
2-Port 10/100 Managed Ethernet Switch with 8/16-Bit Non-PCI CPU Interface

Highlights

- High performance 2-port switch with VLAN, QoS packet prioritization, rate limiting, IGMP monitoring and management functions
- Interfaces to most 8/16-bit embedded controllers and 32-bit embedded controllers with an 8/16-bit bus
- Integrated Ethernet PHYs with HP Auto-MDIX
- Compliant with Energy Efficient Ethernet 802.3az
- Wake on LAN (WoL) support
- Integrated IEEE 1588v2 hardware time stamp unit
- Cable diagnostic support
- 1.8V to 3.3V variable voltage I/O
- Integrated 1.2V regulator for single 3.3V operation

Target Applications

- Cable, satellite, and IP set-top boxes
- Digital televisions & video recorders
- VoIP/Video phone systems, home gateways
- Test/Masurement equipment, industrial automation

Key Benefits

- Ethernet Switch Fabric
 - 32K buffer RAM, 512 entry forwarding table
 - Port based IEEE 802.1Q VLAN support (16 groups)
 - Programmable IEEE 802.1Q tag insertion/removal
 - IEEE 802.1D spanning tree protocol support
 - 4 separate transmit queues available per port
 - Fixed or weighted egress priority servicing
 - QoS/CoS Packet prioritization
 - Input priority determined by VLAN tag, DA lookup, TOS, DIFFSERV or port default value
 - Programmable Traffic Class map based on input priority on per port basis
 - Remapping of 802.1Q priority field on per port basis
 - Programmable rate limiting at the ingress with coloring and random early discard, per port / priority
 - Programmable rate limiting at the egress with leaky bucket algorithm, per port / priority
 - IGMP v1/v2/v3 monitoring for Multicast packet filtering
 - Programmable broadcast storm protection with global % control and enable per port
 - Programmable buffer usage limits
 - Dynamic queues on internal memory
 - Programmable filter by MAC address
- Switch Management
 - Port mirroring/monitoring/sniffing: ingress and/or egress traffic on any port or port pair
 - Fully compliant statistics (MIB) gathering counters

- Ports
 - 2 internal 10/100 PHYs with HP Auto-MDIX support
 - Fully compliant with IEEE 802.3 standards
 - 10BASE-T and 100BASE-TX support
 - 100BASE-FX support via external fiber transceiver
 - Full and half duplex support, full duplex flow control
 - Backpressure (forced collision) half duplex flow control
 - Automatic flow control based on programmable levels
 - Automatic 32-bit CRC generation and checking
 - Programmable interframe gap, flow control pause value
 - Auto-negotiation, polarity correction & MDI/MDI-X
- 8/16-Bit Host Bus Interface
 - Indexed register or multiplexed bus
 - SPI / Quad SPI support
- IEEE 1588v2 hardware time stamp unit
 - Global 64-bit tunable clock
 - Boundary clock: master / slave, one-step / two-step, end-to-end / peer-to-peer delay
 - Transparent Clock with Ordinary Clock: master / slave, one-step / two-step, end-to-end / peer-to-peer delay
 - Fully programmable timestamp on TX or RX, timestamp on GPIO
 - 64-bit timer comparator event generation (GPIO or IRQ)
- Comprehensive power management features
 - 3 power-down levels
 - Wake on link status change (energy detect)
 - Magic packet wakeup, Wake on LAN (WoL), wake on broadcast, wake on perfect DA
 - Wakeup indicator event signal
- Power and I/O
 - Integrated power-on reset circuit
 - Latch-up performance exceeds 150mA per EIA/JESD78, Class II
 - JEDEC Class 3A ESD performance
 - Single 3.3V power supply (integrated 1.2V regulator)
- Additional Features
 - Multifunction GPIOs
 - Ability to use low cost 25MHz crystal for reduced BOM
- Packaging
 - Pb-free RoHS compliant 72-pin QFN or 80-pin TQFP-EP
- Available in commercial and industrial temp. ranges

TO OUR VALUED CUSTOMERS

It is our intention to provide our valued customers with the best documentation possible to ensure successful use of your Microchip products. To this end, we will continue to improve our publications to better suit your needs. Our publications will be refined and enhanced as new volumes and updates are introduced.

If you have any questions or comments regarding this publication, please contact the Marketing Communications Department via E-mail at docerrors@microchip.com. We welcome your feedback.

Most Current Documentation

To obtain the most up-to-date version of this documentation, please register at our Worldwide Web site at:

<http://www.microchip.com>

You can determine the version of a data sheet by examining its literature number found on the bottom outside corner of any page. The last character of the literature number is the version number, (e.g., DS30000000A is version A of document DS30000000).

Errata

An errata sheet, describing minor operational differences from the data sheet and recommended workarounds, may exist for current devices. As device/documentation issues become known to us, we will publish an errata sheet. The errata will specify the revision of silicon and revision of document to which it applies.

To determine if an errata sheet exists for a particular device, please check with one of the following:

- Microchip's Worldwide Web site; <http://www.microchip.com>
- Your local Microchip sales office (see last page)

When contacting a sales office, please specify which device, revision of silicon and data sheet (include -literature number) you are using.

Customer Notification System

Register on our web site at www.microchip.com to receive the most current information on all of our products.

1.0 Preface	4
2.0 General Description	8
3.0 Pin Descriptions and Configuration	10
4.0 Power Connections	30
5.0 Register Map	33
6.0 Clocks, Resets, and Power Management	41
7.0 Configuration Straps	60
8.0 System Interrupts	73
9.0 Host Bus Interface	86
10.0 SPI/SQI Slave	134
11.0 Host MAC	152
12.0 Ethernet PHYs	218
13.0 Switch Fabric	320
14.0 I2C Master EEPROM Controller	457
15.0 IEEE 1588	473
16.0 General Purpose Timer & Free-Running Clock	559
17.0 GPIO/LED Controller	563
18.0 Miscellaneous	572
19.0 JTAG	577
20.0 Operational Characteristics	579
21.0 Package Outlines	594
22.0 Revision History	597

LAN9352

1.0 PREFACE

1.1 General Terms

TABLE 1-1: GENERAL TERMS

Term	Description
10BASE-T	10 Mbps Ethernet, IEEE 802.3 compliant
100BASE-TX	100 Mbps Fast Ethernet, IEEE802.3u compliant
ADC	Analog-to-Digital Converter
ALR	Address Logic Resolution
AN	Auto-Negotiation
BLW	Baseline Wander
BM	Buffer Manager - Part of the switch fabric
BPDU	Bridge Protocol Data Unit - Messages which carry the Spanning Tree Protocol information
Byte	8 bits
CSMA/CD	Carrier Sense Multiple Access/Collision Detect
CSR	Control and Status Registers
CTR	Counter
DA	Destination Address
DWORD	32 bits
EPC	EEPROM Controller
FCS	Frame Check Sequence - The extra checksum characters added to the end of an Ethernet frame, used for error detection and correction.
FIFO	First In First Out buffer
FSM	Finite State Machine
GPIO	General Purpose I/O
Host	External system (Includes processor, application software, etc.)
IGMP	Internet Group Management Protocol
Inbound	Refers to data input to the device from the host
Level-Triggered Sticky Bit	This type of status bit is set whenever the condition that it represents is asserted. The bit remains set until the condition is no longer true and the status bit is cleared by writing a zero.
lsb	Least Significant Bit
LSB	Least Significant Byte
LVDS	Low Voltage Differential Signaling
MDI	Medium Dependent Interface
MDIX	Media Independent Interface with Crossover
MII	Media Independent Interface
MIIM	Media Independent Interface Management
MIL	MAC Interface Layer
MLD	Multicast Listening Discovery
MLT-3	Multi-Level Transmission Encoding (3-Levels). A tri-level encoding method where a change in the logic level represents a code bit "1" and the logic output remaining at the same level represents a code bit "0".
msb	Most Significant Bit
MSB	Most Significant Byte

TABLE 1-1: GENERAL TERMS (CONTINUED)

Term	Description
NRZI	Non Return to Zero Inverted. This encoding method inverts the signal for a “1” and leaves the signal unchanged for a “0”
N/A	Not Applicable
NC	No Connect
OUI	Organizationally Unique Identifier
Outbound	Refers to data output from the device to the host
PISO	Parallel In Serial Out
PLL	Phase Locked Loop
PTP	Precision Time Protocol
RESERVED	Refers to a reserved bit field or address. Unless otherwise noted, reserved bits must always be zero for write operations. Unless otherwise noted, values are not guaranteed when reading reserved bits. Unless otherwise noted, do not read or write to reserved addresses.
RTC	Real-Time Clock
SA	Source Address
SFD	Start of Frame Delimiter - The 8-bit value indicating the end of the preamble of an Ethernet frame.
SIPO	Serial In Parallel Out
SMI	Serial Management Interface
SQE	Signal Quality Error (also known as “heartbeat”)
SSD	Start of Stream Delimiter
UDP	User Datagram Protocol - A connectionless protocol run on top of IP networks
UUID	Universally Unique Identifier
WORD	16 bits

1.2 Buffer Types

TABLE 1-2: BUFFER TYPES

Buffer Type	Description
IS	Schmitt-triggered input
VIS	Variable voltage Schmitt-triggered input
VO8	Variable voltage output with 8 mA sink and 8 mA source
VOD8	Variable voltage open-drain output with 8 mA sink
VO12	Variable voltage output with 12 mA sink and 12 mA source
VOD12	Variable voltage open-drain output with 12 mA sink
VOS12	Variable voltage open-source output with 12 mA source
VO16	Variable voltage output with 16 mA sink and 16 mA source
PU	50 μ A (typical) internal pull-up. Unless otherwise noted in the pin description, internal pull-ups are always enabled. Internal pull-up resistors prevent unconnected inputs from floating. Do not rely on internal resistors to drive signals external to the device. When connected to a load that must be pulled high, an external resistor must be added.
PD	50 μ A (typical) internal pull-down. Unless otherwise noted in the pin description, internal pull-downs are always enabled. Internal pull-down resistors prevent unconnected inputs from floating. Do not rely on internal resistors to drive signals external to the device. When connected to a load that must be pulled low, an external resistor must be added.
AI	Analog input
AIO	Analog bidirectional
ICLK	Crystal oscillator input pin
OCLK	Crystal oscillator output pin
ILVPECL	Low voltage PECL input pin
OLVPECL	Low voltage PECL output pin
P	Power pin

1.3 Register Nomenclature

TABLE 1-3: REGISTER NOMENCLATURE

Register Bit Type Notation	Register Bit Description
R	Read: A register or bit with this attribute can be read.
W	Read: A register or bit with this attribute can be written.
RO	Read only: Read only. Writes have no effect.
WO	Write only: If a register or bit is write-only, reads will return unspecified data.
WC	Write One to Clear: Writing a one clears the value. Writing a zero has no effect
WAC	Write Anything to Clear: Writing anything clears the value.
RC	Read to Clear: Contents is cleared after the read. Writes have no effect.
LL	Latch Low: Clear on read of register.
LH	Latch High: Clear on read of register.
SC	Self-Clearing: Contents are self-cleared after the being set. Writes of zero have no effect. Contents can be read.
SS	Self-Setting: Contents are self-setting after being cleared. Writes of one have no effect. Contents can be read.
RO/LH	Read Only, Latch High: Bits with this attribute will stay high until the bit is read. After it is read, the bit will either remain high if the high condition remains, or will go low if the high condition has been removed. If the bit has not been read, the bit will remain high regardless of a change to the high condition. This mode is used in some Ethernet PHY registers.
NASR	Not Affected by Software Reset. The state of NASR bits do not change on assertion of a software reset.
RESERVED	Reserved Field: Reserved fields must be written with zeros to ensure future compatibility. The value of reserved bits is not guaranteed on a read.

2.0 GENERAL DESCRIPTION

The LAN9352 is a full featured, 2 port 10/100 managed Ethernet switch designed for embedded applications where performance, flexibility, ease of integration and system cost control are required. The LAN9352 combines all the functions of a 10/100 switch system, including the Switch Fabric, packet buffers, Buffer Manager, Media Access Controllers (MACs), PHY transceivers, and host bus interface. IEEE 1588v2 is supported via the integrated IEEE 1588v2 hardware time stamp unit, which supports end-to-end and peer-to-peer transparent clocks. The LAN9352 complies with the IEEE 802.3 (full/half-duplex 10BASE-T and 100BASE-TX) Ethernet protocol, IEEE 802.3az Energy Efficient Ethernet (EEE) (100Mbps only), and 802.1D/802.1Q network management protocol specifications, enabling compatibility with industry standard Ethernet and Fast Ethernet applications. 100BASE-FX is supported via an external fiber transceiver.

At the core of the device is the high performance, high efficiency 3 port Ethernet Switch Fabric. The Switch Fabric contains a 3 port VLAN layer 2 Switch Engine that supports untagged, VLAN tagged, and priority tagged frames. The Switch Fabric provides an extensive feature set which includes spanning tree protocol support, multicast packet filtering and Quality of Service (QoS) packet prioritization by VLAN tag, destination address, port default value or DIFFSERV/TOS, allowing for a range of prioritization implementations. 32K of buffer RAM allows for the storage of multiple packets while forwarding operations are completed, and a 512 entry forwarding table provides ample room for MAC address forwarding tables. Each port is allocated a cluster of 4 dynamic QoS queues which allow each queue size to grow and shrink with traffic, effectively utilizing all available memory. This memory is managed dynamically via the Buffer Manager block within the Switch Fabric. All aspects of the Switch Fabric are managed via the Switch Fabric configuration and status registers, which are indirectly accessible via the system control and status registers.

The LAN9352 provides 2 switched ports. Each port is fully compliant with the IEEE 802.3 standard and all internal MACs and PHYs support full/half duplex 10BASE-T and 100BASE-TX operation. The LAN9352 provides 2 on-chip PHYs, 1 Virtual PHY and 3 MACs. The Virtual PHY and the Host MAC are used to connect the LAN9352 switch fabric to the host bus interface. All ports support automatic or manual full duplex flow control or half duplex backpressure (forced collision) flow control. Automatic 32-bit CRC generation/checking and automatic payload padding are supported to further reduce CPU overhead. 2K jumbo packet (2048 byte) support allows for oversized packet transfers, effectively increasing throughput while decreasing CPU load. All MAC and PHY related settings are fully configurable via their respective registers within the device.

Two user selectable host bus interface options are available:

- **Indexed register access**

This implementation provides three index/data register banks, each with independent Byte/WORD to DWORD conversion. Internal registers are accessed by first writing one of the three index registers, followed by reading or writing the corresponding data register. Three index/data register banks support up to 3 independent driver threads without access conflicts. Each thread can write its assigned index register without the issue of another thread overwriting it. Two 16-bit cycles or four 8-bit cycles are required within the same 32-bit index/data register - however, these access can be interleaved. Direct (non-indexed) read and write accesses are supported to the packet data FIFOs. The direct FIFO access provides independent Byte/WORD to DWORD conversion, supporting interleaved accesses with the index/data registers. Direct FIFO access also supports burst reading of the data FIFO.

- **Multiplexed address/data bus**

This implementation provides a multiplexed address and data bus with both single phase and dual phase address support. The address is loaded with an address strobe followed by data access using a read or write strobe. Two back to back 16-bit data cycles or 4 back to back 8-bit data cycles are required within the same 32-bit DWORD. These accesses must be sequential without any interleaved accesses to other registers. Burst read and write accesses are supported to the packet data and status FIFOs by performing one address cycle followed by multiple read or write data cycles.

The HBI supports 8/16-bit operation with big, little, and mixed endian operations. Four separate FIFO mechanisms (TX/RX Data FIFO's, TX/RX Status FIFO's) interface the HBI to the Host MAC and facilitate the transferring of packet data and status information between the host CPU and the switch fabric. A configurable host interrupt pin allows the device to inform the host CPU of any internal interrupts.

An SPI / Quad SPI slave controller provides a low pin count synchronous slave interface that facilitates communication between the device and a host system. The SPI / Quad SPI slave allows access to the System CSRs, internal FIFOs and memories. It supports single and multiple register read and write commands with incrementing, decrementing and static addressing. Single, Dual and Quad bit lanes are supported with a clock rate of up to 80 MHz.

The LAN9352 supports numerous power management and wakeup features. The LAN9352 can be placed in a reduced power mode and can be programmed to issue an external wake signal (PME) via several methods, including “Magic Packet”, “Wake on LAN”, wake on broadcast, wake on perfect DA, and “Link Status Change”. This signal is ideal for triggering system power-up using remote Ethernet wakeup events. The device can be removed from the low power state via a host processor command or one of the wake events.

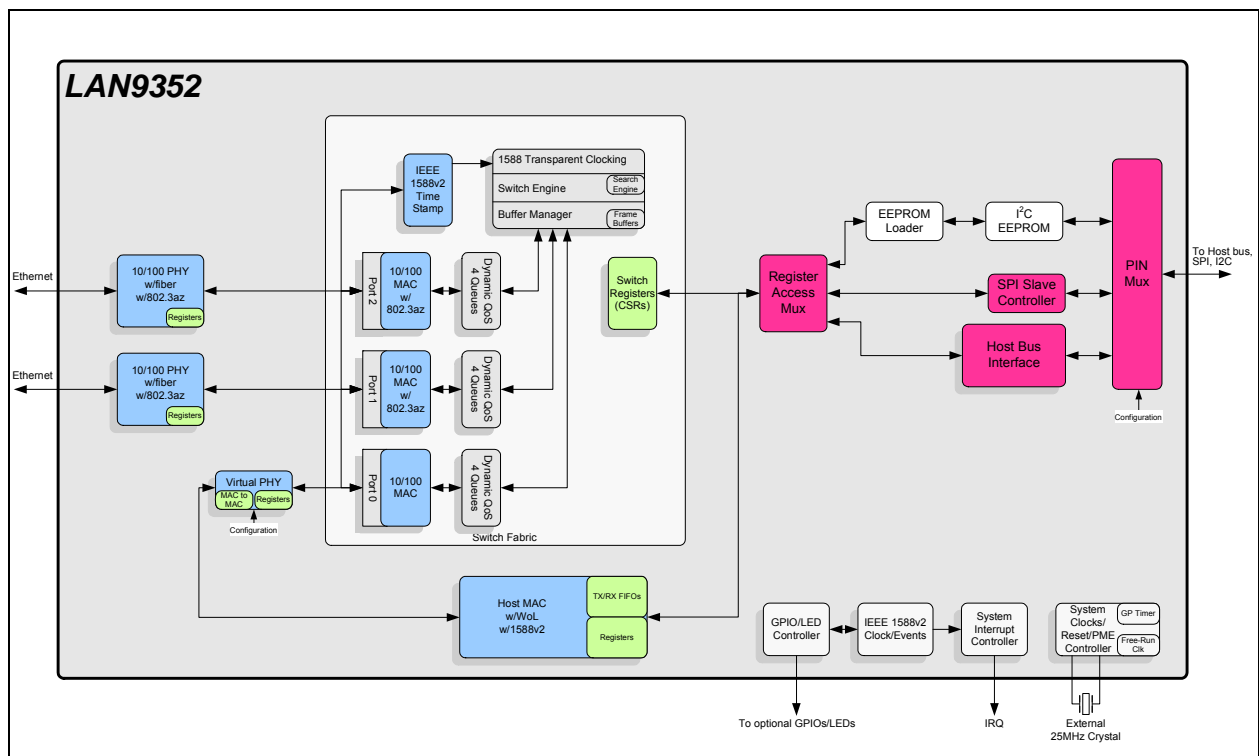
The LAN9352 contains an I²C master EEPROM controller for connection to an optional EEPROM. This allows for the storage and retrieval of static data. The internal EEPROM Loader can be optionally configured to automatically load stored configuration settings from the EEPROM into the device at reset.

In addition to the primary functionality described above, the LAN9352 provides additional features designed for extended functionality. These include a configurable 16-bit General Purpose Timer (GPT), a 32-bit 25MHz free running counter, a configurable GPIO/LED interface, and IEEE 1588 time stamping on all ports and all GPIOs. The IEEE time stamp unit provides a 64-bit tunable clock for accurate PTP timing and a timer comparator to allow time based interrupt generation.

The LAN9352 can be configured to operate via a single 3.3V supply utilizing an integrated 3.3V to 1.2V linear regulator. The linear regulator may be optionally disabled, allowing usage of a high efficiency external regulator for lower system power dissipation.

The LAN9352 is available in commercial and industrial temperature ranges. [Figure 2-1](#) provides an internal block diagram of the LAN9352.

FIGURE 2-1: INTERNAL BLOCK DIAGRAM




LAN9352

3.0 PIN DESCRIPTIONS AND CONFIGURATION

3.1 72-QFN Pin Assignments

FIGURE 3-1: 72-QFN PIN ASSIGNMENTS (TOP VIEW)



Note: When a “#” is used at the end of the signal name, it indicates that the signal is active low. For example, **RST#** indicates that the reset signal is active low.

The buffer type for each signal is indicated in the “Buffer Type” column of the pin description tables in [Section 3.3, “Pin Descriptions”](#). A description of the buffer types is provided in [Section 1.2, “Buffer Types”](#).

Table 3-1 details the 72-QFN package pin assignments in table format. As shown, select pin functions may change based on the device's mode of operation. For modes where a specific pin has no function, the table cell will be marked with "-".

TABLE 3-1: 72-QFN PACKAGE PIN ASSIGNMENTS

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
1		OSCI	
2		OSCO	
3		OSCVDD12	
4		OSCVSS	
5		VDD33	
6		VDDCR	
7		REG_EN	
8		<u>FXLOSEN</u>	
9		FXSDA/FXLOSA/ <u>FXSDENA</u>	
10		FXSDB/FXLOSB/ <u>FXSDENB</u>	
11		RST#	
12		GPIO7	
13		GPIO6	
14	D2	AD2	SIO2
15	D1	AD1	SO/SIO1
16		VDDIO	
17	D14	AD14	-
18	D13	AD13	-
19	D0	AD0	SI/SIO0
20		PME	
21		LED5/GPIO5/ <u>PHYADD</u>	
22		LED4/GPIO4/ <u>1588EN</u>	
23	D9	AD9	SCK
24		VDDIO	
25	D12	AD12	-
26	D11	AD11	-
27	D10	AD10	-
28		VDDCR	
29	A1	ALELO	-
30	A3	<u>MNGT2</u>	-
31	A4	<u>MNGT3</u>	-
32		CS	-

LAN9352

TABLE 3-1: 72-QFN PACKAGE PIN ASSIGNMENTS (CONTINUED)

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
33	A2	ALEHI	-
34	WR/ENB		-
35	RD/RD_WR		-
36	VDDIO		
37	A0/D15	AD15	-
38	LED3/GPIO3/EEEEEN		
39	FIFOSEL	-	-
40	D3	AD3	SIO3
41	D6	AD6	-
42	VDDIO		
43	VDDCR		
44	D7	AD7	-
45	D8	AD8	-
46	TESTMODE		
47	EESDA/TMS		
48	EESCL/TCK		
49	IRQ		
50	LED2/GPIO2/E2PSIZE		
51	LED1/GPIO1/TDI/MNGT1		
52	VDDIO		
53	LED0/GPIO0/TDO/MNGT0		
54	NC		
55	D4	AD4	-
56	D5	AD5	SCS#
57	NC		
58	NC		
59	VDD33TXRX1		
60	TXNA		
61	TXPA		
62	RXNA		
63	RXPA		
64	VDD12TX1		
65	RBIAS		
66	VDD33BIAS		
67	VDD12TX2		


TABLE 3-1: 72-QFN PACKAGE PIN ASSIGNMENTS (CONTINUED)

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
68	RXPB		
69	RXNB		
70	TXPB		
71	TXNB		
72	VDD33TXRX2		
Exposed Pad	VSS		

LAN9352

3.2 80-TQFP-EP Pin Assignments

FIGURE 3-2: 80-TQFP-EP PIN ASSIGNMENTS (TOP VIEW)



Note: When a “#” is used at the end of the signal name, it indicates that the signal is active low. For example, RST# indicates that the reset signal is active low.

The buffer type for each signal is indicated in the “Buffer Type” column of the pin description tables in [Section 3.3, "Pin Descriptions"](#). A description of the buffer types is provided in [Section 1.2, "Buffer Types"](#).

Table 3-2 details the 80-TQFP-EP package pin assignments in table format. As shown, select pin functions may change based on the device's mode of operation. For modes where a specific pin has no function, the table cell will be marked with "-".

TABLE 3-2: 80-TQFP-EP PACKAGE PIN ASSIGNMENTS

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
1	OSCI		
2	OSCO		
3	OSCVDD12		
4	OSCVSS		
5	VDD33		
6	VDDCR		
7	REG_EN		
8	<u>FXLOSEN</u>		
9	FXSDA/FXLOSA/ <u>FXSDENA</u>		
10	FXSDB/FXLOSB/ <u>FXSDENB</u>		
11	RST#		
12	GPIO7		
13	GPIO6		
14	D2	AD2	SIO2
15	D1	AD1	SO/SIO1
16	VDDIO		
17	NC		
18	NC		
19	D14	AD14	-
20	D13	AD13	-
21	D0	AD0	SI/SIO0
22	PME		
23	NC		
24	LED5/GPIO5/ <u>PHYADD</u>		
25	LED4/GPIO4/ <u>1588EN</u>		
26	D9	AD9	SCK
27	VDDIO		
28	D12	AD12	-
29	D11	AD11	-
30	D10	AD10	-
31	VDDCR		
32	A1	ALELO	-

LAN9352

TABLE 3-2: 80-TQFP-EP PACKAGE PIN ASSIGNMENTS (CONTINUED)

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
33	A3	<u>MNGT2</u>	-
34	A4	<u>MNGT3</u>	-
35	CS		-
36	A2	ALEHI	-
37	WR/ENB		-
38	RD/RD_WR		-
39	VDDIO		
40	NC		
41	A0/D15	AD15	-
42	LED3/GPIO3/ <u>EEEEEN</u>		
43	FIFOSEL	-	-
44	NC		
45	NC		
46	NC		
47	NC		
48	D3	AD3	SIO3
49	D6	AD6	-
50	VDDIO		
51	VDDCR		
52	D7	AD7	-
53	D8	AD8	-
54	TESTMODE		
55	EESDA/TMS		
56	EESCL/TCK		
57	IRQ		
58	LED2/GPIO2/ <u>E2PSIZE</u>		
59	LED1/GPIO1/TDI/ <u>MNGT1</u>		
60	VDDIO		
61	LED0/GPIO0/TDO/ <u>MNGT0</u>		
62	D4	AD4	-
63	D5	AD5	SCS#
64	NC		
65	NC		
66	NC		
67	VDD33TXRX1		

TABLE 3-2: 80-TQFP-EP PACKAGE PIN ASSIGNMENTS (CONTINUED)

Pin Number	HBI Indexed Mode Pin Name	HBI Multiplexed Mode Pin Name	SPI Mode Pin Name
68		TXNA	
69		TXPA	
70		RXNA	
71		RXPA	
72		VDD12TX1	
73		RBIAS	
74		VDD33BIAS	
75		VDD12TX2	
76		RXPB	
77		RXNB	
78		TXPB	
79		TXNB	
80		VDD33TXRX2	
Exposed Pad		VSS	

3.3 Pin Descriptions

This section contains descriptions of the various LAN9352 pins. The pin descriptions have been broken into functional groups as follows:

- [LAN Port A Pin Descriptions](#)
- [LAN Port B Pin Descriptions](#)
- [LAN Port A & B Power and Common Pin Descriptions](#)
- [Host Bus Pin Descriptions](#)
- [SPI/SQI Pin Descriptions](#)
- [EEPROM Pin Descriptions](#)
- [GPIO, LED & Configuration Strap Pin Descriptions](#)
- [Miscellaneous Pin Descriptions](#)
- [JTAG Pin Descriptions](#)
- [Core and I/O Power Pin Descriptions](#)

TABLE 3-3: LAN PORT A PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Port A TP TX/RX Positive Channel 1	TXPA	AIO	Port A Twisted Pair Transmit/Receive Positive Channel 1. See Note 1
	Port A FX TX Positive		OLVPECL	Port A Fiber Transmit Positive.
1	Port A TP TX/RX Negative Channel 1	TXNA	AIO	Port A Twisted Pair Transmit/Receive Negative Channel 1. See Note 1 .
	Port A FX TX Negative		OLVPECL	Port A Fiber Transmit Negative.
1	Port A TP TX/RX Positive Channel 2	RXPA	AIO	Port A Twisted Pair Transmit/Receive Positive Channel 2. See Note 1 .
	Port A FX RX Positive		AI	Port A Fiber Receive Positive.
1	Port A TP TX/RX Negative Channel 2	RXNA	AIO	Port A Twisted Pair Transmit/Receive Negative Channel 2. See Note 1 .
	Port A FX RX Negative		AI	Port A Fiber Receive Negative.

TABLE 3-3: LAN PORT A PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	Port A FX Signal Detect (SD)	FXSDA	ILVPECL	Port A Fiber Signal Detect. When FX-LOS mode is not selected, this pin functions as the Signal Detect input from the external transceiver. A level above 2 V (typ.) indicates valid signal. When FX-LOS mode is selected, the input buffer is disabled.
	Port A FX Loss Of Signal (LOS)	FXLOSA	IS (PU)	Port A Fiber Loss of Signal. When FX-LOS mode is selected (via fx_los_strap_1), this pin functions as the Loss of Signal input from the external transceiver. A high indicates LOS while a low indicates valid signal. When FX-LOS mode is not selected, the input buffer and pull-up are disabled.
	Port A FX-SD Enable Strap	<u>FXSDENA</u>	AI	Port A FX-SD Enable. When FX-LOS mode is not selected, this strap input selects between FX-SD and copper twisted pair mode. A level above 1 V (typ.) selects FX-SD. When FX-LOS mode is selected, the input buffer is disabled. See Note 2 .

Note 1: In copper mode, either channel 1 or 2 may function as the transmit pair while the other channel functions as the receive pair. The pin name symbols for the twisted pair pins apply to a normal connection. If HP Auto-MDIX is enabled and a reverse connection is detected or manually selected, the RX and TX pins will be swapped internally.

Note 2: Configuration strap pins are identified by an underlined symbol name. Configuration strap values are latched on power-on reset or **RST#** de-assertion. Refer to [Section 7.0, "Configuration Straps,"](#) on page 60 for more information.

Note: Port A is connected to the Switch Fabric port 1.

TABLE 3-4: LAN PORT B PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Port B TP TX/RX Positive Channel 1	TXPB	AIO	Port B Twisted Pair Transmit/Receive Positive Channel 1. See Note 3
	Port B FX TX Positive		OLVPECL	Port B Fiber Transmit Positive.
1	Port B TP TX/RX Negative Channel 1	TXNB	AIO	Port B Twisted Pair Transmit/Receive Negative Channel 1. See Note 3 .
	Port B FX TX Negative		OLVPECL	Port B Fiber Transmit Negative.

TABLE 3-4: LAN PORT B PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	Port BTP TX/RX Positive Channel 2	RXPB	AIO	Port B Twisted Pair Transmit/Receive Positive Channel 2. See Note 3 .
	Port B FX RX Positive		AI	Port B Fiber Receive Positive.
1	Port B TP TX/RX Negative Channel 2	RXNB	AIO	Port B Twisted Pair Transmit/Receive Negative Channel 2. See Note 3 .
	Port B FX RX Negative		AI	Port B Fiber Receive Negative.
1	Port B FX Signal Detect (SD)	FXSDB	ILVPECL	Port B Fiber Signal Detect. When FX-LOS mode is not selected, this pin functions as the Signal Detect input from the external transceiver. A level above 2 V (typ.) indicates valid signal. When FX-LOS mode is selected, the input buffer is disabled.
	Port B FX Loss Of Signal (LOS)	FXLOSB	IS (PU)	Port B Fiber Loss of Signal. When FX-LOS mode is selected (via fx_los_strap_2), this pin functions as the Loss of Signal input from the external transceiver. A high indicates LOS while a low indicates valid signal. When FX-LOS mode is not selected, the input buffer and pull-up are disabled.
	Port B FX-SD Enable Strap	<u>FXSDENB</u>	AI	Port B FX-SD Enable. When FX-LOS mode is not selected, this strap input selects between FX-SD and copper twisted pair mode. A level above 1 V (typ.) selects FX-SD. When FX-LOS mode is selected, the input buffer is disabled. See Note 4 .

Note 3: In copper mode, either channel 1 or 2 may function as the transmit pair while the other channel functions as the receive pair. The pin name symbols for the twisted pair pins apply to a normal connection. If HP Auto-MDIX is enabled and a reverse connection is detected or manually selected, the RX and TX pins will be swapped internally.

Note 4: Configuration strap pins are identified by an underlined symbol name. Configuration strap values are latched on power-on reset or **RST#** de-assertion. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information.

Note: Port B is connected to Switch Fabric port 2.

TABLE 3-5: LAN PORT A & B POWER AND COMMON PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Bias Reference	RBIAS	AI	Used for internal bias circuits. Connect to an external 12.1 k Ω , 1% resistor to ground. Refer to the device reference schematic for connection information. Note: The nominal voltage is 1.2 V and the resistor will dissipate approximately 1 mW of power.
1	Port A and B FX-LOS Enable Strap	<u>FXLOSEN</u>	AI	Port A and B FX-LOS Enable. This 3 level strap input selects between FX-LOS and FX-SD / copper twisted pair mode. A level below 1 V (typ.) selects FX-SD / copper twisted pair for ports A and B, further determined by FXSDENA and FXSDENB . A level of 1.5 V selects FX-LOS for port A and FX-SD / copper twisted pair for port B, further determined by FXSDENB . A level above 2 V (typ.) selects FX-LOS for ports A and B.
1	+3.3 V Port A Analog Power Supply	VDD33TXRX1	P	See Note 5 .
1	+3.3 V Port B Analog Power Supply	VDD33TXRX2	P	See Note 5 .
1	+3.3 V Master Bias Power Supply	VDD33BIAS	P	See Note 5 .
1	Port A Transmitter +1.2 V Power Supply	VDD12TX1	P	This pin is supplied from either an external 1.2 V supply or from the device's internal regulator via the PCB. This pin must be tied to the VDD12TX2 pin for proper operation. See Note 5 .
1	Port B Transmitter +1.2 V Power Supply	VDD12TX2	P	This pin is supplied from either an external 1.2 V supply or from the device's internal regulator via the PCB. This pin must be tied to the VDD12TX1 pin for proper operation. See Note 5 .

Note 5: Refer to [Section 4.0, "Power Connections,"](#) on page 30, the device reference schematics, and the device LANCheck schematic checklist for additional connection information.

TABLE 3-6: HOST BUS PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Read	RD	VIS	This pin is the host bus read strobe. Normally active low, the polarity can be changed via the HBI_rd_rdw_r_polarity_strap .
	Read or Write	RD_WR	VIS	This pin is the host bus direction control. Used in conjunction with the ENB pin, it indicates a read or write operation. The normal polarity is read when 1, write when 0 (R/nW) but can be changed via the HBI_rd_rdw_r_polarity_strap .
1	Write	WR	VIS	This pin is the host bus write strobe. Normally active low, the polarity can be changed via the HBI_wr_en_polarity_strap .
	Enable	ENB	VIS	This pin is the host bus data enable strobe. Used in conjunction with the RD_WR pin it indicates the data phase of the operation. Normally active low, the polarity can be changed via the HBI_wr_en_polarity_strap .
1	Chip Select	CS	VIS	This pin is the host bus chip select and indicates that the device is selected for the current transfer. Normally active low, the polarity can be changed via the HBI_cs_polarity_strap .
1	FIFO Select	FIFOSEL	VIS	This input directly selects the Host MAC TX and RX Data FIFOs for non-multiplexed address mode.
5	Address	A[4:0]	VIS	These pins provide the address for non-multiplexed address mode. In 16-bit data mode, bit 0 is not used.
16	Data	D[15:0]	VIS/VO8	These pins are the host bus data bus for non-multiplexed address mode. In 8-bit data mode, bits 15-8 are not used and their input and output drivers are disabled.
	Address & Data	AD[15:0]	VIS/VO8	These pins are the host bus address / data bus for multiplexed address mode. Bits 15-8 provide the upper byte of address for single phase multiplexed address mode. Bits 7-0 provide the lower byte of address for single phase multiplexed address mode and both bytes of address for dual phase multiplexed address mode. In 8-bit data dual phase multiplexed address mode, bits 15-8 are not used and their input and output drivers are disabled.

TABLE 3-6: HOST BUS PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	Address Latch Enable High	ALEHI	VIS	<p>This pin indicates the address phase for multiplexed address modes. It is used to load the higher address byte in dual phase multiplexed address mode.</p> <p>Normally active low (address saved on rising edge), the polarity can be changed via the HBI_ale_polarity_strap.</p>
1	Address Latch Enable Low	ALELO	VIS	<p>This pin indicates the address phase for multiplexed address modes. It is used to load both address bytes in single phase multiplexed address mode and the lower address byte in dual phase multiplexed address mode.</p> <p>Normally active low (address saved on rising edge), the polarity can be changed via the HBI_ale_polarity_strap.</p>

TABLE 3-7: SPI/SQI PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	SPI/SQI Slave Chip Select	SCS#	VIS (PU)	This pin is the SPI/SQI slave chip select input. When low, the SPI/SQI slave is selected for SPI/SQI transfers. When high, the SPI/SQI serial data output(s) is(are) 3-stated.
1	SPI/SQI Slave Serial Clock	SCK	VIS (PU)	This pin is the SPI/SQI slave serial clock input.
4	SPI/SQI Slave Serial Data Input/Output	SIO[3:0]	VIS/VO8 (PU)	These pins are the SPI/SQI slave data input and output for multiple bit I/O.
	SPI Slave Serial Data Input	SI	VIS (PU)	This pin is the SPI slave serial data input. SI is shared with the SIO0 pin.
	SPI Slave Serial Data Output	SO	VO8 (PU) Note 6	This pin is the SPI slave serial data output. SO is shared with the SIO1 pin.

Note 6: Although this pin is an output for SPI instructions, it includes a pull-up since it is also SIO bit 1.

TABLE 3-8: EEPROM PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	EEPROM I ² C Serial Data Input/Output	EESDA	VIS/VOD8	When the device is accessing an external EEPROM this pin is the I ² C serial data input/open-drain output. Note: This pin must be pulled-up by an external resistor at all times.
1	EEPROM I ² C Serial Clock	EESCL	VIS/VOD8	When the device is accessing an external EEPROM this pin is the I ² C clock input/open-drain output. Note: This pin must be pulled-up by an external resistor at all times.

TABLE 3-9: GPIO, LED & CONFIGURATION STRAP PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	General Purpose I/O 7	GPIO7	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
1	General Purpose I/O 6	GPIO6	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .

TABLE 3-9: GPIO, LED & CONFIGURATION STRAP PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	LED 5	LED5	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 5 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the PHYADD strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 5	GPIO5	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 5 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	PHY Address Configuration Strap	<u>PHYADD</u>	VIS (PU)	This strap configures the default value of the Switch PHY Address Select soft-strap. See Note 7 .
1	LED 4	LED4	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 4 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the 1588EN strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 4	GPIO4	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 4 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	1588 Enable Configuration Strap	<u>1588EN</u>	VIS (PU)	This strap configures the default value of the 1588 Enable soft-strap. See Note 7 .

TABLE 3-9: GPIO, LED & CONFIGURATION STRAP PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	LED 3	LED3	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 3 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the EEEEEN strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 3	GPIO3	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 3 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	Energy Efficient Ethernet Enable Configuration Strap	EEEEEN	VIS (PU)	This strap configures the default value of the EEE Enable 2-1 soft-straps. See Note 7 .
1	LED 2	LED2	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 2 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the E2PSIZE strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 2	GPIO2	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 2 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	EEPROM Size Configuration Strap	E2PSIZE	VIS (PU)	This strap configures the value of the EEPROM size hard-strap. See Note 7 . A low selects 1K bits (128 x 8) through 16K bits (2K x 8). A high selects 32K bits (4K x 8) through 512K bits (64K x 8).

TABLE 3-9: GPIO, LED & CONFIGURATION STRAP PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	LED 1	LED1	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 1 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the MNGT1 strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 1	GPIO1	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 1 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	Host Interface Configuration Strap 1	<u>MNGT1</u>	VIS (PU)	This strap, along with <u>MNGT0</u> , <u>MNGT2</u> , and <u>MNGT3</u> configures the host interface mode. See Note 7 . See Table 7-3, "HBI Strap Mapping," on page 72 for the host interface strap settings.
1	LED 0	LED0	VO12/ VOD12/ VOS12	This pin is configured to operate as an LED when the LED 0 Enable bit of the LED Configuration Register (LED_CFG) is set. The buffer type depends on the setting of the LED Function 2-0 (LED_FUN[2:0]) field in the LED Configuration Register (LED_CFG) and is configured to be either a push-pull or open-drain/open-source output. When selected as an open-drain/open-source output, the polarity of this pin depends upon the MNGT0 strap value sampled at reset. Note: Refer to Section 17.3, "LED Operation," on page 564 to additional information.
	General Purpose I/O 0	GPIO0	VIS/VO12/ VOD12 (PU)	This pin is configured to operate as a GPIO when the LED 0 Enable bit of the LED Configuration Register (LED_CFG) is clear. The pin is fully programmable as either a push-pull output, an open-drain output or a Schmitt-triggered input by writing the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) .
	Host Interface Configuration Strap 0	<u>MNGT0</u>	VIS (PU)	This strap, along with <u>MNGT1</u> , <u>MNGT2</u> , and <u>MNGT3</u> configures the host mode. See Note 7 . See Table 7-3, "HBI Strap Mapping," on page 72 for the host interface strap settings.

TABLE 3-9: GPIO, LED & CONFIGURATION STRAP PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
1	Host Interface Configuration Strap 3	<u>MNGT3</u>	VIS (PU)	This strap, along with <u>MNGT0</u> , <u>MNGT1</u> , and <u>MNGT2</u> configures the host mode. See Note 7 . See Table 7-3, "HBI Strap Mapping," on page 72 for the host interface strap settings.
1	Host Interface Configuration Strap 2	<u>MNGT2</u>	VIS (PU)	This strap, along with <u>MNGT0</u> , <u>MNGT1</u> , and <u>MNGT3</u> configures the host mode. See Note 7 . See Table 7-3, "HBI Strap Mapping," on page 72 for the host interface strap settings.

Note 7: Configuration strap pins are identified by an underlined symbol name. Configuration strap values are latched on power-on reset or RST# de-assertion. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information.

TABLE 3-10: MISCELLANEOUS PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Power Management Event Output	PME	VO8/VOD8	When programmed accordingly this signal is asserted upon detection of a wakeup event. The polarity and buffer type of this signal is programmable via the PME Enable (PME_EN) bit of the Power Management Control Register (PMT_CTRL) . Refer to Section 6.0, "Clocks, Resets, and Power Management," on page 41 for additional information on the power management features.
1	Interrupt Output	IRQ	VO8/VOD8	Interrupt request output. The polarity, source and buffer type of this signal is programmable via the Interrupt Configuration Register (IRQ_CFG) . For more information, refer to Section 8.0, "System Interrupts," on page 73 .
1	System Reset Input	RST#	VIS (PU)	As an input, this active low signal allows external hardware to reset the device. The device also contains an internal power-on reset circuit. Thus this signal may be left unconnected if an external hardware reset is not needed. When used this signal must adhere to the reset timing requirements as detailed in the Section 20.0, "Operational Characteristics," on page 579 .
1	Regulator Enable	REG_EN	AI	When tied to 3.3 V, the internal 1.2 V regulators are enabled.
1	Test Mode	TESTMODE	VIS (PD)	This pin must be tied to VSS for proper operation.
1	Crystal Input	OSCI	ICLK	External 25 MHz crystal input. This signal can also be driven by a single-ended clock oscillator. When this method is used, OSC0 should be left unconnected.
1	Crystal Output	OSCO	OCLK	External 25 MHz crystal output.
1	Crystal +1.2 V Power Supply	OSCVDD12	P	Supplied by the on-chip regulator unless configured for regulator off mode via REG_EN .
1	Crystal Ground	OSCVSS	P	Crystal ground.

TABLE 3-10: MISCELLANEOUS PIN DESCRIPTIONS (CONTINUED)

Num Pins	Name	Symbol	Buffer Type	Description
Note 8	No Connect	NC	-	This pin must be left unconnected for proper operation.

Note 8: 3 NC pins for the QFN package, 11 NC pins for the TQFP-EP package.

TABLE 3-11: JTAG PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	JTAG Test Mux Select	TMS	VIS	JTAG test mode select
1	JTAG Test Clock	TCK	VIS	JTAG test clock
1	JTAG Test Data Input	TDI	VIS	JTAG data input
1	JTAG Test Data Output	TDO	VO12	JTAG data output

TABLE 3-12: CORE AND I/O POWER PIN DESCRIPTIONS

Num Pins	Name	Symbol	Buffer Type	Description
1	Regulator +3.3 V Power Supply	VDD33	P	+3.3 V power supply for internal regulators. See Note 9 . Note: +3.3 V must be supplied to this pin even if the internal regulators are disabled.
5	+1.8 V to +3.3 V Variable I/O Power	VDDIO	P	+1.8 V to +3.3 V variable I/O power. See Note 9 .
3	+1.2 V Digital Core Power Supply	VDDCR	P	Supplied by the on-chip regulator unless configured for regulator off mode via REG_EN . 1 μ F and 470 pF decoupling capacitors in parallel to ground should be used on pin 6. See Note 9 .
1 pad	Ground	VSS	P	Common ground. This exposed pad must be connected to the ground plane with a via array.

Note 9: Refer to [Section 4.0, "Power Connections," on page 30](#), the device reference schematic, and the device LANCheck schematic checklist for additional connection information.

4.0 POWER CONNECTIONS

Figure 4-1 and Figure 4-2 illustrate the device power connections for regulator enabled and disabled cases, respectively. Refer to the device reference schematic and the device LANCheck schematic checklist for additional information. Section 4.1 provides additional information on the devices internal voltage regulators.

FIGURE 4-1: POWER CONNECTIONS - REGULATORS ENABLED





FIGURE 4-2: POWER CONNECTIONS - REGULATORS DISABLED



4.1 Internal Voltage Regulators

The device contains two internal 1.2 V regulators:

- [1.2 V Core Regulator](#)
- [1.2 V Crystal Oscillator Regulator](#)

4.1.1 1.2 V CORE REGULATOR

The core regulator supplies 1.2 V volts to the main core digital logic, the I/O pads, and the PHYs' digital logic and can be used to supply the 1.2 V power to the PHY analog sections (via an external connection).

When the **REG_EN** input pin is connected to 3.3 V, the core regulator is enabled and receives 3.3 V on the **VDD33** pin. A 1.0 uF 0.1 Ω ESR capacitor must be connected to the **VDDCR** pin associated with the regulator.

When the **REG_EN** input pin is connected to **VSS**, the core regulator is disabled. However, 3.3 V must still be supplied to the **VDD33** pin. The 1.2 V core voltage must then be externally input into the **VDDCR** pins.

4.1.2 1.2 V CRYSTAL OSCILLATOR REGULATOR

The crystal oscillator regulator supplies 1.2 V volts to the crystal oscillator. When the **REG_EN** input pin is connected to 3.3 V, the crystal oscillator regulator is enabled and receives 3.3 V on the **VDD33** pin. An external capacitor is not required.

When the **REG_EN** input pin is connected to **VSS**, the crystal oscillator regulator is disabled. However, 3.3 V must still be supplied to the **VDD33** pin. The 1.2 V crystal oscillator voltage must then be externally input into the **OSCVDD12** pin.

5.0 REGISTER MAP

This chapter details the device register map and summarizes the various directly addressable System Control and Status Registers (CSRs). Detailed descriptions of the System CSRs are provided in the chapters corresponding to their function. Additional indirectly addressable registers are available in the various sub-blocks of the device. These registers are also detailed in their corresponding chapters.

Directly Addressable Registers

- [Section 11.10.1, "TX/RX FIFOs," on page 170](#)
- [Section 5.1, "System Control and Status Registers," on page 35](#)

Indirectly Addressable Registers


- [Section 11.14, "Host MAC Control and Status Registers," on page 202](#)
- [Section 12.2.19, "Physical PHY Registers," on page 242](#)
- [Section 13.7, "Switch Fabric Control and Status Registers," on page 363](#)

[Figure 5-1](#) contains an overall base register memory map of the device. This memory map is not drawn to scale, and should be used for general reference only. [Table 5-1](#) provides a summary of all directly addressable CSRs and their corresponding addresses.

Note: Register bit type definitions are provided in [Section 1.3, "Register Nomenclature," on page 7](#).

Not all device registers are memory mapped or directly addressable. For details on the accessibility of the various device registers, refer the register sub-sections listed above.

FIGURE 5-1: REGISTER ADDRESS MAP



5.1 System Control and Status Registers

The System CSRs are directly addressable memory mapped registers with a base address offset range of 050h to 2F8h. These registers are addressable by the Host via the Host Bus Interface (HBI) or SPI/SQI. For more information on the various device modes and their corresponding address configurations, see [Section 2.0, "General Description," on page 8](#).

[Table 5-1](#) lists the System CSRs and their corresponding addresses in order. All system CSRs are reset to their default value on the assertion of a chip-level reset.

The System CSRs can be divided into the following sub-categories. Each of these sub-categories is located in the corresponding chapter and contains the System CSR descriptions of the associated registers. The register descriptions are categorized as follows:

- [Section 6.2.3, "Reset Registers," on page 47](#)
- [Section 6.3.5, "Power Management Registers," on page 54](#)
- [Section 8.3, "Interrupt Registers," on page 77](#)
- [Section 11.13, "Host MAC & FIFO Interface Registers," on page 188](#)
- [Section 17.4, "GPIO/LED Registers," on page 566](#)
- [Section 14.5, "I2C Master EEPROM Controller Registers," on page 469](#)
- [Section 15.8, "1588 Registers," on page 494](#)
- [Section 13.6, "Switch Fabric Interface Logic Registers," on page 348](#)
- [Section 12.3.3, "Virtual PHY Registers," on page 305](#)
- [Section 18.1, "Miscellaneous System Configuration & Status Registers," on page 572](#)

Note: Unlisted registers are reserved for future use.

TABLE 5-1: SYSTEM CONTROL AND STATUS REGISTERS

Address	Register Name (Symbol)
000h - 04Ch	TX/RX FIFOs
050h	Chip ID and Revision (ID_REV)
054h	Interrupt Configuration Register (IRQ_CFG)
058h	Interrupt Status Register (INT_STS)
05Ch	Interrupt Enable Register (INT_EN)
064h	Byte Order Test Register (BYTE_TEST)
068h	FIFO Level Interrupt Register (FIFO_INT)
06Ch	Receive Configuration Register (RX_CFG)
070h	Transmit Configuration Register (TX_CFG)
074h	Hardware Configuration Register (HW_CFG)
078h	Receive Datapath Control Register (RX_DP_CTRL)
07Ch	RX FIFO Information Register (RX_FIFO_INF)
080h	TX FIFO Information Register (TX_FIFO_INF)
084h	Power Management Control Register (PMT_CTRL)
08Ch	General Purpose Timer Configuration Register (GPT_CFG)
090h	General Purpose Timer Count Register (GPT_CNT)
09Ch	Free Running 25MHz Counter Register (FREE_RUN)
0A0h	Host MAC RX Dropped Frames Counter Register (RX_DROP)
0A4h	Host MAC CSR Interface Command Register (MAC_CSR_CMD)
0A8h	Host MAC CSR Interface Data Register (MAC_CSR_DATA)
0ACh	Host MAC Automatic Flow Control Configuration Register (AFC_CFG)
1588 Registers	
100h	1588 Command and Control Register (1588_CMD_CTL)

TABLE 5-1: SYSTEM CONTROL AND STATUS REGISTERS (CONTINUED)

Address	Register Name (Symbol)
104h	1588 General Configuration Register (1588_GENERAL_CONFIG)
108h	1588 Interrupt Status Register (1588_INT_STS)
10Ch	1588 Interrupt Enable Register (1588_INT_EN)
110h	1588 Clock Seconds Register (1588_CLOCK_SEC)
114h	1588 Clock NanoSeconds Register (1588_CLOCK_NS)
118h	1588 Clock Sub-NanoSeconds Register (1588_CLOCK_SUBNS)
11Ch	1588 Clock Rate Adjustment Register (1588_CLOCK_RATE_ADJ)
120h	1588 Clock Temporary Rate Adjustment Register (1588_CLOCK_TEMP_RATE_ADJ)
124h	1588 Clock Temporary Rate Duration Register (1588_CLOCK_TEMP_RATE_DURATION)
128h	1588 Clock Step Adjustment Register (1588_CLOCK_STEP_ADJ)
12Ch	1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) x=A
130h	1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) x=A
134h	1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) x=A
138h	1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=A
13Ch	1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) x=B
140h	1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) x=B
144h	1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) x=B
148h	1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=B
14Ch	1588 User MAC Address High-WORD Register (1588_USER_MAC_HI)
150h	1588 User MAC Address Low-DWORD Register (1588_USER_MAC_LO)
154h	1588 Bank Port GPIO Select Register (1588_BANK_PORT_GPIO_SEL)
158h	1588 Port x Latency Register (1588_LATENCY_x)
158h	1588 Port x RX Parsing Configuration Register (1588_RX_PARSE_CONFIG_x)
158h	1588 Port x TX Parsing Configuration Register (1588_TX_PARSE_CONFIG_x)
15Ch	1588 Port x Asymmetry and Peer Delay Register (1588_ASYM_PEERDLY_x)
15Ch	1588 Port x RX Timestamp Configuration Register (1588_RX_TIMESTAMP_CONFIG_x)
15Ch	1588 Port x TX Timestamp Configuration Register (1588_TX_TIMESTAMP_CONFIG_x)
15Ch	1588 GPIO Capture Configuration Register (1588_GPIO_CAP_CONFIG)
160h	1588 Port x Capture Information Register (1588_CAP_INFO_x)
160h	1588 Port x RX Timestamp Insertion Configuration Register (1588_RX_TS_INSERT_CONFIG_x)
164h	1588 Port x RX Correction Field Modification Register (1588_RX_CF_MOD_x)
164h	1588 Port x TX Modification Register (1588_TX_MOD_x)
168h	1588 Port x RX Filter Configuration Register (1588_RX_FILTER_CONFIG_x)
168h	1588 Port x TX Modification Register 2 (1588_TX_MOD2_x)
16Ch	1588 Port x RX Ingress Time Seconds Register (1588_RX_INGRESS_SEC_x)
16Ch	1588 Port x TX Egress Time Seconds Register (1588_TX_EGRESS_SEC_x)
16Ch	1588 GPIO x Rising Edge Clock Seconds Capture Register (1588_GPIO_RE_CLOCK_SEC_CAP_x)
170h	1588 Port x RX Ingress Time NanoSeconds Register (1588_RX_INGRESS_NS_x)
170h	1588 Port x TX Egress Time NanoSeconds Register (1588_TX_EGRESS_NS_x)
170h	1588 GPIO x Rising Edge Clock NanoSeconds Capture Register (1588_GPIO_RE_CLOCK_NS_CAP_x)

TABLE 5-1: SYSTEM CONTROL AND STATUS REGISTERS (CONTINUED)

Address	Register Name (Symbol)
174h	1588 Port x RX Message Header Register (1588_RX_MSG_HEADER_x)
174h	1588 Port x TX Message Header Register (1588_TX_MSG_HEADER_x)
178h	1588 Port x RX Pdelay_Req Ingress Time Seconds Register (1588_RX_PDREQ_SEC_x)
178h	1588 Port x TX Delay_Req Egress Time Seconds Register (1588_TX_DREQ_SEC_x)
178h	1588 GPIO x Falling Edge Clock Seconds Capture Register (1588_GPIO_FE_CLOCK_SEC_CAP_x)
17Ch	1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register (1588_RX_PDREQ_NS_x)
17Ch	1588 Port x TX Delay_Req Egress Time NanoSeconds Register (1588_TX_DREQ_NS_x)
17Ch	1588 GPIO x Falling Edge Clock NanoSeconds Capture Register (1588_GPIO_FE_CLOCK_NS_CAP_x)
180h	1588 Port x RX Pdelay_Req Ingress Correction Field High Register (1588_RX_PDREQ_CF_HI_x)
180h	1588 TX One-Step Sync Upper Seconds Register (1588_TX_ONE_STEP_SYNC_SEC)
184h	1588 Port x RX Pdelay_Req Ingress Correction Field Low Register (1588_RX_PDREQ_CF_LOW_x)
188h	1588 Port x RX Checksum Dropped Count Register (1588_RX_CHKSUM_DROPPED_CNT_x)
18Ch	1588 Port x RX Filtered Count Register (1588_RX_FILTERED_CNT_x)
Switch Registers	
1A0h	Port 1 Manual Flow Control Register (MANUAL_FC_1)
1A4h	Port 2 Manual Flow Control Register (MANUAL_FC_2)
1A8h	Port 0 Manual Flow Control Register (MANUAL_FC_0)
1ACh	Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA)
1B0h	Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD)
EEPROM/LED Registers	
1B4h	EEPROM Command Register (E2P_CMD)
1B8h	EEPROM Data Register (E2P_DATA)
1BCh	LED Configuration Register (LED_CFG)
Virtual PHY Registers	
1C0h	Virtual PHY Basic Control Register (VPHY_BASIC_CTRL)
1C4h	Virtual PHY Basic Status Register (VPHY_BASIC_STATUS)
1C8h	Virtual PHY Identification MSB Register (VPHY_ID_MSB)
1CCh	Virtual PHY Identification LSB Register (VPHY_ID_LSB)
1D0h	Virtual PHY Auto-Negotiation Advertisement Register (VPHY_AN_ADV)
1D4h	Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register (VPHY_AN_LP_BASE_ABILITY)
1D8h	Virtual PHY Auto-Negotiation Expansion Register (VPHY_AN_EXP)
1DCh	Virtual PHY Special Control/Status Register (VPHY_SPECIAL_CONTROL_STATUS)
GPIO Registers	
1E0h	General Purpose I/O Configuration Register (GPIO_CFG)
1E4h	General Purpose I/O Data & Direction Register (GPIO_DATA_DIR)
1E8h	General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)
Switch Fabric MAC Address Registers	
1F0h	Switch Fabric MAC Address High Register (SWITCH_MAC_ADDRH)
1F4h	Switch Fabric MAC Address Low Register (SWITCH_MAC_ADDRL)
Reset Register	
1F8h	Reset Control Register (RESET_CTL)
Switch Fabric CSR Interface Direct Data Registers	
200h-2F8h	Switch Fabric CSR Interface Direct Data Registers (SWITCH_CSR_DIRECT_DATA)

5.2 Special Restrictions on Back-to-Back Cycles

5.2.1 BACK-TO-BACK WRITE-READ CYCLES

It is important to note that there are specific restrictions on the timing of back-to-back host write-read operations. These restrictions concern reading registers after any write cycle that may affect the register. In all cases there is a delay between writing to a register and the new value becoming available to be read. In other cases, there is a delay between writing to a register and the subsequent side effect on other registers.

In order to prevent the host from reading stale data after a write operation, minimum wait periods have been established. These periods are specified in [Table 5-2](#). The host processor is required to wait the specified period of time after writing to the indicated register before reading the resource specified in the table. Note that the required wait period is dependent upon the register being read after the write.

Performing “dummy” reads of the [Byte Order Test Register \(BYTE_TEST\)](#) register is a convenient way to guarantee that the minimum write-to-read timing restriction is met. [Table 5-2](#) shows the number of dummy reads that are required before reading the register indicated. The number of BYTE_TEST reads in this table is based on the minimum cycle timing of 45ns. For microprocessors with slower busses the number of reads may be reduced as long as the total time is equal to, or greater than the time specified in the table. Note that dummy reads of the BYTE_TEST register are not required as long as the minimum time period is met.

Note that depending on the host interface mode in use, the basic host interface cycle may naturally provide sufficient time between writes and read. It is required of the system design and register access mechanisms to ensure the proper timing. For example, a write and read to the same register may occur faster than a write and read to different registers.

For 8 and 16-bit write cycles, the wait time for the back-to-back write-read operation applies only to the writing of the last BYTE or WORD of the register, which completes a single DWORD transfer.

For Indexed Address mode HBI operation, the wait time for the back-to-back write-read operation applies only to access to the internal registers and FIFOs. It does not apply to the Host Bus Interface Index Registers or the Host Bus Interface Configuration Register.

TABLE 5-2: READ AFTER WRITE TIMING RULES

After Writing...	wait for this many nanoseconds...	or Perform this many Reads of BYTE_TEST... (assuming T _{cyc} of 45ns)	before reading...
any register	45	1	the same register or any other register affected by the write
Host MAC TX Data FIFO	135	3	TX FIFO Information Register (TX_FIFO_INF)
Interrupt Configuration Register (IRQ_CFG)	60	2	Interrupt Configuration Register (IRQ_CFG)
Interrupt Enable Register (INT_EN)	90	2	Interrupt Configuration Register (IRQ_CFG)
	60	2	Interrupt Status Register (INT_STS)
Interrupt Status Register (INT_STS)	180	4	Interrupt Configuration Register (IRQ_CFG)
	170	4	Interrupt Status Register (INT_STS)
FIFO Level Interrupt Register (FIFO_INT)	90	2	Interrupt Configuration Register (IRQ_CFG)
	80	2	Interrupt Status Register (INT_STS)

TABLE 5-2: READ AFTER WRITE TIMING RULES (CONTINUED)

After Writing...	wait for this many nanoseconds...	or Perform this many Reads of BYTE_TEST... (assuming T _{cyc} of 45ns)	before reading...
Receive Configuration Register (RX_CFG)	80	2	Interrupt Configuration Register (IRQ_CFG)
	60	2	Interrupt Status Register (INT_STS)
	50	2	Interrupt Configuration Register (IRQ_CFG)
Power Management Control Register (PMT_CTRL)	165	4	Power Management Control Register (PMT_CTRL)
	170	4	Interrupt Configuration Register (IRQ_CFG)
	160	4	Interrupt Status Register (INT_STS)
General Purpose Timer Configuration Register (GPT_CFG)	55	2	General Purpose Timer Configuration Register (GPT_CFG)
	170	4	General Purpose Timer Count Register (GPT_CNT)
1588 Command and Control Register (1588_CMD_CTL)	70	2	Interrupt Configuration Register (IRQ_CFG)
	50	2	Interrupt Status Register (INT_STS)
	50	2	1588 Interrupt Status Register (1588_INT_STS)
1588 Interrupt Status Register (1588_INT_STS)	60	2	Interrupt Configuration Register (IRQ_CFG)
1588 Interrupt Enable Register (1588_INT_EN)	60	2	Interrupt Configuration Register (IRQ_CFG)
Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA)	50	2	Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD) Note 10
General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)	60	2	Interrupt Configuration Register (IRQ_CFG)

Note 10: This timing applies only to the auto-increment and auto-decrement modes of Switch Fabric CSR register access.

5.2.2 BACK-TO-BACK READ CYCLES

There are also restrictions on specific back-to-back host read operations. These restrictions concern reading specific registers after reading a resource that has side effects. In many cases there is a delay between reading the device, and the subsequent indication of the expected change in the control and status register values.

In order to prevent the host from reading stale data on back-to-back reads, minimum wait periods have been established. These periods are specified in [Table 5-3](#). The host processor is required to wait the specified period of time between read operations of specific combinations of resources. The wait period is dependent upon the combination of registers being read.

Performing “dummy” reads of the [Byte Order Test Register \(BYTE_TEST\)](#) register is a convenient way to guarantee that the minimum wait time restriction is met. [Table 5-3](#) below also shows the number of dummy reads that are required for back-to-back read operations. The number of BYTE_TEST reads in this table is based on the minimum timing for T_{cyc} (45ns). For microprocessors with slower busses the number of reads may be reduced as long as the total time is equal to, or greater than the time specified in the table. Dummy reads of the BYTE_TEST register are not required as long as the minimum time period is met.

Note that depending on the host interface mode in use, the basic host interface cycle may naturally provide sufficient time between reads. It is required of the system design and register access mechanisms to ensure the proper timing. For example, multiple reads to the same register may occur faster than reads to different registers.

For 8 and 16-bit read cycles, the wait time for the back-to-back read operation is required only after the reading of the last BYTE or WORD of the register, which completes a single DWORD transfer. There is no wait requirement between the BYTE or WORD accesses within the DWORD transfer.

TABLE 5-3: READ AFTER READ TIMING RULES

After reading...	wait for this many nanoseconds...	or Perform this many Reads of BYTE_TEST... (assuming T_{cyc} of 45ns)	before reading...
Host MAC RX Data FIFO	135	3	RX FIFO Information Register (RX_FIFO_INF)
Host MAC RX Status FIFO	135	3	RX FIFO Information Register (RX_FIFO_INF)
Host MAC TX Status FIFO	135	3	TX FIFO Information Register (TX_FIFO_INF)
Host MAC RX Dropped Frames Counter Register (RX_DROP)	180	4	Host MAC RX Dropped Frames Counter Register (RX_DROP)
Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA)	50	2	Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD) Note 11
Virtual PHY Auto-Negotiation Expansion Register (VPHY_AN_EXP)	40	1	Virtual PHY Auto-Negotiation Expansion Register (VPHY_AN_EXP)

Note 11: This timing applies only to the auto-increment and auto-decrement modes of Switch Fabric CSR register access.

6.0 CLOCKS, RESETS, AND POWER MANAGEMENT

6.1 Clocks

The device provides generation of all system clocks as required by the various sub-modules of the device. The clocking sub-system is comprised of the following:

- [Crystal Oscillator](#)
- [PHY PLL](#)

6.1.1 CRYSTAL OSCILLATOR

The device requires a fixed-frequency 25 MHz clock source for use by the internal clock oscillator and PLL. This is typically provided by attaching a 25 MHz crystal to the **OSCI** and **OSCO** pins as specified in [Section 20.7, "Clock Circuit," on page 593](#). Optionally, this clock can be provided by driving the **OSCI** input pin with a single-ended 25 MHz clock source. If a single-ended source is selected, the clock input must run continuously for normal device operation. Power savings modes allow for the oscillator or external clock input to be halted.

The crystal oscillator can be disabled as describe in [Section 6.3.4, "Chip Level Power Management," on page 52](#).

For system level verification, the crystal oscillator output can be enabled onto the **IRQ** pin. See [Section 8.2.10, "Clock Output Test Mode," on page 77](#).

Power for the crystal oscillator is provided by a dedicated regulator or separate input pin. See [Section 4.1.2, "1.2 V Crystal Oscillator Regulator," on page 32](#).

<p>Note: Crystal specifications are provided in Table 20-13, "Crystal Specifications," on page 593.</p>
--

6.1.2 PHY PLL

The PHY module receives the 25 MHz reference clock and, in addition to its internal clock usage, outputs a main system clock that is used to derive device sub-system clocks.

The PHY PLL can be disabled as describe in [Section 6.3.4, "Chip Level Power Management," on page 52](#). The PHY PLL will be disabled only when requested *and* if the PHY ports are in a power down mode.

Power for PHY PLL is provided by an external input pin, usually sourced by the device's 1.2V core regulator. See [Section 4.0, "Power Connections," on page 30](#).

6.2 Resets

The device provides multiple hardware and software reset sources, which allow varying levels of the device to be reset. All resets can be categorized into three reset types as described in the following sections:

- [Chip-Level Resets](#)
 - [Power-On Reset \(POR\)](#)
 - [RST# Pin Reset](#)
- [Multi-Module Resets](#)
 - [DIGITAL RESET \(DIGITAL_RST\)](#)
- [Single-Module Resets](#)
 - [Port A PHY Reset](#)
 - [Port B PHY Reset](#)
 - [Virtual PHY Reset](#)
 - [Host MAC Sub-System Reset](#)
 - [Switch Reset](#)
 - [1588 Reset](#)

The device supports the use of configuration straps to allow automatic custom configurations of various device parameters. These configuration strap values are set upon de-assertion of all chip-level resets and can be used to easily set the default parameters of the chip at power-on or pin (RST#) reset. Refer to [Section 6.3, "Power Management," on page 49](#) for detailed information on the usage of these straps.

[Table 6-1](#) summarizes the effect of the various reset sources on the device. Refer to the following sections for detailed information on each of these reset types.

TABLE 6-1: RESET SOURCES AND AFFECTED DEVICE FUNCTIONALITY

Module/ Functionality	POR	RST# Pin	Digital Reset
25 MHz Oscillator	(1)		
Voltage Regulators	(2)		
Host MAC sub-system	X	X	X
Switch Fabric	X	X	X
Switch Logic	X	X	X
Switch Registers	X	X	X
Switch MAC 0	X	X	X
Switch MAC 1	X	X	X
Switch MAC 2	X	X	X
PHY A	X	X	
PHY B	X	X	
PHY Common	(3)		
Voltage Supervision	(3)		
PLL	(3)		
Virtual PHY	X	X	
1588 Clock / Event Gen.	X	X	X
1588 Timestamp Unit 0	X	X	X
1588 Timestamp Unit 1	X	X	X
1588 Timestamp Unit 2	X	X	X
SPI/SQI Slave	X	X	X
Host Bus Interface	X	X	X
Note 1: POR is performed by the XTAL voltage regulator, not at the system level 2: POR is performed internal to the voltage regulators 3: POR is performed internal to the PHY 4: Strap inputs are not re-latched, however Soft-straps are returned to their previously latched pin defaults before they are potentially updated by the EEPROM values. 5: Part of EEPROM loading 6: Only those output pins that are used for straps			

TABLE 6-1: RESET SOURCES AND AFFECTED DEVICE FUNCTIONALITY (CONTINUED)

Module/ Functionality	POR	RST# Pin	Digital Reset
Power Management	X	X	X
Device EEPROM Loader	X	X	X
I2C Master	X	X	X
GPIO/LED Controller	X	X	X
General Purpose Timer	X	X	X
Free Running Counter	X	X	X
System CSR	X	X	X
Config. Straps Latched	YES	YES	NO ⁽⁴⁾
EEPROM Loader Run	YES	YES	YES
Reload Host MAC Addr.	⁽⁵⁾	⁽⁵⁾	⁽⁵⁾
Tristate Output Pins ⁽⁶⁾	YES	YES	
RST# Pin Driven Low			
Note 1: POR is performed by the XTAL voltage regulator, not at the system level 2: POR is performed internal to the voltage regulators 3: POR is performed internal to the PHY 4: Strap inputs are not re-latched, however Soft-straps are returned to their previously latched pin defaults before they are potentially updated by the EEPROM values. 5: Part of EEPROM loading 6: Only those output pins that are used for straps			

6.2.1 CHIP-LEVEL RESETS

A chip-level reset event activates all internal resets, effectively resetting the entire device. A chip-level reset is initiated by assertion of any of the following input events:

- [Power-On Reset \(POR\)](#)
- [RST# Pin Reset](#)

Chip-level reset/configuration completion can be determined by first polling the [Byte Order Test Register \(BYTE_TEST\)](#). The returned data will be invalid until the Host interface resets are complete. Once the returned data is the correct byte ordering value, the Host interface resets have completed.

The completion of the entire chip-level reset must be determined by polling the READY bit of the [Hardware Configuration Register \(HW_CFG\)](#) or [Power Management Control Register \(PMT_CTRL\)](#) until it is set. When set, the READY bit indicates that the reset has completed and the device is ready to be accessed.

With the exception of the [Hardware Configuration Register \(HW_CFG\)](#), [Power Management Control Register \(PMT_CTRL\)](#), [Byte Order Test Register \(BYTE_TEST\)](#), and [Reset Control Register \(RESET_CTL\)](#), read access to any internal resources should not be done by S/W while the READY bit is cleared. Writes to any address are invalid until the READY bit is set.

A chip-level reset involves tuning of the variable output level pads, latching of configuration straps and generation of the master reset.

CONFIGURATION STRAPS LATCHING

During POR or RST# pin reset, the latches for the straps are open. Following the release of POR or RST# pin reset, the latches for the straps are closed.

VARIABLE LEVEL I/O PAD TUNING

Following the release of the POR or RST# pin resets, a 1 uS pulse (active low), is sent into the VO tuning circuit. 2 uS later, the output pins are enabled. The 2 uS delay allows time for the variable output level pins to tune before enabling the outputs and also provides input hold time for strap pins that are shared with output pins.

MASTER RESET AND CLOCK GENERATION RESET

Following the enabling of the output pins, the reset is synchronized to the main system clock to become the master reset. Master reset is used to generate the local resets and to reset the clocks generation.

LAN9352

6.2.1.1 Power-On Reset (POR)

A power-on reset occurs whenever power is initially applied to the device or if the power is removed and reapplied to the device. This event resets all circuitry within the device. Configuration straps are latched and EEPROM loading is performed as a result of this reset. The POR is used to trigger the tuning of the Variable Level I/O Pads as well as a chip-level reset.

Following valid voltage levels, a POR reset typically takes approximately 21 ms, plus any additional time (91 us per byte) for data loaded from the EEPROM. A full 64KB EEPROM load would complete in approximately 6 seconds.

6.2.1.2 RST# Pin Reset

Driving the **RST#** input pin low initiates a chip-level reset. This event resets all circuitry within the device. Use of this reset input is optional, but when used, it must be driven for the period of time specified in [Section 20.6.3, "Reset and Configuration Strap Timing," on page 590](#). Configuration straps are latched, and EEPROM loading is performed as a result of this reset.

A **RST#** pin reset typically takes approximately 760 μ s plus any additional time (91 us per byte) for data loaded from the EEPROM. A full 64KB EEPROM load would complete in approximately 6 seconds.

Note: The **RST#** pin is pulled-high internally. If unused, this signal can be left unconnected. Do not rely on internal pull-up resistors to drive signals external to the device.

Please refer to [Table 3-10, "Miscellaneous Pin Descriptions," on page 28](#) for a description of the **RST#** pin.

6.2.2 BLOCK-LEVEL RESETS

The block level resets contain an assortment of reset register bit inputs and generate resets for the various blocks. Block level resets can affect one or multiple modules.

6.2.2.1 Multi-Module Resets

Multi-module resets activate multiple internal resets, but do not reset the entire chip. Configuration straps are *not* latched upon multi-module resets. A multi-module reset is initiated by assertion of the following:

- **DIGITAL RESET (DIGITAL_RST)**

Multi-module reset/configuration completion can be determined by first polling the [Byte Order Test Register \(BYTE_TEST\)](#). The returned data will be invalid until the Host interface resets are complete. Once the returned data is the correct byte ordering value, the Host interface resets have completed.

The completion of the entire chip-level reset must be determined by polling the READY bit of the [Hardware Configuration Register \(HW_CFG\)](#) or [Power Management Control Register \(PMT_CTRL\)](#) until it is set. When set, the READY bit indicates that the reset has completed and the device is ready to be accessed.

With the exception of the [Hardware Configuration Register \(HW_CFG\)](#), [Power Management Control Register \(PMT_CTRL\)](#), [Byte Order Test Register \(BYTE_TEST\)](#), and [Reset Control Register \(RESET_CTL\)](#), read access to any internal resources should not be done by S/W while the READY bit is cleared. Writes to any address are invalid until the READY bit is set.

Note: The digital reset does not reset register bits designated as NASR.

DIGITAL RESET (DIGITAL_RST)

A digital reset is performed by setting the DIGITAL_RST bit of the [Reset Control Register \(RESET_CTL\)](#). A digital reset will reset all device sub-modules except the Ethernet PHYs. EEPROM loading is performed following this reset. Configuration straps are *not* latched as a result of a digital reset. However, soft straps are first returned to their previously latched pin values and register bits that default to strap values are reloaded.

A digital reset typically takes approximately 760 μ s plus any additional time (91 uS per byte) for data loaded from the EEPROM. A full 64KB EEPROM load would complete in approximately 6 seconds.

6.2.2.2 Single-Module Resets

A single-module reset will reset only the specified module. Single-module resets do *not* latch the configuration straps or initiate the EEPROM Loader. A single-module reset is initiated by assertion of the following:

- [Port A PHY Reset](#)
- [Port B PHY Reset](#)
- [Virtual PHY Reset](#)
- [Host MAC Sub-System Reset](#)
- [Switch Reset](#)
- [1588 Reset](#)

Port A PHY Reset

A Port A PHY reset is performed by setting the PHY_A_RST bit of the [Reset Control Register \(RESET_CTL\)](#) or the Soft Reset bit in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). Upon completion of the Port A PHY reset, the PHY_A_RST and Soft Reset bits are automatically cleared. No other modules of the device are affected by this reset.

Port A PHY reset completion can be determined by polling the PHY_A_RST bit in the [Reset Control Register \(RESET_CTL\)](#) or the Soft Reset bit in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) until it clears. Under normal conditions, the PHY_A_RST and Soft Reset bit will clear approximately 102 μ s after the Port A PHY reset occurrence.

Note: When using the Soft Reset bit to reset the Port A PHY, register bits designated as NASR are not reset.

In addition to the methods above, the Port A PHY is automatically reset after returning from a PHY power-down mode. This reset differs in that the PHY power-down mode reset does not reload or reset any of the PHY registers. Refer to [Section 12.2.10, "PHY Power-Down Modes," on page 231](#) for additional information.

Refer to [Section 12.2.13, "Resets," on page 236](#) for additional information on Port A PHY resets.

Port B PHY Reset

A Port B PHY reset is performed by setting the PHY_B_RST bit of the [Reset Control Register \(RESET_CTL\)](#) or the Soft Reset bit in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). Upon completion of the Port B PHY reset, the PHY_B_RST and Soft Reset bits are automatically cleared. No other modules of the device are affected by this reset.

Port B PHY reset completion can be determined by polling the PHY_B_RST bit in the [Reset Control Register \(RESET_CTL\)](#) or the Soft Reset bit in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) until it clears. Under normal conditions, the PHY_B_RST and Soft Reset bit will clear approximately 102 μ s after the Port B PHY reset occurrence.

Note: When using the Soft Reset bit to reset the Port B PHY, register bits designated as NASR are not reset.

In addition to the methods above, the Port B PHY is automatically reset after returning from a PHY power-down mode. This reset differs in that the PHY power-down mode reset does not reload or reset any of the PHY registers. Refer to [Section 12.2.10, "PHY Power-Down Modes," on page 231](#) for additional information.

Refer to [Section 12.2.13, "Resets," on page 236](#) for additional information on Port B PHY resets.

Virtual PHY Reset

A Virtual PHY reset is performed by setting the **Virtual PHY Reset (VPHY_RST)** bit of the [Reset Control Register \(RESET_CTL\)](#) or Reset in the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#). No other modules of the device are affected by this reset.

Virtual PHY reset completion can be determined by polling the VPHY_0_RST bit in the [Reset Control Register \(RESET_CTL\)](#) or the Reset bit in the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#) until it clears. Under normal conditions, the VPHY_0_RST and Reset bit will clear approximately 1 μ s after the Virtual PHY reset occurrence.

Refer to [Section 12.3.2, "Virtual PHY Resets," on page 304](#) for additional information on Virtual PHY resets.

Host MAC Sub-System Reset

A Host MAC sub-system reset is performed by setting the HMAC_RST bit in the [Reset Control Register \(RESET_CTL\)](#). In addition, the MAC address of the Host MAC is reloaded from the EEPROM, using the device EEPROM loader.

This will reset the Host MAC and FIFOs, including:

- MAC
- Address Filtering
- Wake-On-LAN
- RX Checksum Offload
- TX Checksum Offload
- Energy Efficient Ethernet Control and Counters
- FIFOs
- Flow Control Logic

The following registers and register fields will be reset:

- All registers described in [Section 11.13, "Host MAC & FIFO Interface Registers," on page 188](#)

Note: The HBI register locks associated with these register are also reset.

- All registers described in [Section 11.14, "Host MAC Control and Status Registers," on page 202](#)

- The [EEPROM Controller Busy \(EPC_BUSY\)](#) and [Configuration Loaded \(CFG_LOADED\)](#) bits in the [EEPROM Command Register \(E2P_CMD\)](#) (set and cleared respectively)

Note: The bits in the [Interrupt Status Register \(INT_STS\)](#) listed in [Section 8.2.5, "Host MAC Interrupts"](#) are not reset.

Note: Host MAC and FIFO related bits in the [Hardware Configuration Register \(HW_CFG\)](#) are not reset.

Host MAC reset completion can be determined by polling the HMAC_RST bit in the [Reset Control Register \(RESET_CTL\)](#) until it clears.

Switch Reset

A reset of the Switch Fabric, including its MACs, is performed by setting the SW_RESET bit in the [Switch Reset Register \(SW_RESET\)](#). The bit must then be manually cleared.

The registers described in [Section 13.7, "Switch Fabric Control and Status Registers," on page 363](#) are reset. The functionality described in [Section 13.5, "Switch Fabric Interface Logic," on page 343](#) and the registers described in [Section 13.6, "Switch Fabric Interface Logic Registers," on page 348](#) are *not* reset.

No other modules of the device are affected by this reset.

1588 Reset

A reset of all 1588 related logic, including the clock/event generation and 1588 TSUs, is performed by setting the [1588 Reset \(1588_RESET\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#).

The registers described in [Section 15.0, "IEEE 1588," on page 473](#) are reset.

No other modules of the device are affected by this reset.

1588 reset completion can be determined by polling the [1588 Reset \(1588_RESET\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) until it clears.

6.2.3 RESET REGISTERS

6.2.3.1 Reset Control Register (RESET_CTL)

Offset: 1F8h Size: 32 bits

This register contains software controlled resets.

Note: This register can be read while the device is in the reset or not ready / power savings states without leaving the host interface in an intermediate state. If the host interface is in a reset state, returned data may be invalid.

It is not necessary to read all four bytes of this register. DWORD access rules do not apply to this register.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6	RESERVED	RO	-
5	Host MAC Reset (HMAC_RST) Setting this bit resets the Host MAC sub-system. When the Host MAC sub-system is released from reset, this bit is automatically cleared. All writes to this bit are ignored while this bit is set. Note: This bit is not accessible via the EEPROM Loader's register initialization function (Section 14.4.5).	R/W SC	0b
4	RESERVED	RO	-
3	Virtual PHY Reset (VPHY_RST) Setting this bit resets the Virtual PHY. When the Virtual PHY is released from reset, this bit is automatically cleared. All writes to this bit are ignored while this bit is set. Note: This bit is not accessible via the EEPROM Loader's register initialization function (Section 14.4.5).	R/W SC	0b

Bits	Description	Type	Default
2	<p>Port B PHY Reset (PHY_B_RST) Setting this bit resets the Port B PHY. The internal logic automatically holds the PHY reset for a minimum of 102uS. When the Port B PHY is released from reset, this bit is automatically cleared. All writes to this bit are ignored while this bit is set.</p> <p>Note: This bit is not accessible via the EEPROM Loader's register initialization function (Section 14.4.5).</p>	R/W SC	0b
1	<p>Port A PHY Reset (PHY_A_RST) Setting this bit resets the Port A PHY. The internal logic automatically holds the PHY reset for a minimum of 102uS. When the Port A PHY is released from reset, this bit is automatically cleared. All writes to this bit are ignored while this bit is set.</p> <p>Note: This bit is not accessible via the EEPROM Loader's register initialization function (Section 14.4.5).</p>	R/W SC	0b
0	<p>Digital Reset (DIGITAL_RST) Setting this bit resets the complete chip except the PLL, Virtual PHY, Port B PHY and Port A PHY. All system CSRs are reset except for any NASR type bits. Any in progress EEPROM commands (including RELOAD) are terminated.</p> <p>The EEPROM Loader will automatically reload the configuration following this reset, but will not reset the Virtual PHY, Port B PHY or Port A PHY. If desired, the above PHY resets can be issued once the device is configured.</p> <p>When the chip is released from reset, this bit is automatically cleared. All writes to this bit are ignored while this bit is set.</p> <p>Note: This bit is not accessible via the EEPROM Loader's register initialization function (Section 14.4.5).</p>	R/W SC	0b

6.3 Power Management

The device supports several block and chip level power management features as well as wake-up event detection and notification.

6.3.1 WAKE-UP EVENT DETECTION

6.3.1.1 Host MAC Wake on LAN (WoL)

The Host MAC provides the following Wake-on-LAN detection modes:

Perfect DA (Destination Address): This mode, enabled by the [Perfect DA Wakeup Enable \(PFDA_EN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), will trigger a WoL event when an incoming frame has a destination address field that exactly matches the MAC address programmed into the MAC. The [Perfect DA Frame Received \(PFDA_FR\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) will be set when PFDA_EN is set, and a Perfect DA event occurs.

Broadcast: This mode, enabled by the [Broadcast Wakeup Enable \(BCST_EN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), will trigger a WoL event when an incoming frame is a broadcast frame. The [Broadcast Frame Received \(BCAST_FR\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) will be set when BCST_EN is set, and a Broadcast event occurs.

Remote Wake-up Frame: This mode, enabled by the [Wake-Up Frame Enable \(WUEN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), will trigger a WoL event when an incoming frame is accepted based on the Wake-Up Filter registers in the Host MAC. The [Remote Wake-Up Frame Received \(WUFR\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) will be set when WUEN is set, and a Remove Wake-up Frame event occurs.

Magic Packet: This mode, enabled by the [Magic Packet Enable \(MPEN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), will trigger a WoL event when an incoming frame is accepted and is a Magic Packet. The [Magic Packet Received \(MPR\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) will be set when MPEN is set, and a Magic Packet event occurs.

If any of the PFDA_FR, BCAST_FR, WUFR or MPR bits are set, the [Wake On Status \(WOL_STS\)](#) bit of the [Power Management Control Register \(PMT_CTRL\)](#) will be set. The [Wake-On-Enable \(WOL_EN\)](#) enables this bit as a PME event.

In addition, when the [Power Management Wakeup \(PM_WAKE\)](#) bit in the [Power Management Control Register \(PMT_CTRL\)](#) is set, these events can wake up the chip.

Refer to [Section 11.6.1, "Perfect DA Detection," on page 157](#), [Section 11.6.2, "Broadcast Detection," on page 158](#), [Section 11.6.3, "Wake-up Frame Detection," on page 158](#) and [Section 11.6.4, "Magic Packet Detection," on page 163](#) for additional details on these features.

6.3.1.2 PHY A & B Energy Detect

Energy Detect Power Down mode reduces PHY power consumption. In energy-detect power-down mode, the PHY will resume from power-down when energy is seen on the cable (typically from link pulses) and set the ENERGYON interrupt bit in the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#).

Refer to [Section 12.2.10.2, "Energy Detect Power-Down," on page 231](#) for details on the operation and configuration of the PHY energy-detect power-down mode.

Note: If a carrier is present when Energy Detect Power Down is enabled, then detection will occur immediately.

If enabled, via the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#), the PHY will generate an interrupt. This interrupt is reflected in the [Interrupt Status Register \(INT_STS\)](#), bit 26 (PHY_INT_A) for PHY A and bit 27 (PHY_INT_B) for PHY B. The INT_STS register bits will trigger the IRQ interrupt output pin if enabled, as described in [Section 8.2.3, "Ethernet PHY Interrupts," on page 75](#).

LAN9352

The energy-detect PHY interrupts will also set the appropriate **Energy-Detect / WoL Status Port A (ED_WOL_STS_A)** or **Energy-Detect / WoL Status Port B (ED_WOL_STS_B)** bit of the **Power Management Control Register (PMT_CTRL)**. The **Energy-Detect / WoL Enable Port A (ED_WOL_EN_A)** and **Energy-Detect / WoL Enable Port B (ED_WOL_EN_B)** bits will enable the corresponding status bits as a PME event.

Note: Any PHY interrupt will set the above status bits. The Host should only enable the appropriate PHY interrupt source in the **PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)**.

6.3.1.3 PHY A & B Wake on LAN (WoL)

PHY A and B provide WoL event detection of Perfect DA, Broadcast, Magic Packet, and Wakeup frames. This is in addition to any WoL functionality provided by the Host MAC.

When enabled, the PHY will detect WoL events and set the WoL interrupt bit in the **PHY x Interrupt Source Flags Register (PHY_INTERRUPT_SOURCE_x)**. If enabled via the **PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)**, the PHY will generate an interrupt. This interrupt is reflected in the **Interrupt Status Register (INT_STS)**, bit 26 (PHY_INT_A) for PHY A and bit 27 (PHY_INT_B) for PHY B. The INT_STS register bits will trigger the IRQ interrupt output pin if enabled, as described in [Section 8.2.3, "Ethernet PHY Interrupts," on page 75](#).

Refer to [Section 12.2.12, "Wake on LAN \(WoL\)," on page 232](#) for details on the operation and configuration of the PHY WoL.

The WoL PHY interrupts will also set the appropriate **Energy-Detect / WoL Status Port A (ED_WOL_STS_A)** or **Energy-Detect / WoL Status Port B (ED_WOL_STS_B)** bit of the **Power Management Control Register (PMT_CTRL)**. The **Energy-Detect / WoL Enable Port A (ED_WOL_EN_A)** and **Energy-Detect / WoL Enable Port B (ED_WOL_EN_B)** bits enable the corresponding status bits as a PME event.

Note: Any PHY interrupt will set the above status bits. The Host should only enable the appropriate PHY interrupt source in the **PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)**.

6.3.2 WAKE-UP (PME) NOTIFICATION

A simplified diagram of the logic that controls the **PME** output pin and PME interrupt can be seen in [Figure 6-1](#).

The PME module handles the latching of the Host MAC **Wake On Status (WOL_STS)** bit, the PHY B **Energy-Detect / WoL Status Port B (ED_WOL_STS_B)** bit and the PHY A **Energy-Detect / WoL Status Port A (ED_WOL_STS_A)** bit in the **Power Management Control Register (PMT_CTRL)**.

This module also masks the status bits with the corresponding enable bits (**Wake-On-Enable (WOL_EN)**, **Energy-Detect / WoL Enable Port B (ED_WOL_EN_B)** and **Energy-Detect / WoL Enable Port A (ED_WOL_EN_A)**) and combines the results together to generate the **Power Management Interrupt Event (PME_INT)** status bit in the **Interrupt Status Register (INT_STS)**. The PME_INT status bit is then masked with the **Power Management Event Interrupt Enable (PME_INT_EN)** bit and combined with the other interrupt sources to drive the **IRQ** output pin.

Note: The PME interrupt status bit (PME_INT) in the INT_STS register is set regardless of the setting of PME_INT_EN.

In addition to generating interrupt events, the PME event can also drive the **PME** output pin to indicate wake-up events exclusively. The PME event is enabled with the **PME Enable (PME_EN)** in the **Power Management Control Register (PMT_CTRL)**. The **PME** output pin characteristics can be configured via the **PME Buffer Type (PME_TYPE)**, **PME Indication (PME_IND)** and **PME Polarity (PME_POL)** bits of the **Power Management Control Register (PMT_CTRL)**. These bits allow the **PME** output pin to be open-drain, active high push-pull or active-low push-pull and configure the output to be continuous, or pulse for 50 ms.

In system configurations where the **PME** output pin is shared among multiple devices (wired ORed), the WOL_STS, ED_WOL_STS_B and ED_WOL_STS_A bits within the PMT_CTRL register can be read to determine which device is driving the PME signal.

When the PM_WAKE bit of the **Power Management Control Register (PMT_CTRL)** is set, the PME event will automatically wake up the system in certain chip level power modes, as described in [Section 6.3.4.2, "Exiting Low Power Modes," on page 53](#). This is done independent from the values of the PME_EN, PME_POL, PME_IND and PME_TYPE register bits.

6.3.3 BLOCK LEVEL POWER MANAGEMENT

The device supports software controlled clock disabling of various modules in order to reduce power consumption.

Note: Disabling individual blocks does not automatically reset the block, it only places it into a static non-operational state in order to reduce the power consumption of the device. If a block reset is not performed before re-enabling the block, then care must be taken to ensure that the block is in a state where it can be disabled and then re-enabled.

6.3.3.1 Disabling The Host MAC

The entire Host MAC may be disabled by setting the HMAC_DIS bit in the [Power Management Control Register \(PMT_CTRL\)](#). As a safety precaution, in order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.

The Host MAC WoL detection can be left functioning by using the HMAC_SYS_ONLY_DIS bit instead, which keeps the RX and TX clocks enabled. As a safety precaution, in order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.

6.3.3.2 Disabling The Switch Fabric

The entire Switch Fabric may be disabled by setting the SWITCH_DIS bit in the [Power Management Control Register \(PMT_CTRL\)](#). As a safety precaution, in order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.

6.3.3.3 Disabling The 1588 Unit

The entire 1588 Unit, including the CSRs, may be disabled by setting the 1588_DIS bit in the [Power Management Control Register \(PMT_CTRL\)](#). As a safety precaution, in order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.

Individual Timestamp Units, including their local CSRs, may be disabled by setting the appropriate 1588_TSU_x_DIS bit in the [Power Management Control Register \(PMT_CTRL\)](#). As a safety precaution, in order for a bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.

6.3.3.4 PHY Power Down

A PHY may be placed into power-down as described in [Section 12.2.10, "PHY Power-Down Modes," on page 231](#).

6.3.3.5 LED Pins Power Down

All LED outputs may be disabled by setting the LED_DIS bit in the [Power Management Control Register \(PMT_CTRL\)](#). Open-drain / open-source LEDs are un-driven. Push-pull LEDs are still driven but are set to their inactive state.

APPLICATION NOTE: Individual LEDs can be disabled by setting them open-drain GPIO outputs with a data value of 1.

6.3.4 CHIP LEVEL POWER MANAGEMENT

The device supports power-down modes to allow applications to minimize power consumption.

Power is reduced by disabling the clocks as outlined in [Table 6-2, "Power Management States"](#). All configuration data is saved when in any power state. Register contents are not affected unless specifically indicated in the register description.

There is one normal operating power state, D0, and three power saving states: D1, D2 and D3. Although appropriate for various wake-up detection functions, the power states do not directly enable and are not enforced by these functions.

D0: Normal Mode - This is the normal mode of operation of this device. In this mode, all functionality is available. This mode is entered automatically on any chip-level reset (POR, **RST#** pin reset).

D1: System Clocks Disabled, XTAL, PLL and network clocks enabled - In this low power mode, all clocks derived from the PLL clock are disabled. The network clocks remain enabled if supplied by the PHYs. The crystal oscillator and the PLL remain enabled. Exit from this mode may be done manually or automatically.

This mode could be used for PHY General Power Down mode, PHY WoL mode and PHY Energy Detect Power Down mode.

D2: System Clocks Disabled, PLL disable requested, XTAL enabled - In this low power mode, all clocks derived from the PLL clock are disabled. The PLL is allowed to be disabled (and will disable if both of the PHYs are in either Energy Detect or General Power Down). The network clocks remain enabled if supplied by the PHYs. The crystal oscillator remains enabled. Exit from this mode may be done manually or automatically.

This mode is useful for PHY Energy Detect Power Down mode and PHY WoL mode. This mode could be used for PHY General Power Down mode.

D3: System Clocks Disabled, PLL disabled, XTAL disabled - In this low power mode, all clocks derived from the PLL clock are disabled. The PLL will be disabled. The crystal oscillator is disabled. Exit from this mode may be only be done manually.

This mode is useful for PHY General Power Down mode.

The Host must place the PHYs into General Power Down mode by setting the [Power Down \(PHY_PWR_DWN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) before setting this power state.

TABLE 6-2: POWER MANAGEMENT STATES

Clock Source	D0	D1	D2	D3
25 MHz Crystal Oscillator	ON	ON	ON	OFF
PLL	ON	ON	OFF(2)	OFF
system clocks (100 MHz, 50 MHz, 25 MHz and others)	ON	OFF	OFF	OFF
network clocks	available(1)	available(1)	available(1)	OFF(3)
Note 1: If supplied by the PHYs 2: PLL is requested to be turned off and will disable if both of the PHYs are in either Energy Detect or General Power Down 3: PHY clocks are off				

6.3.4.1 Entering Low Power Modes

To enter any of the low power modes (D1 - D3) from normal mode (D0), follow these steps:

1. Write the PM_MODE and PM_WAKE fields in the [Power Management Control Register \(PMT_CTRL\)](#) to their desired values
2. Set the wake-up detection desired per [Section 6.3.1, "Wake-Up Event Detection"](#).
3. Set the appropriate wake-up notification per [Section 6.3.2, "Wake-Up \(PME\) Notification"](#).
4. Ensure that the device is in a state where it can safely be placed into a low power mode (all packets transmitted, receivers disabled, packets processed / flushed, etc.)
5. Set the PM_SLEEP_EN bit in the [Power Management Control Register \(PMT_CTRL\)](#).

Note: The PM_MODE field cannot be changed at the same time as the PM_SLEEP_EN bit is set and the PM_SLEEP_EN bit cannot be set at the same time that the PM_MODE field is changed.

Note: The EEPROM Loader Register Data burst sequence ([Section 14.4.5](#)) can be used to achieve an initial power down state without the need of software by:

- First setting the PHYs into General Purpose Power Down by setting the PHY_PWR_DWN bit in PHY_BASIC_CONTROL_1/2 indirectly via the HMAC_MII_DATA / HMAC_MII_ACC via the MAC_CSR_CMD / MAC_CSR_DATA registers.
- Setting the PM_MODE and PM_SLEEP_EN bits in the [Power Management Control Register \(PMT_CTRL\)](#).

Upon entering any low power mode, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) and the [Power Management Control Register \(PMT_CTRL\)](#) is forced low.

Note: Upon entry into any of the power saving states the host interfaces are not functional.

6.3.4.2 Exiting Low Power Modes

Exiting from a low power mode can be done manually or automatically.

LAN9352

An automatic wake-up will occur based on the events described in [Section 6.3.2, "Wake-Up \(PME\) Notification"](#). Automatic wake-up is enabled with the [Power Management Wakeup \(PM_WAKE\)](#) bit in the [Power Management Control Register \(PMT_CTRL\)](#).

A manual wake-up is initiated by the host when:

- an HBI write (CS and **WR** or CS, **RD_WR** and **ENB**) is performed to the device. Although all writes are ignored until the device has been woken and a read performed, the host should direct the write to the [Byte Order Test Register \(BYTE_TEST\)](#). Writes to any other addresses should not be attempted until the device is awake.
- an SPI/SQI cycle (SCS# low and SCK high) is performed to the device. Although all reads and writes are ignored until the device has been woken, the host should direct the use a read of the [Byte Order Test Register \(BYTE_TEST\)](#) to wake the device. Reads and writes to any other addresses should not be attempted until the device is awake.

To determine when the host interface is functional, the [Byte Order Test Register \(BYTE_TEST\)](#) should be polled. Once the correct pattern is read, the interface can be considered functional. At this point, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) or the [Power Management Control Register \(PMT_CTRL\)](#) can be polled to determine when the device is fully awake.

For both automatic and manual wake-up, the [Device Ready \(READY\)](#) bit will go high once the device is returned to power savings state D0 and the PLL has re-stabilized. The PM_MODE and PM_SLEEP_EN fields in the [Power Management Control Register \(PMT_CTRL\)](#) will also clear at this point.

Under normal conditions, the device will wake-up within 2 ms.

6.3.5 POWER MANAGEMENT REGISTERS

6.3.5.1 Power Management Control Register (PMT_CTRL)

Offset: [084h](#) Size: 32 bits

This read-write register controls the power management features and the PME pin of the device. The ready state of the device be determined via the Device Ready (READY) bit of this register.

Note: This register can be read while the device is in the reset or not ready / power savings states without leaving the host interface in an intermediate state. If the host interface is in a reset state, returned data may be invalid.

It is not necessary to read all four bytes of this register. DWORD access rules do not apply to this register.

Bits	Description	Type	Default
31:29	<p>Power Management Mode (PM_MODE) This register field determines the chip level power management mode that will be entered when the Power Management Sleep Enable (PM_SLEEP_EN) bit is set.</p> <p>000: D0 001: D1 010: D2 011: D3 100: Reserved 101: Reserved 110: Reserved 111: Reserved</p> <p>Writes to this field are ignored if Power Management Sleep Enable (PM_SLEEP_EN) is also being written with a 1.</p> <p>This field is cleared when the device wakes up.</p>	R/W/SC	000b
28	<p>Power Management Sleep Enable (PM_SLEEP_EN) Setting this bit enters the chip level power management mode specified with the Power Management Mode (PM_MODE) field.</p> <p>0: Device is not in a low power sleep state 1: Device is in a low power sleep state</p> <p>This bit can <u>not</u> be written at the same time as the PM_MODE register field. The PM_MODE field must be set, and then this bit must be set for proper device operation.</p> <p>Writes to this bit with a value of 1 are ignored if Power Management Mode (PM_MODE) is being written with a new value.</p> <p>Note: Although not prevented by H/W, this bit should not be written with a value of 1 while Power Management Mode (PM_MODE) has a value of "D0".</p> <p>This field is cleared when the device wakes up.</p>	R/W/SC	0b
27	<p>Power Management Wakeup (PM_WAKE) When set, this bit enables automatic wake-up based on PME events.</p> <p>0: Manual Wakeup only 1: Auto Wakeup enabled</p>	R/W	0b
26	<p>LED Disable (LED_DIS) This bit disables LED outputs. Open-drain / open-source LEDs are un-driven. Push-pull LEDs are still driven but are set to their inactive state.</p> <p>0: LEDs are enabled 1: LEDs are disabled</p>	R/W	0b

Bits	Description	Type	Default
25	1588 Clock Disable (1588_DIS) This bit disables the clocks for the entire 1588 Unit. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
24	1588 Timestamp Unit 2 Clock Disable (1588_TSU_2_DIS) This bit disables the clocks for 1588 timestamp unit 2. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
23	1588 Timestamp Unit 1 Clock Disable (1588_TSU_1_DIS) This bit disables the clocks for 1588 timestamp unit 1. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
22	1588 Timestamp Unit 0 Clock Disable (1588_TSU_0_DIS) This bit disables the clocks for 1588 timestamp unit 0. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
21	RESERVED	RO	-
20	Switch Fabric Clock Disable (SWITCH_DIS) This bit disables the clocks for the Switch Fabric. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
19	Host MAC Clock Disable (HMAC_DIS) This bit disables the 25 and 100 MHz, RX and TX clocks to the MAC. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b

Bits	Description	Type	Default
18	Host MAC System Clock Only Disable (HMAC_SYS_ONLY_DIS) This bit disables the 25 and 100 MHz clocks to the MAC but leaves the RX and TX clocks active. 0: Clocks are enabled 1: Clocks are disabled In order for this bit to be set, it must be written as a 1 two consecutive times. A write of a 0 will reset the count.	R/W	0b
17	Energy-Detect / WoL Status Port B (ED_WOL_STS_B) This bit indicates an energy detect or WoL event occurred on the Port B PHY. In order to clear this bit, it is required that the event in the PHY be cleared as well. The event sources are described in Section 6.3, "Power Management," on page 49 .	R/WC	0b
16	Energy-Detect / WoL Status Port A (ED_WOL_STS_A) This bit indicates an energy detect or WoL event occurred on the Port A PHY. In order to clear this bit, it is required that the event in the PHY be cleared as well. The event sources are described in Section 6.3, "Power Management," on page 49 .	R/WC	0b
15	Energy-Detect / WoL Enable Port B (ED_WOL_EN_B) When set, the PME_INT bit in the Interrupt Status Register (INT_STS) will be asserted upon an energy-detect or WoL event from Port B. When set, the PME output pin (if enabled via the PME_EN bit) will be also asserted in accordance with the PME_IND bit upon an energy-detect or WoL event from Port B.	R/W	0b
14	Energy-Detect / WoL Enable Port A (ED_WOL_EN_A) When set, the PME_INT bit in the Interrupt Status Register (INT_STS) will be asserted upon an energy-detect or WoL event from Port A. When set, the PME output pin (if enabled via the PME_EN bit) will also be asserted in accordance with the PME_IND bit upon an energy-detect or WoL event from Port A.	R/W	0b
13:10	RESERVED	RO	-
9	Wake-On-Enable (WOL_EN) When set, the PME_INT bit in the Interrupt Status Register (INT_STS) will be asserted upon a Host MAC WOL event. When set, the PME output pin (if enabled via the PME_EN bit) will also be asserted in accordance with the PME_IND bit upon a Host MAC WOL event.	R/W	0b
8:7	RESERVED	RO	-
6	PME Buffer Type (PME_TYPE) When this bit is cleared, the PME output pin functions as an open-drain buffer for use in a wired-or configuration. When set, the PME output pin is a push-pull driver. When the PME output pin is configured as an open-drain output, the PME_POL field of this register is ignored and the output is always active low. 0: PME pin open-drain output 1: PME pin push-pull driver	R/W NASR	0b

Bits	Description	Type	Default
5	Wake On Status (WOL_STS) This bit indicates that a Wake-Up, Magic Packet, Perfect DA, or Broadcast frame was detected by the Host MAC. In order to clear this bit, it is required that the event in the Host MAC be cleared as well. The event sources are described in Section 6.3, "Power Management," on page 49 .	R/WC	0b
4	RESERVED	RO	-
3	PME Indication (PME_IND) The PME signal can be configured as a pulsed output or a static signal, which is asserted upon detection of a wake-up event. When set, the PME signal will pulse active for 50mS upon detection of a wake-up event. When cleared, the PME signal is driven continuously upon detection of a wake-up event. 0: PME driven continuously on detection of event 1: PME 50mS pulse on detection of event The PME signal can be deactivated by clearing the above status bit(s) or by clearing the appropriate enable(s).	R/W	0b
2	PME Polarity (PME_POL) This bit controls the polarity of the PME signal. When set, the PME output is an active high signal. When cleared, it is active low. Note: When PME is configured as an open-drain output, this field is ignored and the output is always active low. 0: PME active low 1: PME active high	R/W NASR	0b
1	PME Enable (PME_EN) When set, this bit enables the external PME signal pin. When cleared, the external PME signal is disabled. Note: This bit does not affect the PME_INT interrupt bit of the Interrupt Status Register (INT_STS) . 0: PME pin disabled 1: PME pin enabled	R/W	0b
0	Device Ready (READY) When set, this bit indicates that the device is ready to be accessed. Upon power-up, RST# reset, return from power savings states, Host MAC module level reset or digital reset, the host processor may interrogate this field as an indication that the device has stabilized and is fully active. This rising edge of this bit will assert the Device Ready (READY) bit in INT_STS and can cause an interrupt if enabled. Note: With the exception of the HW_CFG, PMT_CTRL, BYTE_TEST, and RESET_CTL registers, read access to any internal resources is forbidden while the READY bit is cleared. Writes to any address are invalid until this bit is set. Note: This bit is identical to bit 27 of the Hardware Configuration Register (HW_CFG) .	RO	0b

6.4 Device Ready Operation

The device supports a Ready status register bit that indicates to the Host software when the device is fully ready for operation. This bit may be read via the [Power Management Control Register \(PMT_CTRL\)](#) or the [Hardware Configuration Register \(HW_CFG\)](#).

Following power-up reset, **RST#** reset, or digital reset (see [Section 6.2, "Resets"](#)), the [Device Ready \(READY\)](#) bit indicates that the device has read, and is configured from, the contents of the EEPROM.

An EEPROM RELOAD command, via the [EEPROM Command Register \(E2P_CMD\)](#), will restart the EEPROM Loader, temporarily causing the [Device Ready \(READY\)](#) to be low.

A Host MAC reset, via the [Reset Control Register \(RESET_CTL\)](#), will utilize the EEPROM Loader, temporarily causing the [Device Ready \(READY\)](#) to be low.

Entry into any power savings state (see [Section 6.3.4, "Chip Level Power Management"](#)) other than D0 will cause [Device Ready \(READY\)](#) to be low. Upon wake-up, the [Device Ready \(READY\)](#) bit will go high once the device is returned to power savings state D0 and the PLL has re-stabilized.

7.0 CONFIGURATION STRAPS

Configuration straps allow various features of the device to be automatically configured to user defined values. Configuration straps can be organized into two main categories: [Hard-Straps](#) and [Soft-Straps](#). Both hard-straps and soft-straps are latched upon Power-On Reset (POR), or pin reset (**RST#**). The primary difference between these strap types is that soft-strap default values can be overridden by the EEPROM Loader, while hard-straps cannot.

Configuration straps which have a corresponding external pin include internal resistors in order to prevent the signal from floating when unconnected. If a particular configuration strap is connected to a load, an external pull-up or pull-down resistor should be used to augment the internal resistor to ensure that it reaches the required voltage level prior to latching. The internal resistor can also be overridden by the addition of an external resistor.

Note: The system designer must guarantee that configuration strap pins meet the timing requirements specified in [Section 20.6.3, "Reset and Configuration Strap Timing"](#). If configuration strap pins are not at the correct voltage level prior to being latched, the device may capture incorrect strap values.

