = Not advertised.
14	Ack	RO	Always 0b	Always 0b	Must be 0b.
13	Remote Fault	R/W	0b	Update	Write a 1b to set remote fault.
12	Reserved	R/W	0b	Update	Reserved
11	Asymmetric Pause	R/W	See Descr.	Update	Write a 1b to set asymmetric pause.
10	Pause	R/W	See Descr.	Update	Write a 1b to set pause.
9	100BASE-T4	R/W	0b	Retain	Not able to perform 100 Base-T4.
8	100BASE-TX Full-Duplex	R/W	1b	Update	Write a 1b to advertise.
7	100BASE-TX Half-Duplex	R/W	1b	Update	Write a 1b to advertise.
6	10BASE-TX Full-Duplex	R/W	1b	Update	Write a 1b to advertise.
5	10BASE-TX Half-Duplex	R/W	1b	Update	Write a 1b to advertise.
4:0	Selector Field	R/W	0x01	Retain	Selector Field mode 00001b = 802.3.

Table 15. Link Partner Ability (Base Page) Base PHY Address 02, Page Any, Register 5

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Next Page	RO	0b	0b	Received Code Word Bit 15 1b = Link partner capable of next page. 0b = Link partner not capable of next page.
14	Acknowledge	RO	0b	0b	Acknowledge Received Code Word Bit 14 1b = Link partner received link code word. 0b = Link partner does not have Next Page ability.
13	Remote Fault	RO	0b	0b	Remote Fault Received Code Word Bit 13 1b = Link partner detected remote fault. 0b = Link partner has not detected remote fault.
12	Technology Ability Field	RO	0b	0b	Received Code Word Bit 12
11	Asymmetric Pause	RO	0b	0b	Received Code Word Bit 11 1b = Link partner requests asymmetric pause. 0b = Link partner does not request asymmetric pause.
10	Pause Capable	RO	0b	0b	Received Code Word Bit 10 1b = Link partner is capable of pause operation. 0b = Link partner is not capable of pause operation.
9	100BASE-T4 Capability	RO	0b	0b	Received Code Word Bit 9 1b = Link partner is 100BASE-T4 capable. 0b = Link partner is not 100BASE-T4 capable.



Table 15. Link Partner Ability (Base Page) Base PHY Address 02, Page Any, Register 5

Bits	Field	Mode	HW Rst	SW Rst	Description
8	100BASE-TX Full-Duplex Capability	RO	0b	0b	Received Code Word Bit 8 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	0b	Received Code Word Bit 7 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	0b	Received Code Word Bit 6 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	0b	Received Code Word Bit 5 1b = Link partner is 10BASE-T half-duplex capable. 0b = Link partner is not 10BASE-T half-duplex capable.
4:0	Selector Field	RO	0x00	0x00	Selector Field Received Code Word Bit 4:0

Table 16. Auto-Negotiation Expansion PHY Address 02, Page Any, Register 6

Bits	Field	Mode	HW Rst	SW Rst	Description
15:5	Reserved	RO	0x000	0x000	Reserved, must be set to 0x000.
4	Parallel Detection Fault	RO,LH	0b	0b	Software resets this bit to 0b; clears it after a read. 1b = A fault has been detected. 0b = A fault has not been detected.
3	Link Partner Next page Able	RO	0b	0b	Software resets this bit to 0b; clears it after a read. 1b = Link partner is next page able. 0b = Link partner is not next page able.
2	Local Next Page Able	RO	1b	1b	The software reset value is determined by bit [15] of the Auto-Negotiation Advertisement register and by bits [9:8] of the 1000 Base-T Control register. 1b = Local device is next page able. 0b = Local device is not next page able.
1	Page Received	RO, LH	0b	0b	On software reset, this bit value is reserved; LH; cleared after a read. 1b = A new page has been received. 0b = A new page has not been received.
0	Link Partner Auto-Negotiation Able	RO	0b	0b	Software resets this bit to 0b. 1b = Link partner is auto-negotiation able. 0b = Link partner is not auto-negotiation able.

Table 17. Next Page Transmit PHY Address 02, Page Any, Register 7

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Next Page	R/W	0b	0b	Transmit code word bit 15.
14	Reserved	RO	0b	0b	Transmit code word bit 14.
13	Message Page Mode	R/W	1b	1b	Transmit code word bit 13.



Table 17. Next Page Transmit PHY Address 02, Page Any, Register 7

Bits	Field	Mode	HW Rst	SW Rst	Description
12	Acknowledge2	R/W	0b	0b	Transmit code word bit Bit 12.
11	Toggle	RO	0b	0b	Transmit code word bit 11.
10:0	Message/ Unformatted Field	R/W	0x001	0x001	Transmit code word bit 10:0.

Table 18. Link Partner Next Page PHY Address 02, Page Any, Register 8

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Next Page	RO	0b	0b	Received code word bit 15.
14	Acknowledge	RO	0b	0b	Received code word bit 14.
13	Message Page	RO	0b	0b	Received code word bit 13.
12	Acknowledge2	RO	0b	0b	Received code word bit 12.
11	Toggle	RO	0b	0b	Received code word bit 11.
10:0	Message Unformatted Field	RO	0x000	0x000	Received code word bit 10:0.

Table 19. 1000BASE-T Control PHY Address 02, Page Any, Register 9

Bits	Field	Mode	HW Rst	SW Rst	Description
15:13	Test Mode	R/W	000b	Retain	The TX_TCLK is provided by the RX_CLK pin for jitter testing in test modes 2 and 3. Hardware reset or software reset (see RESET (bit [15]) of the register Control) should be issued to ensure normal operation after exiting the test mode. 000b = Normal mode. 001b = Test Mode 1 - Transmit waveform test. 010b = Test Mode 2 - Transmit jitter test (master mode). 011b = Test Mode 3 - Transmit jitter test (salve mode). 100b = Test Mode 4 - Transmit distortion test. 101b, 110b, 111b = Reserved.
12	Master/Slave Manual Configuration Enable	R/W	0b	Update	Master/Slave Configuration Control 1b = Manual master/slave configuration. 0b = Automatic master/slave configuration.
11	Master/Slave Configuration Value	R/W	See Descr.	Update	Ignored if bit [12] equals 0b. 1b = Manual configure as master. 0b = Manual configure as salve.
10	Port Type	R/W	See Descr.	Update	Ignored if bit [12] equals 1b. 1b = Prefer multi-port device (master). 0b = Prefer single port device (slave).
9	1000BASE-T Full-Duplex	R/W	1b	Update	Write a 1b to advertise.
8	1000BASE-T Half-Duplex	R/W	See Descr.	Update	Write a 1b to advertise.
7:0	Reserved	R/W	0x00	Retain	Set these bits to 0x00.



Table 20. 1000BASE-T Status PHY Address 02, Page Any, Register 10

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Master/Slave Configuration Fault	RO,LH	0b	0b	This register bit gets cleared on a read. 1b = Master/slave configuration fault detected. 0b = No master/slave configuration fault detected.
14	Master/Slave Configuration Resolution	RO	0b	0b	This register bit is not valid until the <i>PAGE_RECEIVED</i> (bit [1]) of register Auto-Negotiation Expansion is set to 1b. 1b = Local PHY configuration resolved to the master. 0b = Local PHY configuration resolved to the slave.
13	Local Receiver Status	RO	0b	0b	1b = Local receiver is correct. 0b = Local receiver is incorrect.
12	Remote Receiver Status	RO	0b	0b	1b = Remote receiver is correct. 0b = Remote receiver is incorrect.
11	Link Partner 1000BASE-T Full-Duplex Capability	RO	0b	0b	This register bit is not valid until <i>PAGE_RECEIVED</i> (bit [1]) of register Auto-Negotiation Expansion is set to 1b. 1b = Link partner is capable of 1000BASE-T full-duplex. 0b = Link partner is not capable of 1000BASE-T full duplex.
10	Link Partner 1000BASE-T Half-Duplex Capability	RO	0b	0b	This register bit is not valid until <i>PAGE_RECEIVED</i> (bit [1]) of register Auto-Negotiation Expansion is set to 1b. 1b = Link partner is capable of 1000BASE-T half-duplex 0b = Link partner is not capable of 1000BASE-T half duplex
9:8	Reserved	RO	00b	00b	Reserved
7:0	Idle Error Count	RO, SC	0x00	0x00	Reports the idle error count since the last time this register was read. The counter stops at 11111111 and does not roll over. These bits are cleared on a read.

Table 21. Extended Status PHY Address 02, Page Any, Register 15

Bits	Field	Mode	HW Rst	SW Rst	Description
15	1000BASE-X Full-Duplex	RO	Always 0b	Always 0b	PHY not able to perform 1000BASE-X full duplex.
14	1000BASE-X Half-Duplex	RO	Always 0b	Always 0b	PHY not able to perform 1000BASE-X half duplex.
13	1000BASE-T Full-Duplex	RO	Always 1b	Always 1b	PHY able to perform 1000BASE-T full duplex.
12	1000BASE-T Half-Duplex	RO	Always 0b	Always 0b	PHY not able to perform 1000BASE-T half duplex.
11:0	Reserved	RO	0x000	0x000	Reserved



Table 22. Function Control 1 PHY Address 02, Page 0, Register 16

Bits	Field	Mode	HW Rst	SW Rst	Description
15:12	Reserved	R/W	0x0	Update	Reserved
11	Assert_CRS_On_Transmit	R/W	0b	Update	This bit has no effect in full-duplex. 0b = Never assert on transmit. 1b = Assert on transmit.
10	Reserved	R/W	0b	Retain	Reserved
9:8	Energy Detect	R/W	See Descr.	Update	00b = Off. 10b = Sense only on receive (energy detect). 11b = Sense and periodically transmit NLP.
7	Reserved	R/W	0x0	Retain	Reserved
6:5	MDI Crossover Mode	R/W	0x3	Update	Changes to these bits are disruptive to normal operation. As a result, any changes to these registers must be followed by a software reset to take effect. 00b = Manual MDI configuration. 01b = Manual MDIX configuration. 10b = Reserved. 11b = Enable automatic crossover for all modes.
4:3	Reserved	R/W	0x0	Retain	Reserved
2	SQE_TEST	R/W	0b	Retain	SQE test is automatically disabled in full-duplex mode regardless of the state of this bit. 0b = SQE test disabled. 1b = SQE test enabled.
1	Polarity Reversal Disable	R/W	0b	Retain	If polarity is disabled, then polarity is forced to be normal in 10BASE-T. 1b = Polarity reversal disabled. 0b = Polarity reversal enabled.
0	Disable Jabber	R/W	0b	Retain	1b = Disable jabber function. 0b = Enable jabber function.

Table 23. PHY-Specific Status 1 PHY Address 02, Page 0, Register 17

Bits	Field	Mode	HW Rst	SW Rst	Description
15:14	Speed	RO	0x2	Retain	Valid only after resolved bit [11] of this register = 1b. The resolved bit is set when auto-negotiation completed or auto-negotiation is disabled. 11b = Reserved. 10b = 1000 Mb/s. 01b = 100 Mb/s. 00b = 10 Mb/s.
13	Duplex	RO	0b	Retain	Valid only after resolved bit [11] of this register = 1b. The resolved bit is set when auto-negotiation completed or auto-negotiation is disabled. 1 = Full-duplex. 0 = Half-duplex.
12	Page Received	RO, LH	0b	0b	1b = Page received. 0 b= Page not received.
11	Speed and Duplex Resolved	RO	0b	0b	When auto-negotiation is disabled, this bit = 1b for force speed. 1b = Resolved. 0b = Not resolved.



Table 23. PHY-Specific Status 1 PHY Address 02, Page 0, Register 17

Bits	Field	Mode	HW Rst	SW Rst	Description
10	Copper Link (real time)	RO	0b	0b	1b = Link up. 0b = Link down.
9:7	Reserved	RO	000b	000b	Reserved, always set to 000b.
6	MDI Crossover Status	RO	1b	Retain	Valid only after resolved bit [11] of this register = 1b. The resolved bit is set when auto-negotiation completed or auto-negotiation is disabled. This bit is 0b or 1b depending on what is written to bits [6:5] of register Function Control in manual configuration mode. Function Control bits [6:5] are updated with a software reset. 1b = MDI-X. 0b = MDI.
5	Smartspeed_Downgrade	RO	0b	0b	1b = Smartspeed downgrade occurs. 0b = Smartspeed downgrade does not occur.
4	Energy Detect Status	RO	0b	0b	1b = Sleep. 0b = Active.
3	Transmit_Pause_Enabled	RO	0b	0b	Valid only after resolved bit [11] of this register = 1b. The resolved bit is set when auto-negotiation completed or disabled. A reflection of the MAC pause resolution. 0b = Transmit pause disabled. 1b = Transmit pause enabled.
2	Receive_Pause_Enabled	RO	0b	0b	A reflection of the MAC pause resolution. This status bit is valid only after resolved bit [11] of this register = 1b. The resolved bit is set when auto-negotiation completed or is disabled. 0b = Receive pause disabled. 1b = Receive pause enabled.
1	Polarity (real time)	RO	0b	0b	1b = Reversed. 0b = Normal.
0	Jabber (real time)	RO	0b	0b	1b = Jabber. 0b = No jabber.

Table 24. Interrupt Enable PHY Address 02, Page 0, Register 18

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Auto-Negotiation Error Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
14	Speed Changed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
13	Duplex Changed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
12	Page Received Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
11	Auto-Negotiation Completed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
10	Link Status Changed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
9	Symbol Error Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.



Table 24. Interrupt Enable PHY Address 02, Page 0, Register 18

Bits	Field	Mode	HW Rst	SW Rst	Description
8	False Carrier Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
7	FIFO Over/ Underflow Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
6	MDI Crossover Changed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
5	Smartspeed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
4	Energy Detect Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
3:2	Reserved	R/W	00b	00b	Reserved, always set to 00b.
1	Polarity Changed Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.
0	Jabber Interrupt Enable	R/W	0b	Retain	1b = Interrupt enable. 0b = Interrupt disable.

Table 25. Interrupt Status PHY Address 02, Page 0, Register 19

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Auto- Negotiation Error	RO,LH	0b	0b	An error occurs if the master/slave does not resolve, parallel detect fault, no common HCD, or link does not come up after negotiation completed. 0b = No auto-negotiation error 1b = Auto-negotiation error.
14	Speed Changed	RO,LH	0b	0b	1b = Speed changed. 0b = Speed not changed.
13	Duplex Changed	RO,LH	0b	0b	1b = Duplex changed. 0b = Duplex not changed.
12	Page Received	RO,LH	0b	0b	1b = Page received. 0b = Page not received.
11	Auto-Negotiation Completed	RO,LH	0b	0b	1b = Auto-negotiation completed. 0b = Auto-negotiation not completed.
10	Link Status Changed	RO,LH	0b	0b	1b = Link status changed. 0b = Link status not changed.
9	Symbol Error	RO,LH	0b	0b	1b = Symbol error. 0b = No symbol error.
8	False Carrier	RO,LH	0b	0b	1b = False carrier. 0b = No false carrier.
7	FIFO_Over/_Underflow	RO	0b	0b	0b = No FIFO error 1b = Over/underflow error.
6	MDI Crossover Changed	RO,LH	0b	0b	1b = Crossover changed. 0b = Crossover not changed.
5	Smartspeed_Interrupt	RO,LH	0b	0b	0b = No smartspeed interrupt detected. 1b = Smartspeed interrupt detected.
4	Energy Detect Changed	RO,LH	0b	0b	1b = Energy detect state changed. 0b = No energy detect state change detected.



Table 25. Interrupt Status PHY Address 02, Page 0, Register 19

Bits	Field	Mode	HW Rst	SW Rst	Description
3:2	Reserved	RO	0b	0b	
1	Polarity Changed	RO,LH	0b	0b	1b = Polarity changed. 0b = Polarity not changed.
0	Jabber	RO,LH	0b	0b	1b = Jabber. 0b = No jabber.

Table 26. Extended PHY-Specific Control PHY Address 02, Page 0, Register 20

Bits	Field	Mode	HW Rst	SW Rst	Description
15:11	Reserved	R/W	0x000	Retain	Reserved, must be set to 0x000.
10	aneg_now_qual	R/W	0b	Retain	Setting this bit to 1b causes the PHY to restart auto-negotiation. This bit is self clearing.
9	Rev_aneg_qual	R/W	0b	0b	Makes the PHY auto-negotiate in reversed mode. This bit takes its value from the input pin rev_aneg by the following: <ul style="list-style-type: none"> • Hardware reset (fall of rst_dsp_i). • PHY software reset. • Rise of aneg_now.
8	Giga_dis_qual	R/W	0b	0b	Makes the PHY disable GbE mode. This bit takes its value from the input pin giga_dis by the following: <ul style="list-style-type: none"> • Hardware reset (fall of rst_dsp_i). • PHY software reset. • Rise of aneg_now.
7	Cfg_dis_qual	R/W	0b	0b	If this bit is set to 1b, then the auto-negotiation arbitration FSM bypasses the LINK_STATUS_CHECK state when the 10BASE-T/100BASE-T ready signal is asserted. The default value is 0b.
6	Mr_Itdis	R/W	0b	0b	If this bit is set to 1b, then the NLP receive link integrity test FSM stays at the NLP_TEST_PASS state.
5	Smartspeed_En	R/W	1b	1b	The default value is 1b if this bit is set to 1b and the cable inhibits completion of the training phase, then after a few failed attempts, the 82578 automatically adjusts the highest ability to the next lower speed: from 1000 Mb/s to 100 Mb/s to 10 Mb/s.



Table 26. Extended PHY-Specific Control PHY Address 02, Page 0, Register 20

Bits	Field	Mode	HW Rst	SW Rst	Description
4:2	Smartspeed_Retry_Limit	R/W	011b	011b	The default value is three; if set to three, then the device attempts five times before adjusting; the number of attempts can be changed through setting these bits. 000b = 2 retries. 001b = 3 retries. 010b = 4 retries. 011b = 5 retries (default) 100 = 6 retries. 101b = 7 retries. 110b = 8 retries. 111b = 9 retries.
1	Bypass_Smartspeed_Timer	R/W	0b	0b	0b = The stable link condition is determined 2.5 seconds after the link is established (default). 1b = The stable link condition is determined as soon as the link is established.
0	Reserved	R/W	0b	0b	Reserved. Must be set to 0b.

Table 27. Receive Error Counter PHY Address 02, Page 0, Register 21

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	Receive Error Count	RO	0x00	00x00	Counter reaches its maximum at 0xFFFF and does not roll over (when rx_dv is valid, count rx_er numbers). In this version, only for 100BASE-T and 1000BASE-T).