7.1 Soft-Straps

Soft-strap values are latched on the release of POR or **RST#** and are overridden by values from the EEPROM Loader (when an EEPROM is present). These straps are used as direct configuration values or as defaults for CPU registers. Some, but not all, soft-straps have an associated pin. Those that do not have an associated pin, have a tie off default value. All soft-strap values can be overridden by the EEPROM Loader. Refer to [Section 14.4, "EEPROM Loader," on page 465](#) for information on the operation of the EEPROM Loader and the loading of strap values. [Table 14-4, "EEPROM Configuration Bits," on page 467](#) defines the soft-strap EEPROM bit map.

Straps which have an associated pin are also fully defined in [Section 3.0, "Pin Descriptions and Configuration," on page 10](#).

[Table 7-1](#) provides a list of all soft-straps and their associated pin or default value.

Note: The use of the term "configures" in the "Description" section of [Table 7-1](#) indicates the register bit is loaded with the strap value, while the term "Affects" means the value of the register bit is determined by the strap value and some other condition(s).

Upon setting the [Digital Reset \(DIGITAL_RST\)](#) bit in the [Reset Control Register \(RESET_CTL\)](#) or upon issuing a RELOAD command via the [EEPROM Command Register \(E2P_CMD\)](#), these straps return to their original latched (non-overridden) values if an EEPROM is no longer attached or has been erased. The associated pins are not re-sampled (i.e. the value latched on the pin during the last POR or **RST#** will be used, not the value on the pin during the digital reset or RELOAD command issuance). If it is desired to re-latch the current configuration strap pin values, a POR or **RST#** must be issued.

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS

Strap Name	Description	Pin / Default Value
LED_en_strap[5:0]	LED Enable Straps: Configures the default value for the LED Enable 5-0 (LED_EN[5:0]) bits of the LED Configuration Register (LED_CFG) .	111111b
LED_fun_strap[2:0]	LED Function Straps: Configures the default value for the LED Function 2-0 (LED_FUN[2:0]) bits of the LED Configuration Register (LED_CFG) .	000b
HBI_ale_qualification_strap	HBI ALE Qualification Strap: Configures the HBI interface to qualify the ALEHI and ALELO signals with the CS signal. 0 = address input is latched with ALEHI and ALELO 1 = address input is latched with ALEHI and ALELO only when CS is active	1b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
HBI_rw_mode_strap	<p>HBI Read / Write Mode Strap: Configures the HBI interface for separate read & write signals or direction and enable signals.</p> <p>0 = read & write 1 = direction & enable</p>	0b
HBI_cs_polarity_strap	<p>HBI Chip Select Polarity Strap: Configures the polarity of the HBI interface chip select signal.</p> <p>0 = active low 1 = active high</p>	0b
HBI_rd_rdwr_polarity_strap	<p>HBI Read, Read / Write Polarity Strap: Configures the polarity of the HBI interface read signal.</p> <p>0 = active low read 1 = active high read</p> <p>Configures the polarity of the HBI interface read / write signal.</p> <p>0 = read when 1, write when 0 (R/nW) 1 = write when 1, read when 0 (W/nR)</p>	0b
HBI_wr_en_polarity_strap	<p>HBI Write, Enable Polarity Strap: Configures the polarity of the HBI interface write signal.</p> <p>0 = active low write 1 = active high write</p> <p>Configures the polarity of the HBI interface enable signal.</p> <p>0 = active low enable 1 = active high enable</p>	0b
HBI_ale_polarity_strap	<p>HBI ALE Polarity Strap: Configures the polarity of the HBI interface ALEHI and ALELO signals.</p> <p>0 = active low strobe (address saved on rising edge) 1 = active high strobe (address saved on falling edge)</p>	1b
1588_enable_strap	<p>1588 Enable Strap: Configures the default value of the 1588 Enable (1588_ENABLE) bit in the 1588 Command and Control Register (1588_CMD_CTL).</p> <p>Note: The defaults of the 1588 register set are such that the device will perform End-to-End Transparent Clock functionality without further configuration.</p>	<u>1588EN</u>

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
auto_mdix_strap_1	<p>PHY A Auto-MDIX Enable Strap: Configures the default value of the AMDIX_EN Strap State Port A bit of the Hardware Configuration Register (HW_CFG).</p> <p>This strap is also used in conjunction with manual_mdix_strap_1 to configure PHY A Auto-MDIX functionality when the Auto-MDIX Control (AMDIXCTRL) bit in the (x=A) PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) indicates the strap settings should be used for auto-MDIX configuration.</p> <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
manual_mdix_strap_1	<p>PHY A Manual MDIX Strap: Configures MDI(0) or MDIX(1) for PHY A when the auto_mdix_strap_1 is low and the Auto-MDIX Control (AMDIXCTRL) bit in the (x=A) PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) indicates the strap settings are to be used for auto-MDIX configuration.</p> <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	0b
autoneg_strap_1	<p>PHY A Auto Negotiation Enable Strap: Configures the default value of the Auto-Negotiation Enable (PHY_AN) enable bit in the (x=A) PHY x Basic Control Register (PHY_BASIC_CONTROL_x).</p> <p>This strap also affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> Speed Select LSB (PHY_SPEED_SEL_LSB) and Duplex Mode (PHY_DUPLEX) bits of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) 10BASE-T Full Duplex and 10BASE-T Half Duplex bits of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
speed_strap_1	<p>PHY A Speed Select Strap: This strap affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> Speed Select LSB (PHY_SPEED_SEL_LSB) bit of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) 10BASE-T Full Duplex and 10BASE-T Half Duplex bits of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
duplex_strap_1	<p>PHY A Duplex Select Strap: This strap affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> Duplex Mode (PHY_DUPLEX) bit of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) 10BASE-T Full Duplex bit of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Refer to the respective register definition sections for additional information.</p>	1b
BP_EN_strap_1	<p>Switch Port 1 Backpressure Enable Strap: Configures the default value for the Port 1 Backpressure Enable (BP_EN_1) bit of the Port 1 Manual Flow Control Register (MANUAL_FC_1).</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
FD_FC_strap_1	<p>Switch Port 1 Full-Duplex Flow Control Enable Strap: This strap is used to configure the default value of the following register bits:</p> <ul style="list-style-type: none"> Port 1 Full-Duplex Transmit Flow Control Enable (TX_FC_1) and Port 1 Full-Duplex Receive Flow Control Enable (RX_FC_1) bits of the Port 1 Manual Flow Control Register (MANUAL_FC_1) <p>Refer to the respective register definition sections for additional information.</p>	1b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
FD_FC_strap_1 (cont.)	<p>PHY A Full-Duplex Flow Control Enable Strap: This strap affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> Asymmetric Pause bit of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
manual_FC_strap_1	<p>Switch Port 1 Manual Flow Control Enable Strap: Configures the default value of the Port 1 Full-Duplex Manual Flow Control Select (MANUAL_FC_1) bit in the Port 1 Manual Flow Control Register (MANUAL_FC_1).</p> <p>Refer to the respective register definition sections for additional information.</p>	0b
manual_FC_strap_1 (cont.)	<p>PHY A Manual Flow Control Enable Strap: This strap affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> Asymmetric Pause and Symmetric Pause bit of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	0b
EEE_enable_strap_1	<p>Switch Port 1 Energy Efficient Ethernet Enable Strap: Configures the default value of the Energy Efficient Ethernet (EEE_ENABLE) bit in the (x=1) Port x MAC Transmit Configuration Register (MAC_TX_CFG_x).</p> <p>Note: This has no effect when in Port 1 internal PHY mode when the PHY is in 100BASE-FX mode (lack of EEE auto-negotiation results disables EEE).</p> <p>Refer to the respective register definition sections for additional information.</p>	<u>EEEEEN</u>

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
EEE_enable_strap_1 (cont.)	<p>PHY A Energy Efficient Ethernet Enable Strap: This strap affects the default value of the following register bits (x=A):</p> <ul style="list-style-type: none"> PHY Energy Efficient Ethernet Enable (PHYEEEEEN) bit of the PHY x EDPD NLP / Crossover Time / EEE Configuration Register (PHY_EDPD_CFG_x) 100BASE-TX EEE bit of the PHY x EEE Capability Register (PHY_EEE_CAP_x) 100BASE-TX EEE bit of the PHY x EEE Advertisement Register (PHY_EEE_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	<u>EEEEEN</u>
auto_mdix_strap_2	<p>PHY B Auto-MDIX Enable Strap: Configures the default value of the AMDIX_EN Strap State Port B bit of the Hardware Configuration Register (HW_CFG).</p> <p>This strap is also used in conjunction with manual_mdix_strap_2 to configure PHY B Auto-MDIX functionality when the Auto-MDIX Control (AMDIXCTRL) bit in the (x=B) PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) indicates the strap settings should be used for auto-MDIX configuration.</p> <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
manual_mdix_strap_2	<p>PHY B Manual MDIX Strap: Configures MDI(0) or MDIX(1) for Port 2 when the auto_mdix_strap_2 is low and the Auto-MDIX Control (AMDIXCTRL) bit in the (x=B) PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) indicates the strap settings are to be used for auto-MDIX configuration.</p> <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	0b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
autoneg_strap_2	<p>PHY B Auto Negotiation Enable Strap: Configures the default value of the Auto-Negotiation Enable (PHY_AN) enable bit in the (x=B) PHY x Basic Control Register (PHY_BASIC_CONTROL_x).</p> <p>This strap also affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> • Speed Select LSB (PHY_SPEED_SEL_LSB) and Duplex Mode (PHY_DUPLEX) bits of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) • 10BASE-T Full Duplex and 10BASE-T Half Duplex bits of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) • PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
speed_strap_2	<p>PHY B Speed Select Strap: This strap affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> • Speed Select LSB (PHY_SPEED_SEL_LSB) bit of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) • 10BASE-T Full Duplex and 10BASE-T Half Duplex bits of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) • PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
duplex_strap_2	<p>PHY B Duplex Select Strap: This strap affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> • Duplex Mode (PHY_DUPLEX) bit of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) • 10BASE-T Full Duplex bit of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) • PHY Mode (MODE[2:0]) bits of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) <p>Refer to the respective register definition sections for additional information.</p>	1b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
BP_EN_strap_2	<p>Switch Port 2 Backpressure Enable Strap: Configures the default value for the Port 2 Backpressure Enable (BP_EN_2) bit of the Port 2 Manual Flow Control Register (MANUAL_FC_2).</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
FD_FC_strap_2	<p>Switch Port 2 Full-Duplex Flow Control Enable Strap: This strap is used to configure the default value of the following register bits:</p> <ul style="list-style-type: none"> Port 2 Full-Duplex Transmit Flow Control Enable (TX_FC_2) and Port 2 Full-Duplex Receive Flow Control Enable (RX_FC_2) bits of the Port 2 Manual Flow Control Register (MANUAL_FC_2). <p>Refer to the respective register definition sections for additional information.</p>	1b
FD_FC_strap_2 (cont.)	<p>PHY B Full-Duplex Flow Control Enable Strap: This strap also affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> Asymmetric Pause bit of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
manual_FC_strap_2	<p>Switch Port 2 Manual Flow Control Enable Strap: Configures the default value of the Port 2 Full-Duplex Manual Flow Control Select (MANUAL_FC_2) bit in the Port 2 Manual Flow Control Register (MANUAL_FC_2).</p> <p>Refer to the respective register definition sections for additional information.</p>	0b
manual_FC_strap_2 (cont.)	<p>PHY B Manual Flow Control Enable Strap: This strap affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> Asymmetric Pause and Symmetric Pause bits of the PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	0b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
EEE_enable_strap_2	<p>Switch Port 2 Energy Efficient Ethernet Enable Strap: Configures the default value of the Energy Efficient Ethernet (EEE_ENABLE) bit in the (x=2) Port x MAC Transmit Configuration Register (MAC_TX_CFG_x).</p> <p>Note: This has no effect when the PHY is in 100BASE-FX mode (lack of EEE auto-negotiation results disables EEE).</p> <p>Refer to the respective register definition sections for additional information.</p>	<u>EEEEEN</u>
EEE_enable_strap_2 (cont.)	<p>PHY B Energy Efficient Ethernet Enable Strap: This strap affects the default value of the following register bits (x=B):</p> <ul style="list-style-type: none"> • PHY Energy Efficient Ethernet Enable (PHYEEEEEN) bit of the PHY x EDPD NLP / Crossover Time / EEE Configuration Register (PHY_EDPD_CFG_x) • 100BASE-TX EEE bit of the PHY x EEE Capability Register (PHY_EEE_CAP_x) • 100BASE-TX EEE bit of the PHY x EEE Advertisement Register (PHY_EEE_ADV_x) <p>Note: This has no effect when the PHY is in 100BASE-FX mode.</p> <p>Refer to the respective register definition sections for additional information.</p>	<u>EEEEEN</u>
SQE_test_disable_strap_0	<p>Port 0 Virtual PHY SQE Heartbeat Disable Strap: Configures the default value of the SQEOFF bit in the Virtual PHY Special Control/Status Register (VPHY_SPECIAL_CONTROL_STATUS).</p> <p>Refer to the respective register definition sections for additional information.</p>	0b
BP_EN_strap_0	<p>Switch Port 0 Backpressure Enable Strap: Configures the default value of the Port 0 Backpressure Enable (BP_EN_0) bit of the Port 0 Manual Flow Control Register (MANUAL_FC_0).</p> <p>Refer to the respective register definition sections for additional information.</p>	1b
FD_FC_strap_0	<p>Switch Port 0 Full-Duplex Flow Control Enable Strap: This strap is used to configure the default value of the following register bits:</p> <ul style="list-style-type: none"> • Port 0 Full-Duplex Transmit Flow Control Enable (TX_FC_0) and Port 0 Full-Duplex Receive Flow Control Enable (RX_FC_0) bits of the Port 0 Manual Flow Control Register (MANUAL_FC_0) <p>Refer to the respective register definition sections for additional information.</p>	1b

TABLE 7-1: SOFT-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pin / Default Value
FD_FC_strap_0 (cont.)	Port 0 Virtual PHY Full-Duplex Flow Control Enable Strap: This strap affects the default value of the following register bits: <ul style="list-style-type: none"> Asymmetric Pause and Pause bits of the Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register (VPHY_AN_LP_BASE_ABILITY) Refer to the respective register definition sections for additional information.	1b
manual_FC_strap_0	Port 0 Virtual PHY Manual Flow Control Enable Strap: This strap affects the default value of the following register bits: <ul style="list-style-type: none"> Asymmetric Pause and Symmetric Pause bits of the Virtual PHY Auto-Negotiation Advertisement Register (VPHY_AN_ADV) Refer to the respective register definition sections for additional information.	0b

7.2 Hard-Straps

Hard-straps are latched upon Power-On Reset (POR) or pin reset (**RST#**) only. Unlike soft-straps, hard-straps always have an associated pin and cannot be overridden by the EEPROM Loader. These straps are used as either direct configuration values or as register defaults. [Table 7-2](#) provides a list of all hard-straps and their associated pin. These straps, along with their pin assignments are also fully defined in [Section 3.0, "Pin Descriptions and Configuration," on page 10](#).

TABLE 7-2: HARD-STRAP CONFIGURATION STRAP DEFINITIONS

Strap Name	Description	Pins
eeeprom_size_strap	EEPROM Size Strap: Configures the EEPROM size range. A low selects 1K bits (128 x 8) through 16K bits (2K x 8). A high selects 32K bits (4K x 8) through 512K bits (64K x 8).	<u>E2PSIZE</u>

TABLE 7-2: HARD-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pins						
host_intf_mode_strap	<p>Host Interface Mode Strap: Configures the host management mode.</p> <p>0 = SPI Mode 1 = HBI Mode</p> <p>The operating mode results from the following mapping:</p> <table><tr><th><u>MNGT1 : MNGT0</u></th><th>mngt_mode_strap</th></tr><tr><td>00</td><td>0 (SPI)</td></tr><tr><td>01, 10 or 11</td><td>1 (HBI)</td></tr></table> <p>See Table 7-3 for the combined host interface strapping.</p> <p>Note: Refer to Section 2.0, "General Description," on page 8 for additional information on the various modes of the device.</p>	<u>MNGT1 : MNGT0</u>	mngt_mode_strap	00	0 (SPI)	01, 10 or 11	1 (HBI)	<u>MNGT1 : MNGT0</u>
<u>MNGT1 : MNGT0</u>	mngt_mode_strap							
00	0 (SPI)							
01, 10 or 11	1 (HBI)							
HBI_addr_mode_strap	<p>HBI Address Mode Strap: Configures the HBI interface for multiplexed or non-multiplexed indexed addressing modes.</p> <p>0 = multiplexed 1 = non-multiplexed indexed</p> <p>See Table 7-3 for the combined host interface strapping.</p> <p>Note: Refer to Section 2.0, "General Description," on page 8 for additional information on the various modes of the device.</p>	<u>MNGT1</u>						
HBI_addr_phase_strap	<p>HBI Address Phase Strap: Configures the number of address cycles for the HBI interface when in multiplexed address mode.</p> <p>0 = single phase multiplexed 1 = dual phase multiplexed</p> <p>See Table 7-3 for the combined host interface strapping.</p> <p>Note: Refer to Section 2.0, "General Description," on page 8 for additional information on the various modes of the device.</p>	<u>MNGT3</u>						

TABLE 7-2: HARD-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pins												
HBI_data_mode_strap[1:0]	<p>HBI Data Mode Straps: Configures the data width of the HBI interface.</p> <p>00 = 8-bit 01 = 16-bit 1X = RESERVED</p> <p>The data mode results from the following mapping:</p> <table><tr><th><u>MNGT2</u> : <u>MNGT1</u> : <u>MNGT0</u></th><th>HBI_data_mode_strap</th></tr><tr><td>X00 (SPI)</td><td>reserved</td></tr><tr><td>001 (HBI multiplexed)</td><td>00 (8-bit)</td></tr><tr><td>101 (HBI multiplexed)</td><td>01 (16-bit)</td></tr><tr><td>X10 (HBI non-multi- plexed indexed)</td><td>00 (8-bit)</td></tr><tr><td>X11 (HBI non-multi- plexed indexed)</td><td>01 (16-bit)</td></tr></table> <p>Note: Effectively, the data mode is determined by <u>MNGT2</u> in multiplexed mode and by <u>MNGT0</u> in non-multiplexed indexed mode.</p> <p>See Table 7-3 for the combined host interface strapping.</p> <p>Note: Refer to Section 2.0, "General Description," on page 8 for additional information on the various modes of the device.</p>	<u>MNGT2</u> : <u>MNGT1</u> : <u>MNGT0</u>	HBI_data_mode_strap	X00 (SPI)	reserved	001 (HBI multiplexed)	00 (8-bit)	101 (HBI multiplexed)	01 (16-bit)	X10 (HBI non-multi- plexed indexed)	00 (8-bit)	X11 (HBI non-multi- plexed indexed)	01 (16-bit)	<u>MNGT2</u> : <u>MNGT1</u> : <u>MNGT0</u>
<u>MNGT2</u> : <u>MNGT1</u> : <u>MNGT0</u>	HBI_data_mode_strap													
X00 (SPI)	reserved													
001 (HBI multiplexed)	00 (8-bit)													
101 (HBI multiplexed)	01 (16-bit)													
X10 (HBI non-multi- plexed indexed)	00 (8-bit)													
X11 (HBI non-multi- plexed indexed)	01 (16-bit)													
phy_addr_sel_strap	<p>Switch PHY Address Select Strap: Configures the default MII management address values for the PHYs and Virtual PHY as detailed in Section 12.1.1, "PHY Addressing," on page 218.</p>	<u>PHYADD</u>												
fx_mode_strap_1	<p>PHY A FX Mode Strap: Selects FX mode for PHY A.</p> <p>This strap is set high when <u>FXLOSEN</u> is above 1 V (typ.) or <u>FXSDENA</u> is above 1 V (typ.).</p>	<u>FXLOSEN</u> : <u>FXSDENA</u>												
fx_mode_strap_2	<p>PHY B FX Mode Strap: Selects FX mode for PHY B.</p> <p>This strap is set high when <u>FXLOSEN</u> is above 2 V (typ.) or <u>FXSDENB</u> is above 1 V (typ.).</p>	<u>FXLOSEN</u> : <u>FXSDENB</u>												

TABLE 7-2: HARD-STRAP CONFIGURATION STRAP DEFINITIONS (CONTINUED)

Strap Name	Description	Pins
fx_los_strap_1	PHY A FX-LOS Select Strap: Selects Loss of Signal mode for PHY A. This strap is set high when FXLOSEN is above 1 V (typ.).	FXLOSEN
fx_los_strap_2	PHY B FX-LOS Select Strap: Selects Loss of Signal mode for PHY B. This strap is set high when FXLOSEN is above 2 V (typ.).	FXLOSEN

Note 1: The combined host interface strap chart is as follows:

TABLE 7-3: HBI STRAP MAPPING

<u>MNGT1</u>	<u>MNGT0</u>	<u>MNGT3</u>	<u>MNGT2</u>	Host Mode
0	0	X	X	SPI
0	1	0	0	HBI Multiplexed 1 Phase 8-bit
0	1	0	1	HBI Multiplexed 1 Phase 16-bit
0	1	1	0	HBI Multiplexed 2 Phase 8-bit
0	1	1	1	HBI Multiplexed 2 Phase 16-bit
1	0	X	X	HBI Indexed 8-bit
1	1	X	X	HBI Indexed 16-bit

8.0 SYSTEM INTERRUPTS

8.1 Functional Overview

This chapter describes the system interrupt structure of the device. The device provides a multi-tier programmable interrupt structure which is controlled by the System Interrupt Controller. The programmable system interrupts are generated internally by the various device sub-modules and can be configured to generate a single external host interrupt via the **IRQ** interrupt output pin. The programmable nature of the host interrupt provides the user with the ability to optimize performance dependent upon the application requirements. The **IRQ** interrupt buffer type, polarity and de-assertion interval are modifiable. The **IRQ** interrupt can be configured as an open-drain output to facilitate the sharing of interrupts with other devices. All internal interrupts are maskable and capable of triggering the **IRQ** interrupt.

8.2 Interrupt Sources

The device is capable of generating the following interrupt types:

- [1588 Interrupts](#)
- [Switch Fabric Interrupts](#) (Buffer Manager, Switch Engine and Port 2,1,0 MACs)
- [Ethernet PHY Interrupts](#)
- [GPIO Interrupts](#)
- [Host MAC Interrupts](#) (FIFOs)
- [Power Management Interrupts](#)
- [General Purpose Timer Interrupt](#) (GPT)
- [Software Interrupt](#) (General Purpose)
- [Device Ready Interrupt](#)
- [Clock Output Test Mode](#)

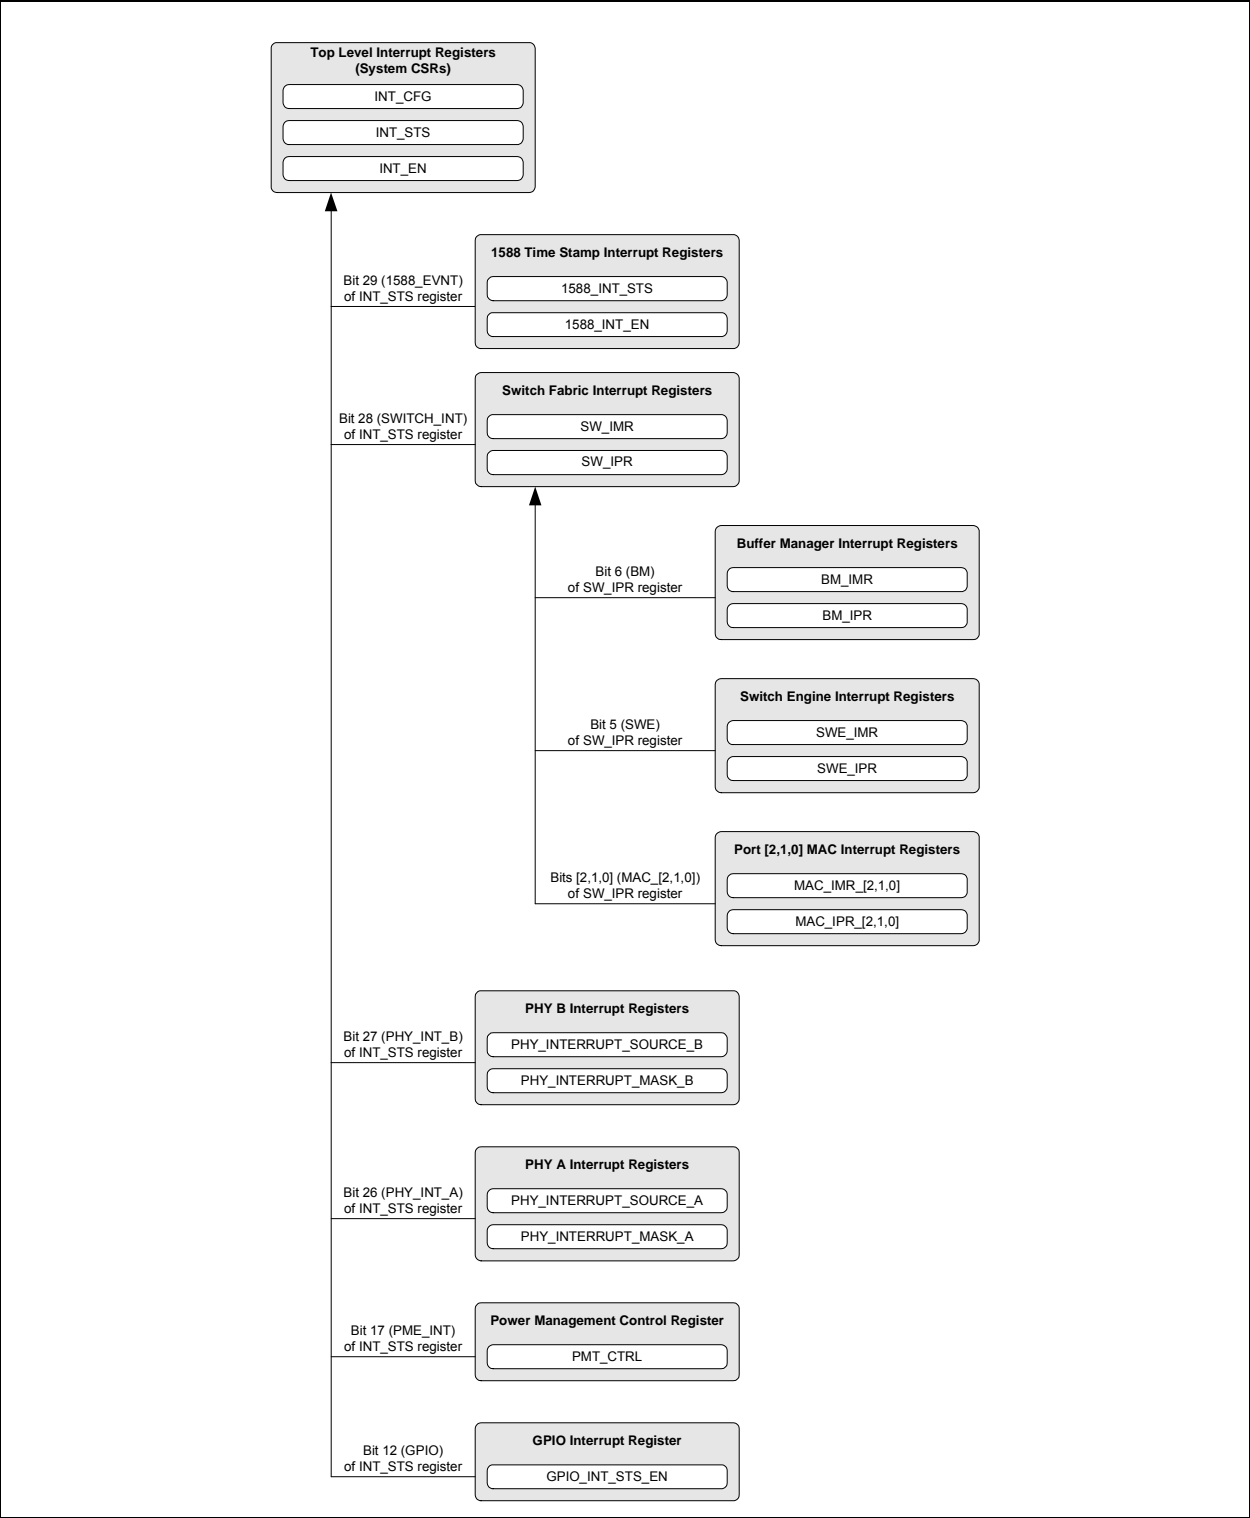
All interrupts are accessed and configured via registers arranged into a multi-tier, branch-like structure, as shown in [Figure 8-1](#). At the top level of the device interrupt structure are the [Interrupt Status Register \(INT_STS\)](#), [Interrupt Enable Register \(INT_EN\)](#) and [Interrupt Configuration Register \(IRQ_CFG\)](#).

The [Interrupt Status Register \(INT_STS\)](#) and [Interrupt Enable Register \(INT_EN\)](#) aggregate and enable/disable all interrupts from the various device sub-modules, combining them together to create the **IRQ** interrupt. These registers provide direct interrupt access/configuration to the Host MAC, General Purpose Timer, software and device ready interrupts. These interrupts can be monitored, enabled/disabled and cleared, directly within these two registers. In addition, event indications are provided for the 1588, Switch Fabric, Power Management, GPIO and Ethernet PHY interrupts. These interrupts differ in that the interrupt sources are generated and cleared in other sub-block registers. The INT_STS register does not provide details on what specific event within the sub-module caused the interrupt and requires the software to poll an additional sub-module interrupt register (as shown in [Figure 8-1](#)) to determine the exact interrupt source and clear it. For interrupts which involve multiple registers, only after the interrupt has been serviced and cleared at its source will it be cleared in the INT_STS register.

The [Interrupt Configuration Register \(IRQ_CFG\)](#) is responsible for enabling/disabling the **IRQ** interrupt output pin as well as configuring its properties. The IRQ_CFG register allows the modification of the **IRQ** pin buffer type, polarity and de-assertion interval. The de-assertion timer guarantees a minimum interrupt de-assertion period for the **IRQ** output and is programmable via the [Interrupt De-assertion Interval \(INT_DEAS\)](#) field of the [Interrupt Configuration Register \(IRQ_CFG\)](#). A setting of all zeros disables the de-assertion timer. The de-assertion interval starts when the **IRQ** pin de-asserts, regardless of the reason.

Note: The de-assertion timer does not apply to the PME interrupt. The PME interrupt is ORed into the IRQ logic following the deassertion timer gating. Assertion of the PME interrupt does not affect the de-assertion timer.
--

FIGURE 8-1: FUNCTIONAL INTERRUPT HIERARCHY



The following sections detail each category of interrupts and their related registers. Refer to the corresponding function's chapter for bit-level definitions of all interrupt registers.

8.2.1 1588 INTERRUPTS

Multiple 1588 Time Stamp interrupt sources are provided by the device. The top-level [1588 Interrupt Event \(1588_EVNT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) provides indication that a 1588 interrupt event occurred in the [1588 Interrupt Status Register \(1588_INT_STS\)](#).

The [1588 Interrupt Enable Register \(1588_INT_EN\)](#) provides enabling/disabling of all 1588 interrupt conditions. The [1588 Interrupt Status Register \(1588_INT_STS\)](#) provides the status of all 1588 interrupts. These include TX/RX 1588 clock capture indication on Ports 2,1,0, 1588 clock capture for GPIO events, as well as 1588 timer interrupt indication.

In order for a 1588 interrupt event to trigger the external IRQ interrupt pin, the desired 1588 interrupt event must be enabled in the [1588 Interrupt Enable Register \(1588_INT_EN\)](#), bit 29 ([1588_EVNT_EN](#)) of the [Interrupt Enable Register \(INT_EN\)](#) must be set and the **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

For additional details on the 1588 Time Stamp interrupts, refer to [Section 15.0, "IEEE 1588," on page 473](#).

8.2.2 SWITCH FABRIC INTERRUPTS

Multiple Switch Fabric interrupt sources are provided by the device in a three-tiered register structure as shown in [Figure 8-1](#). The top-level [Switch Fabric Interrupt Event \(SWITCH_INT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) provides indication that a Switch Fabric interrupt event occurred in the [Switch Global Interrupt Pending Register \(SW_IPR\)](#).

The [Switch Global Interrupt Pending Register \(SW_IPR\)](#) and [Switch Global Interrupt Mask Register \(SW_IMR\)](#) provide status and enabling/disabling of all Switch Fabric sub-modules interrupts (Buffer Manager, Switch Engine and Port 2,1,0 MACs).

The low-level Switch Fabric sub-module interrupt pending and mask registers of the Buffer Manager, Switch Engine and Port 2,1,0 MACs provide multiple interrupt sources from their respective sub-modules. These low-level registers provide the following interrupt sources:

- **Buffer Manager** ([Buffer Manager Interrupt Mask Register \(BM_IMR\)](#) and [Buffer Manager Interrupt Pending Register \(BM_IPR\)](#))
 - Status B Pending
 - Status A Pending
- **Switch Engine** ([Switch Engine Interrupt Mask Register \(SWE_IMR\)](#) and [Switch Engine Interrupt Pending Register \(SWE_IPR\)](#))
 - Interrupt Pending
- **Port 2,1,0 MACs** ([Port x MAC Interrupt Mask Register \(MAC_IMR_x\)](#) and [Port x MAC Interrupt Pending Register \(MAC_IPR_x\)](#))
 - No currently supported interrupt sources. These registers are reserved for future use.

In order for a Switch Fabric interrupt event to trigger the external IRQ interrupt pin, the following must be configured:

- The desired Switch Fabric sub-module interrupt event must be enabled in the corresponding mask register ([Buffer Manager Interrupt Mask Register \(BM_IMR\)](#) for the Buffer Manager, [Switch Engine Interrupt Mask Register \(SWE_IMR\)](#) for the Switch Engine and/or [Port x MAC Interrupt Mask Register \(MAC_IMR_x\)](#) for the Port 2,1,0 MACs)
- The desired Switch Fabric sub-module interrupt event must be enabled in the [Switch Global Interrupt Mask Register \(SW_IMR\)](#)
- The [Switch Engine Interrupt Event Enable \(SWITCH_INT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must be set
- The **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#)

8.2.3 ETHERNET PHY INTERRUPTS

The Ethernet PHYs each provide a set of identical interrupt sources. The top-level [Physical PHY A Interrupt Event \(PHY_INT_A\)](#) and [Physical PHY B Interrupt Event \(PHY_INT_B\)](#) bits of the [Interrupt Status Register \(INT_STS\)](#) provide indication that a PHY interrupt event occurred in the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#).

PHY interrupts are enabled/disabled via their respective [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#). The source of a PHY interrupt can be determined and cleared via the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#). Unique interrupts are generated based on the following events:

- ENERGYON Activated
- Auto-Negotiation Complete
- Remote Fault Detected
- Link Down (Link Status Negated)
- Link Up (Link Status Asserted)
- Auto-Negotiation LP Acknowledge
- Parallel Detection Fault
- Auto-Negotiation Page Received
- Wake-on-LAN Event Detected

In order for an interrupt event to trigger the external **IRQ** interrupt pin, the desired PHY interrupt event must be enabled in the corresponding **PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)**, the **Physical PHY A Interrupt Event Enable (PHY_INT_A_EN)** and/or **Physical PHY B Interrupt Event Enable (PHY_INT_B_EN)** bits of the **Interrupt Enable Register (INT_EN)** must be set and the **IRQ** output must be enabled via the **IRQ Enable (IRQ_EN)** bit of the **Interrupt Configuration Register (IRQ_CFG)**.

For additional details on the Ethernet PHY interrupts, refer to [Section 12.2.9, "PHY Interrupts," on page 228](#).

8.2.4 GPIO INTERRUPTS

Each GPIO of the device is provided with its own interrupt. The top-level **GPIO Interrupt Event (GPIO)** bit of the **Interrupt Status Register (INT_STS)** provides indication that a GPIO interrupt event occurred in the **General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)**. The **General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)** provides enabling/disabling and status of each GPIO interrupt.

In order for a GPIO interrupt event to trigger the external **IRQ** interrupt pin, the desired GPIO interrupt must be enabled in the **General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)**, the **GPIO Interrupt Event Enable (GPIO_EN)** bit of the **Interrupt Enable Register (INT_EN)** must be set and the **IRQ** output must be enabled via the **IRQ Enable (IRQ_EN)** bit of the **Interrupt Configuration Register (IRQ_CFG)**.

For additional details on the GPIO interrupts, refer to [Section 17.2.1, "GPIO Interrupts," on page 563](#).

8.2.5 HOST MAC INTERRUPTS

The top-level **Interrupt Status Register (INT_STS)** and **Interrupt Enable Register (INT_EN)** provide the status and enabling/disabling of multiple Host MAC related interrupts. All Host MAC interrupts are monitored and configured directly within these two registers. The following Host MAC related interrupt events are supported:

- TX Stopped
- RX Stopped
- RX Dropped Frame Counter Halfway
- TX IOC
- RX DMA
- TX Status FIFO Overflow
- Receive Watchdog Time-Out
- Receiver Error
- Transmitter Error
- TX Data FIFO Overrun
- TX Data FIFO Available
- TX Status FIFO Full
- TX Status FIFO Level
- RX Dropped Frame
- RX Status FIFO Full
- RX Status FIFO Level

In order for a Host MAC interrupt event to trigger the external **IRQ** interrupt pin, the desired Host MAC interrupt event must be enabled in the **Interrupt Enable Register (INT_EN)** and the **IRQ** output must be enabled via the **IRQ Enable (IRQ_EN)** bit of the **Interrupt Configuration Register (IRQ_CFG)**.

Refer to the [Interrupt Status Register \(INT_STS\)](#) on page 81 and [Section 11.0, "Host MAC," on page 152](#) for additional information on bit definitions and Host MAC operation.

8.2.6 POWER MANAGEMENT INTERRUPTS

Multiple Power Management Event interrupt sources are provided by the device. The top-level [Power Management Interrupt Event \(PME_INT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) provides indication that a Power Management interrupt event occurred in the [Power Management Control Register \(PMT_CTRL\)](#).

The [Power Management Control Register \(PMT_CTRL\)](#) provides enabling/disabling and status of all Power Management conditions. These include energy-detect on the PHYs and Wake-On-LAN (Perfect DA, Broadcast, Wake-up frame or Magic Packet) detection by the Host MAC and PHYs A&B.

In order for a Power Management interrupt event to trigger the external **IRQ** interrupt pin, the desired Power Management interrupt event must be enabled in the [Power Management Control Register \(PMT_CTRL\)](#), the [Power Management Event Interrupt Enable \(PME_INT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must be set and the **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit 8 of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

The power management interrupts are only a portion of the power management features of the device. For additional details on power management, refer to [Section 6.3, "Power Management," on page 49](#).

8.2.7 GENERAL PURPOSE TIMER INTERRUPT

A [GP Timer \(GPT_INT\)](#) interrupt is provided in the top-level [Interrupt Status Register \(INT_STS\)](#) and [Interrupt Enable Register \(INT_EN\)](#). This interrupt is issued when the [General Purpose Timer Count Register \(GPT_CNT\)](#) wraps past zero to FFFFh and is cleared when the [GP Timer \(GPT_INT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) is written with 1.

In order for a General Purpose Timer interrupt event to trigger the external **IRQ** interrupt pin, the GPT must be enabled via the [General Purpose Timer Enable \(TIMER_EN\)](#) bit in the [General Purpose Timer Configuration Register \(GPT_CFG\)](#), the [GP Timer Interrupt Enable \(GPT_INT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must be set and the **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

For additional details on the General Purpose Timer, refer to [Section 16.1, "General Purpose Timer," on page 559](#).

8.2.8 SOFTWARE INTERRUPT

A general purpose software interrupt is provided in the top level [Interrupt Status Register \(INT_STS\)](#) and [Interrupt Enable Register \(INT_EN\)](#). The [Software Interrupt \(SW_INT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) is generated when the [Software Interrupt Enable \(SW_INT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) changes from cleared to set (i.e. on the rising edge of the enable). This interrupt provides an easy way for software to generate an interrupt and is designed for general software usage.

In order for a Software interrupt event to trigger the external **IRQ** interrupt pin, the **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

8.2.9 DEVICE READY INTERRUPT

A device ready interrupt is provided in the top-level [Interrupt Status Register \(INT_STS\)](#) and [Interrupt Enable Register \(INT_EN\)](#). The [Device Ready \(READY\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#) indicates that the device is ready to be accessed after a power-up or reset condition. Writing a 1 to this bit in the [Interrupt Status Register \(INT_STS\)](#) will clear it.

In order for a device ready interrupt event to trigger the external **IRQ** interrupt pin, the [Device Ready Enable \(READY_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must be set and the **IRQ** output must be enabled via the [IRQ Enable \(IRQ_EN\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

8.2.10 CLOCK OUTPUT TEST MODE

In order to facilitate system level debug, the crystal clock can be enabled onto the **IRQ** pin by setting the [IRQ Clock Select \(IRQ_CLK_SELECT\)](#) bit of the [Interrupt Configuration Register \(IRQ_CFG\)](#).

The **IRQ** pin should be set to a push-pull driver by using the [IRQ Buffer Type \(IRQ_TYPE\)](#) bit for the best result.

8.3 Interrupt Registers

This section details the directly addressable interrupt related System CSRs. These registers control, configure and monitor the **IRQ** interrupt output pin and the various device interrupt sources. For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map," on page 33](#).

TABLE 8-1: INTERRUPT REGISTERS

ADDRESS	REGISTER NAME (SYMBOL)
054h	Interrupt Configuration Register (IRQ_CFG)
058h	Interrupt Status Register (INT_STS)
05Ch	Interrupt Enable Register (INT_EN)

8.3.1 INTERRUPT CONFIGURATION REGISTER (IRQ_CFG)

Offset: 054h Size: 32 bits

This read/write register configures and indicates the state of the IRQ signal.

Bits	Description	Type	Default
31:24	Interrupt De-assertion Interval (INT_DEAS) This field determines the Interrupt Request De-assertion Interval in multiples of 10 microseconds. Setting this field to zero causes the device to disable the INT_DEAS Interval, reset the interval counter and issue any pending interrupts. If a new, non-zero value is written to this field, any subsequent interrupts will obey the new setting. This field does not apply to the PME_INT interrupt.	R/W	00h
23:15	RESERVED	RO	-
14	Interrupt De-assertion Interval Clear (INT_DEAS_CLR) Writing a 1 to this register clears the de-assertion counter in the Interrupt Controller, thus causing a new de-assertion interval to begin (regardless of whether or not the Interrupt Controller is currently in an active de-assertion interval). 0: Normal operation 1: Clear de-assertion counter	R/W SC	0h
13	Interrupt De-assertion Status (INT_DEAS_STS) When set, this bit indicates that the interrupt controller is currently in a de-assertion interval and potential interrupts will not be sent to the IRQ pin. When this bit is clear, the interrupt controller is not currently in a de-assertion interval and interrupts will be sent to the IRQ pin. 0: Interrupt controller not in de-assertion interval 1: Interrupt controller in de-assertion interval	RO	0b
12	Master Interrupt (IRQ_INT) This read-only bit indicates the state of the internal IRQ line, regardless of the setting of the IRQ_EN bit, or the state of the interrupt de-assertion function. When this bit is set, one of the enabled interrupts is currently active. 0: No enabled interrupts active 1: One or more enabled interrupts active	RO	0b
11:9	RESERVED	RO	-
8	IRQ Enable (IRQ_EN) This bit controls the final interrupt output to the IRQ pin. When clear, the IRQ output is disabled and permanently de-asserted. This bit has no effect on any internal interrupt status bits. 0: Disable output on IRQ pin 1: Enable output on IRQ pin	R/W	0b
7:5	RESERVED	RO	-

LAN9352

Bits	Description	Type	Default
4	IRQ Polarity (IRQ_POL) When cleared, this bit enables the IRQ line to function as an active low output. When set, the IRQ output is active high. When the IRQ is configured as an open-drain output (via the IRQ_TYPE bit), this bit is ignored and the interrupt is always active low. 0: IRQ active low output 1: IRQ active high output	R/W NASR Note 1	0b
3:2	RESERVED	RO	-
1	IRQ Clock Select (IRQ_CLK_SELECT) When this bit is set, the crystal clock may be output on the IRQ pin. This is intended to be used for system debug purposes in order to observe the clock and not for any functional purpose. Note: When using this bit, the IRQ pin should be set to a push-pull driver.	R/W	0b
0	IRQ Buffer Type (IRQ_TYPE) When this bit is cleared, the IRQ pin functions as an open-drain output for use in a wired-or interrupt configuration. When set, the IRQ is a push-pull driver. Note: When configured as an open-drain output, the IRQ_POL bit is ignored and the interrupt output is always active low. 0: IRQ pin open-drain output 1: IRQ pin push-pull driver	R/W NASR Note 1	0b

Note 1: Register bits designated as NASR are not reset when the DIGITAL_RST bit in the [Reset Control Register \(RESET_CTL\)](#) is set.

8.3.2 INTERRUPT STATUS REGISTER (INT_STS)

Offset: 058h Size: 32 bits

This register contains the current status of the generated interrupts. A value of 1 indicates the corresponding interrupt conditions have been met, while a value of 0 indicates the interrupt conditions have not been met. The bits of this register reflect the status of the interrupt source regardless of whether the source has been enabled as an interrupt in the [Interrupt Enable Register \(INT_EN\)](#). Where indicated as R/WC, writing a 1 to the corresponding bits acknowledges and clears the interrupt.

Bits	Description	Type	Default
31	Software Interrupt (SW_INT) This interrupt is generated when the Software Interrupt Enable (SW_INT_EN) bit of the Interrupt Enable Register (INT_EN) is set high. Writing a one clears this interrupt.	R/WC	0b
30	Device Ready (READY) This interrupt indicates that the device is ready to be accessed after a power-up or reset condition.	R/WC	0b
29	1588 Interrupt Event (1588_EVNT) This bit indicates an interrupt event from the IEEE 1588 module. This bit should be used in conjunction with the 1588 Interrupt Status Register (1588_INT_STS) to determine the source of the interrupt event within the 1588 module.	RO	0b
28	Switch Fabric Interrupt Event (SWITCH_INT) This bit indicates an interrupt event from the Switch Fabric. This bit should be used in conjunction with the Switch Global Interrupt Pending Register (SW_IPR) to determine the source of the interrupt event within the Switch Fabric.	RO	0b
27	Physical PHY B Interrupt Event (PHY_INT_B) This bit indicates an interrupt event from the Physical PHY B. The source of the interrupt can be determined by polling the PHY x Interrupt Source Flags Register (PHY_INTERRUPT_SOURCE_x) .	RO	0b
26	Physical PHY A Interrupt Event (PHY_INT_A) This bit indicates an interrupt event from the Physical PHY A. The source of the interrupt can be determined by polling the PHY x Interrupt Source Flags Register (PHY_INTERRUPT_SOURCE_x) .	RO	0b
25	TX Stopped (TXSTOP_INT) This interrupt is issued when the Stop Transmitter (STOP_TX) bit in Transmit Configuration Register (TX_CFG) is set and the Host MAC transmitter is halted.	R/WC	0b
24	RX Stopped (RXSTOP_INT) This interrupt is issued when the Receiver Enable (RXEN) bit in Host MAC Control Register (HMAC_CR) is cleared and the Host MAC receiver is halted.	R/WC	0b
23	RX Dropped Frame Counter Halfway (RXDFH_INT) This interrupt is issued when the Host MAC RX Dropped Frames Counter Register (RX_DROP) counts past its halfway point (7FFFFFFh to 80000000h).	R/WC	0b
22	RESERVED	RO	-

Bits	Description	Type	Default
21	TX IOC Interrupt (TX_IOC) This interrupt is generated when a buffer with the IOC flag set has been fully loaded into the TX Data FIFO.	R/WC	0b
20	RX DMA Interrupt (RXD_INT) This interrupt is issued when the amount of data programmed in the RX DMA Count (RX_DMA_CNT) field of the Receive Configuration Register (RX_CFG) has been transferred out of the RX Data FIFO.	R/WC	0b
19	GP Timer (GPT_INT) This interrupt is issued when the General Purpose Timer Count Register (GPT_CNT) wraps past zero to FFFFh.	R/WC	0b
18	RESERVED	RO	-
17	Power Management Interrupt Event (PME_INT) This interrupt is issued when a Power Management Event is detected as configured in the Power Management Control Register (PMT_CTRL) . This interrupt functions independent of the PME signal and will still function if the PME signal is disabled. Writing a '1' clears this bit regardless of the state of the PME hardware signal. In order to clear this bit, all unmasked bits in the Power Management Control Register (PMT_CTRL) must first be cleared. Note: The Interrupt De-assertion interval does not apply to the PME interrupt.	R/WC	0b
16	TX Status FIFO Overflow (TXSO) This interrupt is generated when the TX Status FIFO overflows.	R/WC	0b
15	Receive Watchdog Time-out (RWT) This interrupt is generated when a frame greater than or equal to 2048 bytes has been received by the Host MAC. Frames greater than or equal to 2049 bytes are truncated to 2048 bytes. Note: This can occur when the switch engine adds a tag to a non-tagged jumbo packet that is originally greater than or equal to 2044 bytes.	R/WC	0b
14	Receiver Error (RXE) Indicates that the Host MAC receiver has encountered an error. Please refer to Section 11.12.5, "Receiver Errors," on page 187 for a description of the conditions that will cause an RXE.	R/WC	0b
13	Transmitter Error (TXE) When generated, indicates that the Host MAC transmitter has encountered an error. Please refer to Section 11.11.7, "Transmitter Errors," on page 183 for a description of the conditions that will cause a TXE.	R/WC	0b
12	GPIO Interrupt Event (GPIO) This bit indicates an interrupt event from the General Purpose I/O. The source of the interrupt can be determined by polling the General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)	RO	0b
11	RESERVED	RO	-
10	TX Data FIFO Overrun Interrupt (TDFO) This interrupt is generated when the TX Data FIFO is full and another write is attempted.	R/WC	0b

Bits	Description	Type	Default
9	TX Data FIFO Available Interrupt (TDFA) This interrupt is generated when the TX Data FIFO available space is greater than the programmed level in the TX Data Available Level field of the FIFO Level Interrupt Register (FIFO_INT) .	R/WC	0b
8	TX Status FIFO Full Interrupt (TSFF) This interrupt is generated when the TX Status FIFO is full.	R/WC	0b
7	TX Status FIFO Level Interrupt (TSFL) This interrupt is generated when the TX Status FIFO reaches the programmed level in the TX Status Level field of the FIFO Level Interrupt Register (FIFO_INT) .	R/WC	0b
6	RX Dropped Frame Interrupt (RXDF_INT) This interrupt is issued whenever a receive frame is dropped by the Host MAC.	R/WC	0b
5	RESERVED	RO	-
4	RX Status FIFO Full Interrupt (RSFF) This interrupt is generated when the RX Status FIFO is full and another status write is attempted by the device.	R/WC	0b
3	RX Status FIFO Level Interrupt (RSFL) This interrupt is generated when the RX Status FIFO reaches the programmed level in the RX Status Level field of the FIFO Level Interrupt Register (FIFO_INT) .	R/WC	0b
2:1	RESERVED	RO	-
0	RESERVED	RO	-

8.3.3 INTERRUPT ENABLE REGISTER (INT_EN)

Offset: 05Ch Size: 32 bits

This register contains the interrupt enables for the **IRQ** output pin. Writing 1 to any of the bits enables the corresponding interrupt as a source for **IRQ**. Bits in the [Interrupt Status Register \(INT_STS\)](#) register will still reflect the status of the interrupt source regardless of whether the source is enabled as an interrupt in this register (with the exception of [Software Interrupt Enable \(SW_INT_EN\)](#)). For descriptions of each interrupt, refer to the [Interrupt Status Register \(INT_STS\)](#) bits, which mimic the layout of this register.

Bits	Description	Type	Default
31	Software Interrupt Enable (SW_INT_EN)	R/W	0b
30	Device Ready Enable (READY_EN)	R/W	0b
29	1588 Interrupt Event Enable (1588_EVNT_EN)	R/W	0b
28	Switch Engine Interrupt Event Enable (SWITCH_INT_EN)	R/W	0b
27	Physical PHY B Interrupt Event Enable (PHY_INT_B_EN)	R/W	0b
26	Physical PHY A Interrupt Event Enable (PHY_INT_A_EN)	R/W	0b
25	TX Stopped Interrupt Enable (TXSTOP_INT_EN)	R/W	0b
24	RX Stopped Interrupt Enable (RXSTOP_INT_EN)	R/W	0b
23	RX Dropped Frame Counter Halfway Interrupt Enable (RXDFH_INT_EN)	R/W	0b
22	RESERVED	RO	-
21	TX IOC Interrupt Enable (TIOC_INT_EN)	R/W	0b
20	RX DMA Interrupt Enable (RXD_INT_EN)	R/W	0b
19	GP Timer Interrupt Enable (GPT_INT_EN)	R/W	0b
18	RESERVED	RO	-
17	Power Management Event Interrupt Enable (PME_INT_EN)	R/W	0b
16	TX Status FIFO Overflow Interrupt Enable (TXSO_EN)	R/W	0b
15	Receive Watchdog Time-out Interrupt Enable (RWT_INT_EN)	R/W	0b
14	Receiver Error Interrupt Enable (RXE_INT_EN)	R/W	0b
13	Transmitter Error Interrupt Enable (TXE_INT_EN)	R/W	0b
12	GPIO Interrupt Event Enable (GPIO_EN)	R/W	0b
11	RESERVED	RO	-
10	TX Data FIFO Overrun Interrupt Enable (TDFO_EN)	R/W	0b
9	TX Data FIFO Available Interrupt Enable (TDFA_EN)	R/W	0b
8	TX Status FIFO Full Interrupt Enable (TSFF_EN)	R/W	0b
7	TX Status FIFO Level Interrupt Enable (TSFL_EN)	R/W	0b

Bits	Description	Type	Default
6	RX Dropped Frame Interrupt Enable (RXDF_INT_EN)	R/W	0b
5	RESERVED	RO	-
4	RX Status FIFO Full Interrupt Enable (RSFF_EN)	R/W	0b
3	RX Status FIFO Level Interrupt Enable (RSFL_EN)	R/W	0b
2:1	RESERVED	RO	-
0	RESERVED	RO	-

9.0 HOST BUS INTERFACE

9.1 Functional Overview

The Host Bus Interface (HBI) module provides a high-speed asynchronous slave interface that facilitates communication between the device and a host system. The HBI allows access to the System CSRs and internal FIFOs and memories and handles byte swapping based on the endianness select.

The following is an overview of the functions provided by the HBI:

- **Address bus input:** Two addressing modes are supported. These are a multiplexed address / data bus and a demultiplexed address bus with address index register accesses. The mode selection is done through a configuration input.
- **Selectable data bus width:** The host data bus width is selectable. 16 and 8-bit data modes are supported. This selection is done through a configuration input. The HBI performs BYTE and WORD to DWORD assembly on write data and keeps track of the BYTE / WORD count for reads. Individual BYTE access in 16-bit mode is not supported.
- **Selectable read / write control modes:** Two control modes are available. Separate read and write pins or an enable and direction pin. The mode selection is done through a configuration input.
- **Selectable control line polarity:** The polarity of the chip select, read / write and address latch signals is selectable through configuration inputs.
- **Dynamic Endianness control:** The HBI supports the selection of big and little endian host byte ordering based on the endianness signal. This highly flexible interface provides mixed endian access for registers and memory. Depending on the addressing mode of the device, this signal is either configuration register controlled or as part of the strobed address input.
- **Direct FIFO access:** A FIFO direct select signal directs all host write operations to the TX Data FIFO and all host read operations from the RX Data FIFO. Depending on the addressing mode of the device, this signal is either directly provided by the host or is strobed as part of the address input.

When used with the Indexed Addressing mode, burst read access is supported by toggling the lower address bits.

9.2 Read / Write Control Signals

The device supports two distinct read / write signal methods:

- read (**RD**) and write (**WR**) strobes are input on separate pins.
- read and write signals are decoded from an enable input (**ENB**) and a direction input (**RD_WR**).

9.3 Control Line Polarity

The device supports polarity control on the following:

- chip select input (**CS**)
- read strobe (**RD**) / direction input (**RD_WR**)
- write strobe (**WR**) / enable input (**ENB**)
- address latch control (**ALELO** and **ALEHI**)

9.4 Multiplexed Address / Data Mode

In Multiplexed Address / Data mode, the address, FIFO Direct Select and endianness select inputs are shared with the data bus. Two methods are supported, a single phase address, utilizing up to 16 address / data pins and a dual phase address, utilizing only the lower 8 data bits.

9.4.1 ADDRESS LATCH CYCLES

9.4.1.1 Single Phase Address Latching

In Single Phase mode, all address bits, the FIFO Direct Select signal and the endianness select are strobed into the device using the trailing edge of the **ALELO** signal. The address latch is implemented on all 16 address / data pins. In 8-bit data mode, where pins **AD[15:8]** are used exclusively for addressing, it is not necessary to drive these upper address lines with a valid address continually through read and write operations. However, this operation, referred to as Partial Address Multiplexing, is acceptable since the device will never drive these pins.

Qualification of the **ALELO** signal with the **CS** signal is selectable. When qualification is enabled, **CS** must be active during **ALELO** in order to strobe the address inputs. When qualification is not enabled, **CS** is a don't care during the address phase.

The address is retained for all future read and write operations. It is retained until either a reset event occurs or a new address is loaded. This allows multiple read and write requests to take place to the same address, without requiring multiple address latching operations.

9.4.1.2 Dual Phase Address Latching

In Dual Phase mode, the lower 8 address bits are strobed into the device using the inactive going edge of the **ALELO** signal and the remaining upper address bits, the FIFO Direct Select signals and the endianness select are strobed into the device using the trailing edge of the **ALEHI** signal. The strobes can be in either order. In 8-bit data mode, pins **AD[15:8]** are not used. In 16-bit data mode, pins **D[15:8]** are used only for data.

Qualification of the **ALELO** and **ALEHI** signals with the **CS** signal is selectable. When qualification is enabled, **CS** must be active during **ALELO** and **ALEHI** in order to strobe the address inputs. When qualification is not enabled, **CS** is a don't care during the address phase.

The address is retained for all future read and write operations. It is retained until either a reset event occurs or a new address is loaded. This allows multiple read and write requests to take place to the same address, without requiring multiple address latching operations.

9.4.1.3 Address Bit to Address / Data Pin Mapping

In 8-bit data mode, address bit 0 is multiplexed onto pin **AD[0]**, address bit 1 onto pin **AD[1]**, etc. The highest address bit is bit 9 and is multiplexed onto pin **AD[9]** (single phase) or **AD[1]** (dual phase). The address latched into the device is considered a BYTE address and covers 1K bytes (0 to 3FFh).

In 16-bit data mode, address bit 1 is multiplexed onto pin **AD[0]**, address bit 2 onto pin **AD[1]**, etc. The highest address bit is bit 9 and is multiplexed onto pin **AD[8]** (single phase) or **AD[0]** (dual phase). The address latched into the device is considered a WORD address and covers 512 words (0 to 1FFh).

When the address is sent to the rest of the device, it is converted to a BYTE address.

9.4.1.4 Endianness Select to Address / Data Pin Mapping

The endianness select is included into the multiplexed address to allow the host system to dynamically select the endianness based on the memory address used. This allows for mixed endian access for registers and memory.

The endianness selection is multiplexed to the data pin one bit above the last address bit.

9.4.1.5 FIFO Direct Select to Address / Data Pin Mapping

The FIFO Direct Select signal is included into the multiplexed address to allow the host system to address the TX and RX Data FIFOs as if they were a large flat address space.

The FIFO Direct Select signal is multiplexed to the data pin two bits above the last address bit.

9.4.2 DATA CYCLES

The host data bus can be 16 or 8-bits wide while all internal registers are 32 bits wide. The Host Bus Interface performs the conversion from WORDs or BYTEs to DWORD, while in 8 or 16-bit data mode. Two or four contiguous accesses within the same DWORD are required in order to perform a write or read.

9.4.2.1 Write Cycles

A write cycle occurs when **CS** and **WR** are active (or when **ENB** is active with **RD_WR** indicating write). The host address and endianness were already captured during the address latch cycle.

On the trailing edge of the write cycle (either **WR** or **CS** or **ENB** going inactive), the host data is captured into registers in the HBI. Depending on the bus width, either a WORD or a BYTE is captured. For 8 or 16-bit data modes, this functions as the DWORD assembly with the affected WORD or BYTE determined by the lower address inputs. BYTE swapping is also done at this point based on the endianness.

WRITES FOLLOWING INITIALIZATION

Following device initialization, writes from the Host Bus are ignored until after a read cycle is performed.

WRITES DURING AND FOLLOWING POWER MANAGEMENT

LAN9352

During and following any power management mode other than D0, writes from the Host Bus are ignored until after a read cycle is performed.

8 AND 16-BIT ACCESS

While in 8 or 16-bit data mode, the host is required to perform two or four, 16 or 8-bit writes to complete a single DWORD transfer. No ordering requirements exist. The host can access either the low or high WORD or BYTE first, as long as the other write(s) is(are) performed to the remaining WORD or BYTES.

Note: Writing the same WORD or BYTES in the same DWORD assemble cycle may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

A write BYTE / WORD counter keeps track of the number of writes. At the trailing edge of the write cycle, the counter is incremented. Once all writes occur, a 32-bit write is performed to the internal register.

The write BYTE / WORD counter is reset if the power management mode is set to anything other than D0.

9.4.2.2 Read Cycles

A read cycle occurs when CS and RD are active (or when ENB is active with RD_WR indicating read). The host address and endianness were already captured during the address latch cycle.

At the beginning of the read cycle, the appropriate register is selected and its data is driven onto the data pins. Depending on the bus width, either a WORD or a BYTE is read. For 8 or 16-bit data modes, the returned BYTE or WORD is determined by the endianness and the lower address inputs.

POLLING FOR INITIALIZATION COMPLETE

Before device initialization, the HBI will not return valid data. To determine when the HBI is functional, the [Byte Order Test Register \(BYTE_TEST\)](#) should be polled. Each poll should consist of an address latch cycle(s) and a data cycle. Once the correct pattern is read, the interface can be considered functional. At this point, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) can be polled to determine when the device is fully configured.

READS DURING AND FOLLOWING POWER MANAGEMENT

During any power management mode other than D0, reads from the Host Bus are ignored. If the power management mode changes back to D0 during an active read cycle, the tail end of the read cycle is ignored. Internal registers are not affected and the state of the HBI does not change.

8 AND 16-BIT ACCESS

For certain register accesses, the host is required to perform two or four consecutive 16 or 8-bit reads to complete a single DWORD transfer. No ordering requirements exist. The host can access either the low or high WORD or BYTE first, as long as the other read(s) is(are) performed from the remaining WORD or BYTES.

Note: Reading the same WORD or BYTES from the same DWORD may cause undefined or undesirable operation. The HBI hardware does not protect against this operation. The HBI simply counts that four BYTES have been read.

A read BYTE / WORD counter keeps track of the number of reads. This counter is separate from the write counter above. At the trailing edge of the read cycle, the counter is incremented. On the last read for the DWORD, an internal read is performed to update any Change on Read CSRs.

The read BYTE / WORD counter is reset if the power management mode is set to anything other than D0.

SPECIAL CSR HANDLING

Live Bits

Any register bit that is updated by a H/W event is held at the beginning of the read cycle to prevent it from changing during the read cycle.

Multiple BYTE / WORD Live Registers in 16 or 8-Bit Modes

Some registers have “live” fields or related fields that span across multiple BYTES or WORDs. For 16 and 8-bit data reads, it is possible for the value of these fields to change between host read cycles. In order to prevent reading intermediate values, these registers are locked when the first byte or word is read and unlocked when the last byte or word is read.

The registers are unlocked if the power management mode is set to anything other than D0.

Change on Read Registers and FIFOs

FIFOs or “Change on Read” registers, are updated at the end of the read cycle.

For 16 and 8-bit modes, only one internal read cycle is indicated and occurs for the last byte or word.

Change on Read Live Register Bits

As described above, registers with live bits are held starting at the beginning of the read cycle and those that have multiple bits that span across BYTES or WORDS are also locked for 16 and 8-bit accesses. Although a H/W event that occurs during the hold or lock time would still update the live bit(s), the live bit(s) will be affected (cleared, etc.) at the end of the read cycle and the H/W event would be lost.

In order to prevent this, the individual CSRs defer the H/W event update until after the read or multiple reads.

Register Polling During Reset Or Initialization

Some registers support polling during reset or device initialization to determine when the device is accessible. For these registers, only one read may be performed without the need to read the other WORD or BYTES. The same BYTE or WORD of the register may be re-read repeatedly.

A register that is 16 or 8-bit readable or readable during reset or device initialization, is noted in its register description.

9.4.2.3 Host Endianness

The device supports big and little endian host byte ordering based upon the endianness select that is latched during the address latch cycle. When the endianness select is low, host access is little endian and when high, host access is big endian. In a typical application the endianness select is connected to a high-order address line, making endian selection address-based. This highly flexible interface provides mixed endian access for registers and memory for both PIO and host DMA access.

All internal busses are 32-bit with little endian byte ordering. Logic within the Host Bus Interface re-orders bytes based on the appropriate endianness bit, and the state of the least significant address bits.

Data path operations for the supported endian configurations and data bus sizes are illustrated in [FIGURE 9-1: Little Endian Ordering on page 90](#) and [FIGURE 9-2: Big Endian Ordering on page 91](#).

FIGURE 9-1: LITTLE ENDIAN ORDERING





FIGURE 9-2: BIG ENDIAN ORDERING



9.4.3 TX AND RX FIFO ACCESS

9.4.3.1 TX and RX Status FIFO Peek Address Access

Normal read access to the TX or RX Status FIFO causes the FIFO to advance to its next entry.

For access to the TX and RX Status FIFO Peek addresses, the FIFO does not advance to its next entry.

9.4.3.2 FIFO Direct Select Access

A FIFO Direct Select signal is provided allows the host system to address the TX and RX Data FIFOs as if they were a large flat address space. When the FIFO Direct Select signal, which was latched during the address latch cycle, is active all host write operations are to the TX Data FIFO and all host read operations are from the RX Data FIFO. Only the lower latched address signals are decoded in order to select the proper BYTE or WORD. All other address inputs are ignored in this mode. All other operations are the same (DWORD assembly, FIFO popping, etc.).

The endianness of FIFO Direct Select accesses is determined by the endianness select that was latched during the address latch cycle.

Burst access when reading the RX Data FIFO is not supported. However, since the FIFO Direct Select signal is retained until either a reset event occurs or a new address is loaded, multiple read or write requests can occur without requiring multiple address latching operations.

9.4.4 MULTIPLEXED ADDRESSING MODE FUNCTIONAL TIMING DIAGRAMS

The following timing diagrams illustrate example multiplexed addressing mode read and write cycles for various address/data configurations and bus sizes. These diagrams do not cover every supported host bus permutation, but are selected to detail the main configuration differences (bus size, dual/single phase address latching) within the multiplexed addressing mode of operation.

The following should be noted for the timing diagrams in this section:

- The diagrams in this section depict active-high **ALEHI/ALELO**, **CS**, **RD**, and **WR** signals. The polarities of these signals are selectable via the [HBI_ale_polarity_strap](#), [HBI_cs_polarity_strap](#), [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#), respectively. Refer to [Section 9.3, "Control Line Polarity," on page 86](#) for additional details.
- The diagrams in this section depict little endian byte ordering. However, dynamic big and little endianness are supported via the endianness signal. Endianness changes only the order of the bytes involved, and not the overall timing requirements. Refer to [Section 9.4.1.4, "Endianness Select to Address / Data Pin Mapping," on page 87](#) for additional information.
- The diagrams in [Section 9.4.4.1, "Dual Phase Address Latching"](#) and [Section 9.4.4.2, "Single Phase Address Latching"](#) utilize **RD** and **WR** signals. Alternative **RD_WR** and **ENB** signaling is also supported, as shown in [Section 9.4.4.3, "RD_WR / ENB Control Mode Examples"](#). The HBI read/write mode is selectable via the [HBI_rw_mode_strap](#). The polarities of the **RD_WR** and **ENB** signals are selectable via the [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#).
- Qualification of the **ALELO** and/or **ALEHI** with the **CS** signal is selectable via the [HBI_ale_qualification_strap](#). Refer to [Section 9.4.1.1, "Single Phase Address Latching," on page 86](#) and [Section 9.4.1.2, "Dual Phase Address Latching," on page 87](#) for additional information.
- In dual phase address latching mode, the **ALEHI** and **ALELO** cycles can be in any order. Either or both **ALELO** and **ALEHI** cycles maybe skipped and the device retains the last latched address.
- In single phase address latching mode, the **ALELO** cycle maybe skipped and the device retains the last latched address.

Note:	In 8 and 16-bit modes, the ALELO cycle is normally not skipped since sequential BYTES or WORDs are accessed in order to satisfy a complete DWORD cycle. However, there are registers for which a single BYTE or WORD access is allowed, in which case multiple accesses to these registers may be performed without the need to re-latch the repeated address.
--------------	---

- For 16 and 8-bit modes, consecutive address cycles must be within the same DWORD until the DWORD is completely accessed (with the register exceptions noted above). Although BYTES and WORDs can be accessed in any order, the diagrams in this section depict accessing the lower address BYTE or WORD first.


9.4.4.1 Dual Phase Address Latching

The figures in this section detail read and write operations in multiplexed addressing mode with dual phase address latching for 16 and 8-bit modes.

16-BIT READ

The address is latched sequentially from AD[7:0]. AD[15:8] is not used or driven for the address phase. A read on AD[15:0] follows. The cycle is repeated for the other 16-bits of the DWORD.


FIGURE 9-3: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 16-BIT READ



16-BIT READ WITH SUPPRESSED ALEHI

The address is latched sequentially from AD[7:0]. AD[15:8] is not used or driven for the address phase. A read on AD[15:0] follows. The lower address is then updated to access the opposite WORD.

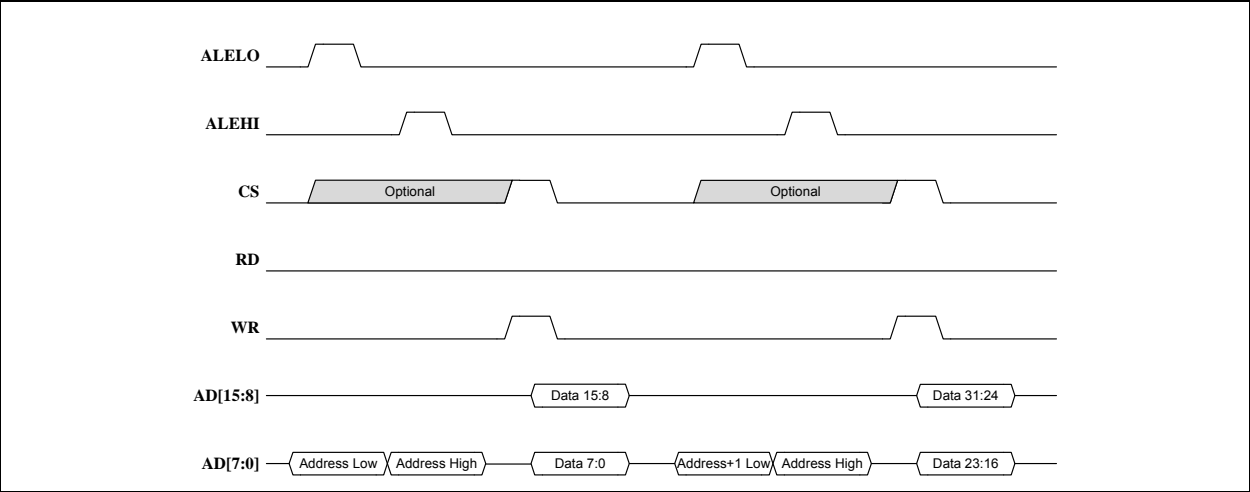
FIGURE 9-4: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 16-BIT READ WITHOUT ALEHI



16-BIT WRITE

The address is latched sequentially from AD[7:0]. AD[15:8] is not used or driven for the address phase. A write on AD[15:0] follows. The cycle is repeated for the other 16-bits of the DWORD.


FIGURE 9-5: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 16-BIT WRITE



16-BIT WRITE WITH SUPPRESSED ALEHI

The address is latched sequentially from AD[7:0]. AD[15:8] is not used or driven for the address phase. A write on AD[15:0] follows. The lower address is then updated to access the opposite WORD.

FIGURE 9-6: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 16-BIT WRITE WITHOUT ALEHI




16-BIT READS AND WRITES TO CONSTANT ADDRESS

The address is latched sequentially from AD[7:0]. AD[15:8] is not used or driven for the address phase. A mix of reads and writes on AD[15:0] follows.

Note: Generally, two 16-bit reads to opposite WORDs of the same DWORD are required, with at least the lower address changing using ALELO. 16-bit reads and writes to the same WORD is a special case.


FIGURE 9-7: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 16-BIT READS AND WRITES CONSTANT ADDRESS



8-BIT READ

The address is latched sequentially from AD[7:0]. A read on AD[7:0] follows. AD[15:8] pins are not used or driven. The cycle is repeated for the other BYTES of the DWORD.


FIGURE 9-8: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 8-BIT READS



8-BIT READ WITH SUPPRESSED ALEHI

The address is latched sequentially from AD[7:0]. A read on AD[7:0] follows. AD[15:8] pins are not used or driven. The lower address is then updated to access the other BYTES.

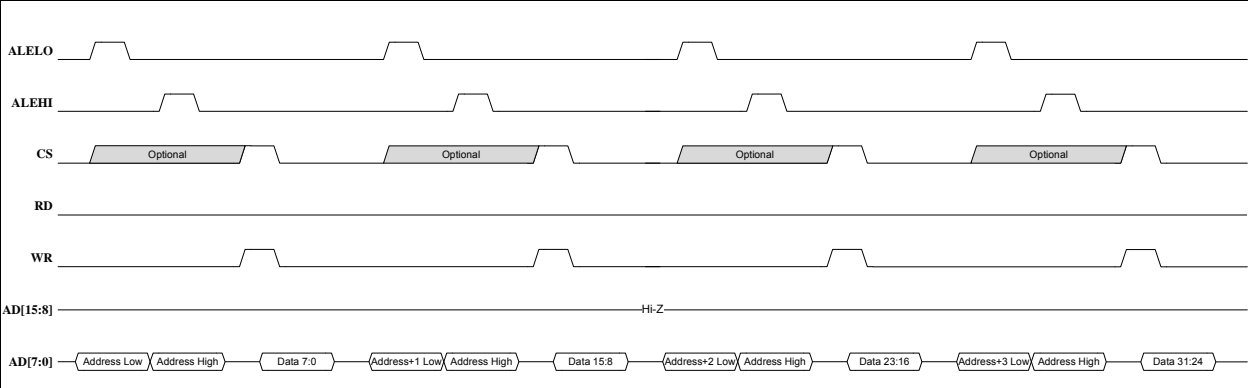
FIGURE 9-9: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 8-BIT READS WITHOUT ALEHI



8-BIT WRITE

The address is latched sequentially from AD[7:0]. A write on AD[7:0] follows. AD[15:8] pins are not used or driven. The cycle is repeated for the other BYTES of the DWORD.


FIGURE 9-10: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 8-BIT WRITE



8-BIT WRITE WITH SUPPRESSED ALEHI

The address is latched sequentially from **AD[7:0]**. A write on **AD[7:0]** follows. **AD[15:8]** pins are not used or driven. The lower address is then updated to access the other BYTES.

FIGURE 9-11: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 8-BIT WRITE WITHOUT ALEHI




8-BIT READS AND WRITES TO CONSTANT ADDRESS

The address is latched sequentially from **AD[7:0]**. A mix of reads and writes on **AD[7:0]** follows. **AD[15:8]** pins are not used or driven.

Note: Generally, four 8-bit reads to opposite BYTES of the same DWORD are required, with at least the lower address changing using **ALELO**. 8-bit reads and writes to the same BYTE is a special case.

FIGURE 9-12: MULTIPLEXED ADDRESSING WITH DUAL PHASE LATCHING - 8-BIT READS AND WRITES CONSTANT ADDRESS




9.4.4.2 Single Phase Address Latching

The figures in this section detail multiplexed addressing mode with single phase addressing for 16 and 8-bit modes of operation.

16-BIT READ

The address is latched simultaneously from **AD[7:0]** and **AD[15:8]**. A read on **AD[15:0]** follows. The cycle is repeated for the other 16-bits of the DWORD.

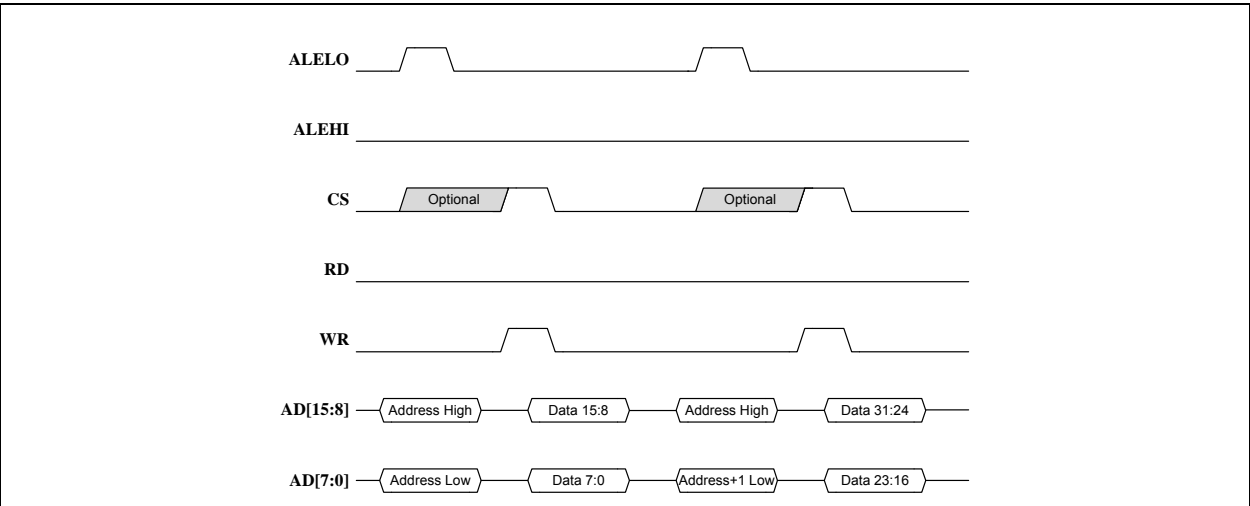
FIGURE 9-13: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 16-BIT READ



16-BIT WRITE

The address is latched simultaneously from **AD[7:0]** and **AD[15:8]**. A write on **AD[15:0]** follows. The cycle is repeated for the other 16-bits of the DWORD.

FIGURE 9-14: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 16-BIT WRITE

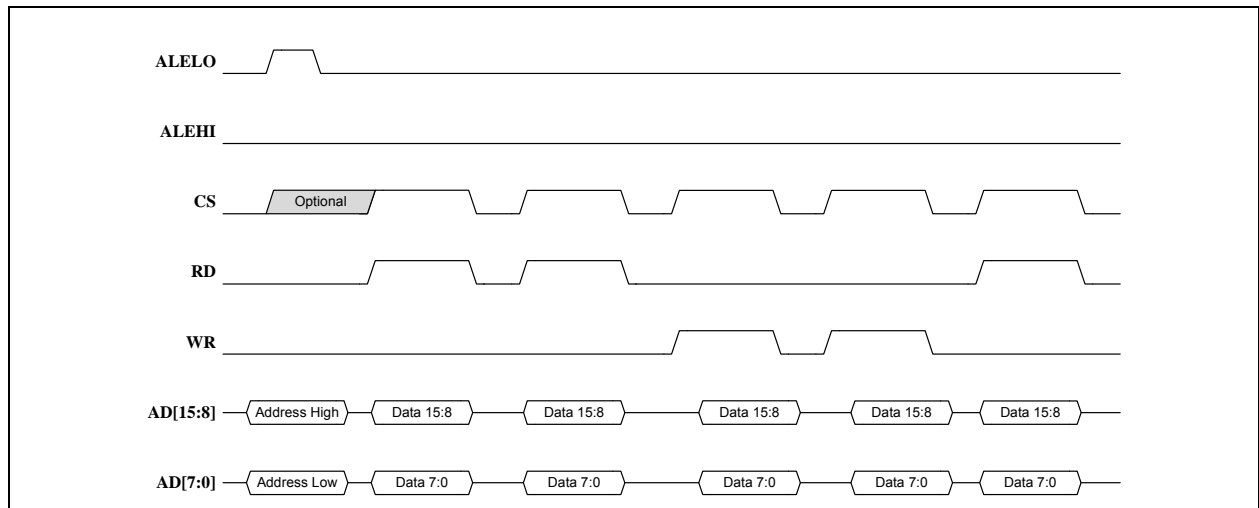


16-BIT READS AND WRITES TO CONSTANT ADDRESS

The address is latched simultaneously from AD[7:0] and AD[15:8]. A mix of reads and writes on AD[15:0] follows.

Note: Generally, two 16-bit reads to opposite WORDs of the same DWORD are required. 16-bit reads and writes to the same WORD is a special case.


FIGURE 9-15: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 16-BIT READS AND WRITES CONSTANT ADDRESS



8-BIT READ

The address is latched simultaneously from AD[7:0] and AD[15:8]. A read on AD[7:0] follows. AD[15:8] pins are not used or driven for the data phase as the host could potentially continue to drive the upper address on these signals. The cycle is repeated for the other BYTES of the DWORD.


FIGURE 9-16: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 8-BIT READ



8-BIT WRITE

The address is latched simultaneously from AD[7:0] and AD[15:8]. A write on AD[7:0] follows. AD[15:8] pins are not used or driven for the data phase as the host could potentially continue to drive the upper address on these signals. The cycle is repeated for the other BYTES of the DWORD.

FIGURE 9-17: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 8-BIT WRITE




8-BIT READS AND WRITES TO CONSTANT ADDRESS

The address is latched simultaneously from AD[7:0] and AD[15:8]. A mix of reads and writes on AD[7:0] follows. AD[15:8] pins are not used or driven for the data phase as the host could potentially continue to drive the upper address on these signals.

Note: Generally, four 8-bit reads to opposite BYTEs of the same DWORD are required. 8-bit reads and writes to the same BYTE is a special case.

FIGURE 9-18: MULTIPLEXED ADDRESSING WITH SINGLE PHASE LATCHING - 8-BIT READS AND WRITES CONSTANT ADDRESS



9.4.4.3 RD_WR / ENB Control Mode Examples

The figures in this section detail read and write operations utilizing the alternative **RD_WR** and **ENB** signaling. The HBI read/write mode is selectable via the [HBI_rw_mode_strap](#).

Note: The examples in this section detail 16-bit mode with dual phase latching. However, the **RD_WR** and **ENB** signaling can be used identically in all other multiplexed addressing modes of operation.

The examples in this section show the **ENB** signal active-high and the **RD_WR** signal low for read and high for write. The polarities of the **RD_WR** and **ENB** signals are selectable via the [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#).

16-BIT

FIGURE 9-19: MULTIPLEXED ADDRESSING RD_WR / ENB CONTROL MODE EXAMPLE - 16-BIT READ





FIGURE 9-20: MULTIPLEXED ADDRESSING RD_WR / ENB CONTROL MODE EXAMPLE - 16-BIT WRITE



The following figures and tables specify the timing requirements during Multiplexed Address / Data mode. Since timing requirements are similar across the multitude of operations (e.g. dual vs. single phase, 8 vs. 16-bit), many timing requirements are illustrated onto the same figures and do not necessarily represent any particular functional operation.

- The diagrams in this section depict active-high **ALEHI/ALELO**, **CS**, **RD**, **WR**, **RD_WR** and **ENB** signals. The polarities of these signals are selectable via the [HBI_ale_polarity_strap](#), [HBI_cs_polarity_strap](#), [HBI_rd_rdwr_polarity_strap](#), and [HBI_wr_en_polarity_strap](#), respectively. Refer to [Section 9.3, "Control Line Polarity,"](#) on page 86 for additional details.
- Qualification of the **ALELO** and/or **ALEHI** with the **CS** signal is selectable via the [HBI_ale_qualification_strap](#). This is shown as a dashed line. Timing requirements between **ALELO** / **ALEHI** and **CS** only apply when this mode is active.
- In dual phase address latching mode, the **ALEHI** and **ALELO** cycles can be in any order. **ALEHI** first is depicted in solid line. **ALELO** first is depicted in dashed line.
- A read cycle maybe followed by followed by an address cycle, a write cycle or another read cycle. A write cycle maybe followed by followed by a read cycle or another write cycle. These are shown in dashed line.

If **RD** and **WR** signaling is used, a host read cycle begins when **RD** is asserted with **CS** active. The cycle ends when **RD** is de-asserted. **CS** maybe asserted and de-asserted along with **RD** but not during **RD** active.

Please refer to [Section 9.4.4, "Multiplexed Addressing Mode Functional Timing Diagrams,"](#) on page 92 for functional descriptions.

TABLE 9-1: MULTIPLEXED ADDRESSING READ CYCLE TIMING VALUES

Symbol	Description	Min	Typ	max	units
t_{csale}	CS Setup to ALELO, ALEHI Active Note 3 , Note 2	0			ns
t_{csrd}	CS Setup to RD or ENB Active	0			ns
t_{rdcs}	CS Hold from RD or ENB Inactive	0			ns
t_{wale}	ALELO, ALEHI Pulse Width	10			ns
t_{adrs}	Address Setup to ALELO, ALEHI Inactive	10			ns
t_{adrh}	Address Hold from ALELO, ALEHI Inactive	5			ns
t_{aleale}	ALELO Inactive to ALEHI Active ALEHI Inactive to ALELO Active Note 1 , Note 2	0			ns
t_{alerd}	ALELO, ALEHI Inactive to RD or ENB Active Note 2	5			ns
t_{rdwrs}	RD_WR Setup to ENB Active Note 4	5			ns
t_{rdwrh}	RD_WR Hold from ENB Inactive Note 4	5			ns
t_{rdon}	RD or ENB to Data Buffer Turn On	0			ns
t_{rddv}	RD or ENB Active to Data Valid			30	ns
t_{rddh}	Data Output Hold Time from RD or ENB Inactive	0			ns
t_{rddz}	Data Buffer Turn Off Time from RD or ENB Inactive			9	ns
t_{cson}	CS to Data Buffer Turn On	0			ns
t_{csdv}	CS Active to Data Valid			30	ns
t_{csdh}	Data Output Hold Time from CS Inactive	0			ns
t_{csdz}	Data Buffer Turn Off Time from CS Inactive			9	ns
t_{aledv}	ALELO, ALEHI Inactive to Data Valid Note 2			35	ns
t_{rd}	RD or ENB Active Time	32			ns
t_{rdcyc}	RD or ENB Cycle Time	45			ns
t_{rdale}	RD or ENB De-assertion Time before Address Phase	13			ns
t_{rdrd}	RD or ENB De-assertion Time before Next RD or ENB Note 5	13			ns
t_{rdwr}	RD De-assertion Time before Next WR Note 5 , Note 6	13			ns

Note 1: Dual Phase Addressing

Note 2: Depends on ALEHI / ALELO order.

LAN9352

Note 3: ALELO and/or ALEHI qualified with the CS.

Note 4: RD_WR and ENB signaling.

Note 5: No interposed address phase.

Note 6: RD and WR signaling.

Note: Timing values are with respect to an equivalent test load of 25 pF.

9.4.5.2 Write Timing Requirements

If RD and WR signaling is used, a host write cycle begins when WR is asserted with CS active. The cycle ends when WR is de-asserted. CS may be asserted and de-asserted along with WR but not during WR active.

Alternatively, if RD_WR and ENB signaling is used, a host write cycle begins when ENB is asserted with CS active and RD_WR indicating a write. The cycle ends when ENB is de-asserted. CS may be asserted and de-asserted along with ENB but not during ENB active.

Please refer to [Section 9.4.4, "Multiplexed Addressing Mode Functional Timing Diagrams," on page 92](#) for functional descriptions.

FIGURE 9-22: MULTIPLEXED ADDRESSING WRITE CYCLE TIMING




TABLE 9-2: MULTIPLEXED ADDRESSING WRITE CYCLE TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{csale}	CS Setup to ALELO , ALEHI Active Note 9 , Note 8	0			ns
t_{cswr}	CS Setup to WR or ENB Active	0			ns
$t_{wr cs}$	CS Hold from WR or ENB Inactive	0			ns
t_{wale}	ALELO , ALEHI Pulse Width	10			ns
t_{adrs}	Address Setup to ALELO , ALEHI Inactive	10			ns
t_{adrh}	Address Hold from ALELO , ALEHI Inactive	5			ns
t_{aleale}	ALELO Inactive to ALEHI Active ALEHI Inactive to ALELO Active Note 7 , Note 8	0			ns
t_{alewr}	ALELO , ALEHI Inactive to WR or ENB Active Note 8	5			ns
$t_{rd wrs}$	RD_WR Setup to ENB Active Note 10	5			ns
$t_{rd wrh}$	RD_WR Hold from ENB Inactive Note 10	5			ns
t_{ds}	Data Setup to WR or ENB Inactive	7			ns
t_{dh}	Data Hold from WR or ENB Inactive	0			ns
t_{wr}	WR or ENB Active Time	32			ns
t_{wrcyc}	WR or ENB Cycle Time	45			ns
$t_{wr ale}$	WR or ENB De-assertion Time before Address Phase	13			ns
$t_{wr wr}$	WR or ENB De-assertion Time before Next WR or ENB Note 11	13			ns
$t_{wr rd}$	WR De-assertion Time before Next RD Note 11 , Note 12	13			ns

Note 7: Dual Phase Addressing

Note 8: Depends on **ALEHI** / **ALELO** order.

Note 9: **ALELO** and/or **ALEHI** qualified with the CS.

Note 10: **RD_WR** and **ENB** signaling.

Note 11: No interposed address phase.

Note 12: **RD** and **WR** signaling.

9.5 Indexed Address Mode

In Indexed Address mode, access to the internal registers and memory of the device are indirectly mapped using Index and Data registers. The desired internal address is written into the device at a particular offset. The value written is then used as the internal address when the associate Data register address is accessed. Three Index / Data register sets are provided allowing for multi-threaded operation without the concern of one thread corrupting the Index set by another thread. Endianness can be configured per Index / Data pair. Another Data register is provided for access to the FIFOs.

The host address register map is given below. In 8-bit data mode, the host address input (ADDR[4:0]) is a BYTE address. In 16-bit data mode, ADDR0 is not provided and the host address input (ADDR[4:1]) is a WORD address.

As discussed below in [Section 9.5.5.2, "Index Register Bypass FIFO Access"](#), the TX and RX Data FIFOs are accessed when reading or writing at address 18h-1Bh.

As discussed below in [Section 9.5.5.3, "FIFO Direct Select Access"](#), when the FIFOSEL input is active, all access is to or from the TX and RX Data FIFOs.

TABLE 9-3: HOST BUS INTERFACE INDEXED ADDRESS MODE REGISTER MAP

FIFOSEL	BYTE ADDRESS	SYMBOL	REGISTER NAME
0	00h-03h	HBI_IDX_0	Host Bus Interface Index Register 0
0	04h-07h	HBI_DATA_0	Host Bus Interface Data Register 0
0	08h-0Bh	HBI_IDX_1	Host Bus Interface Index Register 1
0	0Ch-0Fh	HBI_DATA_1	Host Bus Interface Data Register 1
0	10h-13h	HBI_IDX_2	Host Bus Interface Index Register 2
0	14h-17h	HBI_DATA_2	Host Bus Interface Data Register 2
0	18h-1Bh	TX_DATA_FIFO RX_DATA_FIFO	TX Data FIFO RX Data FIFO
0	1Ch-1Fh	HBI_CFG	Host Bus Interface Configuration Register
1	na	TX_DATA_FIFO RX_DATA_FIFO	TX Data FIFO RX Data FIFO

9.5.1 HOST BUS INTERFACE INDEX REGISTER

The Index registers are writable as WORDs or as BYTEs, depending upon the data mode. There is no concern about DWORD assembly rules when writing these registers. The Index registers are formatted as follows:

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	Internal Address The address used when the corresponding Data register is accessed. Note: The internal address provided by each Index register is always considered to be a BYTE address.	R/W	1234h Note 13

Note 13: The default may be used to help determine the endianness of the register.

9.5.2 HOST BUS INTERFACE CONFIGURATION REGISTER

The HBI Configuration register is used to specify the endianness of the interface. Endianness for each Index / Data pair and for FIFO accesses can be individually specified.

The endianness of this register is irrelevant since each byte is shadowed into 4 positions.

The HBI Configuration register is writable as WORDs or as BYTEs, depending upon the data mode. There is no concern about DWORD assembly rules when writing this register. The Configuration register is formatted as follows:

Bits	Description	Type	Default
31:28	RESERVED	RO	-
27	FIFO Endianness Shadow 3 This bit is a shadow of bit 3.	R/W	0b
26	Host Bus Interface Index / Data Register 2 Endianness Shadow 3 This bit is a shadow of bit 2.	R/W	0b
25	Host Bus Interface Index / Data Register 1 Endianness Shadow 3 This bit is a shadow of bit 1.	R/W	0b
24	Host Bus Interface Index / Data Register 0 Endianness Shadow 3 This bit is a shadow of bit 0.	R/W	0b
23:20	RESERVED	RO	-
19	FIFO Endianness Shadow 2 This bit is a shadow of bit 3.	R/W	0b
18	Host Bus Interface Index / Data Register 2 Endianness Shadow 2 This bit is a shadow of bit 2.	R/W	0b
17	Host Bus Interface Index / Data Register 1 Endianness Shadow 2 This bit is a shadow of bit 1.	R/W	0b
16	Host Bus Interface Index / Data Register 0 Endianness Shadow 2 This bit is a shadow of bit 0.	R/W	0b
15:12	RESERVED	RO	-
11	FIFO Endianness Shadow 1 This bit is a shadow of bit 3.	R/W	0b
10	Host Bus Interface Index / Data Register 2 Endianness Shadow 1 This bit is a shadow of bit 2.	R/W	0b

Bits	Description	Type	Default
9	Host Bus Interface Index / Data Register 1 Endianness Shadow 1 This bit is a shadow of bit 1.	R/W	0b
8	Host Bus Interface Index / Data Register 0 Endianness Shadow 1 This bit is a shadow of bit 0.	R/W	0b
7:4	RESERVED	RO	-
3	FIFO Endianness This bit specifies the endianness of FIFO accesses when they are accessed by means other than the Index / Data Register method. 0 = Little Endian 1 = Big Endian Note: In order to avoid any ambiguity with the endianness of this register, bits 3, 11, 19 and 27 are shadowed. If any of these bits are set during a write, all of the bits will be set.	R/W	0b
2	Host Bus Interface Index / Data Register 2 Endianness This bit specifies the endianness of the Index and Data register set 2. 0 = Little Endian 1 = Big Endian Note: In order to avoid any ambiguity with the endianness of this register, bits 2, 10, 18 and 26 are shadowed. If any of these bits are set during a write, all of the bits will be set.	R/W	0b
1	Host Bus Interface Index / Data Register 1 Endianness This bit specifies the endianness of the Index and Data register set 1. 0 = Little Endian 1 = Big Endian Note: In order to avoid any ambiguity with the endianness of this register, bits 1, 9, 17 and 25 are shadowed. If any of these bits are set during a write, all of the bits will be set.	R/W	0b
0	Host Bus Interface Index / Data Register 0 Endianness This bit specifies the endianness of the Index and Data register set 0. 0 = Little Endian 1 = Big Endian Note: In order to avoid any ambiguity with the endianness of this register, bits 0, 8, 16 and 24 are shadowed. If any of these bits are set during a write, all of the bits will be set.	R/W	0b

9.5.3 INDEX AND CONFIGURATION REGISTER DATA ACCESS

The host data bus can be 16 or 8-bits wide. The HBI Index registers and the HBI Configuration register are 32-bits wide and are writable as WORDs or as BYTEs, depending upon the data mode. They do not have nor do they require WORDs or BYTEs to DWORD conversion.

9.5.3.1 Write Cycles

A write cycle occurs when CS and WR are active (or when ENB is active with RD_WR indicating write).

On the trailing edge of the write cycle (either WR or CS or ENB going inactive), the host data is captured into the Configuration register or one for the Index registers.

Depending on the bus width, either a WORD or a BYTE is written. The affected WORD or BYTE is determined by the endianness of the register (specified in the [Host Bus Interface Configuration Register](#)) and the lower address inputs. Individual BYTE (in 16-bit data mode) access is not supported.

WRITES FOLLOWING INITIALIZATION

Following device initialization, writes from the Host Bus are ignored until after a read cycle is performed.

WRITES DURING AND FOLLOWING POWER MANAGEMENT

During and following any power management mode other than D0, writes from the Host Bus are ignored until after a read cycle is performed.

9.5.3.2 Read Cycles

A read cycle occurs when CS and RD are active (or when ENB is active with RD_WR indicating read). The host address is used directly from the Host Bus.

At the beginning of the read cycle, the appropriate register is selected and its data is driven onto the data pins. Depending on the bus width, either a WORD or a BYTE is read. For 8 or 16-bit data modes, the returned BYTE or WORD is determined by the endianness of the register (specified in the [Host Bus Interface Configuration Register](#)) and the lower host address inputs.

9.5.4 INTERNAL REGISTER DATA ACCESS

The host data bus can be 16 or 8-bits wide while all internal registers are 32 bits wide. The Host Bus Interface performs the conversion from WORDs or BYTEs to DWORD, while in 8 or 16-bit data mode. Two or four accesses within the same DWORD are required in order to perform a write or read.

Each Data register, along with the FIFO direct address access, has a separate WORD or BYTE to DWORD conversion. Accesses may be mixed among these (and the HBI Index and Configuration registers) without concern of data corruption.

9.5.4.1 Write Cycles

A write cycle occurs when CS and WR are active (or when ENB is active with RD_WR indicating write). The host address from the Host Bus selects the contents of one of the Index registers. The result of this operation is captured on the leading edge of the write cycle.

The host address inputs and the FIFO Direct Select signal from the Host Bus are also captured on the leading edge of the write cycle. These are used to increment the appropriate write BYTE / WORD counter (for 8 or 16-bit data mode described below) as well as to select the correct DWORD assembly register.

On the trailing edge of the write cycle (either WR or CS or ENB going inactive), the host data is captured into one of the Data registers. Depending on the bus width, either a WORD or a BYTE is captured. For 8 or 16-bit data modes, this functions as the DWORD assembly with the affected WORD or BYTE determined by the lower host address inputs. BYTE swapping is also done at this point based on the endianness of the register (specified in the [Host Bus Interface Configuration Register](#)).

Note: There are separate write BYTE / WORD counters and DWORD assembly registers for each of the three Data Registers as well as for FIFO access.

WRITES FOLLOWING INITIALIZATION

Following device initialization, writes from the Host Bus are ignored until after a read cycle is performed.

WRITES DURING AND FOLLOWING POWER MANAGEMENT

During and following any power management mode other than D0, writes from the Host Bus are ignored until after a read cycle is performed.

8 AND 16-BIT ACCESS

While in 8 or 16-bit data mode, the host is required to perform two or four, 16 or 8-bit writes to complete a single DWORD transfer. No ordering requirements exist. The host can access either the low or high WORD or BYTE first, as long as the other write(s) is(are) performed to the remaining WORD or BYTES.

Note: Writing the same WORD or BYTES into the same DWORD may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

Accessing the same internal register using two Index / Data register pairs may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

Mixing reads and writes into the same Data register may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

A write BYTE / WORD counter keeps track of the number of writes. Each Data Register has its own BYTE / WORD counter. At the trailing edge of the write cycle, the appropriate counter (based on the captured host address from above) is incremented. Once all writes occur, a 32-bit write is performed to the internal register selected by the captured address from above. The data that is written is selected from one of the three DWORD assembly registers based on the captured host address from above.

All of the write BYTE / WORD counters are reset if the power management mode is set to anything other than D0.

9.5.4.2 Read Cycles

A read cycle occurs when CS and RD are active (or when ENB is active with RD_WR indicating read). The host address from the Host Bus selects the contents of one of the Index registers. The result of this operation is used to select the internal register to be read and also is captured on the leading edge of the read cycle.

The host address inputs and the FIFO Direct Select signal from the Host Bus are also captured on the leading edge of the read cycle. These are used to increment the appropriate read BYTE / WORD counter (for 8 or 16-bit data mode described below).

At the beginning of the read cycle, the appropriate register is selected and its data is driven onto the data pins. Depending on the bus width, either a WORD or a BYTE is read. For 8 or 16-bit data modes, the returned BYTE or WORD is determined by the endianness of the Data register (specified in the [Host Bus Interface Configuration Register](#)) and the lower host address inputs.

Note: There are separate read BYTE / WORD counters for each of the three Data Registers as well as for FIFO access.

POLLING FOR INITIALIZATION COMPLETE

Before device initialization, the HBI will not return valid data. To determine when the HBI is functional, first the Host Bus Interface Index Register 0 should be polled, then the [Byte Order Test Register \(BYTE_TEST\)](#) should be polled. Once the correct pattern is read, the interface can be considered functional. At this point, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) can be polled to determine when the device is fully configured.

READS DURING AND FOLLOWING POWER MANAGEMENT

During any power management mode other than D0, reads from the Host Bus are ignored. If the power management mode changes back to D0 during an active read cycle, the tail end of the read cycle is ignored. Internal registers are not affected and the state of the HBI does not change.

8 AND 16-BIT ACCESS

For certain register accesses, the host is required to perform two or four consecutive 16 or 8-bit reads to complete a single DWORD transfer. No ordering requirements exist. The host can access either the low or high WORD or BYTE first, as long as the other read(s) is(are) performed from the remaining WORD or BYTES.

Note: Reading the same WORD or BYTES from the same DWORD may cause undefined or undesirable operation. The HBI hardware does not protect against this operation. The HBI simply counts that four BYTES have been read.

Accessing the same internal register using two Index / Data register pairs may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

Mixing reads and writes into the same Data register may cause undefined or undesirable operation. The HBI hardware does not protect against this operation.

A read BYTE / WORD counter keeps track of the number of reads. Each Data Register has its own BYTE / WORD counter. These counters are separate from the write counters above. At the trailing edge of the read cycle, the appropriate counter (based on the captured host address from above) is incremented. On the last read for the DWORD, an internal read is performed to update any Change on Read CSRs.

All of the read BYTE / WORD counters are reset if the power management mode is set to anything other than D0.

SPECIAL CSR HANDLING

Live Bits

Any register bit that is updated by a H/W event is held at the beginning of the read cycle to prevent it from changing during the read cycle.

Multiple BYTE / WORD Live Registers in 16 or 8-Bit Modes

Some internal registers have fields or related fields that span across multiple BYTES or WORDs. For 16 and 8-bit data reads, it is possible that the value of these fields change between host read cycles. In order to prevent reading intermediate values, these registers are locked when the first byte or word is read and unlocked when the last byte or word is read.

The registers are unlocked if the power management mode is set to anything other than D0.

Change on Read Registers and FIFOs

FIFOs or "Change on Read" registers, are updated at the end of the read cycle.

For 16 and 8-bit modes, only one internal read cycle is indicated and occurs for the last byte or word.

Change on Read Live Register Bits

As described above, registers with live bits are held starting at the beginning of the read cycle and those that have multiple bits that span across BYTES or WORDs are also locked for 16 and 8-bit accesses. Although a H/W event that occurs during the hold or lock time would still update the live bit(s), the live bit(s) will be affected (cleared, etc.) at the end of the read cycle and the H/W event would be lost.

In order to prevent this, the individual CSRs defer the H/W event update until after the read or multiple reads.

Registers Polling During Reset or Initialization

Some registers support polling during reset or device initialization to determine when the device is accessible. For these registers, only one read may be performed without the need to read the other WORD or BYTES. The same BYTE or WORD of the register may be re-read repeatedly.

A register that is 16 or 8-bit readable or readable during reset or device initialization, is noted in its register description.

9.5.4.3 Host Endianness

The device supports big and little endian host byte ordering based upon the endianness bits in the [Host Bus Interface Configuration Register](#). When the appropriate endianness bit is low, host access is little endian and when high, host access is big endian. Endianness is specified for each Index / Data pair and for FIFO Direct Select accesses.

All internal busses are 32-bit with little endian byte ordering. Logic within the Host Bus Interface re-orders bytes based on the appropriate endianness bit, and the state of the least significant address lines (ADDR[1:0]).

Data path operations for the supported endian configurations and data bus sizes are illustrated in [FIGURE 9-23: Little Endian Ordering on page 112](#) and [FIGURE 9-24: Big Endian Ordering on page 113](#).

FIGURE 9-23: LITTLE ENDIAN ORDERING





FIGURE 9-24: BIG ENDIAN ORDERING



9.5.5 TX AND RX FIFO ACCESS

9.5.5.1 TX and RX Status FIFO Peek Address Access

Normal read access to the TX or RX Status FIFO causes the FIFO to advance to its next entry.

For access to the TX and RX Status FIFO Peek addresses, FIFO does not advance to its next entry.

9.5.5.2 Index Register Bypass FIFO Access

In addition to the indexed access, the Index Registers can be bypassed and the FIFOs accessed at address 18h-1Bh. At this address, host write operations are to the TX Data FIFO and host read operations are from the RX Data FIFO. There is no associated Index Register.

Index Register Bypass and FIFO Direct Select accesses share the same read and write BYTE / WORD counters and the same write DWORD assembly registers.

The endianness of FIFO accesses using this method is specified by the [FIFO Endianness](#) bit in the [Host Bus Interface Configuration Register](#).

9.5.5.3 FIFO Direct Select Access

In addition to the indexed access, a FIFO Direct Select signal is provided. This allows the host system to access the TX and RX Data FIFOs as if they were a large flat address space. When the **FIFSEL** input is active, all host write operations are to the TX Data FIFO and all host read operations are from the RX Data FIFO. The lower host address signals are decoded in order to select the proper BYTE or WORD and to delimit DWORDs during a burst access.

Burst access is supported when reading the RX Data FIFO. With the **FIFSEL** input active, **CS** and **RD** (or **ENB** with **RD_WR** indicating read) may be left active while the lower address lines toggle. Each change of the lower address bits provides the next WORD or BYTE of data. The HBI performs an internal FIFO pop (read cycle) when it detects that a DWORD boundary has been crossed (**A[2]** has toggled).

The endianness of FIFO Direct Select accesses is specified by the [FIFO Endianness](#) bit in the [Host Bus Interface Configuration Register](#).

9.5.6 INDEXED ADDRESS MODE FUNCTIONAL TIMING DIAGRAMS

The following timing diagrams illustrate example indexed (non-multiplexed) addressing mode read and write cycles for various configurations and bus sizes. These diagrams do not cover every supported host bus permutation, but are selected to detail the main configuration differences (bus size, Configuration/Index/Data/FIFO-Direct cycles) within the indexed addressing mode of operation.

The following should be noted for the timing diagrams in this section:

- The diagrams in this section depict active-high **CS**, **RD**, and **WR** signals. The polarities of these signals are selectable via the [HBI_cs_polarity_strap](#), [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#), respectively. Refer to [Section 9.3, "Control Line Polarity," on page 86](#) for additional details.
- The diagrams in this section depict little endian byte ordering. However, configurable big and little endianness are supported via the endianness bits in the [Host Bus Interface Configuration Register](#). Endianness changes only the order of the bytes involved, and not the overall timing requirements. Refer to [Section 9.5.4.3, "Host Endianness," on page 111](#) for additional information.
- The diagrams in this section utilize **RD** and **WR** signals. Alternative **RD_WR** and **ENB** signaling is also supported, similar to the multiplexed example in [Section 9.4.4.3, "RD_WR / ENB Control Mode Examples"](#). The HBI read/write mode is selectable via the [HBI_rw_mode_strap](#). The polarities of the **RD_WR** and **ENB** signals are selectable via the [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#).
- When accessing internal registers or FIFOs in 16 and 8-bit modes, consecutive address cycles must be within the same DWORD until the DWORD is completely accessed (some internal registers are excluded from this requirement). Although BYTES and WORDs can be accessed in any order, the diagrams in this section depict accessing the lower address BYTE or WORD first.

9.5.6.1 Configuration Register Data Access


The figures in this section detail configuration register read and write operations in indexed address mode for 16 and 8-bit modes.

16-BIT READ AND WRITE

For writes, the address is set to access the lower WORD of the Configuration Register and **FIFOSEL** is held low. Data on **D[15:0]** is written on the trailing edge of **WR**. The cycle repeats for the upper WORD of the Configuration Register, if desired by the host.

For reads, the address is set to access the lower WORD of the Configuration Register and **FIFOSEL** is held low. Read data is driven on **D[15:0]** during **RD** active. The cycle repeats for the upper WORD of the Configuration Register, if desired by the host.

FIGURE 9-25: INDEXED ADDRESSING CONFIGURATION REGISTER ACCESS - 16-BIT WRITE/READ




8-BIT READ AND WRITE

For writes, the address is set to access the lower BYTE of the Configuration Register and **FIFOSEL** is held low. Data on **D[7:0]** is written on the trailing edge of **WR**. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Configuration Register, if desired by the host.

For reads, the address is set to access the lower BYTE of the Configuration Register and **FIFOSEL** is held low. Read data is driven on **D[7:0]** during **RD** active. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Configuration Register, if desired by the host.

FIGURE 9-26: INDEXED ADDRESSING CONFIGURATION REGISTER ACCESS - 8-BIT WRITE/ READ



9.5.6.2 Index Register Data Access

The figures in this section detail index register read and write operations in indexed address mode for 16 and 8-bit modes.


16-BIT READ AND WRITE

For writes, the address is set to access the lower WORD of one of the Index Registers and **FIFOSEL** is held low. Data on **D[15:0]** is written on the trailing edge of **WR**. The cycle repeats for the upper WORD of the Index Register, if desired by the host.

For reads, the address is set to access the lower WORD of one of the Index Registers and **FIFOSEL** is held low. Read data is driven on **D[15:0]** during **RD** active. The cycle repeats for the upper WORD of the Index Register, if desired by the host.

Note: The upper WORD of Index Registers is reserved and don't care. Therefore reads and writes to that WORD are not useful.

FIGURE 9-27: INDEXED ADDRESSING INDEX REGISTER ACCESS - 16-BIT WRITE/READ




8-BIT READ AND WRITE

For writes, the address is set to access the lower BYTE of one of the Index Registers and **FIFOSEL** is held low. Data on **D[7:0]** is written on the trailing edge of **WR**. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Index Register, if desired by the host.

For reads, the address is set to access the lower BYTE of one of the Index Registers and **FIFOSEL** is held low. Read data is driven on **D[7:0]** during **RD** active. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Index Register, if desired by the host.

Note: The upper WORD of Index Registers is reserved and don't care. Therefore reads and writes to those BYTES are not useful.

FIGURE 9-28: INDEXED ADDRESSING INDEX REGISTER ACCESS - 8-BIT WRITE/READ



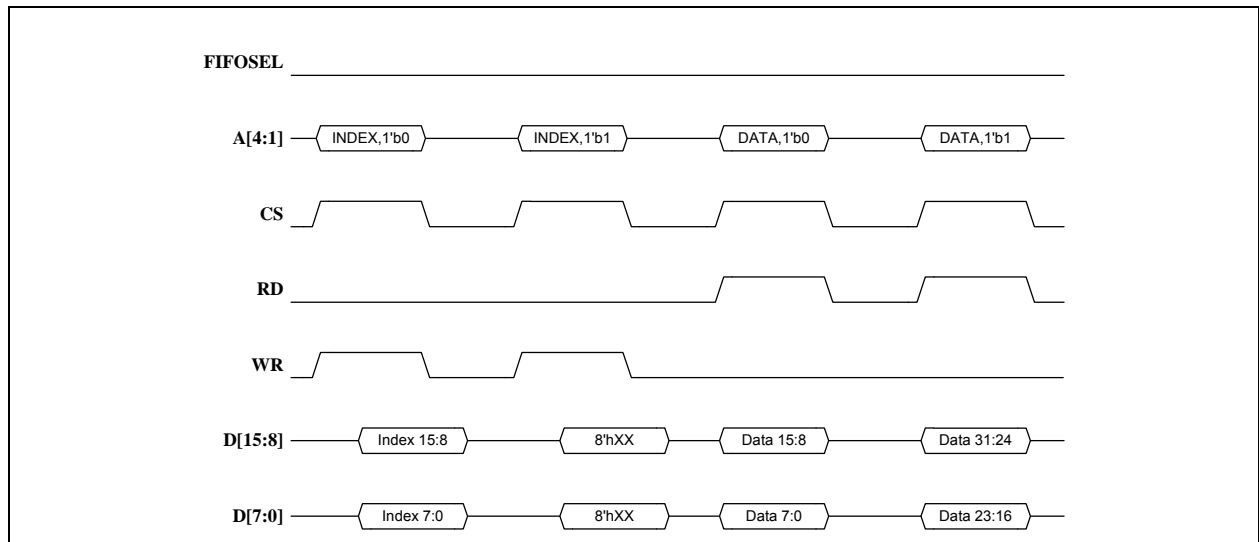
9.5.6.3 Internal Register Data Access

The figures in this section detail typical internal register data read and write cycles in indexed address mode for 16 and 8-bit modes. This includes an index register write followed by either a data read or write.

16-BIT READ

One of the Index Registers is set as described above. The address is then set to access the lower WORD of the corresponding Data Register and **FIFOSEL** is held low. Read data is driven on **D[15:0]** during **RD** active. The cycle repeats for the upper WORD of the Data Register.

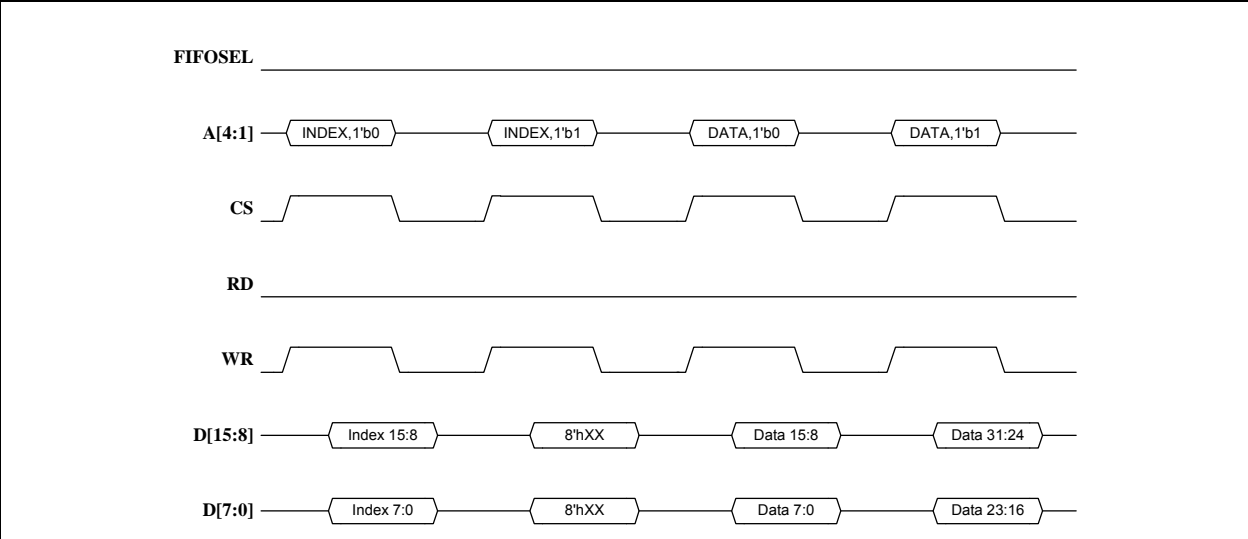
FIGURE 9-29: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 16-BIT READ



16-BIT WRITE

One of the Index Registers is set as described above. The address is then set to access the corresponding Data Register and **FIFOSEL** is held low. Data on **D[15:0]** is written on the trailing edge of **WR**. The cycle repeats for the upper WORD of the Data Register.


FIGURE 9-30: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 16-BIT WRITE



16-BIT READS AND WRITES TO CONSTANT INTERNAL ADDRESS

One of the Index Registers is set as described above. A mix of reads and writes on **D[15:0]** follows, with each read or write consisting of an access to both the lower and upper WORDs of the corresponding Data Register.


FIGURE 9-31: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 16-BIT READS/WRITES CONSTANT ADDRESS



8-BIT READ

One of the Index Registers is set as described above. The address is then set to access the lower BYTE of the corresponding Data Register and **FIFOSEL** is held low. Read data is driven on **D[7:0]** during **RD** active. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Data Register.

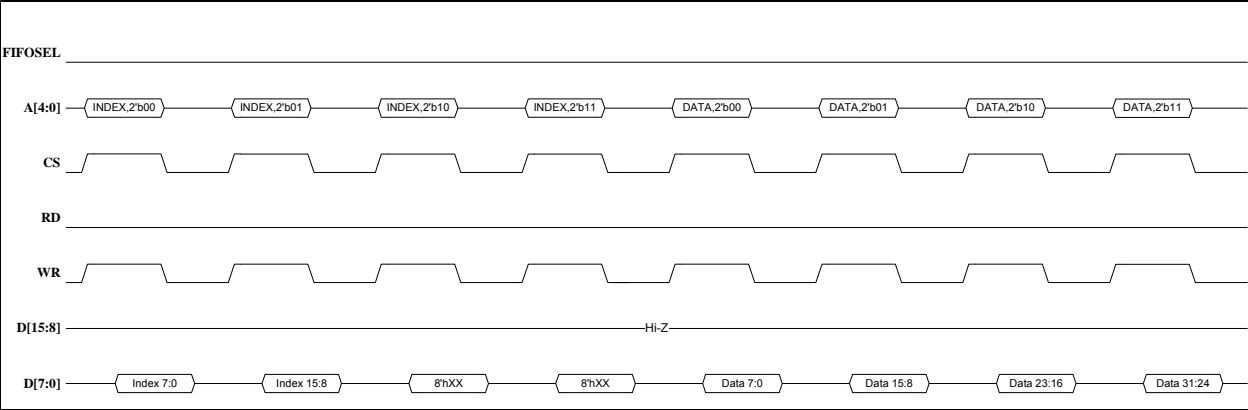
FIGURE 9-32: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 8-BIT READ



8-BIT WRITE

One of the Index Registers is set as described above. The address is then set to access the corresponding Data Register and **FIFSEL** is held low. Data on **D[7:0]** is written on the trailing edge of **WR**. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining BYTES of the Data Register.


FIGURE 9-33: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 8-BIT WRITE



8-BIT READS AND WRITES TO CONSTANT INTERNAL ADDRESS

One of the Index Registers is set as described above. A mix of reads and writes on **D[7:0]** follows, with each read or write consisting of an access to all four BYTES of the corresponding Data Register.

FIGURE 9-34: INDEXED ADDRESSING INTERNAL REGISTER DATA ACCESS - 8-BIT READS/WRITES CONSTANT ADDRESS



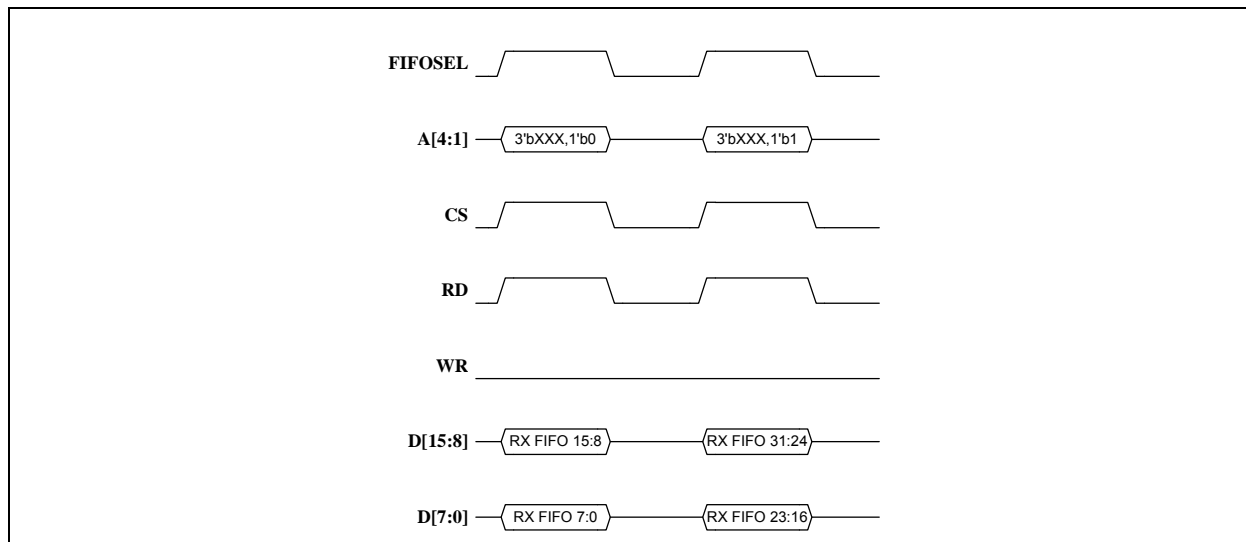
9.5.6.4 FIFO Direct Select Access

The figures in this section detail FIFO Direct Select (**FIFOSEL**) read and write cycles in indexed address mode for 16 and 8-bit modes. Additionally, FIFO Direct Select Burst reads are shown for 16, and 8-bit modes. FIFO Direct Select Burst mode supports up to 8 DWORDs of consecutive reads. FIFO Direct Select (**FIFOSEL**) read and write cycles do not required an index register write.

16-BIT READ

FIFOSEL is set high and address bit 1 is set to access the lower WORD of the FIFO, while address bits 4:2 are don't care. Read data is driven on **D[15:0]** during **RD** active. The cycle repeats for the upper WORD of the FIFO.

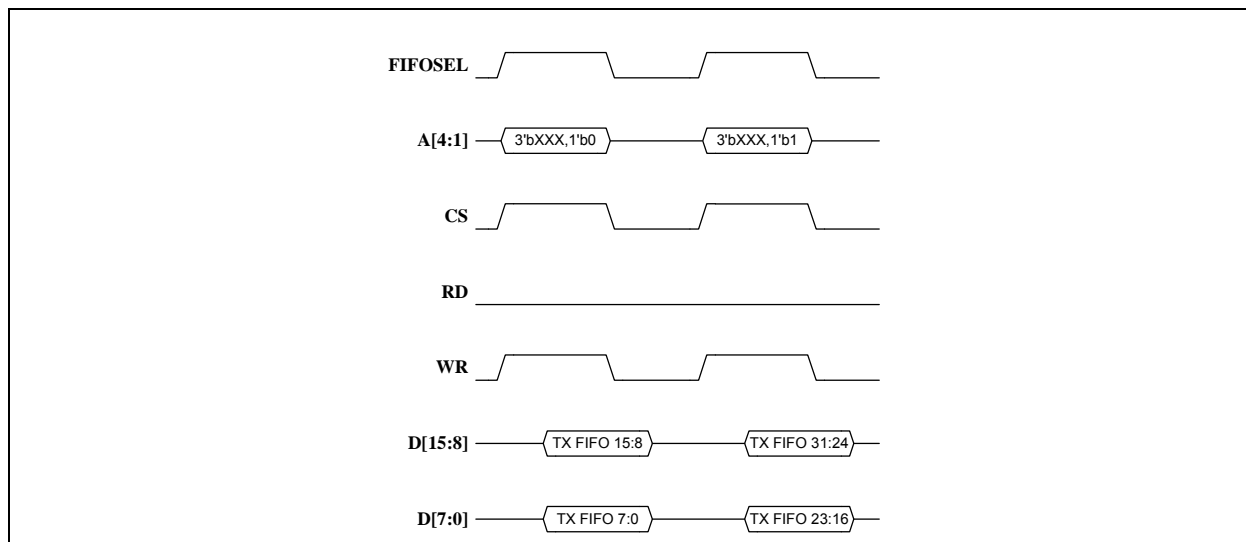
FIGURE 9-35: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 16-BIT READ



16-BIT WRITE

FIFOSEL is set high and address bit 1 is set to access the lower WORD of the FIFO, while address bits 4:2 are don't care. Data on **D[15:0]** is written on the trailing edge of **WR**. The cycle repeats for the upper WORD of the FIFO.

FIGURE 9-36: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 16-BIT WRITE

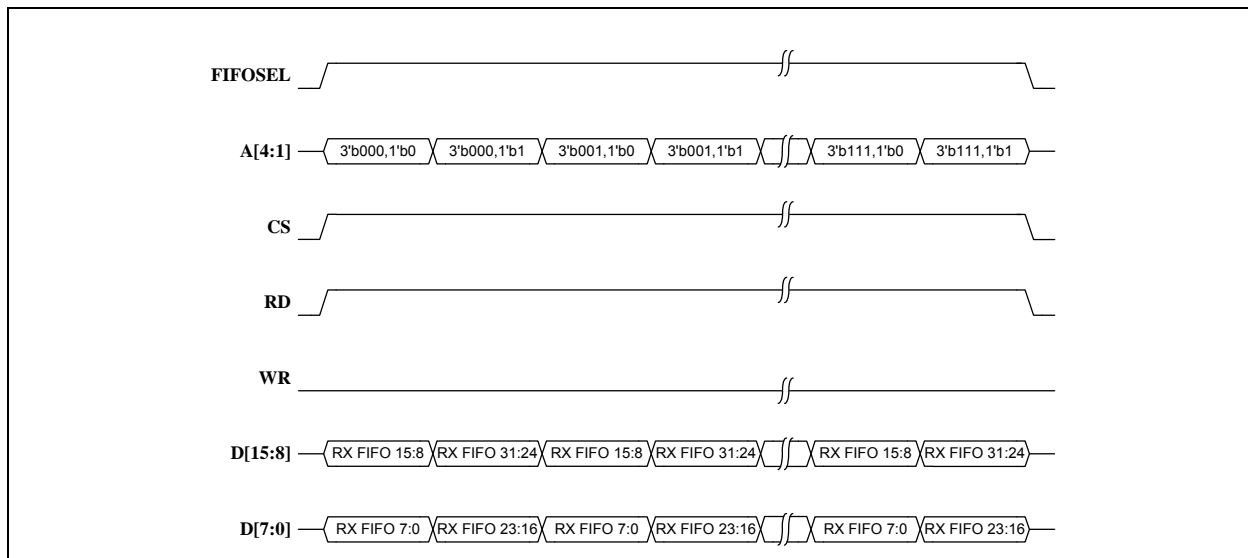


16-BIT BURST READ

FIFOSEL is set high and address bit 1 is set to access the lower WORD of the FIFO, while address bits 4:2 start at 0. Read data is driven on **D[15:0]** during **RD** active. Address bit 1 is then set to access the WORD of the FIFO as **RD** is held active.

While **RD** is held active, address bits 4:2 are cycled from 0 through 7 to access the 8 DWORDs as address bit 1 is toggled to access each WORD. Fresh data is supplied each time **A[1]** toggles. The FIFO is popped when **A[2]** toggles.


FIGURE 9-37: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 16-BIT BURST READ



8-BIT READ

FIFOSEL is set high and address bits 1 and 0 are set to access the lower BYTE of the FIFO, while address bits 4:2 are don't care. Read data is driven on **D[7:0]** during **RD** active. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining 3 BYTES of the FIFO's DWORD.


FIGURE 9-38: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 8-BIT READ



8-BIT WRITE

FIFOSEL is set high and address bits 1 and 0 are set to access the lower BYTE of the FIFO, while address bits 4:2 are don't care. Data on **D[7:0]** is written on the trailing edge of **WR**. **D[15:8]** pins are not used or driven. The cycle repeats for the remaining 3 BYTES of the FIFO's DWORD.

FIGURE 9-39: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 8-BIT WRITE




8-BIT BURST READ

FIFOSEL is set high and address bits 1 and 0 are set to access the lower BYTE of the FIFO, while address bits 4:2 start at 0. Read data is driven on **D[7:0]** during **RD** active. **D[15:8]** pins are not used or driven. Address bits 1 & 0 are then cycled from 0 through 3 to access the next 3 BYTES of the FIFO as **RD** is held active.

While **RD** is held active, address bits 4:2 are cycled from 0 through 7 to access the 8 DWORDs as address bits 1 & 0 are cycled from 0 through 3 to access each BYTE. Fresh data is supplied each time **A[0]** toggles. The FIFO is popped when **A[2]** toggles.

FIGURE 9-40: INDEXED ADDRESSING FIFO DIRECT SELECT ACCESS - 8-BIT BURST READ



9.5.6.5 RD_WR / ENB Control Mode Examples


The figures in this section detail read and write operations utilizing the alternative **RD_WR** and **ENB** signaling. The HBI read/write mode is selectable via the [HBI_rw_mode_strap](#).

Note: The examples in this section detail 16-bit mode with access to an Index Register. However, the **RD_WR** and **ENB** signaling can be used identically for all other accesses including FIFO Direct Select Access.

The examples in this section show the **ENB** signal active-high and the **RD_WR** signal low for read and high for write. The polarities of the **RD_WR** and **ENB** signals are selectable via the [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#).

16-BIT

FIGURE 9-41: INDEXED ADDRESSING RD_WR / ENB CONTROL MODE EXAMPLE - 16-BIT WRITE/READ



9.5.7 INDEXED ADDRESSING MODE TIMING REQUIREMENTS

The following figures and tables specify the timing requirements during Indexed Address mode. Since timing requirements are similar across the multitude of operations (e.g. 8 vs. 16-bit, Index vs. Configuration vs. Data registers, FIFO Direct Select), many timing requirements are illustrated in the same figures and do not necessarily represent any particular functional operation.

The following should be noted for the timing specifications in this section:

- The diagrams in this section depict active-high **CS**, **RD**, **WR**, **RD_WR** and **ENB** signals. The polarities of these signals are selectable via the [HBI_cs_polarity_strap](#), [HBI_rd_rdw_r_polarity_strap](#), and [HBI_wr_en_polarity_strap](#), respectively. Refer to [Section 9.3, "Control Line Polarity,"](#) on page 86 for additional details.
- A read cycle maybe followed by followed by a write cycle or another read cycle. A write cycle maybe followed by followed by a read cycle or another write cycle. These are shown in dashed line.

9.5.7.1 Read Timing Requirements

If **RD** and **WR** signaling is used, a host read cycle begins when **RD** is asserted with **CS** active. The cycle ends when **RD** is de-asserted. **CS** maybe asserted and de-asserted along with **RD** but not during **RD** active.

Alternatively, if **RD_WR** and **ENB** signaling is used, a host read cycle begins when **ENB** is asserted with **CS** active and **RD_WR** indicating a read. The cycle ends when **ENB** is de-asserted. **CS** maybe asserted and de-asserted along with **ENB** but not during **ENB** active.

Please refer to [Section 9.5.6, "Indexed Address Mode Functional Timing Diagrams,"](#) on page 114 for functional descriptions.

FIGURE 9-42: INDEXED ADDRESSING READ CYCLE TIMING




TABLE 9-4: INDEXED ADDRESSING READ CYCLE TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{csrd}	CS Setup to RD or ENB Active	0			ns
t_{rdcs}	CS Hold from RD or ENB Inactive	0			ns
t_{as}	Address, FIFOSEL Setup to RD or ENB Active	0			ns
t_{ah}	Address, FIFOSEL Hold from to RD or ENB Inactive	0			ns
t_{rdwrs}	RD_WR Setup to ENB Active Note 14	5			ns
t_{rdwrh}	RD_WR Hold from ENB Inactive Note 14	5			ns
t_{rdon}	RD or ENB to Data Buffer Turn On	0			ns
t_{rddv}	RD or ENB Active to Data Valid			30	ns
t_{rddh}	Data Output Hold Time from RD or ENB Inactive	0			ns
t_{rddz}	Data Buffer Turn Off Time from RD or ENB Inactive			9	ns
t_{cson}	CS to Data Buffer Turn On	0			ns
t_{csdv}	CS Active to Data Valid			30	ns
t_{csdh}	Data Output Hold Time from CS Inactive	0			ns
t_{csdz}	Data Buffer Turn Off Time from CS Inactive			9	ns
t_{fdv}	FIFOSEL to Data Valid			30	ns
t_{adv}	Address to Data Valid			30	ns
t_{rd}	RD or ENB Active Time	32			ns
t_{rdcyc}	RD or ENB Cycle Time	45			ns
t_{rdrd}	RD or ENB De-assertion Time before Next RD or ENB	13			ns
t_{rdwr}	RD De-assertion Time before Next WR Note 15	13			ns

Note 14: **RD_WR** and **ENB** signaling.

Note 15: **RD** and **WR** signaling.

Note: Timing values are with respect to an equivalent test load of 25 pF.

9.5.7.2 FIFO Direct Select Burst Timing Requirements

If **RD** and **WR** signaling is used, a host burst read from the FIFO begins when **RD** is asserted with **CS** active and **FIFOSEL** high. As the address changes, the next data is read. The cycle ends when **RD** is de-asserted. **CS** maybe asserted and de-asserted along with **RD** but not during **RD** active.

Alternatively, if **RD_WR** and **ENB** signaling is used, a host burst read from the **FIFO** begins when **ENB** is asserted with **CS** active, **RD_WR** indicating a read and **FIFOSEL** high. As the address changes, the next data is read. The cycle ends when **ENB** is de-asserted. **CS** maybe asserted and de-asserted along with **ENB** but not during **ENB** active.

Please refer to [Section 9.5.6, "Indexed Address Mode Functional Timing Diagrams," on page 114](#) for functional descriptions.

FIGURE 9-43: INDEXED ADDRESSING FIFO DIRECT SELECT BURST READ CYCLE TIMING

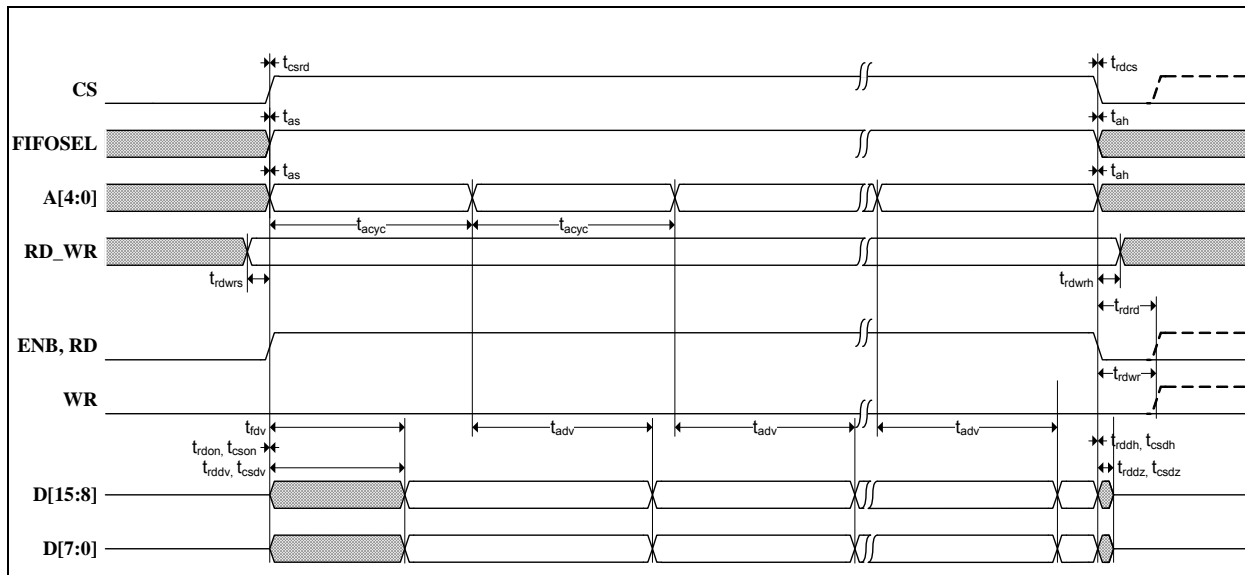


TABLE 9-5: INDEXED ADDRESSING FIFO DIRECT SELECT BURST READ CYCLE TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{csr}	CS Setup to RD or ENB Active	0			ns
t_{rdcs}	CS Hold from RD or ENB Inactive	0			ns
t_{as}	Address, FIFOSEL Setup to RD or ENB Active	0			ns
t_{ah}	Address, FIFOSEL Hold from to RD or ENB Inactive	0			ns
t_{rdwrs}	RD_WR Setup to ENB Active Note 16	5			ns
t_{rdwrh}	RD_WR Hold from ENB Inactive Note 16	5			ns
t_{rdon}	RD or ENB to Data Buffer Turn On	0			ns
t_{rddv}	RD or ENB Active to Data Valid			30	ns
t_{rddh}	Data Output Hold Time from RD or ENB Inactive	0			ns
t_{rddz}	Data Buffer Turn Off Time from RD or ENB Inactive			9	ns

TABLE 9-5: INDEXED ADDRESSING FIFO DIRECT SELECT BURST READ CYCLE TIMING VALUES (CONTINUED)

Symbol	Description	Min	Typ	Max	Units
t_{cson}	CS to Data Buffer Turn On	0			ns
t_{csdv}	CS Active to Data Valid			30	ns
t_{csdh}	Data Output Hold Time from CS Inactive	0			ns
t_{csdz}	Data Buffer Turn Off Time from CS Inactive			9	ns
t_{fdv}	FIFOSEL to Data Valid			30	ns
t_{adv}	Address Change to Next Data Valid			40	ns
t_{acyc}	Address Cycle Time	45			ns
t_{rdrd}	RD or ENB De-assertion Time before Next RD or ENB	13			ns
t_{rdwr}	RD De-assertion Time before Next WR Note 17	13			ns

Note 16: **RD_WR** and **ENB** signaling.

Note 17: **RD** and **WR** signaling.

Note: Timing values are with respect to an equivalent test load of 25 pF.

9.5.7.3 Write Timing Requirements

If **RD** and **WR** signaling is used, a host write cycle begins when **WR** is asserted with **CS** active. The cycle ends when **WR** is de-asserted. **CS** maybe asserted and de-asserted along with **WR** but not during **WR** active.

Alternatively, if **RD_WR** and **ENB** signaling is used, a host write cycle begins when **ENB** is asserted with **CS** active and **RD_WR** indicating a write. The cycle ends when **ENB** is de-asserted. **CS** maybe asserted and de-asserted along with **ENB** but not during **ENB** active.

Please refer to [Section 9.5.6, "Indexed Address Mode Functional Timing Diagrams," on page 114](#) for functional descriptions.

FIGURE 9-44: INDEXED ADDRESSING WRITE CYCLE TIMING

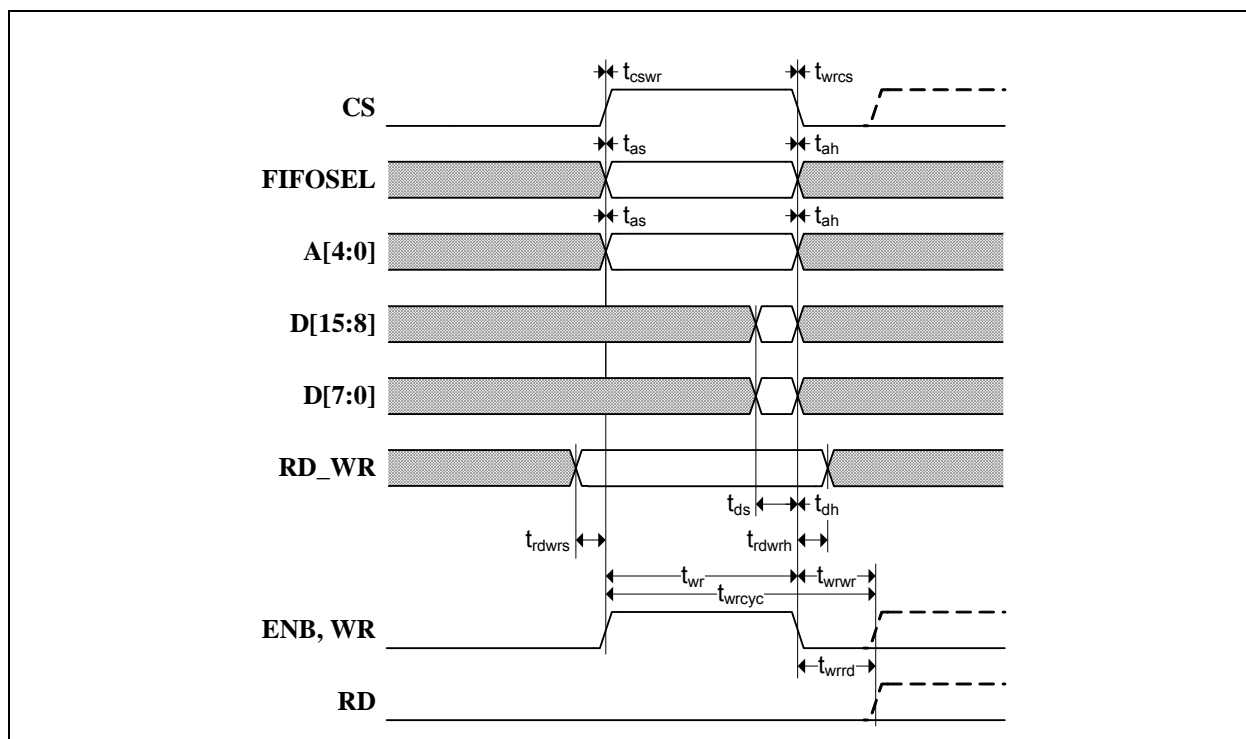


TABLE 9-6: INDEXED ADDRESSING WRITE CYCLE TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{cswr}	CS Setup to WR or ENB Active	0			ns
t_{wrscs}	CS Hold from WR or ENB Inactive	0			ns
t_{as}	Address, FIFOSEL Setup to WR or ENB Active	0			ns
t_{ah}	Address, FIFOSEL Hold from WR or ENB Inactive	0			ns
t_{rdwr}	RD_WR Setup to ENB Active Note 18	5			ns
t_{rdwrh}	RD_WR Hold from ENB Inactive Note 18	5			ns
t_{ds}	Data Setup to WR or ENB Inactive	7			ns
t_{dh}	Data Hold from WR or ENB Inactive	0			ns

TABLE 9-6: INDEXED ADDRESSING WRITE CYCLE TIMING VALUES (CONTINUED)

Symbol	Description	Min	Typ	Max	Units
t_{wr}	WR or ENB Active Time	32			ns
t_{wrcyc}	WR or ENB Cycle Time	45			ns
t_{wrwr}	WR or ENB De-assertion Time before Next WR or ENB	13			ns
t_{wrrd}	WR De-assertion Time before Next RD Note 19	13			ns

Note 18: **RD_WR** and **ENB** signaling.

Note 19: **RD** and **WR** signaling.

10.0 SPI/SQI SLAVE

10.1 Functional Overview

The SPI/SQI Slave module provides a low pin count synchronous slave interface that facilitates communication between the device and a host system. The SPI/SQI Slave allows access to the System CSRs and internal FIFOs and memories. It supports single and multiple register read and write commands with incrementing, decrementing and static addressing. Single, Dual and Quad bit lanes are supported in SPI mode with a clock rate of up to 80 MHz. SQI mode always uses four bit lanes and also operates at up to 80 MHz.

The following is an overview of the functions provided by the SPI/SQI Slave:

- **Serial Read:** 4-wire (clock, select, data in and data out) reads at up to 30 MHz. Serial command, address and data. Single and multiple register reads with incrementing, decrementing or static addressing.
- **Fast Read:** 4-wire (clock, select, data in and data out) reads at up to 80 MHz. Serial command, address and data. Dummy byte(s) for first access. Single and multiple register reads with incrementing, decrementing or static addressing.
- **Dual / Quad Output Read:** 4 or 6-wire (clock, select, data in / out) reads at up to 80 MHz. Serial command and address, parallel data. Dummy byte(s) for first access. Single and multiple register reads with incrementing, decrementing or static addressing.
- **Dual / Quad I/O Read:** 4 or 6-wire (clock, select, data in / out) reads at up to 80 MHz. Serial command, parallel address and data. Dummy byte(s) for first access. Single and multiple register reads with incrementing, decrementing or static addressing.
- **SQI Read:** 6-wire (clock, select, data in / out) writes at up to 80 MHz. Parallel command, address and data. Dummy byte(s) for first access. Single and multiple register reads with incrementing, decrementing or static addressing.
- **Write:** 4-wire (clock, select, data in and data out) writes at up to 80 MHz. Serial command, address and data. Single and multiple register writes with incrementing, decrementing or static addressing.
- **Dual / Quad Data Write:** 4 or 6-wire (clock, select, data in / out) writes at up to 80 MHz. Serial command and address, parallel data. Single and multiple register writes with incrementing, decrementing or static addressing.
- **Dual / Quad Address / Data Write:** 4 or 6-wire (clock, select, data in / out) writes at up to 80 MHz. Serial command, parallel address and data. Single and multiple register writes with incrementing, decrementing or static addressing.
- **SQI Write:** 6-wire (clock, select, data in / out) writes at up to 80 MHz. Parallel command, address and data. Single and multiple register writes with incrementing, decrementing or static addressing.

10.2 SPI/SQI Slave Operation

Input data on the **SIO[3:0]** pins is sampled on the rising edge of the **SCK** input clock. Output data is sourced on the **SIO[3:0]** pins with the falling edge of the clock. The **SCK** input clock can be either an active high pulse or an active low pulse. When the **SCS#** chip select input is high, the **SIO[3:0]** inputs are ignored and the **SIO[3:0]** outputs are three-stated.

In SPI mode, the 8-bit instruction is started on the first rising edge of the input clock after **SCS#** goes active. The instruction is always input serially on **SI/SIO0**.

For read and write instructions, two address bytes follow the instruction byte. Depending on the instruction, the address bytes are input either serially, or 2 or 4 bits per clock. Although all registers are accessed as DWORDs, the address field is considered a byte address. Fourteen address bits specify the address. Bits 15 and 14 of the address field specifies that the address is auto-decremented (10b) or auto-incremented (01b) for continuous accesses.

For some read instructions, dummy byte cycles follow the address bytes. The device does not drive the outputs during the dummy byte cycles. The dummy byte(s) are input either serially, or 2 or 4 bits per clock.

For read and write instructions, one or more 32-bit data fields follow the dummy bytes (if present, else they follow the address bytes). The data is input either serially, or 2 or 4 bits per clock.

SQI mode is entered from SPI with the Enable Quad I/O (EQIO) instruction. Once in SQI mode, all further command, addresses, dummy bytes and data bytes are 4 bits per clock. SQI mode can be exited using the Reset Quad I/O (RSTQIO) instruction.

All instructions, addresses and data are transferred with the most-significant bit (msb) or di-bit (msd) or nibble (msn) first. Addresses are transferred with the most-significant byte (MSB) first. Data is transferred with the least-significant byte (LSB) first (little endian).

The SPI interface supports up to a 80 MHz input clock. Normal (non-high speed) reads instructions are limited to 30 MHz.

The SPI interface supports a minimum time of 50 ns between successive commands (a minimum SCS# inactive time of 50 ns).

The instructions supported in SPI mode are listed in [Table 10-1](#). SQL instructions are listed in [Table 10-2](#). Unsupported instructions are must not be used.

TABLE 10-1: SPI INSTRUCTIONS

Instruction	Description	Bit width Note 1	Inst. code	Addr. Bytes	Dummy Bytes	Data bytes	Max Freq.
Configuration							
EQIO	Enable SQL	1-0-0	38h	0	0	0	80 MHz
RSTQIO	Reset SQL	1-0-0	FFh	0	0	0	80 MHz
Read							
READ	Read	1-1-1	03h	2	0	4 to ∞	30 MHz
FASTREAD	Read at higher speed	1-1-1	0Bh	2	1	4 to ∞	80 MHz
SDOR	SPI Dual Output Read	1-1-2	3Bh	2	1	4 to ∞	80 MHz
SDIOR	SPI Dual I/O Read	1-2-2	BBh	2	2	4 to ∞	80 MHz
SQOR	SPI Quad Out-put Read	1-1-4	6Bh	2	1	4 to ∞	80 MHz
SQIOR	SPI Quad I/O Read	1-4-4	EBh	2	4	4 to ∞	80 MHz
Write							
WRITE	Write	1-1-1	02h	2	0	4 to ∞	80 MHz
SDDW	SPI Dual Data Write	1-1-2	32h	2	0	4 to ∞	80 MHz
SDADW	SPI Dual Address / Data Write	1-2-2	B2h	2	0	4 to ∞	80 MHz
SQDW	SPI Quad Data Write	1-1-4	62h	2	0	4 to ∞	80 MHz
SQADW	SPI Quad Address / Data Write	1-4-4	E2h	2	0	4 to ∞	80 MHz

Note 1: The bit width format is: command bit width, address / dummy bit width, data bit width.

TABLE 10-2: SQI INSTRUCTIONS

Instruction	Description	Bit width Note 2	Inst. code	Addr. Bytes	Dummy Bytes	Data bytes	Max Freq.
Configuration							
RSTQIO	Reset SQI	4-0-0	FFh	0	0	0	80 MHz
Read							
FASTREAD	Read at higher speed	4-4-4	0Bh	2	3	4 to ∞	80 MHz
Write							
WRITE	Write	4-4-4	02h	2	0	4 to ∞	80 MHz

Note 2: The bit width format is: command bit width, address / dummy bit width, data bit width.

10.2.1 DEVICE INITIALIZATION

Until the device has been initialized to the point where the various configuration inputs are valid, the SPI/SQI interface does not respond to and is not affected by any external pin activity.

Once device initialization completes, the SPI/SQI interface will ignore the pins until a rising edge of SCS# is detected.

10.2.1.1 SPI/SQI Slave Read Polling for Initialization Complete

Before device initialization, the SPI/SQI interface will not return valid data. To determine when the SPI/SQI interface is functional, the [Byte Order Test Register \(BYTE_TEST\)](#) should be polled. Once the correct pattern is read, the interface can be considered functional. At this point, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) can be polled to determine when the device is fully configured.

Note: The Host should only use single register reads (one data cycle per SCS# low) while polling the BYTE_TEST register.

10.2.2 ACCESS DURING AND FOLLOWING POWER MANAGEMENT

During any power management mode other than D0, reads and writes are ignored and the SPI/SQI interface does not respond to and is not affected by any external pin activity.

Once the power management mode changes back to D0, the SPI/SQI interface will ignore the pins until a rising edge of SCS# is detected.

To determine when the SPI/SQI interface is functional, the [Byte Order Test Register \(BYTE_TEST\)](#) should be polled. Once the correct pattern is read, the interface can be considered functional. At this point, the [Device Ready \(READY\)](#) bit in the [Hardware Configuration Register \(HW_CFG\)](#) can be polled to determine when the device is fully configured.

Note: The Host should only use single register reads (one data cycle per SCS# low) while polling the BYTE_TEST register.

10.2.3 SPI CONFIGURATION COMMANDS


10.2.3.1 Enable SQI

The [Enable SQI](#) instruction changes the mode of operation to SQI. This instruction is supported in SPI bus protocol only with clock frequencies up to [80 MHz](#). This instruction is not supported in SQI bus protocol.

The SPI slave interface is selected by first bringing SCS# active. The 8-bit [EQIO](#) instruction, [38h](#), is input into the [SI/SIO\[0\]](#) pin one bit per clock. The SCS# input is brought inactive to conclude the cycle.

Figure 10-1 illustrates the Enable SQI instruction.

FIGURE 10-1: ENABLE SQI



10.2.3.2 Reset SQI

The [Reset SQI](#) instruction changes the mode of operation to SPI. This instruction is supported in SPI and SQI bus protocols with clock frequencies up to [80 MHz](#).

The SPI/SQI slave interface is selected by first bringing SCS# active. The 8-bit [RSTQIO](#) instruction, [FFh](#), is input into the SI/SIO[0] pin, one bit per clock, in SPI mode and into the SIO[3:0] pins, four bits per clock, in SQI mode. The SCS# input is brought inactive to conclude the cycle.

[Figure 10-2](#) illustrates the Reset SQI instruction for SPI mode. [Figure 10-3](#) illustrates the Reset SQI instruction for SQI mode.

FIGURE 10-2: SPI MODE RESET SQI





FIGURE 10-3: SQUI MODE RESET SQUI



10.2.4 SPI READ COMMANDS

Various read commands are support by the SPI/SQI slave. The following applies to all read commands.

MULTIPLE READS

Additional reads, beyond the first, are performed by continuing the clock pulses while **SCS#** is active. The upper two bits of the address specify auto-incrementing (address[15:14]=01b) or auto-decrementing (address[15:14]=10b). The internal DWORD address is incremented, decremented, or maintained based on these bits. Maintaining a fixed internal address is useful for register polling.

SPECIAL CSR HANDLING

Live Bits

Since data is read serially, the selected register's value is saved at the beginning of each 32-bit read to prevent the host from reading an intermediate value. The saving occurs multiple times in a multiple read sequence.

Change on Read Registers and FIFOs

Any register that is affected by a read operation (e.g. a clear on read bit or FIFO) is updated once the current data output shift has started. In the event that 32-bits are not read when the **SCS#** is returned high, the register is still affected and any prior data is lost.

Change on Read Live Register Bits

As described above, the current value from a register with live bits (as is the case of any register) is saved before the data is shifted out. Although a H/W event that occurs following the data capture would still update the live bit(s), the live bit(s) will be affected (cleared, etc.) once the output shift has started and the H/W event would be lost. In order to prevent this, the individual CSRs defer the H/W event update until after the read indication.

10.2.4.1 Read

The **Read** instruction inputs the instruction code and address bytes one bit per clock and outputs the data one bit per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **30 MHz**. This instruction is not supported in SQI bus protocol.


The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **READ** instruction, **03h**, is input into the **SI/SIO[0]** pin, followed by the two address bytes. The address bytes specify a BYTE address within the device.

On the falling clock edge following the rising edge of the last address bit, the **SO/SIO[1]** pin is driven starting with the msb of the LSB of the selected register. The remaining register bits are shifted out on subsequent falling clock edges.

The **SCS#** input is brought inactive to conclude the cycle. The **SO/SIO[1]** pin is three-stated at this time.

Figure 10-4 illustrates a typical single and multiple register read.

FIGURE 10-4: SPI READ



10.2.4.2 Fast Read

The [Read at higher speed](#) instruction inputs the instruction code and the address and dummy bytes one bit per clock and outputs the data one bit per clock. In SQI mode, the instruction code and the address and dummy bytes are input four bits per clock and the data is output four bits per clock. This instruction is supported in SPI and SQI bus protocols with clock frequencies up to **80 MHz**.

The SPI/SQI slave interface is selected by first bringing SCS# active. For SPI mode, the 8-bit [FASTREAD](#) instruction, [0Bh](#), is input into the SI/SIO[0] pin, followed by the two address bytes and **1** dummy byte. For SQI mode, the 8-bit [FASTREAD](#) instruction is input into the SIO[3:0] pins, followed by the two address bytes and **3** dummy bytes. The address bytes specify a BYTE address within the device.

On the falling clock edge following the rising edge of the last dummy bit (or nibble), the SO/SIO[1] pin is driven starting with the msb of the LSB of the selected register. For SQI mode, SIO[3:0] are driven starting with the msn of the LSB of the selected register. The remaining register bits are shifted out on subsequent falling clock edges.

The SCS# input is brought inactive to conclude the cycle. The SO/SIO[3:0] pins are three-stated at this time.

Figure 10-5 illustrates a typical single and multiple register fast read for SPI mode. Figure 10-6 illustrates a typical single and multiple register fast read for SQI mode.

FIGURE 10-5: SPI FAST READ


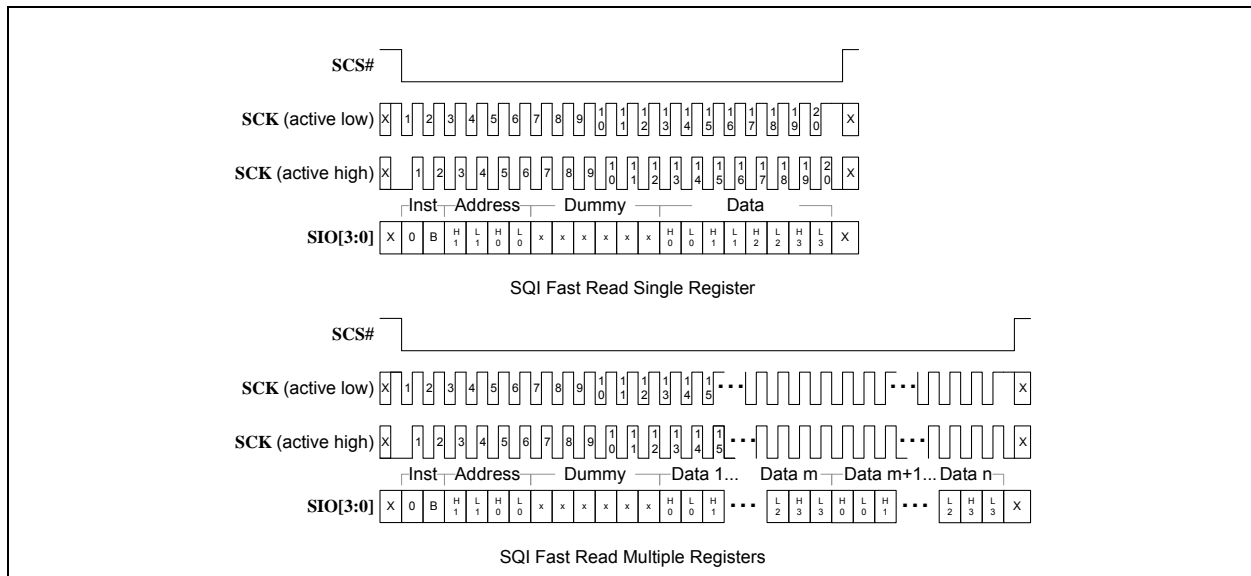


FIGURE 10-6: SQI FAST READ



10.2.4.3 Dual Output Read

The **SPI Dual Output Read** instruction inputs the instruction code and the address and dummy bytes one bit per clock and outputs the data two bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.

The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SDOR** instruction, **3Bh**, is input into the **SIO[0]** pin, followed by the two address bytes and **1** dummy byte. The address bytes specify a BYTE address within the device.

On the falling clock edge following the rising edge of the last dummy di-bit, the **SIO[1:0]** pins are driven starting with the msbs of the LSB of the selected register. The remaining register di-bits are shifted out on subsequent falling clock edges.

The **SCS#** input is brought inactive to conclude the cycle. The **SIO[1:0]** pins are three-stated at this time.

Figure 10-7 illustrates a typical single and multiple register dual output read.

FIGURE 10-7: SPI DUAL OUTPUT READ




Figure 10-8 illustrates a typical single and multiple register quad output read.

SCS#

SCK (active low)

SCK (active high)

Instruction Address Dummy Data

SIO0

SIO1

SIO2

SIO3

SPI Quad Output Read Single Register

SCS#

SCK (active low)

SCK (active high)

Instruction Address Dummy Data 1 ... Data m Data m+1 ... Data n

SIO0

SIO1

SIO2

SIO3

SPI Quad Output Read Multiple Registers

10.2.5.1 Dual I/O Read

The **SPI Dual I/O Read** instruction inputs the instruction code one bit per clock and the address and dummy bytes two bits per clock and outputs the data two bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.

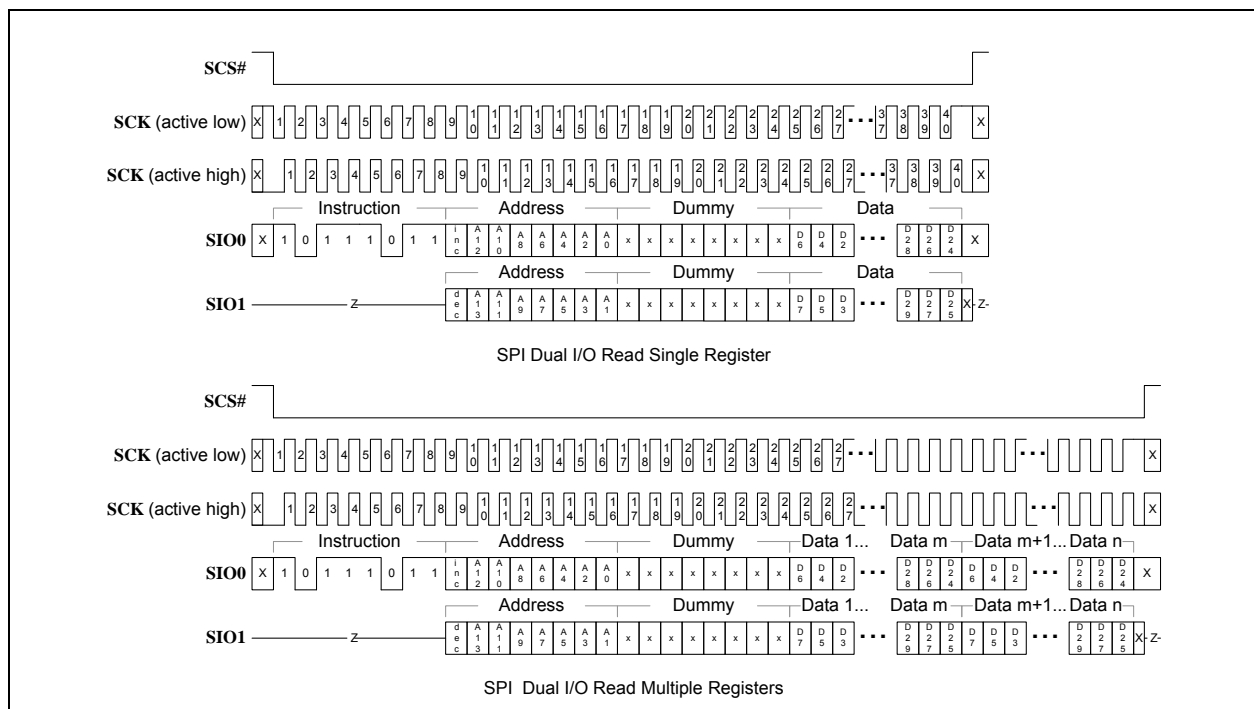
The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SDIOR** instruction, **BBh**, is input into the **SIO[0]** pin, followed by the two address bytes and **2** dummy bytes into the **SIO[1:0]** pins. The address bytes specify a BYTE address within the device.

On the falling clock edge following the rising edge of the last dummy di-bit, the **SIO[1:0]** pins are driven starting with the msbs of the LSB of the selected register. The remaining register di-bits are shifted out on subsequent falling clock edges.

The **SCS#** input is brought inactive to conclude the cycle. The **SIO[1:0]** pins are three-stated at this time.

Figure 10-9 illustrates a typical single and multiple register dual I/O read.

FIGURE 10-9: SPI DUAL I/O READ



10.2.5.2 Quad I/O Read

The [SPI Quad I/O Read](#) instruction inputs the instruction code one bit per clock and the address and dummy bytes four bits per clock and outputs the data four bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.

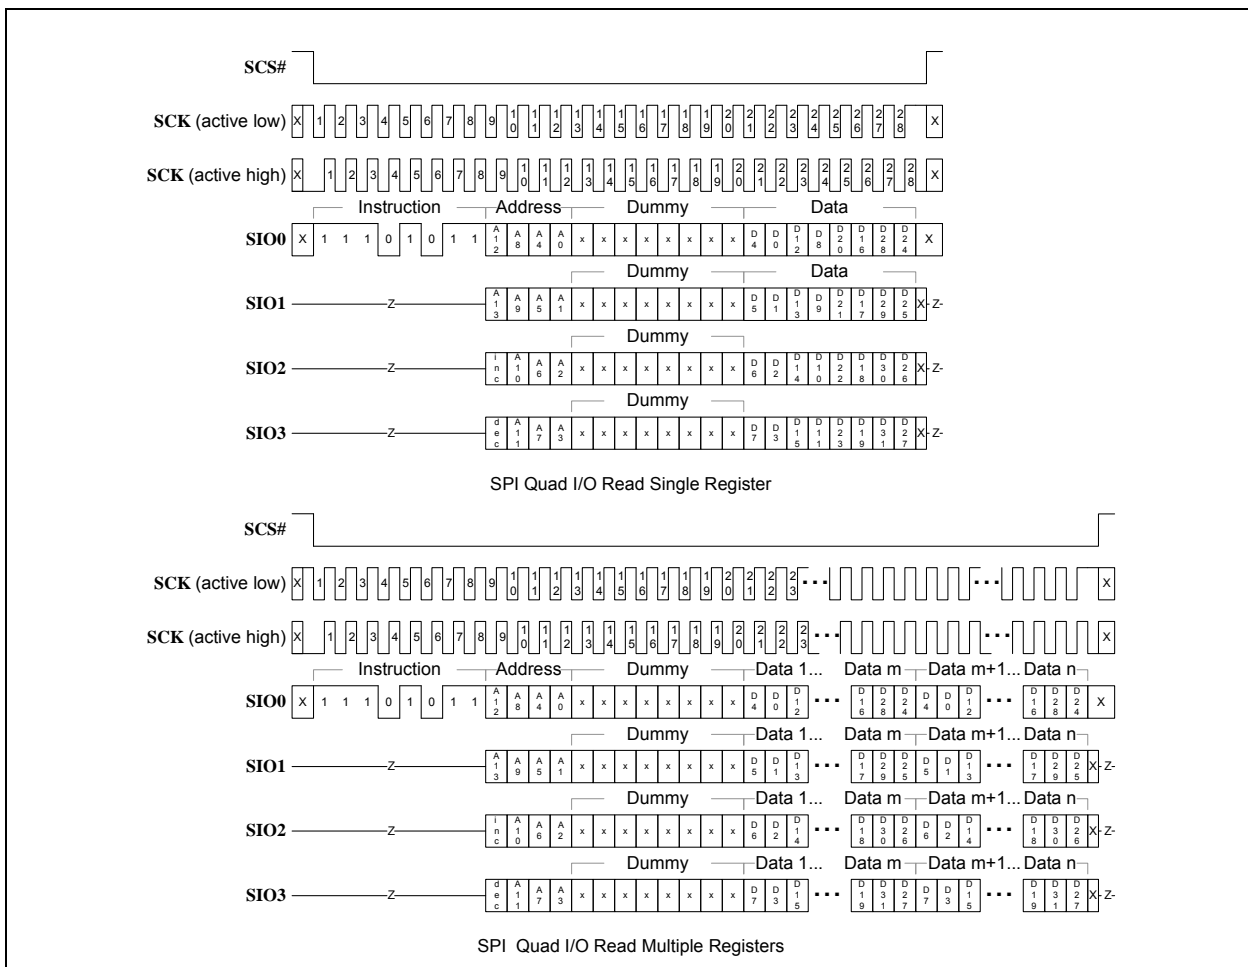
The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SQIOR** instruction, **EBh**, is input into the **SIO[0]** pin, followed by the two address bytes and **4** dummy bytes into the **SIO[3:0]** pins. The address bytes specify a BYTE address within the device.

On the falling clock edge following the rising edge of the last dummy nibble, the **SIO[3:0]** pins are driven starting with the msn of the LSB of the selected register. The remaining register nibbles are shifted out on subsequent falling clock edges.

The **SCS#** input is brought inactive to conclude the cycle. The **SIO[3:0]** pins are three-stated at this time.

[Figure 10-10](#) illustrates a typical single and multiple register quad I/O read.

FIGURE 10-10: SPI QUAD I/O READ



10.2.6 SPI WRITE COMMANDS

Multiple write commands are support by the SPI/SQI slave. The following applies to all write commands.

MULTIPLE WRITES

Multiple reads are performed by continuing the clock pulses and input data while **SCS#** is active. The upper two bits of the address specify auto-incrementing (address[15:14]=01b) or auto-decrementing (address[15:14]=10b). The internal DWORD address is incremented, decremented, or maintained based on these bits. Maintaining a fixed internal address may be useful for register "bit-banging" or other repeated writes.

10.2.6.1 Write

The **Write** instruction inputs the instruction code and address and data bytes one bit per clock. In SQI mode, the instruction code and the address and data bytes are input four bits per clock. This instruction is supported in SPI and SQI bus protocols with clock frequencies up to **80 MHz**.

The SPI/SQI slave interface is selected by first bringing **SCS#** active. For SPI mode, the 8-bit **WRITE** instruction, **02h**, is input into the **SI/SIO[0]** pin, followed by the two address bytes. For SQI mode, the 8-bit **WRITE** instruction, **02h**, is input into the **SIO[3:0]** pins, followed by the two address bytes. The address bytes specify a BYTE address within the device.

The data follows the address bytes. For SPI mode, the data is input into the **SI/SIO[0]** pin starting with the msb of the LSB. For SQI mode the data is input nibble wide using **SIO[3:0]** starting with the msn of the LSB. The remaining bits/nibbles are shifted in on subsequent clock edges. The data write to the register occurs after the 32-bits are input. In the event that 32-bits are not written when the **SCS#** is returned high, the write is considered invalid and the register is not affected.

The **SCS#** input is brought inactive to conclude the cycle.

Figure 10-11 illustrates a typical single and multiple register write for SPI mode. **Figure 10-12** illustrates a typical single and multiple register write for SQI mode.

FIGURE 10-11: SPI WRITE

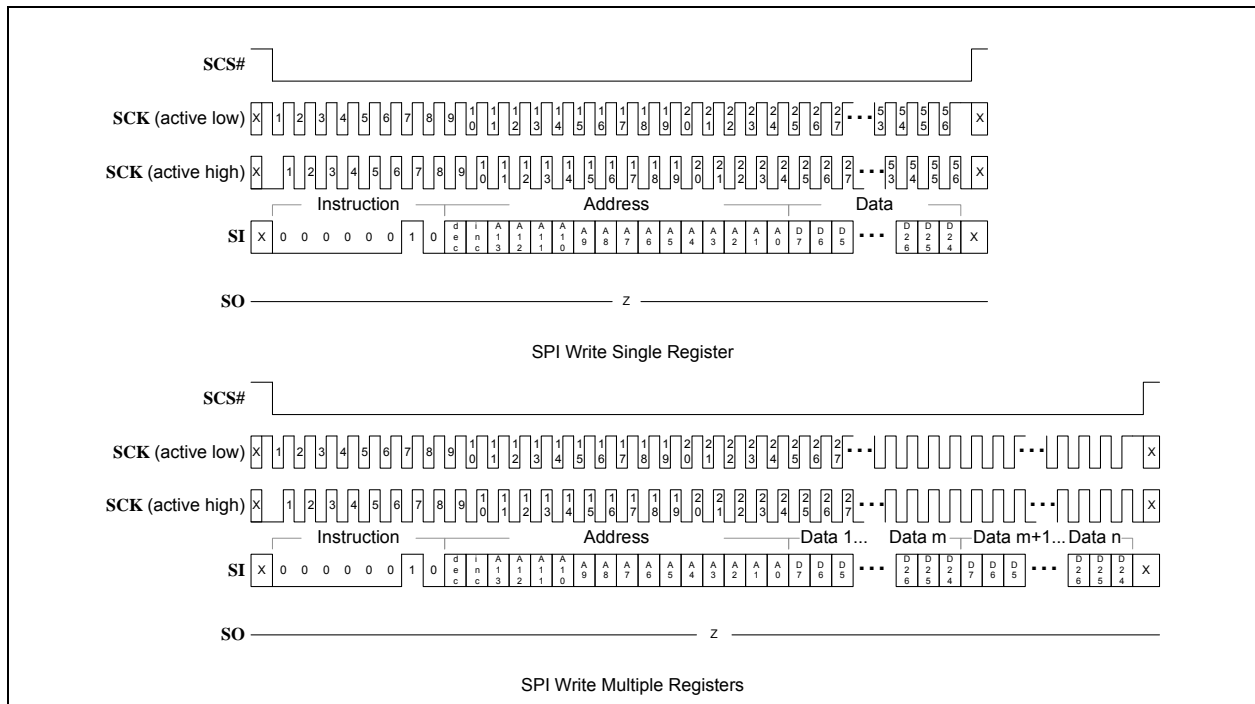



FIGURE 10-12: SQUI WRITE



10.2.6.2 Dual Data Write

The **SPI Dual Data Write** instruction inputs the instruction code and address bytes one bit per clock and inputs the data two bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.

The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SDDW** instruction, **32h**, is input into the **SIO[0]** pin, followed by the two address bytes. The address bytes specify a BYTE address within the device.

The data follows the address bytes. The data is input into the **SIO[1:0]** pins starting with the msbs of the LSB. The remaining di-bits are shifted in on subsequent clock edges. The data write to the register occurs after the 32-bits are input. In the event that 32-bits are not written when the **SCS#** is returned high, the write is considered invalid and the register is not affected.

The **SCS#** input is brought inactive to conclude the cycle.

Figure 10-13 illustrates a typical single and multiple register dual data write.

FIGURE 10-13: SPI DUAL DATA WRITE



10.2.6.3 Quad Data Write

The **SPI Quad Data Write** instruction inputs the instruction code and address bytes one bit per clock and inputs the data four bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.

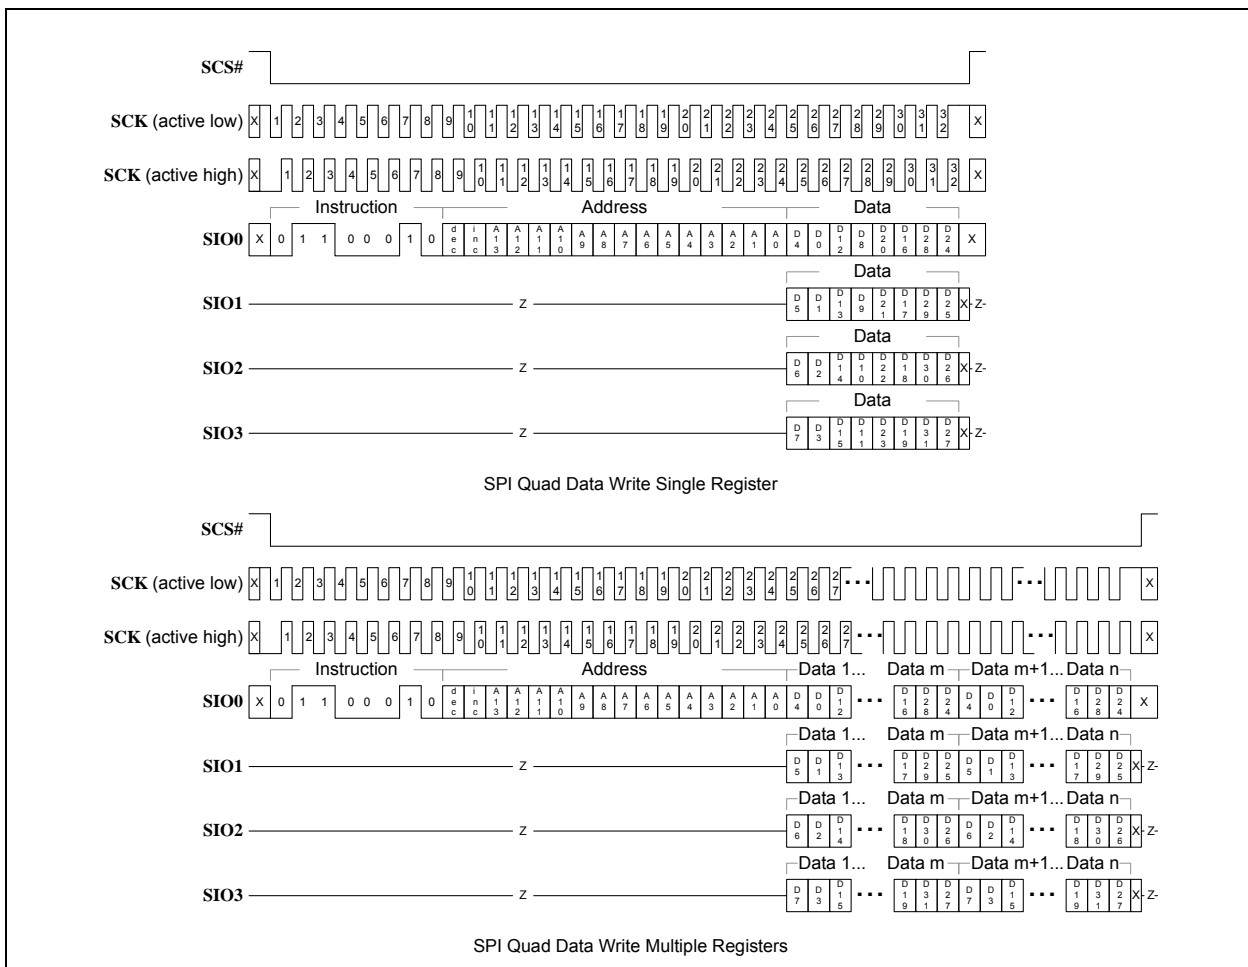
The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SQDW** instruction, **62h**, is input into the **SIO[0]** pin, followed by the two address bytes. The address bytes specify a BYTE address within the device.

The data follows the address bytes. The data is input into the **SIO[3:0]** pins starting with the msn of the LSB. The remaining nibbles are shifted in on subsequent clock edges. The data write to the register occurs after the 32-bits are input. In the event that 32-bits are not written when the **SCS#** is returned high, the write is considered invalid and the register is not affected.

The **SCS#** input is brought inactive to conclude the cycle.

Figure 10-14 illustrates a typical single and multiple register quad data write.

FIGURE 10-14: SPI QUAD DATA WRITE



10.2.6.4 Dual Address / Data Write

The [SPI Dual Address / Data Write](#) instruction inputs the instruction code one bit per clock and the address and data bytes two bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to [80 MHz](#). This instruction is not supported in SQI bus protocol.


The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit [SDADW](#) instruction, [B2h](#), is input into the **SIO[0]** pin, followed by the two address bytes into the **SIO[1:0]** pins. The address bytes specify a BYTE address within the device.

The data follows the address bytes. The data is input into the **SIO[1:0]** pins starting with the msbs of the LSB. The remaining di-bits are shifted in on subsequent clock edges. The data write to the register occurs after the 32-bits are input. In the event that 32-bits are not written when the **SCS#** is returned high, the write is considered invalid and the register is not affected.

The **SCS#** input is brought inactive to conclude the cycle.

[Figure 10-15](#) illustrates a typical single and multiple register dual address / data write.

FIGURE 10-15: SPI DUAL ADDRESS / DATA WRITE



10.2.6.5 Quad Address / Data Write

The **SPI Quad Address / Data Write** instruction inputs the instruction code one bit per clock and the address and data bytes four bits per clock. This instruction is supported in SPI bus protocol only with clock frequencies up to **80 MHz**. This instruction is not supported in SQI bus protocol.


The SPI slave interface is selected by first bringing **SCS#** active. The 8-bit **SQADW** instruction, **E2h**, is input into the **SIO[0]** pin, followed by the two address bytes into the **SIO[3:0]** pins. The address bytes specify a BYTE address within the device.

The data follows the address bytes. The data is input into the **SIO[3:0]** pins starting with the msn of the LSB. The remaining nibbles are shifted in on subsequent clock edges. The data write to the register occurs after the 32-bits are input. In the event that 32-bits are not written when the **SCS#** is returned high, the write is considered invalid and the register is not affected.

The **SCS#** input is brought inactive to conclude the cycle.

Figure 10-16 illustrates a typical single and multiple register dual address / data write.

FIGURE 10-16: SPI QUAD ADDRESS / DATA WRITE



10.3 TX and RX FIFO Access

10.3.0.1 TX and RX Status FIFO Peek Address Access

Normal read access to the TX or RX Status FIFO causes the FIFO to advance to its next entry.

For access to the TX and RX Status FIFO Peek addresses, the FIFO does not advance to its next entry.

10.4 SPI/SQI Timing Requirements

FIGURE 10-17: SPI/SQI INPUT TIMING




FIGURE 10-18: SPI/SQI OUTPUT TIMING




TABLE 10-3: SPI/SQI TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
f_{sck}	SCK clock frequency Note 3			30 / 80	MHz
t_{high}	SCK high time	5.5			ns
t_{low}	SCK low time	5.5			ns
t_{scss}	SCS# setup time to SCK	5			ns
t_{scsh}	SCS# hold time from SCK	5			ns
t_{scshl}	SCS# inactive time	50			ns
t_{su}	Data input setup time to SCK	3			ns
t_{hd}	Data input hold time from SCK	4			ns
t_{on}	Data output turn on time from SCK	0			ns
t_v	Data output valid time from SCK Note 4 , Note 5			11.0/9.0	ns
t_{ho}	Data output hold time from SCK	0			ns
t_{dis}	Data output disable time from SCS# inactive			20	ns

Note 3: The [Read](#) instruction is limited to 30 MHz maximum

Note 4: Depends on loading of 30 pF or 10 pF

Note 5: Depending on the clock frequency and pulse width, data may not be valid until following the next rising edge of SCK. The host SPI controller may need to delay the sampling of the data by either a fixed time or by using the falling edge of SCK.

11.0 HOST MAC

11.1 Functional Overview

The Host MAC incorporates the essential protocol requirements for operating an Ethernet/IEEE 802.3-compliant node and provides an interface between the Host and the Switch Fabric. On the front end, the Host MAC interfaces to the Host via 2 sets of FIFO's (TX Data FIFO, TX Status FIFO, RX Data FIFO, RX Status FIFO). An additional bus is used to access the Host MAC CSRs via the [Host MAC CSR Interface Command Register \(MAC_CSR_CMD\)](#) and [Host MAC CSR Interface Data Register \(MAC_CSR_DATA\)](#) system registers.

The receive and transmit FIFO's allow increased packet buffer storage to the Host MAC. The FIFOs are a conduit between the Host and the Host MAC through which all transmitted and received data and status information is passed. Deep FIFOs allow a high degree of latency tolerance relative to the various transport and OS software stacks reducing and minimizing overrun conditions. Both the Host MAC and the TX/RX FIFOs have separate receive and transmit data paths.

The Host MAC can store up to 250 Ethernet packets utilizing FIFOs, totaling 16KB, with a packet granularity of 4 bytes. This memory is shared by the RX and TX blocks and is configurable in terms of allocation via the [Hardware Configuration Register \(HW_CFG\)](#) register to the ranges described in [Section 11.10.3, "FIFO Memory Allocation Configuration"](#). This depth of buffer storage minimizes or eliminates receive overruns.

On the back end, the Host MAC interfaces with the 10/100 Ethernet PHYs (Virtual PHY, PHY A, PHY B) via an internal SMI (Serial Management Interface) bus. This allows the Host MAC access to the PHY's internal registers via the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#) and [Host MAC MII Data Register \(HMAC_MII_DATA\)](#). The Host MAC interfaces to the Switch Engine Port 0 via an internal MII (Media Independent Interface) connection allowing for incoming and outgoing Ethernet packet transfers.

The Host MAC can operate at either 100Mbps or 10Mbps in both half-duplex or full-duplex modes. When operating in half-duplex mode, the Host MAC complies fully with Section 4 of ISO/IEC 8802-3 (ANSI/IEEE standard) and ANSI/IEEE 802.3 standards. When operating in full-duplex mode, the Host MAC complies with IEEE 802.3 full-duplex operation standard. When connected to the Switch Engine, the option exists to operate at 200Mbps (turbo mode).

The Host MAC provides programmable enhanced features designed to minimize host supervision, bus utilization, and pre- or post-message processing. These features include the ability to disable retries after a collision, dynamic Frame Check Sequence (FCS) generation on a frame-by-frame basis, automatic pad field insertion and deletion to enforce minimum frame size attributes, and automatic retransmission and detection of collision frames. The Host MAC can sustain transmission or reception of minimally-sized back-to-back packets at full line speed with an interpacket gap (IPG) of 9.6 microseconds for 10 Mbps and 0.96/0.48 microseconds for 100/200 Mbps.

The primary attributes of the Host MAC are:

- Transmit and receive message data encapsulation
- Framing (frame boundary delimitation, frame synchronization)
- Error detection (physical medium transmission errors)
- Media access management
- Medium allocation (collision detection, except in full-duplex operation)
- Contention resolution (collision handling, except in full-duplex operation)
- Flow control during full-duplex mode
- Decoding of control frames (PAUSE command) and disabling the transmitter
- Generation of control frames
- Interface between the Host Bus Interface and the Ethernet PHYs/Switch Fabric.

11.2 Flow Control

The Host MAC supports full-duplex flow control using the pause operation and control frame. Half-duplex flow control using back pressure is also supported. The Host MAC flow control is configured via the memory mapped [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#) located in the System CSR space and the [Host MAC Flow Control Register \(HMAC_FLOW\)](#) located in the Host MAC CSR space.

Note: The Host MAC controls the flow between the switch fabric and the Host MAC, not the network flow control. The switch fabric handles the network flow control independently.

11.2.1 FULL-DUPLEX FLOW CONTROL

The pause operation inhibits transmission of data frames for a specified period of time. A pause operation consists of a frame containing the globally assigned multicast address (01-80-C2-00-00-01), the PAUSE opcode, and a parameter indicating the quantum of slot time (512 bit times) to inhibit data transmissions. The PAUSE parameter may range from 0 to 65,535 slot times. The Host MAC logic, upon receiving a frame with the reserved multicast address and PAUSE opcode, inhibits data frame transmissions for the length of time indicated. If a Pause request is received while a transmission is in progress, then the pause will take effect after the transmission is complete. Control frames are received, processed by the Host MAC, and passed on.

The device will automatically transmit pause frames based on the settings of the [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#) and the [Host MAC Flow Control Register \(HMAC_FLOW\)](#). When the RX FIFO reaches the level set in the [Automatic Flow Control High Level \(AFC_HI\)](#) field of [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#), the device will transmit a pause frame. The pause time field that is transmitted is set in the [Pause Time \(FCPT\)](#) field of the [Host MAC Flow Control Register \(HMAC_FLOW\)](#) register. When the RX FIFO drops below the level set in the [Automatic Flow Control Low Level \(AFC_LO\)](#) field of [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#), the device will automatically transmit a pause frame with a pause time of zero. The device will only send another pause frame when the RX FIFO level falls below [Automatic Flow Control Low Level \(AFC_LO\)](#) and then exceeds [Automatic Flow Control High Level \(AFC_HI\)](#) again.

11.2.2 HALF-DUPLEX FLOW CONTROL (BACKPRESSURE)

In half-duplex mode, backpressure is used for flow control. When the RX FIFO reaches the level set in the [Automatic Flow Control High Level \(AFC_HI\)](#) field of [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#), the Host MAC will be enabled to collide with incoming frames. The Host MAC transmit logic enters a state at the end of current transmission (if any), where it waits for the beginning of a receive frame. Based on the settings in [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#), backpressure can be enabled on any frame, a broadcast frame, any multicast frame or frames that match the stations address decoding logic.

In order to avoid any late collisions the Host MAC only generates collision-based backpressure at the start of a new frame. Once a new receive frame starts, the Host MAC intentional transmits which will result in a collision. Upon sensing the collision, the remote station will back off its transmission. Following the transmission of the intentional collision, the Host MAC waits for the next receive frame.

This pattern continues until either the RX FIFO drops below the level set in the [Automatic Flow Control High Level \(AFC_HI\)](#) or until the duration specified by [Backpressure Duration \(BACK_DUR\)](#) in field of [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#) is reached. In either case, the Host MAC will allow one frame to be received before returning to backpressure operation. Note that the Backpressure Duration is timed from when the RX FIFO reaches the level set in [Automatic Flow Control High Level \(AFC_HI\)](#), regardless when or if actual backpressure occurs.

11.2.3 HARDWIRED FLOW CONTROL

While Pause and Backpressure flow control are acceptable means for the Host MAC to flow control the internal switch fabric connection (and visa-versa), superior performance is achieved by using the hardwired flow control connection, since the normal latency of pause flow control packet transmission is avoided.

When enabled with the [Port 0 Hard-wired Flow Control \(HW_FC_0\)](#) bit in the [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#), hardwired flow control will stop the switch fabric transmitter from starting the next frame (the current frame is not affected) when the Host MAC indicates that it short on buffer space. The buffer levels are configured using the [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#). The other methods of flow control (FCANY, FCADD, FCBRD and FCMULT in the AFC_CFG register) should be disabled when using hardwired flow control.


Hardwired flow control also will stop the Host MAC transmitter from starting the next frame (the current frame is not affected) when the switch fabric port indicates that it short on buffer space. The other methods of flow control within the switch (TX_FC_0, RX_FC_0 and BP_EN_0 in the MANUAL_FC_0 register) should be disabled when using hardwired flow control.

11.3 Virtual Local Area Network (VLAN) Support

Virtual Local Area Networks (VLANs), as defined within the IEEE 802.3 standard, provide network administrators a means of grouping nodes within a larger network into broadcast domains. To implement a VLAN, four extra bytes are added to the basic Ethernet packet. As shown in [Figure 11-1](#), the four bytes are inserted after the Source Address Field and before the Type/Length field. The first two bytes of the VLAN tag identify the tag, and by convention are set to the value 0x8100. The last two bytes identify the specific VLAN associated with the packet and provide a priority field.

The device supports VLAN-tagged packets and provides two Host MAC registers, [Host MAC VLAN1 Tag Register \(HMAC_VLAN1\)](#) and [Host MAC VLAN2 Tag Register \(HMAC_VLAN2\)](#), which are used to identify VLAN-tagged packets. The HMAC_VLAN1 register is used to specify the VLAN1 tag which will increase the legal frame length from 1518 to 1522 bytes. The HMAC_VLAN2 register is used to specify the VLAN2 tag which will increase the legal frame length from 1518 to 1538 bytes. If a packet arrives bearing either of these tags in the two bytes succeeding the Source Address field, the controller will recognize the packet as a VLAN-tagged packet, allowing the packet to be received and processed by the host software. If both VLAN1 and VLAN2 tag Identifiers are used, each should be unique. If both are set to the same value, VLAN1 is given higher precedence and the maximum legal frame length is set to 1522.

FIGURE 11-1: VLAN FRAME



11.4 Address Filtering

The Ethernet address fields of an Ethernet packet consist of two 6-byte fields: one for the destination address and one for the source address. The first bit of the destination address signifies whether it is a physical address or a multicast address.

The Host MAC address check logic filters the frame based on the Ethernet receive filter mode that has been enabled. The various filter modes of the Host MAC are specified based on the state of the control bits in the [Host MAC Control Register \(HMAC_CR\)](#), as shown in [TABLE 11-1](#). Please refer to the [Section 11.14.1, "Host MAC Control Register \(HMAC_CR\)," on page 203](#) for more information on this register.

Frames that fail the address filtering are accepted only if the [Receive All Mode \(RXALL\)](#) bit in the [Receive Configuration Register \(RX_CFG\)](#) is set. The Filtering Fail bit in the RX Status will be set for these frames.

Note: This filtering function is performed after any switch fabric filtering functions. The user must ensure the switch filtering is setup properly to allow packets to be passed to the Host MAC for further filtering.

TABLE 11-1: ADDRESS FILTERING MODES

MCPAS	PRMS	INVFILT	HO	HPFILT	Description
0	0	0	0	0	Perfect - MAC address perfect filtering only for all addresses. Broadcast frames accepted if BCAST is low.
0	0	0	0	1	Hash Perfect - MAC address perfect filtering for physical address and hash filtering for multicast addresses. Broadcast frames accepted if BCAST is low.
0	0	0	1	1	Hash Only - Hash Filtering for physical and multicast addresses. Broadcast frames accepted if BCAST is low.
0	0	1	0	0	Inverse Filtering
X	1	0	0	X	Promiscuous
0			X	1	
1	0	0	0	X	Perfect all Multicast - Pass all multicast frames including broadcasts. Frames with physical addresses are perfect-filtered
1	0	0	1	1	Hash Only all Multicast - Pass all multicast frames including broadcasts. Frames with physical addresses are hash-filtered

11.4.1 PERFECT FILTERING

This filtering mode passes only incoming frames whose destination address field exactly matches the value programmed into the [Host MAC Address High Register \(HMAC_ADDRH\)](#) and the [Host MAC Address Low Register \(HMAC_ADDRL\)](#). The MAC address is formed by the concatenation of these two registers.

Broadcast frames are also accepted if [Disable Broadcast Frames \(BCAST\)](#) is low.

Note: If the HMAC_ADDRH and HMAC_ADDRL registers are set to the Broadcast address, Broadcast frames will be accepted regardless of the setting of the BCAST bit.

11.4.2 HASH ONLY FILTERING

This type of filtering checks for incoming receive packets (from switch Port 0) with either multicast or physical destination addresses, and executes an imperfect address filtering against the hash table. The hash table is formed by merging the values in the [Host MAC Multicast Hash Table High Register \(HMAC_HASHH\)](#) and the [Host MAC Multicast Hash Table Low Register \(HMAC_HASHL\)](#) to form a 64-bit hash table.

During imperfect hash filtering, the destination address in the incoming frame is passed through the CRC logic and the upper 6-bits of the CRC register are used to index the contents of the hash table. The most significant bit determines the register to be used (HMAC_HASHH or HMAC_HASHL), while the other five bits determine the bit within the register. A value of 00000 selects Bit 0 of the HMAC_HASHL register and a value of 11111 selects Bit 31 of the HMAC_HASHH register.

Broadcast frames are also accepted if [Disable Broadcast Frames \(BCAST\)](#) is low.

11.4.3 HASH PERFECT FILTERING

Note: If bit 31 of HMAC_HASHH is set, the Broadcast address will cause a hash match and Broadcast frames will be accepted regardless of the setting of the BCAST bit.
--

In hash perfect filtering, if the received frame is a physical address, the Host MAC packet filter will perfect-filter the incoming frame's destination field with the value programmed into the [Host MAC Address High Register \(HMAC_ADDRH\)](#) and the [Host MAC Address Low Register \(HMAC_ADDRL\)](#). However, if the incoming frame is a multicast frame, the Host MAC packet filter function performs an imperfect address filtering against the hash table.

The imperfect filtering against the hash table is the same imperfect filtering process described in [Section 11.4.2, "Hash Only Filtering"](#).

Broadcast frames are also accepted if [Disable Broadcast Frames \(BCAST\)](#) is low.

Note: If bit 31 of HMAC_HASHH is set, the Broadcast address will cause a hash match and Broadcast frames will be accepted regardless of the setting of the BCAST bit.

11.4.4 INVERSE FILTERING

In inverse filtering, the Host MAC packet filter accepts incoming frames (from switch Port 0) with a destination address not matching the perfect address (i.e., the value programmed into the [Host MAC Address High Register \(HMAC_ADDRH\)](#) and the [Host MAC Address Low Register \(HMAC_ADDRL\)](#)) and rejects frames with destination addresses matching the perfect address.

Note: If the HMAC_ADDRH and HMAC_ADDRL registers are set to the Broadcast address, Broadcast frames will be filtered regardless of the setting of the BCAST bit.

11.4.5 PROMISCUOUS

When the [Promiscuous Mode \(PRMS\)](#) bit is set, all frames are accepted regardless of their destination address.

Note: Broadcast frames will be accepted regardless of the setting of the BCAST bit.

11.4.6 PERFECT FILTERING ALL MULTICAST

If the received frame is a physical address, the Host MAC packet filter will perfect-filter the incoming frame's destination field with the value programmed into the [Host MAC Address High Register \(HMAC_ADDRH\)](#) and the [Host MAC Address Low Register \(HMAC_ADDRL\)](#).

With the [Pass All Multicast \(MCPAS\)](#) bit of the [Host MAC Control Register \(HMAC_CR\)](#) set, all multicast frames are accepted. This includes all broadcast frames as well.

11.4.7 HASH ONLY FILTERING ALL MULTICAST

If the received frame is a physical address, the Host MAC packet filter will execute an imperfect address filtering against the hash table. The imperfect filtering against the hash table is the same imperfect filtering process described in [Section 11.4.2, "Hash Only Filtering"](#).

With the [Pass All Multicast \(MCPAS\)](#) bit of the [Host MAC Control Register \(HMAC_CR\)](#) set, all multicast frames are accepted. This includes all broadcast frames as well.

11.5 Frame Filtering

Following the address filtering, frames are accepted or rejected according to the following:

- Good frames that pass the address filtering are accepted. In the RX Status for these frames, the Filtering Fail bit will be clear (since the frame passed the address filter). The Packet Filter bit will be set.
- Good frames that fail the address filtering are accepted if the [Receive All Mode \(RXALL\)](#) bit in the [Receive Configuration Register \(RX_CFG\)](#) is set. In the RX Status for these frames, the Filtering Fail bit will be set (since the frame failed the address filter). The Packet Filter bit will also be set (since the RXALL bit allowed the acceptance of the frame).
- A good Broadcast frame is accepted if it passes the address filtering or if the RXALL bit is set. The [Disable Broadcast Frames \(BCAST\)](#) bit in the [Receive Configuration Register \(RX_CFG\)](#) determines if the Packet Filter bit in the RX Status is set (BCAST=0) or cleared (BCAST=1). Note that [Disable Broadcast Frames \(BCAST\)](#) doesn't cause the frame to be dropped if it passes the address filtering or if the RXALL bit is set. The Filtering Fail bit will indicate if the address filtering passed or failed.
- A good Control frame is accepted if it passes the address filtering or if the RXALL bit is set. The [Pass Control Frames \(FCPASS\)](#) bit in the [Host MAC Flow Control Register \(HMAC_FLOW\)](#) determines if the Packet Filter bit in the RX Status is set (FCPASS=1) or cleared (FCPASS=0). Note that [Pass Control Frames \(FCPASS\)](#) being low doesn't cause the frame to be dropped if it passes the address filtering or if the RXALL bit is set. The Filtering Fail bit will indicate if the address filtering passed or failed.
- A frame that has an error (runt, collision, CRC, too long) and is greater than 60 bytes in length is accepted if it passes the address filtering or if the RXALL bit is set. The [Pass Bad Frames \(PASSBAD\)](#) bit in the [Receive Configuration Register \(RX_CFG\)](#) determines if the Packet Filter bit in the RX Status is set (PASSBAD=1) or cleared (PASSBAD=0). The Filtering Fail bit will indicate if the address filtering passed or failed.
- A frame that has an error (runt, collision, CRC) and is 60 bytes or under in length is accepted if it passes the address filtering or if the RXALL bit is set and the [Pass Bad Frames \(PASSBAD\)](#) bit is set. The Packet Filter bit in the RX Status is set for these frames. The Filtering Fail bit will indicate if the address filtering passed or failed.

11.6 Wake-On-LAN (WOL)

The following bits of the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), when enabled, may allow a WOL event to be asserted:

- [Perfect DA Wakeup Enable \(PFDA_EN\)](#)
- [Broadcast Wakeup Enable \(BCST_EN\)](#)
- [Wake-Up Frame Enable \(WUEN\)](#)
- [Magic Packet Enable \(MPEN\)](#)

The [WoL Wait for Sleep \(WOL_WAIT_SLEEP\)](#) bit in [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) will delay the WoL functions until the host has put the device into a power down mode.

WOL events may be indicated to the power management block via the [Wake On Status \(WOL_STS\)](#) bit of the [Power Management Control Register \(PMT_CTRL\)](#). Each WOL event type is detailed in the following sub-sections.

The WoL feature is part of the broader power management features of the device and can be used to trigger the power management event output pin (**PME**) or general interrupt request pin (**IRQ**). This is accomplished by enabling the desired WoL feature and setting the [Wake-On-Enable \(WOL_EN\)](#) bit of the [Power Management Control Register \(PMT_CTRL\)](#). Refer to [Section 6.3, "Power Management," on page 49](#) for additional information.

11.6.1 PERFECT DA DETECTION

Setting the [Perfect DA Wakeup Enable \(PFDA_EN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) places the MAC in the Perfect DA detection mode. In this mode, normal data reception is disabled, and detection logic within the MAC examines the destination address of each received frame.

When a frame whose destination address matches that specified by the [Host MAC Address High Register \(HMAC_AD_DRH\)](#) and [Host MAC Address Low Register \(HMAC_ADDRL\)](#) is received, the [Perfect DA Frame Received \(PFDA_FR\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) is set. When the host clears the PFDA_EN bit, the Host MAC will resume normal receive operation.

Note: The switch fabric must be configured to pass these packets to the Host MAC for this function to operate properly.

11.6.2 BROADCAST DETECTION

Setting the [Broadcast Wakeup Enable \(BCST_EN\)](#) bit in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) places the MAC in the Broadcast detection mode. In this mode, normal data reception is disabled, and detection logic within the MAC examines the destination address of each received frame.

When a frame whose destination address is FF FF FF FF FF FF is received, the [Broadcast Frame Received \(BCAST_FR\)](#) bit in the WUCSR is set. When the host clears the BCST_EN bit, the Host MAC will resume normal receive operation.

Note: The switch fabric must be configured to pass these packets to the Host MAC for this function to operate properly.

11.6.3 WAKE-UP FRAME DETECTION

Eight programmable wakeup frame filters are supported. Each filter has a 128-bit byte mask that indicates which bytes of the frame should be compared by the MAC. A CRC-16 is calculated over these bytes. The result is then compared with the filter's respective CRC-16 to determine if a match exists.

Setting the [Wake-Up Frame Enable \(WUEN\)](#) in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), places the Host MAC in the wake-up frame detection mode. In this mode, normal data reception is disabled, and detection logic within the Host MAC examines received data for the pre-programmed wake-up frame patterns.

Upon detection, the [Remote Wake-Up Frame Received \(WUFR\)](#) in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) register is set. When the host clears the WUEN bit, the Host MAC will resume normal receive operation.

Before putting the Host MAC into the wake-up frame detection state, the host must provide the detection logic with a list of sample frames and their corresponding byte masks. This information must be written into the [Host MAC Wake-up Frame Filter Register \(HMAC_WUFF\)](#). The wake-up frame filter is configured through this register using an index mechanism. After power-on reset, hardware reset, or soft reset, the Host MAC loads the first value written to the HMAC_WUFF register to the first DWORD in the wake-up frame filter (filter 0 byte mask 0). The second value written to this register is loaded to the second DWORD in the wake-up frame filter (filter 0 byte mask 1) and so on for all 40 DWORDs. The wake-up frame filter functionally is described below.

The Host MAC supports eight programmable 128-bit wake-up filters that support many different receive packet patterns. If remote wake-up mode is enabled, the remote wake-up function receives all frames addressed to the Host MAC. It then checks each frame against the enabled filter and recognizes the frame as a remote wake-up frame if it passes the wakeup frame filter register's address filtering and CRC value match.

In order to determine which bytes of the frames should be checked by the CRC module, the Host MAC uses a programmable byte mask and a programmable pattern offset for each of the eight supported filters.

The pattern's offset defines the location of the first byte that should be checked in the frame. Since the destination address is checked by the address filtering function, the pattern offset is always greater than 12.

The byte mask is a 128-bit field that specifies whether or not each of the 128 contiguous bytes within the frame, beginning in the pattern offset, should be checked. If bit *j* in the byte mask is set, the detection logic checks byte offset *+j* in the frame. In order to load the wake-up frame filter, the host must perform 40 writes to the [Host MAC Wake-up Frame Filter Register \(HMAC_WUFF\)](#). The contents of the [Host MAC Wake-up Frame Filter Register \(HMAC_WUFF\)](#) may be obtained by reading all 40 DWORDs. [Table 11-2](#) shows the wake-up frame filter register's structure.

At the completion of the CRC-16 checking process, the CRC-16 calculated using the pattern offset and byte mask is compared to the expected CRC-16 value associated with the filter. If a match occurs, a remote wakeup event is signaled.

Note: The switch fabric must be configured to pass these packets to the Host MAC for this function to operate properly.

TABLE 11-2: WAKEUP FRAME FILTER REGISTER STRUCTURE

Filter 0 Byte Mask 0
Filter 0 Byte Mask 1
Filter 0 Byte Mask 2
Filter 0 Byte Mask 3
Filter 1 Byte Mask 0
Filter 1 Byte Mask 1
Filter 1 Byte Mask 2
Filter 1 Byte Mask 3
Filter 2 Byte Mask 0
Filter 2 Byte Mask 1
Filter 2 Byte Mask 2
Filter 2 Byte Mask 3
Filter 3 Byte Mask 0
Filter 3 Byte Mask 1
Filter 3 Byte Mask 2
Filter 3 Byte Mask 3
Filter 4 Byte Mask 0
Filter 4 Byte Mask 1
Filter 4 Byte Mask 2
Filter 4 Byte Mask 3
Filter 5 Byte Mask 0
Filter 5 Byte Mask 1
Filter 5 Byte Mask 2
Filter 5 Byte Mask 3
Filter 6 Byte Mask 0
Filter 6 Byte Mask 1
Filter 6 Byte Mask 2
Filter 6 Byte Mask 3
Filter 7 Byte Mask 0
Filter 7 Byte Mask 1
Filter 7 Byte Mask 2
Filter 7 Byte Mask 3

TABLE 11-2: WAKEUP FRAME FILTER REGISTER STRUCTURE (CONTINUED)

Reserved	Filter 3 Command	Reserved	Filter 2 Command	Reserved	Filter 1 Command	Reserved	Filter 0 Command
Reserved	Filter 7 Command	Reserved	Filter 6 Command	Reserved	Filter 5 Command	Reserved	Filter 4 Command
Filter 3 Offset		Filter 2 Offset		Filter 1Offset		Filter 0 Offset	
Filter 7 Offset		Filter 6 Offset		Filter 5 Offset		Filter 4 Offset	
Filter 1 CRC-16				Filter 0 CRC-16			
Filter 3 CRC-16				Filter 2 CRC-16			
Filter 5 CRC-16				Filter 4 CRC-16			
Filter 7 CRC-16				Filter 6 CRC-16			

The Filter i Byte Mask defines which incoming frame bytes Filter i will examine to determine whether or not this is a Wakeup Frame. [Table 11-3](#), describes the byte mask's bit fields.

Filter x Mask 0 corresponds to bits [31:0]. Where the lsb corresponds to the first byte on the wire.

Filter x Mask 1 corresponds to bits [63:32]. Where the lsb corresponds to the first byte on the wire.

Filter x Mask 2 corresponds to bits [95:64]. Where the lsb corresponds to the first byte on the wire.

Filter x Mask 3 corresponds to bits [127:96]. Where the lsb corresponds to the first byte on the wire.

The following tables define elements common to both WUFF register structures.

TABLE 11-3: FILTER I BYTE MASK BIT DEFINITIONS

Filter i Byte Mask Description	
Bits	Description
127:0	Byte Mask: If bit j of the byte mask is set, the CRC machine processes byte <i>pattern-offset + j</i> of the incoming frame. Otherwise, byte <i>pattern-offset + j</i> is ignored.

The Filter i command register controls Filter i operation. [Table 11-4](#) shows the Filter I command register.

TABLE 11-4: FILTER I COMMAND BIT DEFINITIONS

Filter i Commands	
Bits	Description
3:2	Address Type: Defines the destination address type of the pattern. 00 = Pattern applies only to unicast frames. 10 = Pattern applies only to multicast frames. X1 = Pattern applies to all frames that have passed the regular receive filter.
1	RESERVED
0	Enable Filter: When bit is set, Filter i is enabled, otherwise, Filter i is disabled.

The Filter i Offset register defines the offset in the frame's destination address field from which the frames are examined by Filter i. [Table 11-5](#) describes the Filter i Offset bit fields.

TABLE 11-5: FILTER I OFFSET BIT DEFINITIONS

Filter i Offset Description	
Bits	Description
7:0	Pattern Offset: The offset of the first byte in the frame on which CRC is checked for Wakeup Frame recognition. The MAC checks the first offset byte of the frame for CRC and checks to determine whether the frame is a Wakeup Frame. Offset 0 is the first byte of the incoming frame's destination address.

The Filter i CRC-16 register contains the CRC-16 result of the frame that should pass Filter i.

[TABLE 11-6:](#) describes the Filter i CRC-16 bit fields.

The CRC-16 is calculated as follows:

At the start of a frame, CRC-16 is initialized with the value FFFFh. CRC-16 is updated when the pattern offset and mask indicate the received byte is part of the checksum calculation. The following algorithm is used to update the CRC-16 at that time:

Let:

^ denote the exclusive or operator.

Data [7:0] be the received data byte to be included in the checksum.

CRC[15:0] contain the calculated CRC-16 checksum.

F0 ... F7 be intermediate results, calculated when a data byte is determined to be part of the CRC-16.

Calculate:

$$F0 = \text{CRC}[15] \wedge \text{Data}[0]$$

$$F1 = \text{CRC}[14] \wedge F0 \wedge \text{Data}[1]$$

$$F2 = \text{CRC}[13] \wedge F1 \wedge \text{Data}[2]$$

$$F3 = \text{CRC}[12] \wedge F2 \wedge \text{Data}[3]$$

$$F4 = \text{CRC}[11] \wedge F3 \wedge \text{Data}[4]$$

$$F5 = \text{CRC}[10] \wedge F4 \wedge \text{Data}[5]$$

$$F6 = \text{CRC}[09] \wedge F5 \wedge \text{Data}[6]$$

$$F7 = \text{CRC}[08] \wedge F6 \wedge \text{Data}[7]$$

The CRC-16 is updated as follows:

$$\text{CRC}[15] = \text{CRC}[7] \wedge F7$$

$$\text{CRC}[14] = \text{CRC}[6]$$

$$\text{CRC}[13] = \text{CRC}[5]$$

$$\text{CRC}[12] = \text{CRC}[4]$$

$$\text{CRC}[11] = \text{CRC}[3]$$

$$\text{CRC}[10] = \text{CRC}[2]$$

$$\text{CRC}[9] = \text{CRC}[1] \wedge F0$$

$$\text{CRC}[8] = \text{CRC}[0] \wedge F1$$

$$\text{CRC}[7] = F0 \wedge F2$$

$$\text{CRC}[6] = F1 \wedge F3$$

$$\text{CRC}[5] = F2 \wedge F4$$

$$\text{CRC}[4] = F3 \wedge F5$$

CRC[3] = F4 ^ F6

CRC[2] = F5 ^ F7

CRC[1] = F6

CRC[0] = F7

TABLE 11-6: FILTER I CRC-16 BIT DEFINITIONS

Filter i CRC-16 Description	
Bits	Description
15:0	Pattern CRC-16: This field contains the 16-bit CRC value from the pattern and the byte mask programmed to the Wakeup Filter register function. This value is compared against the CRC calculated on the incoming frame, and a match indicates the reception of a Wakeup Frame.

Table 11-7 indicates the cases that produce a wake when the [Wake-Up Frame Enable \(WUEN\)](#) bit of the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) is set. All other cases do not generate a wake.

TABLE 11-7: WAKEUP GENERATION CASES

Filter Enabled (Note 1)	Frame Type	CRC Match (Note 2)	Global Unicast Enabled (Note 3)	Pass Regular Receive Filter	Address Type (Note 4)
Yes	Unicast	Yes	Yes	x	x
Yes	Unicast	Yes	x	Yes	Unicast (=00)
Yes	Multicast	Yes	x	Yes	Multicast (=10)
Yes	Broadcast (Note 5)	Yes	x	x	x
Yes	x	Yes	x	Yes	Passed Receive Filter (=x1b)

Note 1: As determined by bit 0 of Filter i Command.

Note 2: CRC matches Filter i CRC-16 field.

Note 3: As determined by bit 9 of WUCSR.

Note 4: As determined by bits 3:2 of Filter i Command.

Note 5: When wake-up frame detection is enabled via the [Wake-Up Frame Enable \(WUEN\)](#) bit of the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#), a broadcast wake-up frame will wake-up the device despite the state of the Disable Broadcast Frames (BCAST) bit in the [Host MAC Control Register \(HMAC_CR\)](#).

Note: x indicates “don’t care”.

11.6.4 MAGIC PACKET DETECTION

Setting the Magic Packet Enable bit (MPEN) in the [Host MAC Wake-up Control and Status Register \(HMAC_WUCSR\)](#) places the Host MAC in the “Magic Packet” detection mode. In this mode, normal data reception is disabled, and detection logic within the Host MAC examines received data for a Magic Packet.

Upon detection, the Magic Packet Received bit (MPR) in the HMAC_WUCSR register is set. When the host clears the MPEN bit, the Host MAC will resume normal receive operation.

In Magic Packet mode, the Host MAC constantly monitors each frame addressed to the node for a specific Magic Packet pattern. Only packets passing the Address Filtering check of [Section 11.4](#) (see [Note 6](#)) are checked for the Magic Packet requirements. Once the address requirement has been met, the Host MAC checks the received frame for the pattern 48'hFF_FF_FF_FF_FF_FF after the destination and source address field. The Host MAC then looks in the frame for 16 repetitions of the Host MAC address without any breaks or interruptions. In case of a break in the 16 address repetitions, the Host MAC again scans for the 48'hFF_FF_FF_FF_FF_FF pattern in the incoming frame. The 16 repetitions may be anywhere in the frame but must be preceded by the synchronization stream. The device will also accept a multicast frame, as long as it detects the 16 duplications of the Host MAC address.

For example, if the Host MAC address is 00h 11h 22h 33h 44h 55h, then the MAC scans for the following data sequence in an Ethernet frame:

```
Destination Address Source Address .....FF FF FF FF FF FF
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
...CRC
```

Note: The switch fabric must be configured to pass these packets to the Host MAC for this function to operate properly.

Note 6: Normally, for Magic Packet Detection, address filtering should be set for [Perfect Filtering](#) or [Hash Perfect Filtering](#).

11.7 Receive Checksum Offload Engine (RXCOE)


The receive checksum offload engine provides assistance to the Host by calculating a 16-bit checksum for a received Ethernet frame. The RXCOE readily supports the following IEEE802.3 frame formats:

- Type II Ethernet frames
- SNAP encapsulated frames
- Support for up to 2, 802.1q VLAN tags

The resulting checksum value can also be modified by software to support other frame formats.

The RXCOE has two modes of operation. In mode 0, the RXCOE calculates the checksum between the first 14 bytes of the Ethernet frame and the FCS. This is illustrated in [Figure 11-2](#).

FIGURE 11-2: RXCOE CHECKSUM CALCULATION



In mode 1, the RXCOE supports VLAN tags and a SNAP header. In this mode, the RXCOE calculates the checksum at the start of L3 packet. The VLAN1 tag register is used by the RXCOE to indicate what protocol type is to be used to indicate the existence of a VLAN tag. This value is typically 8100h.

Example frame configurations:

FIGURE 11-3: TYPE II ETHERNET FRAMES

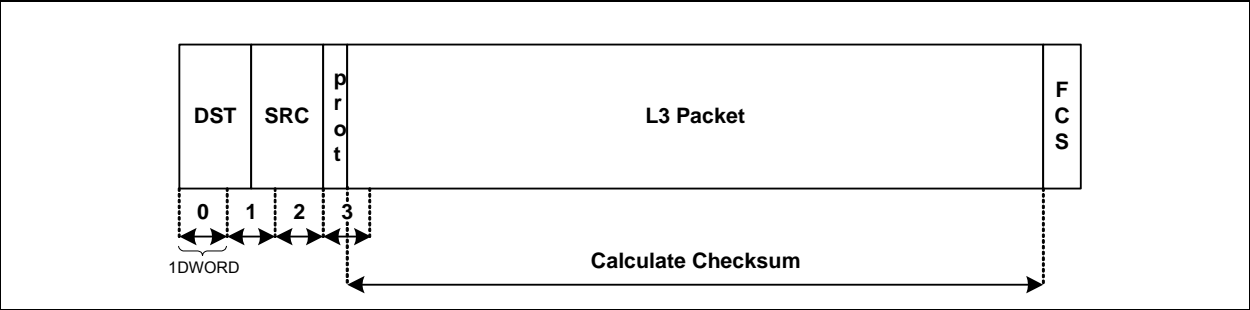


FIGURE 11-4: ETHERNET FRAME WITH VLAN TAG




FIGURE 11-5: ETHERNET FRAME WITH LENGTH FIELD AND SNAP HEADER

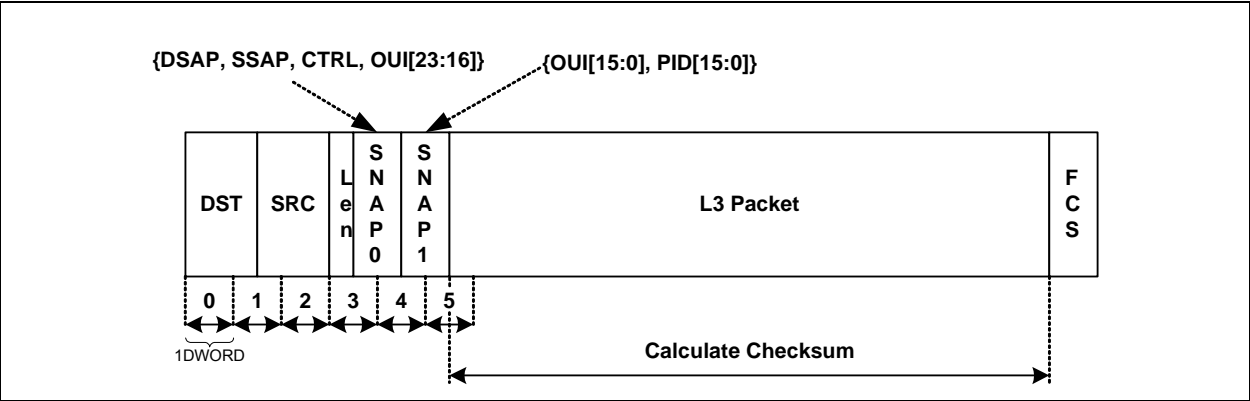


FIGURE 11-6: ETHERNET FRAME WITH VLAN TAG AND SNAP HEADER


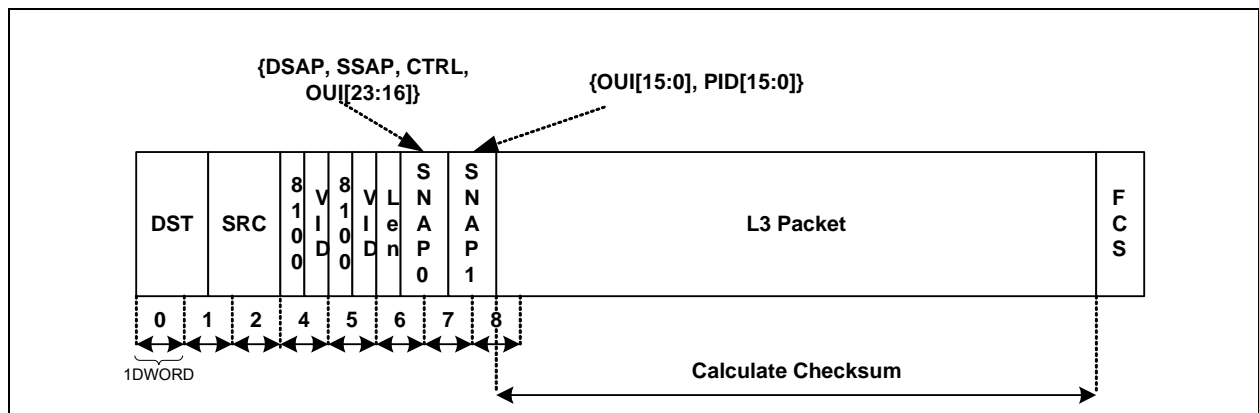


FIGURE 11-7: ETHERNET FRAME WITH MULTIPLE VLAN TAGS AND SNAP HEADER



The RXCOE supports a maximum of two VLAN tags. If there are more than two VLAN tags, the VLAN protocol identifier for the third tag is treated as an Ethernet type field. The checksum calculation will begin immediately after the type field.