Table 28. Cable Defect Tester Control PHY Address 02, Page 0, Register 22

Bits	Field	Mode	HW Rst	SW Rst	Description
15:10	Reserved	RO	0x00	00x00	Reserved
9:8	MDI_PAI_Select	R/W	00b	00b	Cable Defect Tester (CDT) control registers use the cable defect tester control registers to select which MDI pair is shown in the Cable Defect Tester Status register. 00b = MDI[0] pair. 01b = MDI[1] pair. 10b = MDI[2] pair. 11b = MDI[3] pair.
7:1	Reserved	RO	0x00	0x00	Reserved
0	Enable_Test	R/W	0b	0b	When set, hardware automatically disables this bit when CDT completes. 0b = Disable CDT test. 1b = Enable CDT test.



Table 29. LED Control PHY Address 02, Page 0, Register 24

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Disable LED	R/W	0b	Retain	0b = Enable. 1b = Disable.
14:12	LED On Time	R/W	100b	Retain	001b = 10 ms. 010b = 21 ms. 011b = 42 ms. 100b = 84 ms. 101b = 168 ms. 110b to 111b = 42 ms.
11	Force Interrupt	RO	0b	0b	Always 0b.
10:8	LED On Time	R/W	001b	Retain	000b = 21 ms. 001b = 42 ms. 010b = 84 ms. 011b = 168 ms. 100b = 330 ms. 101b to 111b = 168 ms.
7:5	Reserved	RO	000b	000b	Reserved
4:3	LED_LINK Control	R/W	00b	Retain	00b = Direct LED mode. 11b = Master/slave LED mode. 01b, 10b = Combined LED modes.
2	LED_DUPLEX	R/W	0b	Retain	0b = Duplex. 1b = Duplex/collision.
1	LED_RX Control	R/W	0b	Retain	1b = Receive activity/link. 0b = Receive activity.
0	Enable_Test	R/W	0b	0b	When set, hardware automatically disables this bit when CDT completes. 0b = Disable CDT test. 1b = Enable CDT test.



Table 30. Manual LED Override PHY Address 02, Page 0, Register 25

Bits	Field	Mode	HW Rst	SW Rst	Description
15:12	Reserved	RO	0x00	00x00	Reserved
11:10	LED_DUPLEX	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.
9:8	LED_LINK10	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.
7:6	LED_LINK100	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.
5:4	LED_LINK1000	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.
3:2	LED_RX	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.
1:0	LED_TX	R/W	00b	Retain	LED off means LED pin output equals high. LED on means LED pin output equals low. 00b = Normal. 01b = Blink. 10b = LED off. 11b = LED on.



Table 31. Cable Defect Tester Status PHY Address 02, Page 0, Register 28

Bits	Field	Mode	HW Rst	SW Rst	Description
15:10	Reserved	RO	0x00	0x00	Reserved
9:8	Status	RO	00b	00b	The content of this register applies to the cable pair selected in the Cable Defect Tester Control register. 00b = Valid test, normal cable (no short or open in cable). 01b = Valid test, short in cable (impedance < 33 Ω). 10b = Valid test, open in cable (impedance > 333 Ω). 11b = Test failed.
7:0	Delta_Time	RO	0x00	0x00	Delta time to indicate distance

Table 32. Debug Port Address Offset PHY Address 02, Page 0, Register 29

Bits	Field	Mode	HW Rst	SW Rst	Description
15:6	Reserved	RO	0x0000	0x0000	Reserved
5:0	Address_Offset	R/W	0x00	0x00	Address index to access the debug registers.

Table 33. Debug Port Data PHY Address 02, Page 0, Register 30

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	Data	R/W	0x024E		Data contents of the debug registers as addressed by the Debug Port Address Offset register.



8.5.1 Extended Debug Port Registers

The following the 82578 registers are referenced by the Debug Port Address Offset and the Debug Port Data registers.

Table 34. System Mode Control 0, Address Offset = 0x00

Bits	Field	Mode	HW Rst	SW Rst	Description
15:6	Reserved	R/W	0x009		Reserved
5:4	TxDAC Class Select	R/W	00b	Retain	Class AB, Class A Select Bit 1xb = 1000BT/100BT/10BT in class A mode. 01b = 1000BT/100BT/10BT in class AB. 00b = 1000BT in class AB mode, while 100BT/10BT in class A mode.
3:0	Reserved	R/W	0x0	Retain	Reserved

Table 35. System Mode Control 3, Address Offset = 0x03

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	R/W	0b	0	Reserved
14	Tx_romi_sw	R/W	0b	Retain	1b = A frame transmitted within a rising of link_status (register17.10) is never transmitted. 0b = Frames are transmitted when link is up.
13	Phy_pll_on	R/W	1b	Retain	PLL Control Bit Makes an AND connection with the input pin phy_pll_on to control PLL. 1b = PLL is always on, except in iddq mode. 0b = PLL is controlled by the hibernate module.
12:11	Reserved	R/W	11b	Retain	Reserved
10	LED test control	R/W	0b	Retain	1b = After power on reset, the LED does not light. 0b = After power on reset, the LED does light for 2.5 seconds.
9:0	Reserved	R/W	0x3FF	Retain	Reserved

Table 36. Hibernation Mode Control Register, Address Offset = 0x0B

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Ps_hib_en	R/W	1b	Retain	Power Hibernate Control Bit 1b = Hibernate enable. 0b = Hibernate disable.
14	Wake_mode	R/W	0b	Retain	1b = PHY wake up by energy detect or wake up pin. 0b = PHY wake up only by energy detect.
13	Reserved	R/W	1b	Retain	Reserved
12	Hib_pulse_sw	R/W	1b	Retain	1b = When hibernate, PHY sends NLP pulse and detects signal from cable. 0b = When hibernate, PHY doesn't send NLP pulse and only detects signal from cable.
11	Gate_25m_en_sw	R/W	1b	Retain	1b = When hibernate, shut off 25m clock of auto-negotiation. 0b = 25m clock to auto-negotiation is not controlled by hibernate.
10:0	Reserved	R/W	0x400	Retain	Reserved



Table 37. 100BASE-TX Test Mode Register, Address Offset = 0x10

Bits	Field	Mode	HW Rst	SW Rst	Description
15	TM100_ENA	R/W	0b	Retain	Enable dig100 loopback test mode.
14:8	Reserved	R/W	0x00	0b	Reserved
7	Jitter_test	R/W	0b	Retain	100BT jitter test.
6	Os_test	R/W	0b	Retain	100BT over-shoot test.
5	Dcd_test	R/W	0b	Retain	100BT DCD test.
4	PMD_LPBK_2	R/W	0b	0b	PMA loopback, test MLT-3 encoder and MLT-3 decoder.
3	PMD_LPBK_1	R/W	0b	0b	PMD loopback, test scrambler and descrambler.
2	PMA_LPBK_2	R/W	0b	0b	PMA loopback, test carrier detect and link monitor.
1	PMA_LPBK_1	R/W	0b	0b	PMA loopback, test FEF generator and FEF detector.
0	PCS_LPBK	R/W	0b	0b	PCS loopback, test pcs_tx and pcs_rx.

Table 38. 1000BASE-T Test Mode Register, Address Offset = 0x11

Bits	Field	Mode	HW Rst	SW Rst	Description
15:1	Reserved	R/W	0x2AA9	Retain	Reserved
0	Ext_lpbk_1000	R/W	0b	0b	1b = Enable 1000BASE-T external loopback with channel 0 <-> channel 1, channel 2 <-> channel 3.

Table 39. 10BASE-T Test Mode Register, Address Offset = 0x12

Bits	Field	Mode	HW Rst	SW Rst	Description
15:14	Reserved	R/W	0x120	Retain	Reserved
5	Test_mode[2]	R/W	0b	0b	bit 2 of 3-bit test_mode[2:0]. See bit 1:0 of this register.
4:3	Reserved	R/W	00b	0b	Reserved
2	Loopback mode select	R/W	0b	0b	1b = lpbk2—deep in Loopback mode. 0b = lpbk1—shallow in Loopback mode (connect to dig10.test_mode_i[0]).
1:0	Test_mode[1:0]	R/W	00b	0b	Combined with bit 5: 001b = Packet with all ones, 10 MHz sine wave. 010b = Pseudo random. 011b = Normal link pulse only. 100b = 5 MHz sin wave. Others: Normal mode.

Table 40. Power Saving Control, Address Offset = 0x29

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Top_ps_en	RO	1b	Retain	1b = Top level power saving enable. 0b = Top level power saving disable.
14:12	Dac_amp_1000	R/W	0x3	Retain	Control amplitude of transmit signal in 1000BT mode.
11:9	Dac_amp_100	R/W	0x3	Retain	Control amplitude of transmit signal in 100BT mode.



Bits	Field	Mode	HW Rst	SW Rst	Description
8:6	Dac_amp_10	R/W	0x3	Retain	Control amplitude of transmit signal in 100BT mode.
5:1	Reserved	R/W	0x0	0	Reserved
0	ecnc_ps_en	R/W	1b	Retain	1b = ecnc power saving enable. 0b = ecnc power saving disable.



8.6 Port Control Registers (Page 769)

Table 41. Port General Configuration PHY Address 01, Page 769, Register 17

Name	Default	Bits	Description	Type
Tx Gate Wait IFS	01110b	15:11	Determines the size (in nibbles) of non-deferring window from CRS de-assertion.	R/W
BP extension Wait	100b	10:8	Additional waiting byte times after TX Gate Wait IPG expires until the <i>Back Pressure In-band</i> bit is cleared.	R/W
Reserved	0b	7	Reserved	R/W
Active_PD_enable	0b	6	Active Power Down Enable (sD3 Enable) When set to 1b, the Intel® 5 Series Express Chipset needs to enter MAC power down mode.	R/W
Reserved	1b	5	Reserved. This bit is reset by power on reset only.	
Host_WU_Active	0b	4	Enables host wake up from the 82578. This bit is reset by power on reset only.	R/W
Wakeup clocks stop	1b	3	Wake-up clocks are stopped while wake up is disabled.	R/W
MACPD_enable	1b	2	Written as 1b when needs to globally enable the MAC power down feature while the 82578 supports WoL. When set to 1b, pages 800 and 801 are enabled for configuration and <i>Host_WU_Active</i> , <i>ME_WU_Active</i> are not blocked for writes.	R/W
Reserved	00b	1:0	Reserved	RO

Table 42. 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 43. 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 44. 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	1b	8	Enable generation of early preamble based on RX_DV in the receive path.	R/W
rx_en_crs_preamble	0b	7	Enable generation of early preamble based on CRS in the receive path.	R/W
reserved	0b	6	Reserved, write as read.	RWP
rx_flip_bad_sfd	1b	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 45. 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

8.7 Statistics Registers

Table 46. 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 82578 is in half-duplex mode.

Table 47. 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 82578 is in half-duplex mode.

Table 48. 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 82578 is in half-duplex mode.

Table 49. 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 82578 is in half-duplex mode.

Table 50. 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 82578 is in half-duplex mode. This register applies to clear as well as secure traffic.

Table 51. 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 82578 relative to the 82542. For the 82578, this counter does not increment for streaming transmits that are deferred due to TX IPG.

Table 52. 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 82578.

This register counts the number of successful packet transmission in which the CRS input from the 82578 was not asserted within one slot time of start of transmission from the MAC. Start of transmission is defined as the assertion of TX_EN to the 82578.



The 82578 should assert CRS during every transmission. Failure to do so might indicate that the link has failed, or the 82578 has an incorrect link configuration. This register only increments if transmits are enabled. This register is only valid when the 82578 is operating at half duplex.

8.8 PCIe Registers

Table 53. 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 54. PCIe Power Management Control PHY Address 01, Page 770, Register 17

Name	Default	Bits	Description	Type
Burst enable	1b	15	Burst in 10/100 Mb/s Enable 1b = Bursting at 10/100 Mb/s speed is enabled. 0b = Bursting disabled at 10/100 Mb/s.	RW
Reserved	00b	14:13	Reserved	R/W
Reserved	000b	12:10	Reserved.	RO
Reserved	10b	9:8	Reserved.	R/W
Reserved	1b	7	Reserved	R/W
Reserved	00b	6:5	Reserved	R/W
Reserved	0010b	4:1	Reserved	R/W
Reserved	0b	0B	Reserved	R/W



Table 55. In-Band Control PHY Address 01, Page 770, Register 18¹

Name	Default	Bits	Description	Type
Link status transmit timeout	0x5	15:8	Link status retransmission period in tens of microseconds.	R/W
pcie_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 timeouts are multiplied by 1000 while in SMBus mode.

Table 56. PCIe Diagnostic PHY Address 01, Page 770, Register 20¹

Name	Default	Bits	Description	Type
Reserved	0x55	15:8	Reserved, write as read.	R/W
In-band status acknowledge timeout	0x04	7:0	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 57. Timeouts PHY Address 01, Page 770, Register 21¹

Name	Default	Bits	Description	Type
Reserved	0000b	15:12	Reserved, write as read.	R/W
Reserved	010100b	11:6	Reserved	R/W
Reserved	010100b	5:0	Reserved	R/W

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

Table 58. 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 82578 stays in electrical idle state once entered (each bit represents 100 ns).	R/W

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



8.9 General Registers

Table 59. 82578 Capability PHY Address 01, Page 776, Register 19

Name	Default	Bits	Description	Type
Reserved	000000b	15:10	Reserved for future capabilities.	RO
Reserved	0b	9	Reserved	RO
802.1Q & 802.1p	0b	8	802.1Q & 802.1p Enables support for VLAN per 802.1Q & 802.1p.	RO
Receive Side Scaling	0b	7	Receive Side Scaling (RSS) Enables RSS.	RO
2 Tx and 2 Rx Queues	0b	6	Two Tx and 2 Rx Queues When set, enables dual transmit and dual receive queues. When cleared, a single receive and a single transmit queue are enabled.	RO
Energy Detect	0b	5	Energy Detect Enables energy detect capability.	RO
AC/DC Auto Link Speed Connect	0b	4	AC/DC Auto Link Speed Connect Enables different power management policy in AC and battery modes.	RO
Reserved	0b	3	Reserved	RO
Reserved	00b	2:1	Reserved	RO
Ability to initiate a team	0b	0	Ability to initiate a team; enables teaming capability.	RO

The 82578 Capability register is loaded with the set of capabilities that correspond to the selected the 82578 SKU. A change in SKU is reflected in a change in this register. A capability is enabled when its corresponding bit is set to 1b.

Table 60. 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. Allows 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 61. 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 82578 supports PEC on the SMBus.	RW
SMBus Frequency	0b	8	0b = 100 KHz. 1b = Reserved	RW
SMBus Address Valid	0b	7	0b = Address not valid. 1b = SMBus address valid. This bit is written by the MAC when the SMBus Address field is updated. The 82578 cannot send SMBus transactions to the MAC unless this bit is set.	RW
SMBus Address	0x00	6:0	This is the MAC SMBus address. The 82578 uses it for master functionality.	RW

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

Table 62. 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 63. 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.
RW	17:16	X	Address Select (ASEL) Selects how the address is to be used and is decoded as follows: 00b = Destination address (must be set to this in normal mode). 01b = Source address. 10b = Reserved. 11b = Reserved.
RO	30:18	0x00	Reserved, reads as 0b and ignored on writes.
RW	31	See Desc.	Address valid (AV) Cleared after master reset. If the NVM is present, the Address Valid field of Receive Address Register 0 is set to one after a software or PCI reset or NVM read. This bit is cleared by a master (software) reset.



Table 64. 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.

8.9.1 Interrupts

The 82578 maintains status bits (per interrupt cause) to reflect the source of the interrupt request. System software is expected to clear these status bits once the interrupt is being handled.



8.10 Wake Up Registers

8.10.1 Accessing Wake Up Registers Using MDIC

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

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

For the 82578 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 82578'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 82578'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 82578'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 82578, 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 82578'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 82578'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 82578's address from the MDI register
 - d. REGADD = 0x12 (data cycle for write)
 - e. DATA = YYYY (data to be written to the register)

8.10.2 Host Wake Up Control Status Register Description

Table 65. Receive Control – RCTL PHY Address 01, Page 800, Register 0

Attribute	Bit(s)	Initial Value	Description
RW	0	0b	Unicast Promiscuous Enable (UPE) 0b = Disabled. 1b = Enabled.
RW	1	0b	Multicast Promiscuous Enable (MPE) 0b = Disabled. 1b = Enabled.
RW	2	1b	Slave Access Enable 0b = Access disabled, the filters are active. 1b = Access enabled, the filters are not active.
RW	4:3	00b	Multicast Offset (MO) This determines which bits of the incoming multicast address are used in looking up the bit vector. 00b = [47:38]. 01b = [46:37]. 10b = [45:36]. 11b = [43:34].
RW	5	0b	Broadcast Accept Mode (BAM) 0b = Ignore broadcast (unless it matches through exact or imperfect filters) 1b = Accept broadcast packets.
RW	6	0b	Pass MAC Control Frames (PMCF). 0b = Do not (specially) pass MAC control frames. 1b = Pass any MAC control frame (type field value of 0x8808).
RW	7	0b	Receive Flow Control Enable (RFCE) Indicates that the 82578 responds to the reception of flow control packets. If auto-negotiation is enabled, this bit is set to the negotiated duplex value.
RW	8	0b	Reserved
RW	15:9	0x00	Reserved



Note: All wake up registers (page 800-801 except CTRL and IPAV) are not cleared with PHY reset is asserted. It is only cleared when internal power on reset is de-asserted or when cleared by the software device driver.

Note: Access to page 800/801 should be done only in 10 Mb/s and 100 Mb/s.

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 66. 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 82578 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 LNKC bit in the WUFC register. Instead, it is determined by the APM settings in the WUC register.
RO	15:6	0x00	Reserved



Table 67. 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	Reserved
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 3 enable.
RW	10	0b	FLX5 Flexible filter 3 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 68. Wake Up Status – WUS PHY Address 01, Page 800, Register 3

Attribute	Bit(s)	Initial Value	Description
RWC	0	0b	LNKC Link status changed
RWC	1	0b	MAG Magic packet received
RWC	2	0b	EX Directed exact packet received. The packet's address matched one of the 7 pre-programmed exact values in the Receive Address registers.
RWC	3	0b	MC Directed multicast packet received. The packet was a multicast packet that was hashed to a value that corresponded to a 1-bit in the multicast table array.
RWC	4	0b	BC Broadcast packet received.
RWC	5	0b	IPv4 request packet received.
RWC	6	0b	IPV4 Directed IPv4 packet received.
RWC	7	0b	IPV6 Directed IPv6 packet received.
RO	8	0b	Reserved, read as 0b.
RWC	9	0b	FLX4 Flexible filter 4 match.
RWC	10	0b	FLX5 Flexible filter 5 match.
RO	11	0b	Reserved.
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 69. Receive Address Low – RAL PHY Address 01, Page 800, Registers 16-17 + 4*n¹ (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 70. Receive Address High – RAH PHY Address 01, Page 800, Registers 18-19 + 4*n¹ (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.
RW	17:16	X	Address Select (ASEL) Selects how the address is to be used and is decoded as follows: 00b = Destination address (must be set to this in normal mode). 01b = Source address. 10b = Reserved. 11b = Reserved.
RO	30:18	0x00	Reserved, reads as 0b and ignored on writes.
RW	31	See Desc.	Address valid (AV) Cleared after master reset. If the NVM is present, the <i>Address Valid</i> field of Receive Address Register 0 is set to one after a software or PCI reset or NVM read. This bit is cleared by a master (software) reset.