Note that in all cases, the checksum calculation ends just before the frames FCS field. In the case where padding is added to meet the minimum frame length requirement, the checksum calculation will also include the pads byte(s). This may lead to unexpected results if the pad byte(s) are not zero.

The RXCOE resides in the RX path within the MAC. As the RXCOE receives an Ethernet frame, it calculates the 16-bit checksum. The RXCOE passes the Ethernet frame to the RX FIFO with the checksum appended to the end of the frame. The RXCOE inserts the checksum immediately after the last byte of the Ethernet frame and before it transmits the status word. The packet length field in the RX Status Word (refer to [Section 11.12.3](#)) will indicate that the frame size has increased by two bytes to accommodate the checksum.

Note: When enabled, the RXCOE calculates a checksum for every received frame.

Setting the [RX Checksum Offload Engine Enable \(RX_COE_EN\)](#) bit in the [Host MAC Checksum Offload Engine Control Register \(HMAC_COE_CR\)](#) enables the RXCOE, while the [RX Checksum Offload Engine Mode \(RX_COE_MODE\)](#) bit selects the operating mode. When the RXCOE is disabled, the received data is simply passed through the RXCOE unmodified.

Note: Software applications must stop the receiver and flush the RX data path before changing the state of the [RX Checksum Offload Engine Enable \(RX_COE_EN\)](#) or [RX Checksum Offload Engine Mode \(RX_COE_MODE\)](#) bits.

Note: When the RXCOE is enabled, automatic pad stripping must be disabled ([Automatic Pad Stripping \(PAD-STR\)](#) bit of the [Host MAC Control Register \(HMAC_CR\)](#)) and vice versa. These functions cannot be enabled simultaneously.

11.7.1 RX CHECKSUM CALCULATION

The checksum is calculated 16 bits at a time. In the case of an odd sized frame, an extra byte of zero is used to pad up to 16 bits.

Consider the following packet: DA, SA, Type, B0, B1, B2 ... BN, FCS

Let $[A, B] = A \times 256 + B$;

If the packet has an even number of octets then

checksum = $[B1, B0] + C0 + [B3, B2] + C1 + \dots + [BN, BN-1] + CN-1$

Where C0, C1, ... CN-1 are the carry out results of the intermediate sums.

If the packet has an odd number of octets then

checksum = $[B1, B0] + C0 + [B3, B2] + C1 + \dots + [0, BN] + CN-1$

11.8 Transmit Checksum Offload Engine (TXCOE)

The transmit checksum offload engine provides assistance to the CPU by calculating a 16-bit checksum, typically for TCP, for a transmit Ethernet frame. The TXCOE calculates the checksum and inserts the results back into the data stream as it is transferred to the MAC.

To activate the TXCOE and perform a checksum calculation, the Host must first set the [TX Checksum Offload Engine Enable \(TX_COE_EN\)](#) bit in the [Host MAC Checksum Offload Engine Control Register \(HMAC_COE_CR\)](#). The Host then pre-pends a 3 DWORD buffer to the data that will be transmitted. The prepended buffer includes a TX Command A, TX Command B, and a 32-bit TX checksum preamble (refer to [Table 11-8](#)). When the CK bit of the TX Command 'B' is set in conjunction with the FS bit of TX Command 'A' and the [TX Checksum Offload Engine Enable \(TX_COE_EN\)](#) bit of the [Host MAC Checksum Offload Engine Control Register \(HMAC_COE_CR\)](#) register, the TXCOE will perform a checksum calculation on the associated packet. The TX checksum preamble instructs the TXCOE on the handling of the associated packet. The [TXCSSP - TX Checksum Start Pointer](#) field of the TX checksum preamble defines the byte offset at which the data checksum calculation will begin. The checksum calculation will begin at this offset and will continue until the end of the packet. The data checksum calculation must not begin in the MAC header (first 14 bytes) or in the last 4 bytes of the TX packet. When the calculation is complete, the checksum will be inserted into the packet at the byte offset defined by the [TXCSLOC - TX Checksum Location](#) field of the TX checksum preamble. The TX checksum cannot be inserted in the MAC header (first 14 bytes) or in the last 4 bytes of the TX packet. If the CK bit is not set in the first TX Command 'B' of a packet, the packet is passed directly through the TXCOE without modification, regardless if the TXCOE_EN is set. An example of a TX packet with a prepended TX checksum preamble can be found in [Section 11.11.6.3, "TX Example 3"](#). In this example, the Host provides the Ethernet frame to the Ethernet controller in four fragments, the first containing the TX Checksum Preamble. [Figure 11-8](#) shows how these fragments are loaded into the TX Data FIFO. For more information on the TX Command 'A' and TX Command 'B', refer to [Section 11.11.2, "TX Command Format," on page 174](#).

If the TX packet already includes a partial checksum calculation (perhaps inserted by an upper layer protocol), this checksum can be included in the hardware checksum calculation by setting the TXCSSP field in the TX checksum preamble to include the partial checksum. The partial checksum can be replaced by the completed checksum calculation by setting the TXCSLOC pointer to point to the location of the partial checksum.

TABLE 11-8: TX CHECKSUM PREAMBLE

Field	Description
31	TXCSUDP - TX Checksum UDP Frame This bit specifies if a checksum result of 0x0000 should be changed to 0xFFFF.
30:28	RESERVED
27:16	TXCSLOC - TX Checksum Location This field specifies the byte offset where the TX checksum will be inserted in the TX packet. The checksum will replace two bytes of data starting at this offset. The TX checksum cannot be inserted in the MAC header (first 14 bytes) or in the last 4 bytes of the TX packet.
15:12	RESERVED
11:0	TXCSSP - TX Checksum Start Pointer This field indicates start offset, in bytes, where the checksum calculation will begin in the associated TX packet. The data checksum calculation must not begin in the MAC header (first 14 bytes) or in the last 4 bytes of the TX packet.

Note: When the TXCOE is enabled, the third DWORD of the prepended packet is not transmitted. However, 4 bytes must be added to the packet length field in TX Command B.

Note: Software applications must stop the transmitter and flush the TX data path before changing the state of the TXCOE_EN bit. However, the CK bit of TX Command B can be set or cleared on a per-packet basis.

Note: The TXCOE_MODE may only be changed if the TX path is disabled. If it is desired to change this value during run time, it is safe to do so only after the TX Ethernet path is disabled and the TLI is empty.


Note: The TX checksum preamble must be DWORD-aligned.

Note: TX preamble size is accounted for in both the buffer length and packet length.

Note: The first buffer, which contains the TX preamble, may not contain any Ethernet frame data

Figure 11-8 illustrates the use of a prepended checksum preamble when transmitting an Ethernet frame consisting of 3 payload buffers.

FIGURE 11-8: TX EXAMPLE ILLUSTRATING A PREPENDED TX CHECKSUM PREAMBLE



11.8.1 TX CHECKSUM CALCULATION

The TX checksum calculation is performed using the same operation as the RX checksum shown in [Section 11.7.1](#), with the exception that the calculation starts as indicated by the TX checksum preamble and the transmitted checksum is the one's-compliment of the calculated value.

UDP checksums are optional under IPv4, and a checksum value of zero indicates to the receiver that no checksum was calculated. Under IPv6, however, according to RFC 2460, the UDP checksum is not optional. A calculated checksum that yields a result of zero must be changed to FFFFh for insertion into the UDP header. IPv6 receivers discard UDP packets containing a zero checksum. Bit 31 of the checksum preamble specifies if a result of 0x0000 should be changed to 0xFFFF. This allows the choice of checksum usage for UDP and other purposes.

11.9 Host MAC Address

The Host MAC address is configured via the [Host MAC Address Low Register \(HMAC_ADDRL\)](#) and [Host MAC Address High Register \(HMAC_ADDRH\)](#). These registers contain the 48-bit physical address of the Host MAC. The contents of these registers may be loaded directly by the host, or optionally, by the EEPROM Loader from EEPROM at power-on (if a programmed EEPROM is detected). The MAC address value loaded by the EEPROM Loader into the Host MAC address registers (for host packet unicast qualification), is also loaded into the Switch Fabric MAC address registers (for pause packet / flow control): [Switch Fabric MAC Address Low Register \(SWITCH_MAC_ADDRL\)](#) and [Switch Fabric MAC Address High Register \(SWITCH_MAC_ADDRH\)](#). These two sets of registers are loaded simultaneously via the same EEPROM byte addresses.

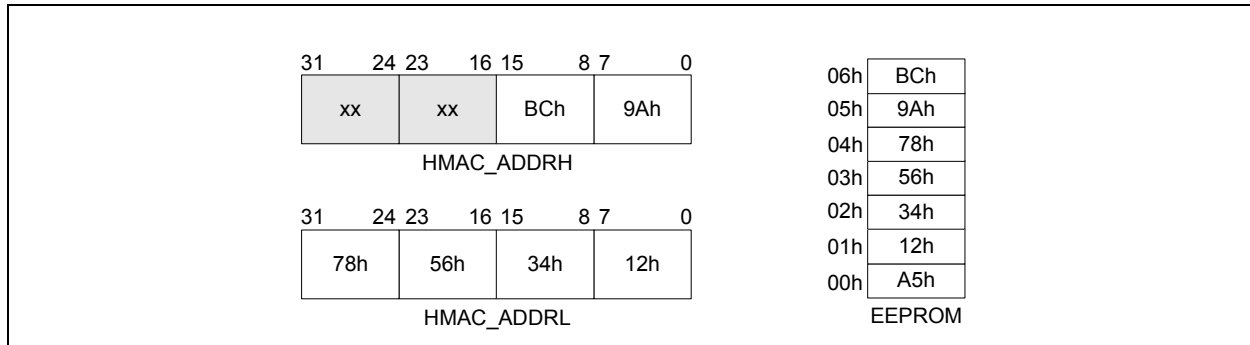
[Table 11-9](#) below illustrates the byte ordering of the HMAC_ADDRL/SWITCH_MAC_ADDRL and HMAC_ADDRH/SWITCH_MAC_ADDRH registers with respect to the reception of the Ethernet physical address. Also shown is the correlation between the EEPROM addresses and HMAC_ADDRL/SWITCH_MAC_ADDRL and HMAC_ADDRH/SWITCH_MAC_ADDRH registers.

TABLE 11-9: EEPROM BYTE ORDERING AND REGISTER CORRELATION

EEPROM Address	Register Locations Written	Order of Reception on Ethernet
01h	HMAC_ADDRL[7:0] SWITCH_MAC_ADDRL[7:0]	1 st
02h	HMAC_ADDRL[15:8] SWITCH_MAC_ADDRL[15:8]	2 nd
03h	HMAC_ADDRL[23:16] SWITCH_MAC_ADDRL[23:16]	3 rd
04h	HMAC_ADDRL[31:24] SWITCH_MAC_ADDRL[31:24]	4 th
05h	HMAC_ADDRH[7:0] SWITCH_MAC_ADDRH[7:0]	5 th
06h	HMAC_ADDRH[15:8] SWITCH_MAC_ADDRH[15:8]	6 th

For example, if the desired Ethernet physical address is 12-34-56-78-9A-BC, the HMAC_ADDRL and HMAC_ADDRH registers would be programmed as shown in [Figure 11-9](#). The values required to automatically load this configuration from the EEPROM are also shown.

FIGURE 11-9: EXAMPLE EEPROM MAC ADDRESS SETUP



Note: By convention, the right nibble of the left most byte of the Ethernet address (in this example, the 2 of the 12h) is the most significant nibble and is transmitted/received first.

For more information on the EEPROM and EEPROM Loader, refer to [Section 14.0, "I2C Master EEPROM Controller," on page 457](#).

11.10 FIFOs

The device contains four host-accessible FIFOs (TX Status, RX Status, TX Data, and RX Data) and two internal inaccessible Host MAC TX/RX MIL FIFO's (TX MIL FIFO, RX MIL FIFO).

11.10.1 TX/RX FIFOs

The TX/RX Data and Status FIFOs store the incoming and outgoing address and data information, acting as a conduit between the host bus interface (HBI) and the Host MAC. The sizes of these FIFOs are configurable via the [Hardware Configuration Register \(HW_CFG\)](#) register to the ranges described in [Table 11-10](#). Refer to [Section 11.10.3, "FIFO Memory Allocation Configuration"](#) for additional information. The the RX and TX FIFOs related register definitions can be found in section [Section 11.13, "Host MAC & FIFO Interface Registers"](#).

The TX and RX Data FIFOs have the base address of 20h and 00h respectively. However, each FIFO is also accessible at seven additional contiguous memory locations, as can be seen in [FIGURE 5-1: Register Address Map on page 34](#). The Host may access the TX or RX Data FIFOs at any of these alias port locations, as they all function identically and contain the same data. This alias port addressing is implemented to allow hosts to burst through sequential addresses.

For HBI access, the TX and RX Data FIFOs may also be accessed using FIFO Direct Selection mode. In this mode, the address input is ignored and all read access are directed to the RX Data FIFO while all write accesses are directed to the TX Data FIFO. See [Section 9.4.3.2, "FIFO Direct Select Access," on page 92](#) and [Section 9.5.5.3, "FIFO Direct Select Access," on page 114](#).

The TX and RX Status FIFOs can each be read from two register locations; the Status FIFO Port, and the Status FIFO PEEK. The TX and RX Status FIFO Ports (48h and 40h respectively) will perform a destructive read, popping the data from the TX or RX Status FIFO. The TX and RX Status FIFO PEEK register locations (4Ch and 44h respectively) allow a non-destructive read of the top (oldest) location of the FIFOs.

Proper use of the The TX/RX Data and Status FIFOs, including the correct data formatting is described in detail in [Section 11.11, "TX Data Path Operation," on page 172](#) and [Section 11.12, "RX Data Path Operation," on page 183](#).

11.10.2 MIL FIFOs

The MAC Interface Layer (MIL), within the Host MAC, contains a 2KB transmit and a 128 Byte receive FIFO which are separate from the TX and RX FIFOs. These MIL FIFOs are not directly accessible from the HBI. The differentiation between the TX/RX FIFOs and the TX/RX MIL FIFOs is that once the transmit or receive packets are in the MIL FIFOs, the host no longer can control or access the TX or RX data. The MIL FIFOs are essentially the working buffers of the Host MAC logic. In the case of reception, the data must be moved into the RX FIFOs before the host can access the data. For TX operations, the MIL operates in store-and-forward mode and will queue an entire frame before beginning transmission.

As space in the TX MIL FIFO frees, data is moved into it from the TX Data FIFO. Depending on the size of the frames to be transmitted, the Host MAC can hold up to two Ethernet frames. This is in addition to any TX data that may be queued in the TX Data FIFO.

Conversely, as data is received, it is moved from the Host MAC to the RX MIL FIFO, and then into the RX Data FIFO. When the RX Data FIFO fills up, the current or subsequent RX frames will be lost until room is made in the RX Data FIFO. For each frame of data that is lost, the [Host MAC RX Dropped Frames Counter Register \(RX_DROP\)](#) is incremented.

RX and TX MIL FIFO levels are not visible to the host processor and operate independent of the TX/RX FIFOs. FIFO levels set for the TX/RX Data and Status FIFOs do not take into consideration the MIL FIFOs.

11.10.3 FIFO MEMORY ALLOCATION CONFIGURATION

TX and RX FIFO space is configurable through the [Hardware Configuration Register \(HW_CFG\)](#). The user must select the FIFO allocation by setting the [TX FIFO Size \(TX_FIF_SZ\)](#) field in the [Hardware Configuration Register \(HW_CFG\)](#). The TX_FIF_SZ field selects the total allocation for the TX data path, including the TX Status FIFO size. The TX Status FIFO size is fixed at 512 Bytes (128 TX Status DWORDs). The TX Status FIFO length is subtracted from the total TX FIFO size with the remainder being the TX Data FIFO Size. The minimum size of the TX FIFOs is 2KB (TX Data and TX Status FIFOs combined). Note that TX Data FIFO space includes both commands and payload data.

RX FIFO Size is the remainder of the unallocated FIFO space (16384 bytes – TX FIFO Size). The RX Status FIFO size is always equal to 1/16 of the RX FIFO size. The RX Status FIFO length is subtracted from the total RX FIFO size with the remainder being the RX Data FIFO Size.

For example, if TX_FIF_SZ = 6 then:

Total TX FIFO Size = 6144 Bytes (6KB)

TX Status FIFO Size = 512 Bytes (Fixed)

TX Data FIFO Size = 6144 – 512 = 5632 Bytes

RX FIFO Size = 16384 – 6144 = 10240 Bytes (10KB)

RX Status FIFO Size = 10240 / 16 = 640 Bytes (160 RX Status DWORDs)

RX Data FIFO Size = 10240 – 640 = 9600 Bytes

[Table 11-10](#) contains an overview of the configurable TX/RX FIFO sizes and defaults. [TABLE 11-11](#) shows every valid setting for the TX_FIF_SZ field and the resulting FIFO sizes. Note that settings not shown in this table are reserved and should not be used.

Note: The RX Data FIFO is considered full 4 DWORDs before the length that is specified in the HW_CFG register.

TABLE 11-10: TX/RX FIFO CONFIGURABLE SIZES

FIFO	Size Range	Default
TX Status	512	512
RX Status	128-892	704
TX Data	1536-13824	4608
RX Data	1920-13440	10560

TABLE 11-11: VALID TX/RX FIFO ALLOCATIONS

TX_FIF_SZ	TX DATA FIFO SIZE (bytes)	TX STATUS FIFO SIZE (bytes)	RX DATA FIFO SIZE (bytes)	RX STATUS FIFO SIZE (bytes)
2	1536	512	13440	896
3	2560	512	12480	832
4	3584	512	11520	768
5	4608	512	10560	704

TABLE 11-11: VALID TX/RX FIFO ALLOCATIONS (CONTINUED)

TX_FIF_SZ	TX DATA FIFO SIZE (bytes)	TX STATUS FIFO SIZE (bytes)	RX DATA FIFO SIZE (bytes)	RX STATUS FIFO SIZE (bytes)
6	5632	512	9600	640
7	6656	512	8640	576
8	7680	512	7680	512
9	8704	512	6720	448
10	9728	512	5760	384
11	10752	512	4800	320
12	11776	512	3840	256
13	12800	512	2880	192
14	13824	512	1920	128

11.11 TX Data Path Operation

Data is queued for transmission by writing it into the TX Data FIFO. Each packet to be transmitted may be divided among multiple buffers. Each buffer starts with a two DWORD TX command (TX command 'A' and TX command 'B'). The TX command instructs the device on the handling of the associated buffer. Packet boundaries are delineated using control bits within the TX command.

The host provides a 16-bit Packet Tag field in the TX command. The Packet Tag value is appended to the corresponding TX status DWORD. All Packet Tag fields must have the same value for all buffers in a given packet. If tags differ between buffers in the same packet the TXE error will be asserted. Any value may be chosen for a Packet Tag as long as all tags in the same Packet are identical. Packet Tags also provide a method of synchronization between transmitted packets and their associated status. Software can use unique Packet Tags to assist with validating matching status completions.

Note: The use of Packet Tags is not required by the hardware. This field can be used by the LAN software driver for any application. Packet Tags is only one application example.

The Packet Length field in the TX command specifies the number of bytes in the associated packet. All Packet Length fields must have the same value for all buffers in a given packet. Hardware compares the Packet Length field and the actual amount of data received by the Ethernet controller. If the actual packet length count does not match the Packet Length field as defined in the TX command, the Transmitter Error (TXE) flag is asserted.

The device can be programmed to start payload transmission of a buffer on a byte boundary by setting the "Data Start Offset" field in the TX command. The "Data Start Offset" field points to the actual start of the payload data within the first 8 DWORDs of the buffer. Data before the "Data Start Offset" pointer will be ignored. When a packet is split into multiple buffers, each successive buffer may begin on any arbitrary byte.

The device can be programmed to strip padding from the end of a transmit packet in the event that the end of the packet does not align with the host burst boundary. This feature is necessary when the device is operating in a system that always performs multi-word bursts. In such cases the device must guarantee that it can accept data in multiples of the Burst length regardless of the actual packet length. When configured to do so, the device will accept extra data at the end of the packet and will remove the extra padding before transmitting the packet. The device automatically removes data up to the boundary specified in the Buffer End Alignment field specified in each TX command.


The host can instruct the device to issue an interrupt when the buffer has been fully loaded into the TX FIFO contained in the device and transmitted. This feature is enabled through the TX command 'Interrupt on Completion' field.

Upon completion of transmission, irrespective of success or failure, the status of the transmission is written to the TX Status FIFO. TX status is available to the host and may be read using PIO operations. An interrupt can be optionally enabled by the host to indicate the availability of a programmable number TX status DWORDS.

Before writing the TX command and payload data to the TX FIFO, the host must check the available TX FIFO space by performing a PIO read of the [TX FIFO Information Register \(TX_FIFO_INF\)](#). The host must ensure that it does not over-fill the TX FIFO or the TX Error (TXE) flag will be asserted.

The host proceeds to write the TX command by first writing TX command 'A', then TX command 'B'. After writing the command, the host can then move the payload data into the TX FIFO. TX status DWORDs are stored in the TX Status FIFO to be read by the host at a later time upon completion of the data transmission onto the wire.

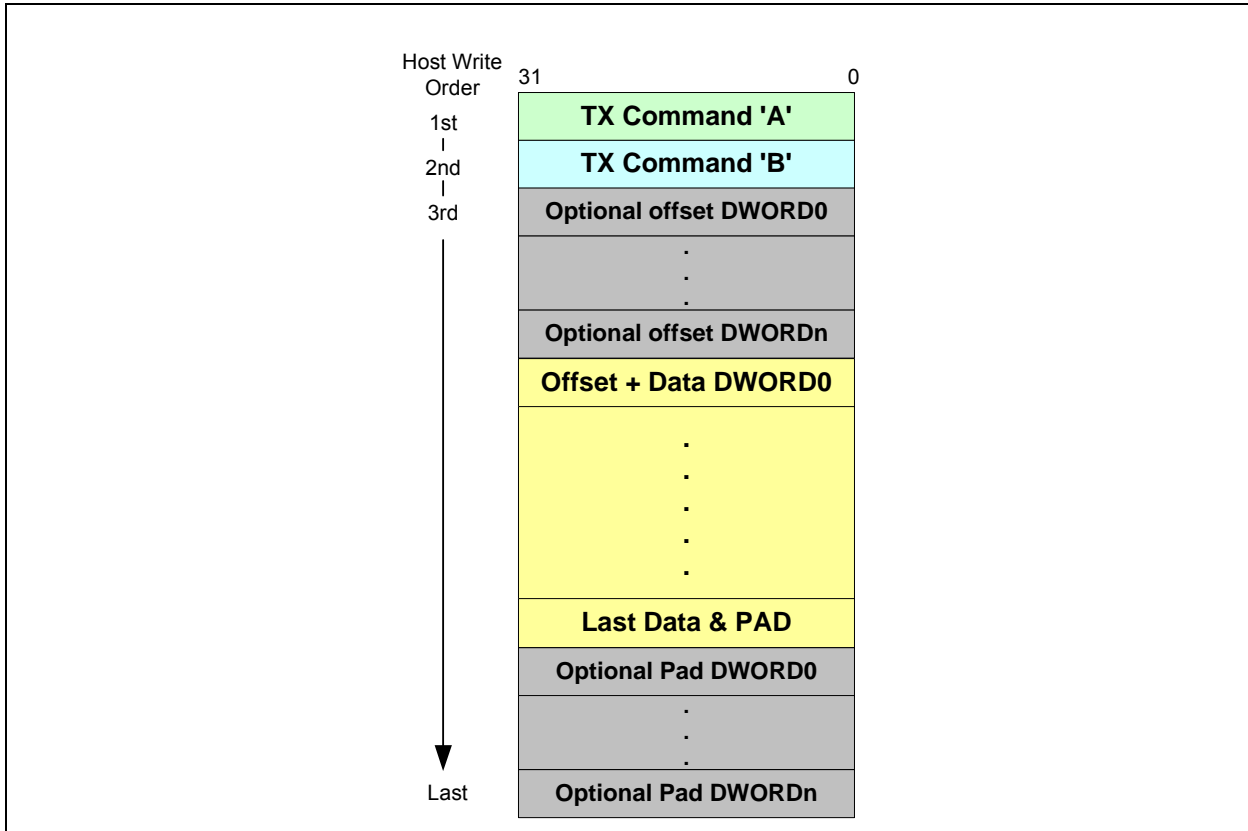
FIGURE 11-10: SIMPLIFIED HOST TX FLOW DIAGRAM



11.11.1 TX BUFFER FORMAT

TX buffers exist in the host's memory in a given format. The host writes a TX command word into the TX data buffer before moving the Ethernet packet data. The TX command A and command B are 32-bit values that are used by the device in the handling and processing of the associated Ethernet packet data buffer. Buffer alignment, segmentation and other packet processing parameters are included in the command structure. The buffer format is illustrated in [Figure 11-11](#).

FIGURE 11-11: TX BUFFER FORMAT



[Figure 11-11](#) shows the TX Buffer as it is written into the device. It should be noted that not all of the data shown in this diagram is actually stored in the TX Data FIFO. This must be taken into account when calculating the actual TX Data FIFO usage. Please refer to [Section 11.11.5, "Calculating Actual TX Data FIFO Usage"](#) for a detailed explanation on calculating the actual TX Data FIFO usage.

11.11.2 TX COMMAND FORMAT

The TX command instructs the TX FIFO controller on handling the subsequent buffer. The command precedes the data to be transmitted. The TX command is divided into two, 32-bit words; TX command 'A' and TX command 'B'.

There is a 16-bit Packet Tag in the TX command 'B' command word. Packet Tags may, if host software desires, be unique for each packet (i.e., an incrementing count). The value of the tag will be returned in the TX status word for the associated packet. The Packet tag can be used by host software to uniquely identify each status word as it is returned to the host.

Both TX command 'A' and TX command 'B' are required for each buffer in a given packet. TX command 'B' must be identical for every buffer in a given packet. If the TX command 'B' words do not match, the Ethernet controller will assert the Transmitter Error (TXE) flag.

11.11.2.1 TX Command 'A'

TABLE 11-12: TX COMMAND 'A' FORMAT

Bits	Description															
31	Interrupt on Completion (IOC). When set, the TX_IOC bit will be asserted in the Interrupt Status Register (INT_STS) when the current buffer has been fully loaded into the TX FIFO.															
30:26	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility.															
25:24	Buffer End Alignment. This field specifies the alignment that must be maintained on the last data transfer of a buffer. The host will add extra DWORDs of data up to the alignment specified in the table below. The device will remove the extra DWORDs. This mechanism can be used to maintain cache line alignment on host processors. <div><table><tr><th>[25]</th><th>[24]</th><th>End Alignment</th></tr><tr><td>0</td><td>0</td><td>4-byte alignment</td></tr><tr><td>0</td><td>1</td><td>16-byte alignment</td></tr><tr><td>1</td><td>0</td><td>32-byte alignment</td></tr><tr><td>1</td><td>1</td><td>Reserved</td></tr></table></div>	[25]	[24]	End Alignment	0	0	4-byte alignment	0	1	16-byte alignment	1	0	32-byte alignment	1	1	Reserved
[25]	[24]	End Alignment														
0	0	4-byte alignment														
0	1	16-byte alignment														
1	0	32-byte alignment														
1	1	Reserved														
23:21	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility															
20:16	Data Start Offset (bytes). This field specifies the offset of the first byte of TX data. The offset value can be anywhere from 0 bytes to a 31 byte offset.															
15:14	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility															
13	First Segment (FS). When set, this bit indicates that the associated buffer is the first segment of the packet.															
12	Last Segment. When set, this bit indicates that the associated buffer is the last segment of the packet															
11	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility.															
10:0	Buffer Size (bytes). This field indicates the number of bytes contained in the buffer following this command. This value, along with the Buffer End Alignment field, is read and checked by the device and used to determine how many extra DWORDs were added to the end of the Buffer. A running count is also maintained in the device of the cumulative buffer sizes for a given packet. This cumulative value is compared against the Packet Length field in the TX command 'B' word and if they do not correlate, the TXE flag is set. The buffer size specified does not include the buffer end alignment padding or data start offset added to a buffer.															

11.11.2.2 TX Command 'B'

TABLE 11-13: TX COMMAND 'B' FORMAT

Bits	Description
31:16	Packet Tag. The host should write a unique packet identifier to this field. This identifier is added to the corresponding TX status word and can be used by the host to correlate TX status words with their corresponding packets. The use of packet tags is not required by the hardware. This field can be used by the LAN software driver for any application. Packet Tags is one application example.
15	Reserved. This bit is reserved. Always write zero to this bit to guarantee future compatibility.
14	TX Checksum Enable (CK). When this bit is set in conjunction with the first segment (FS) bit in TX Command 'A' and the TX Checksum Offload Engine Enable (TX_COE_EN) bit in the Host MAC Checksum Offload Engine Control Register (HMAC_COE_CR) , the TX checksum offload engine (TXCOE) will calculate a L3 checksum for the associated frame.
13	Add CRC Disable. When set, the automatic addition of the CRC is disabled.
12	Disable Ethernet Frame Padding. When set, this bit prevents the automatic addition of padding to an Ethernet frame of less than 64 bytes. The CRC field is also added despite the state of the Add CRC Disable field.
11	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility.
10:0	Packet Length (bytes). This field indicates the total number of bytes in the current packet. This length does not include the offset or padding. If the Packet Length field does not match the actual number of bytes in the packet the Transmitter Error (TXE) flag will be set.

11.11.3 TX DATA FORMAT

The TX data section begins at the third DWORD in the TX buffer (after TX command 'A' and TX command 'B'). The location of the first byte of valid buffer data to be transmitted is specified in the "Data Start Offset" field of the TX command 'A' word. [Table 11-14, "TX DATA Start Offset"](#), shows the correlation between the setting of the LSBs in the "Data Start Offset" field and the byte location of the first valid data byte. Additionally, transmit buffer data can be offset by up to 7 additional DWORDS as indicated by the upper three MSBs (5:2) in the "Data Start Offset" field.

TABLE 11-14: TX DATA START OFFSET

Data Start Offset [1:0]:	11	10	01	00
First TX Data Byte:	D[31:24]	D[23:16]	D[15:8]	D[7:0]

TX data is contiguous until the end of the buffer. The buffer may end on a byte boundary. Unused bytes at the end of the packet will not be sent to the Host MAC Interface Layer for transmission.

The Buffer End Alignment field in TX command 'A' specifies the alignment that must be maintained for the associated buffer. End alignment may be specified as 4-, 16-, or 32-byte. The host processor is responsible for adding the additional data to the end of the buffer. The hardware will automatically remove this extra data.

11.11.3.1 TX Buffer Fragmentation Rules

Transmit buffers must adhere to the following rules:

- Each buffer can start and end on any arbitrary byte alignment
- The first buffer of any transmit packet can be any length
- Middle buffers (i.e., those with First Segment = Last Segment = 0) must be greater than, or equal to 4 bytes in length
- The final buffer of any transmit packet can be any length

The MIL operates in store-and-forward mode and has specific rules with respect to fragmented packets. The total space consumed in the TX MIL FIFO must be limited to no more than 2KB - 3 DWORDs (2,036 bytes total). Any transmit packet that is so highly fragmented that it takes more space than this must be un-fragmented (by copying to a driver-supplied buffer) before the transmit packet can be sent to the device.

One approach to determine whether a packet is too fragmented is to calculate the actual amount of space that it will consume, and check it against 2,036 bytes. Another approach is to check the number of buffers against a worst-case limit of 86 (see explanation below).

For the best case alignment scenario (full DWORDs at the start and the end of each buffer), the absolute largest frame size is 2040 bytes (plus the automatically added FCS if enable).

11.11.3.2 Calculating Worst-Case TX MIL FIFO Usage

The actual space consumed by a buffer in the TX MIL FIFO consists only of any partial DWORD offsets in the first/last DWORD of the buffer, plus all of the whole DWORDs in between. Any whole DWORD offsets and/or alignments are stripped off before the buffer is loaded into the TX Data FIFO, and TX command words are stripped off before the buffer is written to the TX MIL FIFO, so none of those DWORDs count as space consumed. The worst-case overhead for a TX buffer is 6 bytes, which assumes that it started on the high byte of a DWORD and ended on the low byte of a DWORD. A TX packet consisting of 86 such fragments would have an overhead of 516 bytes (6 * 86) which, when added to a 1514-byte max-size transmit packet (1516 bytes, rounded up to the next whole DWORD), would give a total space consumption of 2,032 bytes, leaving 4 bytes to spare; this is the basis for the “86 fragment” rule mentioned above. For more information on the MIL FIFO's refer to [Section 11.10.2, "MIL FIFOs," on page 170](#).

11.11.4 TX STATUS FORMAT

TX status is passed to the host CPU through a separate FIFO mechanism. A status word is returned for each packet transmitted. Data transmission is suspended if the TX Status FIFO becomes full. Data transmission will resume when the host reads the TX status and there is room in the FIFO for more “TX Status” data.

The host can optionally choose to not read the TX status. The TX status can be ignored by setting the “TX Status Discard Allow Overrun Enable” (TXSAO) bit in the [Transmit Configuration Register \(TX_CFG\)](#). Setting this bit high allows the transmitter to continue operation with a full TX Status FIFO. In this mode the status information is still available in the TX Status FIFO, and TX status interrupts still function. In the case of a full FIFO, the TXSUSED counter will stay at its maximum value and no further TX status will be written to the TX Status FIFO, preventing an overrun, until the host frees space by reading TX status. In this mode the host is responsible for re-synchronizing TX status in the case of an overrun.

Note: Though the Host MAC is communicating locally with the switch fabric MAC, the events described in the TX Status word may still occur.

Bits	Description
31:16	Packet TAG. Unique identifier written by the host into the Packet Tag field of the TX command 'B' word. This field can be used by the host to correlate TX status words with the associated TX packets.
15	Error Status (ES). When set, this bit indicates that the Ethernet controller has reported an error. This bit is the logical OR of bits 11, 10, 9, 8, 2, 1 in this status word.
14:12	Reserved. These bits are reserved. Always write zeros to this field to guarantee future compatibility.
11	Loss of Carrier. When set, this bit indicates the loss of carrier during transmission.
10	No Carrier. When set, this bit indicates that the carrier signal from the transceiver was not present during transmission.
9	Late Collision. When set, indicates that the packet transmission was aborted after the collision window of 64 bytes.
8	Excessive Collisions. When set, this bit indicates that the transmission was aborted after 16 collisions while attempting to transmit the current packet.
7	Reserved. This bit is reserved. Always write zeros to this field to guarantee future compatibility.

Bits	Description
6:3	Collision Count. This counter indicates the number of collisions that occurred before the packet was transmitted. It is not valid when excessive collisions (bit 8) is also set.
2	Excessive Deferral. If the deferred bit is set in the control register, the setting of the excessive deferral bit indicates that the transmission has ended because of a deferral of over 24288 bit times during transmission.
1	Reserved. This bit is reserved.
0	Deferred. When set, this bit indicates that the current packet transmission was deferred.

11.11.5 CALCULATING ACTUAL TX DATA FIFO USAGE

The following rules are used to calculate the actual TX Data FIFO space consumed by a TX Packet:

- TX command 'A' is stored in the TX Data FIFO for every TX buffer
- TX command 'B' is written into the TX Data FIFO when the First Segment (FS) bit is set in TX command 'A'
- Any DWORD-long data added as part of the “Data Start Offset” is removed from each buffer before the data is written to the TX Data FIFO. Any data that is less than 1 DWORD is passed to the TX Data FIFO.
- Payload from each buffer within a Packet is written into the TX Data FIFO.
- Any DWORD-long data added as part of the End Padding is removed from each buffer before the data is written to the TX Data FIFO. Any end padding that is less than 1 DWORD is passed to the TX Data FIFO

11.11.6 TRANSMIT EXAMPLES

11.11.6.1 TX Example 1

In this example a single, 111-Byte Ethernet packet will be transmitted. This packet is divided into three buffers. The three buffers are as follows:

Buffer 0:

- 7-Byte “Data Start Offset”
- 79-Bytes of payload data
- 16-Byte “Buffer End Alignment”

Buffer 1:


- 0-Byte “Data Start Offset”
- 15-Bytes of payload data
- 16-Byte “Buffer End Alignment”

Buffer 2:

- 10-Byte “Data Start Offset”
- 17-Bytes of payload data
- 16-Byte “Buffer End Alignment”

Figure 11-12 illustrates the TX command structure for this example, and also shows how data is passed to the TX Data FIFO.

FIGURE 11-12: TX EXAMPLE 1




11.11.6.2 TX Example 2

In this example, a single 183-Byte Ethernet packet will be transmitted. This packet is in a single buffer as follows:

- 2-Byte “Data Start Offset”
- 183-Bytes of payload data
- 4-Byte “Buffer End Alignment”

Figure 11-13 illustrates the TX command structure for this example, and also shows how data is passed to the TX Data FIFO. Note that the packet resides in a single TX Buffer, therefore both the FS and LS bits are set in TX command ‘A’.

FIGURE 11-13: TX EXAMPLE 2



11.11.6.3 TX Example 3

In this example a single, 111-Byte Ethernet packet will be transmitted with a TX checksum. This packet is divided into four buffers. The four buffers are as follows:

Buffer 0:

- 4-Byte "Data Start Offset"
- 4-Byte Checksum Preamble
- 16-Byte "Buffer End Alignment"

Buffer 1:

- 7-Byte "Data Start Offset"
- 79-Bytes of payload data
- 16-Byte "Buffer End Alignment"

Buffer 2:

- 0-Byte "Data Start Offset"
- 15-Bytes of payload data
- 16-Byte "Buffer End Alignment"

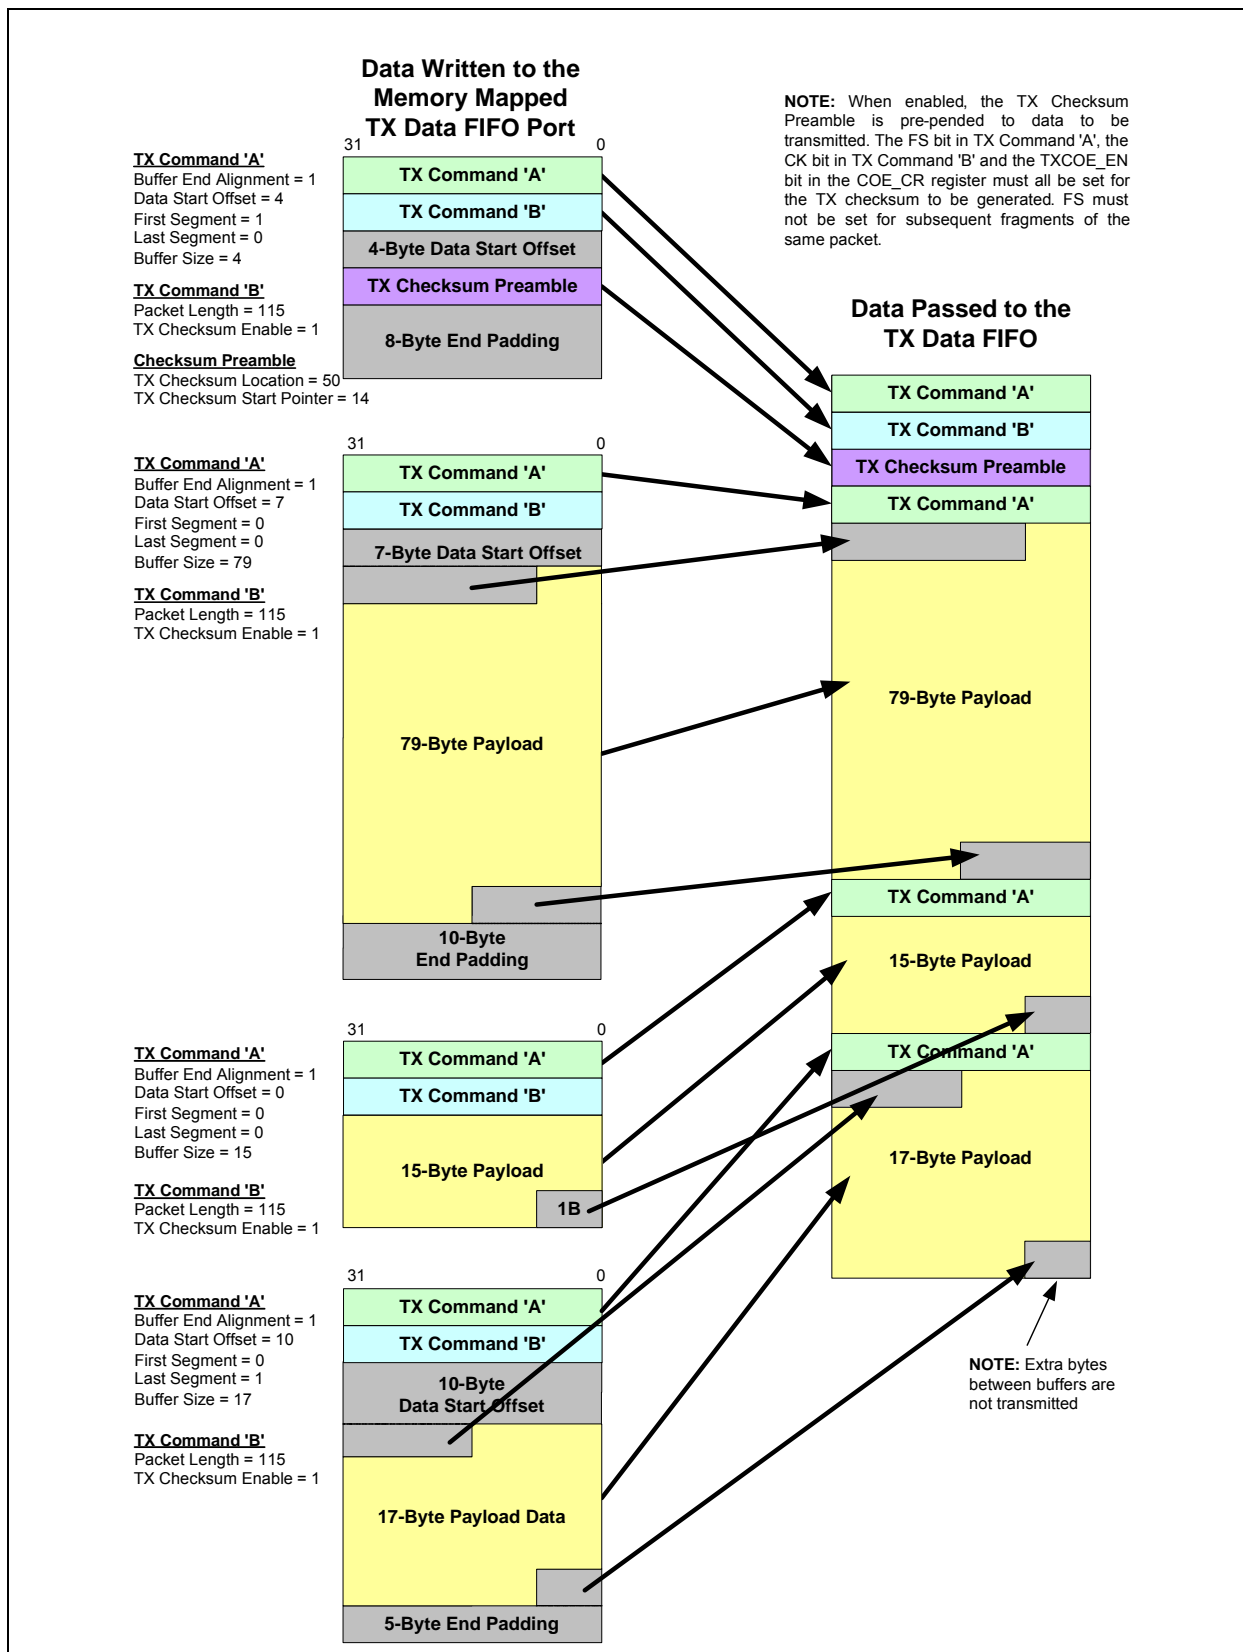
Buffer 3:

- 10-Byte "Data Start Offset"
- 17-Bytes of payload data
- 16-Byte "Buffer End Alignment"

Figure 11-12, "TX Example 1" illustrates the TX command structure for this example, and also shows how data is passed to the TX data FIFO.

Note: In order to perform a TX checksum calculation on the associated packet, bit 14 (CK) of the TX Command 'B' must be set in conjunction with bit 13 (FS) of TX Command 'A' and the [TX Checksum Offload Engine Enable \(TX_COE_EN\)](#) bit of the [Host MAC Checksum Offload Engine Control Register \(HMAC_COE_CR\)](#). For more information, refer to [Section 11.8, "Transmit Checksum Offload Engine \(TXCOE\)"](#).

FIGURE 11-14: TX EXAMPLE 3



11.11.7 TRANSMITTER ERRORS

If the Transmitter Error (TXE) flag is asserted for any reason, the transmitter will continue operation. TX Error (TXE) will be asserted under the following conditions:

- If the actual packet length count does not match the Packet Length field as defined in the TX command.
- Both TX command 'A' and TX command 'B' are required for each buffer in a given packet. TX command 'B' must be identical for every buffer in a given packet. If the TX command 'B' words do not match, the Ethernet controller will assert the Transmitter Error (TXE) flag.
- Host overrun of the TX Data FIFO.

11.11.8 STOPPING AND STARTING THE TRANSMITTER

To halt the transmitter, the host must set the STOP_TX bit in the [Transmit Configuration Register \(TX_CFG\)](#). The transmitter will finish sending the current frame (if there is a frame transmission in progress). When the transmitter has received the TX status for this frame, it will clear the STOP_TX and TX_ON bits, and will pulse the TXSTOP_INT in the [Interrupt Status Register \(INT_STS\)](#).

Once stopped, the host can optionally clear the TX Status and TX Data FIFOs. The host must re-enable the transmitter by setting the TX_ON bit. If there are frames pending in the TX Data FIFO (i.e., TX Data FIFO was not purged), the transmission will resume with this data.

11.12 RX Data Path Operation

When an Ethernet Packet is received, the Host MAC Interface Layer (MIL) first begins to transfer the RX data. This data is loaded into the RX Data FIFO. The RX Data FIFO pointers are updated as data is written into the FIFO.

The last transfer from the MIL is the RX status word. The device implements a separate FIFO for the RX status words. The total available RX data and status queued in the RX FIFO can be read from the [RX FIFO Information Register \(RX_FIFO_INF\)](#). The host may read any number of available RX status words before reading the RX Data FIFO.

The host must use caution when reading the RX data and status. The host must never read more data than what is available in the FIFO's. If this is attempted an underrun condition will occur. If this error occurs, the Ethernet controller will assert the Receiver Error (RXE) interrupt. If an underrun condition occurs, a soft reset is required to regain host synchronization.

A configurable beginning offset is supported in the device. The RX data Offset field in the [Receive Configuration Register \(RX_CFG\)](#) controls the number of bytes that the beginning of the RX data buffer is shifted. The host can set an offset from 0-31 bytes. The offset may be changed in between RX packets, but it must not be changed during an RX packet read.

The device can be programmed to add padding at the end of a receive packet in the event that the end of the packet does not align with the host burst boundary. This feature is necessary when the device is operating in a system that always performs multi-DWORD bursts. In such cases the device must guarantee that it can transfer data in multiples of the Burst length regardless of the actual packet length. When configured to do so, the device will add extra data at the end of the packet to allow the host to perform the necessary number of reads so that the Burst length is not cut short. Once a packet has been padded by the H/W, it is the responsibility of the host to interrogate the packet length field in the RX status and determine how much padding to discard at the end of the packet.

It is possible to read multiple packets out of the RX Data FIFO in one continuous stream. It should be noted that the programmed Offset and Padding will be added to each individual packet in the stream, since packet boundaries are maintained.

11.12.1 RX SLAVE PIO OPERATION

Using PIO mode, the host can either implement a polling or interrupt scheme to empty the received packet out of the RX Data FIFO. The host will remain in the idle state until it receives an indication (interrupt or polling) that data is available in the RX Data FIFO. The host will then read the RX Status FIFO to get the packet status, which will contain the packet length and any other status information. The host should perform the proper number of reads, as indicated by the packet length *plus* the start offset *and* the amount of optional padding added to the end of the frame, from the RX Data FIFO. A typical host receive routine using interrupts can be seen in [Figure 11-15](#), while a typical host receive routine using polling can be seen in [Figure 11-16](#).

FIGURE 11-15: HOST RECEIVE ROUTINE USING INTERRUPTS





FIGURE 11-16: HOST RECEIVE ROUTINE USING POLLING



11.12.1.1 Receive Data FIFO Fast Forward

The RX data path implements an automatic data discard function. Using the RX Data FIFO Fast Forward bit (RX_FFWD) in the [Receive Datapath Control Register \(RX_DP_CTRL\)](#), the host can instruct the device to skip the packet at the head of the RX Data FIFO. The RX Data FIFO pointers are automatically incremented to the beginning of the next RX packet.

When performing a fast-forward, there must be at least 4 DWORDs of data in the RX Data FIFO for the packet being discarded. For cases with less than 4 DWORDs, do not use RX_FFWD. In this case data must be read from the RX Data FIFO and discarded using standard PIO read operations.

After initiating a fast-forward operation, do not perform any reads of the RX Data FIFO until the RX_FFWD bit is cleared. Other resources can be accessed during this time (i.e., any registers and/or the other three FIFO's). Also note that the RX_FFWD will only fast-forward the RX Data FIFO, not the RX Status FIFO. After an RX fast-forward operation the RX status must still be read from the RX Status FIFO.

The receiver does not have to be stopped to perform a fast-forward operation.


11.12.1.2 Force Receiver Discard (Receiver Dump)

In addition to the Receive data Fast Forward feature, device also implements a receiver “dump” feature. This feature allows the host processor to flush the entire contents of the RX Data and RX Status FIFOs. When activated, the read and write pointers for the RX Data and Status FIFO's will be returned to their reset state. To perform a receiver dump, the device receiver must be halted. Once the receiver stop completion is confirmed, the RX_DUMP bit can be set in the [Receive Configuration Register \(RX_CFG\)](#). The RX_DUMP bit is cleared when the dump is complete. For more information on stopping the receiver, please refer to [Section 11.12.4, "Stopping and Starting the Receiver"](#). For more information on the RX_DUMP bit, please refer to [Section 11.13.2, "Receive Configuration Register \(RX_CFG\)," on page 190](#).

11.12.2 RX PACKET FORMAT

The RX status words can be read from the RX Status FIFO port, while the RX data packets can be read from the RX Data FIFO. RX data packets are formatted in a specific manner before the host can read them as shown in [Figure 11-17](#). It is assumed that the host has previously read the associated status word from the RX Status FIFO, to ascertain the data size and any error conditions.

FIGURE 11-17: RX PACKET FORMAT



11.12.3 RX STATUS FORMAT

Note: Though the Host MAC is communicating locally with the switch fabric MAC, the events described in the RX Status word may still occur.

BITS	DESCRIPTION
31	Packet Filter. When set, this bit indicates that the associated frame passed the frame filtering described in Section 11.5 .
30	Filtering Fail. When set, this bit indicates that the associated frame failed the address recognizing filtering described in Section 11.4 .
29:16	Packet Length. The size, in bytes, of the corresponding received frame.
15	Error Status (ES). When set this bit indicates that the Host MAC Interface Layer (MIL) has reported an error. This bit is the Internal logical “or” of bits 11,7,6 and 1.
14	Reserved. These bits are reserved. Reads 0.
13	Broadcast Frame. When set, this bit indicates that the received frame has a Broadcast address.
12	Length Error (LE). When set, this bit indicates that the actual length does not match with the length/type field of the received frame.
11	Runt Frame. When set, this bit indicates that frame was prematurely terminated before the collision window (64 bytes). Runt frames are passed on to the host only if the Pass Bad Frames bit (PASSBAD) of the Host MAC Control Register (HMAC_CR) is set.
10	Multicast Frame. When set, this bit indicates that the received frame has a Multicast address.
9:8	Reserved. These bits are reserved. Reads 0.
7	Frame Too Long. When set, this bit indicates that the frame length exceeds the maximum Ethernet specification of 1518 bytes. This is only a frame too long indication and will not cause the frame reception to be truncated.
6	Collision Seen. When set, this bit indicates that the frame has seen a collision after the collision window. This indicates that a late collision has occurred.
5	Frame Type. When set, this bit indicates that the frame is an Ethernet-type frame (Length/Type field in the frame is greater than 1500). When reset, it indicates the incoming frame was an 802.3 type frame. This bit is not set for Runt frames less than 14 bytes.
4	Receive Watchdog time-out. When set, this bit indicates that the incoming frame was greater than or equal to 2048 bytes, therefore expiring the Receive Watchdog Timer. Frames greater than or equal to 2049 bytes are truncated to 2048 bytes and would most likely have a CRC error as a result.
3	MII Error. When set, this bit indicates that a receive error was detected during frame reception.
2	Dribbling Bit. When set, this bit indicates that the frame contained a non-integer multiple of 8 bits. This error is reported only if the number of dribbling bits in the last byte is at least 3 in the 10 Mbps operating mode. This bit will not be set when the collision seen bit[6] is set. If set and the CRC error bit is [1] reset, then the packet is considered to be valid.
1	CRC Error. When set, this bit indicates that a CRC error was detected. This bit is also set when the RX_ER pin is asserted during the reception of a frame even though the CRC may be correct. This bit is not valid if the received frame is a Runt frame, or a late collision was detected.
0	Reserved. These bits are reserved. Reads 0

11.12.4 STOPPING AND STARTING THE RECEIVER

To stop the receiver, the host must clear the RXEN bit in the [Host MAC Control Register \(HMAC_CR\)](#). When the receiver is halted, the RXSTOP_INT will be pulsed and reflected in the [Interrupt Status Register \(INT_STS\)](#). Once stopped, the host can optionally clear the RX Status and RX Data FIFOs. The host must re-enable the receiver by setting the RXEN bit.

11.12.5 RECEIVER ERRORS

If the Receiver Error (RXE) flag is asserted in the [Interrupt Status Register \(INT_STS\)](#) for any reason, the receiver will continue operation. RX Error (RXE) will be asserted under the following conditions:

- A host underrun of RX Data FIFO
- A host underrun of the RX Status FIFO
- An overrun of the RX Status FIFO ([RX Status FIFO Full Interrupt \(RSFF\)](#))

It is the duty of the host to identify and resolve any error conditions.

11.13 Host MAC & FIFO Interface Registers

This section details the directly addressable Host MAC and TX/RX FIFO related System CSRs. These registers allow for the configuration of the TX/RX FIFO's, Host MAC and indirect access to the complete set of Host MAC CSRs. The Host MAC CSRs are accessible through via the [Host MAC CSR Interface Command Register \(MAC_CSR_CMD\)](#) and [Host MAC CSR Interface Data Register \(MAC_CSR_DATA\)](#).

Note: For more information on the TX/RX FIFO's, refer to [Section 11.10, "FIFOs"](#).

Note: The full list of indirectly addressable Host MAC CSRs are described in [Section 11.14, "Host MAC Control and Status Registers," on page 202](#).

TABLE 11-15: HOST MAC & FIFO INTERFACE LOGIC REGISTERS

Address	Register Name (SYMBOL)
068h	FIFO Level Interrupt Register (FIFO_INT)
06Ch	Receive Configuration Register (RX_CFG)
070h	Transmit Configuration Register (TX_CFG)
078h	Receive Datapath Control Register (RX_DP_CTRL)
07Ch	RX FIFO Information Register (RX_FIFO_INF)
080h	TX FIFO Information Register (TX_FIFO_INF)
0A0h	Host MAC RX Dropped Frames Counter Register (RX_DROP)
0A4h	Host MAC CSR Interface Command Register (MAC_CSR_CMD)
0A8h	Host MAC CSR Interface Data Register (MAC_CSR_DATA)
0ACh	Host MAC Automatic Flow Control Configuration Register (AFC_CFG)

11.13.1 FIFO LEVEL INTERRUPT REGISTER (FIFO_INT)

Offset: 068h Size: 32 bits

This read/write register configures the limits where the RX/TX Data and Status FIFO's will generate system interrupts.

Bits	Description	Type	Default
31:24	TX Data Available Level The value in this field sets the level, in number of 64 Byte blocks, at which the TX Data FIFO Available Interrupt (TDFA) will be generated. When the TX Data FIFO free space is greater than this value, a TX Data FIFO Available Interrupt (TDFA) will be generated in the Interrupt Status Register (INT_STS) .	R/W	48h
23:16	TX Status Level The value in this field sets the level, in number of DWORDs, at which the TX Status FIFO Level Interrupt (TSFL) will be generated. When the TX Status FIFO used space is greater than this value, a TX Status FIFO Level Interrupt (TSFL) will be generated in the Interrupt Status Register (INT_STS) .	R/W	00h
15:8	RESERVED	RO	-
7:0	RX Status Level The value in this field sets the level, in number of DWORDs, at which the RX Status FIFO Level Interrupt (RSFL) will be generated. When the RX Status FIFO used space is greater than this value, a RX Status FIFO Level Interrupt (RSFL) will be generated in the Interrupt Status Register (INT_STS) .	R/W	00h

11.13.2 RECEIVE CONFIGURATION REGISTER (RX_CFG)

Offset: 06Ch Size: 32 bits

This register controls the Host MAC receive engine.

Bits	Description	Type	Default										
31:30	<p>RX End Alignment (RX_EA) This field specifies the alignment that must be maintained on the last data transfer of a buffer. The device will add extra DWORDs of data up to the alignment specified in the table below. The host is responsible for removing these extra DWORDs. This mechanism can be used to maintain cache line alignment on host processors.</p> <table border="1"><thead><tr><th>Bit Values [31:30]</th><th>End Alignment</th></tr></thead><tbody><tr><td>00</td><td>4-Byte Alignment</td></tr><tr><td>01</td><td>16-Byte Alignment</td></tr><tr><td>10</td><td>32-Byte Alignment</td></tr><tr><td>11</td><td>RESERVED</td></tr></tbody></table> <p>Note: The desired RX End Alignment must be set before reading a packet. The RX End Alignment can be changed between reading receive packets, but must not be changed if the packet is partially read.</p>	Bit Values [31:30]	End Alignment	00	4-Byte Alignment	01	16-Byte Alignment	10	32-Byte Alignment	11	RESERVED	R/W	00b
Bit Values [31:30]	End Alignment												
00	4-Byte Alignment												
01	16-Byte Alignment												
10	32-Byte Alignment												
11	RESERVED												
29:28	RESERVED	RO	-										
27:16	<p>RX DMA Count (RX_DMA_CNT) This 12-bit field indicates the amount of data, in DWORDs, to be transferred out of the RX Data FIFO before asserting the RX DMA Interrupt (RXD_INT). After being set, this field is decremented for each DWORD of data that is read from the RX Data FIFO. This field can be overwritten with a new value before it reaches zero.</p>	R/W	000h										
15	<p>Force RX Discard (RX_DUMP) When a 1 is written to this bit, the RX Data and Status FIFO's are cleared of all pending data and the RX data and status pointers are cleared to zero.</p> <p>Note: Please refer to Section 11.12.1.2, "Force Receiver Discard (Receiver Dump)," on page 185 for a detailed description regarding the use of RX_DUMP.</p>	WO SC	0b										
14:13	RESERVED	RO	-										

Bits	Description	Type	Default
12:8	RX Data Offset (RXDOFF) This field controls the offset value, in bytes, that is added to the beginning of an RX data packet. The start of the valid data will be shifted by the number of bytes specified in this field. An offset of 0-31 bytes is a valid number of offset bytes. Note: The two LSBs of this field (D[9:8]) must not be modified while the RX is running. The receiver must be halted, and all data purged before these two bits can be modified. The upper three bits (DWORD offset) may be modified while the receiver is running. Modifications to the upper bits will take affect on the next DWORD read.	R/W	00000b
7:0	RESERVED	RO	-

11.13.3 TRANSMIT CONFIGURATION REGISTER (TX_CFG)

Offset: 070h Size: 32 bits

This register controls the Host MAC transmit functions.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15	Force TX Status Discard (TXS_DUMP) When a 1 is written to this bit, the TX Status FIFO is cleared of all pending status DWORDs and the TX status pointers are cleared to zero.	WO SC	0b
14	Force TX Data Discard (TXD_DUMP) When a 1 is written to this bit, the TX Data FIFO is cleared of all pending data and the TX data pointers are cleared to zero.	WO SC	0b
13:3	RESERVED	RO	-
2	TX Status Allow Overrun (TXSAO) When this bit is cleared, Host MAC data transmission is suspended if the TX Status FIFO becomes full. Setting this bit high allows the transmitter to continue operation with a full TX Status FIFO. Note: This bit does not affect the operation of the TX Status FIFO Full Interrupt (TSFF) .	R/W	0b
1	Transmitter Enable (TX_ON) When this bit is set, the Host MAC transmitter is enabled. Any data in the TX Data FIFO will be sent. This bit is cleared automatically when the STOP_TX bit is set and the transmitter is halted.	R/W	0b
0	Stop Transmitter (STOP_TX) When this bit is set, the Host MAC transmitter will finish the current frame, and will then stop transmitting. When the transmitter has stopped this bit will clear. All writes to this bit are ignored while this bit is high.	R/W SC	0b

11.13.4 RECEIVE DATAPATH CONTROL REGISTER (RX_DP_CTRL)

Offset: 078h Size: 32 bits

This register is used to discard unwanted receive frames.

Bits	Description	Type	Default
31	RX Data FIFO Fast Forward (RX_FFWD) Writing a 1 to this bit causes the RX Data FIFO to fast-forward to the start of the next frame. This bit will remain high until the RX Data FIFO fast-forward operation has completed. No reads should be issued to the RX Data FIFO while this bit is high. Note: Please refer to section Section 11.12.1.1, "Receive Data FIFO Fast Forward," on page 185 for detailed information regarding the use of RX_FFWD.	R/W SC	0h
30:0	RESERVED	RO	-

LAN9352

11.13.5 RX FIFO INFORMATION REGISTER (RX_FIFO_INF)

Offset: 07Ch Size: 32 bits

This register contains the indication of used space in the RX FIFO's.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:16	RX Status FIFO Used Space (RXSUSED) This field indicates the amount of space, in DWORDs, currently used in the RX Status FIFO.	RO	0b
15:0	RX Data FIFO Used Space (RXDUSED) This field indicates the amount of space, in bytes, used in the RX Data FIFO. For each receive frame, the field is incremented by the length of the receive data. In cases where the payload does not end on a DWORD boundary, the total will be rounded up to the nearest DWORD.	RO	0b

11.13.6 TX FIFO INFORMATION REGISTER (TX_FIFO_INF)

Offset: 080h Size: 32 bits

This register contains the indication of free space in the TX Data FIFO and the used space in the TX Status FIFO.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:16	TX Status FIFO Used Space (TXSUSED) This field indicates the amount of space, in DWORDs, currently used in the TX Status FIFO.	RO	0b
15:0	TX Data FIFO Free Space (TXFREE) This field indicates the amount of space, in bytes, available in the TX Data FIFO. The application should never write more than is available, as indicated by this value.	RO	1200h

11.13.7 HOST MAC RX DROPPED FRAMES COUNTER REGISTER (RX_DROP)

Offset: 0A0h Size: 32 bits

This register indicates the number of receive frames that have been dropped by the Host MAC.

Bits	Description	Type	Default
31:0	RX Dropped Frame Counter (RX_DFC) This counter is incremented every time a receive frame is dropped by the Host MAC. RX_DFC is cleared on any read of this register. Note: The interrupt RXDFH_INT (bit 23 of the Interrupt Status Register (INT_STS)) can be issued when this counter passes through its halfway point (7FFFFFFh to 80000000h).	RC	00000000h

11.13.8 HOST MAC CSR INTERFACE COMMAND REGISTER (MAC_CSR_CMD)

Offset: 0A4h Size: 32 bits

This read-write register is used to control the read and write operations to/from the Host MAC. This register is used in conjunction with the [Host MAC CSR Interface Data Register \(MAC_CSR_DATA\)](#) to indirectly access the Host MAC CSRs.

Note: The full list of Host MAC CSRs are described in [Section 11.14, "Host MAC Control and Status Registers," on page 202](#).

Bits	Description	Type	Default
31	Host MAC CSR Busy (HMAC_CSR_BUSY) When a 1 is written into this bit, the read or write operation is performed to the specified Host MAC CSR. This bit will remain set until the operation is complete. In the case of a read, this indicates that the host can read valid data from the Host MAC CSR Interface Data Register (MAC_CSR_DATA) . Note: The MAC_CSR_CMD and MAC_CSR_DATA registers must not be modified until this bit is cleared.	R/W SC	0b
30	R/nW When set, this bit indicates that the host is requesting a read operation. When clear, the host is performing a write. 0: Host MAC CSR Write Operation 1: Host MAC CSR Read Operation	R/W	0b
29:8	RESERVED	RO	-
7:0	CSR Address The 8-bit value in this field selects which Host MAC CSR will be accessed by the read or write operation. The index of each Host MAC CSR is defined in Section 11.14, "Host MAC Control and Status Registers," on page 202 .	R/W	00h

11.13.9 HOST MAC CSR INTERFACE DATA REGISTER (MAC_CSR_DATA)

Offset: 0A8h Size: 32 bits

This read-write register is used in conjunction with the [Host MAC CSR Interface Command Register \(MAC_CSR_CMD\)](#) to indirectly access the Host MAC CSRs.

Note: The full list of Host MAC CSRs are described in [Section 11.14, "Host MAC Control and Status Registers," on page 202](#).

Bits	Description	Type	Default
31:0	<p>Host MAC CSR Data</p> <p>This field contains the value read from or written to the Host MAC CSR as specified in the Host MAC CSR Interface Command Register (MAC_CSR_CMD). Upon a read, the value returned depends on the R/nW bit in the MAC_CSR_CMD register. If R/nW is a 1, the data in this register is from the Host MAC. If R/nW is 0, the data is the value that was last written into this register.</p> <p>Note: The MAC_CSR_CMD and MAC_CSR_DATA registers must not be modified until the CSR Busy bit is cleared in the MAC_CSR_CMD register.</p>	R/W	00000000h

11.13.10 HOST MAC AUTOMATIC FLOW CONTROL CONFIGURATION REGISTER (AFC_CFG)

Offset: 0ACh Size: 32 bits

This read/write register configures the mechanism that controls the automatic and software-initiated transmission of pause frames and back pressure from the Host MAC to the switch fabric. This register is used in conjunction with the [Host MAC Flow Control Register \(HMAC_FLOW\)](#) in the Host MAC CSR space. Pause frames and backpressure are sent to the switch fabric to stop it from sending packets to the Host MAC. Network data into the switch fabric is affected only if the switch fabric buffering fills.

Note: The Host MAC will not transmit pause frames or assert back pressure if the transmitter is disabled. This register controls only the Host MAC flow control and not the Switch Engine MAC's flow control.

Refer to section [Section 11.2, "Flow Control," on page 152](#) for additional information.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:16	<p>Automatic Flow Control High Level (AFC_HI) This field specifies, in multiples of 64 bytes, the level at which flow control will trigger. When this limit is reached, the chip will apply back pressure or will transmit a pause frame as programmed in bits [3:0] of this register.</p> <p>During full-duplex operation only a single pause frame is transmitted when this level is reached. The pause time transmitted in this frame is programmed in the FCPT field of the Host MAC Flow Control Register (HMAC_FLOW) in the Host MAC CSR space.</p> <p>During half-duplex operation each incoming frame that matches the criteria in bits [3:0] of this register will be jammed for the period set in the BACK_DUR field.</p> <p>Note: This level is also used for hard-wired flow control when HW_FC_EN is set in the Port 0 Manual Flow Control Register (MANUAL_FC_0).</p>	R/W	00h
15:8	<p>Automatic Flow Control Low Level (AFC_LO) This field specifies, in multiples of 64 bytes, the level at which a pause frame is transmitted with a pause time setting of zero. When the amount of data in the RX Data FIFO falls below this level the pause frame is transmitted. A pause time value of zero instructs the other transmitting device to immediately resume transmission. The zero time pause frame will only be transmitted if the RX Data FIFO had reached the AFC_HI level and a pause frame was sent. A zero pause time frame is sent whenever automatic flow control is enabled in bits [3:0] of this register.</p> <p>Note: When automatic flow control is enabled the AFC_LO setting must always be less than the AFC_HI setting.</p> <p>Note: This level is also used for hard-wired flow control when HW_FC_EN is set in the Port 0 Manual Flow Control Register (MANUAL_FC_0).</p>	R/W	00h
7:4	<p>Backpressure Duration (BACK_DUR) When the Host MAC automatically asserts back pressure, it will be asserted for this period of time. In full-duplex mode, this field has no function and is not used. Please refer to Table 11-16, describing Backpressure Duration bit mapping for more information.</p>	R/W	0h

Bits	Description	Type	Default
3	Flow Control on Multicast Frame (FCMULT) When this bit is set, the Host MAC will assert back pressure when the AFC level is reached and a multicast frame is received. This field has no function in full-duplex mode. 0: Flow Control on Multicast Frame Disabled 1: Flow Control on Multicast Frame Enabled	R/W	0b
2	Flow Control on Broadcast Frame (FCBRD) When this bit is set, the Host MAC will assert back pressure when the AFC level is reached and a broadcast frame is received. This field has no function in full-duplex mode. 0: Flow Control on Broadcast Frame Disabled 1: Flow Control on Broadcast Frame Enabled	R/W	0b
1	Flow Control on Address Decode (FCADD) When this bit is set, the Host MAC will assert back pressure when the AFC level is reached and a frame addressed to the Host MAC is received. This field has no function in full-duplex mode. 0: Flow Control on Address Decode Disabled 1: Flow Control on Address Decode Enabled	R/W	0b
0	Flow Control on Any Frame (FCANY) When this bit is set, the Host MAC will assert back pressure, or transmit a pause frame when the AFC level is reached and any frame is received. Setting this bit enables full-duplex flow control when the Host MAC is operating in full-duplex mode. When this mode is enabled during half-duplex operation, the Flow Controller does not decode the Host MAC address and will send a JAM upon receipt of a valid preamble (i.e., immediately at the beginning of the next frame after the RX Data FIFO level is reached). When this mode is enabled during full-duplex operation, the Flow Controller will immediately instruct the Host MAC to send a pause frame when the RX Data FIFO level is reached. The MAC will queue the pause frame transmission for the next available window. Setting this bit overrides bits [3:1] of this register.	R/W	0b

TABLE 11-16: BACKPRESSURE DURATION BIT MAPPING

	Backpressure Duration	
[7:4]	100Mbps Mode	10Mbps Mode
0h	5 us	7.2 us
1h	10 us	12.2 us
2h	15 us	17.2 us
3h	25 us	27.2 us
4h	50 us	52.2 us
5h	100 us	102.2 us
6h	150 us	152.2 us
7h	200 us	202.2 us
8h	250 us	252.2 us
9h	300 us	302.2 us
Ah	350 us	352.2 us
Bh	400 us	402.2 us
Ch	450 us	452.2 us
Dh	500 us	502.2 us
Eh	550 us	552.2 us
Fh	600 us	602.2 us

Note: Backpressure Duration is timed from when the RX FIFO reaches the level set in [Automatic Flow Control High Level \(AFC_HI\)](#), regardless when actual backpressure occurs.

11.14 Host MAC Control and Status Registers

This section details the indirectly addressable Host MAC System CSRs. These registers are located in the Host MAC and are accessed indirectly via the system CSRs. [Table 11-17](#) lists Host MAC registers that are accessible through the indexing method using the [Host MAC CSR Interface Command Register \(MAC_CSR_CMD\)](#) and [Host MAC CSR Interface Data Register \(MAC_CSR_DATA\)](#).

The Host MAC registers allow configuration of the various Host MAC parameters including the Host MAC address, flow control, multicast hash table, and wake-up configuration. The Host MAC CSRs also provide serial access to the PHYs via the registers HMAC_MII_ACC and HMAC_MII_DATA. These registers allow access to the 10/100 Ethernet PHY registers and the switch engine Port 0 Virtual PHY.

TABLE 11-17: HOST MAC ADDRESSABLE REGISTERS

Address (Indirect)	Register Name (SYMBOL)
00h	Reserved for Future Use (RESERVED)
01h	Host MAC Control Register (HMAC_CR)
02h	Host MAC Address High Register (HMAC_ADDRH)
03h	Host MAC Address Low Register (HMAC_ADDRL)
04h	Host MAC Multicast Hash Table High Register (HMAC_HASHH)
05h	Host MAC Multicast Hash Table Low Register (HMAC_HASHL)
06h	Host MAC MII Access Register (HMAC_MII_ACC)
07h	Host MAC MII Data Register (HMAC_MII_DATA)
08h	Host MAC Flow Control Register (HMAC_FLOW)
09h	Host MAC VLAN1 Tag Register (HMAC_VLAN1)
0Ah	Host MAC VLAN2 Tag Register (HMAC_VLAN2)
0Bh	Host MAC Wake-up Frame Filter Register (HMAC_WUFF)
0Ch	Host MAC Wake-up Control and Status Register (HMAC_WUCSR)
0Dh	Host MAC Checksum Offload Engine Control Register (HMAC_COE_CR)
0Eh-0Fh	Reserved for Future Use (RESERVED)
10h-FFh	Reserved for Future Use (RESERVED)

11.14.1 HOST MAC CONTROL REGISTER (HMAC_CR)

Offset: 01h Size: 32 bits

This read/write register establishes the RX and TX operation modes and controls for address filtering and packet filtering.

Bits 19-15, 13, and 11 determine if the Host MAC accepts the packets from the switch fabric. The switch fabric address table and configuration determine which packets get sent to the Host MAC.

Bits	Description	Type	Default
31	Receive All Mode (RXALL) When set, all incoming packets will be received and passed on to the address filtering function for processing of the selected filtering mode on the received frame. Address filtering then occurs and is reported in Receive Status. When cleared, only frames that pass Destination Address filtering will be sent to the application.	R/W	0b
30:26	RESERVED	RO	-
25	RESERVED	RO	-
24	RESERVED	RO	0b
23	Disable Receive Own (RCVOWN) When set, the Host MAC disables the reception of frames when it is transmitting (TXEN output is asserted). When cleared, the Host MAC receives all packets, including those transmitted by the Host MAC. This bit has no effect when the Full Duplex Mode (FDPX) bit is set.	R/W	0b
22	RESERVED	RO	-
21	Loopback operation Mode (LOOPBK) Selects the loop back operation modes for the Host MAC. This field is only valid for full duplex mode. In internal loopback mode, the TX frame is received by the internal MII interface, and sent back to the Host MAC without being sent to the network. 0: Normal Operation. Loopback disabled. 1: Loopback enabled Note: When enabling or disabling the loopback mode it can take up to 10 μ s for the mode change to occur. The transmitter and receiver must be stopped and disabled when modifying the LOOPBK bit. The transmitter or receiver should not be enabled within 10 μ s of modifying the LOOPBK bit.	R/W	0b
20	Full Duplex Mode (FDPX) When set, the Host MAC operates in Full-Duplex mode, in which it can transmit and receive simultaneously.	R/W	0b
19	Pass All Multicast (MCPAS) When set, indicates that all incoming frames with a Multicast destination address (first bit in the destination address field is 1) are received. Incoming frames with physical address (Individual Address/Unicast) destinations are filtered and received only if the address matches the Host MAC Address.	R/W	0b

Bits	Description	Type	Default
18	Promiscuous Mode (PRMS) When set, indicates that any incoming frame is received regardless of its destination address.	R/W	1b
17	Inverse filtering (INVFLT) When set, the address check function operates in inverse filtering mode. This is valid only during Perfect filtering mode. Refer to Section 11.4.4, "Inverse Filtering," on page 156 for additional information.	R/W	0b
16	Pass Bad Frames (PASSBAD) When set, all incoming frames that passed address filtering are received, including runt frames and collided frames. Refer to Section 11.4, "Address Filtering," on page 155 for additional information.	R/W	0b
15	Hash Only Filtering mode (HO) When set, the address check Function operates in the Imperfect address filtering mode for both physical and multicast addresses. Refer to Section 11.4.2, "Hash Only Filtering," on page 156 for additional information.	R/W	0b
14	RESERVED	RO	-
13	Hash/Perfect Filtering Mode (HPFLT) When cleared (0), the device will implement a perfect address filter on incoming frames according to the address specified in the Host MAC address registers (Host MAC Address High Register (HMAC_ADDRH) and Host MAC Address Low Register (HMAC_ADDRL)). When set (1), the address check function performs imperfect address filtering of multicast incoming frames according to the hash table specified in the multicast hash table register. If the Hash Only Filtering mode (HO) bit 15 is set, then the physical (IA) addresses are also imperfect filtered. If the Hash Only Filtering mode (HO) bit is cleared, then the IA addresses are perfect address filtered according to the MAC Address register Refer to Section 11.4.3, "Hash Perfect Filtering," on page 156 for additional information.	R/W	0b
12	RESERVED	RO	0b
11	Disable Broadcast Frames (BCAST) When set, disables the reception of broadcast frames. When cleared, forwards all broadcast frames to the application. Note: When wake-up frame detection is enabled via the WUEN bit of the Host MAC Wake-up Control and Status Register (HMAC_WUCSR) , a broadcast wake-up frame will wake-up the device despite the state of this bit.	R/W	0b
10	Disable Retry (DISRTY) When set, the Host MAC attempts only one transmission. When a collision is seen on the network, the Host MAC ignores the current frame and goes to the next frame and a retry error is reported in the Transmit status. When reset, the Host MAC attempts 16 transmissions before signaling a retry error.	R/W	0b
9	RESERVED	RO	-

Bits	Description	Type	Default										
8	Automatic Pad Stripping (PADSTR) When set, the Host MAC strips the pad field on all incoming frames, if the length field is less than 46 bytes. The FCS field is also stripped, since it is computed at the transmitting station based on the data and pad field characters, and is invalid for a received frame that has had the pad characters stripped. Receive frames with a 46-byte or greater length field are passed to the application unmodified (FCS is not stripped). When cleared, the Host MAC passes all incoming frames to the host unmodified.	R/W	0b										
7:6	BackOff Limit (BOLMT) The BOLMT bits allow the user to set the back-off limit in a relaxed or aggressive mode. According to IEEE 802.3, the Host MAC has to wait for a random number [r] of slot-times (see note) after it detects a collision, where: (eq.1) $0 < r < 2^K$ The exponent K is dependent on how many times the current frame to be transmitted has been retried, as follows: (eq.2) $K = \min(n, 10)$ where n is the current number of retries. If a frame has been retried three times, then K = 3 and r = 8 slot-times maximum. If it has been retried 12 times, then K = 10, and r = 1024 slot-times maximum. An LFSR (linear feedback shift register) 20-bit counter emulates a 20 bit random number generator, from which r is obtained. Once a collision is detected, the number of the current retry of the current frame is used to obtain K (eq.2). This value of K translates into the number of bits to use from the LFSR counter. If the value of K is 3, the Host MAC takes the value in the first three bits of the LFSR counter and uses it to count down to zero on every slot-time. This effectively causes the Host MAC to wait eight slot-times. To give the user more flexibility, the BOLMT value forces the number of bits to be used from the LFSR counter to a predetermined value as in the table below. <table><tr><th>BOLMT Value</th><th># Bits Used from LFSR Counter</th></tr><tr><td>00b</td><td>10</td></tr><tr><td>01b</td><td>8</td></tr><tr><td>10b</td><td>4</td></tr><tr><td>11b</td><td>1</td></tr></table> Thus, if the value of K = 10, the Host MAC will look at the BOLMT if it is 00b, then use the lower ten bits of the LFSR counter for the wait countdown. If the BOLMT is 10b, then it will only use the value in the first four bits for the wait countdown, etc. Note: Slot-time = 512 bit times. (See IEEE 802.3 Spec., sections 4.2.3.25 and 4.4.2.1)	BOLMT Value	# Bits Used from LFSR Counter	00b	10	01b	8	10b	4	11b	1	R/W	0b
BOLMT Value	# Bits Used from LFSR Counter												
00b	10												
01b	8												
10b	4												
11b	1												
5	Deferral Check (DFCHK) When set, enables the deferral check in the Host MAC. The Host MAC will abort the transmission attempt if it has deferred for more than 24,288 bit times. Deferral starts when the transmitter is ready to transmit, but is prevented from doing so because the CRS is active. Deferral time is not cumulative. If the transmitter defers for 10,000 bit times, then transmits, collides, backs off, and then has to defer again after completion of back-off, the deferral timer resets to 0 and restarts. When this bit is cleared, the deferral check is disabled in the Host MAC and the Host MAC defers indefinitely.	R/W	0b										
4	RESERVED	RO	-										

LAN9352

Bits	Description	Type	Default
3	Transmitter enable (TXEN) When set, the Host MAC's transmitter is enabled and it will transmit frames from the buffer. When cleared, the Host MAC's transmitter is disabled and will not transmit any frames.	R/W	0b
2	Receiver Enable (RXEN) When set, the Host MAC's receiver is enabled and will receive frames. When cleared, the MAC's receiver is disabled and will not receive any frames.	R/W	0b
1:0	RESERVED	RO	-

11.14.2 HOST MAC ADDRESS HIGH REGISTER (HMAC_ADDRH)

Offset: 02h Size: 32 bits

This read/write register contains the upper 16-bits of the physical address of the Host MAC. The contents of this register are optionally loaded from the EEPROM at power-on through the EEPROM Loader if a programmed EEPROM is detected. The least significant byte of this register (bits [7:0]) is loaded from address 05h of the EEPROM. The second byte (bits [15:8]) is loaded from address 06h of the EEPROM. [Section 11.9, "Host MAC Address," on page 169](#) details the byte ordering of the HMAC_ADDRH and HMAC_ADDRH registers with respect to the reception of the Ethernet physical address. Please refer to [Section 14.4, "EEPROM Loader," on page 465](#) for more information on the EEPROM Loader.

This register used to specify the address used for Perfect DA, Magic Packet and Wakeup frames, the unicast destination address for received pause frames, and the source address for transmitted pause frames. This register is not used for packet filtering.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	Physical Address [47:32] This field contains the upper 16-bits (47:32) of the Physical Address of the Host MAC. The content of this field is undefined until loaded from the EEPROM at power-on. The host can update the contents of this field after the initialization process has completed.	R/W	FFFFh

11.14.3 HOST MAC ADDRESS LOW REGISTER (HMAC_ADDRL)

Offset: 03h Size: 32 bits

This read/write register contains the lower 32-bits of the physical address of the Host MAC. The contents of this register are optionally loaded from the EEPROM at power-on through the EEPROM Loader if a programmed EEPROM is detected. The least significant byte of this register (bits [7:0]) is loaded from address 01h of the EEPROM. The most significant byte of this register is loaded from address 04h of the EEPROM. [Section 11.9, "Host MAC Address," on page 169](#) details the byte ordering of the HMAC_ADDRL and HMAC_ADDRH registers with respect to the reception of the Ethernet physical address. Please refer to [Section 14.4, "EEPROM Loader," on page 465](#) for more information on the EEPROM Loader.

This register used to specify the address used for Perfect DA, Magic Packet and Wakeup frames, the unicast destination address for received pause frames, and the source address for transmitted pause frames. This register is not used for packet filtering.

BITS	DESCRIPTION	TYPE	DEFAULT
31:0	Physical Address [31:0] This field contains the lower 32-bits (31:0) of the Physical Address of the Host MAC. The content of this field is undefined until loaded from the EEPROM at power-on. The host can update the contents of this field after the initialization process has completed.	R/W	FFFFFFFFh

11.14.4 HOST MAC MULTICAST HASH TABLE HIGH REGISTER (HMAC_HASHH)

Offset: 04h Size: 32 bits

The 64-bit Multicast table is used for group address filtering. For hash filtering, the contents of the destination address in the incoming frame is used to index the contents of the Hash table. The most significant bit determines the register to be used (Hi/Low), while the other five bits determine the bit within the register. A value of 00000 selects Bit 0 of the Multicast Hash Table Lo register and a value of 11111 selects the Bit 31 of the Multicast Hash Table Hi register.

If the corresponding bit is 1, then the multicast frame is accepted. Otherwise, it is rejected. If the "Pass All Multicast" (MCPAS) bit of the [Host MAC Control Register \(HMAC_CR\)](#) is set, then all multicast frames are accepted regardless of the multicast hash values.

The Multicast Hash Table High register contains the higher 32 bits of the hash table and the Multicast Hash Table Low register contains the lower 32 bits of the hash table. Refer to [Section 11.4, "Address Filtering," on page 155](#) for more information on address filtering.

This table determines if the Host MAC accepts the packets from the switch fabric. The switch fabric address table and configuration determine the packets that get sent to the Host MAC.

Bits	Description	Type	Default
31:0	Upper 32-bits of the 64-bit Hash Table	R/W	00000000h

11.14.5 HOST MAC MULTICAST HASH TABLE LOW REGISTER (HMAC_HASHL)

Offset: 05h Size: 32 bits

This read/write register defines the lower 32-bits of the Multicast Hash Table. Please refer to the [Host MAC Multicast Hash Table High Register \(HMAC_HASHH\)](#) and [Section 11.4, "Address Filtering," on page 155](#) for more information.

Bits	Description	Type	Default
31:0	Lower 32-bits of the 64-bit Hash Table	R/W	00000000h

11.14.6 HOST MAC MII ACCESS REGISTER (HMAC_MII_ACC)

Offset: 06h Size: 32 bits

This read/write register is used in conjunction with the [Host MAC MII Data Register \(HMAC_MII_DATA\)](#) to access the internal PHY registers. Refer to [Section 12.2.19, "Physical PHY Registers"](#) and [Section 12.3.3, "Virtual PHY Registers"](#) for a list of accessible PHY registers and PHY address information.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:11	PHY Address (PHY_ADDR) This field must be loaded with the PHY address that the MII access is intended for. A list of default PHY addresses can be seen in Table 12-1 . Refer to Section 12.1.1, "PHY Addressing," on page 218 for additional information on PHY addressing.	R/W	00000b
10:6	MII Register Index (MIIRINDA) These bits select the desired MII register in the PHY.	R/W	00000b
5:2	RESERVED	RO	-
1	MII Write (MIIWnR) Setting this bit tells the PHY that this will be a write operation using the Host MAC MII Data Register (HMAC_MII_DATA) . If this bit is cleared, a read operation will occur, packing the data in the Host MAC MII Data Register (HMAC_MII_DATA) .	R/W	0b
0	MII Busy (MIIBZY) This bit must be polled to determine when the MII register access is complete. This bit must read a logical 0 before writing to this register or the Host MAC MII Data Register (HMAC_MII_DATA) . The LAN driver software must set this bit in order for the device to read or write any of the MII PHY registers. During a MII register access, this bit will be set, signifying a read or write access is in progress. The MII data register must be kept valid until the Host MAC clears this bit during a PHY write operation. The MII data register is invalid until the Host MAC has cleared this bit during a PHY read operation.	RO SC	0b

11.14.7 HOST MAC MII DATA REGISTER (HMAC_MII_DATA)

Offset: 07h Size: 32 bits

This read/write register is used in conjunction with the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#) to access the internal PHY registers. This register contains either the data to be written to the PHY register specified in the HMAC_MII_ACC Register, or the read data from the PHY register whose index is specified in the HMAC_MII_ACC Register.

The [MII Busy \(MIIBZY\)](#) bit in the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#) must be cleared when writing to this register.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	MII Data This field contains the 16-bit value read from the PHY read operation or the 16-bit data value to be written to the PHY before an MII write operation.	R/W	0000h

11.14.8 HOST MAC FLOW CONTROL REGISTER (HMAC_FLOW)

Offset: 08h Size: 32 bits

This read/write register controls the generation and reception of the Control (Pause command) frames by the Host MAC's flow control block. The control frame fields are selected as specified in the 802.3 Specification and the Pause-Time value from this register is used in the "Pause Time" field of the control frame. In full-duplex mode the FCBSY bit is set until the control frame is completely transferred. The host has to make sure that the FCBSY bit is cleared before writing the register. The Pass Control Frame bit (FCPASS) does not affect the sending of the frames, including Control Frames, to the host. The Flow Control Enable (FCEN) bit enables the receive portion of the Flow Control block.

This register is used in conjunction with the [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#) in the System CSRs to configure flow control. Software flow control is initiated using the AFC_CFG register.

The Host MAC will not transmit pause frames or assert back pressure if the transmitter is disabled.

For the Host MAC, flow control/backpressure is to/from the switch fabric, not the external network.

Bits	Description	Type	Default
31:16	Pause Time (FCPT) This field indicates the value to be used in the PAUSE TIME field in the control frame. This field must be initialized before full-duplex automatic flow control is enabled.	R/W	0000h
15:3	RESERVED	RO	-
2	Pass Control Frames (FCPASS) When set, the Host MAC sets the packet filter bit in the receive packet status to indicate to the application that a valid pause frame has been received. The application must accept or discard a received frame based on the packet filter control bit. The Host MAC receives, decodes and performs the Pause function when a valid Pause frame is received in Full-Duplex mode and when flow control is enabled (FCE bit set). When this bit is cleared, the Host MAC resets the Packet Filter bit in the Receive packet status. The Host MAC always passes the data of all frames it receives (including flow control frames) to the application. Frames that do not pass address filtering, as well as frames with errors, are passed to the application. The application must discard or retain the received frame's data based on the received frame's STATUS field. Filtering modes (promiscuous mode, for example) take precedence over the FCPASS bit.	R/W	0b
1	Flow Control Enable (FCEN) When set, enables the Host MAC flow control function. The Host MAC decodes all incoming frames for control frames; if it receives a valid control frame (PAUSE command), it disables the transmitter for a specified time (Decoded pause time x slot time). When this bit is cleared, the Host MAC flow control function is disabled; the MAC does not decode frames for control frames. Note: Flow Control is applicable when the Host MAC is set in full duplex mode. In half-duplex mode, this bit enables the backpressure function to control the flow of received frames to the Host MAC.	R/W	0b

Bits	Description	Type	Default
0	<p>Flow Control Busy (FCBSY) In full-duplex mode, this bit indicates that the Host MAC is in the process of sending a PAUSE control frame. This bit is set by the automatic flow control function.</p> <p>During the transmission of the control frame, this bit continues to be set, signifying that a frame transmission is in progress. After the PAUSE control frame's transmission is complete, the Host MAC resets the bit to 0.</p> <p>The host software should read a logical 0 from this bit before writing to the Host MAC Flow Control (HMAC_FLOW) register.</p>	R/W SC	0b

LAN9352

11.14.9 HOST MAC VLAN1 TAG REGISTER (HMAC_VLAN1)

Offset: 09h Size: 32 bits

This read/write register contains the VLAN tag field to identify VLAN1 frames. When a VLAN1 frame is detected, the legal frame length is increased from 1518 bytes to 1522 bytes. Refer to [Section 11.3, "Virtual Local Area Network \(VLAN\) Support," on page 153](#) for additional information.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	VLAN1 Tag Identifier (VTI1) This field contains the VLAN Tag used to identify VLAN1 frames. This field is compared with the 13th and 14th bytes of the incoming frames for VLAN1 frame detection. Note: If used, this register is typically set to the standard VLAN value of 8100h.	R/W	FFFFh

11.14.10 HOST MAC VLAN2 TAG REGISTER (HMAC_VLAN2)

Offset: 0Ah Size: 32 bits

This read/write register contains the VLAN tag field to identify VLAN2 frames. When a VLAN2 frame is detected, the legal frame length is increased from 1518 bytes to 1538 bytes. Refer to [Section 11.3, "Virtual Local Area Network \(VLAN\) Support," on page 153](#) for additional information.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	VLAN2 Tag Identifier (VTI2) This field contains the VLAN Tag used to identify VLAN2 frames. This field is compared with the 13th and 14th bytes of the incoming frames for VLAN2 frame detection. Note: If used, this register is typically set to the standard VLAN value of 8100h. If both VLAN1 and VLAN2 Tag Identifiers are used, they should be unique. If both are set to the same value, VLAN1 is given higher precedence and the maximum legal frame length is set to 1522.	R/W	FFFFh

11.14.11 HOST MAC WAKE-UP FRAME FILTER REGISTER (HMAC_WUFF)

Offset: 0Bh Size: 32 bits

This write-only register is used to configure the wake-up frame filter. Refer to [Section 11.6.3, "Wake-up Frame Detection," on page 158](#) for additional information.

Bits	Description	Type	Default
31:0	<p>Wake-Up Frame Filter (WFF) The Wake-up frame filter is configured through this register using an indexing mechanism. After power-on reset, digital reset, or Host MAC reset, the Host MAC loads the first value written to this location to the first DWORD in the Wake-up frame filter (filter 0 byte mask 0). The second value written to this location is loaded to the second DWORD in the wake-up frame filter (filter 0 byte mask 1) and so on. Once all 40 DWORDs have been written, the internal pointer will once again point to the first entry and the filter entries can be modified in the same manner. Similarly, 40 DWORDs can be read sequentially to obtain the values stored in the WFF. Please refer to Section 11.6.3, "Wake-up Frame Detection," on page 158 for further information.</p> <p>Note: This register should be read and written using 40 consecutive DWORD operations. Failure to read or write the entire contents of the WFF may cause the internal read/write pointers to be left in a position other than pointing to the first entry. A mechanism for resetting the internal pointers to the beginning of the WFF is available via the WFF Pointer Reset (WFF_PTR_RST) bit of the Host MAC Wake-up Control and Status Register (HMAC_WUCSR). This mechanism enables the application program to re-synchronize with the internal WFF pointers if it has not previously read/written the complete contents of the WFF.</p>	R/W	-

11.14.12 HOST MAC WAKE-UP CONTROL AND STATUS REGISTER (HMAC_WUCSR)

Offset: 0Ch Size: 32 bits

This read/write register contains data and control settings pertaining to the Host MAC's remote wake-up status and capabilities. It is used in conjunction with the [Host MAC Wake-up Frame Filter Register \(HMAC_WUFF\)](#) to fully configure the wake-up frame filter. Refer to [Section 11.6.3, "Wake-up Frame Detection,"](#) on page 158 for additional information.

Bits	Description	Type	Default
31	WFF Pointer Reset (WFF_PTR_RST) This self-clearing bit resets the Wakeup Frame Filter (WFF) internal read and write pointers to the beginning of the WFF.	SC	0b
30:10	RESERVED	RO	-
9	Global Unicast Enable (GUE) When set, the Host MAC wakes up from power-saving mode on receipt of a global unicast frame. This is accomplished by enabling global unicasts as a wakeup frame qualifier. A global unicast frame has the MAC Address [0] bits set to 0. Note: The Wake-Up Frame Enable (WUEN) bit of this register must also be set to enable wakeup.	R/W	0b
8	WoL Wait for Sleep (WOL_WAIT_SLEEP) When set, the WoL functions are not active until the device has entered a sleep state. When clear, WoL functions are active immediately.	R/W	0b
7	Perfect DA Frame Received (PFDA_FR) The MAC sets this bit upon receiving a valid frame with a destination address that matches the physical address.	R/WC	0b
6	Remote Wake-Up Frame Received (WUFR) The Host MAC sets this bit upon receiving a valid Remote Wake-up frame.	R/WC	0b
5	Magic Packet Received (MPR) The Host MAC sets this bit upon receiving a valid Magic Packet	R/WC	0b
4	Broadcast Frame Received (BCAST_FR) The MAC sets this bit upon receiving a valid broadcast frame.	R/WC	0b
3	Perfect DA Wakeup Enable (PFDA_EN) When set, remote wakeup mode is enabled and the MAC is capable of waking up on receipt of a frame with a destination address that matches the physical address of the device. The physical address is stored in the Host MAC Address High Register (HMAC_ADDRH) and Host MAC Address Low Register (HMAC_ADDRL) .	R/W	0b
2	Wake-Up Frame Enable (WUEN) When set, Remote Wake-Up mode is enabled and the Host MAC is capable of detecting wake-up frames as programmed in the Host MAC Wake-up Frame Filter Register (HMAC_WUFF) .	R/W	0b
1	Magic Packet Enable (MPEN) When set, Magic Packet Wake-up mode is enabled.	R/W	0b
0	Broadcast Wakeup Enable (BCST_EN) When set, remote wakeup mode is enabled and the MAC is capable of waking up from a broadcast frame.	R/W	0b

11.14.13 HOST MAC CHECKSUM OFFLOAD ENGINE CONTROL REGISTER (HMAC_COE_CR)

Offset: 0Dh Size: 32 bits

This register controls the RX and TX checksum offload engines.

Bits	Description	Type	Default
31:17	RESERVED	RO	-
16	TX Checksum Offload Engine Enable (TX_COE_EN) TX_COE_EN may only be changed if the TX path is disabled. If it is desired to change this value during run time, it is safe to do so only after the Host MAC is disabled and the TLI is empty. 0 = The TXCOE is bypassed 1 = The TXCOE is enabled	R/W	0b
15:2	RESERVED	RO	-
1	RX Checksum Offload Engine Mode (RX_COE_MODE) This register indicates whether the COE will check for VLAN tags or a SNAP header prior to beginning its checksum calculation. In its default mode, the calculation will always begin 14 bytes into the frame. RX_COE_MODE may only be changed if the RX path is disabled. If it is desired to change this value during run time, it is safe to do so only after the Host MAC is disabled and the TLI is empty. 0 = Begin checksum calculation after first 14 bytes of Ethernet Frame 1 = Begin checksum calculation at start of L3 packet by adjusting for VLAN tags and/or SNAP header.	R/W	0b
0	RX Checksum Offload Engine Enable (RX_COE_EN) RX_COE_EN may only be changed if the RX path is disabled. If it is desired to change this value during run time, it is safe to do so only after the Host MAC is disabled and the TLI is empty. 0 = The RXCOE is bypassed 1 = The RXCOE is enabled	R/W	0b

12.0 ETHERNET PHYS

12.1 Functional Overview

The device contains Physical PHYs A and B, and a Virtual PHY.

The A and B Physical PHYs are identical in functionality. Physical PHY A connects to the Switch Fabric Port 1. Physical PHY B connects to Switch Fabric Port 2. These PHYs interface with their respective MAC via an internal MII interface.

The Virtual PHY provides the virtual functionality of a PHY and allows connection of the Host MAC to Port 0 of the switch fabric as if it was connected to a single port PHY.

The Physical PHYs comply with the IEEE 802.3 Physical Layer for Twisted Pair Ethernet and can be configured for full/half duplex 100 Mbps (100BASE-TX / 100BASE-FX) or 10 Mbps (10BASE-T) Ethernet operation. All PHY registers follow the IEEE 802.3 (clause 22.2.4) specified MII management register set and are fully configurable.

The device Ethernet PHYs are discussed in detail in the following sections:

- [Section 12.2, "Physical PHYs A & B," on page 218](#)
- [Section 12.3, "Virtual PHY," on page 303](#)

12.1.1 PHY ADDRESSING

Each individual PHY is assigned a default PHY address via the [phy_addr_sel_strap](#) configuration strap as shown in [Table 12-1](#).

In addition, the addresses for PHY A and B can be changed via the [PHY Address \(PHYADD\)](#) field in the [PHY x Special Modes Register \(PHY_SPECIAL_MODES_x\)](#). For proper operation, the addresses for the Virtual PHY and Physical PHYs A and B must be unique. No check is performed to assure each PHY is set to a different address.

TABLE 12-1: DEFAULT PHY SERIAL MII ADDRESSING

phy_addr_sel_strap	Virtual PHY Default Address Value	PHY A Default Address Value	PHY B Default Address Value
0	0	1	2
1	1	2	3

12.2 Physical PHYs A & B

The device integrates two IEEE 802.3 PHY functions. The PHYs can be configured for either 100 Mbps copper (100BASE-TX), 100 Mbps fiber (100BASE-FX) or 10 Mbps copper (10BASE-T) Ethernet operation and include Auto-Negotiation and HP Auto-MDIX.

Note: Because the Physical PHYs A and B are functionally identical, this section will describe them as the “Physical PHY x”, or simply “PHY”. Wherever a lowercase “x” has been appended to a port or signal name, it can be replaced with “A” or “B” to indicate the PHY A or PHY B respectively. In some instances, a “1” or a “2” may be appropriate instead. All references to “PHY” in this section can be used interchangeably for both the Physical PHYs A and B.

12.2.1 FUNCTIONAL DESCRIPTION

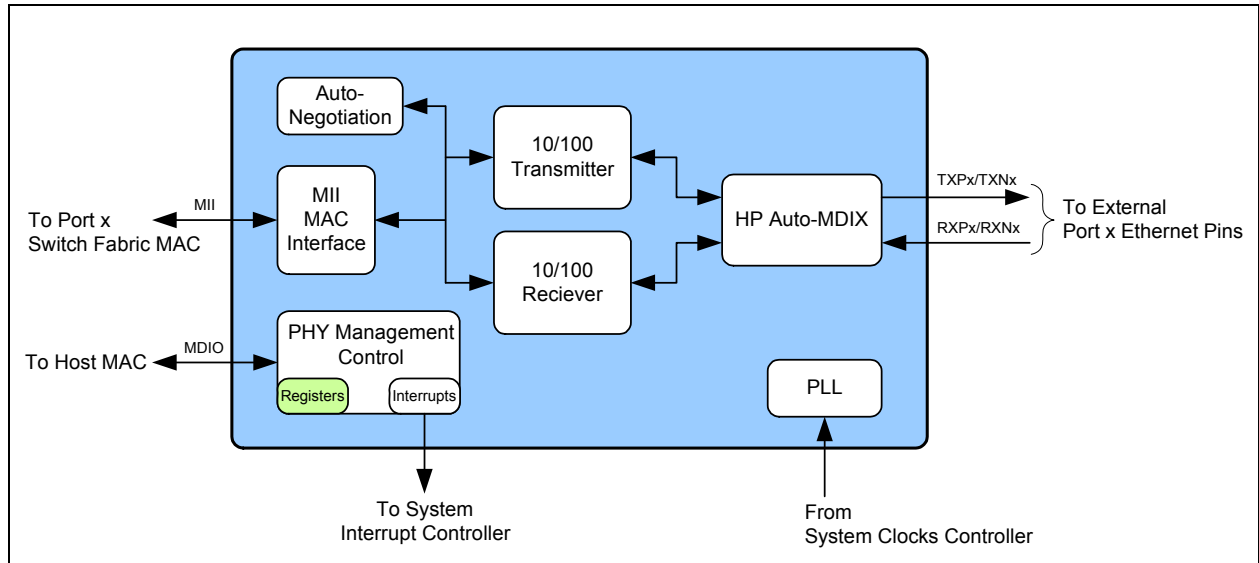
Functionally, each PHY can be divided into the following sections:

- [100BASE-TX Transmit](#) and [100BASE-TX Receive](#)
- [10BASE-T Transmit](#) and [10BASE-T Receive](#)
- [Auto-Negotiation](#)
- [HP Auto-MDIX](#)
- [PHY Management Control](#) and [PHY Interrupts](#)
- [PHY Power-Down Modes](#) and [Energy Efficient Ethernet](#)
- [Wake on LAN \(WoL\)](#)
- [Resets](#)

- [Link Integrity Test](#)
- [Cable Diagnostics](#)
- [Loopback Operation](#)
- [100BASE-FX Far End Fault Indication](#)

A block diagram of the main components of each PHY can be seen in [Figure 12-1](#).

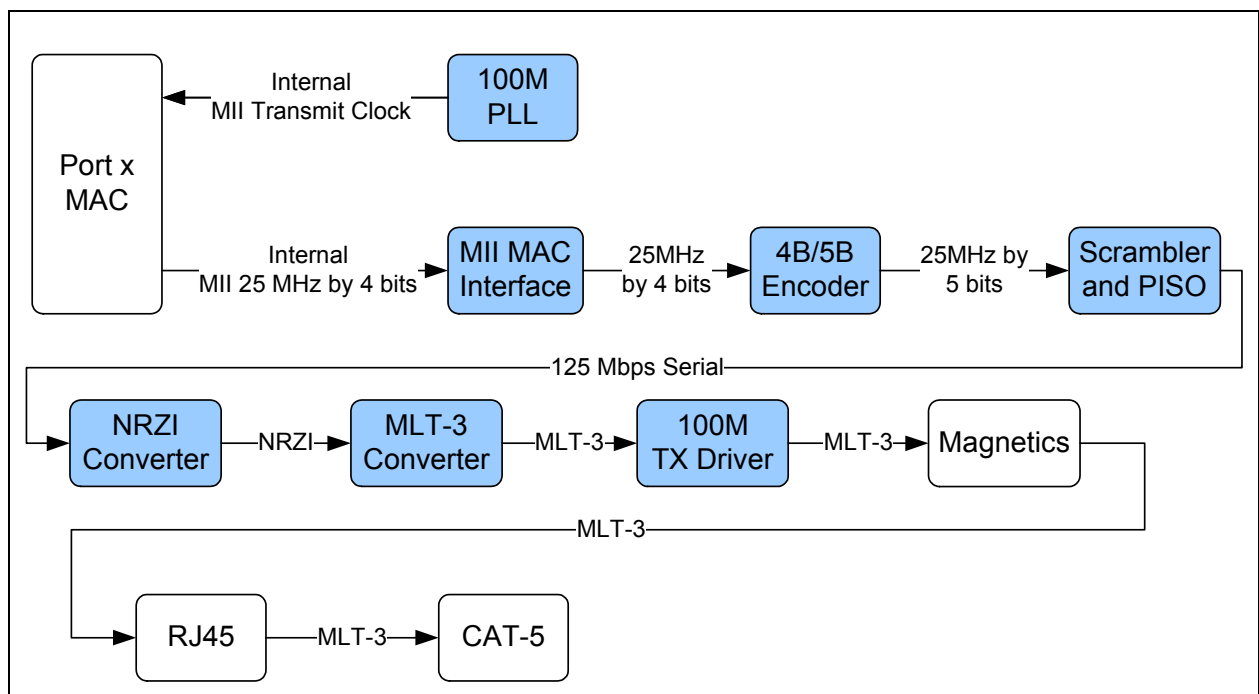
FIGURE 12-1: PHYSICAL PHY BLOCK DIAGRAM



12.2.2 100BASE-TX TRANSMIT

The 100BASE-TX transmit data path is shown in [Figure 12-2](#). Shaded blocks are those which are internal to the PHY. Each major block is explained in the following sections.

FIGURE 12-2: 100BASE-TX TRANSMIT DATA PATH



LAN9352

12.2.2.1 100BASE-TX Transmit Data Across the Internal MII Interface

For a transmission, the Switch Fabric MAC drives the transmit data onto the internal MII TXD bus and asserts the internal MII TXEN to indicate valid data. The data is in the form of 4-bit wide 25 MHz data.