1. While "n" is the exact unicast/Multicast address entry and it is equals to 0,1,..6

AV determines whether this address is compared against the incoming packet. AV is cleared by a master (software) reset.

ASEL enables the 82578 to perform special filtering on receive packets.

Note: RAR0 should always be used to store the individual Ethernet MAC address of the Network Interface Card (NIC).

After reset, if the NVM is present, the first register (Receive Address Register 0) is loaded from the *IA* field in the NVM, its *Address Select* field is 00b, and its *Address Valid* field is 1b. If no NVM is present the *Address Valid* field is 0b. The *Address Valid* field for all of the other registers are 0b.



Table 71. 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 72. 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	17:16	0x00	Address Select (ASEL) Selects how the address is to be used. 00b means that it is used to decode the destination MAC address.
RO	30:18	0x00	Reserved, reads as 0b and is ignored on writes.
RW	31	0b	Address valid (AV) When this bit is set, the relevant RAL and RAH are valid (compared against the incoming packet).

Table 73. 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	17:16	00b	Address Select (ASEL) Selects how the address is to be used. 00b means that it is used to decode the destination MAC address.
RO	29:18	0x00	Reserved, reads as 0x00 and is ignored on writes.
RW	30	0b	All Nodes Multicast Address valid (MAV) The all nodes multicast address (33:33:00:00:00:01) is valid when this bit is set. Note that 0x33 is the first byte on the wire.
RW	31	0b	Address valid (AV) When this bit is set, the relevant address 3 is valid (compared against the incoming packet).



Table 74. 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 75. 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 IPv4 request packets and directed IPv4 packet wake ups. It is a 3-entry table with the following format:

Table 76. 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 can be used by the host.

Table 77. 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 9 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 9 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

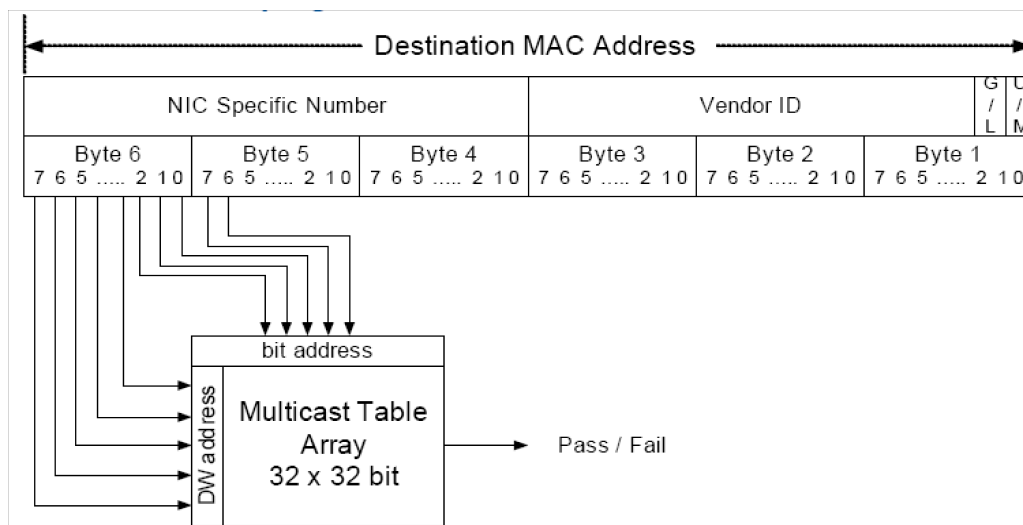


Figure 9. Multicast Table Array Algorithm

Table 78. 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 82578 since each address contains 16 bits, only the least significant bytes are stored in those addresses.

Table 79. 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 82578 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 80. 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 81. 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 3 byte n (n=0, 1... 127).
RW	5	X	Mask 5 Mask for filter 3 byte n (n=0, 1... 127).
RO	15:6	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 82. 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 four 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 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 83. 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.



9.0 Non-Volatile Memory (NVM)

9.1 Introduction

This section is intended for designs using a 10/100/1000 Mb/s MAC that is integrated into an Intel® 5 Series Express Chipset in conjunction with an the 82578 Physical Layer Transceiver (PHY).

There are several LAN clients that might access the NVM such as hardware, LAN driver, and BIOS. Refer to the *Intel® 5 Series Express Chipset Family External Design Specification (Intel® 5 Series Express Chipset EDS)* and the *Intel® 5 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).

9.2 NVM Programming Procedure Overview

The LAN NVM shares space on an SPI Flash device (or devices) along with the BIOS, Manageability Firmware, and a Flash Descriptor Region. It is programmed through the Intel® 5 Series Express Chipset. This combined image is shown in [Figure 10](#). 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® 5 Series Express Chipset Family External Design Specification (Intel® 5 Series Express Chipset EDS)*. This document describes the LAN image contained in the Gigabit Ethernet (GbE) region.




Figure 10. 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® 5 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® 5 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® 5 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® 5 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.

9.3 LAN NVM Format and Contents

Table 11 lists the NVM maps for the LAN region. Each word listed is described in detail in the following sections.

Table 11. 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
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	HW-Shared	LED 0 - 2	
0x19:0x2F	0x32:0x5E	HW-Shared	Reserved	0x0000
LAN Word Offset	NVM Byte Offset	Used By	15	0 Image Value
0x30:0x3E	0x60:0x7C	PXE	PXE Software Region	
0x3F	0x7E	SW	Software Checksum (Bytes 0x00 through 0x7D)	
0x40:0x4A	0x80:0x94	HW	G3 -> S5 LCD 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.

9.3.1 Hardware Accessed Words

This section describes the NVM words that are loaded by the MAC hardware.

9.3.1.1 Ethernet Address (Words 0x00-0x02)

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

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

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



9.3.1.2 PCI Init Control Word (Word 0x0A)

This word contains initialization values that:

- 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
15:13	Reserved	000b	This field is reserved and must be set to 000b.
12	Reserved	1b	Reserved, must be set to 1b.
11:8	Reserved	0000b	These bits are reserved and must be set to 0000b.
7	AUX PWR	1b	Auxiliary Power Indication If set and if PM Ena is set, D3cold wake-up is advertised in the Intel® 5 Series Express Chipset 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 Intel® 5 Series Express Chipset 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 0Bh and 0Ch).
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 0Dh).

9.3.1.3 Subsystem ID (Word 0x0B)

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

9.3.1.4 Subsystem Vendor ID (Word 0x0C)

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

9.3.1.5 Device ID (Word 0x0D)

If the Load Device ID in word 0x0A is set, this word is read in to initialize the Device ID of the 82578 PHY. Default value is 0x10EF.

9.3.1.6 Reserved Words 0x0E and 0x0F



9.3.1.7 LAN Power Consumption (Word 0x10)

This word is meaningful only if the power management is enabled.

Bits	Name	Default	Description
15:8	LAN D0 Power	00001000b	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	0x1	The value in this field is reflected in the PCI Power Management Data register for D3 power consumption and dissipation (<i>Data_Select</i> = 3 or 7). Power is defined in 100 mW units. The power also includes the external logic required for the LAN function. The most significant bits in the Data register that reflects the power values are padded with zeros.

9.3.1.8 Reserved (Word 0x11)

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

9.3.1.9 Reserved (Word 0x12)

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

9.3.1.10 Shared Init Control Word (Word 0x13)

This word controls general initialization values.

Bits	Name	Default	Description
15:14	Sign	00b	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 MAC does not read the rest of the NVM data and default values are used for the device configuration.
13	Reserved	0b	Reserved, set to 1b.
12:10	Reserved	010b	Reserved, set to 001b.
9	PHY PD Ena	1b	Enable PHY Power Down When set, enables PHY power down at DMoff/D3 or Dr and no WoL. This bit is loaded to the <i>PHY Power Down Enable</i> bit in the Extended Device Control (CTRL_EXT) register. 1b = Enable PHY power down. 0b = PHY always powered up.
8	Reserved	1b	Reserved, should be set to 1b.



Bits	Name	Default	Description
7:6	PHYT	10b	PHY Device Type Indicates that the PHY is connected to the MAC and resulted mode of operation of the MAC/PHY link buses. 00b = 82578. 01b = Reserved. 10b = Reserved. 11b = Reserved.
5	Reserved	01	Reserved, should be set to 1b.
4	FRCSPD	0b	Default setting for the <i>Force Speed</i> bit in the Device Control register (CTRL[11]).
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 MAC units. This bit is loaded to the <i>DynCK</i> bit in the CTRL_EXT register.



9.3.1.11 Extended Configuration Word 1 (Word 0x14)

Bits	Name	Default	Description
15:14	Reserved	00b	Reserved, set to 00b.
13	LCD Write Enable	0b	When set, enables loading of the extended LAN connected device configuration area in the 82578. This configuration area also includes the PHY tuning (tuning for IEEE) in the NVM. Since this bit is set to 0b by default, PHY tuning does not take effect until the LAN driver and/or firmware loads. When disabled, the extended LAN connected device configuration area is ignored. Loaded to the EXTCNF_CTRL register.
12	OEM Write Enable	0b	When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. OEM bits include any LED configuration. Since this bit is set to 0b by default, the auto-load of OEM bits do not take effect until the LAN driver and/or firmware loads. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[3]). 1b = OEM bits written to the 82578. 0b = No OEM bits configuration.
11:0	Extended Configuration Pointer	0x208	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 82578. Loaded to the Extended Configuration Control register (EXTCNF_CTRL[27:16]).

9.3.1.12 Extended Configuration Word 2 (Word 0x15)

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

9.3.1.13 Extended Configuration Word 3 (Word 0x16)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved, set to 0x00.
7:0	Reserved	0x00	Reserved, set to 0x00.



9.3.1.14 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 decrease in link speed in non-D0a states when power policy and power management states dictate 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 decrease in link speed in all power states.
8	Reserved	0b	Reserved, set to 0b.
7:0	Reserved	0x0	Reserved.



9.3.1.15 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® 5 Series Family PDG* and the *82578 Reference Schematics* for LED connection details. Also, [Table 12](#) lists mode encodings for LED outputs.

Note: Due to the architecture of the 82578 the customized LEDs settings are written to the 82578 by the LAN driver. As a result, the default LEDs are written during the boot process and when resuming from power states S3, S4, and S5 to normal operation until the LAN driver writes any custom settings. This same behavior is also observed while in S3 and toggling from ac power (wall outlet) to dc power (battery). Once the LAN driver loads after a system boot or when resuming from sleep states, the LEDs function as defined in Word 0x18 of the GbE region of the NVM.

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 12. Mode Encodings for LED Outputs

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

9.3.1.16 Reserved (Word 0x19)

Bits	Name	Default	Description
15:14	Reserved	00b	Reserved, set to 00b.
13	Reserved	0b	Reserved, set to 0b.
12:10	Reserved	010b	Reserved, set to 010b.
9:8	Reserved	11b	Reserved, set to 11b.
7	Reserved	0b	Reserved, set to 0b.
6	Invalid image CSUM	0b	When cleared this bit indicates to the NVM programming tools (EEUPDATE and LANConf) that the image checksum needs to be corrected. When set, the checksum is assumed to be correct.
5:0	Reserved	0x0	Reserved, set to 0x0.

9.3.1.17 Reserved (Word 0x1A)

Bits	Name	Default	Description
15:12	Reserved	0x000	Reserved, set to 0x000.
11	Reserved	1b	Reserved, set to 1b.
10:7	Reserved	0000b	Reserved, set to 0000b.
6	Reserved	1b	Reserved, set to 1b.
5:2	Reserved	0000b	Reserved, set to 0000b.
1	Reserved	1b	Reserved, set to 1b.
0	APM Enable	1b	APM Enable Initial value of Advanced Power Management Wake Up Enable in the Wake Up Control (WUC.APME) register. 1b = Advanced power management enabled. 0b = Advanced power management disabled.



9.3.1.18 Reserved (Word 0x1B)

Bits	Name	Default	Description
15:5	Reserved	0x0	Reserved, set to 0x0.
4	K1_PLL_stop_en	1b	Enable PLL stop in K1.
3	K0s_100_En	0b	Enables K0s mode when PHY link speed is 10/100 Mb/s.
2	K0s_GbE_En	0b	Enables K0s mode when PHY link speed is 1000 Mb/s.
1	Reserved	1b	Reserved, set to 0b.
0	K1_En	1b	When set to 1b enables K1 low power mode.

9.3.1.19 Reserved (Word 0x1C)

Bits	Name	Default	Description
15:0	Reserved	0x10EF	Reserved

9.3.1.20 Reserved (Words 0x1D, 0x1E, 0x1F, and 0x20)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

9.3.1.21 Reserved (Word 0x21)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

9.3.1.22 Reserved (Word 0x22)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved

9.3.1.23 Reserved (Word 0x23)

Bits	Name	Default	Description
15:0	Reserved	0xBAAD	Reserved



9.3.1.24 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.

9.3.1.25 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.

9.3.1.26 Reserved (Word 0x26)

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

9.3.1.27 Reserved (Word 0x27)

Bits	Name	Default	Description
15:0	Reserved	0x00	Reserved

9.3.2 Software Accessed Words

9.3.2.1 PXE Words (Words 0x30 Through 0x3E)

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



9.3.2.1.1 Boot Agent Main Setup Options (Word 0x30)

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

Table 13. Boot Agent Main Setup Options

Bit	Name	Description
15:14	Reserved	Reserved, set to 00b.
13	Reserved	Reserved, must be set to 0b.
12	FDP	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	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	Reserved Set this bit to 0b.
8	DSM	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	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	Reserved
4:3	DBS	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	Reserved
1:0	PS	Protocol Select. These bits select the boot protocol. 00b = PXE (default value). 01b = Reserved. Other values are undefined.



9.3.2.1.2 **Boot Agent Configuration Customization Options (Word 0x31)**

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

Table 14. Boot Agent Configuration Customization Options (Word 0x31)

Bit	Name	Description
15:14	SIG	Signature Set these bits to 11b to indicate valid data.
13:12	Reserved	Reserved, must be set to 00b.
11		Continuous Retry Disabled (0b default).
10:8	MODE	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	Reserved, must be set to 00b.
5	DFU	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.
4	DLWS	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	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.



Bit	Name	Description
2	DPS	<p>Disable Protocol Select</p> <p>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.</p>
1	DTM	<p>Disable Title Message</p> <p>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.</p> <p>The default for this bit is 0b; allow the title message that displays the version of the boot agent and the Control-S message.</p>
0	DSM	<p>Disable Setup Menu</p> <p>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.</p> <p>The default for this bit is 0b; allow invoking the setup menu by pressing Control-S.</p>



9.3.2.1.3 Boot Agent Configuration Customization Options (Word 0x32)

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

Table 15. Boot Agent Configuration Customization Options (Word 0x32)

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

9.3.2.1.4 IBA Capabilities (Word 0x33)

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



Table 16. IBA Capabilities

Bit	Name	Description
15:14	SIG	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	Reserved, must be set to 0x00.
4		iSCSI boot capability not present (0b default).
3	EFI	EFI EBC capability is present in Flash. 0b = The EFI code is not present (default). 1b = The EFI code is present.
2	Reserved	Reserved, set to 1b.
1	UNDI	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	PXE base code is present in Flash. 0b = The PXE base code is present (default). 1b = The PXE base code is not present.

9.3.2.2 Checksum Word Calculation (Word 0x3F)

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

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

9.3.3 Basic Configuration Software Words

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

9.3.3.1 Reserved (Word 0x3)

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



9.3.3.2 Reserved (Words 0x04, 0x06, and 0x07)

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved

9.3.3.3 Image Version Information (Word 0x05)

Note: This is a reserved word and cannot be changed.

Care should be taken to use the correct GbE NVM firmware revision for the stepping combination of the Intel® 5 Series Express Chipset and the 82578. The following table lists the NVM revision that is optimized for use with the silicon stepping combination.

Note: Using a newer revision NVM with an older silicon stepping or older revision NVM with a newer silicon stepping could cause system instability and unpredictable behavior.

Intel® 5 Series Express Chipset Stepping	82578 Stepping/PHY-Ver	NVM Version
B1	C0	0.62

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

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

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

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

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




9.4 Intel® 5 Series Express Chipset/82578 NVM Contents

This section lists the NVM contents for the Intel® 5 Series Express Chipset and the 82578.

Table 17. LAN NVM Contents

Word	Description
0x00:0x02	Ethernet Individual Address
0x03:0x04	Reserved
0x05	Image Version Information
0x06:0x07	Reserved
0x08:0x09	PBA Bytes
0x0A	PCI Init Control Word
0x0B	Subsystem ID
0x0C	Subsystem Vendor ID
0x0D	Device ID
0x0E	Reserved
0x0F	Reserved
0x10	LAN Power Consumption
0x11	Reserved
0x12	Reserved
0x13	Shared Init Control Word
0x14:0x16	Extended Configuration Words
0x17	OEM Configuration Defaults
0x18	LED 0 - 2
0x19:0x2F	Reserved
0x30:0x3E	PXE Region
0x3F	Software Checksum



10.0 Intel® 5 Series Express Chipset MAC Programming Interface

10.1 Register Byte Ordering

This section defines the structure of registers that contain fields carried over the network. For example, L2, L3, L4 fields.

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

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

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

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

The general rule for register ordering is to use host ordering (also called little endian). Using the previous example, a 6-byte fields (such as a MAC address) is stored in a CSR in the following manner:

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

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

	Byte 3	Byte 2	Byte 1	Byte 0
Dword aligned or Word aligned	0x00	0x01
DW address (N + 4)	0x00	0x01



The following exceptions use network ordering:

- All ETherType fields

The normal notation as it appears in text books, etc. is to use network ordering. For example, the following MAC address: 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	Byte 0
Dword address (N)	00	C9	A0	00
Dword address (N + 4)	00	00

10.2 Register Conventions

All registers in the MAC 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 1b by software. Reads from registers containing reserved bits might 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 document should be considered to be reserved, and should not be written to. Writing to reserved or undefined register addresses might cause indeterminate behavior. Reads from reserved or undefined configuration register addresses might 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 initial value. Reading from reserved fields might 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 are listed as such via the text undefined, unknown, or X. Some of these configuration values might need to be set via NVM configuration or via software in order for proper operation to occur. Note that this need is dependent on the function of the bit. Other registers might cite a hardware default that is overridden by a higher-precedence operation. Operations that might supersede hardware defaults might also 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) does not take effect (the write is ignored). Partial reads return all 32 bits of data regardless of the byte enables.

Note: Partial reads to read-on-clear registers (such as 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 (like independent access to each 32-bit field).



10.2.1 PCI Configuration and Status Registers - CSR Space

All configuration registers are listed in [Table 18](#). 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 firmware; read only by software.
- RWC - Read write clear registers. Writing 0b has no affect. Writing 1b 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 write clear registers. Writing 0b has no affect. Writing 1b clears the appropriate fields. Note that a read might also clear the register depending on enablement (see appropriate 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. Generally this would be all zero's (see detailed description of appropriate registers).

Table 18. Register Summary

Offset	Abbreviation	Name	RW	Paragraph
General Register Descriptions				
0x00000	CTRL	Device Control Register	RW	10.2.1.1.1
0x00008	STATUS	Device Status Register	RO	10.2.1.1.2
0x0000C	STRAP	Strapping Option Register	RO	10.2.1.1.3
0x00018	CTRL_EXT	Extended Device Control Register	RW	10.2.1.1.4
0x00020	MDIC	MDI Control Register	RW	10.2.1.1.5
0x00028	FEXTNVM	Future Extended NVM Register	RW	10.2.1.1.6
0x0002C	FEXT	Future Extended Register	RW	10.2.1.1.7
0x00038	BUSNUM	Device and Bus Number	RO	10.2.1.1.8
0x00170	FCTTV	Flow Control Transmit Timer Value	RW	10.2.1.1.9
0x05F40	FCRTV	Flow Control Refresh Threshold Value	RW	10.2.1.1.10
0x00F00	EXTCNF_CTRL	Extended Configuration Control	RW	10.2.1.1.11
0x00F08	EXTCNF_SIZE	Extended Configuration Size	RW	10.2.1.1.12
0x00F10	PHY_CTRL	PHY Control Register	RW	10.2.1.1.13
0x00F18	PCIEANACFG	PCIE Analog Configuration	RW	10.2.1.1.14
0x01000	PBA	Packet Buffer Allocation	RW	10.2.1.1.15
0x01008	PBS	Packet Buffer Size	RW	10.2.1.1.16
0x0100C	PBECCSTS	Packet Buffer ECC Status	RW	10.2.1.1.17
0x01004	PBEEI	Packet Buffer ECC Error Inject	RW	10.2.1.1.18



Offset	Abbreviation	Name	RW	Paragraph
Interrupt Register Descriptions				
0x000C0	ICR	Interrupt Cause Read Register	RC/ WC	10.2.1.2.1
0x000C4	ITR	Interrupt Throttling Register	RW	10.2.1.2.2
0x000C8	ICS	Interrupt Cause Set Register	WO	10.2.1.2.3
0x000D0	IMS	Interrupt Mask Set/Read Register	RW	10.2.1.2.4
0x000D8	IMC	Interrupt Mask Clear Register	WO	10.2.1.2.5
0x000E0	Mask - IAM	Interrupt Acknowledge Auto	RW	10.2.1.2.6
Receive Register Descriptions				
0x00100	RCTL	Receive Control Register	RW	10.2.1.3.1
0x00104	RCTL1	Receive Control Register 1	RW	10.2.1.3.2
0x02008	ERT	Early Receive Threshold	RW	10.2.1.3.3
0x02170	PSRCTL	Packet Split Receive Control Register	RW	10.2.1.3.4
0x02160	FCRTL	Flow Control Receive Threshold Low	RW	10.2.1.3.5
0x02168	FCRTH	Flow Control Receive Threshold High	RW	10.2.1.3.6
0x02800	RDBAL	Receive Descriptor Base Address Low Queue	RW	10.2.1.3.7
0x02804	RDBAH	Receive Descriptor Base Address High Queue	RW	10.2.1.3.8
0x02808	RDLEN	Receive Descriptor Length Queue	RW	10.2.1.3.9
0x02810	RDH	Receive Descriptor Head Queue	RW	10.2.1.3.10
0x02818	RDT	Receive Descriptor Tail Queue	RW	10.2.1.3.11
0x02820	RDTR	Interrupt Delay Timer (Packet Timer)	RW	10.2.1.3.12
0x02828	RXDCTL	Receive Descriptor Control	RW	10.2.1.3.13
0x0282C	RADV	Receive Interrupt Absolute Delay Timer	RW	10.2.1.3.14
0x02C00	RSRPD	Receive Small Packet Detect Interrupt	RW	10.2.1.3.15
0x02C08	RAID	Receive ACK Interrupt Delay Register	RW	10.2.1.3.16
0x05000	RXCSUM	Receive Checksum Control	RW	10.2.1.3.17
0x05008	RFCTL	Receive Filter Control Register	RW	10.2.1.3.18
0x05200-0x0527C	MTA[31:0]	Multicast Table Array	RW	10.2.1.3.19
0x05400 + 8*n (n=0..6)	RAL	Receive Address Low	RW	10.2.1.3.20
0x05404 + 8*n (n=0..6)	RAH	Receive Address High	RW	10.2.1.3.21
0x05438 + 8*n (n=0..3)	SRAL	Shared Receive Address Low	RW	10.2.1.3.22
0x0543C + 8*n (n=0..2)	SRAH	Shared Receive Address High 0..2	RW	10.2.1.3.23
0x05454	SHRAH[3]	Shared Receive Address High 3	RW	10.2.1.3.24
0x05818	MRQC	Multiple Receive Queues Command Register	RW	10.2.1.3.25
0x05C00 + 4*n (n=0..31)	RETA	Redirection Table	RW	10.2.1.3.26
0x05C80 + 4*n (n=0..9)	RSSRK	Random Key Register	RW	10.2.1.3.27



Offset	Abbreviation	Name	RW	Paragraph
Transmit Register Descriptions				
0x00400	TCTL	Transmit Control Register	RW	10.2.1.4.1
0x00410	TIPG	Transmit IPG Register	RW	10.2.1.4.2
0x00458	AIT	Adaptive IFS Throttle	RW	10.2.1.4.3
0x03800	TDBAL	Transmit Descriptor Base Address Low	RW	10.2.1.4.4
0x03804	TDBAH	Transmit Descriptor Base Address High	RW	10.2.1.4.5
0x03808	TDLEN	Transmit Descriptor Length	RW	10.2.1.4.6
0x03810	TDH	Transmit Descriptor Head	RW	10.2.1.4.7
0x03818	TDT	Transmit Descriptor Tail	RW	10.2.1.4.8
0x03840	TARC	Transmit Arbitration Count	RW	10.2.1.4.9
0x03820	TIDV	Transmit Interrupt Delay Value	RW	10.2.1.4.9
0x03828	TXDCTL	Transmit Descriptor Control	RW	10.2.1.4.10
0x0382C	TADV	Transmit Absolute Interrupt Delay Value	RW	10.2.1.4.11
Management Register Descriptions				
0x05800	WUC	Wake Up Control Register	RW	10.2.1.5.1
0x05808	WUFC	Wake Up Filter Control Register	RW	10.2.1.5.2
0x05810	WUS	Wake Up Status Register	RW	10.2.1.5.3
0x5838	IPAV	IP Address Valid	RW	10.2.1.5.4
0x05840 + 8*n (n=1...3)	IP4AT	IPv4 Address Table	RW	10.2.1.5.5
0x05880 + 4*n (n=0...3)	IP6AT	IPv6 Address Table	RW	10.2.1.5.6
0x05F00 + 8*n (n=0...35)	FFLT	Flexible Filter Length Table	RW	10.2.1.5.7
0x09000 + 8*n (n=0...127)	FFMT	Flexible Filter Mask Table	RW	10.2.1.5.8
0x09800 + 8*n (n=0...127)	FFVT	Flexible Filter Value Table	RW	10.2.1.5.10
0x09804 + 8*n (n=0...127)	FFVT2	Flexible Filter Value Table	RW	10.2.1.5.10

Note: Certain registers maintain an alias address designed for backward compatibility with software written for the previous devices. Registers that have an alias address can be accessed by software at either the new offset or the alias offset. It is recommended that software that is written solely for the Intel® 5 Series Express Chipset and the 82578 use the new address offset.



10.2.1.1 General Register Descriptions

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

Bit	Type	Reset	Description
0	RW/SN	1b	Full Duplex (FD). 0b = Half duplex. 1b = Full duplex. Controls the MAC duplex setting when explicitly set by software. Loaded from the NVM word 0x13.
1	RO	0b	Reserved. Write as 0b for future compatibility
2	RW	0b	Master Disable. When set, the MAC 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	000b	Reserved. Write as 0b for future compatibility.
6	RO	1b	Reserved.
7	RO	0b	Reserved. Must be set to 0b.
9:8	RW	10b	Speed selection (SPEED). These bits might 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. 0)b = 10 Mb/s. 0)b = 100 Mb/s. 10b = 1000 Mb/s. 11b = Not used.
10	RO	0b	Reserved. Write as 0b for future compatibility.
11	RW/SN	0b	Force Speed (FRCSPEED). This bit is set when software needs to manually configure the MAC speed settings according to the <i>Speed</i> bits (bits 9:8). When using the 82578, note that it 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 0x13 in the NVM. Note that this bit is superseded by the CTRL_EXT.SPD_BYPS bit, which has a similar function.
12	RW	0b	Force Duplex (FRCDPLX). When set to 1b, software might override the duplex indication from the 82578 that is indicated in the FDX to the MAC. Otherwise, the duplex setting is sampled from the 82578 FDX indication into the MAC on the asserting edge of the PHY link signal. When asserted, the CTRL.FD bit sets duplex.
13	RO	0b	Reserved.
14	RW/SN	0b	Reserved.
15	RO	0b	Reserved. Reads as 0.
18:16	RW	0b0	Reserved.
19	RW	0b	Memory Error Handling Enable (MEHE). When set to 1b, the Intel® 5 Series Express Chipset reaction to correctable and uncorrectable memory errors detection are activated.
20		1b	Reserved.
24:21	RO	0x0	Reserved.
25	RW	0b	Reserved.
26	RW/V	0b	Host Software Reset (SWRST). This bit performs a reset to the PCI data path and the relevant shared logic. Writing 1b initiates the reset. This bit is self-clearing.
27	RW	0b	Receive Flow Control Enable (RFCE). Indicates that the MAC responds to receiving flow control packets. If auto-negotiation is enabled, this bit is set to the negotiated duplex value.



Bit	Type	Reset	Description
28	RW	0b	Transmit Flow Control Enable (TFCE). Indicates that the MAC transmits flow control packets (XON and XOFF frames) based on receiver fullness. If auto-negotiation is enabled, this bit is set to the negotiated duplex value.
29	RO	0b	Reserved.
30	RW	0b	VLAN Mode Enable (VME). When set to 1b, all packets transmitted from MAC that have VLE set is 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	0b	LAN Connected Device Reset (LCD_RST). Controls an inband message to the 82578. 0b = Normal operation 1b = Reset to PHY is asserted. The LCD_RST functionality is gated by the FWSM.RSPCI PHY bit. If the FWSM.RSPCI PHY bit is not set to 1b, 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 01b.

This register, as well as the Extended Device Control register (CTRL_EXT), controls the major operational modes for the MAC. While software writes to this register to control MAC settings, several bits (such as FD and SPEED) might be overridden depending on other bit settings and the resultant link configuration is determined by the 82578's auto-negotiation resolution.

The FD (duplex) and SPEED configurations of the MAC are normally determined from the link configuration process. Software might 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 82578 settings.

Manual link configuration is controlled through the 82578's MII management interface.

Host Software Reset (bit 26), might be used to globally reset the entire host data path and shared logic. 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 are 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 the global device reset has fully completed and that the MAC responds to subsequent accesses, programmers must wait approximately 1 ms 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 also reflected at address 0x00004 for legacy reasons. Neither the software driver nor firmware should use it since it might be unsupported in next generations.



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

Bits	Attribute	Reset	Description
0	RO/V	X	Full Duplex (FD). 0b = Half duplex. 1b = Full duplex. Reflects duplex setting of the MAC and/or link.
1	RO/V	X	Link up (LU). 0b = No link established. 1b = Link established. For this to be valid, the <i>Set Link Up</i> bit of the Device Control register (CTRL.SU) must be set.
3:2	RO/V	00b	PHY Type Indication (PHYTYPE). Indicates that the 82578 attached to the MAC and resulted mode of operation of the MAC/82578 Link buses. 00 = 82578. 01 = Reserved. 10 = Reserved. 11 = Reserved. This field is loaded from the Shared Init control word in the NVM.
4	RO/V	X	Transmission Paused (TXOFF). Indication of pause state of the transmit function when symmetrical flow control is enabled.
5	RO/V	1b	PHY Power Up not (PHYPWR). RO bit that indicates the power state of the 82578. 0b = The 82578 is powered on in the active state. 1b = The 82578 is in the power down state. The <i>PHYPWR</i> bit is valid only after PHY reset is asserted. Note: The PHY power up indication reflects the status of the LANPHYPC signaling to the 82578.
7:6	RO/V	X	Link speed setting (SPEED). This bit reflects the speed setting of the MAC and/or link. 00b = 10 Mb/s. 01b = 100 Mb/s. 10b = 1000 Mb/s. 11b = 1000 Mb/s.
8	RO/V	X	Master Read Completions Blocked. This bit is set when the MAC receives a completion with an error (EP = one or status!= successful). It is cleared on PCI reset.
9	RW/V/C	0b	LAN Init Done. This bit is asserted following completion of the LAN initialization from the Flash. Software is expected to clear this field to make it usable for the next initialization done event.
10	RW/V/C	1b	PHY Reset Asserted (PHYRA). This bit is R/W. Hardware sets this bit following the assertion of a 82578 reset (either hardware or in-band). The bit is cleared on writing 0b to it.
18:11	RO	0x0	Reserved.
19	RO/V	1b	Master Enable Status. Cleared by the MAC 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 are issued by this function as long as the <i>Master Disable</i> bit is set.
29:20	RO	0x0	Reserved. Reads as 0.
30	RO	0b	Reserved.
31	RO/SN	1b	Clock Control ¼ (CLK_CNT_1_4). This bit is loaded from the NVM word 0x13 and indicates the MAC 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) might be several 100's of ms. It reflects whether the PHY's link indication is present.

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 1b 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 might 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:

- 00b = 10 Mb/s
- 01b = 100 Mb/s
- 10b = 1000 Mb/s
- 11b = 1000 Mb/s



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

This register reflects the values of the soft strapping options fetched from the NVM descriptor in the Intel® 5 Series Express Chipset space. These signals are sampled by the MAC following LAN_RST# or global reset (PCI reset assertion).

Bit(s)	Type	Reset	Description
0	RO	1b	Reserved.
5:1	RO	0x0	LAN NVM Size (NVMS). LAN NVM space size is indicated in multiples of 4 KB. LAN NVM size might vary from 4 KB to 128 KB (a zero value means 4 KB).
10:6	RO	0x0	Reserved.
15:11	RO	0x0	Reserved.
16	RO	0b	MAC SMBus address enable (LCSMBADDEN).
23:17	RO	0x0	MAC SMBus address (LCSMBADD).
24	RO	0b	PHY SMBus address enable (LCDSMBADDEN).
31:25	RO	0x0	PHY SMBus address (LCDSMBADD).

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

Bits	Type	Reset	Description
11:0	RO	0x0	Reserved.
12	RW/V	1b	Reserved.
14:13		00b	Reserved.
15	RW	0b	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 hardware clock switching circuitry.
18:16		000b	Reserved.
19	RW/SN	0b	Dynamic Clock Gating (DynCK). When set, this bit enables dynamic clock gating of the DMA and MAC units. Refer to the description of the DynWakeCK in this register. This bit is loaded from NVM word 0x13.
20	RW/SN	1b	PHY Power Down Enable (PHYPDEN). When set, this bit enables the 82578 to enter a low-power state when the MAC is at the DMoff / D3 or Dr with no WoL. This bit is loaded from word 0x13 in the NVM.
24:21		0000b	Reserved.
25	RW	0b	DMA Clock Control (DMACKCTL). Controls the DMA clock source in non-GbE mode (10/100 and no Link). In GbE mode, the DMA clock source is always GLCI PLL divided by two. In normal operation, this bit should be in the default state in which the DMA clock source in non-GbE is mosc_clk. In test mode the DMACKCTL and PLLGateDis should be set to 1b and CLK_CNT_1_4 in the NVM should not be set. In this mode, the DMA clock source is GLCI PLL divided by two.
26	RW	0b	Disable Static GLCI PLL Gating (PLLGateDis). By default the PLL is functional only when the GLCI link is required, and inactive when it is not required (at non-GbE mode if LCI is available). When set to 1b the GLCI PLL is always active.
27	RW	0b	Interrupt Acknowledge Auto-Mask Enable (IAME). When this bit is set, a read or write to the ICR register has the side effect of writing the value in the IAM register to the IMC register. When this bit is 0b, this feature is disabled.



Bits	Type	Reset	Description
28	RW	0b	Driver loaded (DRV_LOAD). This bit should be set by the driver after it was loaded and cleared when the driver unloads or after a soft reset. The Manageability Controller (MC) loads this bit to indicate that the driver has loaded.
29	RW	0b	INT_TIMERS_CLEAR_ENA. When set, this bit enables the clear of the interrupt timers following an IMS clear. In this state, successive interrupts occur only after the timers expire again. When cleared, successive interrupts following IMS clear might happen immediately.
30		0b	Reserved.
31	RO	0b	Reserved. Reads as 0.

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

Note: If software uses the EE_RST function and needs 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 after asserting the EE_RST bit.

Note: The EEPROM reset function might 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.



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

Bits	Type	Reset	Description
15:0	RW/V	X	Data (DATA). In a Write command, software places the data bits and the MAC shifts them out to the 82578. In a Read command, the MAC reads these bits serially from the 82578 and software can read them from this location.
20:16	RW/V	0x0	PHY Register address (REGADD). For example, register 0, 1, 2, ... 31.
25:21	RW/V	0x0	PHY Address (PHYADD).
27:26	RW/V	00b	Op-code (OP). 01b = MDI write. 10b = MDI read. Other values are reserved.
28	RW/V	1b	Ready bit (R). Set to 1b by the MAC at the end of the MDI transaction (for example, indicates a read or write completed). It should be reset to 0b by software at the same time the command is written.
29	RW/V	0b	Interrupt Enable (I). When set to 1b by software, it causes an interrupt to be asserted to indicate the end of an MDI cycle.
30	RW/V	0b	Error (E). This bit set is to 1b by hardware when it fails to complete an MDI read. Software should make sure this bit is clear (0b) before making a MDI read or write command.
31	RO	0b	Reserved. Write as 0b for future compatibility.