12.2.2.2 4B/5B Encoder

The transmit data passes from the MII block to the 4B/5B Encoder. This block encodes the data from 4-bit nibbles to 5-bit symbols (known as “code-groups”) according to [Table 12-2](#). Each 4-bit data-nibble is mapped to 16 of the 32 possible code-groups. The remaining 16 code-groups are either used for control information or are not valid.

The first 16 code-groups are referred to by the hexadecimal values of their corresponding data nibbles, 0 through F. The remaining code-groups are given letter designations with slashes on either side. For example, an IDLE code-group is /I/, a transmit error code-group is /H/, etc.

TABLE 12-2: 4B/5B CODE TABLE

Code Group	Sym	Receiver Interpretation			Transmitter Interpretation		
11110	0	0	0000	DATA	0	0000	DATA
01001	1	1	0001		1	0001	
10100	2	2	0010		2	0010	
10101	3	3	0011		3	0011	
01010	4	4	0100		4	0100	
01011	5	5	0101		5	0101	
01110	6	6	0110		6	0110	
01111	7	7	0111		7	0111	
10010	8	8	1000		8	1000	
10011	9	9	1001		9	1001	
10110	A	A	1010		A	1010	
10111	B	B	1011		B	1011	
11010	C	C	1100		C	1100	
11011	D	D	1101		D	1101	
11100	E	E	1110		E	1110	
11101	F	F	1111		F	1111	
11111	/I/	IDLE			Sent after /T/R/ until the MII Transmitter Enable signal (TXEN) is received		
11000	/J/	First nibble of SSD, translated to “0101” following IDLE, else MII Receive Error (RXER)			Sent for rising MII Transmitter Enable signal (TXEN)		
10001	/K/	Second nibble of SSD, translated to “0101” following J, else MII Receive Error (RXER)			Sent for rising MII Transmitter Enable signal (TXEN)		
01101	/T/	First nibble of ESD, causes de-assertion of CRS if followed by /R/, else assertion of MII Receive Error (RXER)			Sent for falling MII Transmitter Enable signal (TXEN)		

TABLE 12-2: 4B/5B CODE TABLE (CONTINUED)

Code Group	Sym	Receiver Interpretation	Transmitter Interpretation
00111	/R/	Second nibble of ESD, causes de-assertion of CRS if following /T/, else assertion of MII Receive Error (RXER)	Sent for falling MII Transmitter Enable signal (TXEN)
00100	/H/	Transmit Error Symbol	Sent for rising MII Transmit Error (TXER)
00110	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
11001	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
00000	/P/	SLEEP, Indicates to receiver that the transmitter will be going to LPI	Sent due to LPI. Used to tell receiver before transmitter goes to LPI. Also used for refresh cycles during LPI.
00001	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
00010	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
00011	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
00101	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
01000	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
01100	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID
10000	/V/	INVALID, MII Receive Error (RXER) if during MII Receive Data Valid (RXDV)	INVALID

12.2.2.3 Scrambler and PISO

Repeated data patterns (especially the IDLE code-group) can have power spectral densities with large narrow-band peaks. Scrambling the data helps eliminate these peaks and spread the signal power more uniformly over the entire channel bandwidth. This uniform spectral density is required by FCC regulations to prevent excessive EMI from being radiated by the physical wiring.

The seed for the scrambler is generated from the PHY address, ensuring that each PHY will have its own scrambler sequence. For more information on PHY addressing, refer to [Section 12.1.1, "PHY Addressing"](#).

The scrambler also performs the Parallel In Serial Out conversion (PISO) of the data.

12.2.2.4 NRZI and MLT-3 Encoding

The scrambler block passes the 5-bit wide parallel data to the NRZI converter where it becomes a serial 125MHz NRZI data stream. The NRZI is then encoded to MLT-3. MLT-3 is a tri-level code where a change in the logic level represents a code bit "1" and the logic output remaining at the same level represents a code bit "0".

12.2.2.5 100M Transmit Driver

The MLT-3 data is then passed to the analog transmitter, which drives the differential MLT-3 signal on output pins TXPx and TXNx, to the twisted pair media across a 1:1 ratio isolation transformer. The 10BASE-T and 100BASE-TX signals pass through the same transformer so that common "magnetics" can be used for both. The transmitter drives into the 100 Ω impedance of the CAT-5 cable. Cable termination and impedance matching require external components.

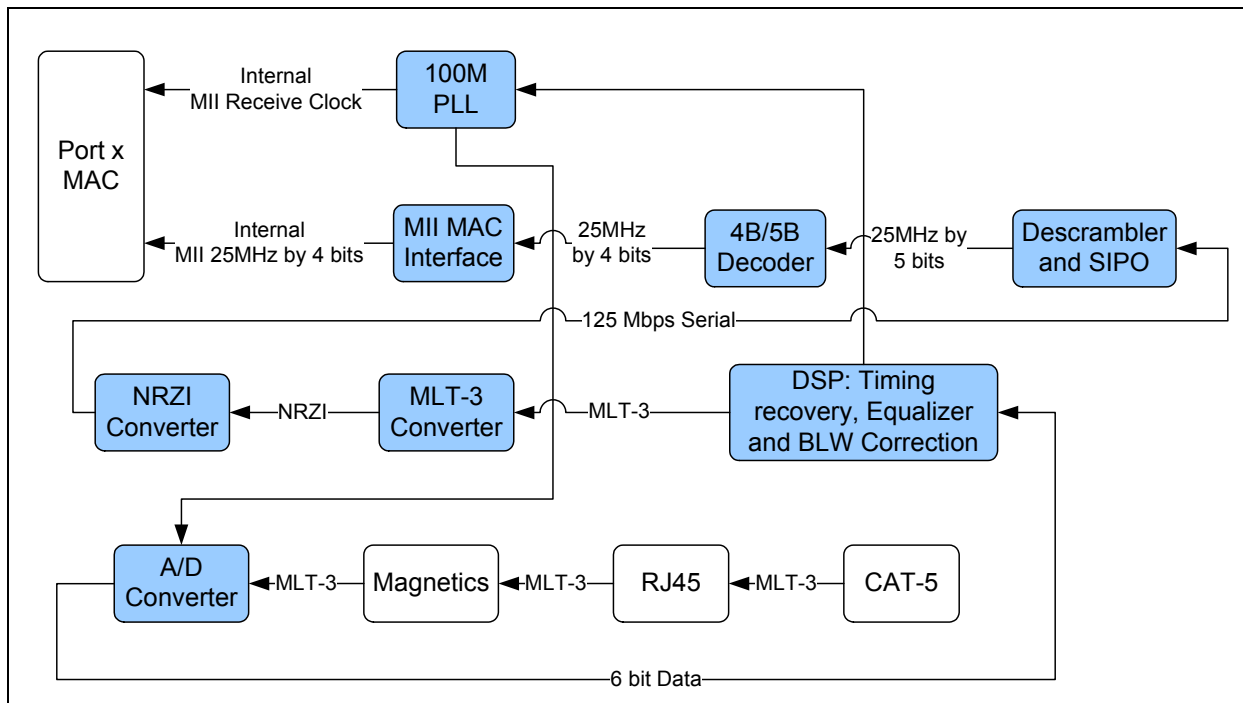
12.2.2.6 100M Phase Lock Loop (PLL)

The 100M PLL locks onto the reference clock and generates the 125 MHz clock used to drive the 125 MHz logic and the 100BASE-TX Transmitter.

12.2.3 100BASE-TX RECEIVE

The 100BASE-TX receive data path is shown in Figure 12-3. Shaded blocks are those which are internal to the PHY. Each major block is explained in the following sections.

FIGURE 12-3: 100BASE-TX RECEIVE DATA PATH



12.2.3.1 100M Receive Input

The MLT-3 data from the cable is fed into the PHY on inputs RXPx and RXNx via a 1:1 ratio transformer. The ADC samples the incoming differential signal at a rate of 125M samples per second. Using a 64-level quantizer, 6 digital bits are generated to represent each sample. The DSP adjusts the gain of the ADC according to the observed signal levels such that the full dynamic range of the ADC can be used.

12.2.3.2 Equalizer, BLW Correction and Clock/Data Recovery

The 6 bits from the ADC are fed into the DSP block. The equalizer in the DSP section compensates for phase and amplitude distortion caused by the physical channel consisting of magnetics, connectors, and CAT-5 cable. The equalizer can restore the signal for any good-quality CAT-5 cable between 1m and 100m.

If the DC content of the signal is such that the low-frequency components fall below the low frequency pole of the isolation transformer, then the droop characteristics of the transformer will become significant and Baseline Wander (BLW) on the received signal will result. To prevent corruption of the received data, the PHY corrects for BLW and can receive the ANSI X3.263-1995 FDDI TP-PMD defined “killer packet” with no bit errors.

The 100M PLL generates multiple phases of the 125MHz clock. A multiplexer, controlled by the timing unit of the DSP, selects the optimum phase for sampling the data. This is used as the received recovered clock. This clock is used to extract the serial data from the received signal.

12.2.3.3 NRZI and MLT-3 Decoding

The DSP generates the MLT-3 recovered levels that are fed to the MLT-3 converter. The MLT-3 is then converted to an NRZI data stream.

12.2.3.4 Descrambler

The descrambler performs an inverse function to the scrambler in the transmitter and also performs the Serial In Parallel Out (SIPO) conversion of the data.

During reception of IDLE (/I/) symbols, the descrambler synchronizes its descrambler key to the incoming stream. Once synchronization is achieved, the descrambler locks on this key and is able to descramble incoming data.

Special logic in the descrambler ensures synchronization with the remote transceiver by searching for IDLE symbols within a window of 4000 bytes (40 us). This window ensures that a maximum packet size of 1514 bytes, allowed by the IEEE 802.3 standard, can be received with no interference. If no IDLE-symbols are detected within this time-period, receive operation is aborted and the descrambler re-starts the synchronization process.

The de-scrambled signal is then aligned into 5-bit code-groups by recognizing the /J/K/ Start-of-Stream Delimiter (SSD) pair at the start of a packet. Once the code-word alignment is determined, it is stored and utilized until the next start of frame.

12.2.3.5 5B/4B Decoding

The 5-bit code-groups are translated into 4-bit data nibbles according to the 4B/5B table. The translated data is presented on the internal MII RXD[3:0] signal lines. The SSD, /J/K/, is translated to "0101 0101" as the first 2 nibbles of the MAC preamble. Reception of the SSD causes the transceiver to assert the receive data valid signal, indicating that valid data is available on the RXD bus. Successive valid code-groups are translated to data nibbles. Reception of either the End of Stream Delimiter (ESD) consisting of the /T/R/ symbols, or at least two /I/ symbols causes the transceiver to de-assert carrier sense and receive data valid signal.

Note: These symbols are not translated into data.

12.2.3.6 Receive Data Valid Signal

The internal MII's Receive Data Valid signal (RXDV) indicates that recovered and decoded nibbles are being presented on the RXD[3:0] outputs synchronous to RXCLK. RXDV becomes active after the /J/K/ delimiter has been recognized and RXD is aligned to nibble boundaries. It remains active until either the /T/R/ delimiter is recognized or link test indicates failure or SIGDET becomes false.

RXDV is asserted when the first nibble of translated /J/K/ is ready for transfer over the Media Independent Interface.

12.2.3.7 Receiver Errors

During a frame, unexpected code-groups are considered receive errors. Expected code groups are the DATA set (0 through F), and the /T/R/ (ESD) symbol pair. When a receive error occurs, the internal MII's RXER signal is asserted and arbitrary data is driven onto the internal MII's RXD[3:0] lines. Should an error be detected during the time that the /J/K/ delimiter is being decoded (bad SSD error), RXER is asserted true and the value 1110b is driven onto the RXD[3:0] lines. Note that the internal MII's data valid signal (RXDV) is not yet asserted when the bad SSD occurs.

12.2.3.8 100M Receive Data Across the Internal MII Interface

For reception, the 4-bit data nibbles are sent to the MII MAC Interface block. These data nibbles are clocked to the controller at a rate of 25 MHz. RXCLK is the output clock for the internal MII bus. It is recovered from the received data to clock the RXD bus. If there is no received signal, it is derived from the system reference clock.

12.2.4 10BASE-T TRANSMIT

The 10BASE-T transmitter receives 4-bit nibbles from the internal MII at a rate of 2.5 MHz and converts them to a 10 Mbps serial data stream. The data stream is then Manchester-encoded and sent to the analog transmitter, which drives a signal onto the twisted pair via the external magnetics.

10BASE-T transmissions use the following blocks:

- MII (digital)
- TX 10M (digital)
- 10M Transmitter (analog)
- 10M PLL (analog)

12.2.4.1 10M Transmit Data Across the Internal MII Interface

For a transmission, the Switch Fabric MAC drives the transmit data onto the internal MII TXD bus and asserts the internal MII TXEN to indicate valid data. The data is in the form of 4-bit wide 2.5 MHz data.

LAN9352

In half-duplex mode the transceiver loops back the transmitted data, on the receive path. This does not confuse the MAC/Controller since the COL signal is not asserted during this time. The transceiver also supports the SQE (Heartbeat) signal.

12.2.4.2 Manchester Encoding

The 4-bit wide data is sent to the 10M TX block. The nibbles are converted to a 10Mbps serial NRZI data stream. The 10M PLL produces a 20MHz clock. This is used to Manchester encode the NRZ data stream. When no data is being transmitted (internal MII TXEN is low), the 10M TX Driver block outputs Normal Link Pulses (NLPs) to maintain communications with the remote link partner.

12.2.4.3 10M Transmit Drivers

The Manchester encoded data is sent to the analog transmitter where it is shaped and filtered before being driven out as a differential signal across the TXPx and TXNx outputs.

12.2.5 10BASE-T RECEIVE

The 10BASE-T receiver gets the Manchester-encoded analog signal from the cable via the magnetics. It recovers the receive clock from the signal and uses this clock to recover the NRZI data stream. This 10M serial data is converted to 4-bit data nibbles which are passed to the controller across the internal MII at a rate of 2.5MHz.

10BASE-T reception uses the following blocks:

- Filter and SQUELCH (analog)
- 10M PLL (analog)
- RX 10M (digital)
- MII (digital)

12.2.5.1 10M Receive Input and Squelch

The Manchester signal from the cable is fed into the transceiver (on inputs RXPx and RXNx) via 1:1 ratio magnetics. It is first filtered to reduce any out-of-band noise. It then passes through a SQUELCH circuit. The SQUELCH is a set of amplitude and timing comparators that normally reject differential voltage levels below 300mV and detect and recognize differential voltages above 585mV.

12.2.5.2 Manchester Decoding

The output of the SQUELCH goes to the 10M RX block where it is validated as Manchester encoded data. The polarity of the signal is also checked. If the polarity is reversed (local RXP is connected to RXN of the remote partner and vice versa), the condition is identified and corrected. The reversed condition is indicated by the [10Base-T Polarity State \(XPOL\)](#) bit in [PHY x Special Control/Status Indication Register \(PHY_SPECIAL_CONTROL_STAT_IND_x\)](#). The 10M PLL is locked onto the received Manchester signal, from which the 20MHz clock is generated. Using this clock, the Manchester encoded data is extracted and converted to a 10MHz NRZI data stream. It is then converted from serial to 4-bit wide parallel data.

The RX10M block also detects valid 10BASE-T IDLE signals - Normal Link Pulses (NLPs) - to maintain the link.

12.2.5.3 10M Receive Data Across the Internal MII Interface

For reception, the 4-bit data nibbles are sent to the MII MAC Interface block. These data nibbles are clocked to the controller at a rate of 2.5 MHz.

12.2.5.4 Jabber Detection

Jabber is a condition in which a station transmits for a period of time longer than the maximum permissible packet length, usually due to a fault condition, which results in holding the internal MII TXEN input for a long period. Special logic is used to detect the jabber state and abort the transmission to the line, within 45 ms. Once TXEN is deasserted, the logic resets the jabber condition.

The [Jabber Detect](#) bit in the [PHY x Basic Status Register \(PHY_BASIC_STATUS_x\)](#) indicates that a jabber condition was detected.

12.2.6 AUTO-NEGOTIATION

The purpose of the Auto-Negotiation function is to automatically configure the transceiver to the optimum link parameters based on the capabilities of its link partner. Auto-Negotiation is a mechanism for exchanging configuration information between two link-partners and automatically selecting the highest performance mode of operation supported by both sides. Auto-Negotiation is fully defined in clause 28 of the IEEE 802.3 specification and is enabled by setting the [Auto-Negotiation Enable \(PHY_AN\)](#) of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#).

Note: Auto-Negotiation is not used for 100BASE-FX mode.

The advertised capabilities of the PHY are stored in the [PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)](#). The PHY contains the ability to advertise 100BASE-TX and 10BASE-T in both full or half-duplex modes. Besides the connection speed, the PHY can advertise remote fault indication and symmetric or asymmetric pause flow control as defined in the IEEE 802.3 specification. The transceiver supports “Next Page” capability which is used to negotiate Energy Efficient Ethernet functionality as well as to support software controlled pages. Many of the default advertised capabilities of the PHY are determined via configuration straps as shown in [Section 12.2.19.5, “PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\),” on page 251](#). Refer to [Section 7.0, “Configuration Straps,” on page 60](#) for additional details on how to use the device configuration straps.

Once Auto-Negotiation has completed, information about the resolved link and the results of the negotiation process are reflected in the [Speed Indication](#) bits in the [PHY x Special Control/Status Register \(PHY_SPECIAL_CONTROL_STATUS_x\)](#), as well as the [PHY x Auto-Negotiation Link Partner Base Page Ability Register \(PHY_AN_LP_BASE_ABILITY_x\)](#). The Auto-Negotiation protocol is a purely physical layer activity and proceeds independently of the MAC controller.

The following blocks are activated during an Auto-Negotiation session:

- Auto-Negotiation (digital)
- 100M ADC (analog)
- 100M PLL (analog)
- 100M equalizer/BLW/clock recovery (DSP)
- 10M SQUELCH (analog)
- 10M PLL (analog)
- 10M Transmitter (analog)

When enabled, Auto-Negotiation is started by the occurrence of any of the following events:

- Power-On Reset (POR)
- Hardware reset (RST#)
- PHY Software reset (via [Reset Control Register \(RESET_CTL\)](#), or bit 15 of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#))
- PHY Power-down reset ([Section 12.2.10, “PHY Power-Down Modes,” on page 231](#))
- PHY Link status down (bit 2 of the [PHY x Basic Status Register \(PHY_BASIC_STATUS_x\)](#) is cleared)
- Setting the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#), bit 9 high (auto-neg restart)
- Digital Reset (via bit 0 of the [Reset Control Register \(RESET_CTL\)](#))
- Issuing an EEPROM Loader RELOAD command ([Section 14.4, “EEPROM Loader,” on page 465](#)) via EEPROM Loader run sequence

Note: Refer to [Section 6.2, “Resets,” on page 42](#) for information on these and other system resets.

On detection of one of these events, the transceiver begins Auto-Negotiation by transmitting bursts of Fast Link Pulses (FLP). These are bursts of link pulses from the 10M TX Driver. They are shaped as Normal Link Pulses and can pass uncorrupted down CAT-3 or CAT-5 cable. A Fast Link Pulse Burst consists of up to 33 pulses. The 17 odd-numbered pulses, which are always present, frame the FLP burst. The 16 even-numbered pulses, which may be present or absent, contain the data word being transmitted. Presence of a data pulse represents a “1”, while absence represents a “0”.

The data transmitted by an FLP burst is known as a “Link Code Word.” These are defined fully in IEEE 802.3 clause 28. In summary, the transceiver advertises 802.3 compliance in its selector field (the first 5 bits of the Link Code Word). It advertises its technology ability according to the bits set in the [PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)](#).

There are 4 possible matches of the technology abilities. In the order of priority these are:

- 100M Full Duplex (Highest priority)
- 100M Half Duplex
- 10M Full Duplex
- 10M Half Duplex (Lowest priority)

If the full capabilities of the transceiver are advertised (100M, full-duplex), and if the link partner is capable of 10M and 100M, then Auto-Negotiation selects 100M as the highest performance mode. If the link partner is capable of half and full-duplex modes, then Auto-Negotiation selects full-duplex as the highest performance mode.

Once a capability match has been determined, the link code words are repeated with the acknowledge bit set. Any difference in the main content of the link code words at this time will cause Auto-Negotiation to re-start. Auto-Negotiation will also re-start if not all of the required FLP bursts are received.

Writing the [PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)](#) bits [8:5] allows software control of the capabilities advertised by the transceiver. Writing the [PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)](#) does not automatically re-start Auto-Negotiation. The [Restart Auto-Negotiation \(PHY_RST_AN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) must be set before the new abilities will be advertised. Auto-Negotiation can also be disabled via software by clearing the [Auto-Negotiation Enable \(PHY_AN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#).

12.2.6.1 Pause Flow Control

The Switch Fabric MACs are capable of generating and receiving pause flow control frames per the IEEE 802.3 specification. The PHY's advertised pause flow control abilities are set via the [Asymmetric Pause](#) and [Symmetric Pause](#) bits of the [PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)](#). This allows the PHY to advertise its flow control abilities and Auto-Negotiate the flow control settings with its link partner. The default values of these bits are determined via configuration straps as defined in [Section 12.2.19.5, "PHY x Auto-Negotiation Advertisement Register \(PHY_AN_ADV_x\)," on page 251](#).

12.2.6.2 Parallel Detection

If the device is connected to a device lacking the ability to Auto-Negotiate (i.e. no FLPs are detected), it is able to determine the speed of the link based on either 100M MLT-3 symbols or 10M Normal Link Pulses. In this case the link is presumed to be half-duplex per the IEEE 802.3 standard. This ability is known as "Parallel Detection." This feature ensures interoperability with legacy link partners. If a link is formed via parallel detection, then the [Link Partner Auto-Negotiation Able](#) bit of the [PHY x Auto-Negotiation Expansion Register \(PHY_AN_EXP_x\)](#) is cleared to indicate that the link partner is not capable of Auto-Negotiation. If a fault occurs during parallel detection, the [Parallel Detection Fault](#) bit of the [PHY x Auto-Negotiation Expansion Register \(PHY_AN_EXP_x\)](#) is set.

The [PHY x Auto-Negotiation Link Partner Base Page Ability Register \(PHY_AN_LP_BASE_ABILITY_x\)](#) is used to store the Link Partner Ability information, which is coded in the received FLPs. If the link partner is not Auto-Negotiation capable, then this register is updated after completion of parallel detection to reflect the speed capability of the link partner.

12.2.6.3 Restarting Auto-Negotiation

Auto-Negotiation can be re-started at any time by setting the [Restart Auto-Negotiation \(PHY_RST_AN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). Auto-Negotiation will also re-start if the link is broken at any time. A broken link is caused by signal loss. This may occur because of a cable break, or because of an interruption in the signal transmitted by the Link Partner. Auto-Negotiation resumes in an attempt to determine the new link configuration.

If the management entity re-starts Auto-Negotiation by setting the [Restart Auto-Negotiation \(PHY_RST_AN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#), the device will respond by stopping all transmission/receiving operations. Once the internal break_link_time is completed in the Auto-Negotiation state-machine (approximately 1200ms), Auto-Negotiation will re-start. In this case, the link partner will have also dropped the link due to lack of a received signal, so it too will resume Auto-Negotiation.

Auto-Negotiation is also restarted after the EEPROM Loader updates the straps.

12.2.6.4 Disabling Auto-Negotiation

Auto-Negotiation can be disabled by clearing the [Auto-Negotiation Enable \(PHY_AN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). The transceiver will then force its speed of operation to reflect the information in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) ([Speed Select LSB \(PHY_SPEED_SEL_LSB\)](#) and [Duplex Mode \(PHY_DUPLEX\)](#)). These bits are ignored when Auto-Negotiation is enabled.

12.2.6.5 Half Vs. Full-Duplex

Half-duplex operation relies on the CSMA/CD (Carrier Sense Multiple Access / Collision Detect) protocol to handle network traffic and collisions. In this mode, the carrier sense signal, CRS, responds to both transmit and receive activity. If data is received while the transceiver is transmitting, a collision results.

In full-duplex mode, the transceiver is able to transmit and receive data simultaneously. In this mode, CRS responds only to receive activity. The CSMA/CD protocol does not apply and collision detection is disabled.

12.2.7 HP AUTO-MDIX

HP Auto-MDIX facilitates the use of CAT-3 (10 BASE-T) or CAT-5 (100 BASE-T) media UTP interconnect cable without consideration of interface wiring scheme. If a user plugs in either a direct connect LAN cable or a cross-over patch cable, as shown in [Figure 12-4](#), the transceiver is capable of configuring the TXPx/TXNx and RXPx/RXNx twisted pair pins for correct transceiver operation.

Note: Auto-MDIX is not used for 100BASE-FX mode.

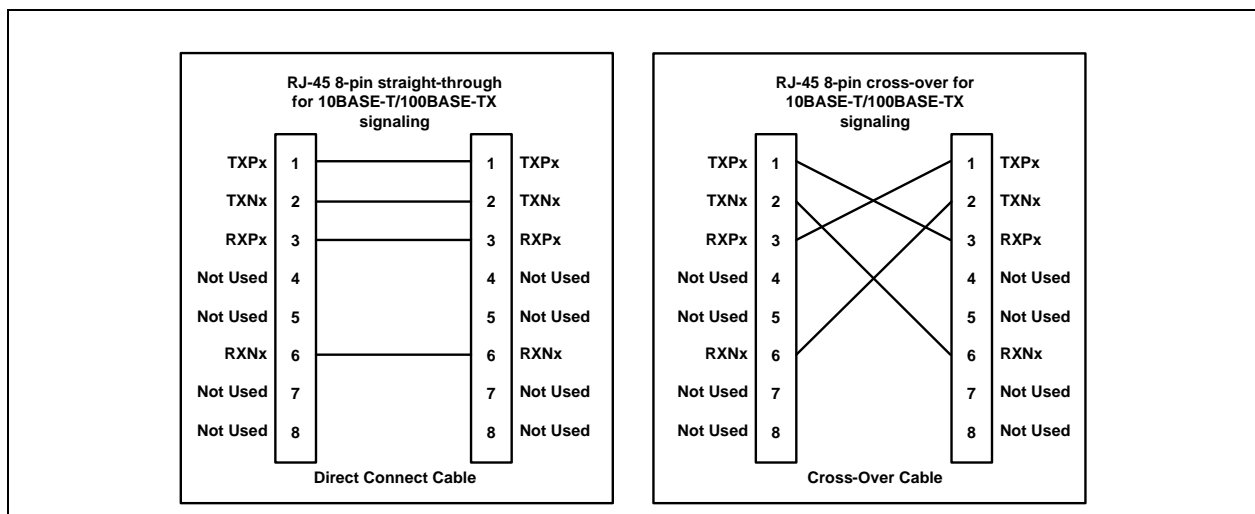
The internal logic of the device detects the TX and RX pins of the connecting device. Since the RX and TX line pairs are interchangeable, special PCB design considerations are needed to accommodate the symmetrical magnetics and termination of an Auto-MDIX design.

The Auto-MDIX function is enabled using the [auto_mdix_strap_1](#) and [auto_mdix_strap_2](#) configuration straps. Manual selection of the cross-over can be set using the [manual_mdix_strap_1](#) and [manual_mdix_strap_2](#) configuration straps. Software based control of the Auto-MDIX function may be performed using the Auto-MDIX Control (AMDIXCTRL) bit of the [PHY x Special Control/Status Indication Register \(PHY_SPECIAL_CONTROL_STAT_IND_x\)](#). When AMDIXCTRL is set to 1, the Auto-MDIX capability is determined by the Auto-MDIX Enable (AMDIXEN) and Auto-MDIX State (AMDIXSTATE) bits of the [PHY x Special Control/Status Indication Register \(PHY_SPECIAL_CONTROL_STAT_IND_x\)](#).

Note: When operating in 10BASE-T or 100BASE-TX manual modes, the Auto-MDIX crossover time can be extended via the [Extend Manual 10/100 Auto-MDIX Crossover Time](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#). Refer to [Section 12.2.19.12, on page 260](#) for additional information.

When Energy Detect Power-Down is enabled, the Auto-MDIX crossover time can be extended via the [EDPD Extend Crossover](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#). Refer to [Section 12.2.19.12, on page 260](#) for additional information

FIGURE 12-4: DIRECT CABLE CONNECTION VS. CROSS-OVER CABLE CONNECTION



12.2.8 PHY MANAGEMENT CONTROL

The PHY Management Control block is responsible for the management functions of the PHY, including register access and interrupt generation. A Serial Management Interface (SMI) is used to support registers as required by the IEEE 802.3 (Clause 22), as well as the vendor specific registers allowed by the specification. The SMI interface consists of the MII Management Data (MDIO) signal and the MII Management Clock (MDC) signal. These signals allow access to all PHY registers. Refer to [Section 12.2.19, "Physical PHY Registers," on page 242](#) for a list of all supported registers and register descriptions. Non-supported registers will be read as FFFFh.

12.2.9 PHY INTERRUPTS

The PHY contains the ability to generate various interrupt events. Reading the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) shows the source of the interrupt. The [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) enables or disables each PHY interrupt.

The PHY Management Control block aggregates the enabled interrupts status into an internal signal which is sent to the System Interrupt Controller and is reflected via the [Physical PHY A Interrupt Event \(PHY_INT_A\)](#) and [Physical PHY B Interrupt Event \(PHY_INT_B\)](#) bits of the [Interrupt Status Register \(INT_STS\)](#). For more information on the device interrupts, refer to [Section 8.0, "System Interrupts," on page 73](#).

The PHY interrupt system provides two modes, a Primary interrupt mode and an Alternative interrupt mode. Both modes will assert the internal interrupt signal sent to the System Interrupt Controller when the corresponding mask bit is set. These modes differ only in how they de-assert the internal interrupt signal. These modes are detailed in the following subsections.

Note: The Primary interrupt mode is the default interrupt mode after a power-up or hard reset. The Alternative interrupt mode requires setup after a power-up or hard reset.

12.2.9.1 Primary Interrupt Mode

The Primary interrupt mode is the default interrupt mode. The Primary interrupt mode is always selected after power-up or hard reset. In this mode, to enable an interrupt, set the corresponding mask bit in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) (see [Table 12-3](#)). When the event to assert an interrupt is true, the internal interrupt signal will be asserted. When the corresponding event to de-assert the interrupt is true, the internal interrupt signal will be de-asserted.

TABLE 12-3: INTERRUPT MANAGEMENT TABLE

Mask	Interrupt Source Flag		Interrupt Source		Event to Assert interrupt	Event to De-assert interrupt
30.9	29.9	Link Up	LINKSTAT See Note 1	Link Status	Rising LINK-STAT	Falling LINKSAT or Reading register 29
30.8	29.8	Wake on LAN	WOL_INT See Note 2	Enabled WOL event	Rising WOL_INT	Falling WOL_INT or Reading register 29
30.7	29.7	ENERGYON	17.1	ENERGYON	Rising 17.1 (Note 3)	Falling 17.1 or Reading register 29
30.6	29.6	Auto-Negotiation complete	1.5	Auto-Negotiate Complete	Rising 1.5	Falling 1.5 or Reading register 29

TABLE 12-3: INTERRUPT MANAGEMENT TABLE (CONTINUED)

30.5	29.5	Remote Fault Detected	1.4	Remote Fault	Rising 1.4	Falling 1.4, or Reading register 1 or Reading register 29
30.4	29.4	Link Down	1.2	Link Status	Falling 1.2	Reading register 1 or Reading register 29
30.3	29.3	Auto-Negotiation LP Acknowledge	5.14	Acknowledge	Rising 5.14	Falling 5.14 or Reading register 29
30.2	29.2	Parallel Detection Fault	6.4	Parallel Detection Fault	Rising 6.4	Falling 6.4 or Reading register 6, or Reading register 29, or Re-Auto Negotiate or Link down
30.1	29.1	Auto-Negotiation Page Received	6.1	Page Received	Rising 6.1	Falling 6.1 or Reading register 6, or Reading register 29, or Re-Auto Negotiate, or Link down.

Note 1: LINKSTAT is the internal link status and is not directly available in any register bit.

Note 2: WOL_INT is defined as bits 7:4 in the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) ANDed with bits 3:0 of the same register, with the resultant 4 bits OR'ed together.

Note 3: If the mask bit is enabled and the internal interrupt signal has been de-asserted while ENERGYON is still high, the internal interrupt signal will assert for 256 ms, approximately one second after ENERGYON goes low when the Cable is unplugged. To prevent an unexpected assertion of the internal interrupt signal, the ENERGYON interrupt mask should always be cleared as part of the ENERGYON interrupt service routine.

Note: The [Energy On \(ENERGYON\)](#) bit in the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#) is defaulted to a '1' at the start of the signal acquisition process, therefore the [INT7](#) bit in the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) will also read as a '1' at power-up. If no signal is present, then both [Energy On \(ENERGYON\)](#) and [INT7](#) will clear within a few milliseconds.

12.2.9.2 Alternate Interrupt Mode

The Alternate interrupt mode is enabled by setting the [ALTINT](#) bit of the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#) to "1". In this mode, to enable an interrupt, set the corresponding bit of the in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) (see [Table 12-4](#)). To clear an interrupt, clear the interrupt source and write a '1' to the corresponding Interrupt Source Flag. Writing a '1' to the Interrupt Source Flag will cause the state machine to check the Interrupt Source to determine if the Interrupt Source Flag should clear or stay as a '1'. If

the condition to de-assert is true, then the Interrupt Source Flag is cleared and the internal interrupt signal is also deasserted. If the condition to de-assert is false, then the Interrupt Source Flag remains set, and the internal interrupt signal remains asserted.

TABLE 12-4: ALTERNATIVE INTERRUPT MODE MANAGEMENT TABLE

Mask	Interrupt Source Flag		Interrupt Source		Event to Assert interrupt	Condition to De-assert	Bit to Clear interrupt
30.9	29.9	Link Up	LINKSTAT See Note 4	Link Status	Rising LINK-STAT	LINKSTAT low	29.9
30.8	29.8	Wake on LAN	WOL_INT See Note 5	Enabled WOL event	Rising WOL_INT	WOL_INT low	29.8
30.7	29.7	ENERGYON	17.1	ENERGYON	Rising 17.1	17.1 low	29.7
30.6	29.6	Auto-Negotiation complete	1.5	Auto-Negotiate Complete	Rising 1.5	1.5 low	29.6
30.5	29.5	Remote Fault Detected	1.4	Remote Fault	Rising 1.4	1.4 low	29.5
30.4	29.4	Link Down	1.2	Link Status	Falling 1.2	1.2 high	29.4
30.3	29.3	Auto-Negotiation LP Acknowledge	5.14	Acknowledge	Rising 5.14	5.14 low	29.3
30.2	29.2	Parallel Detection Fault	6.4	Parallel Detection Fault	Rising 6.4	6.4 low	29.2
30.1	29.1	Auto-Negotiation Page Received	6.1	Page Received	Rising 6.1	6.1 low	29.1

Note 4: LINKSTAT is the internal link status and is not directly available in any register bit.

Note 5: WOL_INT is defined as bits 7:4 in the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) ANDed with bits 3:0 of the same register, with the resultant 4 bits OR'ed together.

Note: The [Energy On \(ENERGYON\)](#) bit in the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#) is defaulted to a '1' at the start of the signal acquisition process, therefore the [INT7](#) bit in the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) will also read as a '1' at power-up. If no signal is present, then both [Energy On \(ENERGYON\)](#) and [INT7](#) will clear within a few milliseconds.

12.2.10 PHY POWER-DOWN MODES

There are two PHY power-down modes: General Power-Down Mode and Energy Detect Power-Down Mode. These modes are described in the following subsections.

Note: For more information on the various power management features of the device, refer to [Section 6.3, "Power Management," on page 49](#).

The power-down modes of each PHY are controlled independently.

The PHY power-down modes do not reload or reset the PHY registers.

12.2.10.1 General Power-Down

This power-down mode is controlled by the [Power Down \(PHY_PWR_DWN\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). In this mode the entire transceiver, except the PHY management control interface, is powered down. The transceiver will remain in this power-down state as long as the [Power Down \(PHY_PWR_DWN\)](#) bit is set. When the [Power Down \(PHY_PWR_DWN\)](#) bit is cleared, the transceiver powers up and is automatically reset.

12.2.10.2 Energy Detect Power-Down

This power-down mode is enabled by setting the [Energy Detect Power-Down \(EDPWRDOWN\)](#) bit of the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#). In this mode, when no energy is present on the line, the entire transceiver is powered down (except for the PHY management control interface, the SQUELCH circuit and the ENERGYON logic). The ENERGYON logic is used to detect the presence of valid energy from 100BASE-TX, 10BASE-T, or Auto-Negotiation signals.

In this mode, when the [Energy On \(ENERGYON\)](#) bit in the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#) signal is low, the transceiver is powered down and nothing is transmitted. When energy is received, via link pulses or packets, the [Energy On \(ENERGYON\)](#) bit goes high, and the transceiver powers up. The transceiver automatically resets itself into the state prior to power-down, and asserts the [INT7](#) bit of the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#). The first and possibly second packet to activate ENERGYON may be lost.

When the [Energy Detect Power-Down \(EDPWRDOWN\)](#) bit of the [PHY x Mode Control/Status Register \(PHY_MODE_CONTROL_STATUS_x\)](#) is low, energy detect power-down is disabled.

When in EDPD mode, the device's NLP characteristics may be modified. The device can be configured to transmit NLPs in EDPD via the [EDPD TX NLP Enable](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#). When enabled, the TX NLP time interval is configurable via the [EDPD TX NLP Interval Timer Select](#) field of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#). When in EDPD mode, the device can also be configured to wake on the reception of one or two NLPs. Setting the [EDPD RX Single NLP Wake Enable](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#) will enable the device to wake on reception of a single NLP. If the [EDPD RX Single NLP Wake Enable](#) bit is cleared, the maximum interval for detecting reception of two NLPs to wake from EDPD is configurable via the [EDPD RX NLP Max Interval Detect Select](#) field of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#).

The energy detect power down feature is part of the broader power management features of the device and can be used to trigger the power management event output pin ([PME](#)) or general interrupt request pin ([IRQ](#)). This is accomplished by enabling the energy detect power-down feature of the PHY as described above, and setting the corresponding energy detect enable (bit 14 for PHY A, bit 15 for PHY B) of the [Power Management Control Register \(PMT_CTRL\)](#). Refer to [Power Management](#) for additional information.

12.2.11 ENERGY EFFICIENT ETHERNET

The PHYs support IEEE 802.3az Energy Efficient Ethernet (EEE). The EEE functionality is enabled/disabled via the [PHY Energy Efficient Ethernet Enable \(PHYEEEEEN\)](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#). Energy Efficient Ethernet is enabled or disabled by default via the [EEE_enable_strap_1](#) and [EEE_enable_strap_2](#) configuration straps. In order for EEE to be utilized, the following conditions must be met:

- EEE functionality must be enabled via the [PHY Energy Efficient Ethernet Enable \(PHYEEEEEN\)](#) bit of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#)
- The [100BASE-TX EEE](#) bit of the [MMD PHY x EEE Advertisement Register \(PHY_EEE_ADV_x\)](#) must be set

- The MAC and link-partner must support and be configured for EEE operation
- The device and link-partner must link in 100BASE-TX full-duplex mode

The value of the [PHY Energy Efficient Ethernet Enable \(PHY_EEE_EN\)](#) bit affects the default values of the following register bits:

- 100BASE-TX EEE bit of the MMD [PHY x EEE Capability Register \(PHY_EEE_CAP_x\)](#)
- 100BASE-TX EEE bit of the MMD [PHY x EEE Advertisement Register \(PHY_EEE_ADV_x\)](#)

Note: Energy Efficient Ethernet is not used for 100BASE-FX mode.

12.2.12 WAKE ON LAN (WoL)

The PHY supports layer WoL event detection of Perfect DA, Broadcast, Magic Packet, and Wakeup frames. This is in addition to any WoL functionality provided by the Host MAC.

Each type of supported wake event (Perfect DA, Broadcast, Magic Packet, or Wakeup frames) may be individually enabled via [Perfect DA Wakeup Enable \(PFDA_EN\)](#), [Broadcast Wakeup Enable \(BCST_EN\)](#), [Magic Packet Enable \(MPEN\)](#), and [Wakeup Frame Enable \(WUEN\)](#) bits of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#), respectively. The WoL event is indicated via the [INT8](#) bit of the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#).

The WoL feature is part of the broader power management features of the device and can be used to trigger the power management event output pin (**PME**) or general interrupt request pin (**IRQ**). This is accomplished by enabling the WoL feature of the PHY as described above, and setting the corresponding WoL enable (bit 14 for PHY A, bit 15 for PHY B) of the [Power Management Control Register \(PMT_CTRL\)](#). Refer to [Section 6.3, "Power Management," on page 49](#) for additional information.

The [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) also provides a [WoL Configured](#) bit, which may be set by software after all WoL registers are configured. Because all WoL related registers are not affected by software resets, software can poll the [WoL Configured](#) bit to ensure all WoL registers are fully configured. This allows the software to skip reprogramming of the WoL registers after reboot due to a WoL event.

The following subsections detail each type of WoL event. For additional information on the main system interrupts, refer to [Section 8.0, "System Interrupts," on page 73](#).

12.2.12.1 Perfect DA (Destination Address) Detection

When enabled, the Perfect DA detection mode allows the detection of a frame with the destination address matching the address stored in the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#). The frame must also pass the FCS and packet length check.

As an example, the Host system must perform the following steps to enable the device to detect a Perfect DA WoL event:

1. Set the desired MAC address to cause the wake event in the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#).
2. Set the [Perfect DA Wakeup Enable \(PFDA_EN\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) to enable Perfect DA detection.
3. Set bit 8 (WoL event indicator) in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) to enable WoL events.

When a match is triggered, bit 8 of the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) will be set, and the [Perfect DA Frame Received \(PFDA_FR\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) will be set.

12.2.12.2 Broadcast Detection

When enabled, the Broadcast detection mode allows the detection of a frame with the destination address value of FF FF FF FF FF. The frame must also pass the FCS and packet length check.

As an example, the Host system must perform the following steps to enable the device to detect a Broadcast WoL event:

1. Set the [Broadcast Wakeup Enable \(BCST_EN\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) to enable Broadcast detection.
2. Set bit 8 (WoL event indicator) in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) to enable

WoL events.

When a match is triggered, bit 8 of the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) will be set, and the [Broadcast Frame Received \(BCAST_FR\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) will be set.

12.2.12.3 Magic Packet Detection

When enabled, the Magic Packet detection mode allows the detection of a Magic Packet frame. A Magic Packet is a frame addressed to the device - either a unicast to the programmed address, or a broadcast - which contains the pattern 48'h FF_FF_FF_FF_FF_FF after the destination and source address field, followed by 16 repetitions of the desired MAC address (loaded into the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#)) without any breaks or interruptions. In case of a break in the 16 address repetitions, the logic scans for the 48'h FF_FF_FF_FF_FF_FF pattern again in the incoming frame. The 16 repetitions may be anywhere in the frame but must be preceded by the synchronization stream. The frame must also pass the FCS check and packet length checking.

As an example, if the desired address is 00h 11h 22h 33h 44h 55h, then the logic scans for the following data sequence in an Ethernet frame:

```
Destination Address Source Address .....FF FF FF FF FF FF
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
00 11 22 33 44 55 00 11 22 33 44 55 00 11 22 33 44 55
...FCS
```

As an example, the Host system must perform the following steps to enable the device to detect a Magic Packet WoL event:

Set the desired MAC address to cause the wake event in the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#).

Set the [Magic Packet Enable \(MPEN\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) to enable Magic Packet detection.

Set bit 8 (WoL event indicator) in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) to enable WoL events.

When a match is triggered, bit 8 of the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#) will be set, and the [Magic Packet Received \(MPR\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) will be set.

12.2.12.4 Wakeup Frame Detection

When enabled, the Wakeup Frame detection mode allows the detection of a pre-programmed Wakeup Frame. Wakeup Frame detection provides a way for system designers to detect a customized pattern within a packet via a programmable wake-up frame filter. The filter has a 128-bit byte mask that indicates which bytes of the frame should be compared by the detection logic. A CRC-16 is calculated over these bytes. The result is then compared with the filter's respective CRC-16 to determine if a match exists. When a wake-up pattern is received, the [Remote Wakeup Frame Received \(WUFR\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) is set.

If enabled, the filter can also include a comparison between the frame's destination address and the address specified in the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#). The specified address can be a unicast or a multicast. If address matching is enabled, only the programmed unicast or multicast address will be considered a match. Non-specific multicast addresses and the broadcast address can be separately enabled. The address matching results are logically OR'd (i.e., specific address match result OR any multicast result OR broadcast result).

Whether or not the filter is enabled and whether the destination address is checked is determined by configuring the [PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#). Before enabling the filter, the application program must provide the detection logic with the sample frame and corresponding byte mask. This information is provided by writing the [PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#), [PHY x Wakeup Filter Configuration](#)

Register B (PHY_WUF_CFGB_x), and PHY x Wakeup Filter Byte Mask Registers (PHY_WUF_MASK_x). The starting offset within the frame and the expected CRC-16 for the filter is determined by the Filter Pattern Offset and Filter CRC-16 fields, respectively.

If remote wakeup mode is enabled, the remote wakeup function checks each frame against the filter and recognizes the frame as a remote wakeup frame if it passes the filter's address filtering and CRC value match.

The pattern offset defines the location of the first byte that should be checked in the frame. The byte mask is a 128-bit field that specifies whether or not each of the 128 contiguous bytes within the frame, beginning with the pattern offset, should be checked. If bit j in the byte mask is set, the detection logic checks the byte (pattern offset + j) in the frame, otherwise byte (pattern offset + j) is ignored.

At the completion of the CRC-16 checking process, the CRC-16 calculated using the pattern offset and byte mask is compared to the expected CRC-16 value associated with the filter. If a match occurs, a remote wake-up event is signaled. The frame must also pass the FCS check and packet length checking.

Table 12-5 indicates the cases that produce a wake-up event. All other cases do not generate a wake-up event.

TABLE 12-5: WAKEUP GENERATION CASES

Filter Enabled	Frame Type	CRC Matches	Address Match Enabled	Any Mcast Enabled	Bcast Enabled	Frame Address Matches
Yes	Unicast	Yes	No	X	X	X
Yes	Unicast	Yes	Yes	X	X	Yes
Yes	Multicast	Yes	X	Yes	X	X
Yes	Multicast	Yes	Yes	No	X	Yes
Yes	Broadcast	Yes	X	X	Yes	X

As an example, the Host system must perform the following steps to enable the device to detect a Wakeup Frame WoL event:

Declare Pattern:

1. Update the PHY x Wakeup Filter Byte Mask Registers (PHY_WUF_MASK_x) to indicate the valid bytes to match.
2. Calculate the CRC-16 value of valid bytes offline and update the PHY x Wakeup Filter Configuration Register B (PHY_WUF_CFGB_x). CRC-16 is calculated as follows:

At the start of a frame, CRC-16 is initialized with the value FFFFh. CRC-16 is updated when the pattern offset and mask indicate the received byte is part of the checksum calculation. The following algorithm is used to update the CRC-16 at that time:

Let:

^ denote the exclusive or operator.

Data [7:0] be the received data byte to be included in the checksum.

CRC[15:0] contain the calculated CRC-16 checksum.

F0 ... F7 be intermediate results, calculated when a data byte is determined to be part of the CRC-16.

Calculate:

$$F0 = CRC[15] \wedge Data[0]$$

$$F1 = CRC[14] \wedge F0 \wedge Data[1]$$

$$F2 = CRC[13] \wedge F1 \wedge Data[2]$$

$$F3 = CRC[12] \wedge F2 \wedge Data[3]$$

$$F4 = CRC[11] \wedge F3 \wedge Data[4]$$

$$F5 = CRC[10] \wedge F4 \wedge Data[5]$$

$$F6 = \text{CRC}[09] \wedge F5 \wedge \text{Data}[6]$$

$$F7 = \text{CRC}[08] \wedge F6 \wedge \text{Data}[7]$$

The CRC-32 is updated as follows:

$$\text{CRC}[15] = \text{CRC}[7] \wedge F7$$

$$\text{CRC}[14] = \text{CRC}[6]$$

$$\text{CRC}[13] = \text{CRC}[5]$$

$$\text{CRC}[12] = \text{CRC}[4]$$

$$\text{CRC}[11] = \text{CRC}[3]$$

$$\text{CRC}[10] = \text{CRC}[2]$$

$$\text{CRC}[9] = \text{CRC}[1] \wedge F0$$

$$\text{CRC}[8] = \text{CRC}[0] \wedge F1$$

$$\text{CRC}[7] = F0 \wedge F2$$

$$\text{CRC}[6] = F1 \wedge F3$$

$$\text{CRC}[5] = F2 \wedge F4$$

$$\text{CRC}[4] = F3 \wedge F5$$

$$\text{CRC}[3] = F4 \wedge F6$$

$$\text{CRC}[2] = F5 \wedge F7$$

$$\text{CRC}[1] = F6$$

$$\text{CRC}[0] = F7$$

3. Determine the offset pattern with offset 0 being the first byte of the destination address. Update the offset in the [Filter Pattern Offset field of the PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#).

Determine Address Matching Conditions:

4. Determine the address matching scheme based on [Table 12-5](#) and update the [Filter Broadcast Enable](#), [Filter Any Multicast Enable](#), and [Address Match Enable](#) bits of the [PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#) accordingly.
5. If necessary (see step 4), set the desired MAC address to cause the wake event in the [PHY x MAC Receive Address A Register \(PHY_RX_ADDRA_x\)](#), [PHY x MAC Receive Address B Register \(PHY_RX_ADDRB_x\)](#), and [PHY x MAC Receive Address C Register \(PHY_RX_ADDR_C_x\)](#).
6. Set the [Filter Enable](#) bit of the [PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#) to enable the filter.

Enable Wakeup Frame Detection:

7. Set the [Wakeup Frame Enable \(WUEN\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) to enable Wakeup Frame detection.
8. Set bit 8 (WoL event indicator) in the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#) to enable WoL events.

When a match is triggered, the [Remote Wakeup Frame Received \(WUFR\)](#) bit of the [PHY x Wakeup Control and Status Register \(PHY_WUCSR_x\)](#) will be set. To provide additional visibility to software, the [Filter Triggered](#) bit of the [PHY x Wakeup Filter Configuration Register A \(PHY_WUF_CFGA_x\)](#) will be set.

12.2.13 RESETS

In addition to the chip-level hardware reset (**RST#**) and Power-On Reset (POR), the PHY supports three block specific resets. These are discussed in the following sections. For detailed information on all device resets and the reset sequence refer to [Section 6.2, "Resets," on page 42](#).

Note: Only a hardware reset (**RST#**) or Power-On Reset (POR) will automatically reload the configuration strap values into the PHY registers.

The [Digital Reset \(DIGITAL_RST\)](#) bit in the [Reset Control Register \(RESET_CTL\)](#) does not reset the PHYs. The [Digital Reset \(DIGITAL_RST\)](#) bit will cause the EEPROM Loader to reload the configuration strap values into the PHY registers and to reset all other PHY registers to their default values. An EEPROM RELOAD command via the [EEPROM Command Register \(E2P_CMD\)](#) also has the same effect.

For all other PHY resets, PHY registers will need to be manually configured via software.

12.2.13.1 PHY Software Reset via RESET_CTL

The PHYs can be reset via the [Reset Control Register \(RESET_CTL\)](#). These bits are self clearing after approximately 102 us. This reset does not reload the configuration strap values into the PHY registers.

12.2.13.2 PHY Software Reset via PHY_BASIC_CTRL_x

The PHY can also be reset by setting the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). This bit is self clearing and will return to 0 after the reset is complete. This reset does not reload the configuration strap values into the PHY registers.

12.2.13.3 PHY Power-Down Reset

After the PHY has returned from a power-down state, a reset of the PHY is automatically generated. The PHY power-down modes do not reload or reset the PHY registers. Refer to [Section 12.2.10, "PHY Power-Down Modes," on page 231](#) for additional information.

12.2.14 LINK INTEGRITY TEST

The device performs the link integrity test as outlined in the IEEE 802.3u (clause 24-15) Link Monitor state diagram. The link status is multiplexed with the 10 Mbps link status to form the [Link Status](#) bit in the [PHY x Basic Status Register \(PHY_BASIC_STATUS_x\)](#) and to drive the LINK LED functions.

The DSP indicates a valid MLT-3 waveform present on the RXPx and RXNx signals as defined by the ANSI X3.263 TP-PMD standard, to the Link Monitor state-machine, using the internal DATA_VALID signal. When DATA_VALID is asserted, the control logic moves into a Link-Ready state and waits for an enable from the auto-negotiation block. When received, the Link-Up state is entered, and the Transmit and Receive logic blocks become active. Should auto-negotiation be disabled, the link integrity logic moves immediately to the Link-Up state when the DATA_VALID is asserted.

To allow the line to stabilize, the link integrity logic will wait a minimum of 330 ms from the time DATA_VALID is asserted until the Link-Ready state is entered. Should the DATA_VALID input be negated at any time, this logic will immediately negate the Link signal and enter the Link-Down state.

When the 10/100 digital block is in 10BASE-T mode, the link status is derived from the 10BASE-T receiver logic.

12.2.15 CABLE DIAGNOSTICS

The PHYs provide cable diagnostics which allow for open/short and length detection of the Ethernet cable. The cable diagnostics consist of two primary modes of operation:

- [Time Domain Reflectometry \(TDR\) Cable Diagnostics](#)
TDR cable diagnostics enable the detection of open or shorted cabling on the TX or RX pair, as well as cable length estimation to the open/short fault.
- [Matched Cable Diagnostics](#)
Matched cable diagnostics enable cable length estimation on 100 Mbps-linked cables.


Refer to the following sub-sections for details on proper operation of each cable diagnostics mode.

Note: Cable diagnostics are not used for 100BASE-FX mode.

12.2.15.1 Time Domain Reflectometry (TDR) Cable Diagnostics

The PHYs provide TDR cable diagnostics which enable the detection of open or shorted cabling on the TX or RX pair, as well as cable length estimation to the open/short fault. To utilize the TDR cable diagnostics, Auto-MDIX and Auto Negotiation must be disabled, and the PHY must be forced to 100 Mbps full-duplex mode. These actions must be performed before setting the **TDR Enable** bit in the **PHY x TDR Control/Status Register (PHY_TDR_CONTROL_STAT_x)**. With Auto-MDIX disabled, the TDR will test the TX or RX pair selected by register bit 27.13 (**Auto-MDIX State (AMDIX-STATE)**). Proper cable testing should include a test of each pair. TDR cable diagnostics is not appropriate for 100BASE-FX mode. When TDR testing is complete, prior register settings may be restored. [Figure 12-5](#) provides a flow diagram of proper TDR usage.

FIGURE 12-5: TDR USAGE FLOW DIAGRAM



LAN9352

The TDR operates by transmitting pulses on the selected twisted pair within the Ethernet cable (TX in MDI mode, RX in MDIX mode). If the pair being tested is open or shorted, the resulting impedance discontinuity results in a reflected signal that can be detected by the PHY. The PHY measures the time between the transmitted signal and received reflection and indicates the results in the **TDR Channel Length** field of the **PHY x TDR Control/Status Register (PHY_TDR_CONTROL_STAT_x)**. The **TDR Channel Length** field indicates the “electrical” length of the cable, and can be multiplied by the appropriate propagation constant in [Table 12-6](#) to determine the approximate physical distance to the fault.

Note: The TDR function is typically used when the link is inoperable. However, an active link will drop when operating the TDR.

Since the TDR relies on the reflected signal of an improperly terminated cable, there are several factors that can affect the accuracy of the physical length estimate. These include:

1. **Cable Type (CAT 5, CAT5e, CAT6):** The electrical length of each cable type is slightly different due to the twists-per-meter of the internal signal pairs and differences in signal propagation speeds. If the cable type is known, the length estimate can be calculated more accurately by using the propagation constant appropriate for the cable type (see [Table 12-6](#)). In many real-world applications the cable type is unknown, or may be a mix of different cable types and lengths. In this case, use the propagation constant for the “unknown” cable type.
2. **TX and RX Pair:** For each cable type, the EIA standards specify different twist rates (twists-per-meter) for each signal pair within the Ethernet cable. This results in different measurements for the RX and TX pair.
3. **Actual Cable Length:** The difference between the estimated cable length and actual cable length grows as the physical cable length increases, with the most accurate results at less than approximately 100 m.
4. **Open/Short Case:** The Open and Shorted cases will return different TDR Channel Length values (electrical lengths) for the same physical distance to the fault. Compensation for this is achieved by using different propagation constants to calculate the physical length of the cable.

For the Open case, the estimated distance to the fault can be calculated as follows:

Distance to Open fault in meters \cong **TDR Channel Length** * P_{OPEN}
Where: P_{OPEN} is the propagation constant selected from [Table 12-6](#)

For the Shorted case, the estimated distance to the fault can be calculated as follows:

Distance to Open fault in meters \cong **TDR Channel Length** * P_{SHORT}
Where: P_{SHORT} is the propagation constant selected from [Table 12-6](#)

TABLE 12-6: TDR PROPAGATION CONSTANTS

TDR Propagation Constant	Cable Type			
	Unknown	CAT 6	CAT 5E	CAT 5
P _{OPEN}	0.769	0.745	0.76	0.85
P _{SHORT}	0.793	0.759	0.788	0.873

The typical cable length measurement margin of error for Open and Shorted cases is dependent on the selected cable type and the distance of the open/short from the device. [Table 12-7](#) and [Table 12-8](#) detail the typical measurement error for Open and Shorted cases, respectively.

TABLE 12-7: TYPICAL MEASUREMENT ERROR FOR OPEN CABLE (+/- METERS)

Physical Distance to Fault	Selected Propagation Constant			
	P _{OPEN} = Unknown	P _{OPEN} = CAT 6	P _{OPEN} = CAT 5E	P _{OPEN} = CAT 5
CAT 6 Cable, 0-100 m	9	6		
CAT 5E Cable, 0-100 m	5		5	

TABLE 12-7: TYPICAL MEASUREMENT ERROR FOR OPEN CABLE (+/- METERS)

CAT 5 Cable, 0-100 m	13			3
CAT 6 Cable, 101-160 m	14	6		
CAT 5E Cable, 101-160 m	8		6	
CAT 5 Cable, 101-160 m	20			6

TABLE 12-8: TYPICAL MEASUREMENT ERROR FOR SHORTED CABLE (+/- METERS)

PHYSICAL DISTANCE TO FAULT	SELECTED PROPAGATION CONSTANT			
	P _{SHORT} = Unknown	P _{SHORT} = CAT 6	P _{SHORT} = CAT 5E	P _{SHORT} = CAT 5
CAT 6 Cable, 0-100 m	8	5		
CAT 5E Cable, 0-100 m	5		5	
CAT 5 Cable, 0-100 m	11			2
CAT 6 Cable, 101-160 m	14	6		
CAT 5E Cable, 101-160 m	7		6	
CAT 5 Cable, 101-160 m	11			3

12.2.15.2 Matched Cable Diagnostics

Matched cable diagnostics enable cable length estimation on 100 Mbps-linked cables of up to 120 meters. If there is an active 100 Mb link, the approximate distance to the link partner can be estimated using the [PHY x Cable Length Register \(PHY_CABLE_LEN_x\)](#). If the cable is properly terminated, but there is no active 100 Mb link (the link partner is disabled, nonfunctional, the link is at 10 Mb, etc.), the cable length cannot be estimated and the [PHY x Cable Length Register \(PHY_CABLE_LEN_x\)](#) should be ignored. The estimated distance to the link partner can be determined via the [Cable Length \(CBLN\)](#) field of the [PHY x Cable Length Register \(PHY_CABLE_LEN_x\)](#) using the lookup table provided in [Table 12-9](#). The typical cable length measurement margin of error for a matched cable case is +/- 20 m. The matched cable length margin of error is consistent for all cable types from 0 to 120 m.

TABLE 12-9: MATCH CASE ESTIMATED CABLE LENGTH (CBLN) LOOKUP

CBLN Field Value	Estimated Cable Length
0 - 3	0
4	6
5	17
6	27
7	38
8	49
9	59
10	70
11	81

TABLE 12-9: MATCH CASE ESTIMATED CABLE LENGTH (CBLN) LOOKUP

12	91
13	102
14	113
15	123

Note: For a properly terminated cable (Match case), there is no reflected signal. In this case, the [TDR Channel Length](#) field is invalid and should be ignored.


12.2.16 LOOPBACK OPERATION

The PHYs may be configured for near-end loopback and connector loopback. These loopback modes are detailed in the following subsections.

12.2.16.1 Near-end Loopback

Near-end loopback mode sends the digital transmit data back out the receive data signals for testing purposes, as indicated by the blue arrows in [Figure 12-6](#). The near-end loopback mode is enabled by setting the [Loopback \(PHY_LOOPBACK\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) to “1”. A large percentage of the digital circuitry is operational in near-end loopback mode because data is routed through the PCS and PMA layers into the PMD sublayer before it is looped back. The COL signal will be inactive in this mode, unless [Collision Test Mode \(PHY_COL_TEST\)](#) is enabled in the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#). The transmitters are powered down regardless of the state of the internal MII TXEN signal.


FIGURE 12-6: NEAR-END LOOPBACK BLOCK DIAGRAM



12.2.16.2 Connector Loopback

The device maintains reliable transmission over very short cables and can be tested in a connector loopback as shown in Figure 12-7. An RJ45 loopback cable can be used to route the transmit signals from the output of the transformer back to the receiver inputs. The loopback works at both 10 and 100 Mbps.

FIGURE 12-7: CONNECTION LOOPBACK BLOCK DIAGRAM



12.2.17 100BASE-FX OPERATION

When set for 100BASE-FX operation, the scrambler and MTL-3 blocks are disabled and the analog RX and TX pins are changed to differential LVPECL pins and connect through external terminations to the external Fiber transceiver. The differential LVPECL pins support a signal voltage range compatible with SFF (LVPECL) and SFP (reduced LVPECL) type transceivers.

While in 100BASE-FX operation, the quality of the receive signal is provided by the external transceiver as either an open-drain, CMOS level, Loss of Signal (SFP) or a LVPECL Signal Detect (SFF).

12.2.17.1 100BASE-FX Far End Fault Indication

Since Auto-Negotiation is not specified for 100BASE-FX, its Remote Fault capability is unavailable. Instead, 100BASE-FX provides an optional Far-End Fault function.

When no signal is being received, the Far-End Fault feature transmits a special Far-End Fault Indication to its far-end peer. The Far-End Fault Indication is sent only when a physical error condition is sensed on the receive channel.

The Far-End Fault Indication is comprised of three or more repeating cycles, each of 84 ONEs followed by a single ZERO. This signal is sent in-band and is readily detectable but is constructed so as to not satisfy the 100BASE-X carrier sense criterion.

Far-End Fault is implemented through the Far-End Fault Generate, Far-End Fault Detect, and the Link Monitor processes. The Far-End Fault Generate process is responsible for sensing a receive channel failure (signal_status=OFF) and transmitting the Far-End Fault Indication in response. The transmission of the Far-End Fault Indication may start or stop at any time depending only on signal_status. The Far-End Fault Detect process continuously monitors the RX process for the Far-End Fault Indication. Detection of the Far-End Fault Indication disables the station by causing the Link Monitor process to de-assert link_status, which in turn causes the station to source IDLEs.

Far-End Fault is enabled by default while in 100BASE-FX mode via the [Far End Fault Indication Enable \(FEFI_EN\)](#) of the [PHY x Special Control/Status Indication Register \(PHY_SPECIAL_CONTROL_STAT_IND_x\)](#).

12.2.17.2 100BASE-FX Enable and LOS/SD Selection

100BASE-FX operation is enabled by the use of the FX mode straps ([fx_mode_strap_1](#) and [fx_mode_strap_2](#)) and is reflected in the [100BASE-FX Mode \(FX_MODE\)](#) bit in the [PHY x Special Modes Register \(PHY_SPECIAL_MODES_x\)](#).

Loss of Signal mode is selected for both PHYs by the three level [FXLOSEN](#) strap input pin. The three levels correspond to Loss of Signal mode for a) neither PHY (less than 1 V (typ.)), b) PHY A (greater than 1 V (typ.) but less than 2 V (typ.)) or c) both PHYs (greater than 2 V (typ.)). It is not possible to select Loss of Signal mode for only PHY B.

If Loss of Signal mode is not selected, then Signal Detect mode is selected, independently, by the **FXSDENA** or **FXSDENB** strap input pin. When greater than 1 V (typ.), Signal Detect mode is enabled, when less than 1 V (typ.), copper twisted pair is enabled.

Note: The **FXSDENA** strap input pin is shared with the **FXSDA** pin and the **FXSDENB** strap input pin is shared with the **FXSDB** pin. As such, the LVPECL levels ensure that the input is greater than 1 V (typ.) and that Signal Detect mode is selected. When TP copper is desired, the Signal Detect input function is not required and the pin should be set to 0 V.

Care must be taken such that an non-powered or disabled transceiver does not load the Signal Detect input below the valid LVPECL level.

Table 12-10 and Table 12-11 summarize the selections.

TABLE 12-10: 100BASE-FX LOS, SD AND TP COPPER SELECTION PHY A

FXLOSEN	FXSDENA	PHY Mode
<1 V (typ.)	<1 V (typ.)	TP copper
	>1 V (typ.)	100BASE-FX Signal Detect
>1 V (typ.)	n/a	100BASE-FX LOS

TABLE 12-11: 100BASE-FX LOS, SD AND TP COPPER SELECTION PHY B

FXLOSEN	FXSDENB	PHY Mode
<1 V (typ.)	<1 V (typ.)	TP copper
	>1 V (typ.)	100BASE-FX Signal Detect
>2 V (typ.)	n/a	100BASE-FX LOS

12.2.18 REQUIRED ETHERNET MAGNETICS (100BASE-TX AND 10BASE-T)

The magnetics selected for use with the device should be an Auto-MDIX style magnetic, which is widely available from several vendors. Please review the SMSC/Microchip Application note 8.13 “Suggested Magnetics” for the latest qualified and suggested magnetics. A list of vendors and part numbers are provided within the application note.

12.2.19 PHYSICAL PHY REGISTERS

The Physical PHYs A and B are comparable in functionality and have an identical set of non-memory mapped registers. These registers are indirectly accessed through the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#) and [Host MAC MII Data Register \(HMAC_MII_DATA\)](#).

Because Physical PHY A and B registers are functionally identical, their register descriptions have been consolidated. A lowercase “x” has been appended to the end of each PHY register name in this section, where “x” hold be replaced with “A” or “B” for the PHY A or PHY B registers respectively. In some instances, a “1” or a “2” may be appropriate instead.

A list of the MII serial accessible Control and Status registers and their corresponding register index numbers is included in [Table 12-12](#). Each individual PHY is assigned a unique PHY address as detailed in [Section 12.1.1, “PHY Addressing,” on page 218](#).

In addition to the MII serial accessible Control and Status registers, a set of indirectly accessible registers provides support for the *IEEE 802.3 Section 45.2 MDIO Manageable Device (MMD) Registers*. A list of these registers and their corresponding register index numbers is included in [Table 12-18](#).

Note: The [Digital Reset \(DIGITAL_RST\)](#) bit will cause the EEPROM Loader to reload the configuration strap values into the PHY registers and to reset all other PHY registers to their default values. An EEPROM RELOAD command via the [EEPROM Command Register \(E2P_CMD\)](#) also has the same effect.

Control and Status Registers

[Table 12-12](#) provides a list of supported registers. Register details, including bit definitions, are provided in the following subsections.

Unless otherwise specified, reserved fields must be written with zeros if the register is written.

TABLE 12-12: PHYSICAL PHY A AND B MII SERIALLY ACCESSIBLE CONTROL AND STATUS REGISTERS

Index	Register Name (SYMBOL)	Group
0	PHY x Basic Control Register (PHY_BASIC_CONTROL_x)	Basic
1	PHY x Basic Status Register (PHY_BASIC_STATUS_x)	Basic
2	PHY x Identification MSB Register (PHY_ID_MSB_x)	Extended
3	PHY x Identification LSB Register (PHY_ID_LSB_x)	Extended
4	PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x)	Extended
5	PHY x Auto-Negotiation Link Partner Base Page Ability Register (PHY_AN_LP_BASE_ABILITY_x)	Extended
6	PHY x Auto-Negotiation Expansion Register (PHY_AN_EXP_x)	Extended
7	PHY x Auto Negotiation Next Page TX Register (PHY_AN_NP_TX_x)	Extended
8	PHY x Auto Negotiation Next Page RX Register (PHY_AN_NP_RX_x)	Extended
13	PHY x MMD Access Control Register (PHY_MMD_ACCESS)	Extended
14	PHY x MMD Access Address/Data Register (PHY_MMD_ADDR_DATA)	Extended
16	PHY x EDPD NLP / Crossover Time / EEE Configuration Register (PHY_EDPD_CFG_x)	Vendor-specific
17	PHY x Mode Control/Status Register (PHY_MODE_CONTROL_STATUS_x)	Vendor-specific
18	PHY x Special Modes Register (PHY_SPECIAL_MODES_x)	Vendor-specific
24	PHY x TDR Patterns/Delay Control Register (PHY_TDR_PAT_DELAY_x)	Vendor-specific
25	PHY x TDR Control/Status Register (PHY_TDR_CONTROL_STAT_x)	Vendor-specific
26	PHY x Symbol Error Counter Register	Vendor-specific
27	PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x)	Vendor-specific
28	PHY x Cable Length Register (PHY_CABLE_LEN_x)	Vendor-specific
29	PHY x Interrupt Source Flags Register (PHY_INTERRUPT_SOURCE_x)	Vendor-specific

TABLE 12-12: PHYSICAL PHY A AND B MII SERIALY ACCESSIBLE CONTROL AND STATUS REGISTERS (CONTINUED)

Index	Register Name (SYMBOL)	Group
30	PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)	Vendor-specific
31	PHY x Special Control/Status Register (PHY_SPECIAL_CONTROL_STATUS_x)	Vendor-specific

12.2.19.1 PHY x Basic Control Register (PHY_BASIC_CONTROL_x)

Index (decimal): 0

Size: 16 bits

This read/write register is used to configure the PHY.

Bits	Description	Type	Default
15	Soft Reset (PHY_SRST) When set, this bit resets all the PHY registers to their default state, except those marked as NASR type. This bit is self clearing. 0: Normal operation 1: Reset	R/W SC	0b
14	Loopback (PHY_LOOPBACK) This bit enables/disables the loopback mode. When enabled, transmissions are not sent to network. Instead, they are looped back into the PHY. 0: Loopback mode disabled (normal operation) 1: Loopback mode enabled	R/W	0b
13	Speed Select LSB (PHY_SPEED_SEL_LSB) This bit is used to set the speed of the PHY when the Auto-Negotiation Enable (PHY_AN) bit is disabled. 0: 10 Mbps 1: 100 Mbps	R/W	Note 6
12	Auto-Negotiation Enable (PHY_AN) This bit enables/disables Auto-Negotiation. When enabled, the Speed Select LSB (PHY_SPEED_SEL_LSB) and Duplex Mode (PHY_DUPLEX) bits are overridden. This bit is forced to a 0 if the 100BASE-FX Mode (FX_MODE) bit of the PHY x Special Modes Register (PHY_SPECIAL_MODES_x) is a high. 0: Auto-Negotiation disabled 1: Auto-Negotiation enabled	R/W	Note 7
11	Power Down (PHY_PWR_DWN) This bit controls the power down mode of the PHY. 0: Normal operation 1: General power down mode	R/W	0b
10	RESERVED	RO	-
9	Restart Auto-Negotiation (PHY_RST_AN) When set, this bit restarts the Auto-Negotiation process. 0: Normal operation 1: Auto-Negotiation restarted	R/W SC	0b

Bits	Description	Type	Default
8	Duplex Mode (PHY_DUPLEX) This bit is used to set the duplex when the Auto-Negotiation Enable (PHY_AN) bit is disabled. 0: Half Duplex 1: Full Duplex	R/W	Note 8
7	Collision Test Mode (PHY_COL_TEST) This bit enables/disables the collision test mode of the PHY. When set, the collision signal is active during transmission. It is recommended that this feature be used only in loopback mode. 0: Collision test mode disabled 1: Collision test mode enabled	R/W	0b
6:0	RESERVED	RO	-

Note 6: The default value of this bit is determined by the logical OR of the Auto-Negotiation strap ([autoneg_strap_1](#) for PHY A, [autoneg_strap_2](#) for PHY B) and the speed select strap ([speed_strap_1](#) for PHY A, [speed_strap_2](#) for PHY B). Essentially, if the Auto-Negotiation strap is set, the default value is 1, otherwise the default is determined by the value of the speed select strap. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information. In 100BASE-FX mode, the default value of this bit is a 1.

Note 7: The default is the value of the Auto-Negotiation strap ([autoneg_strap_1](#) for PHY A, [autoneg_strap_2](#) for PHY B). Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information. In 100BASE-FX mode, the default value of this bit is a 0.

Note 8: The default value of this bit is determined by the logical AND of the negation of the Auto-Negotiation strap ([autoneg_strap_1](#) for PHY A, [autoneg_strap_2](#) for PHY B) and the duplex select strap ([duplex_strap_1](#) for PHY A, [duplex_strap_2](#) for PHY B). Essentially, if the Auto-Negotiation strap is set, the default value is 0, otherwise the default is determined by the value of the duplex select strap. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information.

In 100BASE-FX mode, the Auto-Negotiation strap is not considered and the default of this bit is the value of the duplex select strap.

12.2.19.2 PHY x Basic Status Register (PHY_BASIC_STATUS_x)

Index (decimal): 1

Size: 16 bits

This register is used to monitor the status of the PHY.

Bits	Description	Type	Default
15	100BASE-T4 This bit displays the status of 100BASE-T4 compatibility. 0: PHY not able to perform 100BASE-T4 1: PHY able to perform 100BASE-T4	RO	0b
14	100BASE-X Full Duplex This bit displays the status of 100BASE-X full duplex compatibility. 0: PHY not able to perform 100BASE-X full duplex 1: PHY able to perform 100BASE-X full duplex	RO	1b
13	100BASE-X Half Duplex This bit displays the status of 100BASE-X half duplex compatibility. 0: PHY not able to perform 100BASE-X half duplex 1: PHY able to perform 100BASE-X half duplex	RO	1b
12	10BASE-T Full Duplex This bit displays the status of 10BASE-T full duplex compatibility. 0: PHY not able to perform 10BASE-T full duplex 1: PHY able to perform 10BASE-T full duplex	RO	1b
11	10BASE-T Half Duplex (typ.) This bit displays the status of 10BASE-T half duplex compatibility. 0: PHY not able to perform 10BASE-T half duplex 1: PHY able to perform 10BASE-T half duplex	RO	1b
10	100BASE-T2 Full Duplex This bit displays the status of 100BASE-T2 full duplex compatibility. 0: PHY not able to perform 100BASE-T2 full duplex 1: PHY able to perform 100BASE-T2 full duplex	RO	0b
9	100BASE-T2 Half Duplex This bit displays the status of 100BASE-T2 half duplex compatibility. 0: PHY not able to perform 100BASE-T2 half duplex 1: PHY able to perform 100BASE-T2 half duplex	RO	0b
8	Extended Status This bit displays whether extended status information is in register 15 (per IEEE 802.3 clause 22.2.4). 0: No extended status information in Register 15 1: Extended status information in Register 15	RO	0b

Bits	Description	Type	Default
7	Unidirectional Ability This bit indicates whether the PHY is able to transmit regardless of whether the PHY has determined that a valid link has been established. 0: Can only transmit when a valid link has been established 1: Can transmit regardless	RO	0b
6	MF Preamble Suppression This bit indicates whether the PHY accepts management frames with the preamble suppressed. 0: Management frames with preamble suppressed not accepted 1: Management frames with preamble suppressed accepted	RO	0b
5	Auto-Negotiation Complete This bit indicates the status of the Auto-Negotiation process. 0: Auto-Negotiation process not completed 1: Auto-Negotiation process completed	RO	0b
4	Remote Fault This bit indicates if a remote fault condition has been detected. 0: No remote fault condition detected 1: Remote fault condition detected	RO/LH	0b
3	Auto-Negotiation Ability This bit indicates the PHY's Auto-Negotiation ability. 0: PHY is unable to perform Auto-Negotiation 1: PHY is able to perform Auto-Negotiation	RO	1b
2	Link Status This bit indicates the status of the link. 0: Link is down 1: Link is up	RO/LL	0b
1	Jabber Detect This bit indicates the status of the jabber condition. 0: No jabber condition detected 1: Jabber condition detected	RO/LH	0b
0	Extended Capability This bit indicates whether extended register capability is supported. 0: Basic register set capabilities only 1: Extended register set capabilities	RO	1b

12.2.19.3 PHY x Identification MSB Register (PHY_ID_MSB_x)

Index (decimal): 2

Size: 16 bits

This read/write register contains the MSB of the Organizationally Unique Identifier (OUI) for the PHY. The LSB of the PHY OUI is contained in the [PHY x Identification LSB Register \(PHY_ID_LSB_x\)](#).

Bits	Description	Type	Default
15:0	PHY ID This field is assigned to the 3rd through 18th bits of the OUI, respectively (OUI = 00800Fh).	R/W	0007h

12.2.19.4 PHY x Identification LSB Register (PHY_ID_LSB_x)

Index (decimal): 3

Size: 16 bits

This read/write register contains the LSB of the Organizationally Unique Identifier (OUI) for the PHY. The MSB of the PHY OUI is contained in the [PHY x Identification MSB Register \(PHY_ID_MSB_x\)](#).

Bits	Description	Type	Default
15:10	PHY ID This field is assigned to the 19th through 24th bits of the PHY OUI, respectively. (OUI = 00800Fh).	R/W	C140h
9:4	Model Number This field contains the 6-bit manufacturer's model number of the PHY.	R/W	
3:0	Revision Number This field contain the 4-bit manufacturer's revision number of the PHY.	R/W	

Note: The default value of the Revision Number field may vary dependent on the silicon revision number.

12.2.19.5 PHY x Auto-Negotiation Advertisement Register (PHY_AN_ADV_x)

Index (decimal): 4

Size: 16 bits

This read/write register contains the advertised ability of the PHY and is used in the Auto-Negotiation process with the link partner.

Bits	Description	Type	Default
15	Next Page 0 = No next page ability 1 = Next page capable	R/W	0b
14	RESERVED	RO	-
13	Remote Fault This bit determines if remote fault indication will be advertised to the link partner. 0: Remote fault indication not advertised 1: Remote fault indication advertised	R/W	0b
12	Extended Next Page Note: This bit should be written as 0.	R/W	0b
11	Asymmetric Pause This bit determines the advertised asymmetric pause capability. 0: No Asymmetric PAUSE toward link partner advertised 1: Asymmetric PAUSE toward link partner advertised	R/W	Note 9
10	Symmetric Pause This bit determines the advertised symmetric pause capability. 0: No Symmetric PAUSE toward link partner advertised 1: Symmetric PAUSE toward link partner advertised	R/W	Note 9
9	RESERVED	RO	-
8	100BASE-X Full Duplex This bit determines the advertised 100BASE-X full duplex capability. 0: 100BASE-X full duplex ability not advertised 1: 100BASE-X full duplex ability advertised	R/W	1b
7	100BASE-X Half Duplex This bit determines the advertised 100BASE-X half duplex capability. 0: 100BASE-X half duplex ability not advertised 1: 100BASE-X half duplex ability advertised	R/W	1b
6	10BASE-T Full Duplex This bit determines the advertised 10BASE-T full duplex capability. 0: 10BASE-T full duplex ability not advertised 1: 10BASE-T full duplex ability advertised	R/W	Note 10 Table 12-13

Bits	Description	Type	Default
5	10BASE-T Half Duplex This bit determines the advertised 10BASE-T half duplex capability. 0: 10BASE-T half duplex ability not advertised 1: 10BASE-T half duplex ability advertised	R/W	Note 11 Table 12-14
4:0	Selector Field This field identifies the type of message being sent by Auto-Negotiation. 00001: IEEE 802.3	R/W	00001b

Note 9: The default values of the Asymmetric Pause and Symmetric Pause bits are determined by the Manual Flow Control Enable Strap ([manual_FC_strap_1](#) for PHY A, [manual_FC_strap_2](#) for PHY B). When the Manual Flow Control Enable Strap is 0, the Symmetric Pause bit defaults to 1 and the Asymmetric Pause bit defaults to the setting of the Full-Duplex Flow Control Enable Strap ([FD_FC_strap_1](#) for PHY A, [FD_FC_strap_2](#) for PHY B). When the Manual Flow Control Enable Strap is 1, both bits default to 0. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information. In 100BASE-FX mode, the default value of these bits is 0.

Note 10: The default value of this bit is determined by the logical OR of the Auto-Negotiation Enable strap ([autoneg_strap_1](#) for PHY A, [autoneg_strap_2](#) for PHY B) with the logical AND of the negated Speed Select strap ([speed_strap_1](#) for PHY A, [speed_strap_2](#) for PHY B) and the Duplex Select Strap ([duplex_strap_1](#) for PHY A, [duplex_strap_2](#) for PHY B). [Table 12-13](#) defines the default behavior of this bit. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information. In 100BASE-FX mode, the default value of this bit is a 0.

TABLE 12-13: 10BASE-T FULL DUPLEX ADVERTISEMENT DEFAULT VALUE

autoneg_strap_x	speed_strap_x	duplex_strap_x	Default 10BASE-T Full Duplex (Bit 6) Value
0	0	0	0
0	0	1	1
0	1	0	0
0	1	1	0
1	x	x	1

Note 11: The default value of this bit is determined by the logical OR of the Auto-Negotiation strap ([autoneg_strap_1](#) for PHY A, [autoneg_strap_2](#) for PHY B) and the negated Speed Select strap ([speed_strap_1](#) for PHY A, [speed_strap_2](#) for PHY B). [Table 12-14](#) defines the default behavior of this bit. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information. In 100BASE-FX mode, the default value of this bit is a 0.

TABLE 12-14: 10BASE-T HALF DUPLEX ADVERTISEMENT BIT DEFAULT VALUE

autoneg_strap_x	speed_strap_x	Default 10BASE-T Half Duplex (Bit 5) Value
0	0	1
0	1	0
1	0	1
1	1	1

12.2.19.6 PHY x Auto-Negotiation Link Partner Base Page Ability Register (PHY_AN_LP_BASE_ABILITY_x)

Index (decimal): 5

Size: 16 bits

This read-only register contains the advertised ability of the link partner's PHY and is used in the Auto-Negotiation process between the link partner and the PHY.

Bits	Description	Type	Default
15	Next Page This bit indicates the link partner PHY page capability. 0: Link partner PHY does not advertise next page capability 1: Link partner PHY advertises next page capability	RO	0b
14	Acknowledge This bit indicates whether the link code word has been received from the partner. 0: Link code word not yet received from partner 1: Link code word received from partner	RO	0b
13	Remote Fault This bit indicates whether a remote fault has been detected. 0: No remote fault 1: Remote fault detected	RO	0b
12	Extended Next Page 0: Link partner PHY does not advertise extended next page capability 1: Link partner PHY advertises extended next page capability	RO	0b
11	Asymmetric Pause This bit indicates the link partner PHY asymmetric pause capability. 0: No Asymmetric PAUSE toward link partner 1: Asymmetric PAUSE toward link partner	RO	0b
10	Pause This bit indicates the link partner PHY symmetric pause capability. 0: No Symmetric PAUSE toward link partner 1: Symmetric PAUSE toward link partner	RO	0b
9	100BASE-T4 This bit indicates the link partner PHY 100BASE-T4 capability. 0: 100BASE-T4 ability not supported 1: 100BASE-T4 ability supported	RO	0b
8	100BASE-X Full Duplex This bit indicates the link partner PHY 100BASE-X full duplex capability. 0: 100BASE-X full duplex ability not supported 1: 100BASE-X full duplex ability supported	RO	0b

LAN9352

Bits	Description	Type	Default
7	100BASE-X Half Duplex This bit indicates the link partner PHY 100BASE-X half duplex capability. 0: 100BASE-X half duplex ability not supported 1: 100BASE-X half duplex ability supported	RO	0b
6	10BASE-T Full Duplex This bit indicates the link partner PHY 10BASE-T full duplex capability. 0: 10BASE-T full duplex ability not supported 1: 10BASE-T full duplex ability supported	RO	0b
5	10BASE-T Half Duplex This bit indicates the link partner PHY 10BASE-T half duplex capability. 0: 10BASE-T half duplex ability not supported 1: 10BASE-T half duplex ability supported	RO	0b
4:0	Selector Field This field identifies the type of message being sent by Auto-Negotiation. 00001: IEEE 802.3	RO	00001b

12.2.19.7 PHY x Auto-Negotiation Expansion Register (PHY_AN_EXP_x)

Index (decimal): 6

Size: 16 bits

This read/write register is used in the Auto-Negotiation process between the link partner and the PHY.

Bits	Description	Type	Default
15:7	RESERVED	RO	-
6	Receive Next Page Location Able 0 = Received next page storage location is not specified by bit 6.5 1 = Received next page storage location is specified by bit 6.5	RO	1b
5	Received Next Page Storage Location 0 = Link partner next pages are stored in the PHY x Auto-Negotiation Link Partner Base Page Ability Register (PHY_AN_LP_BASE_ABILITY_x) (PHY register 5) 1 = Link partner next pages are stored in the PHY x Auto Negotiation Next Page RX Register (PHY_AN_NP_RX_x) (PHY register 8)	RO	1b
4	Parallel Detection Fault This bit indicates whether a Parallel Detection Fault has been detected. 0: A fault hasn't been detected via the Parallel Detection function 1: A fault has been detected via the Parallel Detection function	RO/LH	0b
3	Link Partner Next Page Able This bit indicates whether the link partner has next page ability. 0: Link partner does not contain next page capability 1: Link partner contains next page capability	RO	0b
2	Next Page Able This bit indicates whether the local device has next page ability. 0: Local device does not contain next page capability 1: Local device contains next page capability	RO	1b
1	Page Received This bit indicates the reception of a new page. 0: A new page has not been received 1: A new page has been received	RO/LH	0b
0	Link Partner Auto-Negotiation Able This bit indicates the Auto-Negotiation ability of the link partner. 0: Link partner is not Auto-Negotiation able 1: Link partner is Auto-Negotiation able	RO	0b

12.2.19.8 PHY x Auto Negotiation Next Page TX Register (PHY_AN_NP_TX_x)

Index (In Decimal): 7

Size: 16 bits

Bits	Description	Type	Default
15	Next Page 0 = No next page ability 1 = Next page capable	R/W	0b
14	RESERVED	RO	-
13	Message Page 0 = Unformatted page 1 = Message page	R/W	1b
12	Acknowledge 2 0 = Device cannot comply with message. 1 = Device will comply with message.	R/W	0b
11	Toggle 0 = Previous value was HIGH. 1 = Previous value was LOW.	RO	0b
10:0	Message Code Message/Unformatted Code Field	R/W	000 0000 0001b

12.2.19.9 PHY x Auto Negotiation Next Page RX Register (PHY_AN_NP_RX_x)

Index (In Decimal): 8

Size: 16 bits

BITS	DESCRIPTION	TYPE	DEFAULT
15	Next Page 0 = No next page ability 1 = Next page capable	RO	0b
14	Acknowledge 0 = Link code word not yet received from partner 1 = Link code word received from partner	RO	0b
13	Message Page 0 = Unformatted page 1 = Message page	RO	0b
12	Acknowledge 2 0 = Device cannot comply with message. 1 = Device will comply with message.	RO	0b
11	Toggle 0 = Previous value was HIGH. 1 = Previous value was LOW.	RO	0b
10:0	Message Code Message/Unformatted Code Field	RO	000 0000 0000b

12.2.19.10 PHY x MMD Access Control Register (PHY_MMD_ACCESS)

Index (In Decimal): 13

Size: 16 bits

This register in conjunction with the [PHY x MMD Access Address/Data Register \(PHY_MMD_ADDR_DATA\)](#) provides indirect access to the MDIO Manageable Device (MMD) registers. Refer to the [MDIO Manageable Device \(MMD\) Registers on page 276](#) for additional details.

Bits	Description	Type	Default
15:14	MMD Function This field is used to select the desired MMD function: 00 = Address 01 = Data, no post increment 10 = RESERVED 11 = RESERVED	R/W	00b
13:5	RESERVED	RO	-
4:0	MMD Device Address (DEVAD) This field is used to select the desired MMD device address. (3 = PCS, 7 = auto-negotiation)	R/W	0h

12.2.19.11 PHY x MMD Access Address/Data Register (PHY_MMD_ADDR_DATA)

Index (In Decimal): 14

Size: 16 bits

This register in conjunction with the [PHY x MMD Access Control Register \(PHY_MMD_ACCESS\)](#) provides indirect access to the MDIO Manageable Device (MMD) registers. Refer to the [MDIO Manageable Device \(MMD\) Registers on page 276](#) for additional details.

Bits	Description	Type	Default
15:0	MMD Register Address/Data If the MMD Function field of the PHY x MMD Access Control Register (PHY_MMD_ACCESS) is "00", this field is used to indicate the MMD register address to read/write of the device specified in the MMD Device Address (DEVAD) field. Otherwise, this register is used to read/write data from/to the previously specified MMD address.	R/W	0000h

12.2.19.12 PHY x EDPD NLP / Crossover Time / EEE Configuration Register (PHY_EDPD_CFG_x)

Index (decimal): 16

Size: 16 bits

This register is used to Enable EEE functionality and control NLP pulse generation and the Auto-MDIX Crossover Time of the PHY.

Bits	Description	Type	Default
15	EDPD TX NLP Enable Enables the generation of a Normal Link Pulse (NLP) with a selectable interval while in Energy Detect Power-Down. 0=disabled, 1=enabled. The Energy Detect Power-Down (EDPWRDOWN) bit in the PHY x Mode Control/Status Register (PHY_MODE_CONTROL_STATUS_x) needs to be set in order to enter Energy Detect Power-Down mode and the PHY needs to be in the Energy Detect Power-Down state in order for this bit to generate the NLP. The EDPD TX NLP Independent Mode bit of this register also needs to be set when setting this bit.	R/W NASR Note 12	0b
14:13	EDPD TX NLP Interval Timer Select Specifies how often a NLP is transmitted while in the Energy Detect Power-Down state. 00b: 1 s 01b: 768 ms 10b: 512 ms 11b: 256 ms	R/W NASR Note 12	00b
12	EDPD RX Single NLP Wake Enable When set, the PHY will wake upon the reception of a single Normal Link Pulse. When clear, the PHY requires two link pluses, within the interval specified below, in order to wake up. Single NLP Wake Mode is recommended when connecting to “Green” network devices.	R/W NASR Note 12	0b
11:10	EDPD RX NLP Max Interval Detect Select These bits specify the maximum time between two consecutive Normal Link Pulses in order for them to be considered a valid wake up signal. 00b: 64 ms 01b: 256 ms 10b: 512 ms 11b: 1 s	R/W NASR Note 12	00b
9:4	RESERVED	RO	-
3	EDPD TX NLP Independent Mode When set, each PHY port independently detects power down for purposes of the EDPD TX NLP function (via the EDPD TX NLP Enable bit of this register). When cleared, both ports need to be in a power-down state in order to generate TX NLPs during energy detect power-down. Normally set this bit when setting EDPD TX NLP Enable .	R/W NASR Note 12	0b

Bits	Description	Type	Default
2	PHY Energy Efficient Ethernet Enable (PHYEEEEEN) When set, enables Energy Efficient Ethernet (EEE) operation in the PHY. When cleared, EEE operation is disabled. Refer to Section 12.2.11, "Energy Efficient Ethernet," on page 231 for additional information.	R/W NASR Note 12	Note 13
1	EDPD Extend Crossover When in Energy Detect Power-Down (EDPD) mode (Energy Detect Power-Down (EDPWRDOWN) = 1), setting this bit to 1 extends the crossover time by 2976 ms. 0 = Crossover time extension disabled 1 = Crossover time extension enabled (2976 ms)	R/W NASR Note 12	0b
0	Extend Manual 10/100 Auto-MDIX Crossover Time When Auto-Negotiation is disabled, setting this bit extends the Auto-MDIX crossover time by 32 sample times (32 * 62 ms = 1984 ms). This allows the link to be established with a partner PHY that has Auto-Negotiation enabled. When Auto-Negotiation is enabled, this bit has no affect. It is recommended that this bit is set when disabling AN with Auto-MDIX enabled.	R/W NASR Note 12	1b

Note 12: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

Note 13: The default value of this bit is a 0 if in 100BASE-FX mode, otherwise the default value of this bit is determined by the Energy Efficient Ethernet Enable Strap ([EEE_enable_strap_1](#) for PHY A, [EEE_enable_strap_2](#) for PHY B). Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information.

12.2.19.13 PHY x Mode Control/Status Register (PHY_MODE_CONTROL_STATUS_x)

Index (decimal): 17

Size: 16 bits

This read/write register is used to control and monitor various PHY configuration options.

Bits	Description	Type	Default
15:14	RESERVED	RO	-
13	Energy Detect Power-Down (EDPWRDOWN) This bit controls the Energy Detect Power-Down mode. 0: Energy Detect Power-Down is disabled 1: Energy Detect Power-Down is enabled Note: When in EDPD mode, the device's NLP characteristics can be modified via the PHY x EDPD NLP / Crossover Time / EEE Configuration Register (PHY_EDPD_CFG_x) .	R/W	0b
12:7	RESERVED	RO	-
6	ALTINT Alternate Interrupt Mode: 0 = Primary interrupt system enabled (Default) 1 = Alternate interrupt system enabled Refer to Section 12.2.9, "PHY Interrupts," on page 228 for additional information.	R/W NASR Note 14	0b
5:2	RESERVED	RO	-
1	Energy On (ENERGYON) Indicates whether energy is detected. This bit transitions to "0" if no valid energy is detected within 256 ms (1500 ms if auto-negotiation is enabled). It is reset to "1" by a hardware reset and by a software reset if auto-negotiation was enabled or will be enabled via strapping. Refer to Section 12.2.10.2, "Energy Detect Power-Down," on page 231 for additional information.	RO	1b
0	RESERVED	RO	-

Note 14: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.14 PHY x Special Modes Register (PHY_SPECIAL_MODES_x)

Index (decimal): 18

Size: 16 bits

This read/write register is used to control the special modes of the PHY.

Bits	Description	Type	Default
15:11	RESERVED	RO	-
10	100BASE-FX Mode (FX_MODE) This bit enables 100BASE-FX Mode Note: FX_MODE cannot properly be changed with this bit. This bit must always be written with its current value. Device strapping must be used to set the desired mode.	R/W NASR Note 15	Note 16
9:8	RESERVED	RO	-
7:5	PHY Mode (MODE[2:0]) This field controls the PHY mode of operation. Refer to Table 12-15 for a definition of each mode. Note: This field should be written with its read value.	R/W NASR Note 15	Note 17
4:0	PHY Address (PHYADD) The PHY Address field determines the MMI address to which the PHY will respond and is also used for initialization of the cipher (scrambler) key. Each PHY must have a unique address. Refer to Section 12.1.1, "PHY Addressing," on page 218 for additional information. Note: No check is performed to ensure that this address is unique from the other PHY addresses (PHY A, PHY B, and the Virtual PHY).	R/W NASR Note 15	Note 18

Note 15: Register bits designated as NASR are reset when the PHY Reset is generated via the Reset Control Register (RESET_CTL). The NASR designation is only applicable when the Soft Reset (PHY_SRST) bit of the PHY x Basic Control Register (PHY_BASIC_CONTROL_x) is set.

Note 16: The default value of this bit is determined by the Fiber Enable strap (fx_mode_strap_1 for PHY A, fx_mode_strap_2 for PHY B).

Note 17: The default value of this field is determined by a combination of the configuration straps autoneg_strap_x, speed_strap_x, and duplex_strap_x. If the autoneg_strap_x is 1, then the default MODE[2:0] value is 111b. Else, the default value of this field is determined by the remaining straps. MODE[2]=0, MODE[1]=(speed_strap_1 for PHY A, speed_strap_2 for PHY B), and MODE[0]=(duplex_strap_1 for PHY A, duplex_strap_2 for PHY B). Refer to Section 7.0, "Configuration Straps," on page 60 for more information. In 100BASE-FX mode, the default value of these bits is 010b or 011b, depending on the duplex configuration strap.

Note 18: The default value of this field is determined per Section 12.1.1, "PHY Addressing," on page 218.

TABLE 12-15: MODE[2:0] DEFINITIONS

MODE[2:0]	Mode Definitions
000	10BASE-T Half Duplex. Auto-Negotiation disabled.
001	10BASE-T Full Duplex. Auto-Negotiation disabled.

TABLE 12-15: MODE[2:0] DEFINITIONS (CONTINUED)

MODE[2:0]	Mode Definitions
010	100BASE-TX or 100BASE-FX Half Duplex. Auto-Negotiation disabled. CRS is active during Transmit & Receive.
011	100BASE-TX or 100BASE-FX_Full Duplex. Auto-Negotiation disabled. CRS is active during Receive.
100	100BASE-TX Full Duplex is advertised. Auto-Negotiation enabled. CRS is active during Receive.
101	RESERVED
110	Power Down mode.
111	All capable. Auto-Negotiation enabled.

12.2.19.15 PHY x TDR Patterns/Delay Control Register (PHY_TDR_PAT_DELAY_x)

Index (In Decimal): 24

Size: 16 bits

Bits	Description	Type	Default
15	TDR Delay In 0 = Line break time is 2 ms. 1 = The device uses TDR Line Break Counter to increase the line break time before starting TDR.	R/W NASR Note 19	1b
14:12	TDR Line Break Counter When TDR Delay In is 1, this field specifies the increase in line break time in increments of 256 ms, up to 2 seconds.	R/W NASR Note 19	001b
11:6	TDR Pattern High This field specifies the data pattern sent in TDR mode for the high cycle.	R/W NASR Note 19	101110b
5:0	TDR Pattern Low This field specifies the data pattern sent in TDR mode for the low cycle.	R/W NASR Note 19	011101b

Note 19: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.16 PHY x TDR Control/Status Register (PHY_TDR_CONTROL_STAT_x)

Index (In Decimal): 25

Size: 16 bits

Bits	Description	Type	Default
15	TDR Enable 0 = TDR mode disabled 1 = TDR mode enabled Note: This bit self clears when TDR completes (TDR Channel Status goes high)	R/W NASR SC Note 20	0b
14	TDR Analog to Digital Filter Enable 0 = TDR analog to digital filter disabled 1 = TDR analog to digital filter enabled (reduces noise spikes during TDR pulses)	R/W NASR Note 20	0b
13:11	RESERVED	RO	-
10:9	TDR Channel Cable Type Indicates the cable type determined by the TDR test. 00 = Default 01 = Shorted cable condition 10 = Open cable condition 11 = Match cable condition	R/W NASR Note 20	00b
8	TDR Channel Status When high, this bit indicates that the TDR operation has completed. This bit will stay high until reset or the TDR operation is restarted (TDR Enable = 1)	R/W NASR Note 20	0b
7:0	TDR Channel Length This eight bit value indicates the TDR channel length during a short or open cable condition. Refer to Section 12.2.15.1, "Time Domain Reflectometry (TDR) Cable Diagnostics," on page 237 for additional information on the usage of this field. Note: This field is not valid during a match cable condition. The PHY x Cable Length Register (PHY_CABLE_LEN_x) must be used to determine cable length during a non-open/short (match) condition. Refer to Section 12.2.15, "Cable Diagnostics," on page 236 for additional information.	R/W NASR Note 20	00h

Note 20: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.17 PHY x Symbol Error Counter Register

Index (In Decimal): 26

Size: 16 bits

Bits	Description	Type	Default
15:0	<p>Symbol Error Counter (SYM_ERR_CNT) This 100BASE-TX receiver-based error counter increments when an invalid code symbol is received, including IDLE symbols. The counter is incremented only once per packet, even when the received packet contains more than one symbol error. This field counts up to 65,536 and rolls over to 0 if incremented beyond its maximum value.</p> <p>Note: This register is cleared on reset, but is not cleared by reading the register. It does not increment in 10BASE-T mode.</p>	RO	0000h

12.2.19.18 PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x)

Index (decimal): 27

Size:

16 bits

This read/write register is used to control various options of the PHY.

Bits	Description	Type	Default
15	<p>Auto-MDIX Control (AMDIXCTRL) This bit is responsible for determining the source of Auto-MDIX control for Port x. When set, the Manual MDIX and Auto MDIX straps (manual_mdix_strap_1/auto_mdix_strap_1 for PHY A, manual_mdix_strap_2/auto_mdix_strap_2 for PHY B) are overridden, and Auto-MDIX functions are controlled using the AMDIXEN and AMDIXSTATE bits of this register. When cleared, Auto-MDIX functionality is controlled by the Manual MDIX and Auto MDIX straps by default. Refer to Section 7.0, "Configuration Straps," on page 60 for configuration strap definitions.</p> <p>0: Port x Auto-MDIX determined by strap inputs (Table 12-17) 1: Port x Auto-MDIX determined by bits 14 and 13</p> <p>Note: The values of auto_mdix_strap_1 and auto_mdix_strap_2 are indicated in the AMDIX_EN Strap State Port A and the AMDIX_EN Strap State Port B bits of the Hardware Configuration Register (HW_CFG).</p>	R/W NASR Note 21	0b
14	<p>Auto-MDIX Enable (AMDIXEN) When the AMDIXCTRL bit of this register is set, this bit is used in conjunction with the AMDIXSTATE bit to control the Port x Auto-MDIX functionality as shown in Table 12-16.</p> <p>Auto-MDIX is not appropriate and should not be enabled for 100BASE-FX mode.</p>	R/W NASR Note 21	0b
13	<p>Auto-MDIX State (AMDIXSTATE) When the AMDIXCTRL bit of this register is set, this bit is used in conjunction with the AMDIXEN bit to control the Port x Auto-MDIX functionality as shown in Table 12-16.</p>	R/W NASR Note 21	0b
12	RESERVED	RO	-
11	<p>SQE Test Disable (SQEOFF) This bit controls the disabling of the SQE test (Heartbeat). SQE test is enabled by default.</p> <p>0: SQE test enabled 1: SQE test disabled</p>	R/W NASR Note 21	0b
10:6	RESERVED	RO	-
5	<p>Far End Fault Indication Enable (FEFI_EN) This bit enables Far End Fault Generation and Detection. See Section 12.2.17.1, "100BASE-FX Far End Fault Indication," on page 241 for more information.</p>	R/W	Note 22

Bits	Description	Type	Default
4	10Base-T Polarity State (XPOL) This bit shows the polarity state of the 10Base-T. 0: Normal Polarity 1: Reversed Polarity	RO	0b
3:0	RESERVED	RO	-

Note 21: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

Note 22: The default value of this bit is a 1 if in 100BASE-FX mode, otherwise the default is a 0.

TABLE 12-16: AUTO-MDIX ENABLE AND AUTO-MDIX STATE BIT FUNCTIONALITY

Auto-MDIX Enable	Auto-MDIX State	Mode
0	0	Manual mode, no crossover
0	1	Manual mode, crossover
1	0	Auto-MDIX mode
1	1	RESERVED (do not use this state)

TABLE 12-17: MDIX STRAP FUNCTIONALITY

auto_mdix_strap_x	manual_mdix_strap_x	Mode
0	0	Manual mode, no crossover
0	1	Manual mode, crossover
1	x	Auto-MDIX mode

12.2.19.19 PHY x Cable Length Register (PHY_CABLE_LEN_x)

Index (In Decimal): 28

Size: 16 bits

Bits	Description	Type	Default
15:12	Cable Length (CBLN) This four bit value indicates the cable length. Refer to Section 12.2.15.2, "Matched Cable Diagnostics," on page 239 for additional information on the usage of this field. Note: This field indicates cable length for 100BASE-TX linked devices that do not have an open/short on the cable. To determine the open/short status of the cable, the PHY x TDR Patterns/Delay Control Register (PHY_TDR_PAT_DELAY_x) and PHY x TDR Control/Status Register (PHY_TDR_CONTROL_STAT_x) must be used. Cable length is not supported for 10BASE-T links. Refer to Section 12.2.15, "Cable Diagnostics," on page 236 for additional information.	RO	0000b
11:0	RESERVED - Write as 100000000000b, ignore on read	R/W	-

12.2.19.20 PHY x Interrupt Source Flags Register (PHY_INTERRUPT_SOURCE_x)

Index (decimal): 29

Size: 16 bits

This read-only register is used to determine the source of various PHY interrupts. All interrupt source bits in this register are read-only and latch high upon detection of the corresponding interrupt (if enabled). A read of this register clears the interrupts. These interrupts are enabled or masked via the [PHY x Interrupt Mask Register \(PHY_INTERRUPT_MASK_x\)](#).

Bits	Description	Type	Default
15:9	RESERVED	RO	-
9	INT9 This interrupt source bit indicates a Link Up (link status asserted). 0: Not source of interrupt 1: Link Up (link status asserted)	RO/LH	0b
8	INT8 0: Not source of interrupt 1: Wake on LAN (WoL) event detected	RO/LH	0b
7	INT7 This interrupt source bit indicates when the Energy On (ENERGYON) bit of the PHY x Mode Control/Status Register (PHY_MODE_CONTROL_STATUS_x) has been set. 0: Not source of interrupt 1: ENERGYON generated	RO/LH	0b
6	INT6 This interrupt source bit indicates Auto-Negotiation is complete. 0: Not source of interrupt 1: Auto-Negotiation complete	RO/LH	0b
5	INT5 This interrupt source bit indicates a remote fault has been detected. 0: Not source of interrupt 1: Remote fault detected	RO/LH	0b
4	INT4 This interrupt source bit indicates a Link Down (link status negated). 0: Not source of interrupt 1: Link Down (link status negated)	RO/LH	0b
3	INT3 This interrupt source bit indicates an Auto-Negotiation LP acknowledge. 0: Not source of interrupt 1: Auto-Negotiation LP acknowledge	RO/LH	0b

LAN9352

Bits	Description	Type	Default
2	INT2 This interrupt source bit indicates a Parallel Detection fault. 0: Not source of interrupt 1: Parallel Detection fault	RO/LH	0b
1	INT1 This interrupt source bit indicates an Auto-Negotiation page received. 0: Not source of interrupt 1: Auto-Negotiation page received	RO/LH	0b
0	RESERVED	RO	-

12.2.19.21 PHY x Interrupt Mask Register (PHY_INTERRUPT_MASK_x)

Index (decimal): 30

Size: 16 bits

This read/write register is used to enable or mask the various PHY interrupts and is used in conjunction with the [PHY x Interrupt Source Flags Register \(PHY_INTERRUPT_SOURCE_x\)](#).

Bits	Description	Type	Default
15:10	RESERVED	RO	-
9	INT9_MASK This interrupt mask bit enables/masks the Link Up (link status asserted) interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
8	INT8_MASK This interrupt mask bit enables/masks the WoL interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
7	INT7_MASK This interrupt mask bit enables/masks the ENERGYON interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
6	INT6_MASK This interrupt mask bit enables/masks the Auto-Negotiation interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
5	INT5_MASK This interrupt mask bit enables/masks the remote fault interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
4	INT4_MASK This interrupt mask bit enables/masks the Link Down (link status negated) interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
3	INT3_MASK This interrupt mask bit enables/masks the Auto-Negotiation LP acknowledge interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b

LAN9352

Bits	Description	Type	Default
2	INT2_MASK This interrupt mask bit enables/masks the Parallel Detection fault interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
1	INT1_MASK This interrupt mask bit enables/masks the Auto-Negotiation page received interrupt. 0: Interrupt source is masked 1: Interrupt source is enabled	R/W	0b
0	RESERVED	RO	-

12.2.19.22 PHY x Special Control/Status Register (PHY_SPECIAL_CONTROL_STATUS_x)

Index (decimal): 31

Size: 16 bits

This read/write register is used to control and monitor various options of the PHY.

Bits	Description	Type	Default																		
15:13	RESERVED	RO	-																		
12	Autodone This bit indicates the status of the Auto-Negotiation on the PHY. 0: Auto-Negotiation is not completed, is disabled, or is not active 1: Auto-Negotiation is completed	RO	0b																		
11:5	RESERVED - Write as 0000010b, ignore on read	R/W	0000010b																		
4:2	Speed Indication This field indicates the current PHY speed configuration. <table border="1"><thead><tr><th>STATE</th><th>DESCRIPTION</th></tr></thead><tbody><tr><td>000</td><td>RESERVED</td></tr><tr><td>001</td><td>10BASE-T Half-duplex</td></tr><tr><td>010</td><td>100BASE-TX Half-duplex</td></tr><tr><td>011</td><td>RESERVED</td></tr><tr><td>100</td><td>RESERVED</td></tr><tr><td>101</td><td>10BASE-T Full-duplex</td></tr><tr><td>110</td><td>100BASE-TX Full-duplex</td></tr><tr><td>111</td><td>RESERVED</td></tr></tbody></table>	STATE	DESCRIPTION	000	RESERVED	001	10BASE-T Half-duplex	010	100BASE-TX Half-duplex	011	RESERVED	100	RESERVED	101	10BASE-T Full-duplex	110	100BASE-TX Full-duplex	111	RESERVED	RO	XXXb
STATE	DESCRIPTION																				
000	RESERVED																				
001	10BASE-T Half-duplex																				
010	100BASE-TX Half-duplex																				
011	RESERVED																				
100	RESERVED																				
101	10BASE-T Full-duplex																				
110	100BASE-TX Full-duplex																				
111	RESERVED																				
1:0	RESERVED	RO	0b																		

MDIO Manageable Device (MMD) Registers

The device MMD registers adhere to the *IEEE 802.3-2008 45.2 MDIO Interface Registers* specification. The MMD registers are not memory mapped. These registers are accessed indirectly via the [PHY x MMD Access Control Register \(PHY_MMD_ACCESS\)](#) and [PHY x MMD Access Address/Data Register \(PHY_MMD_ADDR_DATA\)](#). The supported MMD device addresses are 3 (PCS), 7 (Auto-Negotiation), and 30 (Vendor Specific). [Table 12-18, "MMD Registers"](#) details the supported registers within each MMD device.

TABLE 12-18: MMD REGISTERS

MMD DEVICE ADDRESS (IN DECIMAL)	INDEX (IN DECIMAL)	REGISTER NAME
3 (PCS)	0	PHY x PCS Control 1 Register (PHY_PCS_CTL1_x)
	1	PHY x PCS Status 1 Register (PHY_PCS_STAT1_x)
	5	PHY x PCS MMD Devices Present 1 Register (PHY_PCS_MMD_PRESENT1_x)
	6	PHY x PCS MMD Devices Present 2 Register (PHY_PCS_MMD_PRESENT2_x)
	20	PHY x EEE Capability Register (PHY_EEE_CAP_x)
	22	PHY x EEE Wake Error Register (PHY_EEE_WAKE_ERR_x)
	32784	PHY x Wakeup Control and Status Register (PHY_WUCSR_x)
	32785	PHY x Wakeup Filter Configuration Register A (PHY_WUF_CFGA_x)
	32786	PHY x Wakeup Filter Configuration Register B (PHY_WUF_CFGB_x)
	32801	PHY x Wakeup Filter Byte Mask Registers (PHY_WUF_MASK_x)
	32802	
	32803	
	32804	
	32805	
	32806	
	32807	
	32808	
	32865	PHY x MAC Receive Address A Register (PHY_RX_ADDRA_x)
	32866	PHY x MAC Receive Address B Register (PHY_RX_ADDRB_x)
	32867	PHY x MAC Receive Address C Register (PHY_RX_ADDRC_x)
7 (Auto-Negotiation)	5	PHY x Auto-Negotiation MMD Devices Present 1 Register (PHY_AN_MMD_PRESENT1_x)
	6	PHY x Auto-Negotiation MMD Devices Present 2 Register (PHY_AN_MMD_PRESENT2_x)
	60	PHY x EEE Advertisement Register (PHY_EEE_ADV_x)
	61	PHY x EEE Link Partner Advertisement Register (PHY_EEE_LP_ADV_x)

TABLE 12-18: MMD REGISTERS (CONTINUED)

MMD DEVICE ADDRESS (IN DECIMAL)	INDEX (IN DECIMAL)	REGISTER NAME
30 (Vendor Specific)	2	PHY x Vendor Specific MMD 1 Device ID 1 Register (PHY_VEND_SPEC_MMD1_DEVID1_x)
	3	PHY x Vendor Specific MMD 1 Device ID 2 Register (PHY_VEND_SPEC_MMD1_DEVID2_x)
	5	PHY x Vendor Specific MMD 1 Devices Present 1 Register (PHY_VEND_SPEC_MMD1_PRESENT1_x)
	6	PHY x Vendor Specific MMD 1 Devices Present 2 Register (PHY_VEND_SPEC_MMD1_PRESENT2_x)
	8	PHY x Vendor Specific MMD 1 Status Register (PHY_VEND_SPEC_MMD1_STAT_x)
	14	PHY x Vendor Specific MMD 1 Package ID 1 Register (PHY_VEND_SPEC_MMD1_PKG_ID1_x)
	15	PHY x Vendor Specific MMD 1 package ID 2 Register (PHY_VEND_SPEC_MMD1_PKG_ID2_x)

To read or write an MMD register, the following procedure must be observed:

1. Write the [PHY x MMD Access Control Register \(PHY_MMD_ACCESS\)](#) with 00b (address) for the [MMD Function](#) field and the desired MMD device (3 for PCS, 7 for Auto-Negotiation) for the [MMD Device Address \(DEVAD\)](#) field.
2. Write the [PHY x MMD Access Address/Data Register \(PHY_MMD_ADDR_DATA\)](#) with the 16-bit address of the desired MMD register to read/write within the previously selected MMD device (PCS or Auto-Negotiation).
3. Write the [PHY x MMD Access Control Register \(PHY_MMD_ACCESS\)](#) with 01b (data) for the [MMD Function](#) field and choose the previously selected MMD device (3 for PCS, 7 for Auto-Negotiation) for the [MMD Device Address \(DEVAD\)](#) field.
4. If reading, read the [PHY x MMD Access Address/Data Register \(PHY_MMD_ADDR_DATA\)](#), which contains the selected MMD register contents. If writing, write the [PHY x MMD Access Address/Data Register \(PHY_MMD_ADDR_DATA\)](#) with the register contents intended for the previously selected MMD register.

Unless otherwise specified, reserved fields must be written with zeros if the register is written.

LAN9352

12.2.19.23 PHY x PCS Control 1 Register (PHY_PCS_CTL1_x)

Index (In Decimal): 3.0

Size: 16 bits

Bits	Description	Type	Default
15:11	RESERVED	RO	-
10	Clock Stop Enable 0 = The PHY cannot stop the clock during Low Power Idle (LPI) 1 = The PHY may stop the clock during LPI Note: This bit has no affect since the device does not support this mode.	R/W	0b
9:0	RESERVED	RO	-

12.2.19.24 PHY x PCS Status 1 Register (PHY_PCS_STAT1_x)

Index (In Decimal): 3.1

Size: 16 bits

Bits	Description	Type	Default
15:12	RESERVED	RO	-
11	TX LPI Received 0 = TX PCS has not received LPI 1 = TX PCS has received LPI	RO/LH	0b
10	RX LPI Received 0 = RX PCS has not received LPI 1 = RX PCS has received LPI	RO/LH	0b
9	TX LPI Indication 0 = TX PCS is not currently receiving LPI 1 = TX PCS is currently receiving LPI	RO	0b
8	RX LPI Indication 0 = RX PCS is not currently receiving LPI 1 = RX PCS is currently receiving LPI	RO	0b
7	RESERVED	RO	-
6	Clock Stop Capable 0 = The MAC cannot stop the clock during Low Power Idle (LPI) 1 = The MAC may stop the clock during LPI Note: The device does not support this mode.	RO	0b
5:0	RESERVED	RO	-

12.2.19.25 PHY x PCS MMD Devices Present 1 Register (PHY_PCS_MMD_PRESENT1_x)

Index (In Decimal): 3.5

Size: 16 bits

Bits	Description	Type	Default
15:8	RESERVED	RO	-
7	Auto-Negotiation Present 0 = Auto-negotiation not present in package 1 = Auto-negotiation present in package	RO	1b
6	TC Present 0 = TC not present in package 1 = TC present in package	RO	0b
5	DTE XS Present 0 = DTE XS not present in package 1 = DTE XS present in package	RO	0b
4	PHY XS Present 0 = PHY XS not present in package 1 = PHY XS present in package	RO	0b
3	PCS Present 0 = PCS not present in package 1 = PCS present in package	RO	1b
2	WIS Present 0 = WIS not present in package 1 = WIS present in package	RO	0b
1	PMD/PMA Present 0 = PMD/PMA not present in package 1 = PMD/PMA present in package	RO	0b
0	Clause 22 Registers Present 0 = Clause 22 registers not present in package 1 = Clause 22 registers present in package	RO	0b

12.2.19.26 PHY x PCS MMD Devices Present 2 Register (PHY_PCS_MMD_PRESENT2_x)

Index (In Decimal): 3.6

Size: 16 bits

Bits	Description	Type	Default
15	Vendor Specific Device 2 Present 0 = Vendor specific device 2 not present in package 1 = Vendor specific device 2 present in package	RO	0b
14	Vendor Specific Device 1 Present 0 = Vendor specific device 1 not present in package 1 = Vendor specific device 1 present in package	RO	1b
13	Clause 22 Extension Present 0 = Clause 22 extension not present in package 1 = Clause 22 extension present in package	RO	0b
12:0	RESERVED	RO	-

LAN9352

12.2.19.27 PHY x EEE Capability Register (PHY_EEE_CAP_x)

Index (In Decimal): 3.20

Size: 16 bits

Bits	Description	Type	Default
15:7	RESERVED	RO	-
6	10GBASE-KR EEE 0 = EEE is not supported for 10GBASE-KR 1 = EEE is supported for 10GBASE-KR Note: The device does not support this mode.	RO	0b
5	10GBASE-KX4 EEE 0 = EEE is not supported for 10GBASE-KX4 1 = EEE is supported for 10GBASE-KX4 Note: The device does not support this mode.	RO	0b
4	10GBASE-KX EEE 0 = EEE is not supported for 10GBASE-KX 1 = EEE is supported for 10GBASE-KX Note: The device does not support this mode.	RO	0b
3	10GBASE-T EEE 0 = EEE is not supported for 10GBASE-T 1 = EEE is supported for 10GBASE-T Note: The device does not support this mode.	RO	0b
2	1000BASE-T EEE 0 = EEE is not supported for 1000BASE-T 1 = EEE is supported for 1000BASE-T Note: The device does not support this mode.	RO	0b
1	100BASE-TX EEE 0 = EEE is not supported for 100BASE-TX 1 = EEE is supported for 100BASE-TX	RO	Note 23
0	RESERVED	RO	-

Note 23: The default value of this field is determined by the value of the [PHY Energy Efficient Ethernet Enable \(PHY_EEEEN\)](#) of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#) on [page 260](#). If [PHY Energy Efficient Ethernet Enable \(PHY_EEEEN\)](#) is 0b, this field is 0b and 100BASE-TX EEE capability is not supported. If [PHY Energy Efficient Ethernet Enable \(PHY_EEEEN\)](#) is 1b, then this field is 1b and 100BASE-TX EEE capability is supported.

12.2.19.28 PHY x EEE Wake Error Register (PHY_EEE_WAKE_ERR_x)

Index (In Decimal): 3.22

Size: 16 bits

Bits	Description	Type	Default
15:0	EEE Wake Error Counter This counter is cleared to zeros on read and is held to all ones on overflow.	RO/RC	0000h

12.2.19.29 PHY x Wakeup Control and Status Register (PHY_WUCSR_x)

Index (In Decimal): 3.32784

Size:

16 bits

Bits	Description	Type	Default
15:9	RESERVED	RO	-
8	WoL Configured This bit may be set by software after the WoL registers are configured. This sticky bit (and all other WoL related register bits) is reset only via a power cycle or a pin reset, allowing software to skip programming of the WoL registers in response to a WoL event. Note: Refer to Section 12.2.12, "Wake on LAN (WoL)," on page 232 for additional information.	R/W/ NASR Note 24	0b
7	Perfect DA Frame Received (PFDA_FR) The MAC sets this bit upon receiving a valid frame with a destination address that matches the physical address.	R/WC/ NASR Note 24	0b
6	Remote Wakeup Frame Received (WUFR) The MAC sets this bit upon receiving a valid remote Wakeup Frame.	R/WC/ NASR Note 24	0b
5	Magic Packet Received (MPR) The MAC sets this bit upon receiving a valid Magic Packet.	R/WC/ NASR Note 24	0b
4	Broadcast Frame Received (BCAST_FR) The MAC Sets this bit upon receiving a valid broadcast frame.	R/WC/ NASR Note 24	0b
3	Perfect DA Wakeup Enable (PFDA_EN) When set, remote wakeup mode is enabled and the MAC is capable of waking up on receipt of a frame with a destination address that matches the physical address of the device. The physical address is stored in the PHY x MAC Receive Address A Register (PHY_RX_ADDRA_x) , PHY x MAC Receive Address B Register (PHY_RX_ADDRB_x) and PHY x MAC Receive Address C Register (PHY_RX_ADDRC_x) .	R/W/ NASR Note 24	0b
2	Wakeup Frame Enable (WUEN) When set, remote wakeup mode is enabled and the MAC is capable of detecting Wakeup Frames as programmed in the Wakeup Filter.	R/W/ NASR Note 24	0b
1	Magic Packet Enable (MPEN) When set, Magic Packet wakeup mode is enabled.	R/W/ NASR Note 24	0b
0	Broadcast Wakeup Enable (BCST_EN) When set, remote wakeup mode is enabled and the MAC is capable of waking up from a broadcast frame.	R/W/ NASR Note 24	0b

Note 24: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.30 PHY x Wakeup Filter Configuration Register A (PHY_WUF_CFGA_x)

Index (In Decimal): 3.32785

Size:

16 bits

Bits	Description	Type	Default
15	Filter Enable 0 = Filter disabled 1 = Filter enabled	R/W/ NASR Note 25	0b
14	Filter Triggered 0 = Filter not triggered 1 = Filter triggered	R/W/ NASR Note 25	0b
13:11	RESERVED	RO	-
10	Address Match Enable When set, the destination address must match the programmed address. When cleared, any unicast packet is accepted. Refer to Section 12.2.12.4, "Wakeup Frame Detection," on page 233 for additional information.	R/W/ NASR Note 25	0b
9	Filter Any Multicast Enable When set, any multicast packet other than a broadcast will cause an address match. Refer to Section 12.2.12.4, "Wakeup Frame Detection," on page 233 for additional information. Note: This bit has priority over bit 10 of this register.	R/W/ NASR Note 25	0b
8	Filter Broadcast Enable When set, any broadcast frame will cause an address match. Refer to Section 12.2.12.4, "Wakeup Frame Detection," on page 233 for additional information. Note: This bit has priority over bit 10 of this register.	R/W/ NASR Note 25	0b
7:0	Filter Pattern Offset Specifies the offset of the first byte in the frame on which CRC checking begins for Wakeup Frame recognition. Offset 0 is the first byte of the incoming frame's destination address.	R/W/ NASR Note 25	00h

Note 25: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.31 PHY x Wakeup Filter Configuration Register B (PHY_WUF_CFGB_x)

Index (In Decimal): 3.32786 Size: 16 bits

Bits	Description	Type	Default
15:0	Filter CRC-16 This field specifies the expected 16-bit CRC value for the filter that should be obtained by using the pattern offset and the byte mask programmed for the filter. This value is compared against the CRC calculated on the incoming frame, and a match indicates the reception of a Wakeup Frame.	R/W/ NASR Note 26	0000h

Note 26: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.32 PHY x Wakeup Filter Byte Mask Registers (PHY_WUF_MASK_x)

Index (In Decimal): 3.32801

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [127:112]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32802

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [111:96]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32803

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [95:80]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32804

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [79:64]	R/W/ NASR Note 27	0000h

LAN9352

Index (In Decimal): 3.32805

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [63:48]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32806

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [47:32]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32807

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [31:16]	R/W/ NASR Note 27	0000h

Index (In Decimal): 3.32808

Size: 16 bits

Bits	Description	Type	Default
15:0	Wakeup Filter Byte Mask [15:0]	R/W/ NASR Note 27	0000h

Note 27: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.33 PHY x MAC Receive Address A Register (PHY_RX_ADDRA_x)

Index (In Decimal): 3.32865

Size: 16 bits

Bits	Description	Type	Default
15:0	Physical Address [47:32]	R/W/ NASR Note 28	FFFFh

Note 28: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.34 PHY x MAC Receive Address B Register (PHY_RX_ADDRB_x)

Index (In Decimal): 3.32866

Size: 16 bits

Bits	Description	Type	Default
15:0	Physical Address [31:16]	R/W/ NASR Note 29	FFFFh

Note 29: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.35 PHY x MAC Receive Address C Register (PHY_RX_ADDRC_x)

Index (In Decimal): 3.32867

Size: 16 bits

Bits	Description	Type	Default
15:0	Physical Address [15:0]	R/W/ NASR Note 30	FFFFh

Note 30: Register bits designated as NASR are reset when the PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Soft Reset \(PHY_SRST\)](#) bit of the [PHY x Basic Control Register \(PHY_BASIC_CONTROL_x\)](#) is set.

12.2.19.36 PHY x Auto-Negotiation MMD Devices Present 1 Register (PHY_AN_MMD_PRESENT1_x)

Index (In Decimal): 7.5

Size: 16 bits

Bits	Description	Type	Default
15:8	RESERVED	RO	-
7	Auto-Negotiation Present 0 = Auto-negotiation not present in package 1 = Auto-negotiation present in package	RO	1b
6	TC Present 0 = TC not present in package 1 = TC present in package	RO	0b
5	DTE XS Present 0 = DTE XS not present in package 1 = DTE XS present in package	RO	0b
4	PHY XS Present 0 = PHY XS not present in package 1 = PHY XS present in package	RO	0b
3	PCS Present 0 = PCS not present in package 1 = PCS present in package	RO	1b
2	WIS Present 0 = WIS not present in package 1 = WIS present in package	RO	0b
1	PMD/PMA Present 0 = PMD/PMA not present in package 1 = PMD/PMA present in package	RO	0b
0	Clause 22 Registers Present 0 = Clause 22 registers not present in package 1 = Clause 22 registers present in package	RO	0b

12.2.19.37 PHY x Auto-Negotiation MMD Devices Present 2 Register (PHY_AN_MMD_PRESENT2_x)

Index (In Decimal): 7.6

Size: 16 bits

Bits	Description	Type	Default
15	Vendor Specific Device 2 Present 0 = Vendor specific device 2 not present in package 1 = Vendor specific device 2 present in package	RO	0b
14	Vendor Specific Device 1 Present 0 = Vendor specific device 1 not present in package 1 = Vendor specific device 1 present in package	RO	1b
13	Clause 22 Extension Present 0 = Clause 22 extension not present in package 1 = Clause 22 extension present in package	RO	0b
12:0	RESERVED	RO	-

LAN9352

12.2.19.38 PHY x EEE Advertisement Register (PHY_EEE_ADV_x)

Index (In Decimal): 7.60

Size: 16 bits

BITS	DESCRIPTION	TYPE	DEFAULT
15:2	RESERVED	RO	-
1	100BASE-TX EEE 0 = Do not advertise EEE capability for 100BASE-TX. 1 = Advertise EEE capability for 100BASE-TX.	Note 31	Note 32
0	RESERVED	RO	-

Note 31: This bit is read/write (R/W). However, the user must not set this bit if EEE is disabled.

Note 32: The default value of this field is determined by the value of the [PHY Energy Efficient Ethernet Enable \(PHY-EEEN\)](#) of the [PHY x EDPD NLP / Crossover Time / EEE Configuration Register \(PHY_EDPD_CFG_x\)](#) on [page 260](#). If [PHY Energy Efficient Ethernet Enable \(PHYEEEN\)](#) is 0b, this field is 0b and 100BASE-TX EEE capability is not advertised. If [PHY Energy Efficient Ethernet Enable \(PHYEEEN\)](#) is 1b, then this field is 1b and 100BASE-TX EEE capability is advertised.

12.2.19.39 PHY x EEE Link Partner Advertisement Register (PHY_EEE_LP_ADV_x)

Index (In Decimal): 7.61

Size: 16 bits

BITS	DESCRIPTION	TYPE	DEFAULT
15:7	RESERVED	RO	-
6	10GBASE-KR EEE 0 = Link partner does not advertise EEE capability for 10GBASE-KR. 1 = Link partner advertises EEE capability for 10GBASE-KR. Note: This device does not support this mode.	RO	0b
5	10GBASE-KX4 EEE 0 = Link partner does not advertise EEE capability for 10GBASE-KX4. 1 = Link partner advertises EEE capability for 10GBASE-KX4. Note: This device does not support this mode.	RO	0b
4	10GBASE-KX EEE 0 = Link partner does not advertise EEE capability for 10GBASE-KX. 1 = Link partner advertises EEE capability for 10GBASE-KX. Note: This device does not support this mode.	RO	0b
3	10GBASE-T EEE 0 = Link partner does not advertise EEE capability for 10GBASE-T. 1 = Link partner advertises EEE capability for 10GBASE-T. Note: This device does not support this mode.	RO	0b
2	1000BASE-T EEE 0 = Link partner does not advertise EEE capability for 1000BASE-T. 1 = Link partner advertises EEE capability for 1000BASE-T. Note: This device does not support this mode.	RO	0b
1	100BASE-TX EEE 0 = Link partner does not advertise EEE capability for 100BASE-TX. 1 = Link partner advertises EEE capability for 100BASE-TX.	RO	0b
0	RESERVED	RO	-

LAN9352

12.2.19.40 PHY x Vendor Specific MMD 1 Device ID 1 Register (PHY_VEND_SPEC_MMD1_DEVID1_x)

Index (In Decimal): 30.2

Size: 16 bits

Bits	Description	Type	Default
15:0	RESERVED	RO	0000h

12.2.19.41 PHY x Vendor Specific MMD 1 Device ID 2 Register (PHY_VEND_SPEC_MMD1_DEVID2_x)

Index (In Decimal): 30.3

Size: 16 bits

Bits	Description	Type	Default
15:0	RESERVED	RO	0000h

12.2.19.42 PHY x Vendor Specific MMD 1 Devices Present 1 Register (PHY_VEND_SPEC_MMD1_PRESENT1_x)

Index (In Decimal): 30.5

Size: 16 bits

Bits	Description	Type	Default
15:8	RESERVED	RO	-
7	Auto-Negotiation Present 0 = Auto-negotiation not present in package 1 = Auto-negotiation present in package	RO	1b
6	TC Present 0 = TC not present in package 1 = TC present in package	RO	0b
5	DTE XS Present 0 = DTE XS not present in package 1 = DTE XS present in package	RO	0b
4	PHY XS Present 0 = PHY XS not present in package 1 = PHY XS present in package	RO	0b
3	PCS Present 0 = PCS not present in package 1 = PCS present in package	RO	1b
2	WIS Present 0 = WIS not present in package 1 = WIS present in package	RO	0b
1	PMD/PMA Present 0 = PMD/PMA not present in package 1 = PMD/PMA present in package	RO	0b
0	Clause 22 Registers Present 0 = Clause 22 registers not present in package 1 = Clause 22 registers present in package	RO	0b

12.2.19.43 PHY x Vendor Specific MMD 1 Devices Present 2 Register (PHY_VEND_SPEC_MMD1_PRESENT2_x)

Index (In Decimal): 30.6

Size: 16 bits

Bits	Description	Type	Default
15	Vendor Specific Device 2 Present 0 = Vendor specific device 2 not present in package 1 = Vendor specific device 2 present in package	RO	0b
14	Vendor Specific Device 1 Present 0 = Vendor specific device 1 not present in package 1 = Vendor specific device 1 present in package	RO	1b
13	Clause 22 Extension Present 0 = Clause 22 extension not present in package 1 = Clause 22 extension present in package	RO	0b
12:0	RESERVED	RO	-

LAN9352

12.2.19.44 PHY x Vendor Specific MMD 1 Status Register (PHY_VEND_SPEC_MMD1_STAT_x)

Index (In Decimal): 30.8

Size: 16 bits

Bits	Description	Type	Default
15:14	Device Present 00 = No device responding at this address 01 = No device responding at this address 10 = Device responding at this address 11 = No device responding at this address	RO	10b
13:0	RESERVED	RO	-

12.2.19.45 PHY x Vendor Specific MMD 1 Package ID 1 Register (PHY_VEND_SPEC_MMD1_PKG_ID1_x)

Index (In Decimal): 30.14

Size: 16 bits

Bits	Description	Type	Default
15:0	RESERVED	RO	0000h

LAN9352

12.2.19.46 PHY x Vendor Specific MMD 1 package ID 2 Register
(PHY_VEND_SPEC_MMD1_PKG_ID2_x)

Index (In Decimal): 30.15

Size: 16 bits

Bits	Description	Type	Default
15:0	RESERVED	RO	0000h

12.3 Virtual PHY

The Virtual PHY provides a basic MII management interface (MDIO) per IEEE 802.3 (clause 22) so that a MAC with an unmodified driver can be supported as if it was attached to a single port PHY. This functionality is designed to allow easy and quick integration of the device into designs with minimal driver modifications. The Virtual PHY provides a full bank of registers which comply with the IEEE 802.3 specification. This enables the Virtual PHY to provide various status and control bits similar to those provided by a real PHY. These include the output of speed selection, duplex, loopback, isolate, collision test, and Auto-Negotiation status. For a list of all Virtual PHY registers and related bit descriptions, refer to [Section 12.3.3, "Virtual PHY Registers," on page 305](#).

12.3.1 VIRTUAL PHY AUTO-NEGOTIATION

The purpose of the Auto-Negotiation function is to automatically configure the Virtual PHY to the optimum link parameters based on the capabilities of its link partner. Because the Virtual PHY has no actual link partner, the Auto-Negotiation process is emulated with deterministic results.

Auto-Negotiation is enabled by setting the [Auto-Negotiation \(VPHY_AN\)](#) bit of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#) and is restarted by the occurrence of any of the following events:

- Power-On Reset (POR)
- Hardware reset (**RST#**)
- PHY Software reset (via the **Virtual PHY Reset (VPHY_RST)** bit of the [Reset Control Register \(RESET_CTL\)](#) or the [Reset \(VPHY_RST\)](#) bit of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#))
- Setting the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#), [Restart Auto-Negotiation \(VPHY_RST_AN\)](#) bit high
- Digital Reset (via the [Digital Reset \(DIGITAL_RST\)](#) bit of the [Reset Control Register \(RESET_CTL\)](#))
- Issuing an EEPROM Loader RELOAD command ([Section 14.4, "EEPROM Loader," on page 465](#))

Note: Auto-Negotiation is also restarted after the EEPROM Loader updates the straps.

The emulated Auto-Negotiation process is much simpler than the real process and can be categorized into three steps:

1. The [Auto-Negotiation Complete](#) bit is set in the [Virtual PHY Basic Status Register \(VPHY_BASIC_STATUS\)](#).
2. The [Page Received](#) bit is set in the [Virtual PHY Auto-Negotiation Expansion Register \(VPHY_AN_EXP\)](#).
3. The Auto-Negotiation result (speed, duplex and pause) is determined and registered.

The Auto-Negotiation result (speed and duplex) is determined using the Highest Common Denominator (HCD) of the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#) and [Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register \(VPHY_AN_LP_BASE_ABILITY\)](#) as specified in the IEEE 802.3 standard. The technology ability bits of these registers are ANDed, and if there are multiple bits in common, the priority is determined as follows:

- 100Mbps Full Duplex (highest priority)
- 100Mbps Half Duplex
- 10Mbps Full Duplex
- 10Mbps Half Duplex (lowest priority)

For example, if the full capabilities of the Virtual PHY are advertised (100Mbps, Full Duplex), and if the link partner is capable of 10Mbps and 100Mbps, then Auto-Negotiation selects 100Mbps as the highest performance mode. If the link partner is capable of half and full-duplex modes, then Auto-Negotiation selects full-duplex as the highest performance operation. In the event that there are no bits in common, an emulated [Parallel Detection](#) is used.

The [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#) defaults to having all four ability bits set. These values can be reconfigured via software. Once the Auto-Negotiation is complete, any change to the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#) only takes affect when the Auto-Negotiation process is re-run.

The emulated link partner always advertises all four abilities (100BASE-X full duplex, 100BASE-X half duplex, 10BASE-T full duplex, and 10BASE-T half duplex) in the [Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register \(VPHY_AN_LP_BASE_ABILITY\)](#).

Neither the Virtual PHY or the emulated link partner support next page capability, remote faults, or 100BASE-T4.

If there is at least one common selection between the emulated link partner and the Virtual PHY advertised abilities, then the Auto-Negotiation succeeds, the [Link Partner Auto-Negotiation Able](#) bit of the [Virtual PHY Auto-Negotiation Expansion Register \(VPHY_AN_EXP\)](#) is set, and the technology ability bits in the [Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register \(VPHY_AN_LP_BASE_ABILITY\)](#) are set to indicate the emulated link partners abilities.

Note: For the Virtual PHY, the Auto-Negotiation register bits (and management of such) are used by the MAC driver, so the perception of local and link partner is reversed. The local device is the MAC, while the link partner is the switch fabric. This is consistent with the intention of the Virtual PHY.

12.3.1.1 Parallel Detection

In the event that there are no common bits between the advertised ability and the emulated link partners ability, Auto-Negotiation fails and emulated parallel detect is used. In this case, the [Link Partner Auto-Negotiation Able](#) bit in the [Virtual PHY Auto-Negotiation Expansion Register \(VPHY_AN_EXP\)](#) will be cleared, and the communication set to half-duplex. The speed is set to 100Mbps. Only one of the technology ability bits in the [Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register \(VPHY_AN_LP_BASE_ABILITY\)](#) will be set, indicating the emulated parallel detect result.

12.3.1.2 Disabling Auto-Negotiation

Auto-Negotiation can be disabled in the Virtual PHY by clearing the [Auto-Negotiation \(VPHY_AN\)](#) bit of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#). The Virtual PHY will then force its speed of operation to reflect the speed ([Speed Select LSB \(VPHY_SPEED_SEL_LSB\)](#)) and duplex ([Duplex Mode \(VPHY_DUPLEX\)](#)) of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#). The speed and duplex bits in the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#) are ignored when Auto-Negotiation is enabled.

12.3.1.3 Virtual PHY Pause Flow Control

The Virtual PHY supports pause flow control per the IEEE 802.3 specification. The Virtual PHY's advertised pause flow control abilities are set via the [Symmetric Pause](#) and [Asymmetric Pause](#) bits of the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#). This allows the Virtual PHY to advertise its flow control abilities and Auto-Negotiate the flow control settings with the emulated link partner. The default values of these bits are as shown in [Section 12.3.3.5, "Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)," on page 312](#).

The symmetric/asymmetric pause ability of the emulated link partner is based upon the advertised pause flow control abilities of the Virtual PHY as indicated in the [Symmetric Pause](#) and [Asymmetric Pause](#) bits of the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#). Thus, the emulated link partner always accommodates the asymmetric/symmetric pause ability settings requested by the Virtual PHY, as shown in [Table 12-20, "Emulated Link Partner Pause Flow Control Ability Default Values," on page 315](#).

The pause flow control settings may also be manually set via the [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#). This register allows the Switch Fabric port flow control settings to be manually set when Auto-Negotiation is disabled or the [Port 0 Full-Duplex Manual Flow Control Select \(MANUAL_FC_0\)](#) is set. The currently enabled duplex and flow control settings can also be monitored via this register. The flow control values in the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#) are not affected by the values of the manual flow control register. Refer to [Section 13.5.1, "Flow Control Enable Logic," on page 343](#) for additional information.

12.3.2 VIRTUAL PHY RESETS

In addition to the chip-level hardware reset (**RST#**) and Power-On Reset (POR), block specific resets are supported. These are discussed in the following sections. For detailed information on all device resets, refer to [Section 6.2, "Resets," on page 42](#).

12.3.2.1 Virtual PHY Software Reset via RESET_CTL

The Virtual PHY can be reset via the [Reset Control Register \(RESET_CTL\)](#) by setting the **Virtual PHY Reset (VPHY_RST)** bit. This bit is self clearing after approximately 102 us.

12.3.2.2 Virtual PHY Software Reset via VPHY_BASIC_CTRL

The Virtual PHY can also be reset by setting the [Reset \(VPHY_RST\)](#) bit 15 of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#). This bit is self clearing and will return to 0 after the reset is complete.

12.3.3 VIRTUAL PHY REGISTERS

This section details the Virtual PHY System CSRs. These registers provide status and control information similar to that of a real PHY while maintaining IEEE 802.3 compatibility. The Virtual PHY registers are addressable via the memory map, as described in [Table 5-1, "System Control and Status Registers," on page 35](#), as well as serially via the MII management protocol (IEEE 802.3 clause 22). When accessed serially, these registers are accessed indirectly through the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#) and [Host MAC MII Data Register \(HMAC_MII_DATA\)](#) via the MII serial management protocol specified in IEEE 802.3 clause 22.

When being accessed serially, the Virtual PHY will respond when the PHY address equals the address assigned by the [phy_addr_sel_strap](#) configuration strap, as defined in [Section 12.1.1, "PHY Addressing," on page 218](#). A list of all Virtual PHY register indexes for serial access can be seen in [Table 12-19](#). For Virtual PHY functionality and operation information, see [Section 12.3, "Virtual PHY," on page 303](#).

Note: All Virtual PHY registers follow the IEEE 802.3 (clause 22.2.4) specified MII management register set. All functionality and bit definitions comply with these standards. The IEEE 802.3 specified register index (in decimal) is included under the memory mapped offset of each Virtual PHY register as a reference. For additional information, refer to the IEEE 802.3 Specification.

Note: When serially accessed, the Virtual PHY registers are only 16-bits wide, as is standard for MII management of PHYs.

TABLE 12-19: VIRTUAL PHY MII SERIALY ADDRESSABLE REGISTER INDEX

ADDRESS (DIRECT)	INDEX # (INDIRECT)	Register Name (SYMBOL)
1C0h	0	Virtual PHY Basic Control Register (VPHY_BASIC_CTRL)
1C4h	1	Virtual PHY Basic Status Register (VPHY_BASIC_STATUS)
1C8h	2	Virtual PHY Identification MSB Register (VPHY_ID_MSB)
1CCh	3	Virtual PHY Identification LSB Register (VPHY_ID_LSB)
1D0h	4	Virtual PHY Auto-Negotiation Advertisement Register (VPHY_AN_ADV)
1D4h	5	Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register (VPHY_AN_LP_BASE_ABILITY)
1D8h	6	Virtual PHY Auto-Negotiation Expansion Register (VPHY_AN_EXP)
1DCh	31	Virtual PHY Special Control/Status Register (VPHY_SPECIAL_CONTROL_STATUS)

LAN9352

12.3.3.1 Virtual PHY Basic Control Register (VPHY_BASIC_CTRL)

Offset: 1C0h Size: 32 bits
Index (decimal): 0 16 bits

This read/write register is used to configure the Virtual PHY.

Bits	Description	Type	Default
31:16	RESERVED (See Note 33)	RO	-
15	Reset (VPHY_RST) When set, this bit resets the Virtual PHY registers to their default state. This bit is self clearing. 0: Normal Operation 1: Reset	R/W SC	0b
14	Loopback (VPHY_LOOPBACK) This bit enables/disables the loopback mode. When enabled, transmissions from the Host MAC are not sent to the Switch Fabric. Instead, they are looped back onto the receive path. 0: Loopback mode disabled (normal operation) 1: Loopback mode enabled	R/W	0b
13	Speed Select LSB (VPHY_SPEED_SEL_LSB) This bit is used to set the speed of the Virtual PHY when the Auto-Negotiation (VPHY_AN) bit is disabled. 0: 10 Mbps 1: 100/200 Mbps	R/W	0b
12	Auto-Negotiation (VPHY_AN) This bit enables/disables Auto-Negotiation. When enabled, the Speed Select LSB (VPHY_SPEED_SEL_LSB) and Duplex Mode (VPHY_DUPLEX) bits are overridden. 0: Auto-Negotiation disabled 1: Auto-Negotiation enabled	R/W	1b
11	Power Down (VPHY_PWR_DWN) This bit is not used by the Virtual PHY and has no effect.	R/W	0b
10	Isolate (VPHY_ISO) This bit is not used by the Virtual PHY and has no effect.	R/W	0b
9	Restart Auto-Negotiation (VPHY_RST_AN) When set, this bit updates the emulated Auto-Negotiation results. 0: Normal operation 1: Auto-Negotiation restarted	R/W SC	0b

Bits	Description	Type	Default
8	Duplex Mode (VPHY_DUPLEX) This bit is used to set the duplex when the Auto-Negotiation (VPHY_AN) bit is disabled. 0: Half Duplex 1: Full Duplex	R/W	0b
7	Collision Test (VPHY_COL_TEST) This bit enables/disables the collision test mode. When set, the collision signal to the Host MAC is active during transmission from the MAC. Note: It is recommended that this bit be used only when in loopback mode. 0: Collision test mode disabled 1: Collision test mode enabled	R/W	0b
6	Speed Select MSB (VPHY_SPEED_SEL_MSB) This bit is not used by the Virtual PHY and has no effect. The value returned is always 0.	RO	0b
5:0	RESERVED	RO	-

Note 33: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

12.3.3.2 Virtual PHY Basic Status Register (VPHY_BASIC_STATUS)

Offset: 1C4h Size: 32 bits
Index (decimal): 1 16 bits

This register is used to monitor the status of the Virtual PHY.

Bits	Description	Type	Default
31:16	RESERVED (See Note 34)	RO	-
15	100BASE-T4 This bit displays the status of 100BASE-T4 compatibility. 0: PHY not able to perform 100BASE-T4 1: PHY able to perform 100BASE-T4	RO	0b Note 35
14	100BASE-X Full Duplex This bit displays the status of 100BASE-X full duplex compatibility. 0: PHY not able to perform 100BASE-X full duplex 1: PHY able to perform 100BASE-X full duplex	RO	1b
13	100BASE-X Half Duplex This bit displays the status of 100BASE-X half duplex compatibility. 0: PHY not able to perform 100BASE-X half duplex 1: PHY able to perform 100BASE-X half duplex	RO	1b
12	10BASE-T Full Duplex This bit displays the status of 10BASE-T full duplex compatibility. 0: PHY not able to perform 10BASE-T full duplex 1: PHY able to perform 10BASE-T full duplex	RO	1b
11	10BASE-T Half Duplex This bit displays the status of 10BASE-T half duplex compatibility. 0: PHY not able to perform 10BASE-T half duplex 1: PHY able to perform 10BASE-T half duplex	RO	1b
10	100BASE-T2 Full Duplex This bit displays the status of 100BASE-T2 full duplex compatibility. 0: PHY not able to perform 100BASE-T2 full duplex 1: PHY able to perform 100BASE-T2 full duplex	RO	0b Note 35
9	100BASE-T2 Half Duplex This bit displays the status of 100BASE-T2 half duplex compatibility. 0: PHY not able to perform 100BASE-T2 half duplex 1: PHY able to perform 100BASE-T2 half duplex	RO	0b Note 35
8	Extended Status This bit displays whether extended status information is in register 15 (per IEEE 802.3 clause 22.2.4). 0: No extended status information in Register 15 1: Extended status information in Register 15	RO	0b Note 36

Bits	Description	Type	Default
7	RESERVED	RO	-
6	MF Preamble Suppression This bit indicates whether the Virtual PHY accepts management frames with the preamble suppressed. 0: Management frames with preamble suppressed not accepted 1: Management frames with preamble suppressed accepted	RO	0b
5	Auto-Negotiation Complete This bit indicates the status of the Auto-Negotiation process. 0: Auto-Negotiation process not completed 1: Auto-Negotiation process completed	RO	1b Note 37
4	Remote Fault This bit indicates if a remote fault condition has been detected. 0: No remote fault condition detected 1: Remote fault condition detected	RO	0b Note 38
3	Auto-Negotiation Ability This bit indicates the status of the Virtual PHY's Auto-Negotiation. 0: Virtual PHY is unable to perform Auto-Negotiation 1: Virtual PHY is able to perform Auto-Negotiation	RO	1b
2	Link Status This bit indicates the status of the link. 0: Link is down 1: Link is up	RO	1b Note 38
1	Jabber Detect This bit indicates the status of the jabber condition. 0: No jabber condition detected 1: Jabber condition detected	RO	0b Note 38
0	Extended Capability This bit indicates whether extended register capability is supported. 0: Basic register set capabilities only 1: Extended register set capabilities	RO	1b Note 39

Note 34: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 35: The Virtual PHY supports 100BASE-X (half and full duplex) and 10BASE-T (half and full duplex) only. All other modes will always return as 0 (unable to perform).

Note 36: The Virtual PHY does not support Register 15 or 1000 Mb/s operation. Thus this bit is always returned as 0.

Note 37: The [Auto-Negotiation Complete](#) bit is first cleared on a reset, but set shortly after (when the Auto-Negotiation process is run). Refer to [Section 12.3.1, "Virtual PHY Auto-Negotiation," on page 303](#) for additional details.

Note 38: The Virtual PHY never has remote faults, its link is always up, and does not detect jabber.

Note 39: The Virtual PHY supports basic and some extended register capability. The Virtual PHY supports Registers 0-6 (per the IEEE 802.3 specification).

12.3.3.3 Virtual PHY Identification MSB Register (VPHY_ID_MSB)

Offset: 1C8h Size: 32 bits
Index (decimal): 2 16 bits

This read/write register contains the MSB of the Virtual PHY Organizationally Unique Identifier (OUI). The LSB of the Virtual PHY OUI is contained in the [Virtual PHY Identification LSB Register \(VPHY_ID_LSB\)](#).

Bits	Description	Type	Default
31:16	RESERVED (See Note 40)	RO	-
15:0	PHY ID This field contains the MSB of the Virtual PHY OUI (Note 41).	R/W	0000h

Note 40: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 41: IEEE allows a value of zero in each of the 32-bits of the PHY Identifier.

12.3.3.4 Virtual PHY Identification LSB Register (VPHY_ID_LSB)

Offset: 1CCh Size: 32 bits
Index (decimal): 3 16 bits

This read/write register contains the LSB of the Virtual PHY Organizationally Unique Identifier (OUI). The MSB of the Virtual PHY OUI is contained in the [Virtual PHY Identification MSB Register \(VPHY_ID_MSB\)](#).

BITS	DESCRIPTION	TYPE	DEFAULT
31:16	RESERVED (See Note 42)	RO	-
15:10	PHY ID This field contains the lower 6-bits of the Virtual PHY OUI (Note 43).	R/W	000000b
9:4	Model Number This field contains the 6-bit manufacturer's model number of the Virtual PHY (Note 43).	R/W	000000b
3:0	Revision Number This field contain the 4-bit manufacturer's revision number of the Virtual PHY (Note 43).	R/W	0000b

Note 42: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 43: IEEE allows a value of zero in each of the 32-bits of the PHY Identifier.

12.3.3.5 Virtual PHY Auto-Negotiation Advertisement Register (VPHY_AN_ADV)

Offset: 1D0h Size: 32 bits
Index (decimal): 4 16 bits

This read/write register contains the advertised ability of the Virtual PHY and is used in the Auto-Negotiation process with the link partner.

Bits	Description	Type	Default
31:16	RESERVED (See Note 44)	RO	-
15	Next Page This bit determines the advertised next page capability and is always 0. 0: Virtual PHY does not advertise next page capability 1: Virtual PHY advertises next page capability	RO	0b Note 45
14	RESERVED	RO	-
13	Remote Fault This bit is not used since there is no physical link partner.	RO	0b Note 46
12	RESERVED	RO	-
11	Asymmetric Pause This bit determines the advertised asymmetric pause capability. 0: No Asymmetric PAUSE toward link partner advertised 1: Asymmetric PAUSE toward link partner advertised	R/W	Note 47
10	Symmetric Pause This bit determines the advertised symmetric pause capability. 0: No Symmetric PAUSE toward link partner advertised 1: Symmetric PAUSE toward link partner advertised	R/W	Note 47
9	100BASE-T4 This bit determines the advertised 100BASE-T4 capability and is always 0. 0: 100BASE-T4 ability not advertised 1: 100BASE-T4 ability advertised	RO	0b Note 48
8	100BASE-X Full Duplex This bit determines the advertised 100BASE-X full duplex capability. 0: 100BASE-X full duplex ability not advertised 1: 100BASE-X full duplex ability advertised	R/W	1b
7	100BASE-X Half Duplex This bit determines the advertised 100BASE-X half duplex capability. 0: 100BASE-X half duplex ability not advertised 1: 100BASE-X half duplex ability advertised	R/W	1b

Bits	Description	Type	Default
6	10BASE-T Full Duplex This bit determines the advertised 10BASE-T full duplex capability. 0: 10BASE-T full duplex ability not advertised 1: 10BASE-T full duplex ability advertised	R/W	1b
5	10BASE-T Half Duplex This bit determines the advertised 10BASE-T half duplex capability. 0: 10BASE-T half duplex ability not advertised 1: 10BASE-T half duplex ability advertised	R/W	1b
4:0	Selector Field This field identifies the type of message being sent by Auto-Negotiation. 00001: IEEE 802.3	R/W	00001b Note 49

Note 44: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 45: The Virtual PHY does not support next page capability. This bit value will always be 0.

Note 46: The [Remote Fault](#) bit is not useful since there is no actual link partner to send a fault to.

Note 47: The [Symmetric Pause](#) and [Asymmetric Pause](#) bits default to 1 if the manual_FC_strap_0 configuration strap is low (both Symmetric and Asymmetric are advertised), and 0 if the manual_FC_strap_0 configuration strap is high.

Note 48: Virtual 100BASE-T4 is not supported.

Note 49: The Virtual PHY supports only IEEE 802.3. Only a value of 00001b should be used in this field.

12.3.3.6 Virtual PHY Auto-Negotiation Link Partner Base Page Ability Register (VPHY_AN_LP_BASE_ABILITY)

Offset: 1D4h Size: 32 bits
Index (decimal): 5 16 bits

This read-only register contains the advertised ability of the link partner's PHY and is used in the Auto-Negotiation process with the Virtual PHY. Because the Virtual PHY does not physically connect to an actual link partner, the values in this register are emulated as described below.

Bits	Description	Type	Default
31:16	RESERVED (See Note 50)	RO	-
15	Next Page This bit indicates the emulated link partner PHY next page capability and is always 0. 0: Link partner PHY does not advertise next page capability 1: Link partner PHY advertises next page capability	RO	0b Note 51
14	Acknowledge This bit indicates whether the link code word has been received from the partner and is always 1. 0: Link code word not yet received from partner 1: Link code word received from partner	RO	1b Note 51
13	Remote Fault Since there is no physical link partner, this bit is not used and is always returned as 0.	RO	0b Note 51
12	RESERVED	RO	-
11	Asymmetric Pause This bit indicates the emulated link partner PHY asymmetric pause capability. 0: No Asymmetric PAUSE toward link partner 1: Asymmetric PAUSE toward link partner	RO	Note 52
10	Pause This bit indicates the emulated link partner PHY symmetric pause capability. 0: No Symmetric PAUSE toward link partner 1: Symmetric PAUSE toward link partner	RO	Note 52
9	100BASE-T4 This bit indicates the emulated link partner PHY 100BASE-T4 capability. This bit is always 0. 0: 100BASE-T4 ability not supported 1: 100BASE-T4 ability supported	RO	0b Note 51
8	100BASE-X Full Duplex This bit indicates the emulated link partner PHY 100BASE-X full duplex capability. 0: 100BASE-X full duplex ability not supported 1: 100BASE-X full duplex ability supported	RO	Note 53

Bits	Description	Type	Default
7	100BASE-X Half Duplex This bit indicates the emulated link partner PHY 100BASE-X half duplex capability. 0: 100BASE-X half duplex ability not supported 1: 100BASE-X half duplex ability supported	RO	Note 53
6	10BASE-T Full Duplex This bit indicates the emulated link partner PHY 10BASE-T full duplex capability. 0: 10BASE-T full duplex ability not supported 1: 10BASE-T full duplex ability supported	RO	Note 53
5	10BASE-T Half Duplex This bit indicates the emulated link partner PHY 10BASE-T half duplex capability. 0: 10BASE-T half duplex ability not supported 1: 10BASE-T half duplex ability supported	RO	Note 53
4:0	Selector Field This field identifies the type of message being sent by Auto-Negotiation. 00001: IEEE 802.3	RO	00001b

Note 50: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 51: The emulated link partner does not support next page, always instantly sends its link code word, never sends a fault, and does not support 100BASE-T4.

Note 52: The emulated link partner's asymmetric/symmetric pause ability is based upon the values of the [Asymmetric Pause](#) and [Symmetric Pause](#) bits of the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#). Thus the emulated link partner always accommodates the request of the Virtual PHY, as shown in [Table 12-20](#).

The link partner pause ability bits are determined when Auto-Negotiation is complete. Changing the [Virtual PHY Auto-Negotiation Advertisement Register \(VPHY_AN_ADV\)](#) will have no affect until the Auto-Negotiation process is re-run.

If the local device advertises both Symmetric and Asymmetric Pause, the result is determined based on the FD_FC_strap_0 configuration strap. This allows the user the choice of network emulation. If FD_FC_strap_0 = 1, then the result is Symmetrical, else Asymmetrical. See [Section 12.3.1, "Virtual PHY Auto-Negotiation," on page 303](#) for additional information.

TABLE 12-20: EMULATED LINK PARTNER PAUSE FLOW CONTROL ABILITY DEFAULT VALUES

	VPHY Symmetric Pause (register 4.10)	VPHY Asymmetric Pause (register 4.11)	FD_FC_strap_0	Link Partner Symmetric Pause (register 5.10)	Link Partner Asymmetric Pause (register 5.11)
No Flow Control Enabled	0	0	x	0	0
Symmetric Pause	1	0	x	1	0

TABLE 12-20: EMULATED LINK PARTNER PAUSE FLOW CONTROL ABILITY DEFAULT VALUES

	VPHY Symmetric Pause (register 4.10)	VPHY Asymmetric Pause (register 4.11)	FD_FC_strap_0	Link Partner Symmetric Pause (register 5.10)	Link Partner Asymmetric Pause (register 5.11)
Asymmetric Pause Towards Switch	0	1	x	1	1
Asymmetric Pause Towards MAC	1	1	0	0	1
Symmetric Pause	1	1	1	1	1

Note 53: The emulated link partner always has the following capabilities: 100BASE-X full duplex, 100BASE-X half duplex, 10BASE-T full duplex, and 10BASE-T half duplex. For more information on the Virtual PHY Auto-Negotiation, see [Section 12.3.1, "Virtual PHY Auto-Negotiation," on page 303](#).

12.3.3.7 Virtual PHY Auto-Negotiation Expansion Register (VPHY_AN_EXP)

Offset: 1D8h Size: 32 bits
Index (decimal): 6 16 bits

This register is used in the Auto-Negotiation process.

Bits	Description	Type	Default
31:16	RESERVED (See Note 54)	RO	-
15:5	RESERVED	RO	-
4	Parallel Detection Fault This bit indicates whether a Parallel Detection Fault has been detected. This bit is always 0. 0: A fault hasn't been detected via the Parallel Detection function 1: A fault has been detected via the Parallel Detection function	RO	0b Note 55
3	Link Partner Next Page Able This bit indicates whether the link partner has next page ability. This bit is always 0. 0: Link partner does not contain next page capability 1: Link partner contains next page capability	RO	0b Note 56
2	Local Device Next Page Able This bit indicates whether the local device has next page ability. This bit is always 0. 0: Local device does not contain next page capability 1: Local device contains next page capability	RO	0b Note 56
1	Page Received This bit indicates the reception of a new page. 0: A new page has not been received 1: A new page has been received	RO/LH	1b Note 57
0	Link Partner Auto-Negotiation Able This bit indicates the Auto-Negotiation ability of the link partner. 0: Link partner is not Auto-Negotiation able 1: Link partner is Auto-Negotiation able	RO	1b Note 58

Note 54: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 55: Since the Virtual PHY link partner is emulated, there is never a Parallel Detection Fault and this bit is always 0.

Note 56: Next page ability is not supported by the Virtual PHY or emulated link partner.

Note 57: The [Page Received](#) bit is clear when read. It is first cleared on reset, but set shortly thereafter when the Auto-Negotiation process is run.

Note 58: The emulated link partner will show Auto-Negotiation able unless Auto-Negotiation fails (no common bits between the advertised ability and the link partner ability).

12.3.3.8 Virtual PHY Special Control/Status Register (VPHY_SPECIAL_CONTROL_STATUS)

Offset: 1DCh Size: 32 bits
Index (decimal): 31 16 bits

This read/write register contains a current link speed/duplex indicator and SQE control.

Bits	Description	Type	Default
31:16	RESERVED (See Note 59)	RO	-
15	RESERVED	RO	-
14	Switch Loopback When set, transmissions from the switch fabric MAC are not sent to the Host MAC. Instead, they are looped back into the switch engine. From the MAC viewpoint, this is effectively a FAR LOOPBACK. If loopback is enabled during half-duplex operation, then the Enable Receive Own Transmit bit in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) must be set for the port. Otherwise, the switch fabric will ignore receive activity when transmitting in half-duplex mode. Note: This mode works even if the Isolate (VPHY_ISO) bit of the Virtual PHY Basic Control Register (VPHY_BASIC_CTRL) is set.	R/W	0b
13:11	RESERVED	RO	-
10	Turbo Mode Enable When set, this bit changes the 100 Mbps data rate to 200 Mbps. The normal Virtual PHY selection mechanism that chooses between 10 and 100 Mbps will instead choose between 10 Mbps and 200 Mbps.	R/W	Note 60
9:8	RESERVED	RO	-
7	Switch Collision Test When set, the collision signal to the switch fabric is active during transmission from the switch engine. Note: It is recommended that this bit be used only when using loopback mode.	R/W	0b
6:5	RESERVED	RO	-

Bits	Description	Type	Default																																													
4:2	Current Speed/Duplex Indication This field indicates the current speed and duplex of the Virtual PHY link.	RO	Note 61																																													
	<table><tr><td>[4]</td><td>[3]</td><td>[2]</td><td>Speed</td><td>Duplex</td></tr><tr><td>0</td><td>0</td><td>0</td><td colspan="2">RESERVED</td></tr><tr><td>0</td><td>0</td><td>1</td><td>10Mbps</td><td>half-duplex</td></tr><tr><td>0</td><td>1</td><td>0</td><td>100/200Mbps</td><td>half-duplex</td></tr><tr><td>0</td><td>1</td><td>1</td><td colspan="2">RESERVED</td></tr><tr><td>1</td><td>0</td><td>0</td><td colspan="2">RESERVED</td></tr><tr><td>1</td><td>0</td><td>1</td><td>10Mbps</td><td>full-duplex</td></tr><tr><td>1</td><td>1</td><td>0</td><td>100/200Mbps</td><td>full-duplex</td></tr><tr><td>1</td><td>1</td><td>1</td><td colspan="2">RESERVED</td></tr></table>			[4]	[3]	[2]	Speed	Duplex	0	0	0	RESERVED		0	0	1	10Mbps	half-duplex	0	1	0	100/200Mbps	half-duplex	0	1	1	RESERVED		1	0	0	RESERVED		1	0	1	10Mbps	full-duplex	1	1	0	100/200Mbps	full-duplex	1	1	1	RESERVED	
	[4]			[3]	[2]	Speed	Duplex																																									
	0			0	0	RESERVED																																										
	0			0	1	10Mbps	half-duplex																																									
	0			1	0	100/200Mbps	half-duplex																																									
	0			1	1	RESERVED																																										
	1			0	0	RESERVED																																										
	1			0	1	10Mbps	full-duplex																																									
	1			1	0	100/200Mbps	full-duplex																																									
1	1	1	RESERVED																																													
1	RESERVED	RO	-																																													
0	SQEOFF This bit enables/disables the Signal Quality Error (Heartbeat) test. 0: SQE test enabled 1: SQE test disabled	R/W NASR Note 62	Note 63																																													

Note 59: The reserved bits 31-16 are used to pad the register to 32-bits so that each register is on a DWORD boundary. When accessed serially (through the MII management protocol), the register is 16-bits wide.

Note 60: The default value of this field is a 0.

Note 61: The default value of this field is the result of the Auto-Negotiation process if the [Auto-Negotiation \(VPHY_AN\)](#) bit of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#) is set. Otherwise, this field reflects the [Speed Select LSB \(VPHY_SPEED_SEL_LSB\)](#) and [Duplex Mode \(VPHY_DUPLEX\)](#) bit settings of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#). Refer to [Section 12.3.1, "Virtual PHY Auto-Negotiation," on page 303](#) for information on the Auto-Negotiation determination process of the Virtual PHY.

Note 62: Register bits designated as NASR are reset when the Virtual PHY Reset is generated via the [Reset Control Register \(RESET_CTL\)](#). The NASR designation is only applicable when the [Reset \(VPHY_RST\)](#) bit of the [Virtual PHY Basic Control Register \(VPHY_BASIC_CTRL\)](#) is set.

Note 63: The default value of this field is determined via the [SQE_test_disable_strap_0](#) configuration strap. Refer to [Section 7.1, "Soft-Straps," on page 60](#) for additional information.

13.0 SWITCH FABRIC

13.1 Functional Overview

The Switch Fabric contains a 3-port VLAN layer 2 Switch Engine that supports untagged, VLAN tagged and priority tagged frames. The Switch Fabric provides an extensive feature set which includes spanning tree protocol support, multicast packet filtering and Quality of Service (QoS) packet prioritization by VLAN tag, destination address, port default value or DIFFSERV/TOS, allowing for a range of prioritization implementations. 32k of buffer RAM allows for the storage of multiple packets while forwarding operations are completed and a 512 entry forwarding table provides room for MAC address forwarding tables. Each port is allocated a cluster of 4 dynamic QoS queues which allow each queue size to grow and shrink with traffic, effectively utilizing all available memory. This memory is managed dynamically via the Buffer Manager block within the Switch Fabric. All aspects of the Switch Fabric are managed via the Switch Fabric configuration and status registers (CSR), which are indirectly accessible via the system control and status registers.

The Switch Fabric consists of these major blocks:

- [Switch Fabric Control and Status Registers](#) - These registers provide access to various Switch Fabric parameters for configuration and monitoring.
- [10/100 Ethernet MAC](#) - A total of three MACs are included in the Switch Fabric which provide basic 10/100 Ethernet functionality for each Switch Fabric port. Note: one port connects internally to the Host MAC.
- [Switch Engine \(SWE\)](#) - This block is the core of the Switch Fabric and provides VLAN layer 2 switching for all three switch ports.
- [Buffer Manager \(BM\)](#) - This block provides control of the free buffer space, transmit queues and scheduling.
- [Switch Fabric Interface Logic](#) - This block provides some auxiliary registers and interfaces the [Switch Fabric Control and Status Registers](#) to the rest of the device. It also enables the flow control functions based on various settings and port conditions.

Refer to [FIGURE 2-1: Internal Block Diagram on page 9](#) for details on the interconnection of the Switch Fabric blocks within the device.

13.2 10/100 Ethernet MAC

The Switch Fabric contains three 10/100 MAC blocks, one for each switch port (0,1,2). The 10/100 MAC provides the basic 10/100 Ethernet functionality, including transmission deferral and collision back-off/retry, receive/transmit FCS checking and generation, receive/transmit pause flow control and transmit back pressure. The 10/100 MAC also includes RX and TX FIFOs and per port statistic counters.

13.2.1 RECEIVE MAC

The receive MAC (IEEE 802.3) sublayer decomposes Ethernet packets acquired via the internal MII interface by stripping off the preamble sequence and Start of Frame Delimiter (SFD). The receive MAC checks the FCS, the MAC Control Type and the byte count against the drop conditions. The packet is stored in the RX FIFO as it is received.

The receive MAC determines the validity of each received packet by checking the Type field, FCS and oversize or undersize conditions. All bad packets will be either immediately dropped or marked (at the end) as bad packets.

Oversized packets are normally truncated at 1519 or 1523 (VLAN tagged) octets and marked as erroneous. The MAC can be configured to accept packets up to 2048 octets (inclusive), in which case the oversize packets are truncated at 2048 bytes and marked as erroneous.

Undersized packets are defined as packets with a length less than the minimum packet size. The minimum packet size is defined to be 64 bytes, exclusive of preamble sequence and SFD, regardless of the occurrence of a VLAN tag.

The FCS and length/type fields of the frame are checked to detect if the packet has a valid MAC control frame. When the MAC receives a MAC control frame with a valid FCS and determines the operation code is a pause command (Flow Control frame), the MAC will load its internal pause counter with the Number_of_Slots variable from the MAC control frame just received. Anytime the internal pause counter is zero, the transmit MAC will be allowed to transmit (XON). If the internal pause counter is not zero, the receive MAC will not allow the transmit MAC to transmit (XOFF). When the transmit MAC detects an XOFF condition it will continue to transmit the current packet, terminating transmission after the current packet has been transmitted until receiving the XON condition from the receive MAC. The pause counter will begin to decrement at the end of the current transmission or immediately if no transmission is underway. If another pause command is received while the transmitter is already in pause, the new pause time indicated by the Flow Control

packet will be loaded into the pause counter. The pause function is enabled by either Auto-Negotiation or manually as discussed in [Section 13.5.1, "Flow Control Enable Logic," on page 343](#). Pause frames are consumed by the MAC and are not sent to the Switch Engine. Non-pause control frames are optionally filtered or forwarded.

Note: To meet the IEEE 802.1 Filtering Database requirements, the MAC address of 01-80-C2-00-00-01 should be added into the ALR address table as filtering entries by either EEPROM sequence or by software.

When the receive FIFO is full and additional data continues to be received, an overrun condition occurs and the frame is discarded (FIFO space recovered) or marked as a bad frame.

The receive MAC can be disabled from receiving all frames by clearing the [RX Enable \(RXEN\)](#) bit of the [Port x MAC Receive Configuration Register \(MAC_RX_CFG_x\)](#).

For information on MAC EEE functionality, refer to [Section 13.2.3, "IEEE 802.3az Energy Efficient Ethernet," on page 322](#).

13.2.1.1 Receive Counters

The receive MAC gathers statistics on each packet and increments the related counter registers. The following receive counters are supported for each Switch Fabric port. Refer to [Table 13-9, "Indirectly Accessible Switch Control and Status Registers," on page 363](#) and [Section 13.7.2.3 through Section 13.7.2.22](#) for detailed descriptions of these counters.

- Total undersized packets ([Section 13.7.2.3, on page 378](#))
- Total packets 64 bytes in size ([Section 13.7.2.4, on page 378](#))
- Total packets 65 through 127 bytes in size ([Section 13.7.2.5, on page 379](#))
- Total packets 128 through 255 bytes in size ([Section 13.7.2.6, on page 379](#))
- Total packets 256 through 511 bytes in size ([Section 13.7.2.7, on page 380](#))
- Total packets 512 through 1023 bytes in size ([Section 13.7.2.8, on page 380](#))
- Total packets 1024 through maximum bytes in size ([Section 13.7.2.9, on page 381](#))
- Total oversized packets ([Section 13.7.2.10, on page 381](#))
- Total OK packets ([Section 13.7.2.11, on page 382](#))
- Total packets with CRC errors ([Section 13.7.2.12, on page 382](#))
- Total multicast packets ([Section 13.7.2.13, on page 383](#))
- Total broadcast packets ([Section 13.7.2.14, on page 383](#))
- Total MAC Pause packets ([Section 13.7.2.15, on page 384](#))
- Total fragment packets ([Section 13.7.2.16, on page 384](#))
- Total jabber packets ([Section 13.7.2.17, on page 385](#))
- Total alignment errors ([Section 13.7.2.18, on page 385](#))
- Total bytes received from all packets ([Section 13.7.2.19, on page 386](#))
- Total bytes received from good packets ([Section 13.7.2.20, on page 386](#))
- Total packets with a symbol error ([Section 13.7.2.21, on page 387](#))
- Total MAC control packets ([Section 13.7.2.22, on page 387](#))
- Total number of RX LPIs received ([Section 13.7.2.23, on page 388](#))
- Total time in RX LPI state ([Section 13.7.2.24, on page 388](#))

13.2.2 TRANSMIT MAC

The transmit MAC generates an Ethernet MAC frame from TX FIFO data. This includes generating the preamble and SFD, calculating and appending the frame checksum value, optionally padding undersize packets to meet the minimum packet requirement size (64 bytes) and maintaining a standard inter-frame gap time during transmit.

The transmit MAC can operate at 10/100Mbps, half or full-duplex and with or without flow control depending on the state of the transmission. In half-duplex mode, the transmit MAC meets CSMA/CD IEEE 802.3 requirements. The transmit MAC will re-transmit if collisions occur during the first 64 bytes (normal collisions) or will discard the packet if collisions occur after the first 64 bytes (late collisions). The transmit MAC follows the standard truncated binary exponential back-off algorithm, collision and jamming procedures.

The transmit MAC pre-pends the standard preamble and SFD to every packet from the FIFO. The transmit MAC also follows, as default, the standard Inter-Frame Gap (IFG). The default IFG is 96 bit times and can be adjusted via the [IFG Config](#) field of the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#).

Packet padding and cyclic redundant code (FCS) calculation may be optionally performed by the transmit MAC. The auto-padding process automatically adds enough zeros to packets shorter than 64 bytes. The auto-padding and FCS generation is controlled via the [TX Pad Enable](#) bit of the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#).

When in full-duplex mode, the transmit MAC uses the flow-control algorithm specified in IEEE 802.3. MAC pause frames are used primarily for flow control packets, which pass signaling information between stations. MAC pause frames have a unique type of 8808h and a pause op-code of 0001h. The MAC pause frame contains the pause value in the data field. The flow control manager will auto-adapt the procedure based on traffic volume and speed to avoid packet loss and unnecessary pause periods.

When in half-duplex mode, the MAC uses a back pressure algorithm. The back pressure algorithm is based on a forced collision and an aggressive back-off algorithm.

For information on MAC EEE functionality, refer to [Section 13.2.3, "IEEE 802.3az Energy Efficient Ethernet," on page 322](#).

13.2.2.1 Transmit Counters

The transmit MAC gathers statistics on each packet and increments the related counter registers. The following transmit counters are supported for each Switch Fabric port. Refer to [Table 13-9, "Indirectly Accessible Switch Control and Status Registers," on page 363](#) and [Section 13.7.2.29 through Section 13.7.2.46](#) for detailed descriptions of these counters.

- Total packets deferred ([Section 13.7.2.29, on page 393](#))
- Total pause packets ([Section 13.7.2.30, on page 393](#))
- Total OK packets ([Section 13.7.2.31, on page 394](#))
- Total packets 64 bytes in size ([Section 13.7.2.32, on page 394](#))
- Total packets 65 through 127 bytes in size ([Section 13.7.2.33, on page 395](#))
- Total packets 128 through 255 bytes in size ([Section 13.7.2.34, on page 395](#))
- Total packets 256 through 511 bytes in size ([Section 13.7.2.35, on page 396](#))
- Total packets 512 through 1023 bytes in size ([Section 13.7.2.36, on page 396](#))
- Total packets 1024 through maximum bytes in size ([Section 13.7.2.37, on page 397](#))
- Total undersized packets ([Section 13.7.2.38, on page 397](#))
- Total bytes transmitted from all packets ([Section 13.7.2.39, on page 398](#))
- Total broadcast packets ([Section 13.7.2.40, on page 398](#))
- Total multicast packets ([Section 13.7.2.41, on page 399](#))
- Total packets with a late collision ([Section 13.7.2.42, on page 399](#))
- Total packets with excessive collisions ([Section 13.7.2.43, on page 399](#))
- Total packets with a single collision ([Section 13.7.2.44, on page 400](#))
- Total packets with multiple collisions ([Section 13.7.2.45, on page 400](#))
- Total collision count ([Section 13.7.2.46, on page 400](#))
- Total number of TX LPIs Generated ([Section 13.7.2.47, on page 401](#))
- Total time in TX LPI state ([Section 13.7.2.48, on page 401](#))

13.2.3 IEEE 802.3AZ ENERGY EFFICIENT ETHERNET

The device supports Energy Efficient Ethernet (EEE) in 100 Mbps mode as defined in the most recent version of the IEEE 802.3az standard. EEE functions are not used on Port 0 since this port is connected internally to the Host MAC.

13.2.3.1 TX LPI Generation

The process of when the MAC should indicate LPI requests to the PHY is divided into two sections:

- [CLIENT LPI REQUESTS TO MAC](#)
- [MAC LPI REQUEST TO PHY](#)

CLIENT LPI REQUESTS TO MAC

When the TX FIFO is empty for a time (in microseconds) specified in [Port x EEE TX LPI Request Delay Register \(EEE_TX_LPI_REQ_DELAY_x\)](#), a TX LPI request is asserted to the MAC. A setting of 0 us is possible for this time. If the TX FIFO becomes not empty while the timer is running, the timer is reset (i.e. empty time is not cumulative). Once TX LPI is requested and the TX FIFO becomes not empty, the TX LPI request is negated.

The TX FIFO empty timer is reset if [Energy Efficient Ethernet \(EEE_ENABLE\)](#) in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is cleared.

TX LPI requests are asserted only if the [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bit is set, and when appropriate, if the current speed is 100 Mbps, the current duplex is full and the auto-negotiation result indicates that both the local and partner device support EEE 100 Mbps. In order to prevent an unstable link condition, the PHY link status also must indicate “up” for one second before LPI is requested.

These tests for the allowance of TX LPI are done in the [Switch Fabric Interface Logic](#) block. See [Section 13.5.2, “EEE Enable Logic,” on page 345](#) for further details.

TX LPI requests are asserted even if the [TX Enable \(TXEN\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is cleared.

MAC LPI REQUEST TO PHY

Lower Power Idle (LPI) is requested by the MAC to the PHY using the MII value of TXEN=0, TXER=1, TXD[3:0]=4'b0001.

The MAC always finishes the current packet before signaling TX LPI to the PHY.

The MAC will generate TX LPI requests to the PHY even if the [TX Enable \(TXEN\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is cleared.

802.3az specifies the usage of a simplified full duplex MAC with carrier sense deferral. Basically this means that once the TX LPI request to the PHY is de-asserted, the MAC will defer the time specified in [Port x EEE Time Wait TX System Register \(EEE_TW_TX_SYS_x\)](#) in addition to the normal IPG before sending a frame.

TX LPI COUNTERS

The MAC maintains a counter, [EEE TX LPI Transitions](#), that counts the number of times that TX LPI request to the PHY changes from de-asserted to asserted. The counter is not writable and does not clear on read. The counter is reset if the [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is low.

The MAC maintains a counter, [EEE TX LPI Time](#), that counts (in microseconds) the amount of time that TX LPI request to the PHY is asserted. Note that this counter does not include the time specified in the [Port x EEE Time Wait TX System Register \(EEE_TW_TX_SYS_x\)](#). The counter is not writable and does not clear on read. The counter is reset if the [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is low.

13.2.3.2 RX LPI Detection

Receive Lower Power Idle (LPI) is indicated by the PHY to the MAC using the MII value of RXDV=0, RXER=1, TXD[3:0]=4'b0001.

DECODING LPI

The MAC will decode the LPI indication only when [Energy Efficient Ethernet \(EEE_ENABLE\)](#) is set in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#), and when appropriate, the current speed is 100Mbps, the current duplex is full and the auto-negotiation result indicates that both the local and partner device supports EEE at 100Mbps. In order to prevent an unstable link condition, the PHY link status also must indicate “up” for one second before LPI is decoded.

These tests for the allowance of TX LPI are done in the [Switch Fabric Interface Logic](#) block. See [Section 13.5.2, “EEE Enable Logic,” on page 345](#) for further details.

The MAC will decode the LPI indication even if [RX Enable \(RXEN\)](#) in the [Port x MAC Receive Configuration Register \(MAC_RX_CFG_x\)](#) is cleared.

RX LPI COUNTERS

LAN9352

The MAC maintains a counter, [EEE RX LPI Transitions](#), that counts the number of times that the LPI indication from the PHY changes from de-asserted to asserted. The counter is not writable and does not clear on read. The counter is reset if the [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is low.

The MAC maintains a counter, [EEE RX LPI Time](#), that counts (in microseconds) the amount of time that the PHY indicates LPI. The counter is not writable and does not clear on read. The counter is reset if the [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bit in the [Port x MAC Transmit Configuration Register \(MAC_TX_CFG_x\)](#) is low.

13.2.4 HOST MAC FLOW CONTROL

As described in [Section 13.2.1, "Receive MAC"](#), pause frames from the Host MAC will temporarily pause the switch fabric Port 0 to Host MAC operation until the internal pause timer expires. In addition to this, an out of band flow control signal is connected directly between port 0 and the Host MAC's receive buffer status. This may be used for a low latency indication and is enabled via the [Port 0 Hard-wired Flow Control \(HW_FC_0\)](#) field in the [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#) and configured via the [Automatic Flow Control High Level \(AFC_HI\)](#) and [Automatic Flow Control Low Level \(AFC_LO\)](#) fields in the [Host MAC Automatic Flow Control Configuration Register \(AFC_CFG\)](#).

In the reverse direction (Host MAC to port 0), a similar out of band flow control signal is connected directly between the [Buffer Manager \(BM\)](#) the Host MAC's transmitter. This is likewise enabled via the [Port 0 Hard-wired Flow Control \(HW_FC_0\)](#) field and is configured via the [Pause Level Low](#) and [Pause Level High](#) fields in the [Buffer Manager Flow Control Pause Level Register \(BM_FC_PAUSE_LVL\)](#) and the [Resume Level Low](#) and [Resume Level High](#) fields in the [Buffer Manager Flow Control Resume Level Register \(BM_FC_RESUME_LVL\)](#).

Note: This out of band flow control is independent of the flow control enable discussed in [Section 13.5.1, "Flow Control Enable Logic,"](#) on page 343.

13.3 Switch Engine (SWE)

The Switch Engine (SWE) is a VLAN layer 2 (link layer) switching engine supporting 3 ports. The SWE supports the following types of frame formats: untagged frames, VLAN tagged frames and priority tagged frames. The SWE supports both the 802.3 and Ethernet II frame formats.

The SWE provides the control for all forwarding/filtering rules. It handles the address learning and aging and the destination port resolution based upon the MAC address and VLAN of the packet. The SWE implements the standard bridge port states for spanning tree and provides packet metering for input rate control. It also implements port mirroring, broadcast throttling and multicast pruning and filtering. Packet priorities are supported based on the IPv4 TOS bits and IPv6 Traffic Class bits using a DIFFSERV Table mapping, the non-DIFFSERV mapped IPv4 precedence bits, VLAN priority using a per port Priority Regeneration Table, DA based static priority and Traffic Class mapping to one of 4 QoS transmit priority queues.

The following sections detail the various features of the Switch Engine.

13.3.1 MAC ADDRESS LOOKUP TABLE

The Address Logic Resolution (ALR) maintains a 512 entry MAC Address Table. The ALR searches the table for the destination MAC address. If the search finds a match, the associated data is returned indicating the destination port or ports, whether to filter the packet, the packet's priority (used if enabled) and whether to override the ingress and egress spanning tree port state. [Figure 13-1](#) displays the ALR table entry structure. Refer to the [Switch Engine ALR Write Data 0 Register \(SWE_ALR_WR_DAT_0\)](#) and the [Switch Engine ALR Write Data 1 Register \(SWE_ALR_WR_DAT_1\)](#) for detailed descriptions of these bits.

FIGURE 13-1: ALR TABLE ENTRY STRUCTURE

Bit	58	57	56	55	54	53	52	51	50	49	48	47	0
	Valid	Age 1 / Override	Static	Age 0 / Filter	Priority Enable	Priority			Port			MAC Address	

13.3.1.1 Learning/Aging/Migration

The ALR adds new MAC addresses upon ingress along with the associated receive port.

If the source MAC address already exists, the entry is refreshed. This action serves two purposes. First, if the source port has changed due to a network reconfiguration (migration), it is updated. Second, each instance the entry is refreshed, the age status bits are set, keeping the entry active. Learning can be disabled per port via the [Enable Learning on Ingress](#) field of the [Switch Engine Port Ingress Configuration Register \(SWE_PORT INGRSS_CFG\)](#).

During each aging period, the ALR scans the learned MAC addresses. For entries which have an age status greater than 0, the ALR decrements the age. As mentioned above, if a MAC address is subsequently refreshed, the age status bits will be set again and the process would repeat. If a learned entry already had its age status bits decremented to 0 (by previous scans), the ALR will instead remove the learned entry. Four scans need to occur for a MAC address to be aged and removed. Since the first scan could occur immediately following the add or refresh of an entry, an entry will be aged and removed after a minimum of 3 age periods and a maximum of 4 age periods.

The minimum aging time is programmable using the [Aging Time](#) field of the [Switch Engine ALR Configuration Register \(SWE_ALR_CFG\)](#) in 1 second increments from 1 second to approximately 69 minutes. The maximum aging time is 33% higher.

The [ALR Age Test](#) bit in the [Switch Engine ALR Configuration Register \(SWE_ALR_CFG\)](#) changes the [Aging Time](#) from seconds to milliseconds.

Aging can be disabled by clearing the [ALR Age Enable](#) field of the [Switch Engine ALR Configuration Register \(SWE_ALR_CFG\)](#).

13.3.1.2 Static Entries

If a MAC address entry is manually added by the host CPU, it can be (and typically is) marked as [Static](#). Static entries are not subjected to the aging process. Static entries also cannot be changed by the learning process (including migration).

13.3.1.3 Multicast Pruning

The destination port that is returned as a result of a destination MAC address lookup may be a single port or any combination of ports. The latter is used to setup multicast address groups. An entry with a multicast MAC address would be entered manually by the host CPU with the appropriate destination port(s). Typically, the [Static](#) bit should also be set to prevent automatic aging of the entry.

13.3.1.4 Broadcast Entries

If desired, the host CPU can manually add the broadcast address of 0xFFFFFFFF. This feature is enabled by setting the [Allow Broadcast Entries](#) field of the [Switch Engine ALR Configuration Register \(SWE_ALR_CFG\)](#). Typically, the [Static](#) bit should also be set in the ALR entry to prevent automatic aging of the entry.

13.3.1.5 Address Filtering

Filtering can be performed on a destination MAC address. Such an entry would be entered manually by the host CPU with the Filter bit active. Typically, the [Static](#) bit should also be set to prevent automatic aging of the entry.

Note: To meet the IEEE 802.1 Filtering Database requirements, the MAC addresses of 01-80-C2-00-00-01 through 01-80-C2-00-00-0F should be added into the ALR address table as filtering entries by either EEPROM sequence or by software. The MAC address of 01-80-C2-00-00-00 is typically added as a forwarding entry to direct BPDU frames to the host CPU.

13.3.1.6 Spanning Tree Port State Override

A special spanning tree port state override setting can be applied to MAC address entries. When the host CPU manually adds an entry with both the [Static](#) and [Age 1/Override](#) bits set, packets with a matching destination address will bypass the spanning tree port state (except the Disabled state) and will be forwarded. This feature is typically used to allow the reception of the BPDU packets while a port is in the non-forwarding state. Refer to [Section 13.3.5, "Spanning Tree Support," on page 331](#) for additional details.

13.3.1.7 MAC Destination Address Lookup Priority

If enabled, globally via the [DA Highest Priority](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#) along with the, per entry, Priority Enable bit, the transmit priority for MAC address entries is taken from the associated data of that entry.

13.3.1.8 ALR Result Override

Results from the ALR Destination MAC lookup can be overridden on a per port basis. This feature is enabled by setting the appropriate ALR Override Enable bit in the [Switch Engine ALR Override Register \(SWE_ALR_OVERRIDE\)](#). When enabled, the destination port from the ALR Destination MAC Address lookup is replaced with the appropriate ALR Override Destination field in the [Switch Engine ALR Override Register \(SWE_ALR_OVERRIDE\)](#).

The ALR Spanning Tree Override, Static, Filter and Priority results for the Destination MAC Address are still used.

Note: Forwarding rules described in [Section 13.3.2](#) are still followed.

13.3.1.9 Host Access

Note: Refer to [Section 13.7.3.1, on page 403](#) through [Section 13.7.3.6, on page 410](#) for detailed definitions of the registers.

ADD, DELETE, AND MODIFY ENTRIES

The ALR contains a learning engine that is used by the host CPU to add, delete and modify the MAC Address Table. This engine is accessed by using the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#), the [Switch Engine ALR Command Status Register \(SWE_ALR_CMD_STS\)](#), the [Switch Engine ALR Write Data 0 Register \(SWE_ALR_WR_DAT_0\)](#) and the [Switch Engine ALR Write Data 1 Register \(SWE_ALR_WR_DAT_1\)](#).

The following procedure should be followed in order to add, delete and modify the ALR entries:

1. Write the [Switch Engine ALR Write Data 0 Register \(SWE_ALR_WR_DAT_0\)](#) and the [Switch Engine ALR Write Data 1 Register \(SWE_ALR_WR_DAT_1\)](#) with the desired MAC address and control bits.
An entry can be deleted by setting the [Valid](#) bit to 0.
2. Set the [Make Entry](#) bit in the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#).
3. Poll the [Operation Pending](#) bit in the [Switch Engine ALR Command Status Register \(SWE_ALR_CMD_STS\)](#) until it is cleared.

READ ENTRIES

The ALR contains a search engine that is used by the host to read the MAC Address Table. This engine is accessed by using the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#), the [Switch Engine ALR Command Status Register \(SWE_ALR_CMD_STS\)](#), [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) and the [Switch Engine ALR Read Data 1 Register \(SWE_ALR_RD_DAT_1\)](#).

Note: The entries read are not necessarily in the same order as they were learned or manually added.

The following procedure should be followed in order to read the ALR entries:

1. Set the [Get First Entry](#) bit in the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#).
2. Poll the [Operation Pending](#) bit in the [Switch Engine ALR Command Status Register \(SWE_ALR_CMD_STS\)](#) until it is cleared.
3. If the [Valid](#) bit in the [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) is set, then the entry is valid and the data from the [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) and the [Switch Engine ALR Read Data 1 Register \(SWE_ALR_RD_DAT_1\)](#) can be stored.
4. If the [End of Table](#) bit in the [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) is set, then exit.
5. Set the [Get Next Entry](#) bit in the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#).
6. Go to step 3.

13.3.2 FORWARDING RULES

Upon ingress, packets are filtered or forwarded based on the following rules:

- If the destination port equals the source port (local traffic), the packet is filtered. (This rule is for a destination MAC address which is found in the ALR table and the ALR result indicates a single destination port.)
- If the source port is in the Disabled state, via the [Switch Engine Port State Register \(SWE_PORT_STATE\)](#), the packet is filtered.
- If the source port is in the Learning or Listening / Blocking state, via the [Switch Engine Port State Register \(SWE_PORT_STATE\)](#), the packet is filtered (unless the [Spanning Tree Port State Override](#) is in effect).
- If the packet is a multicast packet and it is identified as a IGMP or MLD packet and IGMP/MLD monitoring is enabled (respectively), via the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL_INGRESS_CFG\)](#), the packet is redirected to the IGMP/MLD monitor port(s). This check is not done on special tagged packets from the host CPU port when an ALR lookup is not requested. Refer to [Section 13.3.10.1, "Packets from](#)

the Host CPU," on page 337 for additional information.

- If the destination port is in the disabled state, via the [Switch Engine Port State Register \(SWE_PORT_STATE\)](#), the packet is filtered. (This rule is for a destination MAC address which is found in the ALR table and the ALR result indicates a single destination port. When there are multiple destination ports or when the MAC address is not found, the packet is sent to only those ports that are in the Forwarding state. This rule is also suppressed if ALR Result Override is enabled.)
- If the destination port is in the Learning or Listening / Blocking state, via the [Switch Engine Port State Register \(SWE_PORT_STATE\)](#), the packet is filtered (unless the [Spanning Tree Port State Override](#) is in effect). (This rule is for a destination MAC address which is found in the ALR table and the ALR result indicates a single destination port. When there are multiple destination ports or when the MAC address is not found, the packet is sent to only those ports that are in the Forwarding state. This rule is also suppressed if ALR Result Override is enabled.)
- If the [Age 0/Filter](#) bit for the Destination Address is set in the ALR table, the packet is filtered.
- If the packet has a unicast destination MAC address which is not found in the ALR table and the [Drop Unknown](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#) is set, the packet is filtered.
- If the packet has a multicast destination MAC address which is not found in the ALR table and the [Filter Multicast](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#) is set, the packet is filtered.
- If the packet has a broadcast destination MAC address and the Broadcast Storm Control level has been reached, the packet is discarded.
- If the [Drop on Yellow](#) field in the [Buffer Manager Configuration Register \(BM_CFG\)](#) is set, the packet is colored Yellow and randomly selected, it is discarded.
- If the [Drop on Red](#) field in the [Buffer Manager Configuration Register \(BM_CFG\)](#) is set and the packet is colored Red, it is discarded.
- If the destination address was not found in the ALR table (an unknown or a broadcast) and the Broadcast Buffer Level is exceeded, the packet is discarded.
- If there is insufficient buffer space, the packet is discarded.
- If the destination address was not found in the ALR table (an unknown or a broadcast) or the destination address was found in the ALR table with the ALR result indicating multiple destination ports and the port forward states resulted in zero valid destination ports, the packet is filtered.
- For cases where the packet is not filtered, the ALR Override Enable bit in the [Switch Engine ALR Override Register \(SWE_ALR_OVERRIDE\)](#) is checked for the source port and, if set, the packet is redirected to the specified override destination.

When the switch is enabled for VLAN support, these following rules also apply:

- If the packet is untagged or priority tagged and the [Admit Only VLAN](#) field in the [Switch Engine Admit Only VLAN Register \(SWE_ADMT_ONLY_VLAN\)](#) for the ingress port is set, the packet is filtered.
- If the packet is tagged and has a VID equal to FFFh, it is filtered.
- If [Enable Membership Checking](#) field in the [Switch Engine Port Ingress Configuration Register \(SWE_PORT INGRSS_CFG\)](#) is set, [Admit Non Member](#) field in the [Switch Engine Admit Non Member Register \(SWE_ADMT_N_MEMBER\)](#) is cleared and the source port is not a member of the incoming VLAN, the packet is filtered.
- If [Enable Membership Checking](#) field is set and the destination port is not a member of the incoming VLAN, the packet is filtered. (This rule is for a destination MAC address which is found in the ALR table and the ALR result indicates a single destination port. When there are multiple destination ports or when the MAC address is not found, the packet is sent to only those ports that are members of the VLAN. This rule is also suppressed if ALR Result Override is enabled.)
- If the destination address was not found in the ALR table (an unknown or broadcast) or the destination address was found in the ALR table with the ALR result indicating multiple destination ports and the VLAN broadcast domain containment resulted in zero valid destination ports, the packet is filtered.
- For the last three cases, if the VID is not in the VLAN table, the VLAN is considered foreign and the membership result is NULL. A NULL membership will result in the packet being filtered if [Enable Membership Checking](#) is set. A NULL membership will also result in the packet being filtered if the destination address is not found in the ALR table (since the packet would have no destinations).


13.3.3 TRANSMIT PRIORITY QUEUE SELECTION

The transmit priority queue may be selected from five options. As shown in [Figure 13-2](#), the priority may be based on:

- The static value for the destination address in the ALR table
- The precedence bits in the IPv4 TOS octet
- The DIFFSERV mapping table indexed by the IPv4 TOS octet or the IPv6 Traffic Class octet
- The VLAN tag priority field using the per port Priority Regeneration table
- The port default with separate values for packets with or without a broadcast destination address.


All options are sent through the Traffic Class table which maps the selected priority to one of the four output queues.

FIGURE 13-2: SWITCH ENGINE TRANSMIT QUEUE SELECTION



The transmit queue priority is based on the packet type and device configuration as shown in Figure 13-3. Refer to [Section 13.7.3.17, "Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\),"](#) on page 421 for definitions of the configuration bits.

FIGURE 13-3: SWITCH ENGINE TRANSMIT QUEUE CALCULATION



13.3.3.1 Port Default Priority

As detailed in [Figure 13-3](#), the default priority is based on the ingress port's priority bits in its port VID value. Separate values exist for packets with or without a broadcast destination address. The PVID table is read and written by using the [Switch Engine VLAN Command Register \(SWE_VLAN_CMD\)](#), the [Switch Engine VLAN Write Data Register \(SWE_VLAN_WR_DATA\)](#), the [Switch Engine VLAN Read Data Register \(SWE_VLAN_RD_DATA\)](#) and the [Switch Engine VLAN Command Status Register \(SWE_VLAN_CMD_STS\)](#). Refer to [Section 13.7.3.9, on page 414](#) through [Section 13.7.3.12, on page 419](#) for detailed VLAN register descriptions.

13.3.3.2 IP Precedence Based Priority

The transmit priority queue can be chosen based on the Precedence bits of the IPv4 TOS octet. This is supported for tagged and non-tagged packets for both type field and length field encapsulations. The Precedence bits are the three most significant bits of the IPv4 TOS octet.

13.3.3.3 DIFFSERV Based Priority

The transmit priority queue can be chosen based on the DIFFSERV usage of the IPv4 TOS or IPv6 Traffic Class octet. This is supported for tagged and non-tagged packets for both type field and length field encapsulations.

The DIFFSERV table is used to determine the packet priority from the 6-bit Differentiated Services (DS) field. The DS field is defined as the six most significant bits of the IPv4 TOS octet or the IPv6 Traffic Class octet and is used as an index into the DIFFSERV table. The output of the DIFFSERV table is then used as the priority. This priority is then passed through the Traffic Class table to select the transmit priority queue.

Note: The DIFFSERV table is not initialized upon reset or power-up. If DIFFSERV is enabled, then the full table must be initialized by the host.

The DIFFSERV table is read and written by using the [Switch Engine DIFFSERV Table Command Register \(SWE_DIFFSERV_TBL_CFG\)](#), the [Switch Engine DIFFSERV Table Write Data Register \(SWE_DIFFSERV_TBL_WR_DATA\)](#), the [Switch Engine DIFFSERV Table Read Data Register \(SWE_DIFFSERV_TBL_RD_DATA\)](#) and the [Switch Engine DIFFSERV Table Command Status Register \(SWE_DIFFSERV_TBL_CMD_STS\)](#). Refer to [Section 13.7.3.13, on page 419](#) through [Section 13.7.3.16, on page 420](#) for detailed DIFFSERV register descriptions.

13.3.3.4 VLAN Priority


As detailed in [Figure 13-3](#), the transmit priority queue can be taken from the priority field of the VLAN tag. The VLAN priority is sent through a per port Priority Regeneration table, which is used to map the VLAN priority into a user defined priority.

The Priority Regeneration table is programmed by using the [Switch Engine Port 0 Ingress VLAN Priority Regeneration Table Register \(SWE INGRSS_REGEN_TBL_0\)](#), the [Switch Engine Port 1 Ingress VLAN Priority Regeneration Table Register \(SWE INGRSS_REGEN_TBL_1\)](#) and the [Switch Engine Port 2 Ingress VLAN Priority Regeneration Table Register \(SWE INGRSS_REGEN_TBL_2\)](#). Refer to [Section 13.7.3.34, on page 434](#) through [Section 13.7.3.36, on page 436](#) for detailed descriptions of these registers.

13.3.4 VLAN SUPPORT

The Switch Engine supports 16 active VLANs out of a possible 4096. The VLAN table contains the 16 active VLAN entries, each consisting of the VID, the port membership and un-tagging instructions.

FIGURE 13-4: VLAN TABLE ENTRY STRUCTURE



On ingress, if a packet has a VLAN tag containing a valid VID (not 000h or FFFh), the VID table is searched. If the VID is found, the VLAN is considered active and the membership and un-tag instruction is used. If the VID is not found, the VLAN is considered foreign and the membership result is NULL. A NULL membership will result in the packet being filtered if [Enable Membership Checking](#) is set. A NULL membership will also result in the packet being filtered if the destination address is not found in the ALR table (since the packet would have no destinations).

On ingress, if a packet does not have a VLAN tag or if the VLAN tag contains VID with a value of 0 (priority tag), the packet is assigned a VLAN based on the Port Default VID (PVID) and Priority. The PVID is then used to access the above VLAN table. The usage of the PVID can be forced by setting the [802.1Q VLAN Disable](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL_INGRSS_CFG\)](#), in effect creating port based VLANs.

The VLAN membership of the packet is used for ingress and egress checking and for VLAN broadcast domain containment. The un-tag instructions are used at egress on ports defined as hybrid ports.

Refer to [Section 13.7.3.9, on page 414](#) through [Section 13.7.3.12, on page 419](#) for detailed VLAN register descriptions.

13.3.5 SPANNING TREE SUPPORT

Hardware support for the Spanning Tree Protocol (STP) and the Rapid Spanning Tree Protocol (RSTP) includes a per port state register as well as the override bit in the MAC Address Table entries ([Section 13.3.1.6, on page 325](#)) and the host CPU port special tagging ([Section 13.3.10, on page 337](#)).

The [Switch Engine Port State Register \(SWE_PORT_STATE\)](#) is used to place a port into one of the modes as shown in [Table 13-1](#). Normally only Port 1 and Port 2 are placed into modes other than forwarding. Port 0, which is connected to the host CPU, should normally be left in forwarding mode.

TABLE 13-1: SPANNING TREE STATES

Port State	Hardware Action	Software Action
11 - Disabled	<p>Received packets on the port are always discarded.</p> <p>Transmissions to the port are always blocked.</p> <p>Learning on the port is disabled.</p>	<p>The host CPU may attempt to send packets to the port in this state, but they will not be transmitted.</p>
01 - Blocking	<p>Received packets on the port are discarded unless overridden.</p> <p>Transmissions to the port are blocked unless overridden.</p> <p>Learning on the port is disabled.</p>	<p>The MAC Address Table should be programmed with entries that the host CPU needs to receive (e.g., the BPDU address). The Static and Age 1/Override bits should be set.</p> <p>The host CPU may send packets to the port in this state. Only packets with STP override will be transmitted.</p> <p>There is no hardware distinction between the Blocking and Listening states.</p>
01 - Listening	<p>Received packets on the port are discarded unless overridden.</p> <p>Transmissions to the port are blocked unless overridden.</p> <p>Learning on the port is disabled.</p>	<p>The MAC Address Table should be programmed with entries that the host CPU needs to receive (e.g., the BPDU address). The Static and Age 1/Override bits should be set.</p> <p>The host CPU may send packets to the port in this state. Only packets with STP override will be transmitted.</p>

TABLE 13-1: SPANNING TREE STATES (CONTINUED)

Port State	Hardware Action	Software Action
10 - Learning	<p>Received packets on the port are discarded unless overridden.</p> <p>Transmissions to the port are blocked unless overridden.</p> <p>Learning on the port is enabled.</p>	<p>The MAC Address Table should be programmed with entries that the host CPU needs to receive (e.g., the BPDU address). The Static and Age 1/Override bits should be set.</p> <p>The host CPU may send packets to the port in this state. Only packets with STP override will be transmitted.</p>
00 - Forwarding	<p>Received packets on the port are forwarded normally.</p> <p>Transmissions to the port are sent normally.</p> <p>Learning on the port is enabled.</p>	<p>The MAC Address Table should be programmed with entries that the host CPU needs to receive (e.g., the BPDU address). The Static and Age 1/Override bits should be set.</p> <p>The host CPU may send packets to the port in this state.</p>

13.3.6 INGRESS FLOW METERING AND COLORING

Hardware ingress rate limiting is supported by metering packet streams and marking packets as either Green, Yellow or Red according to three traffic parameters: [Committed Information Rate \(CIR\)](#), [Committed Burst Size \(CBS\)](#) and [Excess Burst Size \(EBS\)](#). A packet is marked Green if it does not exceed the CBS, Yellow if it exceeds to CBS but not the EBS or Red otherwise.

Ingress flow metering and coloring is enabled via the [Ingress Rate Enable](#) field in the [Switch Engine Ingress Rate Configuration Register \(SWE_INGRSS_RATE_CFG\)](#). Once enabled, each incoming packet is classified into a stream. Streams are defined as per port (3 streams), per priority (8 streams) or per port & priority (24 streams) as selected via the [Rate Mode](#) field in the [Switch Engine Ingress Rate Configuration Register \(SWE_INGRSS_RATE_CFG\)](#). Each stream can have a different CIR setting. All streams share common CBS and EBS settings. CIR, CBS and EBS are programmed via the [Switch Engine Ingress Rate Command Register \(SWE_INGRSS_RATE_CMD\)](#) and the [Switch Engine Ingress Rate Write Data Register \(SWE_INGRSS_RATE_WR_DATA\)](#).

Each stream is metered according to RFC 2697. At the rate set by the CIR, two token buckets are credited per stream. First, the Committed Burst bucket is incremented up to the maximum set by the CBS. Once the Committed Burst bucket is full, the Excess Burst bucket is incremented up to the maximum set by the EBS. The CIR rate is specified in time per byte. The value programmed is in approximately 20 ns per byte increments. Typical values are listed in [Table 13-2](#). When a port is receiving at 10 Mbps, any setting faster than 39 has the effect of not limiting the rate.

TABLE 13-2: TYPICAL INGRESS RATE SETTINGS

CIR Setting	Time Per Byte	Bandwidth
0-3	80 ns	100 Mbps
4	100 ns	80 Mbps
5	120 ns	67 Mbps
6	140 ns	57 Mbps
7	160 ns	50 Mbps
9	200 ns	40 Mbps
12	260 ns	31 Mbps
19	400 ns	20 Mbps
39	800 ns	10 Mbps

TABLE 13-2: TYPICAL INGRESS RATE SETTINGS

CIR Setting	Time Per Byte	Bandwidth
79	1600 ns	5 Mbps
160	3220 ns	2.5 Mbps
402	8060 ns	1 Mbps
804	16100 ns	500 kbps
1610	32220 ns	250 kbps
4028	80580 ns	100 kbps
8056	161140 ns	50 kbps

After each packet is received, the bucket is decremented. If the Committed Burst bucket has sufficient tokens, it is debited and the packet is colored Green. If the Committed Burst bucket lacks sufficient tokens for the packet, the Excess Burst bucket is checked. If the Excess Burst bucket has sufficient tokens, it is debited, the packet is colored Yellow and is subjected to random discard. If the Excess Burst bucket lacks sufficient tokens for the packet, the packet is colored Red and is discarded.

Note: All of the token buckets are initialized to the default value of 1536. If lower values are programmed into the CBS and EBS parameters, the token buckets will need to be normally depleted below these values before the values have any effect on limiting the maximum value of the token buckets.


Refer to [Section 13.7.3.26, on page 429](#) through [Section 13.7.3.30, on page 432](#) for detailed register descriptions.

13.3.6.1 Ingress Flow Calculation

Based on the flow monitoring mode, an ingress flow definition can include the ingress priority. This is calculated similarly to the transmit queue with the exception that the Traffic Class table is not used. As shown in [Figure 13-2](#), the priority can be based on:


- The static value for the destination address in the ALR table
- The precedence bits in the IPv4 TOS octet
- The DIFFSERV mapping table indexed by the IPv4 TOS octet or the IPv6 Traffic Class octet
- The VLAN tag priority field using the per port Priority Regeneration table
- The port default with separate values for packets with or without a broadcast destination address.

FIGURE 13-5: SWITCH ENGINE INGRESS FLOW PRIORITY SELECTION



The ingress flow calculation is based on the packet type and the device configuration as shown in [Figure 13-6](#).

FIGURE 13-6: SWITCH ENGINE INGRESS FLOW PRIORITY CALCULATION



13.3.7 BROADCAST STORM CONTROL

In addition to ingress rate limiting, the device supports hardware broadcast storm control on a per port basis. This feature is enabled via the [Switch Engine Broadcast Throttling Register \(SWE_BCST_THROT\)](#). The allowed rate per port is specified as the number of bytes multiplied by 64 allowed to be received every 1.72 ms interval. Packets that exceed this limit are dropped. Typical values are listed in [Table 13-3](#). When a port is receiving at 10 Mbps, any setting above 34 has the effect of not limiting the rate.

TABLE 13-3: TYPICAL BROADCAST RATE SETTINGS

Broadcast Throttle Level	Bandwidth
252	75 Mbps
168	50 Mbps
134	40 Mbps
67	20 Mbps
34	10 Mbps
17	5 Mbps
8	2.4 Mbps
4	1.2 Mbps
3	900 kbps
2	600 kbps
1	300 kbps

In addition to the rate limit, the [Buffer Manager Broadcast Buffer Level Register \(BM_BCST_LVL\)](#) specifies the maximum number of buffers that can be used by broadcasts, multicasts and unknown unicasts.

13.3.8 IPV4 IGMP / IPV6 MLD SUPPORT

The device provides Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) hardware support using two mechanisms: IGMP/MLD monitoring and Multicast Pruning.

On ingress, if the [Enable IGMP Monitoring](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#) is set, IGMP *multicast* packets are trapped and redirected to the [MLD/IGMP Monitor Port](#) (typically set to the port to which the host CPU is connected). IGMP packets are identified as IPv4 packets with a protocol of 2. Both Ethernet and IEEE 802.3 frame formats are supported as are VLAN tagged packets.

On ingress, if the [Enable MLD Monitoring](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#) is set, MLD *multicast* packets are trapped and redirected to the [MLD/IGMP Monitor Port](#) (typically set to the port to which the host CPU is connected). MLD packets are identified as IPv6 packets with a Next Header value or a Hop-by-Hop Next Header value of 58 decimal (ICMPv6). Optionally, via the [Enable Other MLD Next Headers](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#), IPv6 Next Header values or Hop-by-Hop Next Header values of 43 (Routing), 44 (Fragment), 50 (ESP), 51 (AH) and 60 (Destination Options) can be enabled. Optionally, via the [Enable Any MLD Hop-by-Hop Next Header](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#), all Hop-by-Hop Next Header values can be enabled. Both Ethernet and IEEE 802.3 frame formats are supported as are VLAN tagged packets.

Note: There is a limitation with packets using the IEEE 802.3 frame format. For single and double (such as in the case of a CPU tag and VLAN tag) tagged packets, the Hop-by-Hop Next Header value can not be reached within the 64 byte processing limit and therefore would not be detected.

Once the IGMP or MLD packets are received by the host CPU, the host software can decide which port or ports need to be members of the multicast group. This group is then added to the ALR table as detailed in [Section 13.3.1.3, "Multicast Pruning," on page 325](#). The host software should also forward the original IGMP or MLD packet if necessary.

Normally, packets are never transmitted back to the receiving port. For IGMP/MLD monitoring, this may optionally be enabled via the [Allow Monitor Echo](#) field in the [Switch Engine Global Ingress Configuration Register \(SWE_GLOBAL INGRSS_CFG\)](#). This function would be used if the monitoring port wished to participate in the IGMP/MLD group without the need to perform special handling in the transmit portion of the driver software.

Note: Most forwarding rules are skipped when a packet is monitored. However, a packet is still filtered if:

- The source port is in the Disabled state.
- The source port is in the Learning or Listening / Blocking state (unless Spanning Tree Port State Override is in effect).
- VLANs are enabled, the packet is untagged or priority tagged and the Admit Only VLAN bit for the ingress port is set.
- VLANs are enabled and the packet is tagged and had a VID equal to FFFh.
- VLANs are enabled, Enable Membership Checking on Ingress is set, Admit Non Member is cleared and the source port is not a member of the incoming VLAN.

13.3.9 PORT MIRRORING

The device supports port mirroring where packets received or transmitted on a port or ports can also be copied onto another “sniffer” port.

Port mirroring is configured using the [Switch Engine Port Mirroring Register \(SWE_PORT_MIRROR\)](#). Multiple mirrored ports can be defined, but only one sniffer port can be defined.

When receive mirroring is enabled via the [Enable RX Mirroring](#) field, packets that are forwarded from a port designated as a [Mirrored Port](#) are also transmitted by the [Sniffer Port](#). For example, Port 2 is setup to be a mirrored port and Port 0 is setup to be the sniffer port. If a packet is received on Port 2 with a destination of Port 1, it is forwarded to both Port 1 and Port 0.

When transmit mirroring is enabled via the [Enable TX Mirroring](#) field, packets that are forwarded to a port designated as a [Mirrored Port](#) are also transmitted by the [Sniffer Port](#). For example, Port 2 is setup to be a mirrored port and Port 0 is setup to be the sniffer port. If a packet is received on Port 1 with a destination of Port 2, it is forwarded to both Port 2 and Port 0.

A packet will never be transmitted out of the receiving port. A receive packet is not normally mirrored if it is filtered. This can optionally be enabled via the [Enable RX Mirroring Filtered](#) field.

13.3.10 HOST CPU PORT SPECIAL TAGGING

The [Switch Engine Ingress Port Type Register \(SWE_INGRSS_PORT_TYP\)](#) and the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) are used to enable a special VLAN tag that is used by the host CPU. This special tag is used to specify the port(s) where packets from the CPU should be sent and to indicate which port received the packet that was forwarded to the CPU.

13.3.10.1 Packets from the Host CPU

The [Switch Engine Ingress Port Type Register \(SWE_INGRSS_PORT_TYP\)](#) configures the switch to use the special VLAN tag in packets from the host CPU as a destination port indicator. A setting of 11b should be used on the port that is connected to the host CPU (typically Port 0). A setting of 00b should be used on the normal network ports.

The special VLAN tag is a normal VLAN tag where the VID field is used as the destination port indicator.

VID bit 3 indicates a request for an ALR lookup.

If VID bit 3 is zero, then bits 0 and 1 specify the destination port (0, 1, 2) or broadcast (3). Bit 4 is used to specify if the STP port state should be overridden. When set, the packet will be transmitted, even if the destination port(s) is (are) in the Learning or Listening / Blocking state.

If VID bit 3 is one, then the normal ALR lookup is performed and learning is performed on the source address (if enabled in the [Switch Engine Port Ingress Configuration Register \(SWE_PORT INGRSS_CFG\)](#) and the port state for the CPU port is set to Forwarding or Learning). The STP port state override is taken from the ALR entry.

VID bit 5 indicates a request to calculate the packet priority (and egress queue) based on the packet contents.

If VID bit 5 is zero, the PRI field from the VLAN tag is used as the packet priority.

If VID bit 5 is one, the packet priority is calculated from the packet contents. The procedure described in [Section 13.3.3, "Transmit Priority Queue Selection," on page 327](#) is followed with the exception that the special tag is skipped and the VLAN priority is taken from the second VLAN tag, if it exists.

VID bit 6 indicates a request to follow VLAN rules.

If VID bit 6 is zero, a default membership of “all ports” is assumed and no VLAN rules are followed.

If VID bit 6 is one, all ingress and egress VLAN rules are followed. The procedure described in [Section 13.3.2, "Forwarding Rules," on page 326](#) is followed with the exception that the special tag is skipped and the VID is taken from the second VLAN tag if it exists.

Upon egress from the destination port(s), the special tag is removed. If a regular VLAN tag needs to be sent as part of the packet, then it should be part of the packet data from the host CPU following the special tag.

When specifying Port 0 as the destination port, the VID will be set to 0. A VID of 0 is normally considered a priority tagged packet. Such a packet will be filtered if Admit Only VLAN is set on the host CPU port. Either avoid setting Admit Only VLAN on the host CPU port or set an unused bit in the VID field.

Note: The maximum size tagged packet that can normally be sent into a switch port (on port 0) is 1522 bytes. Since the special tag consumes four bytes of the packet length, the outgoing packet is limited to 1518 bytes, even if it contains a regular VLAN tag as part of the packet data. If a larger outgoing packet is required, the Jumbo2K bit in the [Port x MAC Receive Configuration Register \(MAC_RX_CFG_x\)](#) of Port 0 should be set.

13.3.10.2 Packets to the Host CPU

The [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) configures the switch to add the special VLAN tag in packets to the host CPU as a source port indicator. A setting of 11b should be used only on the port that is connected to the host CPU (typically Port 0). Other settings can be used on the normal network ports as needed.

The special VLAN tag is a normal VLAN tag where:

- The priority field indicates the packet's priority as classified on receive.
- Bits 0 and 1 of the VID field specify the source port (0, 1 or 2).
- Bit 3 of the VID field indicates the packet was a monitored IGMP or MLD packet.
- Bit 4 of the VID field indicates STP override was set ([Static](#) and [Age 1/Override](#) bits set) in the ALR entry for the packet's Destination MAC Address.
- Bit 5 of the VID field indicates the [Static](#) bit was set in the ALR entry for the packet's Destination MAC address.
- Bit 6 of the VID field indicates [Priority Enable](#) was set in the ALR entry for the packet's Destination MAC address.
- Bits 7, 8 and 9 of the VID field are the [Priority](#) field in the ALR entry for the packet's Destination MAC address - these can be used as a tag to identify different packet types (PTP, RSTP, etc.) when the host CPU adds MAC address entries.

Note: Bits 4 through 9 of the VID field will be all zero for Destination MAC Addresses that have been learned (i.e., not added by the host) or are not found in the ALR table (i.e., not learned or added by the host).

Upon egress from the host CPU port, the special tag is added. If a regular VLAN tag already exists, it is not deleted. Instead it will follow the special tag.