This register is used by software to read or write Management Data Interface (MDI) registers in the 82578.

Note: Internal logic uses MDIC to communicate with the 82578. All fields in these registers are indicated as "/V" since the internal logic might use them to access the 82578. Since hardware uses this register, all hardware, software and firmware must use semaphore logic (the ownership flags) before accessing the MDIC.

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

- The CPU performs a write cycle to the MII register with:
 - Ready = 0b
 - Interrupt Enable bit set to 1b or 0b
 - Op-Code = 10b (read)
 - PHYADD = The 82578 address from the MDI register
 - REGADD = The register address of the specific register to be accessed (0 through 31)
- The MAC applies the following sequence on the MDIO signal to the 82578:
 - <PREAMBLE><01><10><PHYADD><REGADD><Z>

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

- The 82578 returns the following sequence on the MDIO signal:
 - <0><DATA><IDLE>
- The MAC discards the leading bit and places the following 16 data bits in the MII register.
- The MAC asserts an Interrupt indicating MDI done, if the *Interrupt Enable* bit was set.
- The MAC sets the *Ready* bit in the MII register indicating the read is complete.



7. The CPU might 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 = 0b
 - Interrupt Enable bit set to 1b or 0b
 - Op-Code = 01b (write)
 - PHYADD = The 82578 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 the 82578
2. The MAC applies the following sequence on the MDIO signal to the 82578:
 - <PREAMBLE><01><01><PHYADD><REGADD><10><DATA><IDLE>
3. The MAC asserts an Interrupt indicating MDI done if the *Interrupt Enable* bit was set.
4. The MAC sets the *Ready* bit in the MII register to indicate step 2 has been completed.
5. The CPU might issue a new MDI command.

Note: An MDI read or write might take as long as 64 ms from the CPU write to the *Ready* bit assertion.

If an invalid opcode is written by software, the MAC does not execute any accesses to the 82578 registers.

If the 82578 does not generate a zero as the second bit of the turnaround cycle for reads, the MAC aborts the access, sets the *E* (error) bit, writes 0xFFFF to the data field to indicate an error condition, and sets the ready bit.

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

This register is initialized to a hardware default only at LAN_RST# 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 0s19 and bits 31:16 are loaded from the NVM word 0x1A.

Bits	Type	Reset	Description
0	RW/SN	0b	Reserved
1	RW/SN	0b	dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1b, clk is always ticking. The default value is 0b (hardware and NVM).
2	RW/SN	0b	wake_dma_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1b, clk is always ticking. The default value is 0b (hardware and NVM)
3	RW/SN	0b	gpt_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1b, clk is always ticking. The default value is 0b (hardware and NVM)
4	RW/SN	0b	mac_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1b, clk is always ticking. The default value is 0b (hardware and NVM)
5	RW/SN	0b	m2k_clk_enable_d. Enable dynamic clock stop. When this bit is set to 1b, clk is always ticking. The default value is 0b (hardware and NVM).
6	RW/SN	0b	Invalid Image CSUM. When cleared, this bit indicates to the Intel NVM programming tools (eupdate) that the image CSUM needs to be corrected. When set the CSUM is assumed to be correct.



Bits	Type	Reset	Description
9:7	RW/SN	0x0	Reserved.
10	RW/SN	0b	Enable MDIO Watchdog Timer (MDIOWatchEna). When set to 0b, the 100 ms MDIO watchdog timer is enabled. Default NVM setting = 1b.
11	RW/SN	0b	Update DMA PTR. 0b = The pointer to the packet header is updated at the start of the packet. 1b = The pointer to the packet header is updated at the end of the previous packet (legacy behavior). Default NVM setting = 0b.
12	RW/SN	0b	MAC Synchronization. 1b = In GbE mode, the MAC does not need to wait for synchronization between clock domains (the clock domains are the same) and the synchronization stage is skipped. 0b = The synchronization stage is not skipped. When operating in 10/100 Mb/s, the synchronization is still needed, therefore it is never skipped. Default NVM setting = 0b.
13	RW/SN	0b	Reserved.
14	RW/SN	0b	Auto PHYINT Clear. 0b = Clears the interrupt indication from the 82578 immediately after the ICR is read. Default NVM setting = 0b.
15	RW/SN	0b	Drop Rx Packet. 0b = Causes packet dropping when it comes, if no descriptors while early receive is enabled. Default NVM setting = 0b.
19:16	RW/SN	0x0	Reserved.
20	RW/SN	0b	Disable CLK gate Enable Due to D3hot. When set, disables assertion of bb_clkgaten due to D3hot. Default NVM setting = 0b.
26:21	RW/SN	0x0	Reserved.
27	RW/SN	0b	Software LCD Config Enable. This bit has no impact on hardware but rather influences the software flow. The software should initialize the 82578 using the extended configuration image in the NVM only when both the <i>Software LCD Config Enable</i> bit is set and the <i>LCD Write Enable</i> bit in the EXTCNF_CTRL register is cleared.
31:28	RW/SN	00b	Reserved.



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

This register is initialized to a hardware default only at LAN_RST# reset. Software should not modify these fields to values other than their recommended values.

Bits	Type	Reset	Description
0		0b	Reserved.
1	RO	0b	Reserved.
3:2	RO/V	00b	Reserved.
7:4	RW	0x0	Reserved.
8	RW	0b	Hardware/Software CRC Mismatch Trigger. When set to 1b the MAC generates a trigger signal each time there is a mismatch between the software calculated CRC and hardware calculated CRC. This feature is ignored when CRC calculation is off-loaded to hardware.
9	RW	0b	Write Disable Ghost and DMA RAMs on CRC Mismatch. When set to 1b, disables any writes to the following RAMs in the event of CRC mismatch until reset: <ul style="list-style-type: none"> • Ghost read PCI descriptor • Ghost read PCI data • The four RAMs in the descriptor engine • The packet buffer
10	RW	0b	When set to 1b, enables the data visibility of the ghost read PCI descriptor and PCI data RAMs to the NOA.
11	RW	0b	Visibility in/out read data select. 1b = in. Bit 10 of the FEXT register must be set to 1b.
12	RW	0b	Visibility data/desc read Ram select. 1b = data. Bit 10 of the FEXT register must be set to 1b.
13	RW	0b	When set to 1b, the ghost read RAMs are readable by the slave bus.
17:14	RW	0x0	Must be set to 0x0.
31:18	RW	0x0	Future Extended. Reserved for future setting.

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

Bit	Type	Reset	Description
7:0	RO	0x0	Reserved.
10:8	RO	000b	Function Number. The MAC is a single PCI function being function 0.
15:11	RO	0x19	Device Number. During normal operation, the MAC has a pre-defined device number equal to 25 (0x19).
23:16	RO	0x0	Bus Number. The MAC captures its bus number during host configuration write cycles type 0 aimed at the device. This field is initialized by LAN_RST# reset, PCI reset, and D3 to D0 transition.
31:24	RO	0x0	Reserved.

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

Bit	Type	Reset	Description
15:0	RW	X	Transmit Timer Value (TTV). Included in XOFF frame.
31:16	RO	0x0	Reserved. Read as 0b. Should be written to 0b 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 needs to send an XON frame, it must set *TTV* to zero prior to initiating the PAUSE frame.

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

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

Bit	Type	Reset	Description
15:0	RW	X	Flow Control Refresh Threshold (FCRT). This value indicates the threshold value of the flow control shadow counter. When the counter reaches this value, and the conditions for a pause state are still valid (buffer fullness above low threshold value), a pause (XOFF) frame is sent to the link partner. The FCRTV timer count interval is the same as other flow control timers and counts at slot times of 64 byte times. If this field contains a zero value, the flow control refresh is disabled.
31:16	RO	0x0	Reserved. Read as 0b. Should be written to 0b for future compatibility.

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

Bit	Type	Reset	Description
0	RW/SN	0b	LCD Write Enable. When set, enables the extended PHY configuration area in the MAC. When disabled, the extended PHY configuration area is ignored. Loaded from NVM word 0x14.
2:1	RW/SN	00b	Reserved
3	RW/SN	1b	OEM Write Enable. When set, enables auto load of the OEM bits from the PHY_CTRL register to the PHY. Loaded from NVM word 0x14.
4	RO	0b	Reserved.
5	RW/V	0b	Software 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	0b	MDIO Hardware Ownership. Hardware requests access to MDIO. Part of the arbitration scheme for MDIO access. This is a RO bit.
15:7	RO	0x0	Reserved.
27:16	RW/SN	0x001	Extended Configuration Pointer. Defines the base address (in Dwords) of the extended configuration area in the NVM.
31:28	RW	0b	Reserved.



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

Bit	Type	Reset	Description
31:24	RO	0x0	Reserved.
23:16	RW/SN	0x0	Extended LCD Length. Size (in Dwords) of the extended PHY configuration area loaded from Extended Configuration word 2 in the NVM. If an extended configuration area is disabled by the <i>LCD Write Enable</i> field in word 0x14 in the NVM, this length must be set to zero.
15:0	RW/SN	0x0	Reserved

10.2.1.1.13 PHY Control Register - PHY_CTRL (0x00F10; RW)

This register is initialized to a hardware default at LAN_RST# reset.

Bit	Type	Reset	Description
31:29	RO	0x0	Reserved
28:25	RO	0x0	SKU Read Data. These four bits contain the SKU value read from the 82578 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.
22:20	RW	0x0	Reserved.
19:17	RW	0x2	Reserved.
16:7	RO	0x0	Reserved
6	RW/SN	0b	Global GbE Disable. Prevents the 82578 from auto negotiating 1000 Mb/s link in all power states (including D0a). This bit is initialized by word 0x17, bit 14 in the NVM.
5:4	RO	00b	Reserved.
3	RW/SN	1b	GbE Disable at Non D0a. Prevents the 82578 from auto negotiating 1000 Mb/s link in all power states except D0a (DR, D0u and D3). Bit is initialized by word 0x17, bit 11 in the NVM. This bit must be set since GbE is not supported in Sx by the platform.
2	RW/SN	1b	LPLU in Non D0a. Enables the 82578 to negotiate for slowest possible link (reverse auto negotiate) in all power states except D0a (DR, D0u and D3). This bit is initialized by word 0x17, bit 10 in the NVM.
1	RW/SN	0b	LPLU in D0a. Enables the 82578 to negotiate for the slowest possible link (reverse auto negotiate) in all power states (including D0a). This bit overrides the LPLU in non-D0abit. This bit is initialized by word 0x17, bit 9 in the NVM.
0	RW/SN	0b	Reserved.



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

Bit	Type	Reset	Description
0	RW	0b	Invert Polarity. Indicates to the GP unit to invert bit polarity (only receiver). this bit is set from the NVM.
6:1	RW	0x20	Command Mode Voltage Select.
31:7	RO	0x0	Reserved. Read as 0b ignore on write.

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

Bit	Type	Reset	Description
4:0	RW	0x	Receive packet buffer allocation (RXA). Defines the size of the Rx buffer in K byte units. Default is KB.
15:5	RO	X	Reserved.
20:16	RO	0x	Transmit Packet Buffer Allocation (TXA). Defines the size of the Tx buffer in KB 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.

Note: Programming this register does not automatically re-load or initialize internal packet-buffer RAM pointers. 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 is not reset by assertion of the software reset, but is only reset upon initial hardware power on.

Note: If early receive functionality is not enabled (indicate field/register), the receive packet buffer should be larger than the maximum expected received packet + 32 bytes.

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

Note: Transmit packet buffer size should be configured to be more than 4 KB.

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

Bit	Type	Reset	Description
5:0		0x0	Packet Buffer Size (PBS). Defines the total packet buffer size both for transmit and receive in 1 KB granularity. Software should keep this register at a value of decimal (equals KB).
31:6	RO	0x0	Reserved. Read as zero.



This register sets the on-chip receive and transmit storage allocation size. The allocation value is read/write for the lower six 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. 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 is not reset by assertion of the software reset, but is only reset upon initial hardware power on.

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

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

Bit	Type	Reset	Description
7:0	RC	0x0	Correctable Error Count (Corr_err_cnt). This counter is incremented every time a correctable error is detected. The counter stops counting after reaching 0xFF. Cleared by read.
15:8	RC	0x0	Uncorrectable Error Count (uncorr_err_cnt). This counter is incremented every time an uncorrectable error is detected. The counter stops counting after reaching 0xFF. Cleared by read.
16	RW	0b	ECC enable.
17	RW	0b	Stop on First Error (SOFE). When set, the ECC test captures the failing address into Last Failure Address (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 1b, when there is ECC failure, the LFA register captures the failing address of the failure.

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

Bit	Type	Reset	Description
0	RW	0b	Inject an error on Tx Buffer on header line. When this bit is set, an error is injected in the next write cycle to a header line of the Tx buffer. Auto cleared by hardware when an error is injected if PBECCINJ.ENECCADD is clear (0b).
1	RW	0b	Inject an error on Tx Buffer on data line. When this bit is set, an error is injected in the next write cycle to a data line of the Tx buffer. Auto cleared by hardware when an error is injected if PBECCINJ.ENECCADD is clear (0b).
2	RW	0b	Inject an error on Rx Buffer on header line. When this bit is set, an error is injected in the next write cycle to a header line of the Rx buffer. Auto cleared by hardware when an error is injected if PBECCINJ.ENECCADD is clear (0b).
3	RW	0b	Inject an error on Rx Buffer on data line. When this bit is set, an error is injected in the next write cycle to a data line of the Rx buffer. Auto cleared by hardware when an error is injected if PBECCINJ.ENECCADD is clear (0b).
15:4	RO	0x0	Reserved.
23:16	RW	0x0	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).



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

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

10.2.1.2 Interrupt Register Descriptions

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

This register is RC or WC. If enabled, read access also clears the ICR content after it is posted to software. Otherwise, a write cycle is required to clear the relevant bit fields. Write a 1b clears the written bit while writing 0b has no affect (with the exception of the INT_ASSERTED bit).

Bit	Type	Reset	Description
0	RWC/CR/V	0b	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	0b	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 is issued if one of the queues is empty and is not cleared until all the queues have valid descriptors.
2	RWC/CR/V	0b	Link Status Change (LSC). This bit is set each time the link status changes (either from up to down, or from down to up). This bit is affected by the LINK indication from the 82578.
3	RO	0b	Reserved.
4	RWC/CR/V	0b	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	0b	Disable Software Write Access (DSW). The DSW bit indicates that firmware changed the status of the DISSW or the DISSWLNK bits in the FWSM register.
6	RWC/CR/V	0b	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	0b	Receiver Timer Interrupt (RXT0). Set when the timer expires.
8	RWC/CR/V	0b	LCAPD Exit Interrupt (LCAPD). Set when the Intel® 5 Series Express Chipset takes the MAC out of LCAPD state.
9	RWC/CR/V	0b	MDIO Access Complete (MDAC). Set when the MDIO access completes.
11:10	RO	00b	Reserved.
12	RWC/CR/V	0b	PHY Interrupt (PHYINT). Set when the 82578 generates an interrupt.
13	RO	0b	Reserved.
14	RWC/CR/V	0b	Reserved.
15	RWC/CR/V	0b	Transmit Descriptor Low Threshold hit (TXD_LOW). Indicates that the descriptor ring has reached the threshold specified in the Transmit Descriptor Control register.



Bit	Type	Reset	Description
16	RWC/CR/V	0b	Small Receive Packet Detected (SRPD). Indicates that a packet size < RSRPD.SIZE register has been detected and transferred to host memory. The interrupt is only asserted if RSRPD.SIZE register has a non-zero value.
17	RWC/CR/V	0b	Receive ACK Frame Detected (ACK). Indicates that an ACK frame has been received and the timer in RAID.ACK_DELAY has expired.
21:18	RWC/CR/V	0x0	Reserved.
22	RWC/CR/V	0b	ECC Error (ECCER). Indicates an uncorrectable EEC error occurred.
30:23	RO	0x0	Reserved. Reads as 0b.
31	RWC/CR/V	0b	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 the MAC. Each time an interrupt causing event occurs, the corresponding interrupt bit is set in this register. An interrupt is generated each time one of the bits in this register is set, and the corresponding interrupt is enabled via the Interrupt Mask Set/Read register (see Section 10.2.1.3.5).

Each time 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 is cleared.
 - Case 2 - Interrupt was asserted (ICR.INT_ASSERTED=1) - ICR content is 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 1b to any bit in the register also clears that bit. Writing a 0b to any bit has no effect on that bit. The INT_ASSERTED bit is a special case. Writing a 1b or 0b to this bit has no affect. It is cleared only when all interrupt sources are cleared.

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

Bit	Type	Reset	Description
15:0	RW	0x0	INTERVAL. Minimum inter-interrupt interval. The interval is specified in 256 ns units. Zero disables interrupt throttling logic.
31:16	RO	0x0	Reserved. Should be written with 0b 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 is not interrupted by the network controller for 128 ms 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 10/100 Mb/s and running at ¼ clock the interval time is multiplied by four.

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

Bit	Type	Reset	Description
0	WO	X	TXDW. Sets transmit descriptor written back.
1	WO	X	TXQE. Sets transmit queue empty.
2	WO	X	LSC. Sets link status change.
3	RO	X	Reserved.
4	WO	X	RXDMT. Sets receive descriptor minimum threshold hit.
5	WO	X	DSW. Sets block software 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	LCAPD. Sets LCAPD interrupt.
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	Reserved.
15	WO	X	TXD_LOW. Transmit descriptor low threshold hit.
16	WO	X	Small Receive Packet Detected (SRPD) 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 0b to ensure future compatibility.

Software uses this register to set an interrupt condition. Any bit written with a 1b sets the corresponding interrupt. This results in the corresponding bit being set in the Interrupt Cause Read register (see [Section 10.2.1.3](#)), 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 10.2.1.3.5](#)).

Bits written with 0b are unchanged.