Note: Since the special tag adds four bytes to the length of the packet, it is possible for a normally tagged, maximum size, incoming packet to become 1526 bytes in length. In order for the Host MAC to receive this length packet without indicating a length error, the [Host MAC VLAN2 Tag Register \(HMAC_VLAN2\)](#) in the Host MAC should be set to 8100h and the [Host MAC VLAN1 Tag Register \(HMAC_VLAN1\)](#) should be set to a value other than 8100h. This configuration will allow frames up to 1538 bytes in length to be received.

Note: Since the special tag adds four bytes to the length of the packet, it is possible for a normally tagged, maximum size, incoming jumbo packet to become 2052 bytes in length. This packet will be received by the Host MAC with the following conditions:

- The receive status will indicate Frame Too Long
- Up to four bytes of the end of packet may be truncated (the maximum receive length at the Host MAC is 2048).

13.3.11 COUNTERS

A counter is maintained per port that contains the number of MAC address that were not learned or were overwritten by a different address due to MAC Address Table space limitations. These counters are accessible via the following registers:

- [Switch Engine Port 0 Learn Discard Count Register \(SWE_LRN_DISCRD_CNT_0\)](#)
- [Switch Engine Port 1 Learn Discard Count Register \(SWE_LRN_DISCRD_CNT_1\)](#)
- [Switch Engine Port 2 Learn Discard Count Register \(SWE_LRN_DISCRD_CNT_2\)](#)

A counter is maintained per port that contains the number of packets filtered at ingress. This count includes packets filtered due to broadcast throttling, but does not include packets dropped due to ingress rate limiting. These counters are accessible via the following registers:

- [Switch Engine Port 0 Ingress Filtered Count Register \(SWE_FILTERED_CNT_0\)](#)
- [Switch Engine Port 1 Ingress Filtered Count Register \(SWE_FILTERED_CNT_1\)](#)
- [Switch Engine Port 2 Ingress Filtered Count Register \(SWE_FILTERED_CNT_2\)](#)

13.4 Buffer Manager (BM)

The Buffer Manager (BM) provides control of the free buffer space, the multiple priority transmit queues, transmission scheduling and packet dropping. VLAN tag insertion and removal is also performed by the Buffer Manager. The following sections detail the various features of the Buffer Manager.

13.4.1 PACKET BUFFER ALLOCATION

The packet buffer consists of 32 kB of RAM that is dynamically allocated in 128 byte blocks as packets are received. Up to 16 blocks may be used per packet, depending on the packet length. The blocks are linked together as the packet is received. If a packet is filtered, dropped or contains a receive error, the buffers are reclaimed.

13.4.1.1 Buffer Limits and Flow Control Levels

The BM keeps track of the amount of buffers used per each ingress port. These counts are used to generate flow control (half-duplex backpressure or full-duplex pause frames) and to limit the amount of buffer space that can be used by any individual receiver (hard drop limit). The flow control and drop limit thresholds are dynamic and adapt based on the current buffer usage. Based on the number of active receiving ports, the drop level and flow control pause and resume thresholds adjust between fixed settings and two user programmable levels via the [Buffer Manager Drop Level Register \(BM_DROP_LVL\)](#), the [Buffer Manager Flow Control Pause Level Register \(BM_FC_PAUSE_LVL\)](#) and the [Buffer Manager Flow Control Resume Level Register \(BM_FC_RESUME_LVL\)](#) respectively.

The BM also keeps a count of the number of buffers that are queued for multiple ports (broadcast queue). This count is compared against the [Buffer Manager Broadcast Buffer Level Register \(BM_BCST_LVL\)](#) and if the configured drop level is reached or exceeded, subsequent packets are dropped.

13.4.2 RANDOM EARLY DISCARD (RED)

Based on the ingress flow monitoring detailed in [Section 13.3.6, "Ingress Flow Metering and Coloring," on page 332](#), packets are colored as Green, Yellow or Red. Packets colored Red are always discarded if the [Drop on Red](#) bit in the [Buffer Manager Configuration Register \(BM_CFG\)](#) is set. If the [Drop on Yellow](#) bit in the [Buffer Manager Configuration Register \(BM_CFG\)](#) is set, packets colored Yellow are randomly discarded based on the moving average number of buffers used by the ingress port.

The probability of a discard is programmable into the Random Discard Weight table via the [Buffer Manager Random Discard Table Command Register \(BM_RNDM_DSCRD_TBL_CMD\)](#), the [Buffer Manager Random Discard Table Write Data Register \(BM_RNDM_DSCRD_TBL_WDATA\)](#) and the [Buffer Manager Random Discard Table Read Data Register \(BM_RNDM_DSCRD_TBL_RDATA\)](#). The Random Discard Weight table contains sixteen entries, each 10-bits wide. Each entry corresponds to a range of the average number of buffers used by the ingress port. Entry 0 is for 0 to 15 buffers, entry 1 is for 16 to 31 buffers, etc. The probability for each entry is set in 1/1024. For example, a setting of 1 is 1-in-1024 or approximately 0.1%. A setting of all ones (1023) is 1023-in-1024 or approximately 99.9%.

Refer to [Section 13.7.4.10, "Buffer Manager Random Discard Table Command Register \(BM_RNDM_DSCRD_TBL_CMD\)," on page 445](#) for additional details on writing and reading the Random Discard Weight table.

13.4.3 TRANSMIT QUEUES

Once a packet has been completely received, it is queued for transmit. There are four queues per transmit port, one for each level of transmit priority. Each queue is virtual (if there are no packets for that port/priority, the queue is empty) and dynamic (a queue may have any length if there is enough memory space). When a packet is read from the memory and sent out to the corresponding port, the used buffers are released.

13.4.4 TRANSMIT PRIORITY QUEUE SERVICING

When a transmit queue is non-empty, it is serviced and the packet is read from the buffer RAM and sent to the transmit MAC. If there are multiple queues that require servicing, one of two methods may be used: fixed priority ordering or weighted round-robin ordering. If the [Fixed Priority Queue Servicing](#) bit in the [Buffer Manager Configuration Register](#)

(BM_CFG) is set, a strict order, fixed priority is selected. Transmit queue 3 has the highest priority, followed by 2, 1 and 0. If the [Fixed Priority Queue Servicing](#) bit in the [Buffer Manager Configuration Register \(BM_CFG\)](#) is cleared, a weighted round-robin order is followed. Assuming all four queues are non-empty, the service is weighted with a 9:4:2:1 ratio (queue 3,2,1,0). The servicing is blended to avoid burstiness (e.g., queue 3, then queue 2, then queue 3, etc.).

13.4.5 EGRESS RATE LIMITING (LEAKY BUCKET)

For egress rate limiting, the leaky bucket algorithm is used on each output priority queue. For each output port, the bandwidth that is used by each priority queue can be limited. If any egress queue receives packets faster than the specified egress rate, packets will be accumulated in the packet memory. After the memory is used, packet dropping or flow control will be triggered.

Egress rate limiting occurs before the [Transmit Priority Queue Servicing](#), such that a lower priority queue will be serviced if a higher priority queue is being rate-limited.

The egress limiting is enabled per priority queue. After a packet is selected to be sent, its length is recorded. The switch then waits a programmable amount of time, scaled by the packet length, before servicing that queue once again. The amount of time per byte is programmed into the Buffer Manager Egress Rate registers (refer to [Section 13.7.4.14](#) through [Section 13.7.4.19](#) for detailed register definitions). The value programmed is in approximately 20 ns per byte increments. Typical values are listed in [Table 13-4](#). When a port is transmitting at 10 Mbps, any setting above 39 has the effect of not limiting the rate.

TABLE 13-4: TYPICAL EGRESS RATE SETTINGS

Egress Rate Setting	Time Per Byte	Bandwidth @ 64 Byte Packet	Bandwidth @ 512 Byte Packet	Bandwidth @ 1518 Byte Packet
0-3	80 ns	76 Mbps (Note 1)	96 Mbps (Note 1)	99 Mbps (Note 1)
4	100 ns	66 Mbps	78 Mbps	80 Mbps
5	120 ns	55 Mbps	65 Mbps	67 Mbps
6	140 ns	48 Mbps	56 Mbps	57 Mbps
7	160 ns	42 Mbps	49 Mbps	50 Mbps
9	200 ns	34 Mbps	39 Mbps	40 Mbps
12	260 ns	26 Mbps	30 Mbps	31 Mbps
19	400 ns	17 Mbps	20 Mbps	20 Mbps
39	800 ns	8.6 Mbps	10 Mbps	10 Mbps
78	1580 ns	4.4 Mbps	5 Mbps	5 Mbps
158	3180 ns	2.2 Mbps	2.5 Mbps	2.5 Mbps
396	7940 ns	870 kbps	990 kbps	1 Mbps
794	15900 ns	440 kbps	490 kbps	500 kbps
1589	31800 ns	220 kbps	250 kbps	250 kbps
3973	79480 ns	87 kbps	98 kbps	100 kbps
7947	158960 ns	44 kbps	49 kbps	50 kbps

Note 1: These are the unlimited max. bandwidths when IFG and preamble are taken into account.

13.4.6 ADDING, REMOVING AND CHANGING VLAN TAGS

Based on the port configuration and the received packet format, a VLAN tag can be added to, removed from or modified in a packet. There are four received packet type cases: non-tagged, priority-tagged, normal-tagged and CPU special-tagged. There are also four possible settings for an egress port: dumb, access, hybrid and CPU. In addition, each VLAN table entry can specify the removal of the VLAN tag (the entry's un-tag bit).

The tagging/un-tagging rules are specified as follows:

- **Dumb Port** - This port type generally does not change the tag.

- When a received packet is non-tagged, priority-tagged or normal-tagged the packet passes untouched.
- When a packet is received special-tagged from a CPU port, the special tag is removed.
- **Access Port** - This port type generally does not support tagging.
 - When a received packet is non-tagged, the packet passes untouched.
 - When a received packet is priority-tagged or normal-tagged, the tag is removed.
 - When a received packet is special-tagged from a CPU port, the special tag is removed.
- **CPU Port** - Packets transmitted from this port type generally contain a special tag. Special tags are described in detail in [Section 13.3.10, "Host CPU Port Special Tagging," on page 337](#).
- **Hybrid Port** - Generally, this port type supports a mix of normal-tagged and non-tagged packets. It is the most complex, but most flexible port type.

For clarity, the following details the incoming un-tag instruction. As described in [Section 13.3.4, "VLAN Support," on page 330](#), the un-tag instruction is the three un-tag bits from the applicable entry in the VLAN table. The entry in the VLAN table is either the VLAN from the received packet or the ingress port's default VID.

When a received packet is non-tagged, a new VLAN tag is added if two conditions are met. First, the Insert Tag bit for the egress port in the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) must be set. Second, the un-tag bit, for the egress port, from the un-tag instruction associated with the ingress port's default VID, must be cleared. The VLAN tag that is added will have a VID taken from either the ingress or egress port's default VID. The priority of the VLAN tag is either the priority calculated on ingress or the egress port's default. The choice of ingress or egress is determined by the egress port's VID/Priority Select bit in the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#).


When a received packet is priority-tagged, either the tag is removed or it is modified. If the un-tag bit, for the egress port, from the un-tag instruction associated with the ingress port's default VID is set, then the tag is removed. Otherwise, the tag is modified. The VID of the new VLAN tag is changed to either the ingress or egress port's default VID. If the Change Priority bit in the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) for the egress port is set, then the Priority field of the new VLAN tag is also changed. The priority of the VLAN tag is either the priority calculated on ingress or the egress port's default. The choice of ingress or egress is determined by the egress port's VID/Priority Select bit.

When a received packet is normal-tagged, either the tag is removed, modified or passed unchanged. If the un-tag bit, for the egress port, from the un-tag instruction associated with the VID in the received packet is set, then the tag is removed. Else, if the Change Tag bit in the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) for the egress port is clear, the packet passes untouched. Else, if both the Change VLAN ID and the Change Priority bits in the [Buffer Manager Egress Port Type Register \(BM_EGRSS_PORT_TYPE\)](#) for the egress port are clear, the packet passes untouched. Otherwise, the tag is modified. If the Change VLAN ID bit for the egress port is set, the VID of the new VLAN tag is changed to either the ingress or egress port's default VID. If the Change Priority bit for the egress port is set, the Priority field of the new VLAN tag is changed to either the priority calculated on ingress or the egress port's default. The choice of ingress or egress is determined by the egress port's VID / Priority Select bit.

When a packet is received special-tagged from a CPU port, the special tag is removed.

Hybrid tagging is summarized in [Figure 13-7](#).

FIGURE 13-7: HYBRID PORT TAGGING AND UN-TAGGING



The default VLAN ID and priority of each port may be configured via the following registers:

- [Buffer Manager Port 0 Default VLAN ID and Priority Register \(BM_VLAN_0\)](#)
- [Buffer Manager Port 1 Default VLAN ID and Priority Register \(BM_VLAN_1\)](#)
- [Buffer Manager Port 2 Default VLAN ID and Priority Register \(BM_VLAN_2\)](#)

13.4.7 COUNTERS

A counter is maintained per port that contains the number of packets dropped due to buffer space limits and ingress rate limit discarding (Red and random Yellow dropping). These counters are accessible via the following registers:

- [Buffer Manager Port 0 Drop Count Register \(BM_DRP_CNT_SRC_0\)](#)
- [Buffer Manager Port 1 Drop Count Register \(BM_DRP_CNT_SRC_1\)](#)
- [Buffer Manager Port 2 Drop Count Register \(BM_DRP_CNT_SRC_2\)](#)

A counter is maintained per port that contains the number of packets dropped due solely to ingress rate limit discarding (Red and random Yellow dropping). This count value can be subtracted from the drop counter, as described above, to obtain the drop counts due solely to buffer space limits. The ingress rate drop counters are accessible via the following registers:

- [Buffer Manager Port 0 Ingress Rate Drop Count Register \(BM_RATE_DRP_CNT_SRC_0\)](#)
- [Buffer Manager Port 1 Ingress Rate Drop Count Register \(BM_RATE_DRP_CNT_SRC_1\)](#)
- [Buffer Manager Port 2 Ingress Rate Drop Count Register \(BM_RATE_DRP_CNT_SRC_2\)](#)

13.5 Switch Fabric Interface Logic

13.5.1 FLOW CONTROL ENABLE LOGIC

Each Switch Fabric port (0,1,2) is provided with two flow control enable inputs, one for transmission and one for reception. Flow control on transmission allows the transmitter to generate back pressure in half-duplex mode and pause packets in full-duplex. Flow control in reception enables the reception of pause packets to pause transmissions.

The state of these enables is based on the state of the port's duplex and Auto-Negotiation settings and results, provided by the attached PHY. For port 0, the PHY is the Virtual PHY. For port 1, the PHY is Physical PHY A. For port 2, the PHY is Physical PHY B. The PHYs' advertised pause flow control abilities are set via the Symmetric Pause and Asymmetric Pause bits of the PHYs' Auto-Negotiation Advertisement Register. This allows the PHY to advertise its flow control abilities and auto-negotiate the flow control settings with its link partner. The link partners' advertised pause flow control abilities are returned via the Symmetric Pause and Asymmetric Pause bits of the PHYs' Auto-Negotiation Link Partner Base Ability Register.

The pause flow control settings may also be manually set via the manual flow control registers ([Port 1 Manual Flow Control Register \(MANUAL_FC_1\)](#), [Port 2 Manual Flow Control Register \(MANUAL_FC_2\)](#) or [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#)). [Table 13-5](#) details the Switch Fabric flow control enable logic. These registers allow the Switch Fabric ports flow control settings to be manually set when Auto-Negotiation is disabled or the respective manual flow control select bit is set. The currently enabled duplex and flow control settings can also be monitored via these registers.

When in half-duplex mode, the transmit flow control (back pressure) enable is determined directly by the BP_EN_x bit of the port's manual flow control register. When Auto-Negotiation is disabled or the MANUAL_FC_x bit of the port's manual flow control register is set, the switch port flow control enables during full-duplex are determined by the TX_FC_x and RX_FC_x bits of the port's manual flow control register. When Auto-Negotiation is enabled and the MANUAL_FC_x bit is cleared, the switch port flow control enables during full-duplex are determined by Auto-Negotiation.

Note: The flow control values in the PHYs' Auto-Negotiation Advertisement Register are not affected by the values of the manual flow control register.

TABLE 13-5: SWITCH FABRIC FLOW CONTROL ENABLE LOGIC

Case	Manual_FC_X	AN Enable	AN Complete	LP AN Able	Duplex	AN Pause Advertisement (Note 3)	AN ASYM Pause Advertisement (Note 3)	LP Pause Ability (Note 3)	LP ASYM Pause Ability (Note 3)	RX Flow Control Enable	TX Flow Control Enable
-	1	X	X	X	Half	X	X	X	X	0	BP_EN_x
-	X	0	X	X	Half	X	X	X	X	0	BP_EN_x
-	1	X	X	X	Full	X	X	X	X	RX_FC_x	TX_FC_x
-	X	0	X	X	Full	X	X	X	X	RX_FC_x	TX_FC_x
1	0	1	0	X	X	X	X	X	X	0	0
2	0	1	1	0	Half (Note 2)	X	X	X	X	0	BP_EN_x
3	0	1	1	1	Half	X	X	X	X	0	BP_EN_x
4	0	1	1	1	Full	0	0	X	X	0	0
5	0	1	1	1	Full	0	1	0	X	0	0
6	0	1	1	1	Full	0	1	1	0	0	0
7	0	1	1	1	Full	0	1	1	1	0	1
8	0	1	1	1	Full	1	0	0	X	0	0
9	0	1	1	1	Full	1	X	1	X	1	1
10	0	1	1	1	Full	1	1	0	0	0	0
11	0	1	1	1	Full	1	1	0	1	1	0

Note 2: If Auto-Negotiation is enabled and complete, but the link partner is not Auto-Negotiation capable, half-duplex is forced via the parallel detect function.

Note 3: These are the bits from the PHYs' Auto-Negotiation Advertisement and Auto-Negotiation Link Partner Base Ability Registers. If a Switch Fabric Port is connected to a Virtual PHY, these are the local/partner swapped outputs from the Virtual PHY's Auto-Negotiation Advertisement and Auto-Negotiation Link Partner Base Ability Registers. Refer to the Virtual PHY Auto-Negotiation section for more information.

Per Table 13-5, the following cases are possible:

- **Case 1** - Auto-Negotiation is still in progress. Since the result is not yet established, flow control is disabled.
- **Case 2** - Auto-Negotiation is enabled and unsuccessful (link partner not Auto-Negotiation capable). The link partner ability is undefined, effectively a don't-care value, in this case. The duplex setting will default to half-duplex in this case. Flow control is determined by the BP_EN_x bit.
- **Case 3** - Auto-Negotiation is enabled and successful with half-duplex as a result. The link partner ability is undefined since it only applies to full-duplex operation. Flow control is determined by the BP_EN_x bit.
- **Cases 4-11** - Auto-Negotiation is enabled and successful with full-duplex as the result. In these cases, the advertisement registers and the link partner ability controls the RX and TX enables. These cases match IEEE 802.3 Annex 28B.3.
 - **Cases 4,5,6,8,10** - No flow control enabled
 - **Case 7** - Asymmetric pause towards partner (away from switch port)
 - **Case 9** - Symmetric pause

- **Case 11** - Asymmetric pause from partner (towards switch port)

13.5.2 EEE ENABLE LOGIC

Each Switch Fabric port (0,1,2) is provided with an input which permits the generation and decoding of EEE LPI signaling. These signals are in addition to the Switch Fabric ports' [Energy Efficient Ethernet \(EEE_ENABLE\)](#) bits and are used to check various port conditions such as speed, duplex and mode.

Normally, in order to permit EEE functions, the port must be in internal PHY mode or MII MAC mode, the port speed must be 100 Mbps, the current duplex must be full and the auto-negotiation result must indicate that both the local and partner device support EEE 100 Mbps. In order to prevent an unstable link condition, the PHY link status also must indicate "up" for one second.

13.5.2.1 Port 0

EEE functions are not used since the port is connected internally to the Host MAC.

13.5.2.2 Port 1

The port speed, duplex, link status and auto-negotiation result come from physical PHY A.

13.5.2.3 Port 2

The port speed, duplex, link status and auto-negotiation result come from physical PHY B.

13.5.3 SWITCH FABRIC CSR INTERFACE

The Switch Fabric CSRs provide register level access to the various parameters of the Switch Fabric. Switch Fabric related registers can be classified into two main categories based upon their method of access: direct and indirect.

The directly accessible Switch Fabric registers are part of the main system CSRs and are detailed in [Section 13.6, "Switch Fabric Interface Logic Registers," on page 348](#). These registers provide Switch Fabric manual flow control (Ports 0-2), data/command registers (for access to the indirect Switch Fabric registers) and switch MAC address configuration.

The indirectly accessible Switch Fabric registers reside within the Switch Fabric and must be accessed indirectly via the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#) and the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) or the set of [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#). The indirectly accessible Switch Fabric CSRs provide full access to the many configurable parameters of the Switch Engine, Buffer Manager and each switch port. The Switch Fabric CSRs are detailed in [Section 13.7, "Switch Fabric Control and Status Registers," on page 363](#).

13.5.4 SWITCH FABRIC CSR WRITES

To perform a write to an individual Switch Fabric register, the desired data must first be written into the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#). The write cycle is initiated by performing a single write to the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) with the [CSR Busy \(CSR_BUSY\)](#) bit set, the [CSR Address \(CSR_ADDR\[15:0\]\)](#) field set to the desired register address, the [Read/Write \(R_nW\)](#) bit cleared, the [Auto Increment \(AUTO_INC\)](#) and [Auto Decrement \(AUTO_DEC\)](#) fields cleared and the desired [CSR Byte Enable \(CSR_BE\[3:0\]\)](#) bits selected. The completion of the write cycle is indicated by the clearing of the [CSR Busy \(CSR_BUSY\)](#) bit.


A second write method may be used which utilizes the auto increment/decrement function of the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) for writing sequential register addresses. When using this method, the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) must first be written with the [Auto Increment \(AUTO_INC\)](#) or [Auto Decrement \(AUTO_DEC\)](#) bit set, the [CSR Address \(CSR_ADDR\[15:0\]\)](#) field written with the desired register address, the [Read/Write \(R_nW\)](#) bit cleared and the desired CSR byte enable bits selected (typically all set). The write cycles are then initiated by writing the desired data into the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#). The completion of the write cycle is indicated by the clearing of the [CSR Busy \(CSR_BUSY\)](#) bit, at which time the address in the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) is incremented or decremented accordingly. The user may then initiate a subsequent write cycle by writing the desired data into the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#).

The third write method is to use the direct data range write function. Writes within the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) address range automatically set the appropriate register address, set all four [CSR Byte Enable \(CSR_BE\[3:0\]\)](#) bits, clears the [Read/Write \(R_nW\)](#) bit and set the [CSR Busy \(CSR_BUSY\)](#) bit of the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#). The completion of the write cycle is

indicated by the clearing of the [CSR Busy \(CSR_BUSY\)](#) bit. Since the address range of the Switch Fabric CSRs exceeds that of the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) address range, a sub-set of the Switch Fabric CSRs is mapped to the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) address range as detailed in [Table 13-8, "Switch Fabric CSR to SWITCH_CSR_DIRECT_DATA Address Range Map," on page 360](#).

[Figure 13-8](#) illustrates the process required to perform a Switch Fabric CSR write. The minimum wait periods as specified in [Table 5-2, "Read After Write Timing Rules," on page 38](#) are required where noted.

FIGURE 13-8: SWITCH FABRICS CSR WRITE ACCESS FLOW DIAGRAM



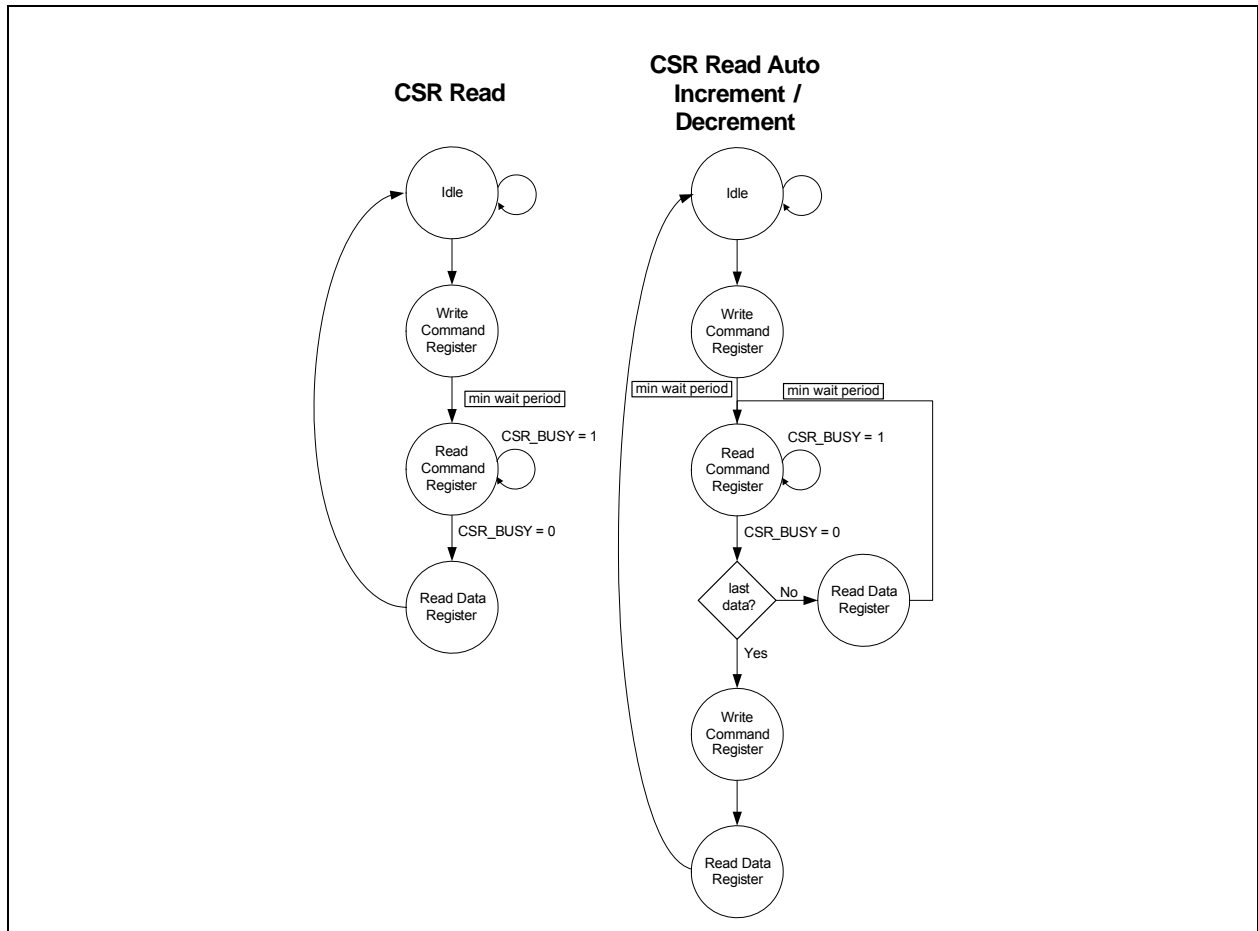
13.5.5 SWITCH FABRIC CSR READS

To perform a read of an individual Switch Fabric register, the read cycle must be initiated by performing a single write to the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) with the [CSR Busy \(CSR_BUSY\)](#) bit set, the [CSR Address \(CSR_ADDR\[15:0\]\)](#) field set to the desired register address, the [Read/Write \(R_nW\)](#) bit set and the [Auto Increment \(AUTO_INC\)](#) and [Auto Decrement \(AUTO_DEC\)](#) fields cleared. Valid data is available for reading when the [CSR Busy \(CSR_BUSY\)](#) bit is cleared, indicating that the data can be read from the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#).

A second read method may be used which utilizes the auto increment/decrement function of the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) for reading sequential register addresses. When using this method, the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) must first be written with the [CSR Busy \(CSR_BUSY\)](#) bit set, the [Auto Increment \(AUTO_INC\)](#) or [Auto Decrement \(AUTO_DEC\)](#) bit set, the [CSR Address \(CSR_ADDR\[15:0\]\)](#) field written with the desired register address and the [Read/Write \(R_nW\)](#) bit set. The completion of a read cycle is indicated by the clearing of the [CSR Busy \(CSR_BUSY\)](#) bit, at which time the data can be read from the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#). When the data is read, the address in the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) is incremented or decremented accordingly and another read cycle is started automatically. The user should clear the [Auto Increment \(AUTO_INC\)](#) and [Auto Decrement \(AUTO_DEC\)](#) bits before reading the last data to avoid an unintended read cycle.

Figure 13-9 illustrates the process required to perform a Switch Fabric CSR read. The minimum wait periods as specified in Table 5-2, "Read After Write Timing Rules," on page 38 and Table 5-3, "Read After Read Timing Rules," on page 40 are required where noted.

FIGURE 13-9: SWITCH FABRICS CSR READ ACCESS FLOW DIAGRAM



13.6 Switch Fabric Interface Logic Registers

This section details the directly addressable System CSRs which are related to the Switch Fabric.

The flow control of all three ports of the Switch Fabric can be configured via the System CSR's [Port 1 Manual Flow Control Register \(MANUAL_FC_1\)](#), [Port 2 Manual Flow Control Register \(MANUAL_FC_2\)](#) and [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#).

The MAC address used by the switch for Pause frames is configured via the [Switch Fabric MAC Address High Register \(SWITCH_MAC_ADDRH\)](#) and the [Switch Fabric MAC Address Low Register \(SWITCH_MAC_ADDRL\)](#).

The [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#), the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#) and the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) serve as an accessible interface to the full range of otherwise inaccessible switch control and status registers. A list of all the Switch Fabric CSRs can be seen in [Table 13-9](#). For detailed descriptions of the Switch Fabric CSRs that are accessible via these interface registers, refer to [Section 13.7, "Switch Fabric Control and Status Registers"](#). For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map,"](#) on page 33.

TABLE 13-6: SWITCH FABRIC INTERFACE LOGIC REGISTERS

ADDRESS	Register Name (SYMBOL)
1A0h	Port 1 Manual Flow Control Register (MANUAL_FC_1)
1A4h	Port 2 Manual Flow Control Register (MANUAL_FC_2)
1A8h	Port 0 Manual Flow Control Register (MANUAL_FC_0)
1ACh	Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA)
1B0h	Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD)
1F0h	Switch Fabric MAC Address High Register (SWITCH_MAC_ADDRH)
1F4h	Switch Fabric MAC Address Low Register (SWITCH_MAC_ADDRL)
200h-2F8h	Switch Fabric CSR Interface Direct Data Registers (SWITCH_CSR_DIRECT_DATA)

13.6.1 PORT 1 MANUAL FLOW CONTROL REGISTER (MANUAL_FC_1)

Offset: 1A0h Size: 32 bits

This read/write register allows for the manual configuration of the switch Port 1 flow control. This register also provides read back of the currently enabled flow control settings, whether set manually or Auto-Negotiated. Refer to [Section 13.5.1, "Flow Control Enable Logic"](#) for additional information.

Note: The flow control values in the PHY's Auto-Negotiation Advertisement Register are not affected by the values of this register.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6	Port 1 Backpressure Enable (BP_EN_1) This bit enables/disables the generation of half-duplex backpressure on switch Port 1. 0: Disable backpressure 1: Enable backpressure	R/W	Note 4
5	Port 1 Current Duplex (CUR_DUP_1) This bit indicates the actual duplex setting of switch Port 1. 0: Full-Duplex 1: Half-Duplex	RO	Note 5
4	Port 1 Current Receive Flow Control Enable (CUR_RX_FC_1) This bit indicates the actual receive flow setting of switch Port 1. 0: Flow control receive is currently disabled 1: Flow control receive is currently enabled	RO	Note 5
3	Port 1 Current Transmit Flow Control Enable (CUR_TX_FC_1) This bit indicates the actual transmit flow setting of switch Port 1. 0: Flow control transmit is currently disabled 1: Flow control transmit is currently enabled	RO	Note 5
2	Port 1 Full-Duplex Receive Flow Control Enable (RX_FC_1) When the MANUAL_FC_1 bit is set or Auto-Negotiation is disabled, this bit enables/disables the detection of full-duplex Pause packets on switch Port 1. 0: Disable flow control receive 1: Enable flow control receive	R/W	Note 6
1	Port 1 Full-Duplex Transmit Flow Control Enable (TX_FC_1) When the MANUAL_FC_1 bit is set or Auto-Negotiation is disabled, this bit enables/disables full-duplex Pause packets to be generated on switch Port 1. 0: Disable flow control transmit 1: Enable flow control transmit	R/W	Note 6

Bits	Description	Type	Default
0	Port 1 Full-Duplex Manual Flow Control Select (MANUAL_FC_1) This bit toggles flow control selection between manual and Auto-Negotiation. 0: If Auto-Negotiation is enabled, the Auto-Negotiation function determines the flow control of switch Port 1 (RX_FC_1 and TX_FC_1 values ignored). If Auto-Negotiation is disabled, the RX_FC_1 and TX_FC_1 values are used. 1: TX_FC_1 and RX_FC_1 bits determine the flow control of switch Port 1 when in full-duplex mode.	R/W	Note 7

Note 4: The default value of this field is determined by the [BP_EN_strap_1](#) configuration strap.

Note 5: The default value of this bit is determined by multiple strap settings.

Note 6: The default value of this field is determined by the [FD_FC_strap_1](#) configuration strap.

Note 7: The default value of this field is determined by the [manual_FC_strap_1](#) configuration strap.

13.6.2 PORT 2 MANUAL FLOW CONTROL REGISTER (MANUAL_FC_2)

Offset: 1A4h Size: 32 bits

This read/write register allows for the manual configuration of the switch Port 2 flow control. This register also provides read back of the currently enabled flow control settings, whether set manually or Auto-Negotiated. Refer to [Section 13.5.1, "Flow Control Enable Logic"](#) for additional information.

Note: The flow control values in the PHY's Auto-Negotiation Advertisement Register are not affected by the values of this register.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6	Port 2 Backpressure Enable (BP_EN_2) This bit enables/disables the generation of half-duplex backpressure on switch Port 2. 0: Disable backpressure 1: Enable backpressure	R/W	Note 8
5	Port 2 Current Duplex (CUR_DUP_2) This bit indicates the actual duplex setting of switch Port 2. 0: Full-Duplex 1: Half-Duplex	RO	Note 9
4	Port 2 Current Receive Flow Control Enable (CUR_RX_FC_2) This bit indicates the actual receive flow setting of switch Port 2. 0: Flow control receive is currently disabled 1: Flow control receive is currently enabled	RO	Note 9
3	Port 2 Current Transmit Flow Control Enable (CUR_TX_FC_2) This bit indicates the actual transmit flow setting of switch Port 2. 0: Flow control transmit is currently disabled 1: Flow control transmit is currently enabled	RO	Note 9
2	Port 2 Full-Duplex Receive Flow Control Enable (RX_FC_2) When the MANUAL_FC_2 bit is set or Auto-Negotiation is disabled, this bit enables/disables the detection of full-duplex Pause packets on switch Port 2. 0: Disable flow control receive 1: Enable flow control receive	R/W	Note 10
1	Port 2 Full-Duplex Transmit Flow Control Enable (TX_FC_2) When the MANUAL_FC_2 bit is set or Auto-Negotiation is disabled, this bit enables/disables full-duplex Pause packets to be generated on switch Port 2. 0: Disable flow control transmit 1: Enable flow control transmit	R/W	Note 10

Bits	Description	Type	Default
0	Port 2 Full-Duplex Manual Flow Control Select (MANUAL_FC_2) This bit toggles flow control selection between manual and Auto-Negotiation. 0: If Auto-Negotiation is enabled, the Auto-Negotiation function determines the flow control of switch Port 2 (RX_FC_2 and TX_FC_2 values ignored). If Auto-Negotiation is disabled, the RX_FC_2 and TX_FC_2 values are used. 1: TX_FC_2 and RX_FC_2 bits determine the flow control of switch Port 2 when in full-duplex mode	R/W	Note 11

Note 8: The default value of this field is determined by the [BP_EN_strap_2](#) configuration strap.

Note 9: The default value of this bit is determined by multiple strap settings.

Note 10: The default value of this field is determined by the [FD_FC_strap_2](#) configuration strap.

Note 11: The default value of this field is determined by the [manual_FC_strap_2](#) configuration strap.

13.6.3 PORT 0 MANUAL FLOW CONTROL REGISTER (MANUAL_FC_0)

Offset: 1A8h Size: 32 bits

This read/write register allows for the manual configuration of the switch Port 0 flow control. This register also provides read back of the currently enabled flow control settings, whether set manually or Auto-Negotiated. Refer to [Section 13.5.1, "Flow Control Enable Logic"](#) for additional information.

Note: The flow control values in the PHY's Auto-Negotiation Advertisement Register are not affected by the values of this register.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7	Port 0 Hard-wired Flow Control (HW_FC_0) When set to "1", the Host MACs RX FIFO level is connected to the switch engine's transmitter and the switch engines RX FIFO level is connected to the Host MACs transmitter. This achieves lower latency flow control. Note: All other flow control methods must be disabled when using this feature. (MANUAL_FC_0 should be set, TX_FC_0, RX_FC_0, and BP_EN_0 should be cleared. FCANY, FCADD, FCBRD, and FCMULT in the AFC_CFG register should be cleared).	R/W	0
6	Port 0 Backpressure Enable (BP_EN_0) This bit enables/disables the generation of half-duplex backpressure on switch Port 0. 0: Disable backpressure 1: Enable backpressure	R/W	Note 12
5	Port 0 Current Duplex (CUR_DUP_0) This bit indicates the actual duplex setting of switch Port 0. 0: Full-Duplex 1: Half-Duplex	RO	Note 13
4	Port 0 Current Receive Flow Control Enable (CUR_RX_0) This bit indicates the actual receive flow setting of switch Port 0 0: Flow control receive is currently disabled 1: Flow control receive is currently enabled	RO	Note 13
3	Port 0 Current Transmit Flow Control Enable (CUR_TX_FC_0) This bit indicates the actual transmit flow setting of switch Port 0. 0: Flow control transmit is currently disabled 1: Flow control transmit is currently enabled	RO	Note 13
2	Port 0 Full-Duplex Receive Flow Control Enable (RX_FC_0) When the MANUAL_FC_0 bit is set or Virtual Auto-Negotiation is disabled, this bit enables/disables the detection of full-duplex Pause packets on switch Port 0. 0: Disable flow control receive 1: Enable flow control receive	R/W	Note 14

Bits	Description	Type	Default
1	Port 0 Full-Duplex Transmit Flow Control Enable (TX_FC_0) When the MANUAL_FC_0 bit is set or Virtual Auto-Negotiation is disabled, this bit enables/disables full-duplex Pause packets to be generated on switch Port 0. 0: Disable flow control transmit 1: Enable flow control transmit	R/W	Note 14
0	Port 0 Full-Duplex Manual Flow Control Select (MANUAL_FC_0) This bit toggles flow control selection between manual and Auto-Negotiation. 0: If Auto-Negotiation is enabled, the Auto-Negotiation function determines the flow control of switch Port 0 (RX_FC_0 and TX_FC_0 values ignored). If Auto-Negotiation is disabled, the RX_FC_0 and TX_FC_0 values are used. 1: TX_FC_0 and RX_FC_0 bits determine the flow control of switch Port 0 when in full-duplex mode.	R/W	Note 15

Note 12: The default value of this field is determined by the [BP_EN_strap_0](#) configuration strap.

Note 13: The default value of this bit is determined by multiple strap settings.

Note 14: The default value of this field is determined by the [FD_FC_strap_0](#) configuration strap.

Note 15: The default value of this field is determined by the [manual_FC_strap_0](#) configuration strap.

13.6.4 SWITCH FABRIC CSR INTERFACE DATA REGISTER (SWITCH_CSR_DATA)

Offset: 1ACh Size: 32 bits

This read/write register is used in conjunction with the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) to perform read and write operations with the Switch Fabric CSRs. Refer to [Section 13.7, "Switch Fabric Control and Status Registers," on page 363](#) for details on the registers indirectly accessible via this register.

Bits	Description	Type	Default
31:0	<p>Switch CSR Data (CSR_DATA) This field contains the value read from or written to the Switch Fabric CSR. The Switch Fabric CSR is selected via the CSR Address (CSR_ADDR[15:0]) bits of the Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD).</p> <p>Upon a read, the value returned depends on the Read/Write (R_nW) bit in the Switch Fabric CSR Interface Command Register (SWITCH_CSR_CMD). If Read/Write (R_nW) is set, the data is from the switch fabric. If Read/Write (R_nW) is cleared, the data is the value that was last written into this register.</p>	R/W	00000000h

13.6.5 SWITCH FABRIC CSR INTERFACE COMMAND REGISTER (SWITCH_CSR_CMD)

Offset: 1B0h Size: 32 bits

This read/write register is used in conjunction with the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#) to control the read and write operations to the various Switch Fabric CSRs. Refer to [Section 13.7, "Switch Fabric Control and Status Registers," on page 363](#) for details on the registers indirectly accessible via this register.

Bits	Description	Type	Default
31	CSR Busy (CSR_BUSY) When a 1 is written to this bit, the read or write operation (as determined by the R_nW bit) is performed to the specified Switch Fabric CSR in CSR Address (CSR_ADDR[15:0]) . This bit will remain set until the operation is complete, at which time the bit will self-clear. In the case of a read, the clearing of this bit indicates to the Host that valid data can be read from the Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA) . The SWITCH_CSR_CMD and SWITCH_CSR_DATA registers should not be modified until this bit self-clears.	R/W SC	0b
30	Read/Write (R_nW) This bit determines whether a read or write operation is performed by the Host to the specified Switch Fabric CSR. 0: Write 1: Read	R/W	0b
29	Auto Increment (AUTO_INC) This bit enables/disables the auto increment feature. When this bit is set, a write to the Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA) will automatically set the CSR Busy (CSR_BUSY) bit. Once the write command is finished, the CSR Address (CSR_ADDR[15:0]) will automatically increment. When this bit is set, a read from the Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA) will automatically increment the CSR Address (CSR_ADDR[15:0]) and set the CSR Busy (CSR_BUSY) bit. This bit should be cleared by software before the last read from the SWITCH_CSR_DATA register. 0: Disable Auto Increment 1: Enable Auto Increment Note: This bit has precedence over the Auto Decrement (AUTO_DEC) bit.	R/W	0b

Bits	Description	Type	Default
28	<p>Auto Decrement (AUTO_DEC) This bit enables/disables the auto decrement feature.</p> <p>When this bit is set, a write to the Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA) will automatically set the CSR Busy (CSR_BUSY) bit. Once the write command is finished, the CSR Address (CSR_ADDR[15:0]) will automatically decrement.</p> <p>When this bit is set, a read from the Switch Fabric CSR Interface Data Register (SWITCH_CSR_DATA) will automatically decrement the CSR Address (CSR_ADDR[15:0]) and set the CSR Busy (CSR_BUSY) bit. This bit should be cleared by software before the last read from the SWITCH_CSR_DATA register.</p> <p>0: Disable Auto Decrement 1: Enable Auto Decrement</p>	R/W	0b
27:20	RESERVED	RO	-
19:16	<p>CSR Byte Enable (CSR_BE[3:0]) This field is a 4-bit byte enable used for selection of valid bytes during write operations. Bytes which are not selected will not be written to the corresponding Switch Fabric CSR.</p> <p>CSR_BE[3] corresponds to register data bits [31:24] CSR_BE[2] corresponds to register data bits [23:16] CSR_BE[1] corresponds to register data bits [15:8] CSR_BE[0] corresponds to register data bits [7:0]</p> <p>Typically all four-byte-enables should be set for auto increment and auto decrement operations.</p>	R/W	0h
15:0	<p>CSR Address (CSR_ADDR[15:0]) This field selects the 16-bit address of the Switch Fabric CSR that will be accessed with a read or write operation. Refer to Table 13-9, “Indirectly Accessible Switch Control and Status Registers,” on page 363 for a list of Switch Fabric CSR addresses.</p>	R/W	00h

13.6.6 SWITCH FABRIC MAC ADDRESS HIGH REGISTER (SWITCH_MAC_ADDRH)

Offset: 1F0h Size: 32 bits

This register contains the upper 16 bits of the MAC address used by the switch for Pause frames. This register is used in conjunction with [Switch Fabric MAC Address Low Register \(SWITCH_MAC_ADDRL\)](#). The contents of this register are optionally loaded from the EEPROM at power-on through the EEPROM Loader if a programmed EEPROM is detected. The least significant byte of this register (bits [7:0]) is loaded from address 05h of the EEPROM. The second byte (bits [15:8]) is loaded from address 06h of the EEPROM. The Host can update the contents of this field after the initialization process has completed.

Refer to [Section 13.6.7, "Switch Fabric MAC Address Low Register \(SWITCH_MAC_ADDRL\)"](#) for information on how this address is loaded by the EEPROM Loader. [Section 14.4, "EEPROM Loader," on page 465](#) contains additional details on using the EEPROM Loader.

Bits	Description	Type	Default
31:23	RESERVED	RO	-
22	DiffPauseAddr When set, each port may have a unique MAC address.	R/W	0b
21:20	Port 2 Physical Address [41:40] When DiffPauseAddr is set, these bits are used as bits 41 and 40 of the MAC Address for Port 2.	R/W	10b
19:18	Port 1 Physical Address [41:40] When DiffPauseAddr is set, these bits are used as bits 41 and 40 of the MAC Address for Port 1.	R/W	01b
17:16	Port 0 Physical Address [41:40] When DiffPauseAddr is set, these bits are used as bits 41 and 40 of the MAC Address for Port 0.	R/W	00b
15:0	Physical Address[47:32] This field contains the upper 16-bits (47:32) of the physical address of the Switch Fabric MACs. Bits 41 and 10 are ignored if DiffPauseAddr is set.	R/W	FFFFh

13.6.7 SWITCH FABRIC MAC ADDRESS LOW REGISTER (SWITCH_MAC_ADDRL)

Offset: 1F4h Size: 32 bits

This register contains the lower 32 bits of the MAC address used by the switch for Pause frames. This register is used in conjunction with [Switch Fabric MAC Address High Register \(SWITCH_MAC_ADDRH\)](#). The contents of this register are optionally loaded from the EEPROM at power-on through the EEPROM Loader if a programmed EEPROM is detected. The least significant byte of this register (bits [7:0]) is loaded from address 01h of the EEPROM. The most significant byte (bits [31:24]) is loaded from address 04h of the EEPROM. The Host can update the contents of this field after the initialization process has completed.

Refer to [Section 14.4, "EEPROM Loader," on page 465](#) for information on using the EEPROM Loader.

Bits	Description	Type	Default
31:0	Physical Address[31:0] This field contains the lower 32 bits (31:0) of the physical address of the Switch Fabric MACs.	R/W	FF0F8000h

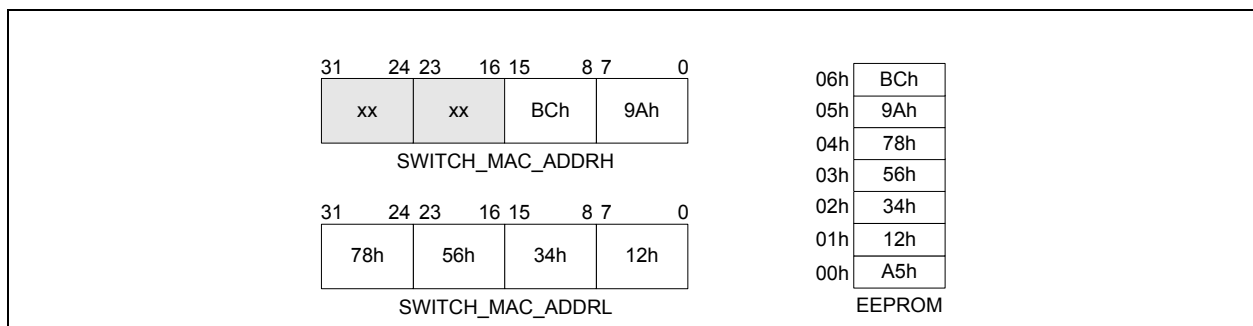
[Table 13-7](#) illustrates the byte ordering of the SWITCH_MAC_ADDRL and SWITCH_MAC_ADDRH registers with respect to the reception of the Ethernet physical address. Also shown is the correlation between the EEPROM addresses and the SWITCH_MAC_ADDRL and SWITCH_MAC_ADDRH registers.

TABLE 13-7: SWITCH_MAC_ADDRL, SWITCH_MAC_ADDRH AND EEPROM BYTE ORDERING

EEPROM Address	Register Location Written	Order of Reception on Ethernet
01h	SWITCH_MAC_ADDRL[7:0]	1 st
02h	SWITCH_MAC_ADDRL[15:8]	2 nd
03h	SWITCH_MAC_ADDRL[23:16]	3 rd
04h	SWITCH_MAC_ADDRL[31:24]	4 th
05h	SWITCH_MAC_ADDRH[7:0]	5 th
06h	SWITCH_MAC_ADDRH[15:8]	6 th

For example, if the desired Ethernet physical address is 12-34-56-78-9A-BC, the SWITCH_MAC_ADDRL and SWITCH_MAC_ADDRH registers would be programmed as shown in [Figure 13-10](#). The values required to automatically load this configuration from the EEPROM are also shown.

FIGURE 13-10: EXAMPLE SWITCH_MAC_ADDL, SWITCH_MAC_ADDRH AND EEPROM SETUP



Note: By convention, the right nibble of the left most byte of the Ethernet address (in this example, the 2 of the 12h) is the most significant nibble and is transmitted/received first.

13.6.8 SWITCH FABRIC CSR INTERFACE DIRECT DATA REGISTERS (SWITCH_CSR_DIRECT_DATA)

Offset: 200h-2F8h Size: 32 bits

This write-only register set is used to perform directly addressed write operations to the Switch Fabric CSRs. Using this set of registers, writes can be directly addressed to select Switch Fabric registers, as specified in [Table 13-8](#).

Writes within the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) address range automatically set the appropriate [CSR Address \(CSR_ADDR\[15:0\]\)](#), set the four [CSR Byte Enable \(CSR_BE\[3:0\]\)](#) bits, clear the [Read/Write \(R_nW\)](#) bit and set the [CSR Busy \(CSR_BUSY\)](#) bit in the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#). The completion of the write cycle is indicated when the [CSR Busy \(CSR_BUSY\)](#) bit self-clears. The address that is set in the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) is mapped via [TABLE 13-8](#). For more information on this method of writing to the Switch Fabric CSRs, refer to [Section 13.5.4, "Switch Fabric CSR Writes," on page 345](#).

Bits	Description	Type	Default
31:0	Switch CSR Data (CSR_DATA) This field contains the value to be written to the corresponding Switch Fabric register.	WO	00000000h

Note: This set of registers is for write operations only. Reads can be performed via the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#) and the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#) only.

TABLE 13-8: SWITCH FABRIC CSR TO SWITCH_CSR_DIRECT_DATA ADDRESS RANGE MAP

REGISTER NAME	SWITCH FABRIC CSR REGISTER #	SWITCH_CSR_DIRECT_DATA ADDRESS
General Switch CSRs		
SW_RESET	0001h	200h
SW_IMR	0004h	204h
Switch Port 0 CSRs		
MAC_RX_CFG_0	0401h	208h
MAC_TX_CFG_0	0440h	20Ch
MAC_TX_FC_SETTINGS_0	0441h	210h
MAC_IMR_0	0480h	214h
Switch Port 1 CSRs		
MAC_RX_CFG_1	0801h	218h
MAC_TX_CFG_1	0840h	21Ch
MAC_TX_FC_SETTINGS_1	0841h	220h
EEE_TW_TX_SYS_1	0842h	2E8h
EEE_TX_LPI_REQ_DELAY_CNT_1	0843h	2ECh
MAC_IMR_1	0880h	224h

TABLE 13-8: SWITCH FABRIC CSR TO SWITCH_CSR_DIRECT_DATA ADDRESS RANGE MAP

REGISTER NAME	SWITCH FABRIC CSR REGISTER #	SWITCH_CSR_DIRECT_DATA ADDRESS
Switch Port 2 CSRs		
MAC_RX_CFG_2	0C01h	228h
MAC_TX_CFG_2	0C40h	22Ch
MAC_TX_FC_SETTINGS_2	0C41h	230h
EEE_TW_TX_SYS_2	0C42h	2F0h
EEE_TX_LPI_REQ_DELAY_CNT_2	0C43h	2F4h
MAC_IMR_2	0C80h	234h
Switch Engine CSRs		
SWE_ALR_CMD	1800h	238h
SWE_ALR_WR_DAT_0	1801h	23Ch
SWE_ALR_WR_DAT_1	1802h	240h
SWE_ALR_CFG	1809h	244h
SWE_ALR_OVERRIDE	180Ah	2F8h
SWE_VLAN_CMD	180Bh	248h
SWE_VLAN_WR_DATA	180Ch	24Ch
SWE_DIFFSERV_TBL_CMD	1811h	250h
SWE_DIFFSERV_TBL_WR_DATA	1812h	254h
SWE_GLB_INGRESS_CFG	1840h	258h
SWE_PORT_INGRESS_CFG	1841h	25Ch
SWE_ADMT_ONLY_VLAN	1842h	260h
SWE_PORT_STATE	1843h	264h
SWE_PRI_TO_QUE	1845h	268h
SWE_PORT_MIRROR	1846h	26Ch
SWE_INGRESS_PORT_TYP	1847h	270h
SWE_BCST_THROT	1848h	274h
SWE_ADMT_N_MEMBER	1849h	278h
SWE_INGRESS_RATE_CFG	184Ah	27Ch
SWE_INGRESS_RATE_CMD	184Bh	280h
SWE_INGRESS_RATE_WR_DATA	184Dh	284h
SWE_INGRESS_REGEN_TBL_0	1855h	288h
SWE_INGRESS_REGEN_TBL_1	1856h	28Ch
SWE_INGRESS_REGEN_TBL_2	1857h	290h

TABLE 13-8: SWITCH FABRIC CSR TO SWITCH_CSR_DIRECT_DATA ADDRESS RANGE MAP

REGISTER NAME	SWITCH FABRIC CSR REGISTER #	SWITCH_CSR_DIRECT_DATA ADDRESS
SWE_IMR	1880h	294h
Buffer Manager (BM) CSRs		
BM_CFG	1C00h	298h
BM_DROP_LVL	1C01h	29Ch
BM_FC_PAUSE_LVL	1C02h	2A0h
BM_FC_RESUME_LVL	1C03h	2A4h
BM_BCST_LVL	1C04h	2A8h
BM_RNDM_DSCRD_TBL_CMD	1C09h	2ACh
BM_RNDM_DSCRD_TBL_WDATA	1C0Ah	2B0h
BM_EGRSS_PORT_TYPE	1C0Ch	2B4h
BM_EGRSS_RATE_00_01	1C0Dh	2B8h
BM_EGRSS_RATE_02_03	1C0Eh	2BCh
BM_EGRSS_RATE_10_11	1C0Fh	2C0h
BM_EGRSS_RATE_12_13	1C10h	2C4h
BM_EGRSS_RATE_20_21	1C11h	2C8h
BM_EGRSS_RATE_22_23	1C12h	2CCh
BM_VLAN_0	1C13h	2D0h
BM_VLAN_1	1C14h	2D4h
BM_VLAN_2	1C15h	2D8h
BM_IMR	1C20h	2DCh

13.7 Switch Fabric Control and Status Registers

This section details the various indirectly addressable switch control and status registers that reside within the Switch Fabric. The switch control and status registers allow configuration of each individual switch port, the Switch Engine and Buffer Manager. Switch Fabric related interrupts and resets are also controlled and monitored via the switch CSRs.

The switch CSRs are not directly mapped into the system address space. All switch CSRs are accessed indirectly via the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#), the [Switch Fabric CSR Interface Data Register \(SWITCH_CSR_DATA\)](#) and the [Switch Fabric CSR Interface Direct Data Registers \(SWITCH_CSR_DIRECT_DATA\)](#) in the system CSR address space. All accesses to the switch CSRs must be performed through these registers. Refer to [Section 13.6, "Switch Fabric Interface Logic Registers"](#) for additional information.

Note: The flow control settings of the switch ports are configured via the [Switch Fabric Interface Logic Registers: Port 1 Manual Flow Control Register \(MANUAL_FC_1\)](#), [Port 2 Manual Flow Control Register \(MANUAL_FC_2\)](#) and [Port 0 Manual Flow Control Register \(MANUAL_FC_0\)](#) located in the system CSR address space.

[Table 13-9](#) lists the Switch CSRs and their corresponding addresses in order. The Switch Fabric registers can be categorized into the following sub-sections:

- [Section 13.7.1, "General Switch CSRs," on page 372](#)
- [Section 13.7.2, "Switch Port 0, Port 1 and Port 2 CSRs," on page 376](#)
- [Section 13.7.3, "Switch Engine CSRs," on page 403](#)
- [Section 13.7.4, "Buffer Manager CSRs," on page 440](#)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
General Switch CSRs	
0000h	Switch Device ID Register (SW_DEV_ID)
0001h	Switch Reset Register (SW_RESET)
0002h-0003h	Reserved for Future Use (RESERVED)
0004h	Switch Global Interrupt Mask Register (SW_IMR)
0005h	Switch Global Interrupt Pending Register (SW_IPR)
0006h-03FFh	Reserved for Future Use (RESERVED)
Switch Port 0 CSRs (x=0)	
0400h	Port x MAC Version ID Register (MAC_VER_ID_x)
0401h	Port x MAC Receive Configuration Register (MAC_RX_CFG_x)
0402h-040Fh	Reserved for Future Use (RESERVED)
0410h	Port x MAC Receive Undersize Count Register (MAC_RX_UNDSIZE_CNT_x)
0411h	Port x MAC Receive 64 Byte Count Register (MAC_RX_64_CNT_x)
0412h	Port x MAC Receive 65 to 127 Byte Count Register (MAC_RX_65_TO_127_CNT_x)
0413h	Port x MAC Receive 128 to 255 Byte Count Register (MAC_RX_128_TO_255_CNT_x)
0414h	Port x MAC Receive 256 to 511 Byte Count Register (MAC_RX_256_TO_511_CNT_x)
0415h	Port x MAC Receive 512 to 1023 Byte Count Register (MAC_RX_512_TO_1023_CNT_x)
0416h	Port x MAC Receive 1024 to Max Byte Count Register (MAC_RX_1024_TO_MAX_CNT_x)
0417h	Port x MAC Receive Oversize Count Register (MAC_RX_OVRSIZE_CNT_x)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
0418h	Port x MAC Receive OK Count Register (MAC_RX_PKTOK_CNT_x)
0419h	Port x MAC Receive CRC Error Count Register (MAC_RX_CRCERR_CNT_x)
041Ah	Port x MAC Receive Multicast Count Register (MAC_RX_MULCST_CNT_x)
041Bh	Port x MAC Receive Broadcast Count Register (MAC_RX_BRDCST_CNT_x)
041Ch	Port x MAC Receive Pause Frame Count Register (MAC_RX_PAUSE_CNT_x)
041Dh	Port x MAC Receive Fragment Error Count Register (MAC_RX_FRAG_CNT_x)
041Eh	Port x MAC Receive Jabber Error Count Register (MAC_RX_JABB_CNT_x)
041Fh	Port x MAC Receive Alignment Error Count Register (MAC_RX_ALIGN_CNT_x)
0420h	Port x MAC Receive Packet Length Count Register (MAC_RX_PKTLEN_CNT_x)
0421h	Port x MAC Receive Good Packet Length Count Register (MAC_RX_GOODPKTLEN_CNT_x)
0422h	Port x MAC Receive Symbol Error Count Register (MAC_RX_SYMBOL_CNT_x)
0423h	Port x MAC Receive Control Frame Count Register (MAC_RX_CTLFRM_CNT_x)
0424h-043Fh	Reserved for Future Use (RESERVED)
0440h	Port x MAC Transmit Configuration Register (MAC_TX_CFG_x)
0441h	Port x MAC Transmit Flow Control Settings Register (MAC_TX_FC_SETTINGS_x)
0442h-0450h	Reserved for Future Use (RESERVED)
0451h	Port x MAC Transmit Deferred Count Register (MAC_TX_DEFER_CNT_x)
0452h	Port x MAC Transmit Pause Count Register (MAC_TX_PAUSE_CNT_x)
0453h	Port x MAC Transmit OK Count Register (MAC_TX_PKTOK_CNT_x)
0454h	Port x MAC Transmit 64 Byte Count Register (MAC_TX_64_CNT_x)
0455h	Port x MAC Transmit 65 to 127 Byte Count Register (MAC_TX_65_TO_127_CNT_x)
0456h	Port x MAC Transmit 128 to 255 Byte Count Register (MAC_TX_128_TO_255_CNT_x)
0457h	Port x MAC Transmit 256 to 511 Byte Count Register (MAC_TX_256_TO_511_CNT_x)
0458h	Port x MAC Transmit 512 to 1023 Byte Count Register (MAC_TX_512_TO_1023_CNT_x)
0459h	Port x MAC Transmit 1024 to Max Byte Count Register (MAC_TX_1024_TO_MAX_CNT_x)
045Ah	Port x MAC Transmit Undersize Count Register (MAC_TX_UNDSZE_CNT_x)
045Bh	Reserved for Future Use (RESERVED)
045Ch	Port x MAC Transmit Packet Length Count Register (MAC_TX_PKTLEN_CNT_x)
045Dh	Port x MAC Transmit Broadcast Count Register (MAC_TX_BRDCST_CNT_x)
045Eh	Port x MAC Transmit Multicast Count Register (MAC_TX_MULCST_CNT_x)
045Fh	Port x MAC Transmit Late Collision Count Register (MAC_TX_LATECOL_CNT_x)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
0460h	Port x MAC Transmit Excessive Collision Count Register (MAC_TX_EXCCOL_CNT_x)
0461h	Port x MAC Transmit Single Collision Count Register (MAC_TX_SNGLECOL_CNT_x)
0462h	Port x MAC Transmit Multiple Collision Count Register (MAC_TX_MULTICOL_CNT_x)
0463h	Port x MAC Transmit Total Collision Count Register (MAC_TX_TOTALCOL_CNT_x)
0464h-047Fh	Reserved for Future Use (RESERVED)
0480h	Port x MAC Interrupt Mask Register (MAC_IMR_x)
0481h	Port x MAC Interrupt Pending Register (MAC_IPR_x)
0482h-07FFh	Reserved for Future Use (RESERVED)
Switch Port 1 CSRs (x=1)	
0800h	Port x MAC Version ID Register (MAC_VER_ID_x)
0801h	Port x MAC Receive Configuration Register (MAC_RX_CFG_x)
0802h-080Fh	Reserved for Future Use (RESERVED)
0810h	Port x MAC Receive Undersize Count Register (MAC_RX_UNDSIZE_CNT_x)
0811h	Port x MAC Receive 64 Byte Count Register (MAC_RX_64_CNT_x)
0812h	Port x MAC Receive 65 to 127 Byte Count Register (MAC_RX_65_TO_127_CNT_x)
0813h	Port x MAC Receive 128 to 255 Byte Count Register (MAC_RX_128_TO_255_CNT_x)
0814h	Port x MAC Receive 256 to 511 Byte Count Register (MAC_RX_256_TO_511_CNT_x)
0815h	Port x MAC Receive 512 to 1023 Byte Count Register (MAC_RX_512_TO_1023_CNT_x)
0816h	Port x MAC Receive 1024 to Max Byte Count Register (MAC_RX_1024_TO_MAX_CNT_x)
0817h	Port x MAC Receive Oversize Count Register (MAC_RX_OVRSIZE_CNT_x)
0818h	Port x MAC Receive OK Count Register (MAC_RX_PKTOK_CNT_x)
0819h	Port x MAC Receive CRC Error Count Register (MAC_RX_CRCERR_CNT_x)
081Ah	Port x MAC Receive Multicast Count Register (MAC_RX_MULCST_CNT_x)
081Bh	Port x MAC Receive Broadcast Count Register (MAC_RX_BRDCST_CNT_x)
081Ch	Port x MAC Receive Pause Frame Count Register (MAC_RX_PAUSE_CNT_x)
081Dh	Port x MAC Receive Fragment Error Count Register (MAC_RX_FRAG_CNT_x)
081Eh	Port x MAC Receive Jabber Error Count Register (MAC_RX_JABB_CNT_x)
081Fh	Port x MAC Receive Alignment Error Count Register (MAC_RX_ALIGN_CNT_x)
0820h	Port x MAC Receive Packet Length Count Register (MAC_RX_PKTLEN_CNT_x)
0821h	Port x MAC Receive Good Packet Length Count Register (MAC_RX_GOODPKTLEN_CNT_x)
0822h	Port x MAC Receive Symbol Error Count Register (MAC_RX_SYMBOL_CNT_x)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
0823h	Port x MAC Receive Control Frame Count Register (MAC_RX_CTLFRM_CNT_x)
0824h	Port x RX LPI Transitions Register (RX_LPI_TRANSITION_x)
0825h	Port x RX LPI Time Register (RX_LPI_TIME_x)
0826h-083Fh	Reserved for Future Use (RESERVED)
0840h	Port x MAC Transmit Configuration Register (MAC_TX_CFG_x)
0841h	Port x MAC Transmit Flow Control Settings Register (MAC_TX_FC_SETTINGS_x)
0842h	Port x EEE Time Wait TX System Register (EEE_TW_TX_SYS_x)
0843h	Port x EEE TX LPI Request Delay Register (EEE_TX_LPI_REQ_DELAY_x)
0844h-0850h	Reserved for Future Use (RESERVED)
0851h	Port x MAC Transmit Deferred Count Register (MAC_TX_DEFER_CNT_x)
0852h	Port x MAC Transmit Pause Count Register (MAC_TX_PAUSE_CNT_x)
0853h	Port x MAC Transmit OK Count Register (MAC_TX_PKTOK_CNT_x)
0854h	Port x MAC Transmit 64 Byte Count Register (MAC_TX_64_CNT_x)
0855h	Port x MAC Transmit 65 to 127 Byte Count Register (MAC_TX_65_TO_127_CNT_x)
0856h	Port x MAC Transmit 128 to 255 Byte Count Register (MAC_TX_128_TO_255_CNT_x)
0857h	Port x MAC Transmit 256 to 511 Byte Count Register (MAC_TX_256_TO_511_CNT_x)
0858h	Port x MAC Transmit 512 to 1023 Byte Count Register (MAC_TX_512_TO_1023_CNT_x)
0859h	Port x MAC Transmit 1024 to Max Byte Count Register (MAC_TX_1024_TO_MAX_CNT_x)
085Ah	Port x MAC Transmit Undersize Count Register (MAC_TX_UNDSZE_CNT_x)
085Bh	Reserved for Future Use (RESERVED)
085Ch	Port x MAC Transmit Packet Length Count Register (MAC_TX_PKTLEN_CNT_x)
085Dh	Port x MAC Transmit Broadcast Count Register (MAC_TX_BRDCST_CNT_x)
085Eh	Port x MAC Transmit Multicast Count Register (MAC_TX_MULCST_CNT_x)
085Fh	Port x MAC Transmit Late Collision Count Register (MAC_TX_LATECOL_CNT_x)
0860h	Port x MAC Transmit Excessive Collision Count Register (MAC_TX_EXCCOL_CNT_x)
0861h	Port x MAC Transmit Single Collision Count Register (MAC_TX_SNGLECOL_CNT_x)
0862h	Port x MAC Transmit Multiple Collision Count Register (MAC_TX_MULTICOL_CNT_x)
0863h	Port x MAC Transmit Total Collision Count Register (MAC_TX_TOTALCOL_CNT_x)
0864h	Port x TX LPI Transitions Register (TX_LPI_TRANSITION_x)
0865h	Port x TX LPI Time Register (TX_LPI_TIME_x)
0866h-087Fh	Reserved for Future Use (RESERVED)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
0880h	Port x MAC Interrupt Mask Register (MAC_IMR_x)
0881h	Port x MAC Interrupt Pending Register (MAC_IPR_x)
0882h-0BFFh	Reserved for Future Use (RESERVED)
Switch Port 2 CSRs (x=2)	
0C00h	Port x MAC Version ID Register (MAC_VER_ID_x)
0C01h	Port x MAC Receive Configuration Register (MAC_RX_CFG_x)
0C02h-0C0Fh	Reserved for Future Use (RESERVED)
0C10h	Port x MAC Receive Undersize Count Register (MAC_RX_UNDSIZE_CNT_x)
0C11h	Port x MAC Receive 64 Byte Count Register (MAC_RX_64_CNT_x)
0C12h	Port x MAC Receive 65 to 127 Byte Count Register (MAC_RX_65_TO_127_CNT_x)
0C13h	Port x MAC Receive 128 to 255 Byte Count Register (MAC_RX_128_TO_255_CNT_x)
0C14h	Port x MAC Receive 256 to 511 Byte Count Register (MAC_RX_256_TO_511_CNT_x)
0C15h	Port x MAC Receive 512 to 1023 Byte Count Register (MAC_RX_512_TO_1023_CNT_x)
0C16h	Port x MAC Receive 1024 to Max Byte Count Register (MAC_RX_1024_TO_MAX_CNT_x)
0C17h	Port x MAC Receive Oversize Count Register (MAC_RX_OVRSIZE_CNT_x)
0C18h	Port x MAC Receive OK Count Register (MAC_RX_PKTOK_CNT_x)
0C19h	Port x MAC Receive CRC Error Count Register (MAC_RX_CRCERR_CNT_x)
0C1Ah	Port x MAC Receive Multicast Count Register (MAC_RX_MULCST_CNT_x)
0C1Bh	Port x MAC Receive Broadcast Count Register (MAC_RX_BRDCST_CNT_x)
0C1Ch	Port x MAC Receive Pause Frame Count Register (MAC_RX_PAUSE_CNT_x)
0C1Dh	Port x MAC Receive Fragment Error Count Register (MAC_RX_FRAG_CNT_x)
0C1Eh	Port x MAC Receive Jabber Error Count Register (MAC_RX_JABB_CNT_x)
0C1Fh	Port x MAC Receive Alignment Error Count Register (MAC_RX_ALIGN_CNT_x)
0C20h	Port x MAC Receive Packet Length Count Register (MAC_RX_PKTLEN_CNT_x)
0C21h	Port x MAC Receive Good Packet Length Count Register (MAC_RX_GOODPKTLEN_CNT_x)
0C22h	Port x MAC Receive Symbol Error Count Register (MAC_RX_SYMBOL_CNT_x)
0C23h	Port x MAC Receive Control Frame Count Register (MAC_RX_CTLFRM_CNT_x)
0C24h	Port x RX LPI Transitions Register (RX_LPI_TRANSITION_x)
0C25h	Port x RX LPI Time Register (RX_LPI_TIME_x)
0C26h-0C3Fh	Reserved for Future Use (RESERVED)
0C40h	Port x MAC Transmit Configuration Register (MAC_TX_CFG_x)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
0C41h	Port x MAC Transmit Flow Control Settings Register (MAC_TX_FC_SETTINGS_x)
0C42h	Port x EEE Time Wait TX System Register (EEE_TW_TX_SYS_x)
0C43h	Port x EEE TX LPI Request Delay Register (EEE_TX_LPI_REQ_DELAY_x)
0C44h-0C50h	Reserved for Future Use (RESERVED)
0C51h	Port x MAC Transmit Deferred Count Register (MAC_TX_DEFER_CNT_x)
0C52h	Port x MAC Transmit Pause Count Register (MAC_TX_PAUSE_CNT_x)
0C53h	Port x MAC Transmit OK Count Register (MAC_TX_PKTOK_CNT_x)
0C54h	Port x MAC Transmit 64 Byte Count Register (MAC_TX_64_CNT_x)
0C55h	Port x MAC Transmit 65 to 127 Byte Count Register (MAC_TX_65_TO_127_CNT_x)
0C56h	Port x MAC Transmit 128 to 255 Byte Count Register (MAC_TX_128_TO_255_CNT_x)
0C57h	Port x MAC Transmit 256 to 511 Byte Count Register (MAC_TX_256_TO_511_CNT_x)
0C58h	Port x MAC Transmit 512 to 1023 Byte Count Register (MAC_TX_512_TO_1023_CNT_x)
0C59h	Port x MAC Transmit 1024 to Max Byte Count Register (MAC_TX_1024_TO_MAX_CNT_x)
0C5Ah	Port x MAC Transmit Undersize Count Register (MAC_TX_UNDSIZE_CNT_x)
0C5Bh	Reserved for Future Use (RESERVED)
0C5Ch	Port x MAC Transmit Packet Length Count Register (MAC_TX_PKTLEN_CNT_x)
0C5Dh	Port x MAC Transmit Broadcast Count Register (MAC_TX_BRDCST_CNT_x)
0C5Eh	Port x MAC Transmit Multicast Count Register (MAC_TX_MULCST_CNT_x)
0C5Fh	Port x MAC Transmit Late Collision Count Register (MAC_TX_LATECOL_CNT_x)
0C60h	Port x MAC Transmit Excessive Collision Count Register (MAC_TX_EXCCOL_CNT_x)
0C61h	Port x MAC Transmit Single Collision Count Register (MAC_TX_SNGLECOL_CNT_x)
0C62h	Port x MAC Transmit Multiple Collision Count Register (MAC_TX_MULTICOL_CNT_x)
0C63h	Port x MAC Transmit Total Collision Count Register (MAC_TX_TOTALCOL_CNT_x)
0C64h	Port x TX LPI Transitions Register (TX_LPI_TRANSITION_x)
0C65h	Port x TX LPI Time Register (TX_LPI_TIME_x)
0C66h-0C7Fh	Reserved for Future Use (RESERVED)
0C80h	Port x MAC Interrupt Mask Register (MAC_IMR_x)
0C81h	Port x MAC Interrupt Pending Register (MAC_IPR_x)
0C82h-17FFh	Reserved for Future Use (RESERVED)
Switch Engine CSRs	
1800h	Switch Engine ALR Command Register (SWE_ALR_CMD)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
1801h	Switch Engine ALR Write Data 0 Register (SWE_ALR_WR_DAT_0)
1802h	Switch Engine ALR Write Data 1 Register (SWE_ALR_WR_DAT_1)
1803h-1804h	Reserved for Future Use (RESERVED)
1805h	Switch Engine ALR Read Data 0 Register (SWE_ALR_RD_DAT_0)
1806h	Switch Engine ALR Read Data 1 Register (SWE_ALR_RD_DAT_1)
1807h	Reserved for Future Use (RESERVED)
1808h	Switch Engine ALR Command Status Register (SWE_ALR_CMD_STS)
1809h	Switch Engine ALR Configuration Register (SWE_ALR_CFG)
180Ah	Switch Engine ALR Override Register (SWE_ALR_OVERRIDE)
180Bh	Switch Engine VLAN Command Register (SWE_VLAN_CMD)
180Ch	Switch Engine VLAN Write Data Register (SWE_VLAN_WR_DATA)
180Dh	Reserved for Future Use (RESERVED)
180Eh	Switch Engine VLAN Read Data Register (SWE_VLAN_RD_DATA)
180Fh	Reserved for Future Use (RESERVED)
1810h	Switch Engine VLAN Command Status Register (SWE_VLAN_CMD_STS)
1811h	Switch Engine DIFFSERV Table Command Register (SWE_DIFFSERV_TBL_CFG)
1812h	Switch Engine DIFFSERV Table Write Data Register (SWE_DIFFSERV_TBL_WR_DATA)
1813h	Switch Engine DIFFSERV Table Read Data Register (SWE_DIFFSERV_TBL_RD_DATA)
1814h	Switch Engine DIFFSERV Table Command Status Register (SWE_DIFFSERV_TBL_CMD_STS)
1815h-183Fh	Reserved for Future Use (RESERVED)
1840h	Switch Engine Global Ingress Configuration Register (SWE_GLOBAL INGRSS_CFG)
1841h	Switch Engine Port Ingress Configuration Register (SWE_PORT INGRSS_CFG)
1842h	Switch Engine Admit Only VLAN Register (SWE_ADMT_ONLY_VLAN)
1843h	Switch Engine Port State Register (SWE_PORT_STATE)
1844h	Reserved for Future Use (RESERVED)
1845h	Switch Engine Priority to Queue Register (SWE_PRI_TO_QUE)
1846h	Switch Engine Port Mirroring Register (SWE_PORT_MIRROR)
1847h	Switch Engine Ingress Port Type Register (SWE INGRSS_PORT_TYP)
1848h	Switch Engine Broadcast Throttling Register (SWE_BCST_THROT)
1849h	Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER)
184Ah	Switch Engine Ingress Rate Configuration Register (SWE INGRSS_RATE_CFG)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
184Bh	Switch Engine Ingress Rate Command Register (SWE_INGRSS_RATE_CMD)
184Ch	Switch Engine Ingress Rate Command Status Register (SWE_INGRSS_RATE_CMD_STS)
184Dh	Switch Engine Ingress Rate Write Data Register (SWE_INGRSS_RATE_WR_DATA)
184Eh	Switch Engine Ingress Rate Read Data Register (SWE_INGRSS_RATE_RD_DATA)
184Fh	Reserved for Future Use (RESERVED)
1850h	Switch Engine Port 0 Ingress Filtered Count Register (SWE_FILTERED_CNT_0)
1851h	Switch Engine Port 1 Ingress Filtered Count Register (SWE_FILTERED_CNT_1)
1852h	Switch Engine Port 2 Ingress Filtered Count Register (SWE_FILTERED_CNT_2)
1853h-1854h	Reserved for Future Use (RESERVED)
1855h	Switch Engine Port 0 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_RE- GEN_TBL_0)
1856h	Switch Engine Port 1 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_RE- GEN_TBL_1)
1857h	Switch Engine Port 2 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_RE- GEN_TBL_2)
1858h	Switch Engine Port 0 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_0)
1859h	Switch Engine Port 1 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_1)
185Ah	Switch Engine Port 2 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_2)
185Bh-187Fh	Reserved for Future Use (RESERVED)
1880h	Switch Engine Interrupt Mask Register (SWE_IMR)
1881h	Switch Engine Interrupt Pending Register (SWE_IPR)
1882h-1BFFh	Reserved for Future Use (RESERVED)
Buffer Manager (BM) CSRs	
1C00h	Buffer Manager Configuration Register (BM_CFG)
1C01h	Buffer Manager Drop Level Register (BM_DROP_LVL)
1C02h	Buffer Manager Flow Control Pause Level Register (BM_FC_PAUSE_LVL)
1C03h	Buffer Manager Flow Control Resume Level Register (BM_FC_RESUME_LVL)
1C04h	Buffer Manager Broadcast Buffer Level Register (BM_BCST_LVL)
1C05h	Buffer Manager Port 0 Drop Count Register (BM_DRP_CNT_SRC_0)
1C06h	Buffer Manager Port 1 Drop Count Register (BM_DRP_CNT_SRC_1)
1C07h	Buffer Manager Port 2 Drop Count Register (BM_DRP_CNT_SRC_2)
1C08h	Buffer Manager Reset Status Register (BM_RST_STS)
1C09h	Buffer Manager Random Discard Table Command Register (BM_RNDM_DSCRD_TBL_CMD)

TABLE 13-9: INDIRECTLY ACCESSIBLE SWITCH CONTROL AND STATUS REGISTERS

Address (INDIRECT)	Register Name (SYMBOL)
1C0Ah	Buffer Manager Random Discard Table Write Data Register (BM_RNDM_DSCRD_TBL_WDATA)
1C0Bh	Buffer Manager Random Discard Table Read Data Register (BM_RNDM_DSCRD_TBL_RDATA)
1C0Ch	Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE)
1C0Dh	Buffer Manager Port 0 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_00_01)
1C0Eh	Buffer Manager Port 0 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_02_03)
1C0Fh	Buffer Manager Port 1 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_10_11)
1C10h	Buffer Manager Port 1 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_12_13)
1C11h	Buffer Manager Port 2 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_20_21)
1C12h	Buffer Manager Port 2 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_22_23)
1C13h	Buffer Manager Port 0 Default VLAN ID and Priority Register (BM_VLAN_0)
1C14h	Buffer Manager Port 1 Default VLAN ID and Priority Register (BM_VLAN_1)
1C15h	Buffer Manager Port 2 Default VLAN ID and Priority Register (BM_VLAN_2)
1C16h	Buffer Manager Port 0 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_0)
1C17h	Buffer Manager Port 1 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_1)
1C18h	Buffer Manager Port 2 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_2)
1C19h-1C1Fh	Reserved for Future Use (RESERVED)
1C20h	Buffer Manager Interrupt Mask Register (BM_IMR)
1C21h	Buffer Manager Interrupt Pending Register (BM_IPR)
1C22h-FFFFh	Reserved for Future Use (RESERVED)

13.7.1 GENERAL SWITCH CSRS

This section details the general Switch Fabric CSRs. These registers control the main reset and interrupt functions of the Switch Fabric. A list of the general switch CSRs and their corresponding register numbers is included in [Table 13-9](#).

13.7.1.1 Switch Device ID Register (SW_DEV_ID)

Register #: [0000h](#) Size: 32 bits

This read-only register contains switch device ID information, including the device type, chip version and revision codes.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:16	Device Type Code (DEVICE_TYPE)	RO	03h
15:8	Chip Version Code (CHIP_VERSION)	RO	06h
7:0	Revision Code (REVISION)	RO	07h

13.7.1.2 Switch Reset Register (SW_RESET)

Register #: 0001h Size: 32 bits

This register contains the Switch Fabric global reset. Refer to the [Switch Reset](#) portion of [Section 6.2, "Resets,"](#) on [page 42](#) for more information.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Switch Fabric Reset (SW_RESET) This bit is the global switch fabric reset. All switch fabric blocks are affected. This bit must be manually cleared by software.	WO	0b

13.7.1.3 Switch Global Interrupt Mask Register (SW_IMR)

Register #: 0004h Size: 32 bits

This read/write register contains the global interrupt mask for the Switch Fabric interrupts. All switch related interrupts in the [Switch Global Interrupt Pending Register \(SW_IPR\)](#) may be masked via this register. An interrupt is masked by setting the corresponding bit of this register. Clearing a bit will unmask the interrupt. When an unmasked Switch Fabric interrupt is generated in the [Switch Global Interrupt Pending Register \(SW_IPR\)](#), the interrupt will trigger the [Switch Fabric Interrupt Event \(SWITCH_INT\)](#) bit in the [Interrupt Status Register \(INT_STS\)](#). Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default
31:9	RESERVED	RO	-
8:7	RESERVED Note: These bits must be written as 11b.	R/W	11b
6	Buffer Manager Interrupt Mask (BM) When set, prevents the generation of Switch Fabric interrupts due to the Buffer Manager via the Buffer Manager Interrupt Pending Register (BM_IPR) . The status bits in the Switch Global Interrupt Pending Register (SW_IPR) are not affected.	R/W	1b
5	Switch Engine Interrupt Mask (SWE) When set, prevents the generation of Switch Fabric interrupts due to the Switch Engine via the Switch Engine Interrupt Pending Register (SWE_IPR) . The status bits in the Switch Global Interrupt Pending Register (SW_IPR) are not affected.	R/W	1b
4:3	RESERVED Note: These bits must be written as 11b.	R/W	11b
2	Port 2 MAC Interrupt Mask (MAC_2) When set, prevents the generation of Switch Fabric interrupts due to the Port 2 MAC via the MAC_IPR_2 register (see Section 13.7.2.50, on page 402). The status bits in the Switch Global Interrupt Pending Register (SW_IPR) are not affected.	R/W	1b
1	Port 1 MAC Interrupt Mask (MAC_1) When set, prevents the generation of Switch Fabric interrupts due to the Port 1 MAC via the MAC_IPR_1 register (see Section 13.7.2.50, on page 402). The status bits in the Switch Global Interrupt Pending Register (SW_IPR) are not affected.	R/W	1b
0	Port 0 MAC Interrupt Mask (MAC_0) When set, prevents the generation of Switch Fabric interrupts due to the Port 0 MAC via the MAC_IPR_0 register (see Section 13.7.2.50, on page 402). The status bits in the Switch Global Interrupt Pending Register (SW_IPR) are not affected.	R/W	1b

13.7.1.4 Switch Global Interrupt Pending Register (SW_IPR)

Register #: 0005h Size: 32 bits

This read-only register contains the pending global interrupts for the Switch Fabric. A set bit indicates an unmasked bit in the corresponding Switch Fabric sub-system has been triggered. All switch-related interrupts in this register may be masked via the [Switch Global Interrupt Mask Register \(SW_IMR\)](#). When an unmasked Switch Fabric interrupt is generated in this register, the interrupt will trigger the [Switch Fabric Interrupt Event \(SWITCH_INT\)](#) bit in the [Interrupt Status Register \(INT_STS\)](#). Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6	Buffer Manager Interrupt (BM) Set when any unmasked bit in the Buffer Manager Interrupt Pending Register (BM_IPR) is triggered. A read of this register clears this bit.	RC	0b
5	Switch Engine Interrupt (SWE) Set when any unmasked bit in the Switch Engine Interrupt Pending Register (SWE_IPR) is triggered. A read of this register clears this bit.	RC	0b
4:3	RESERVED	RO	-
2	Port 2 MAC Interrupt (MAC_2) Set when any unmasked bit in the MAC_IPR_2 register (see Section 13.7.2.50, on page 402) is triggered. A read of this register clears this bit.	RC	0b
1	Port 1 MAC Interrupt (MAC_1) Set when any unmasked bit in the MAC_IPR_1 register (see Section 13.7.2.50, on page 402) is triggered. A read of this register clears this bit.	RC	0b
0	Port 0 MAC Interrupt (MAC_0) Set when any unmasked bit in the MAC_IPR_0 register (see Section 13.7.2.50, on page 402) is triggered. A read of this register clears this bit.	RC	0b

13.7.2 SWITCH PORT 0, PORT 1 AND PORT 2 CSRS

This section details the switch Port 0, Port 1 and Port 2 CSRs. Each port provides a functionally identical set of registers which allow for the configuration of port settings, interrupts and the monitoring of the various packet counters.

Because the Port 0, Port 1 and Port 2 CSRs are functionally identical, their register descriptions have been consolidated. A lowercase “x” has been appended to the end of each switch port register name in this section, where “x” should be replaced with “0”, “1” or “2” for the Port 0, Port 1 or Port 2 registers respectively. A list of the Switch Port 0, Port 1 and Port 2 registers and their corresponding register numbers is included in [TABLE 13-9](#).

13.7.2.1 Port x MAC Version ID Register (MAC_VER_ID_x)

Register #: Port0: 0400h Size: 32 bits
 Port1: 0800h
 Port2: 0C00h

This read-only register contains switch device ID information, including the device type, chip version and revision codes.

Bits	Description	Type	Default
31:12	RESERVED	RO	-
11:8	Device Type Code (DEVICE_TYPE)	RO	5h
7:4	Chip Version Code (CHIP_VERSION)	RO	9h
3:0	Revision Code (REVISION)	RO	3h

13.7.2.2 Port x MAC Receive Configuration Register (MAC_RX_CFG_x)

Register #: Port0: 0401h Size: 32 bits
 Port1: 0801h
 Port2: 0C01h

This read/write register configures the packet type passing parameters of the port.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7	RESERVED Note: This bit must always be written as 0.	R/W	0b
6	RESERVED	RO	-
5	Enable Receive Own Transmit When set, the switch port will receive its own transmission if it is looped back from the PHY. Normally, this function is only used in half-duplex PHY loop-back.	R/W	0b
4	RESERVED	RO	-
3	Jumbo2K When set, the maximum packet size accepted is 2048 bytes. Statistics boundaries are also adjusted.	R/W	0b
2	RESERVED	RO	-
1	Reject MAC Types When set, MAC control frames (packets with a type field of 8808h) are filtered. When cleared, MAC Control frames, other than MAC Control Pause frames, are sent to the forwarding process. MAC Control Pause frames are always consumed by the switch.	R/W	1b
0	RX Enable (RXEN) When set, the receive port is enabled. When cleared, the receive port is disabled.	R/W	1b

13.7.2.3 Port x MAC Receive Undersize Count Register (MAC_RX_UNDSIZE_CNT_x)

Register #: Port0: 0410h Size: 32 bits
Port1: 0810h
Port2: 0C10h

This register provides a counter of undersized packets received by the port. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Undersize Count of packets that have less than 64 byte and a valid FCS. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 115 hours.	RC	00000000h

13.7.2.4 Port x MAC Receive 64 Byte Count Register (MAC_RX_64_CNT_x)

Register #: Port0: 0411h Size: 32 bits
Port1: 0811h
Port2: 0C11h

This register provides a counter of 64 byte packets received by the port. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 64 Bytes Count of packets (including bad packets) that have exactly 64 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes is rounded down to the nearest byte.

13.7.2.5 Port x MAC Receive 65 to 127 Byte Count Register (MAC_RX_65_TO_127_CNT_x)

Register #: Port0: 0412h Size: 32 bits
 Port1: 0812h
 Port2: 0C12h

This register provides a counter of received packets between the size of 65 to 127 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 65 to 127 Bytes Count of packets (including bad packets) that have between 65 and 127 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 487 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes is rounded down to the nearest byte.

13.7.2.6 Port x MAC Receive 128 to 255 Byte Count Register (MAC_RX_128_TO_255_CNT_x)

Register #: Port0: 0413h Size: 32 bits
 Port1: 0813h
 Port2: 0C13h

This register provides a counter of received packets between the size of 128 to 255 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 128 to 255 Bytes Count of packets (including bad packets) that have between 128 and 255 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 848 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes is rounded down to the nearest byte.

13.7.2.7 Port x MAC Receive 256 to 511 Byte Count Register (MAC_RX_256_TO_511_CNT_x)

Register #: Port0: 0414h Size: 32 bits
 Port1: 0814h
 Port2: 0C14h

This register provides a counter of received packets between the size of 256 to 511 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 256 to 511 Bytes Count of packets (including bad packets) that have between 256 and 511 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 1581 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes is rounded down to the nearest byte.

13.7.2.8 Port x MAC Receive 512 to 1023 Byte Count Register (MAC_RX_512_TO_1023_CNT_x)

Register #: Port0: 0415h Size: 32 bits
 Port1: 0815h
 Port2: 0C15h

This register provides a counter of received packets between the size of 512 to 1023 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 512 to 1023 Bytes Count of packets (including bad packets) that have between 512 and 1023 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 3047 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes is rounded down to the nearest byte.

13.7.2.9 Port x MAC Receive 1024 to Max Byte Count Register (MAC_RX_1024_TO_MAX_CNT_x)

Register #: Port0: 0416h Size: 32 bits
 Port1: 0816h
 Port2: 0C16h

This register provides a counter of received packets between the size of 1024 to the maximum allowable number bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX 1024 to Max Bytes Count of packets (including bad packets) that have between 1024 and the maximum allowable number of bytes. The max number of bytes is 1518 for untagged packets and 1522 for tagged packets. If the Jumbo2K bit is set in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) , the max number of bytes is 2048. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 5979 hours.	RC	00000000h

Note: A bad packet is defined as a packet that has an FCS or Symbol error. For this counter, a packet with the maximum number of bytes that is not an integral number of bytes (e.g., a 1518 1/2 byte packet) is counted.

13.7.2.10 Port x MAC Receive Oversize Count Register (MAC_RX_OVRSIZE_CNT_x)

Register #: Port0: 0417h Size: 32 bits
 Port1: 0817h
 Port2: 0C17h

This register provides a counter of received packets with a size greater than the maximum byte size. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Oversize Count of packets that have more than the maximum allowable number of bytes and a valid FCS. The max number of bytes is 1518 for untagged packets and 1522 for tagged packets. If the Jumbo2K bit is set in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) , the max number of bytes is 2048. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 8813 hours.	RC	00000000h

Note: For this counter, a packet with the maximum number of bytes that is not an integral number of bytes (e.g., a 1518 1/2 byte packet) is not considered oversize.

13.7.2.11 Port x MAC Receive OK Count Register (MAC_RX_PKTOK_CNT_x)

Register #: Port0: 0418h Size: 32 bits
Port1: 0818h
Port2: 0C18h

This register provides a counter of received packets that are of proper length and are free of errors. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX OK Count of packets that are of proper length and are free of errors. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is one that has a FCS or Symbol error.

13.7.2.12 Port x MAC Receive CRC Error Count Register (MAC_RX_CRCERR_CNT_x)

Register #: Port0: 0419h Size: 32 bits
Port1: 0819h
Port2: 0C19h

This register provides a counter of received packets that with CRC errors. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX CRC Count of packets that have between 64 and the maximum allowable number of bytes and have a bad FCS, but do not have an extra nibble. The max number of bytes is 1518 for untagged packets and 1522 for tagged packets. If the Jumbo2K bit is set in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) , the max number of bytes is 2048. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 137 hours.	RC	00000000h

13.7.2.13 Port x MAC Receive Multicast Count Register (MAC_RX_MULCST_CNT_x)

Register #: Port0: 041Ah Size: 32 bits
 Port1: 081Ah
 Port2: 0C1Ah

This register provides a counter of valid received packets with a multicast destination address. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Multicast Count of good packets (proper length and free of errors), including MAC control frames, that have a multicast destination address (not including broadcasts). Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is one that has a FCS or Symbol error.

13.7.2.14 Port x MAC Receive Broadcast Count Register (MAC_RX_BRDCST_CNT_x)

Register #: Port0: 041Bh Size: 32 bits
 Port1: 081Bh
 Port2: 0C1Bh

This register provides a counter of valid received packets with a broadcast destination address. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Broadcast Count of valid packets (proper length and free of errors) that have a broadcast destination address. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is one that has a FCS or Symbol error.

LAN9352

13.7.2.15 Port x MAC Receive Pause Frame Count Register (MAC_RX_PAUSE_CNT_x)

Register #: Port0: 041Ch Size: 32 bits
Port1: 081Ch
Port2: 0C1Ch

This register provides a counter of valid received pause frame packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Pause Frame Count of valid packets (proper length and free of errors) that have a type field of 8808h and an op-code of 0001(Pause). Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is one that has a FCS or Symbol error.

13.7.2.16 Port x MAC Receive Fragment Error Count Register (MAC_RX_FRAG_CNT_x)

Register #: Port0: 041Dh Size: 32 bits
Port1: 081Dh
Port2: 0C1Dh

This register provides a counter of received packets of less than 64 bytes and a FCS error. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Fragment Count of packets that have less than 64 bytes and a FCS error. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 115 hours.	RC	00000000h

13.7.2.17 Port x MAC Receive Jabber Error Count Register (MAC_RX_JABB_CNT_x)

Register #: Port0: 041Eh Size: 32 bits
 Port1: 081Eh
 Port2: 0C1Eh

This register provides a counter of received packets with greater than the maximum allowable number of bytes and a FCS error. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Jabber Count of packets that have more than the maximum allowable number of bytes and a FCS error. The max number of bytes is 1518 for untagged packets and 1522 for tagged packets. If the Jumbo2K bit is set in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) , the max number of bytes is 2048. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 8813 hours.	RC	00000000h

Note: For this counter, a packet with the maximum number of bytes that is not an integral number of bytes (e.g. a 1518 1/2 byte packet) and contains a FCS error is not considered jabber and is not counted here.

13.7.2.18 Port x MAC Receive Alignment Error Count Register (MAC_RX_ALIGN_CNT_x)

Register #: Port0: 041Fh Size: 32 bits
 Port1: 081Fh
 Port2: 0C1Fh

This register provides a counter of received packets with 64 bytes to the maximum allowable and a FCS error. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Alignment Count of packets that have between 64 bytes and the maximum allowable number of bytes and are not byte aligned and have a bad FCS. The max number of bytes is 1518 for untagged packets and 1522 for tagged packets. If the Jumbo2K bit is set in the Port x MAC Receive Configuration Register (MAC_RX_CFG_x) , the max number of bytes is 2048. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: For this counter, a packet with the maximum number of bytes that is not an integral number of bytes (e.g. a 1518 1/2 byte packet) and a FCS error is considered an alignment error and is counted.

LAN9352

13.7.2.19 Port x MAC Receive Packet Length Count Register (MAC_RX_PKTLEN_CNT_x)

Register #: Port0: 0420h Size: 32 bits
Port1: 0820h
Port2: 0C20h

This register provides a counter of total bytes received. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Bytes Count of total bytes received (including bad packets). Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 5.8 hours.	RC	00000000h

Note: If necessary, for oversized packets, the packet is either truncated at 1518 bytes (untagged, Jumbo2K=0), 1522 bytes (tagged, Jumbo2K=0) or 2048 bytes (Jumbo2K=1). If this occurs, the byte count recorded is 1518, 1522 or 2048, respectively. The [Jumbo2K](#) bit is located in the [Port x MAC Receive Configuration Register \(MAC_RX_CFG_x\)](#).

Note: A bad packet is one that has an FCS or Symbol error. For this counter, a packet that is not an integral number of bytes (e.g. a 1518 1/2 byte packet) is rounded down to the nearest byte.

13.7.2.20 Port x MAC Receive Good Packet Length Count Register (MAC_RX_GOODPKTLEN_CNT_x)

Register #: Port0: 0421h Size: 32 bits
Port1: 0821h
Port2: 0C21h

This register provides a counter of total bytes received in good packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Good Bytes Count of total bytes received in good packets (proper length and free of errors). Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 5.8 hours.	RC	00000000h

Note: A bad packet is one that has an FCS or Symbol error.

13.7.2.21 Port x MAC Receive Symbol Error Count Register (MAC_RX_SYMBOL_CNT_x)

Register #: Port0: 0422h Size: 32 bits
 Port1: 0822h
 Port2: 0C22h

This register provides a counter of received packets with a symbol error. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Symbol Count of packets that had a receive symbol error. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 115 hours.	RC	00000000h

13.7.2.22 Port x MAC Receive Control Frame Count Register (MAC_RX_CTLFRM_CNT_x)

Register #: Port0: 0423h Size: 32 bits
 Port1: 0823h
 Port2: 0C23h

This register provides a counter of good packets with a type field of 8808h. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	RX Control Frame Count of good packets (proper length and free of errors) that have a type field of 8808h. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

Note: A bad packet is one that has an FCS or Symbol error.

This register indicates the number of times that the RX LPI indication from the PHY changed from de-asserted to asserted.

This register shows the total duration that the PHY has indicated RX LPI.

13.7.2.25 Port x MAC Transmit Configuration Register (MAC_TX_CFG_x)

Register #: Port0: [0440h](#) Size: 32 bits
 Port1: [0840h](#)
 Port2: [0C40h](#)

This read/write register configures the transmit packet parameters of the port.

Bits	Description	Type	Default
31:9	RESERVED	RO	-
8	Energy Efficient Ethernet (EEE_ENABLE) When set, this bit enables EEE operation (both TX LPI and RX LPI) Note: For switch port 0, this bit must always be written as 0.	R/W	Note 16
7	MAC Counter Test When set, TX and RX counters that normally clear to 0 when read, will be set to 7FFF_FFFCh when read with the exception of the Port x MAC Receive Packet Length Count Register (MAC_RX_PKTLEN_CNT_x) , Port x MAC Transmit Packet Length Count Register (MAC_TX_PKTLEN_CNT_x) and Port x MAC Receive Good Packet Length Count Register (MAC_RX_GOOD_PKTLEN_CNT_x) counters which will be set to 7FFF_FF80h.	R/W	0b
6:2	IFG Config These bits control the transmit inter-frame gap. IFG bit times = (IFG Config * 4) + 12 Note: IFG Config values less than 15 are unsupported.	R/W	10101b
1	TX Pad Enable When set, transmit packets shorter than 64 bytes are padded with zeros and will become 64 bytes in length. Note: Padding is used when a VLAN tagged frame of less than 68 bytes is received and has its tag removed (becoming less than 64 bytes in length).	R/W	1b
0	TX Enable (TXEN) When set, the transmit port is enabled. When cleared, the transmit port is disabled.	R/W	1b

Note 16: The default value of this field is determined by the [EEE_enable_strap_1](#) or [EEE_enable_strap_2](#) configuration strap. The default for switch port 0 is 0.

LAN9352

13.7.2.26 Port x MAC Transmit Flow Control Settings Register (MAC_TX_FC_SETTINGS_x)

Register #: Port0: 0441h Size: 32 bits
 Port1: 0841h
 Port2: 0C41h

This read/write register configures the flow control settings of the port.

Bits	Description	Type	Default
31:18	RESERVED	RO	-
17:16	Backoff Reset RX/TX Half-duplex-only. Determines when the truncated binary exponential backoff attempts counter is reset. 00 = Reset on successful transmission (IEEE standard) 01 = Reset on successful reception 1X = Reset on either successful transmission or reception	R/W	00b
15:0	Pause Time Value The value that is inserted into the transmitted pause packet when the switch wants to "XOFF" its link partner.	R/W	FFFFh

13.7.2.27 Port x EEE Time Wait TX System Register (EEE_TW_TX_SYS_x)

Register #: Port1: 0842h
 Port2: 0C42h Size: 32 bits

This register configures the time to wait before starting packet transmission after TX LPI removal.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:0	TX Delay After TX LPI Removal This field configures the time to wait, in microseconds, before starting packet transmission after TX LPI removal. Software should only change this field when the Energy Efficient Ethernet (EEE_ENABLE) bit is cleared. Note: In order to meet the IEEE 802.3 specified requirement, the minimum value of this field should be 00001Eh.	R/W	00001Eh

13.7.2.28 Port x EEE TX LPI Request Delay Register (EEE_TX_LPI_REQ_DELAY_x)

Register #: Port1: 0843h
 Port2: 0C43h Size: 32 bits

This register contains the amount of time, in microseconds, the MAC must wait after the TX FIFO is empty before invoking the LPI protocol.

- Note:** The actual time can be up to 1 us longer than specified.
- Note:** A value of zero is valid and will cause no delay to occur.
- Note:** If the TX FIFO becomes non-empty, the timer is restarted

Bits	Description	Type	Default
31:0	EEE TX LPI Request Delay This field contains the time to wait, in microseconds, before invoking the LPI protocol. Software should only change this field when the Energy Efficient Ethernet (EEE_ENABLE) bit is cleared.	R/W	00000000h

13.7.2.29 Port x MAC Transmit Deferred Count Register (MAC_TX_DEFER_CNT_x)

Register #: Port0: 0451h Size: 32 bits
 Port1: 0851h
 Port2: 0C51h

This register provides a counter deferred packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Deferred Count of packets that were available for transmission but were deferred on the first transmit attempt due to network traffic (either on receive or prior transmission). This counter is not incremented on collisions. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.30 Port x MAC Transmit Pause Count Register (MAC_TX_PAUSE_CNT_x)

Register #: Port0: 0452h Size: 32 bits
 Port1: 0852h
 Port2: 0C52h

This register provides a counter of transmitted pause packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Pause Count of pause packets transmitted. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.31 Port x MAC Transmit OK Count Register (MAC_TX_PKTOK_CNT_x)

Register #: Port0: 0453h Size: 32 bits
 Port1: 0853h
 Port2: 0C53h

This register provides a counter of successful transmissions. The counter is cleared upon being read.

BITS	DESCRIPTION	TYPE	DEFAULT
31:0	TX OK Count of successful transmissions. Undersize packets are not included in this count. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.32 Port x MAC Transmit 64 Byte Count Register (MAC_TX_64_CNT_x)

Register #: Port0: 0454h Size: 32 bits
 Port1: 0854h
 Port2: 0C54h

This register provides a counter of 64 byte packets transmitted by the port. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 64 Bytes Count of packets that have exactly 64 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.33 Port x MAC Transmit 65 to 127 Byte Count Register (MAC_TX_65_TO_127_CNT_x)

Register #: Port0: 0455h Size: 32 bits
 Port1: 0855h
 Port2: 0C55h

This register provides a counter of transmitted packets between the size of 65 to 127 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 65 to 127 Bytes Count of packets that have between 65 and 127 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 487 hours.	RC	00000000h

13.7.2.34 Port x MAC Transmit 128 to 255 Byte Count Register (MAC_TX_128_TO_255_CNT_x)

Register #: Port0: 0456h Size: 32 bits
 Port1: 0856h
 Port2: 0C56h

This register provides a counter of transmitted packets between the size of 128 to 255 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 128 to 255 Bytes Count of packets that have between 128 and 255 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 848 hours.	RC	00000000h

13.7.2.35 Port x MAC Transmit 256 to 511 Byte Count Register (MAC_TX_256_TO_511_CNT_x)

Register #: Port0: 0457h Size: 32 bits
 Port1: 0857h
 Port2: 0C57h

This register provides a counter of transmitted packets between the size of 256 to 511 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 256 to 511 Bytes Count of packets that have between 256 and 511 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 1581 hours.	RC	00000000h

13.7.2.36 Port x MAC Transmit 512 to 1023 Byte Count Register (MAC_TX_512_TO_1023_CNT_x)

Register #: Port0: 0458h Size: 32 bits
 Port1: 0858h
 Port2: 0C58h

This register provides a counter of transmitted packets between the size of 512 to 1023 bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 512 to 1023 Bytes Count of packets that have between 512 and 1023 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 3047 hours.	RC	00000000h

13.7.2.37 Port x MAC Transmit 1024 to Max Byte Count Register (MAC_TX_1024_TO_MAX_CNT_x)

Register #: Port0: 0459h Size: 32 bits
 Port1: 0859h
 Port2: 0C59h

This register provides a counter of transmitted packets between the size of 1024 to the maximum allowable number bytes. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX 1024 to Max Bytes Count of packets that have more than 1024 bytes. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 5979 hours.	RC	00000000h

13.7.2.38 Port x MAC Transmit Undersize Count Register (MAC_TX_UNDSIZE_CNT_x)

Register #: Port0: 045Ah Size: 32 bits
 Port1: 085Ah
 Port2: 0C5Ah

This register provides a counter of undersized packets transmitted by the port. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Undersize Count of packets that have less than 64 bytes. Note: This condition could occur when TX padding is disabled and a tag is removed. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 458 hours.	RC	00000000h

13.7.2.39 Port x MAC Transmit Packet Length Count Register (MAC_TX_PKTLEN_CNT_x)

Register #: Port0: 045Ch Size: 32 bits
 Port1: 085Ch
 Port2: 0C5Ch

This register provides a counter of total bytes transmitted. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Bytes Count of total bytes transmitted (does not include bytes from collisions, but does include bytes from Pause packets). Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 5.8 hours.	RC	00000000h

13.7.2.40 Port x MAC Transmit Broadcast Count Register (MAC_TX_BRDCST_CNT_x)

Register #: Port0: 045Dh Size: 32 bits
 Port1: 085Dh
 Port2: 0C5Dh

This register provides a counter of transmitted broadcast packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Broadcast Count of broadcast packets transmitted. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.41 Port x MAC Transmit Multicast Count Register (MAC_TX_MULCST_CNT_x)

Register #: Port0: 045Eh Size: 32 bits
 Port1: 085Eh
 Port2: 0C5Eh

This register provides a counter of transmitted multicast packets. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Multicast Count of multicast packets transmitted including MAC Control Pause frames. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.42 Port x MAC Transmit Late Collision Count Register (MAC_TX_LATECOL_CNT_x)

Register #: Port0: 045Fh Size: 32 bits
 Port1: 085Fh
 Port2: 0C5Fh

This register provides a counter of transmitted packets which experienced a late collision. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Late Collision Count of transmitted packets that experienced a late collision. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.2.43 Port x MAC Transmit Excessive Collision Count Register (MAC_TX_EXCCOL_CNT_x)

Register #: Port0: 0460h Size: 32 bits
 Port1: 0860h
 Port2: 0C60h

This register provides a counter of transmitted packets which experienced 16 collisions. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Excessive Collision Count of transmitted packets that experienced 16 collisions. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 1466 hours.	RC	00000000h

13.7.2.44 Port x MAC Transmit Single Collision Count Register (MAC_TX_SNGLECOL_CNT_x)

Register #: Port0: 0461h Size: 32 bits
Port1: 0861h
Port2: 0C61h

This register provides a counter of transmitted packets which experienced exactly 1 collision. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Single Collision Count of transmitted packets that experienced exactly 1 collision. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 573 hours.	RC	00000000h

13.7.2.45 Port x MAC Transmit Multiple Collision Count Register (MAC_TX_MULTICOL_CNT_x)

Register #: Port0: 0462h Size: 32 bits
Port1: 0862h
Port2: 0C62h

This register provides a counter of transmitted packets which experienced between 2 and 15 collisions. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Multiple Collision Count of transmitted packets that experienced between 2 and 15 collisions. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 664 hours.	RC	00000000h

13.7.2.46 Port x MAC Transmit Total Collision Count Register (MAC_TX_TOTALCOL_CNT_x)

Register #: Port0: 0463h Size: 32 bits
Port1: 0863h
Port2: 0C63h

This register provides a counter of total collisions including late collisions. The counter is cleared upon being read.