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

Bit	Type	Reset	Description
0	RWS	0b	TXDW. Sets transmit descriptor written back.
1	RWS	0b	TXQE. Sets transmit queue empty.
2	RWS	0b	LSC. Sets link status change.
3	RO	0b	Reserved.
4	RWS	0b	RXDMTO. Sets mask for receive descriptor minimum threshold hit.
5	RWS	0b	DSW. Sets mask for block software write accesses.
6	RWS	0b	RXO. Sets mask for receiver overrun. Set on receive data FIFO overrun.
7	RWS	0b	RXT0. Sets mask for receiver timer interrupt.
8	RWS	0b	LCAPD. Sets mask for LCAPD interrupt. LCAPD mask is set after reset to enable LCAPD interrupt (driven by Intel® 5 Series Express Chipset).
9	RWS	0b	MDAC. Sets mask for MDIO access complete interrupt.
11:10	RO	00b	Reserved.
12	RWS	0b	PHYINT. Sets mask for PHY interrupt.
13	RO	0b	Reserved.
14	RWS	0b	Reserved.
15	RWS	0b	TXD_LOW. Sets the mask for transmit descriptor low threshold hit.
16	RWS	0b	SRPD. Sets mask for small receive packet detection.
17	RWS	0b	ACK. Sets the mask for receive ACK frame detection.
18	RWS	0b	MNG. Sets mask for manageability event interrupt.
19	RWS	0b	Reserved.
20	RWS	0b	Reserved.
21	RO	0b	Reserved.
22	RWS	0b	ECCER Sets mask for uncorrectable EEC error
31:23	RO	0x0	Reserved. Should be written with 0b 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 1b, and disabled if its corresponding mask bit is set to 0b. An interrupt is generated each time 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 10.2.1.3](#)).

A particular interrupt might be enabled by writing a 1b to the corresponding mask bit in this register. Any bits written with a 0b 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 10.2.1.3.6](#)), rather than writing a 0b to a bit in this register.

When the CTRL_EXT.INT_TIMERS_CLEAR_ENA bit is set, then following writing all 1b'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.



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

Bit	Type	Reset	Description
0	WO	0b	TXDW. Sets transmit descriptor written back.
1	WO	0b	TXQE. Sets transmit queue empty.
2	WO	0b	LSC. Sets link status change.
3	RO	0b	Reserved.
4	WO	0b	RXDMT0. Clears mask for receive descriptor minimum threshold hit.
5	WO	0b	DSW. Clears mask for block software Write accesses.
6	WO	0b	RXO. Clears mask for receiver overrun.
7	WO	0b	RXT0. Clears mask for receiver timer interrupt.
8	WO	0b	LCAPD. Clears mask for LCAPD interrupt.
9	WO	0b	MDAC. Clears mask for MDIO access complete interrupt.
11:10	RO	00b	Reserved. Reads as 0b.
12	WO	0b	PHYINT. Clears PHY interrupt.
13	RO	0b	Reserved.
14	WO	0b	Reserved.
15	WO	0b	TXD_LOW. Clears the mask for transmit descriptor low threshold hit.
16	WO	0b	SRPD. Clears mask for small receive packet detect interrupt.
17	WO	0b	ACK. Clears the mask for receive ACK frame detect interrupt.
18	WO	0b	MNG. Clears mask for the manageability event interrupt.
19	WO	0b	Reserved.
20	WO	0b	Reserved.
21	RO	0b	Reserved.
22	WO	0b	ECCER Clears the mask for uncorrectable EEC error.
31:23	RO	0x0	Reserved. Should be written with 0b 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 1b and the cause bit is a 1b. 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 10.2.1.3](#)).

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

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



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

Bit	Type	Reset	Description
31:0	RW	0x0	IAM_VALUE. When the CTRL_EXT.IAME bit is set and the ICR.INT_ASSERTED=1, an ICR read or write has the side effect of writing the contents of this register to the IMC register.

10.2.1.3 Receive Register Descriptions

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

Bit	Type	Reset	Description
0	RO	0b	Reserved. This bit represents 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 0b for future compatibility.
1	RW	0b	Enable (EN). The receiver is enabled when this bit is 1b. Writing this bit to 0b stops reception after receipt of any in progress packets. All subsequent packets are then immediately dropped until this bit is set to 1b. 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	0b	Store bad packets (SBP). 0b = Do not store bad packets. 1b = 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 MAC (even bad packets). Note: Packet errors are not routed to manageability even if this bit is set.
3	RW	0b	Unicast promiscuous enable (UPE). 0b = Disabled. 1b = Enabled.
4	RW	0b	Multicast promiscuous enable (MPE). 0b = Disabled. 1b = Enabled.
5	RW	0b	Long packet enable (LPE). 0b = Disabled. 1b = Enabled.
7:6	RW	00b	Reserved.
9:8	RW	0b	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set each time the fractional number of free descriptors becomes equal to RDMTS. Table 84 lists which fractional values correspond to RDMTS values. See Section 10.2.1.4.8 for details regarding RDLEN.
11:10	RW	00b	Descriptor Type (DTYP). 00b = Legacy or extended descriptor type. 01b = Packet split descriptor type. 10b and 11b = Reserved.
13:12	RW	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.
14	RW	0b	Reserved.



Bit	Type	Reset	Description
15	RW	0b	Broadcast Accept Mode (BAM). 0b = Ignore broadcast (unless it matches through exact or imperfect filters). 1 = Accept broadcast packets.
17:16	RW	00b	Receive Buffer Size (BSIZE). RCTL.BSEX – zero: 00b = 2048 bytes. 01b = 1024 bytes. 10b = 512 bytes. 11b = 256 bytes. RCTL.BSEX – one: 00b = Reserved. 01b = 16384 bytes. 10b = 8192 bytes. 11b = 4096 bytes. BSIZE is only used when DTYP – 00b. When DTYP – 01b, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than zero, in that case, FLXBUF determines the buffer size.
21:18	RO	0x0	Reserved. Should be written with 0b.
22	RW	0b	Reserved.
23	RW	0b	Pass MAC Control Frames (PMCF). 0b = 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	0b	Reserved. Should be written with 0b to ensure future compatibility.
25	RW	0b	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE). 0b = Buffer size is as defined in BSIZE. 1b = Original BSIZE values are multiplied by 16.
26	RW	0b	Strip Ethernet CRC from incoming packet (SECRC). 0b = Does not strip CRC. 1b = 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	0x0	FLXBUF. Determines a flexible buffer size. When this field is 0000b, the buffer size is determined by BSIZE. If this field is different from 0000b, the receive buffer size is the number represented in KB: For example, 0001 = 1 KB (1024 bytes).
31	RO	0b	Reserved. Should be written with 0b to ensure future compatibility.



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

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

Table 84. RDMTS Values

RDMTS	Free Buffer Threshold
00b	1/2
01b	1/4
10b	1/8
11b	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 the 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 = 1b then frames meeting this criteria is DMA'd to host memory.

The SECRC bit controls whether 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.



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

Bit	Type	Reset	Description
7:0	RO	0x0	Reserved. This bit represents 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 0b for future compatibility.
9:8	RW	00b	Receive Descriptor Minimum Threshold Size (RDMTS). The corresponding interrupt is set each time the fractional number of free descriptors becomes equal to RDMTS. Table 84 lists which fractional values correspond to RDMTS values. See Section 10.2.1.4.8 for details regarding RDLEN.
11:10	RW	00b	Descriptor Type (DTYP). 00b = Legacy or Extended descriptor type. 01b = Packet Split descriptor type. 10b and 11b = Reserved. The value of RCTL1.DTYP should be the same as RCTL.DTYP
15:12	RO	0x0	Reserved.
17:16	RW	00b	Receive Buffer Size (BSIZE). RCTL.BSEX - zero: 00b = 2048 Bytes. 01b = 1024 Bytes. 10b = 512 Bytes. 11b = 256 Bytes. RCTL.BSEX - one: 00b = Reserved. 01b = 16384 Bytes. 10b = 8192 Bytes. 11b = 4096 Bytes. BSIZE is only used when DTYP - 00b. When DTYP - 01b, the buffer sizes for the descriptor are controlled by fields in the PSRCTL register. BSIZE is not relevant when the FLXBUF is other than zero, in that case, FLXBUF determines the buffer size.
24:18	RO	0x0	Reserved. Should be written with 0b.
25	RW	0b	Buffer Size Extension (BSEX). Modifies buffer size indication (BSIZE above). 0b = Buffer size is as defined in BSIZE. 1b = Original BSIZE values are multiplied by 16.
26	RW	0b	Reserved. Should be written with 0b.
30:27	RW	0x0	FLXBUF. Determine a flexible buffer size. When this field is 0000b, the buffer size is determined by BSIZE. If this field is different from 0000b, the receive buffer size is the number represented in KB. For example, 0001b = 1 KB (1024 bytes).
31	RO	0b	Reserved. Should be written with 0b to ensure future compatibility.



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

Bit	Type	Reset	Description
12:0	RW	0x0	Receive Threshold Value (RxThreshold). This threshold is in units of eight bytes.
21:13	RO	0x0	Reserved.
31:22	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.

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

This field has a granularity of eight 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 0b, then the MAC starts the PCI transfer only after the entire packet is contained in the MAC's receive packet buffer. the MAC 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 MAC acquires the bus, it attempts 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 causes additional PCI bursts for the packet. In other words, this register enables software to trade-off latency versus bus utilization. Too high a value effectively eliminates the early receive benefits (at least for short packets) and too low a value deteriorates 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 has 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.



10.2.1.3.4 Packet Split Receive Control Register - PSRCTL (0x02170)

Bit	Type	Reset	Description
6:0	RW	0x2	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 KB). Default buffer size is 256 bytes. Software should not program this field to a zero value.
7	RO	0b	Reserved. Should be written with 0b to ensure future compatibility.
13:8	RW	0x4	Receive Buffer Size for Buffer 1 (BSIZE1). The value is in 1 KB resolution. Value can be from 1 KB to 63 KB. Default buffer size is 4 KB. Software should not program this field to a zero value.
15:14	RO	00b	Reserved. Should be written with 0b to ensure future compatibility.
21:16	RW	0x4	Receive Buffer Size for Buffer 2 (BSIZE2). The value is in 1 KB resolution. Value can be from 1 KB to 63 KB. Default buffer size is 4 KB. Software might program this field to any value.
23:22	RO	00b	Reserved. Should be written with 0b to ensure future compatibility.
29:24	RW	0x0	Receive Buffer Size for Buffer 3 (BSIZE3). The value is in 1 KB resolution. Value can be from 1 KB to 63 KB. Default buffer size is 0 KB. Software might program this field to any value.
31:30	RO	00b	Reserved. Should be written with 0b to ensure future compatibility.

Note: If software 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.

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

Bit	Type	Reset	Description
2:0	RO	0x0	Reserved. The underlying bits might not be implemented in all versions of the chip. Must be written with 0b.
15:3	RW	0x0	Receive Threshold Low (RTL). FIFO low water mark for flow control transmission.
30:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.
31	RW	0b	XON Enable (XONE). 0b = Disabled. 1b = Enabled.

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

Note that flow control reception/transmission are negotiated capabilities by the auto-negotiation process. When the MAC is manually configured, flow control operation is determined by the RFCE and TFCE bits of the Device Control register.



10.2.1.3.6 Flow Control Receive Threshold High - FCRTH (0x02168; RW)

Bit	Type	Reset	Description
2:0	RO	0x0	Reserved. Must be written with 0.
15:3	RW	0x0	Receive Threshold High (RTH). FIFO high water mark for flow control transmission.
31:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0b 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 eight bytes less than the maximum number of bytes allocated to the Receive Packet Buffer (PBA, RXA), and the lower three bits must be programmed to zero (8-byte granularity). Each time 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 MAC is manually configured, flow control operation is determined by the RFCE and TFCE bits of the Device Control register.

10.2.1.3.7 Receive Descriptor Base Address Low Queue - RDBAL (0x02800; RW)

Bit	Type	Reset	Description
3:0	RO	0x0	Reserved. Ignored on writes. Returns 0b 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 four bits are always ignored. The receive descriptor base address must point to a 16-byte aligned block of data.

10.2.1.3.8 Receive Descriptor Base Address High Queue - RDBAH (0x02804; RW)

Bits	Type	Reset	Description
31:0	RW	X	Receive Descriptor Base Address [63:32] (RDBAH).

This register contains the upper 32 bits of the 64-bit descriptor base address.



10.2.1.3.9 Receive Descriptor Length Queue- RDLEN (0x02808; RW)

Bits	Type	Reset	Description
6:0	RO	0x0	Reserved. Ignore on write. Reads back as 0b.
19:7	RW	0x0	Descriptor Length (LEN)
31:20	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.

This register sets the number of bytes allocated for descriptors in the circular descriptor buffer. It must be 128-byte aligned.

Note: The descriptor ring must be equal to or larger than eight descriptors.

10.2.1.3.10 Receive Descriptor Head Queue - RDH (0x02810; RW)

Bits	Type	Reset	Description
15:0	RW/V	0x0	Receive Descriptor Head (RDH).
31:16	RO	0x0	Reserved. Should be written with 0b.

This register contains the head pointer for the receive descriptor buffer. The register points to a 16-byte 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 might be invalidated and hardware could become unstable.

10.2.1.3.11 Receive Descriptor Tail Queue - RDT (0x02818; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Receive Descriptor Tail (RDT).
31:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.

This register contains the tail pointer for the receive descriptor buffer. The register points to a 16-byte datum. Software writes the tail register to add receive descriptors for hardware to process.

10.2.1.3.12 Interrupt Delay Timer (Packet Timer) - RDTR (0x02820; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Receive Delay Timer. Receive packet delay timer measured in increments of 1.024 ms.
30:16	RO	0x0	Reserved. Reads as 0b.
31	WO	0b	Flush Partial Descriptor Block (FPD), when set to 1b, ignored otherwise. Reads 0b.

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 0b 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 cancels a pending RDTR interrupt. The RDTR countdown timer is reloaded but stopped, so as to avoid generation of a spurious second interrupt after the RADV has been noted, but might be restarted by a subsequent received packet.

Note: FPD is self-clearing.

10.2.1.3.13 Receive Descriptor Control - RXDCTL (0x02828; RW)

Bits	Type	Reset	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
14	RW	0b	Reserved.
15	RW	0b	Reserved.
21:16	RW	0x01	Write-Back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0b	Granularity (GRAN). Units for the thresholds in this register. 0b = Cache lines. 1b = Descriptors.
31:25	RO	0x00	Reserved.

Note: This register was not fully validated. Software should set it to 0x0000 during normal operation.

This register controls the fetching and write back of receive descriptors. The three threshold values are used to determine when descriptors is read from and written to host memory. The values might be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag. If GRAN=zero (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 = 0b) or (WTHRESH = 1b and GRAN = 1b) only descriptors with the RS bit set is written back.



PTHRESH is used to control when a prefetch of descriptors is 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 considers pre-fetching descriptors from host memory. This fetch does 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 occurs 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 zero, 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 one, 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 the MAC.

10.2.1.3.14 Receive Interrupt Absolute Delay Timer- RADV (0x0282C; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Receive Absolute Delay Timer. Receive absolute delay timer measured in increments of 1.024 ms (0b = disabled).
31:16	RO	0x0	Reserved. Reads as 0b.

If the packet delay timer is used to coalesce receive interrupts, it ensures that when receive traffic abates, an interrupt is generated within a specified interval of no receives. During times when receive traffic is continuous, it might be necessary to ensure that no receive remains unnoticed for too long an interval. This register might be used to ENSURE that a receive interrupt occurs at some pre-defined 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 zero disables the absolute timer mechanism (the RDTR register should be used with a value of zero to cause immediate interrupts for all receive packets).

Receive interrupts due to a Receive Packet Timer (RDTR) expiration cancels 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.

10.2.1.3.15 Receive Small Packet Detect Interrupt- RSRPD (0x02C00; RW)

Bits	Type	Reset	Description
11:0	RW	0x0	SIZE. If the interrupt is enabled any receive packet of size <= SIZE asserts 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.

10.2.1.3.16 Receive ACK Interrupt Delay Register - RAID (0x02C08; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	ACK_DELAY. ACK delay timer measured in increments of 1.024 ms. When the Receive ACK frame detect interrupt is enabled in the IMS register, ACK packets being received uses a unique delay timer to generate an interrupt. When an ACK is received, an absolute timer loads to the value of ACK_DELAY. The interrupt signal is set only when the timer expires. If another ACK packet is received while the timer is counting down, the timer is not reloaded to ACK_DELAY.
31:16	RO	0x0	Reserved.

If an immediate (non-scheduled) interrupt is desired for any received ACK frame, the ACK_DELAY should be set to zero.

10.2.1.3.17 Receive Checksum Control - RXCSUM (0x05000; RW)

Bits	Type	Reset	Description
7:0	RW	0x00	Packet Checksum Start (PCSS).
8	RW	1b	IP Checksum Offload Enable (IPOFL).
9	RW	1b	TCP/UDP Checksum Offload Enable (TUOFL).
11:10	RO	00b	Reserved.
12	RW	0b	IP Payload Checksum Enable (IPPCSE).
13	RW	0b	Packet Checksum Disable (PCSD).
14	RW	0b	Reserved.
31:15	RO	0x0	Reserved.

The Receive Checksum Control register controls the receive checksum offloading features of the MAC. The MAC supports the offloading 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 following table:



RXCSUM.PCSD	0 (Checksum Enable)	1 (Checksum Disable)
Legacy Rx descriptor (RCTL.DTYP = 00b)	Packet checksum is reported in the Rx descriptor	Not supported
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 previously noted, 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 (zero 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 does 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 erroneously calculates IP header checksum or TCP/UDP checksum.

RXCSUM.IPOFLD is used to enable the IP Checksum offloading feature. If RXCSUM.IPOFLD is set to one, the MAC calculates 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, the MAC calculates 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, the MAC calculates 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 offloading only. Supported frame types:

- Ethernet II
- Ethernet SNAP

This register should only be initialized (written) when the receiver is not enabled (only write this register when RCTL.EN = 0).



10.2.1.3.18 Receive Filter Control Register - RFCTL (0x05008; RW)

Bits	Type	Reset	Description
0	RW	0b	iSCSI Disable (ISCSI_DIS). Disable the iSCSI filtering for header split functionality.
5:1	RW	0x0	iSCSI DWord Count (ISCSI_DWC). This field indicated the Dword count of the iSCSI header, which is used for packet split mechanism.
6	RW	0b	NFS Write Disable (NFSW_DIS). Disable filtering of NFS write request headers for header split functionality.
7	RW	0b	NFS Read Disable (NFSR_DIS). Disable filtering of NFS read reply headers for header split functionality.
9:8	RW	00b	NFS Version (NFS_VER). 00b = NFS version 2. 01b = NFS version 3. 10b = NFS version 4. 11b = Reserved for future use.
10	RW	0b	Reserved.
11	RW	0b	Reserved.
12	RW	0b	ACK Accelerate Disable (ACKDIS). When this bit is set the MAC does not accelerate interrupt on TCP ACK packets.
13	RW	0b	ACK data Disable (ACKD_DIS). 1b = MAC recognizes ACK packets according to the ACK bit in the TCP header + No -CP data 0b = MAC recognizes ACK packets according to the ACK bit only. This bit is relevant only if the ACKDIS bit is not set.
14	RW	0b	IP Fragment Split Disable (IPFRSP_DIS). When this bit is set the header of IP fragmented packets are not set.
15	RW	0b	Extended Status Enable (EXSTEN). When the EXSTEN bit is set or when the packet split receive descriptor is used, the MAC writes the extended status to the Rx descriptor.
17:16		0x0	Reserved.
31:18	RO	0x0	Reserved. Should be written with 0b to ensure future compatibility.

10.2.1.3.19 Multicast Table Array - MTA[31:0] (0x05200-0x0527C; RW)

Bits	Type	Reset	Description
31:0	RW	X	Bit Vector. Word wide bit vector specifying 32 bits in the multicast address filter table.

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

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

Figure 19 shows 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 19. Multicast Table Array Algorithm

10.2.1.3.20 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	Reset	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.

10.2.1.3.21 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	Reset	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: 00b = Destination address (must be set to this in normal mode). 01b = Source address. 10b = Reserved. 11b = Reserved.
18	RW	0b	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	0x0	Reserved. Reads as 0b. Ignored on write.
31	RW	See as follows	Address Valid (AV). Cleared after master reset. If the NVM is present, the <i>Address Valid</i> field of the Receive Address Register 0 is set to 1b 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 MAC 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.

After reset, if the NVM is present, the first register (Receive Address register 0) is loaded from the IA field in the NVM, its *Address Select* field is 00b, and its *Address Valid* field is 1b. If no NVM is present the *Address Valid* field is 0b. The *Address Valid* field for all of the other registers is zero.

10.2.1.3.22 Shared Receive Address Low - SHRAL[n] (0x05438 + 8*n (n=0...3); RW)

Bits	Type	Reset	Description
31:0	RW	X	Receive Address Low (RAL). The lower 32 bits of the 48-bit Ethernet address n (n=0...3).

10.2.1.3.23 Shared Receive Address High 0...2 - SHRAH[n] (0x0543C + 8*n (n=0...2); RW)

Bits	Type	Reset	Description
15:0	RW	X	Receive Address High (RAH). The upper 16 bits of the 48-bit Ethernet address n (n=0...3).
17:16	RO	00b	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	0b	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	0x0	Reserved. Reads as 0b. Ignored on write.
31	RW	0b	Address valid (AV). When this bit is set, the relevant RAL,RAH are valid (compared against the incoming packet).



10.2.1.3.24 Shared Receive Address High 3 - SHRAH[3] (0x05454; RW)

Bits	Type	Reset	Description
15:0	RW	X	Receive Address High (RAH). The upper 16 bits of the 48-bit Ethernet address n (n=0...3).
17:16	RO	00b	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	0b	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	0x0	Reserved. Reads as 0b. Ignored on write.
30	RW	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.
31	RW	0b	Address valid (AV). When this bit is set the relevant address 3 is valid (compared against the incoming packet).

10.2.1.3.25 Multiple Receive Queues Command register - MRQC (0x05818; RW)

Bits	Type	Reset	Description
1:0	RW	0x00b	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 Microsoft* 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.
15:2		0x0	Reserved.
21:16	RW	0x0	RSS Field Enable. Each bit, when set, enables a specific field selection to be used by the hash function. Several bits can be set at the same time. Bit[16] = Enable TcpIPv4 hash function. Bit[17] = Enable IPv4 hash function. Bit[18] = Enable TcpIPv6 hash function. Bit[19] = Enable IPv6Ex hash function. Bit[20] = Enable IPv6 hash function. Bit[21] = Reserved.
31:22	RO	0x0	Reserved.



10.2.1.3.26 Redirection Table - RETA (0x05C00 + 4 * n (n=0...31); RW)

The re-direction table is a 32 entry table. Each entry is composed of four tags each 8-bits wide. Only the first or last six bits of each tag are used (five bits for the CPU index and 1 bit for queue index).

Offset	31:24	23:16	15:8	7:0
0x05C00 + n*4	Tag 4*n+3	Tag 4*n+2	Tag 4*n+1	Tag 4*n

Bits	Type	Reset	Description
4:0	RW	X	CPU INDX 0. CPU index for Tag 4*n (n=0,1,...31).
6:5	RO	X	Reserved.
7	RW	X	QUE INDX 0. Queue Index for Tag 4*n (n=0,1,...31).
12:8	RW	X	CPU INDX 1. CPU index for Tag 4*n+1 (n=0,1,...31).
14:13	RO	X	Reserved.
15	RW	X	QUE INDX 1. Queue Index for Tag 4*n+1 (n=0,1,...31).
20:16	RW	X	CPU INDX 2. CPU index for Tag 4*n+2 (n=0,1,...31).
22:21	RO	X	Reserved.
23	RW	X	QUE INDX 2. Queue Index for Tag 4*n+2 (n=0,1,...31).
28:24	RW	X	CPU INDX 3. CPU index for Tag 4*n+3 (n=0,1,...31).
30:29	RO	X	Reserved.
31	RW	X	QUE INDX 3. Queue Index for Tag 4*n+3 (n=0,1,...31).

Note: RETA cannot be read when RSS is enabled.

10.2.1.3.27 Random Key Register - RSSRK (0x05C80 + 4 * n (n=0...9); RW)

The RSS Random Key register stores a 40-byte key (10 Dword entry table) used by the RSS hash function.

Bits	Type	Reset	Description
7:0	RW	0x0	K0. Byte n*4 of the RSS random key (n=0,1,...9).
15:8	RW	0x0	K1. Byte n*4+1 of the RSS random key (n=0,1,...9).
23:16	RW	0x0	K2. Byte n*4+2 of the RSS random key (n=0,1,...9).
31:24	RW	0x0	K3. Byte n*4+3 of the RSS random key (n=0,1,...9).



10.2.1.4 Transmit Register Descriptions

10.2.1.4.1 Transmit Control Register - TCTL (0x00400; RW)

Bits	Type	Reset	Description
0	RW	0b	<p>IP Identification 15 bit (IPID15).</p> <p>When set to 1b, the <i>IP Identification</i> field is incremented and wrapped around on 15-bit base. For example, if IP ID is equal to 0x7FFF then the next value is 0x0000; if IP ID is equal to 0xFFFF then the next value is 0x8000.</p> <p>When set to 0b, the <i>IP Identification</i> field is incremented and wrapped around on 16-bit base. In this case, the value following 0xFFFF is 0x8000, and the value following 0xFFFF is 0x0000.</p> <p>The purpose of this feature is to enable the software to manage two sub-groups of connections.</p>
1	RW	0b	<p>Enable (EN). The transmitter is enabled when this bit is set to 1b. Writing this bit to 0b stops transmission after any in-progress packets are sent. Data remains in the transmit FIFO until the MAC is re-enabled. Software should combine this with reset if the packets in the FIFO should be flushed.</p>
2	RO	0b	<p>Reserved. Reads as 0b. Should be written to 0b for future compatibility.</p>
3	RW	1b	<p>Pad Short Packets (PSP). With valid data, NOT padding symbols.</p> <p>0b = Do not pad 1b = Pad.</p> <p>Padding makes the packet 64 bytes. This is not the same as the minimum collision distance.</p> <p>If padding of short packets is allowed, the value in Tx descriptor length field should be not less than 17 bytes.</p>
11:4	RW	0x0F	<p>Collision Threshold (CT). This determines the number of attempts at re-transmission 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.</p>
21:12	RW	0x3F	<p>Collision Distance (COLD). Specifies the minimum number of byte times that 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 64-byte to 512-byte times.</p>
22	RW/V	0b	<p>Software XOFF Transmission (SWXOFF). When set to a 1b, the MAC schedules 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.</p>
23	RW	0b	<p>Reserved.</p>
24	RW	0b	<p>Re-transmit on Late Collision (RTLCL). Enables the MAC to re-transmit on a late collision event.</p>
27:25	RW	0x0	<p>Reserved. Used to be UNORTX and TXDSCMT in predecessors.</p>
28		1b	<p>Reserved.</p>
30:29	RW	01b	<p>Read Request Threshold (RRTHRESH). These bits 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:</p> <p>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).</p>
31	RO	0b	<p>Reserved. Reads as 0. Should be written to 0 for future compatibility.</p>

Two fields deserve special mention: CT and COLD. Software might choose to abort packet transmission in less than the Ethernet mandated 16 collisions. For this reason, hardware provides CT.



Wire speeds of 1000 Mb/s 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 Mb/s environments to 100 Mb/s environments. Note that 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 MAC. Software should not write a 1b to this bit while the MAC is configured for half-duplex operation.

RTL configures the MAC to perform re-transmission 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 1000 Mb/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.

10.2.1.4.2 Transmit IPG Register - TIPG (0x00410; RW)

Bits	Type	Reset	Description
9:0	RW	0x8	IPG Transmit Time (IPGT). Specifies the IPG length for back-to-back transmissions equal to [(IPGT+4) x 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 x 8) when DJHDX=0 and equals to (IPGR1+2) x 8 when DJHDX=1.
29:20	RW	0x9	IPG Receive Time 2 (IPGR2). Specifies the defer IPG. Equal to (IPGR2+3) x 8 when DJHDX=0 and equal to (IPGR2+5) x 8 when DJHDX=1.
31:30	RO	00b	Reserved. Reads as 0b. Should be written to 0b for future compatibility.

This register controls the Inter Packet Gap (IPG) 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 eight is recommended to achieve a 12-byte time IPG.

IPGR1 specifies the portion of the IPG in which the transmitter defers to receive events. This should be set to 2/3 of the total effective IPG, or eight.

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 seven 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 overrides the IPGT IPG setting in half duplex, resulting in inter packet gaps that are larger than intended by IPGT in that case. Full duplex is unaffected by this, and always relies 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.



10.2.1.4.3 Adaptive IFS Throttle - AIT (0x00458; RW)

Bits	Type	Reset	Description
15:0	RW	0x0000	Adaptive IFS value (AIFS). This value is in units of 8 ns.
31:16	RO	0x0000	Reserved. This field should be written with 0b.

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 zero. However, if additional delay is desired between back-to-back transmits, then this register might 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 10.2.1.5.2](#)). However, it only affects the initial transmission timing, not re-transmission timing.

Note: If the value of the AdaptiveIFS field is less than the IPGTransmitTime field in the Transmit IPG registers then it has no effect, as the chip selects the maximum of the two values.

10.2.1.4.4 Transmit Descriptor Base Address Low - TDBAL (0x03800 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
3:0	RO	0x0	Reserved. Ignored on writes. Returns 0b 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 four bits are ignored. The transmit descriptor base address must point to a 16-byte aligned block of data.

10.2.1.4.5 Transmit Descriptor Base Address High - TDBAH (0x03804 + n*0x100[n=0..1]; RW)

Bits	Type	Reset	Description
31:0	RW	X	Transmit Descriptor Base Address [63:32] (TDBAH).

This register contains the upper 32 bits of the 64-bit descriptor base address.



10.2.1.4.6 Transmit Descriptor Length - TDLEN (0x03808 ; RW)

Bits	Type	Reset	Description
6:0	RO	0x0	Reserved. Ignore on write. Reads back as 0b.
19:7	RW	0x0	Descriptor Length (LEN).
31:20	RO	0x0	Reserved. Reads as 0b. Should be written to 0b.

This register contains the descriptor length and must be 128-byte aligned.

Note: The descriptor ring must be equal to or larger than eight descriptors.

10.2.1.4.7 Transmit Descriptor Head - TDH (0x03810; RW)

Bits	Type	Reset	Description
15:0	RW/V	0x0	Transmit Descriptor Head (TDH).
31:16	RO	0x0	Reserved. Should be written with 0b.

This register contains the head pointer for the transmit descriptor ring. It points to a 16-byte 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 might be invalidated and hardware could become confused.

10.2.1.4.8 Transmit Descriptor Tail - TDT (0x03818; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Transmit Descriptor Tail (TDT).
31:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0 for future compatibility.

Note: This register contains the tail pointer for the transmit descriptor ring. It points to a 16-byte 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.

10.2.1.4.9 Transmit Interrupt Delay Value - TIDV (0x03820; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Interrupt Delay Value (IDV). Counts in units of 1.024 ms. A value of zero is not allowed.
30:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.
31	WO	0b	Flush Partial Descriptor Block (FPD). when set to 1b; ignored otherwise. Reads 0b.

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 zero is not allowed. If an immediate (non-scheduled) interrupt is desired for any transmit descriptor, the descriptor IDE should be set to zero.

The occurrence of either an immediate (non-scheduled) or absolute transmit timer interrupt halts 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=1b, IDE=0b) cancels a pending TIDV interrupt. The TIDV countdown timer is reloaded but halted, though it might be restarted by a processing a subsequent transmit descriptor.

Writing this register with FPD set initiates an immediate expiration of the timer, causing a write back of any consumed transmit descriptors pending write back, and results in a transmit timer interrupt in the ICR.

Note: FPD is self-clearing.



10.2.1.4.10 Transmit Descriptor Control - TXDCTL (0x03828; RW)

Note: This register was not fully validated. Software should set it to 0x0000 during nominal operation.

Bits	Type	Reset	Description
5:0	RW	0x00	Prefetch Threshold (PTHRESH).
7:6	RO	0x00	Reserved.
13:8	RW	0x00	Host Threshold (HTHRESH).
15:14	RO	0x00	Reserved.
21:16	RW	0x00	Write-Back Threshold (WTHRESH).
23:22	RO	0x00	Reserved.
24	RW	0x0	Granularity (GRAN). Units for the thresholds in this register. 0b = Cache lines. 1b = 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 is read from and written to host memory. The values might be in units of cache lines or descriptors (each descriptor is 16 bytes) based on the GRAN flag.

Note: When GRAN = one all descriptors is written back (even if not requested).

PTHRESH is used to control when a prefetch of descriptors is 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 considers pre-fetching descriptors from host memory. This fetch does 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 each time 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 occurs 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 zero - 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 zero, 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 MAC. If the WTHRESH is written to a non-zero value then all of the descriptors are written back consecutively no matter the setting of the *RS* bit.

Since write back of transmit descriptors is optional (under the control of *RS* bit in the descriptor), not all processed descriptors are counted with respect to WTHRESH. Descriptors start accumulating after a descriptor with *RS* is set. Furthermore, with transmit descriptor bursting enabled, all of the descriptors are written back consecutively no matter the setting of the *RS* bit.

LWTHRESH controls the number of pre-fetched transmit descriptors at which a transmit descriptor-low interrupt (ICR.TXD_LOW) is reported. This might enable software to operate more efficiently by maintaining a continuous addition of transmit work, interrupting only when hardware nears completion of all submitted work. LWTHRESH specifies a multiple of eight 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 zero disables this feature.

10.2.1.4.11 Transmit Absolute Interrupt Delay Value-TADV (0x0382C; RW)

Bits	Type	Reset	Description
15:0	RW	0x0	Interrupt Delay Value (IDV). Counts in units of 1.024 ms. (0b = disabled)
31:16	RO	0x0	Reserved. Reads as 0b. Should be written to 0b for future compatibility.

The transmit interrupt delay timer (TIDV) might be used to coalesce transmit interrupts. However, it might 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 might 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 halts the TADV timer and eliminate any spurious second interrupts.

Setting the value to zero 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 zero.



10.2.1.5 Power Management Register Descriptions

10.2.1.5.1 Wake Up Control Register - WUC (0x05800; RW)

Bits	Type	Reset	Description
0	RW/ SN	0b	Advance Power Management Enable (APME). 1b = APM Wakeup is enabled. 0b = APM Wakeup is disabled. Loaded from the NVM word 0x0A.
1	RW/V	0b	PME_En. This read/write bit is used by the driver to access the PME_En bit of the Power Management Control / Status Register (PMCSR) without writing to PCI configuration space.
2	RWC	0b	PME_Status. This bit is set when the MAC receives a wake-up event. It is the same as the PME_Status bit in the PMCSR. Writing a 1b to this bit clears it, and also clears the PME_Status bit in the PMCSR.
3	RW	1b	Assert PME On APM Wakeup (APMPME). If set to 1b, the MAC sets the PMCSR and asserts Host_Wake when APM wake up is enabled and the MAC receives a matching magic packet.
4	RW/ SN	0b	Link Status Change Wake Enable (LSCWE). Enables wake on link status change as part of APM wake capabilities.
5	RW/ SN	0b	Link Status Change Wake Override (LSCWO). If set to 1b, wake on link status change does not depend on the LNK bit in the Wake Up Filter Control (WUFC) register. Instead, it is determined by the APM settings in the WUC register.
7:6	RO	00b	Reserved.
8	RW/ SN	0b	Phy_Wake. This bit indicates if the 82578 connected to the MAC supports wake up. This bit is loaded from NVM word 0x13, bit 8.
29:9	RO	0x0	Reserved. Reads as 0.
31:30	RO	00b	Reserved.

The *PME_Status* bits are cleared in the following conditions:

- If there is VAUX, then the *PME Status* bits should be cleared by:
 - LAN_RST# or PCI reset
 - Explicit software clear
- If there is NO VAUX, then the *PME Status* bits should be cleared by:
 - LAN_RST# or PCI reset
 - PCI reset de-assertion
 - Explicit software clear



10.2.1.5.2 Wake Up Filter Control Register - WUFC (0x05808; RW)

Bits	Type	Reset	Description
0	RW	0b	LNKC. Link Status Change Wake Up Enable.
1	RW	0b	MAG. Magic Packet Wake Up Enable.
2	RW	0b	EX. Directed Exact Wake Up Enable.
3	RW	0b	MC. Directed Multicast Wake Up Enable.
4	RW	0b	BC. Broadcast Wake Up Enable.
5	RW	0b	IPv4. Request Packet Wake Up Enable.
6	RW	0b	IPV4. Directed IPv4 Packet Wake Up Enable.
7	RW	0b	IPV6. Directed IPv6 Packet Wake Up Enable.
14:8	RO	0x0	Reserved.
15	RW	0b	NoTCO. Ignore TCO Packets for TCO. If the <i>NoTCO</i> bit is set, then any packet that passes the manageability packet filtering does not cause a wake up event even if it passes one of the wake up filters.
16	RW	0b	FLX0. Flexible Filter 0 Enable.
17	RW	0b	FLX1. Flexible Filter 1 Enable.
18	RW	0b	FLX2. Flexible Filter 2 Enable.
19	RW	0b	FLX3. Flexible Filter 3 Enable.
20	RW	0b	FLX4. Flexible Filter 4 Enable.
21	RW	0b	FLX5. Flexible Filter 5 Enable.
31:2	RO	0x0	Reserved.

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.

10.2.1.5.3 Wake Up Status Register - WUS (0x05810; RW)

Bits	Type	Reset	Description
0	RW	0b	LNKC. Link Status Changed.
1	RW	0b	MAG. Magic Packet Received.
2	RW	0b	EX. Directed Exact Packet Received. The packet's address matched one of the seven pre-programmed exact values in the Receive Address registers.
3	RW	0b	MC. Directed Multicast Packet Received. The packet was a multicast packet that hashed to a value corresponding to a one bit in the Multicast Table Array.
4	RW	0b	BC. Broadcast Packet Received.
5	RW	0b	IPv4. Request Packet Received.
6	RW	0b	IPV4. Directed IPv4 Packet Received.
7	RW	0b	IPV6. Directed IPv6 Packet Received.
15:8	RO	0x0	Reserved. Read as 0b.
16	RW	0b	FLX0. Flexible Filter 0 Match.
17	RW	0b	FLX1. Flexible Filter 1 Match.
18	RW	0b	FLX2. Flexible Filter 2 Match.
19	RW	0b	FLX3. Flexible Filter 3 Match.



Bits	Type	Reset	Description
20	RW	0b	FLX4. Flexible Filter 4 Match.
21	RW	0b	FLX5. Flexible Filter 5 Match.
31:2	RO	0x0	Reserved.

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

This register is not cleared when PCI_RST_N is asserted. It is only cleared when LAN_RST# is de-asserted or when cleared by the driver.

10.2.1.5.4 IP Address Valid - IPAV (0x5838; RW)

The IP address valid indicates whether the IP addresses in the IP address table are valid:

Bits	Type	Reset	Description
0	RO	0b	Reserved.
1	RW	0b	V41. IPv4 Address 1 Valid.
2	RW	0b	V42. IPv4 Address 2 Valid.
3	RW	0b	V43. IPv4 Address 3 Valid.
15:4	RO	0x00	Reserved.
16	RW	0b	V60. IPv6 Address Valid.
31:17	RO	0x00	Reserved.

10.2.1.5.5 IPv4 Address Table - IP4AT (0x05840 + 8*n (n=1...3); RW)

The IPv4 address table is used to store the three IPv4 addresses for IPv4 request packet and directed IPv4 packet wake up. It is a 4-entry table with the following format:

Bits	Type	Reset	Description
31:0	RW	X	IPADD. IP Address n (n=1, 2, 3).

The register at address 0x5840 (n=0) was used in predecessors and reserved in the Intel® 5 Series Express Chipset.

10.2.1.5.6 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	Reset	Description
31:0	RW	X	IPv6 Address. IPv6 Address bytes n*4...n*4+3 (n=0, 1, 2, 3) while byte 0 is first on the wire and byte 15 is last.



10.2.1.5.7 Flexible Filter Length Table - FFLT (0x05F00 + 8*n (n=0..5); RW)

There are six 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 does not pass that filter. Each flexible filter considers 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 does not check any bytes past that point.

Bits	Type	Reset	Description
10:0	RW	X	LEN. Minimum Length for Flexible Filter n.
31:11	RO	X	Reserved.

All reserved fields read as 0b's and ignore writes.

Note: Before writing to the flexible filter length table the driver must first disable the flexible filters by writing 0b's to the *Flexible Filter Enable* bits of the Wake Up Filter Control register (WUFC.FLXn).

10.2.1.5.8 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 1b, 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.

Bits	Type	Reset	Description
0	RW	X	Mask 0. Mask for filter 0 byte n (n=0, 1... 127).
1	RW	X	Mask 1. Mask for filter 1 byte n (n=0, 1... 127).
2	RW	X	Mask 2. Mask for filter 2 byte n (n=0, 1... 127).
3	RW	X	Mask 3. Mask for filter 3 byte n (n=0, 1... 127).
4	RW	X	Mask 4. Mask for filter 4 byte n (n=0, 1... 127).
5	RW	X	Mask 5. Mask for filter 5 byte n (n=0, 1... 127).
31:	RO	X	Reserved.

Note: The table is organized to permit expansion to eight (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 0b's to the *Flexible Filter Enable* bits of the Wake Up Filter Control register (WUFC.FLXn).

10.2.1.5.9 Flexible Filter Value Table - FFVT (0x09800 + 8*n (n=0...127); RW)

There are 128 filter values. The flexible filter value is used to store the one value for each byte location in a packet for each flexible filter. If the corresponding mask bit is 1b, the flexible filter compares 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).



Bits	Type	Reset	Description
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 0b's to the *Flexible Filter Enable* bits of the Wake Up Filter Control register (WUFC.FLXn).

10.2.1.5.10 Flexible Filter Value Table - FFVT2 (0x09804 + 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 1b, the flexible filter compares the incoming data byte to the values stored in this table.

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

Note: Before writing to the flexible filter value table the driver must first disable the flexible filters by writing 0b's to the *Flexible Filter Enable* bits of the Wake Up Filter Control register (WUFC.FLXn).



Note: This page intentionally left blank.



11.0 Electrical and Timing Specifications

11.1 Introduction

This section describes the 82578's recommended operating conditions, power delivery, DC electrical characteristics, power sequencing and reset requirements, PCIe specifications, reference clock, and packaging information.

11.2 Operating Conditions

11.2.1 Absolute Maximum Ratings

Symbol	Parameter	Min	Max	Units
T_{case}	Case Temperature Under Bias	0	85	°C
$T_{storage}$	Storage Temperature Range	-45	125	°C
V_i/V_o	3.3 Vdc I/O Voltage 2.5 Vdc I/O Voltage 1.2 Analog Vdc I/O Voltage 1.8 Analog Vdc Voltage	$V_{ss} - 0.5$ $V_{ss} - 0.4$ $V_{ss} - 0.2$ $V_{ss} - 0.3$	4.6 3.5 1.68 2.52	Vdc
VCC	3.3 Vdc Periphery DC Supply Voltage	$V_{ss} - 0.5$	4.6	Vdc
VCC	2.5 Vdc Core DC Supply Voltage	$V_{ss} - 0.4$	3.5	Vdc
VCC1p8	1.8 Vdc Supply Voltage	$V_{ss} - 0.3$	2.52	Vdc
VCC1p2	1.2 Vdc Supply Voltage	$V_{ss} - 0.2$	1.68	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).



11.2.2 Recommended Operating Conditions

Symbol	Parameter	Min	Max	Units
Ta	Operating Temperature Range Commercial (Ambient; 0 CFS airflow)	0	85 ¹	°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.

11.2.3 ESD Specifications

Title	Specification
Human body model	JESD22-A114
Charged device model	JESD22-C101
Machine model	JESD22-A115
Cable discharge event	N/A

11.3 Power Delivery

The following power options are available for the 82578:

1. Connecting the 82578 to two external power supplies with nominal voltages of 3.3 Vdc and 1.2 Vdc.
2. Powering the 82578 with an external 3.3 Vdc supply and using an internal power regulator as follows:
 - Regulator mode:
 - 2.5 Vdc is generated from 3.3 Vdc using an internal regulator.
 - 1.8 Vdc is generated from 3.3 Vdc using an internal regulator for central tap voltage only.
 - 1.2 Vdc is generated externally using an external PnP.



11.3.1 Voltage Regulator Power Supply Specifications

Note: These requirements apply when using an external power source.

11.3.1.1 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	3.6	V
Ripple	Maximum voltage ripple (peak to peak) @ 20 MHz BW	N/A	70	mV
Overshoot	Maximum overshoot allowed	N/A	100	mV

11.3.1.2 1.8 Vdc Rail

Title	Description	Min	Max	Units
Operational Range	Voltage range for normal operating conditions	1.71	2.25	Vdc
Ripple	Maximum voltage ripple (peak to peak) @ 20 MHz BW	N/A	50	mV



11.3.1.3 1.2 Vdc Rail

Title	Description	Min	Max	Units
Rise Time	Time from 10% to 90% mark	0.1	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	1.14	1.26	Vdc
Ripple	Maximum voltage ripple (peak to peak) @ 20 MHz BW	N/A	50	mV
Overshoot	Maximum overshoot allowed	N/A	100	mV
Decoupling Capacitance	Capacitance range	10	30	μF
Capacitance ESR	Equivalent series resistance of output capacitance	5	50	$\text{m}\Omega$

11.3.1.4 1.2 Vdc PNP Regulator Power Delivery Schematic

Note: See Figure 4 for the 1.2 Vdc PNP regulator power delivery schematic.

11.3.1.5 PNP Specifications

Title	Description	Min	Max	Units
VCBO		20		Vdc
VCEO		20		Vdc
IC(max)		1		A
IC(peak)		1.2		A
Ptot	Minimum total dissipated power @ 25 °C ambient temperature	1.5		W
hFE	DC current gain @ Vce = -10 Vdc, Ic = 500 mA	85		
hfe	AC current gain @ Ic = 50 mA VCE = -10 Vdc, f = 20 MHz	2.5		
Cc	Collector capacitance @ VCB=-5 Vdc, f = 1 MHz		50	pF
fT	Transition frequency @ Ic = 10 mA, VCE = -5 Vdc, f = 100 MHz	40		MHz
Recommended transistor	BCP69			
Ib		50 μA	4 mA	

Note: Maximum current of 1.2 Vdc is less than 350 mA.



11.3.1.6 External Components

Description	Name	Qty	Electrical Characteristics	Recommended Components		Pkg
				Source	Part #	
1.2 Vdc Regulator PNP Transistor	Q1	1	<ul style="list-style-type: none"> Minimum HFE (Vdc Gain) 85 @ $V_{ce} = 2.5 \text{ Vdc}$ $I = 0.35\text{A}$ $T = 25^\circ\text{C}$ $R_{ja} < 60^\circ\text{C/W}$ 	Philips, OnSemi, Infineon	BCP69	SOT223

11.3.2 Power On/Off Sequence

The 82578 does not require a power on or power off sequence between the 3.3 Vdc and 1.2 Vdc power rails.

Table 85. Power-On Reset Detection Thresholds

Symbol	Parameter	Specifications			Units
		Min	Typ	Max	
V1a	High-threshold for 3.3 Vdc supply	2.35	2.5	2.75	Vdc
V2a	Low-threshold for 3.3 Vdc supply	2.3	2.5	2.7	Vdc
V1b	High-threshold for 1.2 Vdc supply	0.8	0.9	0.95	Vdc
V2b	Low-threshold for 1.2 Vdc supply	0.75	0.8	0.9	Vdc

11.4 I/O DC/AC Parameters

11.4.1 3.3 Vdc DC/IO

Note: All the 3.3 Vdc I/Os are open-drain types.

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.4	0	0.4	Vdc
VIH		2	3.3	3.6	Vdc
VOL		-0.4	0	0.4	Vdc
VOH		2.4	3.3	3.6	Vdc
Ipullup		10	20	30	μA
Ileakage				10	μA
Ci			2	4	pF



Signal Name	Bus Size	Description
CLK_REQ_N ¹	1	Open drain I/O
SMB_CLK	1	Open drain I(H)/O
SMB_DATA	1	Open drain I(H)/O

1. I_{leakage} applies only when the PHY is powered on.

11.4.2 2.5 Vdc/IO

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.4	0	0.4	Vdc
VIH		2	2.6	3.3	Vdc
VOL	I _{OL} = 10 mA VCC = Min	-0.4	0	0.4	Vdc
VOH	I _{OH} = -8 mA VCC = Min	2	2.6	2.8	Vdc
I _{pullup}		10	20	30	μA
I _{leakage}		15 (pull down)	25 (pull down)	35 (pull down)	μA
C _i			2	4	pF
PU			4.7		KΩ
PD			4.7		KΩ



Signal Name	Bus Size	Description
RSVD_VCC3P3	2	I/O, PU
LED[2:0]	3	I/O, PU
JTAG_TDI ¹	1	I/O, PU
JTAG_TMS ¹	1	I/O, PU
JTAG_TDO	1	I/O, PU
JTAG_TCK ¹	1	I/O, PU

1. It might be desirable to drive these input signals from +3.3 Vdc level CMOS drivers.

11.4.3 Input Buffer Only

Parameter	Conditions	Minimum	Typical	Maximum	Unit
VIL		-0.4	0	0.4	Vdc
VIH		2	3.3	3.6	Vdc
Ipullup		10	20	30	μA
Ileakage				10	μA
Ci			2	4	pF

Signal Name	Bus Size	Description
Internal Power On Reset/ LAN_DISABLE_N	1	I(H), PU
TEST_EN	1	I (no PU, no PD)
PE_RST_N	1	I(H), PU



11.4.4 SMBus AC I/O


Refer to the System Management Bus (SMBus) Specification Version 2.0.

11.4.5 PCIe DC/AC Specifications

11.4.5.1 PCIe Specifications (Transmitter)

Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
UI	Unit interval	799.92	800.08	ps	Each UI is 800 pS +/- 100 ppm
$V_{tx-diff-pp}$	Differential peak-to-peak Tx voltage swing	0.8	1.2	Vdc	
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	
$T_{tx-rise-fall}$	Transmitter rise and fall time	0.125		UI	
$RL_{tx-diff}$	Tx package plus silicon differential return loss	10		db	
RL_{tx-cm}	Tx package plus silicon common mode return loss	6		db	
$Z_{tx-diff-dc}$	DC differential Tx impedance	80	120	Ω	
$V_{tx-cm-ac-p}$	Tx Vac 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	Vdc	
$V_{tx-cm-dc-active-idle-delta}$	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	
$T_{tx-idle-min}$	Minimum time spent in electrical idle	20		ns	
$T_{tx-idle-set-to-idle}$	Maximum time to transition to a valid electrical idle after sending an EIOS		8	ns	
$T_{tx-idle-to-diff-data}$	Maximum time to transition to valid differential signaling after leaving electrical idle		8	ns	

Note: Figure 20 is for informational purposes only. Do not use for actual eye comparisons.



Note: Not To Scale

Figure 20. Transmitter Eye Diagram


11.4.5.2 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	Vdc	
$V_{rx-diff-pp-dc}$	Differential peak-to-peak Rx voltage swing for data clock	0.175	1.2	Vdc	
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	
$BW_{rx-pll-hi}$	Maximum Rx PLL bandwidth	N/A	22	MHz	
$BW_{rx-pll-lo-3db}$	Minimum Rx PLL bandwidth for 3 dB peaking	1.5	N/A	MHz	
$RL_{rx-diff}$	Rx differential return loss	10	N/A	dB	
RL_{rx-cm}	Rx CM return loss	6	N/A	dB	
Z_{rx-dc}	Rx CM DC impedance	40	60	Ω	
$Z_{rx-diff-dc}$	Rx differential Vdc impedance	80	120	Ω	
$V_{rx-cm-ac-p}$	Rx Vac CM voltage	N/A	150	mVp	



Symbol	Parameter	1.25 GT/s		Units	Comments
		Min	Max		
$Z_{rx-high-imp-dc-pos}$	DC input CM impedance for $V > 0$	50 K	N/A	Ω	
$Z_{rx-high-imp-dc-neg}$	DC input CM impedance for $V < 0$	1 K	N/A	Ω	
$V_{rx-idle-det-diff-p}$	Electrical idle detect threshold	65	175	mV	
$T_{rx-idle-det-diff-entertime}$	Unexpected electrical idle detect	N/A	10	ms	

Note: Figure 21 is intended to show the difference between the PCIe 1.0 and PCIe-based receiver sensitivity templates. It is for informational purposes only.



Note: Not To Scale

Figure 21. Receiver Eye Diagram



11.5 Discrete/Integrated Magnetics Specifications

Criteria	Condition	Values (Min/Max)
Voltage Isolation	At 50 to 60 Hertz for 60 seconds	1500 Vrms (min)
	For 60 seconds	2250 Vdc (min)
Open Circuit Inductance (OCL) or OCL (alternate)	With 8 mA DC bias at 25 °C	400 μH (min)
	With 8 mA DC bias at 0 °C to 70 °C	350 μH (min)
Insertion Loss	100 kHz through 999 kHz	1 dB (max)
	1.0 MHz through 60 MHz	0.6 dB (max)
	60.1 MHz through 80 MHz	0.8 dB (max)
	80.1 MHz through 100 MHz	1.0 dB (max)
	100.1 MHz through 125 MHz	2.4 dB (max)
Return Loss	1.0 MHz through 40 MHz 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)



11.6 Oscillator/Crystal Specifications

Table 86. 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]		
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 ¹	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

1. Crystal must meet or exceed the specified drive level (D_L). Refer to the crystal design guidelines in the Intel® 5 Series Family PDG for more details.

Table 87. Clock Oscillator Specifications

Parameter Name	Symbol/Parameter	Conditions	Min	Typ	Max	Unit
Frequency	f_o	@25 [°C]		25.0		MHz
Swing	Vp-p		3	3.3	3.6 ¹	V
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]
Coupling Capacitor	C coupling		8	10	12	pF
XTAL_X1 swing in - High	V _{iH}		1.0		1.2	V
XTAL_X1 swing in - Low	V _{iL}		0		0.2	V
TH_XTAL_IN	XTAL_IN High Time		13	20		nS
TL_XTAL_IN	XTAL_IN Low Time		13	20		nS
TJ_XTAL_IN	XTAL_IN Total Jitter				200 ²	pS

1. Conditioning circuit required to limit the voltage swing of V_{iH}/V_{iL} to 1.2 Vdc.
2. Broadband peak-to-peak = 200 pS, Broadband rms = 3 pS, 12 KHz to 20 MHz rms = 1 ps




Figure 22. XTAL Timing Diagram



Note: Peak-to-peak voltage presented at the XTAL1 input cannot exceed 1.8 Vdc.

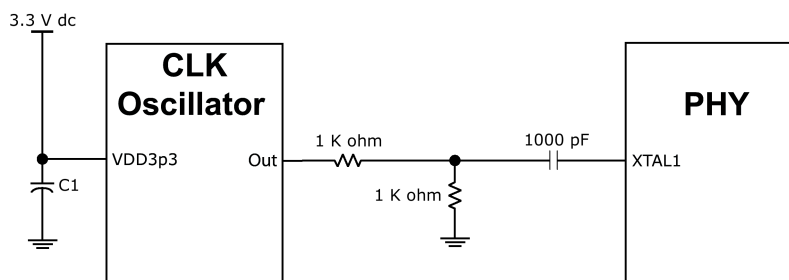


Figure 23. Clock Oscillator Schematic



12.0 Schematic and Board Layout Checklists

The 82578 Design and Board Layout Checklists can be found at www.intel.com.



Note: This page intentionally left blank.



13.0 Reference Schematics

The 82578 reference schematics can be found at www.intel.com.



Note: This page intentionally left blank.



14.0 Models

Contact your Intel Representative for access to the 82578 IBIS model.



Note: This page intentionally left blank.