Bits	Description	Type	Default
31:0	TX Total Collision Total count of collisions including late collisions. This counter is incremented only in half-duplex operation. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 92 hours.	RC	00000000h

13.7.2.47 Port x TX LPI Transitions Register (TX_LPI_TRANSITION_x)

Register #: Size: 32 bits
 Port1: [0864h](#)
 Port2: [0C64h](#)

This register indicates the total number of times TX LPI request to the PHY changed from de-asserted to asserted.

Bits	Description	Type	Default
31:0	EEE TX LPI Transitions Count of total number of times the TX LPI request to the PHY changed from de-asserted to asserted. The counter is reset if the Energy Efficient Ethernet (EEE_ENABLE) bit in the Port x MAC Transmit Configuration Register (MAC_TX_CFG_x) is low.	RO	00000000h

13.7.2.48 Port x TX LPI Time Register (TX_LPI_TIME_x)

Register #: Size: 32 bits
 Port1: [0865h](#)
 Port2: [0C65h](#)

This register shows the total duration that TX LPI request to the PHY has been asserted.

Bits	Description	Type	Default
31:0	EEE TX LPI Time This field shows the total duration, in microseconds, that TX LPI request to the PHY has been asserted. The counter is reset if the Energy Efficient Ethernet (EEE_ENABLE) bit in the Port x MAC Transmit Configuration Register (MAC_TX_CFG_x) is low.	RO	00000000h

13.7.2.49 Port x MAC Interrupt Mask Register (MAC_IMR_x)

Register #: Port0: 0480h Size: 32 bits
Port1: 0880h
Port2: 0C80h

This register contains the Port x interrupt mask. Port x related interrupts in the [Port x MAC Interrupt Pending Register \(MAC_IPR_x\)](#) may be masked via this register. An interrupt is masked by setting the corresponding bit of this register. Clearing a bit will unmask the interrupt. Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Note: There are no possible Port x interrupt conditions available. This register exists for future use and should be configured as indicated for future compatibility.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7:0	RESERVED Note: These bits must be written as 11h.	R/W	11h

13.7.2.50 Port x MAC Interrupt Pending Register (MAC_IPR_x)

Register #: Port0: 0481h Size: 32 bits
Port1: 0881h
Port2: 0C81h

This read-only register contains the pending Port x interrupts. A set bit indicates an interrupt has been triggered. All interrupts in this register may be masked via the [Port x MAC Interrupt Pending Register \(MAC_IPR_x\)](#) register. Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Note: There are no possible Port x interrupt conditions available. This register exists for future use.

Bits	Description	Type	Default
31:0	RESERVED	RO	-

13.7.3 SWITCH ENGINE CSRS

This section details the Switch Engine related CSRs. These registers allow configuration and monitoring of the various Switch Engine components including the ALR, VLAN, Port VID and DIFFSERV tables. A list of the general switch CSRs and their corresponding register numbers is included in [Table 13-9](#).

13.7.3.1 Switch Engine ALR Command Register (SWE_ALR_CMD)

Register #: **1800h** Size: 32 bits

This register is used to manually read and write for MAC addresses from/into the ALR table. Setting any bit in this register will set the [Operation Pending](#) bit in the [Switch Engine ALR Command Status Register \(SWE_ALR_CMD_STS\)](#) and perform the specified command. Only one bit should be set at a time.

For a read accesses (Get commands), the [Operation Pending](#) bit indicates when the command is finished. The [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) and the [Switch Engine ALR Read Data 1 Register \(SWE_ALR_RD_DAT_1\)](#) can then be read.

For write accesses (Make command), the [Switch Engine ALR Write Data 0 Register \(SWE_ALR_WR_DAT_0\)](#) and the [Switch Engine ALR Write Data 1 Register \(SWE_ALR_WR_DAT_1\)](#) should first be written with the MAC address and data. The [Operation Pending](#) bit indicates when the command is finished.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2	<p>Make Entry When set, the contents of SWE_ALR_WR_DAT_0 and SWE_ALR_WR_DAT_1 are written into the ALR table. The ALR logic determines the location where the entry is written. This command can also be used to change or delete a previously written or automatically learned entry.</p> <p>This bit self-clears once the operation is complete as indicated by a low in the Operation Pending bit in the Switch Engine ALR Command Status Register (SWE_ALR_CMD_STS).</p> <p>This bit has no affect when written low.</p>	R/W SC	0b
1	<p>Get First Entry When set, the ALR read pointer is reset to the beginning of the ALR table and the ALR table is searched for the first valid entry, which is loaded into the SWE_ALR_RD_DAT_0 and SWE_ALR_RD_DAT_1 registers.</p> <p>This bit self-clears once the operation is complete as indicated by a low in the Operation Pending bit in the Switch Engine ALR Command Status Register (SWE_ALR_CMD_STS).</p> <p>This bit has no affect when written low.</p>	R/W SC	0b
0	<p>Get Next Entry When set, the next valid entry in the ALR MAC address table is loaded into the SWE_ALR_RD_DAT_0 and SWE_ALR_RD_DAT_1 registers.</p> <p>This bit self-clears once the operation is complete as indicated by a low in the Operation Pending bit in the Switch Engine ALR Command Status Register (SWE_ALR_CMD_STS).</p> <p>This bit has no affect when written low.</p>	R/W SC	0b

13.7.3.2 Switch Engine ALR Write Data 0 Register (SWE_ALR_WR_DAT_0)

Register #: 1801h Size: 32 bits

This register is used in conjunction with the [Switch Engine ALR Write Data 1 Register \(SWE_ALR_WR_DAT_1\)](#) and contains the first 32 bits of ALR data to be written.

Bits	Description	Type	Default
31:0	MAC Address This field contains the first 32 bits of the ALR entry that will be written in the ALR table. These bits correspond to the first 32 bits of the MAC address. Bit 0 holds the LSB of the first byte (the multicast bit).	R/W	00000000h

13.7.3.3 Switch Engine ALR Write Data 1 Register (SWE_ALR_WR_DAT_1)

Register #: 1802h

Size: 32 bits

This register is used in conjunction with the [Switch Engine ALR Write Data 0 Register \(SWE_ALR_WR_DAT_0\)](#) and contains the last 32 bits of ALR data to be written.

Bits	Description	Type	Default
31:27	RESERVED	RO	-
26	Valid When set, this bit makes the entry valid. It can be cleared to invalidate a previous entry that contained the specified MAC address.	R/W	0b
25	Age 1/Override This bit is used by the aging and forwarding processes. If the Static bit of this register is cleared, this bit is the msb of the aging timer. Software should set this bit so that the entry will age in the normal amount of time. If the Static bit is set, this bit is used as a port state override bit. When set, packets received with a destination address that matches the MAC address in the SWE_ALR_WR_DAT_1 and SWE_ALR_WR_DAT_0 registers will be forwarded regardless of the port state (except the Disabled state) of the ingress or egress port(s). This is typically used to allow the reception of BPDU packets in the non-forwarding state.	R/W	0b
24	Static When this bit is set, this entry will not be removed by the aging process and/or be changed by the learning process. When this bit is cleared, this entry will be automatically removed after 5 to 10 minutes of inactivity. Inactivity is defined as no packets being received with a source address that matches this MAC address. Note: This bit is normally set by software when adding manual entries.	R/W	0b
23	Age 0/Filter This bit is used by the aging and forwarding processes. If the Static bit of this register is cleared, this bit is the lsb of the aging timer. Software should set this bit so that the entry will age in the normal amount of time. If the Static bit is set, this bit is used to filter packets. When set, packets with a destination address that matches this MAC address will be filtered.	R/W	0b
22	Priority Enable When set, this bit enables usage of the Priority field for this MAC address entry. When clear, the Priority field is not used.	R/W	0b
21:19	Priority These bits specify the priority that is used for packets with a destination address that matches this MAC address. This priority is only used if both the Priority Enable bit of this register and the DA Highest Priority bit of the Switch Engine Global Ingress Configuration Register (SWE_GLOBAL_INGRSS_CFG) are set.	R/W	000b

LAN9352

Bits	Description	Type	Default																		
18:16	<p>Port These bits indicate the port(s) associated with this MAC address. When bit 18 is cleared, a single port is selected. When bit 18 is set, multiple ports are selected.</p> <table><tr><th>VALUE</th><th>ASSOCIATED PORT(S)</th></tr><tr><td>000</td><td>Port 0</td></tr><tr><td>001</td><td>Port 1</td></tr><tr><td>010</td><td>Port 2</td></tr><tr><td>011</td><td>RESERVED</td></tr><tr><td>100</td><td>Port 0 and Port 1</td></tr><tr><td>101</td><td>Port 0 and Port 2</td></tr><tr><td>110</td><td>Port 1 and Port 2</td></tr><tr><td>111</td><td>Port 0, Port 1 and Port 2</td></tr></table>	VALUE	ASSOCIATED PORT(S)	000	Port 0	001	Port 1	010	Port 2	011	RESERVED	100	Port 0 and Port 1	101	Port 0 and Port 2	110	Port 1 and Port 2	111	Port 0, Port 1 and Port 2	R/W	000b
VALUE	ASSOCIATED PORT(S)																				
000	Port 0																				
001	Port 1																				
010	Port 2																				
011	RESERVED																				
100	Port 0 and Port 1																				
101	Port 0 and Port 2																				
110	Port 1 and Port 2																				
111	Port 0, Port 1 and Port 2																				
15:0	<p>MAC Address These field contains the last 16 bits of the ALR entry that will be written into the ALR table. They correspond to the last 16 bits of the MAC address. Bit 15 holds the msb of the last byte (the last bit on the wire). The first 32 bits of the MAC address are located in the Switch Engine ALR Write Data 0 Register (SWE_ALR_WR_DAT_0).</p>	R/W	0000h																		

13.7.3.4 Switch Engine ALR Read Data 0 Register (SWE_ALR_RD_DAT_0)

Register #: 1805h Size: 32 bits

This register is used in conjunction with the [Switch Engine ALR Read Data 1 Register \(SWE_ALR_RD_DAT_1\)](#) to read the ALR table. It contains the first 32 bits of the ALR entry and is loaded via the [Get First Entry](#) or [Get Next Entry](#) commands in the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#). This register is only valid when either of the [Valid](#) or [End of Table](#) bits in the [Switch Engine ALR Read Data 1 Register \(SWE_ALR_RD_DAT_1\)](#) are set.

Bits	Description	Type	Default
31:0	MAC Address This field contains the first 32 bits of the ALR entry. These bits correspond to the first 32 bits of the MAC address. Bit 0 holds the LSB of the first byte (the multicast bit).	RO	00000000h

13.7.3.5 Switch Engine ALR Read Data 1 Register (SWE_ALR_RD_DAT_1)

Register #: 1806h Size: 32 bits

This register is used in conjunction with the [Switch Engine ALR Read Data 0 Register \(SWE_ALR_RD_DAT_0\)](#) to read the ALR table. It contains the last 32 bits of the ALR entry and is loaded via the [Get First Entry](#) or [Get Next Entry](#) commands in the [Switch Engine ALR Command Register \(SWE_ALR_CMD\)](#). This register is only valid when either of the [Valid](#) or [End of Table](#) bits are set.

Bits	Description	Type	Default
31:28	RESERVED	RO	-
27	End of Table This bit indicates that the end of the ALR table has been reached and further Get Next Entry commands are not required. Note: The Valid bit may or may not be set when the end of the table is reached.	RO	0b
26	Valid This bit clears when the Get First Entry or Get Next Entry bits of the Switch Engine ALR Command Register (SWE_ALR_CMD) are written. This bit sets when a valid entry is found in the ALR table. This bit stays cleared if the top of the ALR table is reached without finding an entry.	RO	0b
25	Age 1/Override This bit is used by the aging and forwarding processes. If the Static bit of this register is cleared, this bit is the msb of the aging timer. Software should set this bit so that the entry will age in the normal amount of time. If the Static bit is set, this bit is used as a port state override bit. When set, packets received with a destination address that matches the MAC address in the SWE_ALR_WR_DAT_1 and SWE_ALR_WR_DAT_0 registers will be forwarded regardless of the port state (except the Disabled state) of the ingress or egress port(s). This is typically used to allow the reception of BPDU packets in the non-forwarding state.	RO	0b
24	Static Indicates that this entry will not be removed by the aging process. When this bit is cleared, this entry will be automatically removed after 5 to 10 minutes of inactivity. Inactivity is defined as no packets being received with a source address that matches this MAC address.	RO	0b
23	Age 0/Filter This bit is used by the aging and forwarding processes. If the Static bit of this register is cleared, this bit is the lsb of the aging timer. Software should set this bit so that the entry will age in the normal amount of time. If the Static bit is set, this bit is used to filter packets. When set, packets with a destination address that matches this MAC address will be filtered.	RO	0b
22	Priority Enable Indicates whether or not the usage of the Priority field is enabled for this MAC address entry.	RO	0b

Bits	Description	Type	Default																		
21:19	Priority These bits specify the priority that is used for packets with a destination address that matches this MAC address. This priority is only used if both the Priority Enable bit of this register and the DA Highest Priority bit in the Switch Engine Global Ingress Configuration Register (SWE_GLOBAL_ING-GRSS_CFG) are set.	RO	000b																		
18:16	Port These bits indicate the port(s) associated with this MAC address. When bit 18 is cleared, a single port is selected. When bit 18 is set, multiple ports are selected. <table><tr><th>VALUE</th><th>ASSOCIATED PORT(S)</th></tr><tr><td>000</td><td>Port 0</td></tr><tr><td>001</td><td>Port 1</td></tr><tr><td>010</td><td>Port 2</td></tr><tr><td>011</td><td>RESERVED</td></tr><tr><td>100</td><td>Port 0 and Port 1</td></tr><tr><td>101</td><td>Port 0 and Port 2</td></tr><tr><td>110</td><td>Port 1 and Port 2</td></tr><tr><td>111</td><td>Port 0, Port 1 and Port 2</td></tr></table>	VALUE	ASSOCIATED PORT(S)	000	Port 0	001	Port 1	010	Port 2	011	RESERVED	100	Port 0 and Port 1	101	Port 0 and Port 2	110	Port 1 and Port 2	111	Port 0, Port 1 and Port 2	RO	000b
VALUE	ASSOCIATED PORT(S)																				
000	Port 0																				
001	Port 1																				
010	Port 2																				
011	RESERVED																				
100	Port 0 and Port 1																				
101	Port 0 and Port 2																				
110	Port 1 and Port 2																				
111	Port 0, Port 1 and Port 2																				
15:0	MAC Address These field contains the last 16 bits of the ALR entry. They correspond to the last 16 bits of the MAC address. Bit 15 holds the MSB of the last byte (the last bit on the wire). The first 32 bits of the MAC address are located in the Switch Engine ALR Read Data 0 Register (SWE_ALR_RD_DAT_0) .	RO	0000h																		

13.7.3.6 Switch Engine ALR Command Status Register (SWE_ALR_CMD_STS)

Register #: 1808h Size: 32 bits

This register indicates the current ALR command status.

Bits	Description	Type	Default
31:2	RESERVED	RO	-
1	ALR Init Done When set, indicates that the ALR table has finished being initialized by the reset process. The initialization is performed upon any reset that resets the Switch Fabric. The initialization takes approximately 20 μ s. During this time, any received packet will be dropped. Software should monitor this bit before writing any of the ALR tables or registers.	RO SS	Note 17
0	Operation Pending When set, indicates that the ALR command is taking place. This bit self-clears once the ALR command has finished.	RO SC	0b

Note 17: The default value of this bit is 0 immediately following any Switch Fabric reset and then self-sets to 1 once the ALR table is initialized.

13.7.3.7 Switch Engine ALR Configuration Register (SWE_ALR_CFG)

Register #: 1809h Size: 32 bits

This register controls the ALR aging timer duration and the allowance of Broadcast entries.

Bits	Description	Type	Default
31:28	RESERVED	RO	-
27:16	Aging Time This field sets the minimum time to age MAC Addresses from the ALR table. The time is specified in 1 second increments plus 1 second. A value of 0 is 1 second, a value of 1 is 2 seconds, etc. The maximum value of FFFh is approximately 69 minutes. The default sets a minimum time of 300 seconds.	R/W	129h
15:3	RESERVED	RO	-
2	Allow Broadcast Entries When set, this bit allows the use of the Broadcast MAC address in the ALR table.	R/W	0b
1	ALR Age Enable When set, this bit enables the aging process.	R/W	1b
0	ALR Age Test When set, this bit changes the aging timer from seconds to milliseconds.	R/W	0b

13.7.3.8 Switch Engine ALR Override Register (SWE_ALR_OVERRIDE)

Register #: **180Ah** Size: 32 bits

This register controls the ALR destination override function.

Bits	Description	Type	Default										
31:11	RESERVED	RO	-										
10:9	ALR Override Destination Port 2 When the ALR Override Enable Port 2 bit is set, packets received on Port 2 are sent to the port(s) specified by this field. <table><tr><th>Value</th><th>Port(s)</th></tr><tr><td>00</td><td>Port 0</td></tr><tr><td>01</td><td>Port 1</td></tr><tr><td>10</td><td>RESERVED</td></tr><tr><td>11</td><td>RESERVED</td></tr></table>	Value	Port(s)	00	Port 0	01	Port 1	10	RESERVED	11	RESERVED	R/W	00b
Value	Port(s)												
00	Port 0												
01	Port 1												
10	RESERVED												
11	RESERVED												
8	ALR Override Enable Port 2 When set, the ALR Destination MAC Address lookup result for packets received on port 2 are ignored and replaced with the value in ALR Override Destination Port 2 . Note: The ALR associated data Age 1/Override , Static , Age 0/Filter , Priority Enable and Priority are still used. Note: Forwarding rules described in Section 13.3.2 are still followed.	R/W	0b										
7	RESERVED	RO	-										
6:5	ALR Override Destination Port 1 When the ALR Override Enable Port 1 bit is set, packets received on Port 1 are sent to the port(s) specified by this field. <table><tr><th>Value</th><th>Port(s)</th></tr><tr><td>00</td><td>Port 0</td></tr><tr><td>01</td><td>RESERVED</td></tr><tr><td>10</td><td>Port 2</td></tr><tr><td>11</td><td>RESERVED</td></tr></table>	Value	Port(s)	00	Port 0	01	RESERVED	10	Port 2	11	RESERVED	R/W	00b
Value	Port(s)												
00	Port 0												
01	RESERVED												
10	Port 2												
11	RESERVED												
4	ALR Override Enable Port 1 When set, the ALR Destination MAC Address lookup result for packets received on port 1 are ignored and replaced with the value in ALR Override Destination Port 1 . Note: The ALR associated data Age 1/Override , Static , Age 0/Filter , Priority Enable and Priority are still used. Note: Forwarding rules described in Section 13.3.2 are still followed.	R/W	0b										
3	RESERVED	RO	-										

Bits	Description	Type	Default										
2:1	ALR Override Destination Port 0 When the ALR Override Enable Port 0 bit is set, packets received on Port 0 are sent to the port(s) specified by this field. <table><tr><th>Value</th><th>Port(s)</th></tr><tr><td>00</td><td>RESERVED</td></tr><tr><td>01</td><td>Port 1</td></tr><tr><td>10</td><td>Port 2</td></tr><tr><td>11</td><td>RESERVED</td></tr></table>	Value	Port(s)	00	RESERVED	01	Port 1	10	Port 2	11	RESERVED	R/W	00b
Value	Port(s)												
00	RESERVED												
01	Port 1												
10	Port 2												
11	RESERVED												
0	ALR Override Enable Port 0 When set, the ALR Destination MAC Address lookup result for packets received on port 0 are ignored and replaced with the value in ALR Override Destination Port 0 . Note: The ALR associated data Age 1/Override , Static , Age 0/Filter , Priority Enable and Priority are still used. Note: Forwarding rules described in Section 13.3.2 are still followed.	R/W	0b										

13.7.3.9 Switch Engine VLAN Command Register (SWE_VLAN_CMD)

Register #: 180Bh Size: 32 bits

This register is used to read and write the VLAN or Port VID tables. A write to this address performs the specified access. For a read access, the [Operation Pending](#) bit in the [Switch Engine VLAN Command Status Register \(SWE_VLAN_CMD_STS\)](#) indicates when the command is finished. The [Switch Engine VLAN Read Data Register \(SWE_VLAN_RD_DATA\)](#) can then be read.

For a write access, the [Switch Engine VLAN Write Data Register \(SWE_VLAN_WR_DATA\)](#) should be written first. The [Operation Pending](#) bit in the [Switch Engine VLAN Command Status Register \(SWE_VLAN_CMD_STS\)](#) indicates when the command is finished.

Bits	Description	Type	Default
31:6	RESERVED	RO	-
5	VLAN RnW This bit specifies a read(1) or a write(0) command.	R/W	0b
4	PVIDnVLAN When set, this bit selects the Port VID table. When cleared, this bit selects the VLAN table.	R/W	0b
3:0	VLAN/Port This field specifies the VLAN(0-15) or port(0-2) to be read or written. Note: Values outside of the valid range may cause unexpected results.	R/W	0h

13.7.3.10 Switch Engine VLAN Write Data Register (SWE_VLAN_WR_DATA)

Register #: 180Ch Size: 32 bits

This register is used write the VLAN or Port VID tables.

Bits	Description	Type	Default
31:18	RESERVED	RO	-

Bits	Description	Type	Default																								
17:0	<p>Port Default VID and Priority When the port VID table is selected (PVIDnVLAN=1 of the Switch Engine VLAN Command Register (SWE_VLAN_CMD)), bits 11:0 of this field specify the default VID for the port, bits 14:12 specify the default priority for packets with a non-broadcast destination MAC address and bits 17:15 specify the default priority for packets with a broadcast destination MAC address. All other bits of this field are reserved. These bits are used when a packet is received without a VLAN tag or with a NULL VLAN ID. The default VID is also used when the 802.1Q VLAN Disable bit is set. The default priority is also used when no other priority choice is selected. By default, the VID for all three ports is 1 and the priorities for all three ports is 0.</p> <p>Note: Values of 0 and FFFh should not be used since they are special VLAN IDs per the IEEE 802.3Q specification.</p> <p>VLAN Data When the VLAN table is selected (PVIDnVLAN=0 of the Switch Engine VLAN Command Register (SWE_VLAN_CMD)), the bits form the VLAN table entry as follows:</p> <table><tr><th>Bits</th><th>Description</th><th>Default</th></tr><tr><td>17</td><td><p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p><p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p><p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p></td><td>0b</td></tr><tr><td>16</td><td><p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p></td><td>0b</td></tr><tr><td>15</td><td><p>Member Port 1 See description for Member Port 2.</p></td><td>0b</td></tr><tr><td>14</td><td><p>Un-Tag Port 1 See description for Un-Tag Port 2.</p></td><td>0b</td></tr><tr><td>13</td><td><p>Member Port 0 See description for Member Port 2.</p></td><td>0b</td></tr><tr><td>12</td><td><p>Un-Tag Port 0 See description for Un-Tag Port 2.</p></td><td>0b</td></tr><tr><td>11:0</td><td><p>VID These bits specify the VLAN ID associated with this VLAN entry.</p><p>To disable a VLAN entry, a value of 0 should be used.</p><p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p><p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p></td><td>000h</td></tr></table>	Bits	Description	Default	17	<p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p> <p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p> <p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p>	0b	16	<p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p>	0b	15	<p>Member Port 1 See description for Member Port 2.</p>	0b	14	<p>Un-Tag Port 1 See description for Un-Tag Port 2.</p>	0b	13	<p>Member Port 0 See description for Member Port 2.</p>	0b	12	<p>Un-Tag Port 0 See description for Un-Tag Port 2.</p>	0b	11:0	<p>VID These bits specify the VLAN ID associated with this VLAN entry.</p> <p>To disable a VLAN entry, a value of 0 should be used.</p> <p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p> <p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p>	000h	R/W	00 0000 0000 0000 0000b
Bits	Description	Default																									
17	<p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p> <p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p> <p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p>	0b																									
16	<p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p>	0b																									
15	<p>Member Port 1 See description for Member Port 2.</p>	0b																									
14	<p>Un-Tag Port 1 See description for Un-Tag Port 2.</p>	0b																									
13	<p>Member Port 0 See description for Member Port 2.</p>	0b																									
12	<p>Un-Tag Port 0 See description for Un-Tag Port 2.</p>	0b																									
11:0	<p>VID These bits specify the VLAN ID associated with this VLAN entry.</p> <p>To disable a VLAN entry, a value of 0 should be used.</p> <p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p> <p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p>	000h																									

13.7.3.11 Switch Engine VLAN Read Data Register (SWE_VLAN_RD_DATA)

Register #: 180Eh Size: 32 bits

This register is used to read the VLAN or Port VID tables.

Bits	Description	Type	Default
31:18	RESERVED	RO	-

Bits	Description	Type	Default																								
17:0	<p>Port Default VID and Priority When the port VID table is selected (PVIDnVLAN=1 of the Switch Engine VLAN Command Register (SWE_VLAN_CMD)), bits 11:0 of this field specify the default VID for the port, bits 14:12 specify the default priority for packets with a non-broadcast destination MAC address and bits 17:15 specify the default priority for packets with a broadcast destination MAC address. All other bits of this field are reserved. These bits are used when a packet is received without a VLAN tag or with a NULL VLAN ID. The default VID is also used when the 802.1Q VLAN Disable bit is set. The default priority is also used when no other priority choice is selected. By default, the VID for all three ports is 1 and the priorities for all three ports is 0.</p> <p>Note: Values of 0 and FFFh should not be used since they are special VLAN IDs per the IEEE 802.3Q specification.</p> <p>VLAN Data When the VLAN table is selected (PVIDnVLAN=0 of the Switch Engine VLAN Command Register (SWE_VLAN_CMD)), the bits form the VLAN table entry as follows:</p> <table><tr><th>Bits</th><th>Description</th><th>Default</th></tr><tr><td>17</td><td><p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p><p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p><p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p></td><td>0b</td></tr><tr><td>16</td><td><p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p></td><td>0b</td></tr><tr><td>15</td><td><p>Member Port 1 See description for Member Port 2.</p></td><td>0b</td></tr><tr><td>14</td><td><p>Un-Tag Port 1 See description for Un-Tag Port 2.</p></td><td>0b</td></tr><tr><td>13</td><td><p>Member Port 0 See description for Member Port 2.</p></td><td>0b</td></tr><tr><td>12</td><td><p>Un-Tag Port 0 See description for Un-Tag Port 2.</p></td><td>0b</td></tr><tr><td>11:0</td><td><p>VID These bits specify the VLAN ID associated with this VLAN entry.</p><p>To disable a VLAN entry, a value of 0 should be used.</p><p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p><p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p></td><td>000h</td></tr></table>	Bits	Description	Default	17	<p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p> <p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p> <p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p>	0b	16	<p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p>	0b	15	<p>Member Port 1 See description for Member Port 2.</p>	0b	14	<p>Un-Tag Port 1 See description for Un-Tag Port 2.</p>	0b	13	<p>Member Port 0 See description for Member Port 2.</p>	0b	12	<p>Un-Tag Port 0 See description for Un-Tag Port 2.</p>	0b	11:0	<p>VID These bits specify the VLAN ID associated with this VLAN entry.</p> <p>To disable a VLAN entry, a value of 0 should be used.</p> <p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p> <p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p>	000h	RO	00 0000 0000 0000 0000b
Bits	Description	Default																									
17	<p>Member Port 2 Indicates the configuration of Port 2 for this VLAN entry.</p> <p>1 = Member - Packets with a VID that matches this entry are allowed on ingress. The port is a member of the broadcast domain on egress.</p> <p>0 = Not a Member - Packets with a VID that matches this entry are filtered on ingress unless the Admit Non Member bit in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER) is set for this port. The port is not a member of the broadcast domain on egress.</p>	0b																									
16	<p>Un-Tag Port 2 When this bit is set, packets with a VID that matches this entry will have their tag removed when re-transmitted on Port 2 when it is designated as a Hybrid port via the Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE).</p>	0b																									
15	<p>Member Port 1 See description for Member Port 2.</p>	0b																									
14	<p>Un-Tag Port 1 See description for Un-Tag Port 2.</p>	0b																									
13	<p>Member Port 0 See description for Member Port 2.</p>	0b																									
12	<p>Un-Tag Port 0 See description for Un-Tag Port 2.</p>	0b																									
11:0	<p>VID These bits specify the VLAN ID associated with this VLAN entry.</p> <p>To disable a VLAN entry, a value of 0 should be used.</p> <p>Note: A value of 0 is considered a NULL VLAN and should not normally be used other than to disable a VLAN entry.</p> <p>Note: A value of 3FFh is considered reserved by IEEE 802.1Q and should not be used.</p>	000h																									

13.7.3.12 Switch Engine VLAN Command Status Register (SWE_VLAN_CMD_STS)

Register #: 1810h Size: 32 bits

This register indicates the current VLAN command status.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Operation Pending When set, this bit indicates that the read or write command is taking place. This bit self-clears once the command has finished.	RO SC	0b

13.7.3.13 Switch Engine DIFFSERV Table Command Register (SWE_DIFFSERV_TBL_CFG)

Register #: 1811h Size: 32 bits

This register is used to read and write the DIFFSERV table. A write to this address performs the specified access. This table is used to map the received IP ToS/CS to a priority.

For a read access, the [Operation Pending](#) bit in the [Switch Engine DIFFSERV Table Command Status Register \(SWE_DIFFSERV_TBL_CMD_STS\)](#) indicates when the command is finished. The [Switch Engine DIFFSERV Table Read Data Register \(SWE_DIFFSERV_TBL_RD_DATA\)](#) can then be read.

For a write access, the [Switch Engine DIFFSERV Table Write Data Register \(SWE_DIFFSERV_TBL_WR_DATA\)](#) should be written first. The [Operation Pending](#) bit in the [Switch Engine DIFFSERV Table Command Status Register \(SWE_DIFFSERV_TBL_CMD_STS\)](#) indicates when the command is finished.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7	DIFFSERV Table RnW This bit specifies a read(1) or a write(0) command.	R/W	0b
6	RESERVED	RO	-
5:0	DIFFSERV Table Index This field specifies the ToS/CS entry that is accessed.	R/W	000000b

13.7.3.14 Switch Engine DIFFSERV Table Write Data Register (SWE_DIFFSERV_TBL_WR_DATA)

Register #: 1812h Size: 32 bits

This register is used to write the DIFFSERV table. The DIFFSERV table is not initialized upon reset on power-up. If DIFFSERV is enabled, the full table should be initialized by the host.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2:0	DIFFSERV Priority These bits specify the assigned receive priority for IP packets with a ToS/CS field that matches this index.	R/W	000b

13.7.3.15 Switch Engine DIFFSERV Table Read Data Register (SWE_DIFFSERV_TBL_RD_DATA)

Register #: 1813h Size: 32 bits

This register is used to read the DIFFSERV table.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2:0	DIFFSERV Priority These bits specify the assigned receive priority for IP packets with a ToS/CS field that matches this index.	RO	000b

13.7.3.16 Switch Engine DIFFSERV Table Command Status Register (SWE_DIFFSERV_TBL_CMD_STS)

Register #: 1814h Size: 32 bits

This register indicates the current DIFFSERV command status.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Operation Pending When set, this bit indicates that the read or write command is taking place. This bit self-clears once the command has finished.	RO SC	0b

13.7.3.17 Switch Engine Global Ingress Configuration Register (SWE_GLOBAL INGRSS_CFG)

Register #: 1840h

Size: 32 bits

This register is used to configure the global ingress rules.

Bits	Description	Type	Default
31:18	RESERVED	RO	-
17	Enable Other MLD Next Headers When set, Next Header values of 43, 44, 50, 51 and 60 are also used when monitoring MLD packets.	R/W	0b
16	Enable Any MLD Hop-by-Hop Next Header When set, the Next Header value in the IPv6 Hop-by-Hop Options header is ignore when monitoring MLD packets.	R/W	0b
15	802.1Q VLAN Disable When set, the VID from the VLAN tag is ignored and the per port default VID (PVID) is used for purposes of VLAN rules. This does not affect the packet tag on egress.	R/W	0b
14	Use Tag When set, the priority from the VLAN tag is enabled as a transmit priority queue choice.	R/W	0b
13	Allow Monitor Echo When set, monitoring packets are allowed to be echoed back to the source port. When cleared, monitoring packets, like other packets, are never sent back to the source port. This bit is useful when the monitor port wishes to receive its own MLD/IGMP packets.	R/W	0b
12:10	MLD/IGMP Monitor Port This field is the port bit map where IPv6 MLD packets and IPv4 IGMP packets are sent.	R/W	0b
9	Use IP When set, the IPv4 TOS or IPv6 SC field is enabled as a transmit priority queue choice.	R/W	0b
8	Enable MLD Monitoring When set, IPv6 Multicast Listening Discovery packets are monitored and sent to the MLD/IGMP monitoring port.	R/W	0b
7	Enable IGMP Monitoring When set, IPv4 IGMP packets are monitored and sent to the MLD/IGMP monitor port.	R/W	0b
6	SWE Counter Test When this bit is set the Switch Engine counters that normally clear to 0 when read will be set to 7FFF_FFFCh when read.	R/W	0b
5	DA Highest Priority When this bit is set and the priority enable bit in the ALR table for the destination MAC address is set, the transmit priority queue that is selected is taken from the ALR Priority bits (see the Switch Engine ALR Read Data 1 Register (SWE_ALR_RD_DAT_1)).	R/W	0b

LAN9352

Bits	Description	Type	Default
4	Filter Multicast When this bit is set, packets with a multicast destination address are filtered if the address is not found in the ALR table. Broadcasts are not included in this filter.	R/W	0b
3	Drop Unknown When this bit is set, packets with a unicast destination address are filtered if the address is not found in the ALR table.	R/W	0b
2	Use Precedence When the priority is taken from an IPV4 packet (enabled via the Use IP bit), this bit selects between precedence bits in the TOS octet or the DIFFSERV table. When set, IPv4 packets will use the precedence bits in the TOS octet to select the transmit priority queue. When cleared, IPv4 packets will use the DIFFSERV table to select the transmit priority queue.	R/W	1b
1	VL Higher Priority When this bit is set and VLAN priority is enabled (via the Use Tag bit), the priority from the VLAN tag has higher priority than the IP TOS/SC field.	R/W	1b
0	VLAN Enable When set, VLAN ingress rules are enabled.	R/W	0b

13.7.3.18 Switch Engine Port Ingress Configuration Register (SWE_PORT INGRSS_CFG)

Register #: 1841h Size: 32 bits

This register is used to configure the per port ingress rules.

Bits	Description	Type	Default
31:6	RESERVED	RO	-
5:3	Enable Learning on Ingress When set, source addresses are learned when a packet is received on the corresponding port and the corresponding Port State in the Switch Engine Port State Register (SWE_PORT_STATE) is set to forwarding or learning. There is one enable bit per ingress port. Bits 5,4,3 correspond to switch ports 2,1,0 respectively.	R/W	111b
2:0	Enable Membership Checking When set, VLAN membership is checked when a packet is received on the corresponding port. The packet will be filtered if the ingress port is not a member of the VLAN (unless the Admit Non Member bit is set for the port in the Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER)). For destination addresses that are found in the ALR table, the packet will be filtered if the egress port is not a member of the VLAN (for destination addresses that are not found in the ALR table only the ingress port is checked for membership). The VLAN Enable bit in the Switch Engine Global Ingress Configuration Register (SWE_GLOBAL INGRSS_CFG) needs to be set for these bits to have an affect. There is one enable bit per ingress port. Bits 2,1,0 correspond to switch ports 2,1,0 respectively.	R/W	000b

LAN9352

13.7.3.19 Switch Engine Admit Only VLAN Register (SWE_ADMT_ONLY_VLAN)

Register #: [1842h](#) Size: 32 bits

This register is used to configure the per port ingress rule for allowing only VLAN tagged packets.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2:0	Admit Only VLAN When set, untagged and priority tagged packets are filtered. The VLAN Enable bit in the Switch Engine Global Ingress Configuration Register (SWE_GLOBAL INGRSS_CFG) needs to be set for these bits to have an affect. There is one enable bit per ingress port. Bits 2,1,0 correspond to switch ports 2,1,0 respectively.	R/W	000b

13.7.3.20 Switch Engine Port State Register (SWE_PORT_STATE)

Register #: 1843h Size: 32 bits

This register is used to configure the per port spanning tree state.

Bits	Description	Type	Default
31:6	RESERVED	RO	-
5:4	Port State Port 2 These bits specify the spanning tree port states for Port 2. 00 = Forwarding 01 = Listening/Blocking 10 = Learning 11 = Disabled	R/W	00b
3:2	Port State Port 1 These bits specify the spanning tree port states for Port 1. 00 = Forwarding 01 = Listening/Blocking 10 = Learning 11 = Disabled	R/W	00b
1:0	Port State Port 0 These bits specify the spanning tree port states for Port 0. 00 = Forwarding 01 = Listening/Blocking 10 = Learning 11 = Disabled Note: Typically, the Host MAC port is kept in the forwarding state, since it is not a true network port.	R/W	00b

13.7.3.21 Switch Engine Priority to Queue Register (SWE_PRI_TO_QUE)

Register #: 1845h Size: 32 bits

This register specifies the Traffic Class table that maps the packet priority into the egress queues.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:14	Priority 7 traffic Class These bits specify the egress queue that is used for packets with a priority of 7.	R/W	11b
13:12	Priority 6 traffic Class These bits specify the egress queue that is used for packets with a priority of 6.	R/W	11b
11:10	Priority 5 traffic Class These bits specify the egress queue that is used for packets with a priority of 5.	R/W	10b
9:8	Priority 4 traffic Class These bits specify the egress queue that is used for packets with a priority of 4.	R/W	10b
7:6	Priority 3 traffic Class These bits specify the egress queue that is used for packets with a priority of 3.	R/W	01b
5:4	Priority 2 traffic Class These bits specify the egress queue that is used for packets with a priority of 2.	R/W	00b
3:2	Priority 1 traffic Class These bits specify the egress queue that is used for packets with a priority of 1.	R/W	00b
1:0	Priority 0 traffic Class These bits specify the egress queue that is used for packets with a priority of 0.	R/W	01b

13.7.3.22 Switch Engine Port Mirroring Register (SWE_PORT_MIRROR)

Register #: 1846h Size: 32 bits

This register is used to configure port mirroring.

Bits	Description	Type	Default
31:9	RESERVED	RO	-
8	Enable RX Mirroring Filtered When set, packets that would normally have been filtered are included in the receive mirroring function and are sent only to the sniffer port. When cleared, filtered packets are not mirrored. Note: The Ingress Filtered Count Registers will still count these packets as filtered and the Switch Engine Interrupt Pending Register (SWE_IPR) will still register a drop interrupt.	R/W	0b
7:5	Sniffer Port These bits specify the sniffer port that transmits packets that are monitored. Bits 7,6,5 correspond to switch ports 2,1,0 respectively. Note: Only one port should be set as the sniffer.	R/W	00b
4:2	Mirrored Port These bits specify if a port is to be mirrored. Bits 4,3,2 correspond to switch ports 2,1,0 respectively. Note: Multiple ports can be set as mirrored.	R/W	00b
1	Enable RX Mirroring This bit enables packets received on the mirrored ports to be also sent to the sniffer port.	R/W	0b
0	Enable TX Mirroring This bit enables packets transmitted on the mirrored ports to be also sent to the sniffer port.	R/W	0b

13.7.3.23 Switch Engine Ingress Port Type Register (SWE_INGRSS_PORT_TYP)

Register #: 1847h Size: 32 bits

This register is used to enable the special tagging mode used to determine the destination port based on the VLAN tag contents.

Bits	Description	Type	Default
31:6	RESERVED	RO	-
5:4	Ingress Port Type Port 2 A setting of 11b enables the usage of the VLAN tag to specify the packet destination. All other values disable this feature.	R/W	00b
3:2	Ingress Port Type Port 1 A setting of 11b enables the usage of the VLAN tag to specify the packet destination. All other values disable this feature.	R/W	00b
1:0	Ingress Port Type Port 0 A setting of 11b enables the usage of the VLAN tag to specify the packet destination. All other values disable this feature.	R/W	00b

13.7.3.24 Switch Engine Broadcast Throttling Register (SWE_BCST_THROT)

Register #: 1848h Size: 32 bits

This register configures the broadcast input rate throttling.

Bits	Description	Type	Default
31:27	RESERVED	RO	-
26	Broadcast Throttle Enable Port 2 This bit enables broadcast input rate throttling on Port 2.	R/W	0b
25:18	Broadcast Throttle Level Port 2 These bits specify the number of bytes x 64 allowed to be received per every 1.72mS interval.	R/W	00000010b
17	Broadcast Throttle Enable Port 1 This bit enables broadcast input rate throttling on Port 1.	R/W	0b
16:9	Broadcast Throttle Level Port 1 These bits specify the number of bytes x 64 allowed to be received per every 1.72 ms interval.	R/W	00000010b
8	Broadcast Throttle Enable Port 0 This bit enables broadcast input rate throttling on Port 0.	R/W	0b
7:0	Broadcast Throttle Level Port 0 These bits specify the number of bytes x 64 allowed to be received per every 1.72 ms interval.	R/W	00000010b

13.7.3.25 Switch Engine Admit Non Member Register (SWE_ADMT_N_MEMBER)

Register #: **1849h** Size: 32 bits

This register is used to allow access to a VLAN even if the ingress port is not a member.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2:0	Admit Non Member When set, a received packet is accepted even if the ingress port is not a member of the destination VLAN. The VLAN still must be active in the switch. There is one bit per ingress port. Bits 2,1,0 correspond to switch ports 2,1,0 respectively.	R/W	000b

13.7.3.26 Switch Engine Ingress Rate Configuration Register (SWE_INGRSS_RATE_CFG)

Register #: **184Ah** Size: 32 bits

This register, along with the settings accessible via the [Switch Engine Ingress Rate Command Register \(SWE_INGRSS_RATE_CMD\)](#), is used to configure the ingress rate metering/coloring.

Bits	Description	Type	Default
31:3	RESERVED	RO	-
2:1	Rate Mode These bits configure the rate metering/coloring mode. 00 = Source Port & Priority 01 = Source Port Only 10 = Priority Only 11 = RESERVED	R/W	00b
0	Ingress Rate Enable When set, ingress rates are metered and packets are colored and dropped if necessary.	R/W	0b

13.7.3.27 Switch Engine Ingress Rate Command Register (SWE_INGRSS_RATE_CMD)

Register #: 184Bh Size: 32 bits

This register is used to indirectly read and write the ingress rate metering/color table registers. A write to this address performs the specified access.

For a read access, the [Operation Pending](#) bit in the [Switch Engine Ingress Rate Command Status Register \(SWE_INGRSS_RATE_CMD_STS\)](#) indicates when the command is finished. The [Switch Engine Ingress Rate Read Data Register \(SWE_INGRSS_RATE_RD_DATA\)](#) can then be read.

For a write access, the [Switch Engine Ingress Rate Write Data Register \(SWE_INGRSS_RATE_WR_DATA\)](#) should be written first. The [Operation Pending](#) bit in the [Switch Engine Ingress Rate Command Status Register \(SWE_INGRSS_RATE_CMD_STS\)](#) indicates when the command is finished.

For details on 16-bit wide Ingress Rate Table registers indirectly accessible by this register, see [INGRESS RATE TABLE REGISTERS](#) below.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7	Ingress Rate RnW These bits specify a read(1) or write(0) command.	R/W	0b
6:5	Type These bits select between the ingress rate metering/color table registers as follows: 00 = RESERVED 01 = Committed Information Rate Registers (uses CIS Address field) 10 = Committed Burst Register 11 = Excess Burst Register	R/W	00b
4:0	CIR Address These bits select one of the 24 Committed Information Rate registers. When Rate Mode is set to Source Port & Priority in the Switch Engine Ingress Rate Configuration Register (SWE_INGRSS_RATE_CFG) , the first set of 8 registers (CIR addresses 0-7) are for Port 0, the second set of 8 registers (CIR addresses 8-15) are for Port 1 and the third set of registers (CIR addresses 16-23) are for Port 2. Priority 0 is the lower register of each set (e.g., 0, 8 and 16). When Rate Mode is set to Source Port Only, the first register (CIR address 0) is for Port 0, the second register (CIR address 1) is for Port 1 and the third register (CIR address 2) is for Port 2. When Rate Mode is set to Priority Only, the first register (CIR address 0) is for priority 0, the second register (CIR address 1) is for priority 1 and so forth up to priority 23. Note: Values outside of the valid range may cause unexpected results.	R/W	00000b

INGRESS RATE TABLE REGISTERS

The ingress rate metering/color table consists of 24 Committed Information Rate (CIR) registers (one per port/priority), a Committed Burst Size register and an Excess Burst Size register. All metering/color table registers are 16-bits in size and are accessed indirectly via the [Switch Engine Ingress Rate Command Register \(SWE_INGRSS_RATE_CMD\)](#). Descriptions of these registers are detailed in [Table 13-10](#) below.

TABLE 13-10: METERING/COLOR TABLE REGISTER DESCRIPTIONS

Description	Type	Default
<p>Excess Burst Size This register specifies the maximum excess burst size in bytes. Bursts larger than this value that exceed the excess data rate are dropped.</p> <p>Note: Either this value or the Committed Burst Size should be set larger than or equal to the largest possible packet expected.</p> <p>Note: All of the Excess Burst token buckets are initialized to this default value. If a lower value is programmed into this register, the token buckets will need to be normally depleted below this value before this value has any affect on limiting the token bucket maximum values.</p> <p>This register is 16-bits wide.</p>	R/W	0600h
<p>Committed Burst Size This register specifies the maximum committed burst size in bytes. Bursts larger than this value that exceed the committed data rate are subjected to random dropping.</p> <p>Note: Either this value or the Excess Burst Size should be set larger than or equal to the largest possible packet expected.</p> <p>Note: All of the Committed Burst token buckets are initialized to this default value. If a lower value is programmed into this register, the token buckets will need to be normally depleted below this value before this value has any affect on limiting the token bucket maximum values.</p> <p>This register is 16-bits wide.</p>	R/W	0600h
<p>Committed Information Rate (CIR) These registers specify the committed data rate for the port/priority pair. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.</p> <p>There are 24 of these registers each 16-bits wide.</p>	R/W	0014h

13.7.3.28 Switch Engine Ingress Rate Command Status Register (SWE_INGRSS_RATE_CMD_STS)

Register #: 184Ch Size: 32 bits

This register indicates the current ingress rate command status.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Operation Pending When set, indicates that the read or write command is taking place. This bit self-clears once the command has finished.	RO SC	0b

13.7.3.29 Switch Engine Ingress Rate Write Data Register (SWE_INGRSS_RATE_WR_DATA)

Register #: 184Dh Size: 32 bits

This register is used to write the ingress rate table registers.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	Data This is the data to be written to the ingress rate table registers as specified in the Switch Engine Ingress Rate Command Register (SWE_INGRSS_RATE_CMD) . Refer to INGRESS RATE TABLE REGISTERS on page 431 for details on these registers.	R/W	0000h

13.7.3.30 Switch Engine Ingress Rate Read Data Register (SWE_INGRSS_RATE_RD_DATA)

Register #: 184Eh Size: 32 bits

This register is used to read the ingress rate table registers.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	Data This is the read data from the ingress rate table registers as specified in the Switch Engine Ingress Rate Command Register (SWE_INGRSS_RATE_CMD) . Refer to INGRESS RATE TABLE REGISTERS on page 431 for details on these registers.	RO	0000h

13.7.3.31 Switch Engine Port 0 Ingress Filtered Count Register (SWE_FILTERED_CNT_0)

Register #: 1850h Size: 32 bits

This register counts the number of packets filtered at ingress on Port 0. This count includes packets filtered due to broadcast throttling but does not include packets dropped due to ingress rate limiting (which are counted separately).

Bits	Description	Type	Default
31:0	Filtered This field is a count of packets filtered at ingress and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.32 Switch Engine Port 1 Ingress Filtered Count Register (SWE_FILTERED_CNT_1)

Register #: 1851h Size: 32 bits

This register counts the number of packets filtered at ingress on Port 1. This count includes packets filtered due to broadcast throttling but does not include packets dropped due to ingress rate limiting (which are counted separately).

Bits	Description	Type	Default
31:0	Filtered This field is a count of packets filtered at ingress and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.33 Switch Engine Port 2 Ingress Filtered Count Register (SWE_FILTERED_CNT_2)

Register #: 1852h Size: 32 bits

This register counts the number of packets filtered at ingress on Port 2. This count includes packets filtered due to broadcast throttling but does not include packets dropped due to ingress rate limiting (which are counted separately).

Bits	Description	Type	Default
31:0	Filtered This field is a count of packets filtered at ingress and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.34 Switch Engine Port 0 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_REGEN_TBL_0)

Register #: 1855h Size: 32 bits

This register provides the ability to map the received VLAN priority to a regenerated priority. The regenerated priority is used in determining the output priority queue. By default, the regenerated priority is identical to the received priority.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:21	Regen7 These bits specify the regenerated priority for received priority 7.	R/W	111b
20:18	Regen6 These bits specify the regenerated priority for received priority 6.	R/W	110b
17:15	Regen5 These bits specify the regenerated priority for received priority 5.	R/W	101b
14:12	Regen4 These bits specify the regenerated priority for received priority 4.	R/W	100b
11:9	Regen3 These bits specify the regenerated priority for received priority 3.	R/W	011b
8:6	Regen2 These bits specify the regenerated priority for received priority 2.	R/W	010b
5:3	Regen1 These bits specify the regenerated priority for received priority 1.	R/W	001b
2:0	Regen0 These bits specify the regenerated priority for received priority 0.	R/W	000b

13.7.3.35 Switch Engine Port 1 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_REGEN_TBL_1)

Register #: 1856h Size: 32 bits

This register provides the ability to map the received VLAN priority to a regenerated priority. The regenerated priority is used in determining the output priority queue. By default, the regenerated priority is identical to the received priority.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:21	Regen7 These bits specify the regenerated priority for received priority 7.	R/W	111b
20:18	Regen6 These bits specify the regenerated priority for received priority 6.	R/W	110b
17:15	Regen5 These bits specify the regenerated priority for received priority 5.	R/W	101b
14:12	Regen4 These bits specify the regenerated priority for received priority 4.	R/W	100b
11:9	Regen3 These bits specify the regenerated priority for received priority 3.	R/W	011b
8:6	Regen2 These bits specify the regenerated priority for received priority 2.	R/W	010b
5:3	Regen1 These bits specify the regenerated priority for received priority 1.	R/W	001b
2:0	Regen0 These bits specify the regenerated priority for received priority 0.	R/W	000b

13.7.3.36 Switch Engine Port 2 Ingress VLAN Priority Regeneration Table Register (SWE_INGRSS_REGEN_TBL_2)

Register #: 1857h Size: 32 bits

This register provides the ability to map the received VLAN priority to a regenerated priority. The regenerated priority is used in determining the output priority queue. By default, the regenerated priority is identical to the received priority.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:21	Regen7 These bits specify the regenerated priority for received priority 7.	R/W	111b
20:18	Regen6 These bits specify the regenerated priority for received priority 6.	R/W	110b
17:15	Regen5 These bits specify the regenerated priority for received priority 5.	R/W	101b
14:12	Regen4 These bits specify the regenerated priority for received priority 4.	R/W	100b
11:9	Regen3 These bits specify the regenerated priority for received priority 3.	R/W	011b
8:6	Regen2 These bits specify the regenerated priority for received priority 2.	R/W	010b
5:3	Regen1 These bits specify the regenerated priority for received priority 1.	R/W	001b
2:0	Regen0 These bits specify the regenerated priority for received priority 0.	R/W	000b

13.7.3.37 Switch Engine Port 0 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_0)

Register #: 1858h Size: 32 bits

This register counts the number of MAC addresses on Port 0 that were not learned or were overwritten by a different address due to address table space limitations.

Bits	Description	Type	Default
31:0	Learn Discard This field is a count of MAC addresses not learned or overwritten and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.38 Switch Engine Port 1 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_1)

Register #: **1859h** Size: 32 bits

This register counts the number of MAC addresses on Port 1 that were not learned or were overwritten by a different address due to address table space limitations.

Bits	Description	Type	Default
31:0	Learn Discard This field is a count of MAC addresses not learned or overwritten and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.39 Switch Engine Port 2 Learn Discard Count Register (SWE_LRN_DISCRD_CNT_2)

Register #: **185Ah** Size: 32 bits

This register counts the number of MAC addresses on Port 2 that were not learned or were overwritten by a different address due to address table space limitations.

Bits	Description	Type	Default
31:0	Learn Discard This field is a count of MAC addresses not learned or overwritten and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.3.40 Switch Engine Interrupt Mask Register (SWE_IMR)

Register #: **1880h** Size: 32 bits

This register contains the Switch Engine interrupt mask, which masks the interrupts in the [Switch Engine Interrupt Pending Register \(SWE_IPR\)](#). All Switch Engine interrupts are masked by setting the Interrupt Mask bit. Clearing this bit will unmask the interrupts. Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Interrupt Mask When set, this bit masks interrupts from the Switch Engine. The status bits in the Switch Engine Interrupt Pending Register (SWE_IPR) are not affected.	R/W	1b

13.7.3.41 Switch Engine Interrupt Pending Register (SWE_IPR)

Register #: 1881h Size: 32 bits

This register contains the Switch Engine interrupt status. The status is double buffered. All interrupts in this register may be masked via the [Switch Engine Interrupt Mask Register \(SWE_IMR\)](#). Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default																																		
31:15	RESERVED	RO	-																																		
14:11	<div><div><div>Drop Reason B</div><div>When the Set B Valid bit is set, these bits indicate the reason a packet was dropped per the table below:</div></div><table><thead><tr><th>Bit Values</th><th>Description</th></tr></thead><tbody><tr><td>0000</td><td>Admit Only VLAN was set and the packet was untagged or priority tagged.</td></tr><tr><td>0001</td><td>The destination address was not in the ALR table (unknown or broadcast), Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.</td></tr><tr><td>0010</td><td>The destination address was found in the ALR table but the source port was not in the forwarding state.</td></tr><tr><td>0011</td><td>The destination address was found in the ALR table but the destination port was not in the forwarding state.</td></tr><tr><td>0100</td><td>The destination address was found in the ALR table but Enable Membership Checking on ingress was set and the destination port was not a member of the incoming VLAN.</td></tr><tr><td>0101</td><td>The destination address was found in the ALR table but the Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.</td></tr><tr><td>0110</td><td>Drop Unknown was set and the destination address was a unicast but not in the ALR table.</td></tr><tr><td>0111</td><td>Filter Multicast was set and the destination address was a multicast and not in the ALR table.</td></tr><tr><td>1000</td><td>The packet was a broadcast but exceeded the Broadcast Throttling limit.</td></tr><tr><td>1001</td><td>The destination address was not in the ALR table (unknown or broadcast) and the source port was not in the forwarding state.</td></tr><tr><td>1010</td><td>The destination address was found in the ALR table but the source and destination ports were the same.</td></tr><tr><td>1011</td><td>The destination address was found in the ALR table and the Filter bit was set for that address.</td></tr><tr><td>1100</td><td>RESERVED</td></tr><tr><td>1101</td><td>RESERVED</td></tr><tr><td>1110</td><td>A packet was received with a VLAN ID of FFFh.</td></tr><tr><td>1111</td><td>RESERVED</td></tr></tbody></table></div>	Bit Values	Description	0000	Admit Only VLAN was set and the packet was untagged or priority tagged.	0001	The destination address was not in the ALR table (unknown or broadcast), Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.	0010	The destination address was found in the ALR table but the source port was not in the forwarding state.	0011	The destination address was found in the ALR table but the destination port was not in the forwarding state.	0100	The destination address was found in the ALR table but Enable Membership Checking on ingress was set and the destination port was not a member of the incoming VLAN.	0101	The destination address was found in the ALR table but the Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.	0110	Drop Unknown was set and the destination address was a unicast but not in the ALR table.	0111	Filter Multicast was set and the destination address was a multicast and not in the ALR table.	1000	The packet was a broadcast but exceeded the Broadcast Throttling limit.	1001	The destination address was not in the ALR table (unknown or broadcast) and the source port was not in the forwarding state.	1010	The destination address was found in the ALR table but the source and destination ports were the same.	1011	The destination address was found in the ALR table and the Filter bit was set for that address.	1100	RESERVED	1101	RESERVED	1110	A packet was received with a VLAN ID of FFFh.	1111	RESERVED	RC	0000b
Bit Values	Description																																				
0000	Admit Only VLAN was set and the packet was untagged or priority tagged.																																				
0001	The destination address was not in the ALR table (unknown or broadcast), Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.																																				
0010	The destination address was found in the ALR table but the source port was not in the forwarding state.																																				
0011	The destination address was found in the ALR table but the destination port was not in the forwarding state.																																				
0100	The destination address was found in the ALR table but Enable Membership Checking on ingress was set and the destination port was not a member of the incoming VLAN.																																				
0101	The destination address was found in the ALR table but the Enable Membership Checking on ingress was set, Admit Non Member was cleared and the source port was not a member of the incoming VLAN.																																				
0110	Drop Unknown was set and the destination address was a unicast but not in the ALR table.																																				
0111	Filter Multicast was set and the destination address was a multicast and not in the ALR table.																																				
1000	The packet was a broadcast but exceeded the Broadcast Throttling limit.																																				
1001	The destination address was not in the ALR table (unknown or broadcast) and the source port was not in the forwarding state.																																				
1010	The destination address was found in the ALR table but the source and destination ports were the same.																																				
1011	The destination address was found in the ALR table and the Filter bit was set for that address.																																				
1100	RESERVED																																				
1101	RESERVED																																				
1110	A packet was received with a VLAN ID of FFFh.																																				
1111	RESERVED																																				

Bits	Description	Type	Default
10:9	Source Port B When the Set B Valid bit is set, these bits indicate the source port on which the packet was dropped. 00 = Port 0 01 = Port 1 10 = Port 2 11 = RESERVED	RC	00b
8	Set B Valid When set, bits 14:9 are valid.	RC	0b
7:4	Drop Reason A When the Set A Valid bit is set, these bits indicate the reason a packet was dropped. See the Drop Reason B description above for definitions of each value of this field.	RC	0000b
3:2	Source port A When the Set A Valid bit is set, these bits indicate the source port on which the packet was dropped. 00 = Port 0 01 = Port 1 10 = Port 2 11 = RESERVED	RC	00b
1	Set A Valid When set, bits 7:2 are valid.	RC	0b
0	Interrupt Pending When set, a packet dropped event(s) is indicated.	RC	0b

13.7.4 BUFFER MANAGER CSRS

This section details the Buffer Manager (BM) registers. These registers allow configuration and monitoring of the switch buffer levels and usage. A list of the general switch CSRs and their corresponding register numbers is included in [Table 13-9](#).

13.7.4.1 Buffer Manager Configuration Register (BM_CFG)

Register #: **1C00h** Size: 32 bits

This register enables egress rate pacing and ingress rate discarding.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6	BM Counter Test When this bit is set, Buffer Manager (BM) counters that normally clear to 0 when read, will be set to 7FFF_FFFC when read.	R/W	0b
5	Fixed Priority Queue Servicing When set, output queues are serviced with a fixed priority ordering. When cleared, output queues are serviced with a weighted round robin ordering.	R/W	0b
4:2	Egress Rate Enable When set, egress rate pacing is enabled. Bits 4,3,2 correspond to switch ports 2,1,0 respectively.	R/W	0b
1	Drop on Yellow When this bit is set, packets that exceed the Ingress Committed Burst Size (colored Yellow) are subjected to random discard. Note: See Section 13.7.3.27, "Switch Engine Ingress Rate Command Register (SWE INGRSS_RATE_CMD)," on page 430 for information on configuring the Ingress Committed Burst Size.	R/W	0b
0	Drop on Red When this bit is set, packets that exceed the Ingress Excess Burst Size (colored Red) are discarded. Note: See Section 13.7.3.27, "Switch Engine Ingress Rate Command Register (SWE INGRSS_RATE_CMD)," on page 430 for information on configuring the Ingress Excess Burst Size.	R/W	0b

13.7.4.2 Buffer Manager Drop Level Register (BM_DROP_LVL)

Register #: 1C01h Size: 32 bits

This register configures the overall buffer usage limits.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:8	Drop Level Low These bits specify the buffer limit that can be used per ingress port during times when 2 or 3 ports are active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	49h
7:0	Drop Level High These bits specify the buffer limit that can be used per ingress port during times when 1 port is active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	64h

13.7.4.3 Buffer Manager Flow Control Pause Level Register (BM_FC_PAUSE_LVL)

Register #: 1C02h Size: 32 bits

This register configures the buffer usage level when a Pause frame or backpressure is sent.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:8	Pause Level Low These bits specify the buffer usage level during times when 2 or 3 ports are active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	21h
7:0	Pause Level High These bits specify the buffer usage level during times when 1 port is active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	3Ch

13.7.4.4 Buffer Manager Flow Control Resume Level Register (BM_FC_RESUME_LVL)

Register #: 1C03h Size: 32 bits

This register configures the buffer usage level when a Pause frame with a pause value of 1 is sent.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:8	Resume Level Low These bits specify the buffer usage level during times when 2 or 3 ports are active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	03h
7:0	Resume Level High These bits specify the buffer usage level during times when 0 or 1 ports are active. Each buffer is 128 bytes. Note: A port is “active” when 36 buffers are in use for that port.	R/W	07h

13.7.4.5 Buffer Manager Broadcast Buffer Level Register (BM_BCST_LVL)

Register #: 1C04h Size: 32 bits

This register configures the buffer usage limits for broadcasts, multicasts and unknown unicasts.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7:0	Broadcast Drop Level These bits specify the maximum number of buffers that can be used by broadcasts, multicasts and unknown unicasts. Each buffer is 128 bytes.	R/W	31h

13.7.4.6 Buffer Manager Port 0 Drop Count Register (BM_DRP_CNT_SRC_0)

Register #: 1C05h Size: 32 bits

This register counts the number of packets dropped by the Buffer Manager that were received on Port 0. This count includes packets dropped due to buffer space limits and ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 0 and is cleared when read. Note: The counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.7 Buffer Manager Port 1 Drop Count Register (BM_DRP_CNT_SRC_1)

Register #: 1C06h Size: 32 bits

This register counts the number of packets dropped by the Buffer Manager that were received on Port 1. This count includes packets dropped due to buffer space limits and ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 1 and is cleared when read. Note: The counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.8 Buffer Manager Port 2 Drop Count Register (BM_DRP_CNT_SRC_2)

Register #: 1C07h Size: 32 bits

This register counts the number of packets dropped by the Buffer Manager that were received on Port 2. This count includes packets dropped due to buffer space limits and ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 2 and is cleared when read. Note: The counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.9 Buffer Manager Reset Status Register (BM_RST_STS)

Register #: 1C08h Size: 32 bits

This register indicates when the Buffer Manager has been initialized by the reset process.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	BM Ready When set, indicates the Buffer Manager tables have finished being initialized by the reset process. The initialization is performed upon any reset that resets the Switch Fabric.	RO SS	Note 18

Note 18: The default value of this bit is 0 immediately following any Switch Fabric reset and then self-sets to 1 once the ALR table is initialized.

13.7.4.10 Buffer Manager Random Discard Table Command Register (BM_RNDM_DSCRD_TBL_CMD)

Register #: **1C09h** Size: 32 bits

This register is used to read and write the Random Discard Weight table. A write to this address performs the specified access. This table is used to set the packet drop probability verses the buffer usage.

For a read access, the [Buffer Manager Random Discard Table Read Data Register \(BM_RNDM_DSCRD_TBL_RDATA\)](#) can be read following a write to this register.

For a write access, the [Buffer Manager Random Discard Table Write Data Register \(BM_RNDM_DSCRD_TBL_W-DATA\)](#) should be written before writing this register.

Bits	Description	Type	Default																																		
31:5	RESERVED	RO	-																																		
4	Random Discard Weight Table RnW Specifies a read (1) or a write (0) command.	R/W	0b																																		
3:0	<p>Random Discard Weight Table Index Specifies the buffer usage range that is accessed.</p> <p>There are a total of 16 probability entries. Each entry corresponds to a range of the number of buffers used by the ingress port. The ranges are structured to give more resolution towards the lower buffer usage end.</p> <table><tr><th>BIT VALUES</th><th>BUFFER USAGE LEVEL</th></tr><tr><td>0000</td><td>0 to 7</td></tr><tr><td>0001</td><td>8 to 15</td></tr><tr><td>0010</td><td>16 to 23</td></tr><tr><td>0011</td><td>24 to 31</td></tr><tr><td>0100</td><td>32 to 39</td></tr><tr><td>0101</td><td>40 to 47</td></tr><tr><td>0110</td><td>48 to 55</td></tr><tr><td>0111</td><td>56 to 63</td></tr><tr><td>1000</td><td>64 to 79</td></tr><tr><td>1001</td><td>80 to 95</td></tr><tr><td>1010</td><td>96 to 111</td></tr><tr><td>1011</td><td>112 to 127</td></tr><tr><td>1100</td><td>128 to 159</td></tr><tr><td>1101</td><td>160 to 191</td></tr><tr><td>1110</td><td>192 to 223</td></tr><tr><td>1111</td><td>224 to 255</td></tr></table>	BIT VALUES	BUFFER USAGE LEVEL	0000	0 to 7	0001	8 to 15	0010	16 to 23	0011	24 to 31	0100	32 to 39	0101	40 to 47	0110	48 to 55	0111	56 to 63	1000	64 to 79	1001	80 to 95	1010	96 to 111	1011	112 to 127	1100	128 to 159	1101	160 to 191	1110	192 to 223	1111	224 to 255	R/W	0h
BIT VALUES	BUFFER USAGE LEVEL																																				
0000	0 to 7																																				
0001	8 to 15																																				
0010	16 to 23																																				
0011	24 to 31																																				
0100	32 to 39																																				
0101	40 to 47																																				
0110	48 to 55																																				
0111	56 to 63																																				
1000	64 to 79																																				
1001	80 to 95																																				
1010	96 to 111																																				
1011	112 to 127																																				
1100	128 to 159																																				
1101	160 to 191																																				
1110	192 to 223																																				
1111	224 to 255																																				

13.7.4.11 Buffer Manager Random Discard Table Write Data Register (BM_RNDM_DSCRD_TBL_WDATA)

Register #: 1C0Ah Size: 32 bits

This register is used to write the Random Discard Weight table.

Note: The Random Discard Weight table is not initialized upon reset or power-up. If a random discard is enabled, the full table should be initialized by the host.

Bits	Description	Type	Default
31:10	RESERVED	RO	-
9:0	Drop Probability These bits specify the discard probability of a packet that has been colored Yellow by the ingress metering. The probability is given in 1/1024's. For example, a setting of 1 is one in 1024 or approximately 0.1%. A setting of all ones (1023) is 1023 in 1024 or approximately 99.9%. There are a total of 16 probability entries. Each entry corresponds to a range of the number of buffers used by the ingress port, as specified in Section 13.7.4.10, "Buffer Manager Random Discard Table Command Register (BM_RNDM_DSCRD_TBL_CMD)" .	R/W	00 0000 0000b

13.7.4.12 Buffer Manager Random Discard Table Read Data Register (BM_RNDM_DSCRD_TBL_RDATA)

Register #: 1C0Bh Size: 32 bits

This register is used to read the Random Discard Weight table.

Bits	Description	Type	Default
31:10	RESERVED	RO	-
9:0	Drop Probability These bits specify the discard probability of a packet that has been colored Yellow by the ingress metering. The probability is given in 1/1024's. For example, a setting of 1 is one in 1024 or approximately 0.1%. A setting of all ones (1023) is 1023 in 1024 or approximately 99.9%. There are a total of 16 probability entries. Each entry corresponds to a range of the number of buffers used by the ingress port, as specified in Section 13.7.4.10, "Buffer Manager Random Discard Table Command Register (BM_RNDM_DSCRD_TBL_CMD)" .	RO	00 0000 0000b

13.7.4.13 Buffer Manager Egress Port Type Register (BM_EGRSS_PORT_TYPE)

Register #: 1C0Ch Size: 32 bits

This register is used to configure the egress VLAN tagging rules. See [Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340](#) for additional details.

Bits	Description	Type	Default
31:23	RESERVED	RO	-
22	VID/Priority Select Port 2 This bit determines the VID and priority in inserted or changed tags. 0: The default VID of the ingress port / priority calculated on ingress. 1: The default VID / priority of the egress port. This is only used when the Egress Port Type is set as Hybrid.	R/W	0b
21	Insert Tag Port 2 When set, untagged packets will have a tag added. The VID and priority is determined by the VID/Priority Select Port 2 bit. The un-tag bit in the VLAN table for the default VLAN ID also needs to be cleared in order for the tag to be inserted. This is only used when the Egress Port Type is set as Hybrid.	R/W	0b
20	Change VLAN ID Port 2 When set, regular tagged packets will have their VLAN ID overwritten with the Default VLAN ID of either the ingress or egress port, as determined by the VID/Priority Select Port 2 bit. The Change Tag bit also needs to be set. The un-tag bit in the VLAN table for the incoming VLAN ID also needs to be cleared, otherwise the tag will be removed instead. Priority tagged packets will have their VLAN ID overwritten with the Default VLAN ID of either the ingress or egress port independent of this bit. This is only used when the Egress Port Type is set as Hybrid.	R/W	0b
19	Change Priority Port 2 When set, regular tagged and priority tagged packets will have their Priority overwritten with the priority determined by the VID/Priority Select Port 2 bit. For regular tagged packets, the Change Tag bit also needs to be set by software. The un-tag bit in the VLAN table for the incoming VLAN ID also needs to be cleared, otherwise the tag would be removed instead. This is only used when the Egress Port Type is set as Hybrid.	R/W	0b

Bits	Description	Type	Default										
18	Change Tag Port 2 When set, allows the Change Tag and Change Priority bits to affect regular tagged packets. This bit has no affect on priority tagged packets. This is only used when the Egress Port Type is set as Hybrid.	R/W	0b										
17:16	Egress Port Type Port 2 These bits set the egress port type which determines the tagging/un-tagging rules. <table border="1"><thead><tr><th>Bit Values</th><th>EGRESS PORT TYPE</th></tr></thead><tbody><tr><td>00</td><td>Dumb Packets from regular ports pass untouched. Special tagged packets from the External MII port have their tagged stripped.</td></tr><tr><td>01</td><td>Access Tagged packets (including special tagged packets from the External MII port) have their tagged stripped.</td></tr><tr><td>10</td><td>Hybrid Supports a mix of tagging, un-tagging and changing tags. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.</td></tr><tr><td>11</td><td>CPU A special tag is added to indicate the source of the packet. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.</td></tr></tbody></table>	Bit Values	EGRESS PORT TYPE	00	Dumb Packets from regular ports pass untouched. Special tagged packets from the External MII port have their tagged stripped.	01	Access Tagged packets (including special tagged packets from the External MII port) have their tagged stripped.	10	Hybrid Supports a mix of tagging, un-tagging and changing tags. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.	11	CPU A special tag is added to indicate the source of the packet. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.	R/W	0b
Bit Values	EGRESS PORT TYPE												
00	Dumb Packets from regular ports pass untouched. Special tagged packets from the External MII port have their tagged stripped.												
01	Access Tagged packets (including special tagged packets from the External MII port) have their tagged stripped.												
10	Hybrid Supports a mix of tagging, un-tagging and changing tags. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.												
11	CPU A special tag is added to indicate the source of the packet. See Section 13.4.6, "Adding, Removing and Changing VLAN Tags," on page 340 for additional details.												
15	RESERVED	RO	-										
14	VID/Priority Select Port 1 Identical to VID/Priority Select Port 2 definition above.	R/W	0b										
13	Insert Tag Port 1 Identical to Insert Tag Port 2 definition above.	R/W	0b										
12	Change VLAN ID Port 1 Identical to Change VLAN ID Port 2 definition above.	R/W	0b										
11	Change Priority Port 1 Identical to Change Priority Port 2 definition above.	R/W	0b										
10	Change Tag Port 1 Identical to Change Tag Port 2 definition above.	R/W	0b										
9:8	Egress Port Type Port 1 Identical to Egress Port Type Port 2 definition above.	R/W	0b										
7	RESERVED	RO	-										
6	VID/Priority Select Port 0 Identical to VID/Priority Select Port 2 definition above.	R/W	0b										
5	Insert Tag Port 0 Identical to Insert Tag Port 2 definition above.	R/W	0b										
4	Change VLAN ID Port 0 Identical to Change VLAN ID Port 2 definition above.	R/W	0b										
3	Change Priority Port 0 Identical to Change Priority Port 2 definition above.	R/W	0b										

Bits	Description	Type	Default
2	Change Tag Port 0 Identical to Change Tag Port 2 definition above.	R/W	0b
1:0	Egress Port Type Port 0 Identical to Egress Port Type Port 2 definition above.	R/W	0b

13.7.4.14 Buffer Manager Port 0 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_00_01)

Register #: **1C0Dh** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 0 Priority Queue 1 These bits specify the egress data rate for the Port 0 priority queue 1. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 0 Priority Queue 0 These bits specify the egress data rate for the Port 0 priority queue 0. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.15 Buffer Manager Port 0 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_02_03)

Register #: **1C0Eh** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 0 Priority Queue 3 These bits specify the egress data rate for the Port 0 priority queue 3. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 0 Priority Queue 2 These bits specify the egress data rate for the Port 0 priority queue 2. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.16 Buffer Manager Port 1 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_10_11)

Register #: **1C0Fh** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 1 Priority Queue 1 These bits specify the egress data rate for the Port 1 priority queue 1. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 1 Priority Queue 0 These bits specify the egress data rate for the Port 1 priority queue 0. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.17 Buffer Manager Port 1 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_12_13)

Register #: **1C10h** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 1 Priority Queue 3 These bits specify the egress data rate for the Port 1 priority queue 3. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 1 Priority Queue 2 These bits specify the egress data rate for the Port 1 priority queue 2. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.18 Buffer Manager Port 2 Egress Rate Priority Queue 0/1 Register (BM_EGRSS_RATE_20_21)

Register #: **1C11h** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 2 Priority Queue 1 These bits specify the egress data rate for the Port 2 priority queue 1. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 2 Priority Queue 0 These bits specify the egress data rate for the Port 2 priority queue 0. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.19 Buffer Manager Port 2 Egress Rate Priority Queue 2/3 Register (BM_EGRSS_RATE_22_23)

Register #: **1C12h** Size: 32 bits

This register, along with the [Buffer Manager Configuration Register \(BM_CFG\)](#), is used to configure the egress rate pacing.

Bits	Description	Type	Default
31:26	RESERVED	RO	-
25:13	Egress Rate Port 2 Priority Queue 3 These bits specify the egress data rate for the Port 2 priority queue 3. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b
12:0	Egress Rate Port 2 Priority Queue 2 These bits specify the egress data rate for the Port 2 priority queue 2. The rate is specified in time per byte. The time is this value plus 1 times 20 ns.	R/W	00000 00000000b

13.7.4.20 Buffer Manager Port 0 Default VLAN ID and Priority Register (BM_VLAN_0)

Register #: **1C13h** Size: 32 bits

This register is used to specify the default VLAN ID and priority of Port 0.

Bits	Description	Type	Default
31:15	RESERVED	RO	-
14:12	Default Priority These bits specify the default priority that is used when a tag is inserted or changed on egress.	R/W	000b
11:0	Default VLAN ID These bits specify the default that is used when a tag is inserted or changed on egress.	R/W	0000 00000001b

13.7.4.21 Buffer Manager Port 1 Default VLAN ID and Priority Register (BM_VLAN_1)

Register #: **1C14h** Size: 32 bits

This register is used to specify the default VLAN ID and priority of Port 1.

Bits	Description	Type	Default
31:15	RESERVED	RO	-
14:12	Default Priority These bits specify the default priority that is used when a tag is inserted or changed on egress.	R/W	000b
11:0	Default VLAN ID These bits specify the default that is used when a tag is inserted or changed on egress.	R/W	0000 00000001b

13.7.4.22 Buffer Manager Port 2 Default VLAN ID and Priority Register (BM_VLAN_2)

Register #: **1C15h** Size: 32 bits

This register is used to specify the default VLAN ID and priority of Port 2.

Bits	Description	Type	Default
31:15	RESERVED	RO	-
14:12	Default Priority These bits specify the default priority that is used when a tag is inserted or changed on egress.	R/W	000b
11:0	Default VLAN ID These bits specify the default that is used when a tag is inserted or changed on egress.	R/W	0000 00000001b

13.7.4.23 Buffer Manager Port 0 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_0)

Register #: **1C16h** Size: 32 bits

This register counts the number of packets received on Port 0 that were dropped by the Buffer Manager due to ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 0 and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.24 Buffer Manager Port 1 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_1)

Register #: **1C17h** Size: 32 bits

This register counts the number of packets received on Port 1 that were dropped by the Buffer Manager due to ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 1 and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.25 Buffer Manager Port 2 Ingress Rate Drop Count Register (BM_RATE_DRP_CNT_SRC_2)

Register #: [1C18h](#) Size: 32 bits

This register counts the number of packets received on Port 2 that were dropped by the Buffer Manager due to ingress rate limit discarding (Red and random Yellow dropping).

Bits	Description	Type	Default
31:0	Dropped Count These bits count the number of dropped packets received on Port 2 and is cleared when read. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	RC	00000000h

13.7.4.26 Buffer Manager Interrupt Mask Register (BM_IMR)

Register #: [1C20h](#) Size: 32 bits

This register contains the Buffer Manager interrupt mask, which masks the interrupts in the [Buffer Manager Interrupt Pending Register \(BM_IPR\)](#). All Buffer Manager interrupts are masked by setting the Interrupt Mask bit. Clearing this bit will unmask the interrupts. Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	Interrupt Mask When set, this bit masks interrupts from the Buffer Manager. The status bits in the Buffer Manager Interrupt Pending Register (BM_IPR) are not affected.	R/W	1b

13.7.4.27 Buffer Manager Interrupt Pending Register (BM_IPR)

Register #: **1C21h** Size: 32 bits

This register contains the Buffer Manager interrupt status. The status is double buffered. All interrupts in this register may be masked via the [Buffer Manager Interrupt Mask Register \(BM_IMR\)](#). Refer to [Section 8.0, "System Interrupts," on page 73](#) for more information.

Bits	Description	Type	Default																																		
31:14	RESERVED	RO	-																																		
13:10	<div><div>Drop Reason B</div><div>When the Status B Pending bit is set, these bits indicate the reason a packet was dropped per the table below:</div><table><thead><tr><th>Bit Values</th><th>Description</th></tr></thead><tbody><tr><td>0000</td><td>The destination address was not in the ALR table (unknown or broadcast) and the Broadcast Buffer Level was exceeded.</td></tr><tr><td>0001</td><td>Drop on Red was set and the packet was colored Red.</td></tr><tr><td>0010</td><td>There were no buffers available.</td></tr><tr><td>0011</td><td>There were no memory descriptors available.</td></tr><tr><td>0100</td><td>The destination address was not in the ALR table (unknown or broadcast) and there were no valid destination ports.</td></tr><tr><td>0101</td><td>The packet had a receive error and was >64 bytes.</td></tr><tr><td>0110</td><td>The Buffer Drop Level was exceeded.</td></tr><tr><td>0111</td><td>RESERVED</td></tr><tr><td>1000</td><td>RESERVED</td></tr><tr><td>1001</td><td>Drop on Yellow was set, the packet was colored Yellow and was randomly selected to be dropped.</td></tr><tr><td>1010</td><td>RESERVED</td></tr><tr><td>1011</td><td>RESERVED</td></tr><tr><td>1100</td><td>RESERVED</td></tr><tr><td>1101</td><td>RESERVED</td></tr><tr><td>1110</td><td>RESERVED</td></tr><tr><td>1111</td><td>RESERVED</td></tr></tbody></table></div>	Bit Values	Description	0000	The destination address was not in the ALR table (unknown or broadcast) and the Broadcast Buffer Level was exceeded.	0001	Drop on Red was set and the packet was colored Red.	0010	There were no buffers available.	0011	There were no memory descriptors available.	0100	The destination address was not in the ALR table (unknown or broadcast) and there were no valid destination ports.	0101	The packet had a receive error and was >64 bytes.	0110	The Buffer Drop Level was exceeded.	0111	RESERVED	1000	RESERVED	1001	Drop on Yellow was set, the packet was colored Yellow and was randomly selected to be dropped.	1010	RESERVED	1011	RESERVED	1100	RESERVED	1101	RESERVED	1110	RESERVED	1111	RESERVED	RC	0000b
Bit Values	Description																																				
0000	The destination address was not in the ALR table (unknown or broadcast) and the Broadcast Buffer Level was exceeded.																																				
0001	Drop on Red was set and the packet was colored Red.																																				
0010	There were no buffers available.																																				
0011	There were no memory descriptors available.																																				
0100	The destination address was not in the ALR table (unknown or broadcast) and there were no valid destination ports.																																				
0101	The packet had a receive error and was >64 bytes.																																				
0110	The Buffer Drop Level was exceeded.																																				
0111	RESERVED																																				
1000	RESERVED																																				
1001	Drop on Yellow was set, the packet was colored Yellow and was randomly selected to be dropped.																																				
1010	RESERVED																																				
1011	RESERVED																																				
1100	RESERVED																																				
1101	RESERVED																																				
1110	RESERVED																																				
1111	RESERVED																																				
9:8	<div><div>Source Port B</div><div>When the Status B Pending bit is set, these bits indicate the source port on which the packet was dropped.</div><div><div>00 = Port 0</div><div>01 = Port 1</div><div>10 = Port 2</div><div>11 = RESERVED</div></div></div>	RC	00b																																		

LAN9352

Bits	Description	Type	Default
7	Status B Pending When set, bits 13:8 are valid.	RC	0b
6:3	Drop Reason A When the Set A Valid bit is set, these bits indicate the reason a packet was dropped. See the Drop Reason B description above for definitions of each value of this field.	RC	0000b
2:1	Source port A When the Set A Valid bit is set, these bits indicate the source port on which the packet was dropped. 00 = Port 0 01 = Port 1 10 = Port 2 11 = RESERVED	RC	00b
0	Set A Valid When set, bits 6:1 are valid.	RC	0b

14.0 I²C MASTER EEPROM CONTROLLER

14.1 Functional Overview

This chapter details the EEPROM I²C master and EEPROM Loader provided by the device. The I²C EEPROM controller is an I²C master module which interfaces an optional external EEPROM with the system register bus and the EEPROM Loader. Multiple sizes of external EEPROMs are supported. Configuration of the EEPROM size is accomplished via the [eeprom_size_strap](#) configuration strap. Various commands are supported for EEPROM access, allowing for the storage and retrieval of static data. The I²C interface conforms to the NXP *I²C-Bus Specification*.

The EEPROM Loader provides the automatic loading of configuration settings from the EEPROM into the device at reset. The EEPROM Loader module interfaces to the EEPROM Controller, Ethernet PHYs and the system CSRs.

14.2 I²C Overview

I²C is a bi-directional 2-wire data protocol. A device that sends data is defined as a transmitter and a device that receives data is defined as a receiver. The bus is controlled by a master which generates the EESCL clock, controls bus access and generates the start and stop conditions. Either a master or slave may operate as a transmitter or receiver as determined by the master.


Both the clock (EESCL) and data (EESDA) signals have digital input filters that reject pulses that are less than 100 ns. The data pin is driven low when either interface sends a low, emulating the wired-AND function of the I²C bus.

The following bus states exist:

- **Idle:** Both EESDA and EESCL are high when the bus is idle.
- **Start & Stop Conditions:** A start condition is defined as a high to low transition on the EESDA line while EESCL is high. A stop condition is defined as a low to high transition on the EESDA line while EESCL is high. The bus is considered to be busy following a start condition and is considered free 4.7 μ s/1.3 μ s (for 100 kHz and 400 kHz operation, respectively) following a stop condition. The bus stays busy following a repeated start condition (instead of a stop condition). Starts and repeated starts are otherwise functionally equivalent.
- **Data Valid:** Data is valid, following the start condition, when EESDA is stable while EESCL is high. Data can only be changed while the clock is low. There is one valid bit per clock pulse. Every byte must be 8 bits long and is transmitted MSB first.
- **Acknowledge:** Each byte of data is followed by an acknowledge bit. The master generates a ninth clock pulse for the acknowledge bit. The transmitter releases EESDA (high). The receiver drives EESDA low so that it remains valid during the high period of the clock, taking into account the setup and hold times. The receiver may be the master or the slave depending on the direction of the data. Typically the receiver acknowledges each byte. If the master is the receiver, it does not generate an acknowledge on the last byte of a transfer. This informs the slave to not drive the next byte of data so that the master may generate a stop or repeated start condition.

Figure 14-1 displays the various bus states of a typical I²C cycle.

FIGURE 14-1: I²C CYCLE



14.3 I²C Master EEPROM Controller

The I²C EEPROM controller supports I²C compatible EEPROMs.

Note: When the EEPROM Loader is running, it has exclusive use of the I²C EEPROM controller. Refer to [Section 14.4, "EEPROM Loader"](#) for more information.

The I²C master implements a low level serial interface (start and stop condition generation, data bit transmission and reception, acknowledge generation and reception) for connection to I²C EEPROMs and consists of a data wire (**EESDA**) and a serial clock (**EESCL**). The serial clock is driven by the master, while the data wire is bi-directional. Both signals are open-drain and require external pull-up resistors.

The I²C master interface runs at the standard-mode rate of 100 kHz. I²C master interface timing information is detailed in [Figure 14-2](#) and [Table 14-1](#).

FIGURE 14-2: I²C MASTER TIMING




TABLE 14-1: I²C MASTER TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
f_{scl}	EESCL clock frequency			100	kHz
t_{high}	EESCL high time	4.0			μ s
t_{low}	EESCL low time	4.7			μ s
t_r	Rise time of EESDA and EESCL			1000	ns
t_f	Fall time of EESDA and EESCL			300	ns
$t_{su,sta}$	Setup time (provided to slave) of EESCL high before EESDA output falling for repeated start condition	5.2 Note 1			μ s
$t_{hd,sta}$	Hold time (provided to slave) of EESCL after EESDA output falling for start or repeated start condition	4.5 Note 1			μ s
$t_{su,dat,in}$	Setup time (from slave) EESDA input before EESCL rising	200 Note 2			ns
$t_{hd,dat,in}$	Hold time (from slave) of EESDA input after EESCL falling	0			ns
$t_{su,dat,out}$	Setup time (provided to slave) EESDA output before EESCL rising	1250 Note 3			ns

TABLE 14-1: I²C MASTER TIMING VALUES (CONTINUED)

Symbol	Description	Min	Typ	Max	Units
$t_{hd;dat;out}$	Hold time (provided to slave) of EESDA output after EESCL falling	1000 Note 3			ns
$t_{su;sto}$	Setup time (provided to slave) of EESCL high before EESDA output rising for stop condition	4.5 Note 1			μs
t_{buf}	Bus free time	4.7			μs
t_{sp}	Input spike suppression on EESCL and EESDA			100	ns

Note 1: These values provide 500 ns of margin compared to the I²C specification.

Note 2: This value provides 50 ns of margin compared to the I²C specification.

Note 3: These values provide 1000 ns of margin compared to the I²C specification.

Based on the [eeprom_size_strap](#) configuration strap, various sized I²C EEPROMs are supported. The varying size ranges are supported by additional bits in the [EEPROM Controller Address \(EPC_ADDRESS\)](#) field of the [EEPROM Command Register \(E2P_CMD\)](#). Within each size range, the largest EEPROM uses all the address bits, while the smaller EEPROMs treat the upper address bits as don't cares. The EEPROM controller drives all the address bits as requested regardless of the actual size of the EEPROM. The supported size ranges for I²C operation are shown in [Table 14-2](#).

TABLE 14-2: I²C EEPROM SIZE RANGES

eeprom_size_strap	# of Address Bytes	EEPROM Size	EEPROM Types
0	1 (Note 4)	128 x 8 through 2048 x 8	24xx01, 24xx02, 24xx04, 24xx08, 24xx16
1	2	4096 x 8 through 65536 x 8	24xx32, 24xx64, 24xx128, 24xx256, 24xx512

Note 4: Bits in the control byte are used as the upper address bits.


14.3.1 I²C EEPROM DEVICE ADDRESSING

The I²C EEPROM is addressed for a read or write operation by first sending a control byte followed by the address byte or bytes. The control byte is preceded by a start condition. The control byte and address byte(s) are each acknowledged by the EEPROM slave. If the EEPROM slave fails to send an acknowledge, then the sequence is aborted (a start condition and a stop condition are sent) and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#) is set.

The control byte consists of a 4 bit control code, 3 bits of chip/block select and one direction bit. The control code is 1010b. For single byte addressing EEPROMs, the chip/block select bits are used for address bits 10, 9 and 8. For double byte addressing EEPROMs, the chip/block select bits are set low. The direction bit is set low to indicate the address is being written.

Figure 14-3 illustrates a typical I²C EEPROM addressing bit order for single and double byte addressing.

FIGURE 14-3: I²C EEPROM ADDRESSING




14.3.2 I²C EEPROM BYTE READ

Following the device addressing, a data byte may be read from the EEPROM by outputting a start condition and control byte with a control code of 1010b, chip/block select bits as described in Section 14.3.1 and the R/~W bit high. The EEPROM will respond with an acknowledge, followed by 8 bits of data. If the EEPROM slave fails to send an acknowledge, then the sequence is aborted (a start condition and a stop condition are sent) and the **EEPROM Controller Timeout (EPC_TIMEOUT)** bit in the **EEPROM Command Register (E2P_CMD)** is set. The I²C master then sends a no-acknowledge, followed by a stop condition.

Figure 14-4 illustrates a typical I²C EEPROM byte read for single and double byte addressing.

FIGURE 14-4: I²C EEPROM BYTE READ




For a register level description of a read operation, refer to Section 14.3.7, "I²C Master EEPROM Controller Operation," on page 463.

14.3.3 I²C EEPROM SEQUENTIAL BYTE READS

Following the device addressing, data bytes may be read sequentially from the EEPROM by outputting a start condition and control byte with a control code of 1010b, chip/block select bits as described in Section 14.3.1 and the R/~W bit high. The EEPROM will respond with an acknowledge, followed by 8 bits of data. If the EEPROM slave fails to send an acknowledge, then the sequence is aborted (a start condition and a stop condition are sent) and the **EEPROM Controller Timeout (EPC_TIMEOUT)** bit in the **EEPROM Command Register (E2P_CMD)** is set. The I²C master then sends an acknowledge and the EEPROM responds with the next 8 bits of data. This continues until the last desired byte is read, at which point the I²C master sends a no-acknowledge (instead of the acknowledge), followed by a stop condition.

Figure 14-5 illustrates a typical I²C EEPROM sequential byte reads for single and double byte addressing.

FIGURE 14-5: I²C EEPROM SEQUENTIAL BYTE READS



Sequential reads are used by the EEPROM Loader. Refer to [Section 14.4, "EEPROM Loader"](#) for additional information. For a register level description of a read operation, refer to [Section 14.3.7, "I²C Master EEPROM Controller Operation,"](#) on page 463.

14.3.4 I²C EEPROM BYTE WRITES


Following the device addressing, a data byte may be written to the EEPROM by outputting the data after receiving the acknowledge from the EEPROM. The data byte is acknowledged by the EEPROM slave and the I²C master finishes the write cycle with a stop condition. If the EEPROM slave fails to send an acknowledge, then the sequence is aborted (a start condition and a stop condition are sent) and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#) is set.

Following the data byte write cycle, the I²C master will poll the EEPROM to determine when the byte write is finished. After meeting the minimum bus free time, a start condition is sent followed by a control byte with a control code of 1010b, chip/block select bits low (since they are don't cares) and the R/W bit low. If the EEPROM is finished with the byte write, it will respond with an acknowledge. Otherwise, it will respond with a no-acknowledge and the I²C master will issue a stop and repeat the poll. If the acknowledge does not occur within 30 ms, a timeout occurs (a start condition and a stop condition are sent) and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#) is set. The check for timeout is only performed following each no-acknowledge, since it may be possible that the EEPROM write finished before the timeout but the 30 ms expired before the poll was performed (due to the bus being used by another master).

Once the I²C master receives the acknowledge, it concludes by sending a start condition, followed by a stop condition, which will place the EEPROM into standby.

Figure 14-6 illustrates a typical I²C EEPROM byte write.

FIGURE 14-6: I²C EEPROM BYTE WRITE



For a register level description of a write operation, refer to [Section 14.3.7, "I2C Master EEPROM Controller Operation," on page 463](#).

14.3.5 WAIT STATE GENERATION

The serial clock is also used as an input as it can be held low by the slave device in order to wait-state the data cycle. Once the slave has data available or is ready to receive, it will release the clock. Assuming the masters clock low time is also expired, the clock will rise and the cycle will continue. If the slave device holds the clock low for more than 30 ms, the current command sequence is aborted (a start condition and a stop condition are not sent since the clock is being held low, instead the clock and data lines are just released) and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#) is set.

14.3.6 I²C BUS ARBITRATION AND CLOCK SYNCHRONIZATION

Since the I²C master and the I²C slave serial interfaces share common pins, there are at least two master I²C devices on the bus (the device and the Host). There exists the potential that both masters try to access the bus at the same time. The I²C specification handles this situation with three mechanisms: bus busy, clock synchronization and bus arbitration.

Note: The timing parameters referred to in the following subsections refer to the detailed timing information presented in the *NXP I²C-Bus Specification*.

14.3.6.1 Bus Busy

A master may start a transfer only if the bus is not busy. The bus is considered to be busy after the START condition and is considered to be free again t_{buf} time after the STOP condition. The standard mode value of 4.7 μ s is used for t_{buf} since the EEPROM master runs at the standard mode rate. Following reset, it is unknown if the bus is actually busy, since the START condition may have been missed. Therefore, following reset, the bus is initially considered busy and is considered free t_{buf} time after the STOP condition or if clock and data are seen high for 4 ms.

14.3.6.2 Clock Synchronization

Clock synchronization is used, since both masters may be generating different clock frequencies. When the clock is driven low by one master, each other active master will restart its low timer and also drive the clock low. Each master will drive the clock low for its minimum low time and then release it. The clock line will not go high until all masters have released it. The slowest master therefore determines the actual low time. Devices with shorter low timers will wait. Once the clock goes high, each master will start its high timer. The first master to reach its high time will once again drive the clock low. The fastest master therefore determines the actual high time. The process then repeats. Clock synchronization is similar to the cycle stretching that can be done by a slave device, with the exception that a slave device can only extend the low time of the clock. It can not cause the falling edge of the clock.

14.3.6.3 Arbitration

Arbitration involves testing the input data vs. the output data, when the clock goes high, to see if they match. Since the data line is wired-AND'ed, a master transmitting a high value will see a mismatch if another master is transmitting a low value. The comparison is not done when receiving bits from the slave. Arbitration starts with the control byte and, if both masters are accessing the same slave, can continue into address and data bits (for writes) or acknowledge bits (for reads). If desired, a master that loses arbitration can continue to generate clock pulses until the end of the losing byte (note that the ACK on a read is considered the end of the byte) but the losing master may no longer drive any data bits.

It is not permitted for another master to access the EEPROM while the device is using it during startup or due to an EEPROM command. The other master should wait sufficient time or poll the device to determine when the EEPROM is available. This restriction simplifies the arbitration and access process since arbitration will always be resolved when transmitting the 8 control bits during the device addressing or during the Poll Cycles.

If arbitration is lost during the device addressing, the I²C master will return to the beginning of the device addressing sequence and wait for the bus to become free.

If arbitration is lost during a Poll Cycle, the I²C master will return to the beginning of the Poll Cycle sequence and wait for the bus to become free. Note that in this case the 30 ms timeout-counter should not be reset. If the 30 ms timeout should expire while waiting for the bus to become free, the sequence should not abort without first completing a final poll (with the exception of the busy / arbitration timeout described in [Section 14.3.6.4](#)).

14.3.6.4 Timeout Due to Busy or Arbitration

It is possible for another master to monopolize the bus (due to a continual bus busy or more successful arbitration). If successful arbitration is not achieved within 1.92 s from the start of the read or write request or from the start of the Poll Cycle, the command sequence or Poll Cycle is aborted and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#) is set. Note that this is a total timeout value and not the timeout for any one portion of the sequence.

14.3.7 I²C MASTER EEPROM CONTROLLER OPERATION

I²C master EEPROM operations are performed using the [EEPROM Command Register \(E2P_CMD\)](#) and [EEPROM Data Register \(E2P_DATA\)](#).

The following operations are supported:

- READ (Read Location)
- WRITE (Write Location)
- RELOAD (EEPROM Loader Reload - See [Section 14.4, "EEPROM Loader"](#))

Note: The EEPROM Loader uses the READ command only.

The supported commands are detailed in [Section 14.5.1, "EEPROM Command Register \(E2P_CMD\)," on page 470](#). Details specific to each operational mode are explained in [Section 14.2, "I2C Overview," on page 457](#) and [Section 14.4, "EEPROM Loader"](#), respectively.

When issuing a WRITE command, the desired data must first be written into the [EEPROM Data Register \(E2P_DATA\)](#). The WRITE command may then be issued by setting the [EEPROM Controller Command \(EPC_COMMAND\)](#) field of the [EEPROM Command Register \(E2P_CMD\)](#) to the desired command value. If the operation is a WRITE, the [EEPROM Controller Address \(EPC_ADDRESS\)](#) field in the [EEPROM Command Register \(E2P_CMD\)](#) must also be set to the desired location. The command is executed when the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#) is set. The completion of the operation is indicated when the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit is cleared.

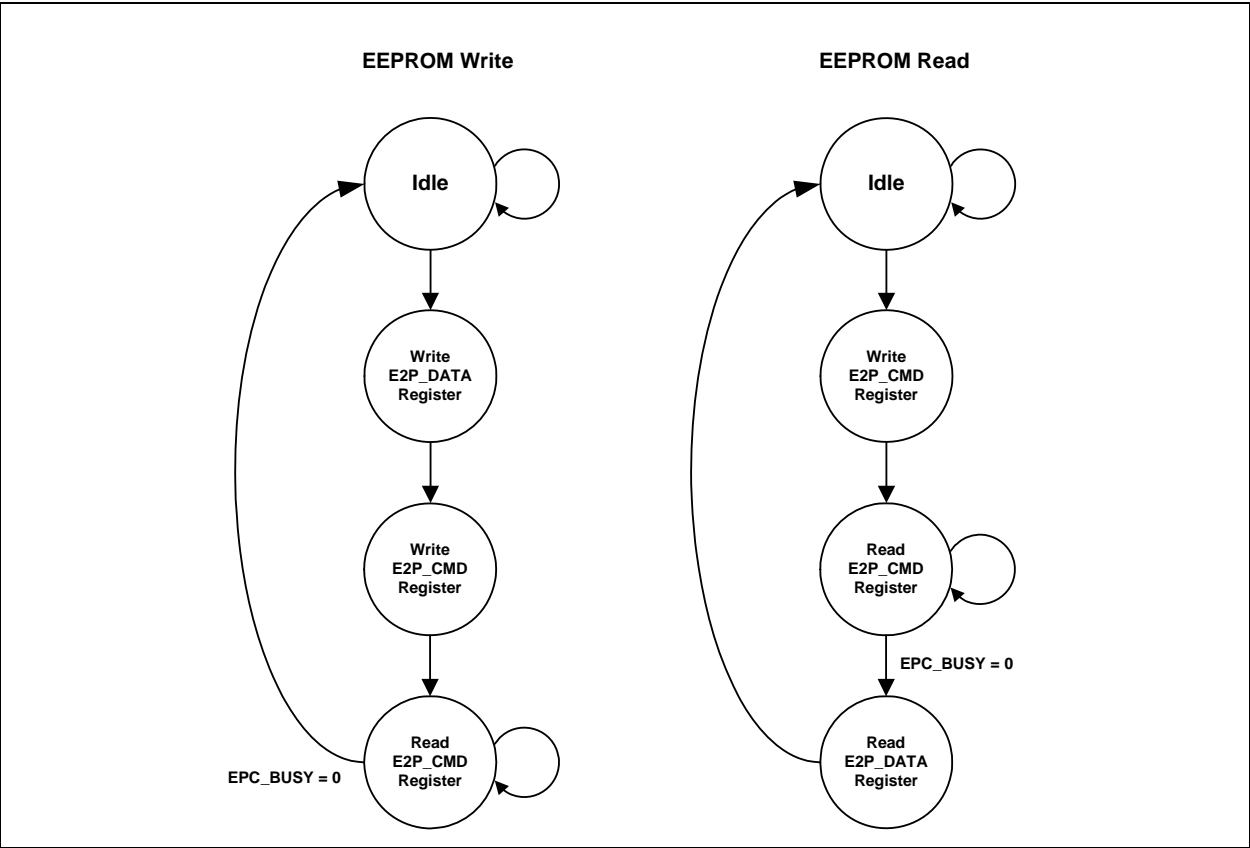
When issuing a READ command, the [EEPROM Controller Command \(EPC_COMMAND\)](#) and [EEPROM Controller Address \(EPC_ADDRESS\)](#) fields of the [EEPROM Command Register \(E2P_CMD\)](#) must be configured with the desired command value and the read address, respectively. The READ command is executed by setting the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#). The completion of the operation is indicated when the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit is cleared, at which time the data from the EEPROM may be read from the [EEPROM Data Register \(E2P_DATA\)](#).

The RELOAD operation is performed by writing the RELOAD command into the [EEPROM Controller Command \(EPC_COMMAND\)](#) field of the [EEPROM Command Register \(E2P_CMD\)](#). The command is executed by setting the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#). In all cases, the software must wait for the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit to clear before modifying the [EEPROM Command Register \(E2P_CMD\)](#).

If an operation is attempted and the EEPROM device does not respond within 30 ms, the device will timeout and the [EEPROM Controller Timeout \(EPC_TIMEOUT\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#) will be set.

Figure 14-7 illustrates the process required to perform an EEPROM read or write operation.

FIGURE 14-7: EEPROM ACCESS FLOW DIAGRAM



14.4 EEPROM Loader

The EEPROM Loader interfaces to the I²C EEPROM controller, the PHYs and to the system CSRs (via the Register Access MUX). All system CSRs are accessible to the EEPROM Loader.

The EEPROM Loader runs upon a pin reset (**RST#**), power-on reset (POR), digital reset ([Digital Reset \(DIGITAL_RST\)](#) bit in the [Reset Control Register \(RESET_CTL\)](#)) or upon the issuance of a RELOAD command via the [EEPROM Command Register \(E2P_CMD\)](#). Refer to [Section 6.2, "Resets," on page 42](#) for additional information on resets.

The EEPROM contents must be loaded in a specific format for use with the EEPROM Loader. An overview of the EEPROM content format is shown in [Table 14-3](#). Each section of EEPROM contents is discussed in detail in the following sections.


TABLE 14-3: EEPROM CONTENTS FORMAT OVERVIEW

EEPROM Address	Description	Value
0	EEPROM Valid Flag	A5h
1	MAC Address Low Word [7:0]	1 st Byte on the Network
2	MAC Address Low Word [15:8]	2 nd Byte on the Network
3	MAC Address Low Word [23:16]	3 rd Byte on the Network
4	MAC Address Low Word [31:24]	4 th Byte on the Network
5	MAC Address High Word [7:0]	5 th Byte on the Network
6	MAC Address High Word [15:8]	6 th Byte on the Network
7	Configuration Strap Values Valid Flag	A5h
8 - 16	Configuration Strap Values	See Table 14-4
17	Burst Sequence Valid Flag	A5h
18	Number of Bursts	See Section 14.4.5, "Register Data"
19 and above	Burst Data	See Section 14.4.5, "Register Data"

14.4.1 EEPROM LOADER OPERATION

Upon a pin reset (**RST#**), power-on reset (POR), digital reset ([Digital Reset \(DIGITAL_RST\)](#) bit in the [Reset Control Register \(RESET_CTL\)](#)) or upon the issuance of a RELOAD command via the [EEPROM Command Register \(E2P_CMD\)](#), the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#) will be set. While the EEPROM Loader is active, the [Device Ready \(READY\)](#) bit of the [Hardware Configuration Register \(HW_CFG\)](#) is cleared and no writes to the device should be attempted. The operational flow of the EEPROM Loader can be seen in [Figure 14-8](#).

FIGURE 14-8: EEPROM LOADER FLOW DIAGRAM



14.4.2 EEPROM VALID FLAG

Following the release of **RST#**, **POR**, **DIGITAL_RST** or a **RELOAD** command, the EEPROM Loader starts by reading the first byte of data from the EEPROM. If the value of A5h is not read from the first byte, the EEPROM Loader will load the current configuration strap values into the registers, restart PHY Auto-negotiation and then terminate, clearing the **EEPROM Controller Busy (EPC_BUSY)** bit in the **EEPROM Command Register (E2P_CMD)**. Otherwise, the EEPROM Loader will continue reading sequential bytes from the EEPROM.

14.4.3 MAC ADDRESS

The next six bytes in the EEPROM, after the EEPROM Valid Flag, are written into the **Host MAC Address High Register (HMAC_ADDRH)** and **Host MAC Address Low Register (HMAC_ADDRL)** registers and the **Switch Fabric MAC Address High Register (SWITCH_MAC_ADDRH)** and **Switch Fabric MAC Address Low Register (SWITCH_MAC_ADDRL)**. The EEPROM bytes are written into the MAC address registers in the order specified in [Table 14-3](#).

14.4.3.1 Host MAC Address Reload

While the EEPROM Loader is in the wait state, if a Host MAC reset is detected (via the **Host MAC Reset (HMAC_RST)** bit in the **Reset Control Register (RESET_CTL)**), the EEPROM Loader will read byte 0. If the byte 0 value is A5h, the EEPROM Loader will read bytes 1 through 6 from the EEPROM and reload the **Host MAC Address High Register (HMAC_ADDRH)** and **Host MAC Address Low Register (HMAC_ADDRL)**. During this time, the **EPC_BUSY** bit in the **EEPROM Command Register (E2P_CMD)** is set and **Device Ready (READY)** bit of the **Hardware Configuration Register (HW_CFG)** is cleared.

Note: The switch MAC address registers are not reloaded due to this condition.

14.4.4 SOFT-STRAPS

The 7th byte of data to be read from the EEPROM is the Configuration Strap Values Valid Flag. If this byte has a value of A5h, the next 9 bytes of data (8-16) are written into the configuration strap registers per the assignments detailed in [Table 14-4](#).

If the flag byte is not A5h, these next 9 bytes are skipped (they are still read to maintain the data burst, but are discarded). However, the current configuration strap values are still loaded into the registers and the PHY Auto-negotiation is still restarted. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for more information on configuration straps.

Note: Bit locations in [Table 14-4](#) that do not define a configuration strap must be written as 0.

TABLE 14-4: EEPROM CONFIGURATION BITS

Byte/Bit	7	6	5	4	3	2	1	0
Byte 8	BP_EN_strap_1	FD_FC_strap_1	manual_FC_strap_1	manual_mdix_strap_1	auto_mdix_strap_1	speed_strap_1	duplex_strap_1	autoneg_strap_1
Byte 9	BP_EN_strap_2	FD_FC_strap_2	manual_FC_strap_2	manual_mdix_strap_2	auto_mdix_strap_2	speed_strap_2	duplex_strap_2	autoneg_strap_2
Byte 10			BP_EN_strap_0	FD_FC_strap_0	manual_F-C_strap_0			SQE_test_dis-able_strap_0
Byte 11	1588_enable_strap		LED_fun_strap[2]	LED_fun_strap[1]	LED_fun_strap[0]	EEE_enable_strap_2	EEE_enable_strap_1	
Byte 12			LED_en_strap[5]	LED_en_strap[4]	LED_en_strap[3]	LED_en_strap[2]	LED_en_strap[1]	LED_en_strap[0]
Byte 13								
Byte 14	HBI_ale_qualification_strap	HBI_rw_mode_strap	HBI_cs_polarity_strap	HBI_rd_rd_wr_polarity_strap	HBI_wr_en_polarity_strap	HBI_ale_polarity_strap		
Byte 15								

TABLE 14-4: EEPROM CONFIGURATION BITS (CONTINUED)

Byte/Bit	7	6	5	4	3	2	1	0
Byte 16								

14.4.5 REGISTER DATA

Optionally following the configuration strap values, the EEPROM data may be formatted to allow access to the device's parallel, directly writable registers. Access to indirectly accessible registers is achievable with an appropriate sequence of writes (at the cost of EEPROM space).

This data is first preceded with a Burst Sequence Valid Flag (EEPROM byte 17). If this byte has a value of A5h, the data that follows is recognized as a sequence of bursts. Otherwise, the EEPROM Loader is finished, will go into a wait state and clear the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#). This can optionally generate an interrupt.

The data at EEPROM byte 18 and above should be formatted in a sequence of bursts. The first byte is the total number of bursts. Following this is a series of bursts, each consisting of a starting address, count and the count x 4 bytes of data. This results in the following formula for formatting register data:

8 bits number_of_bursts

repeat (number_of_bursts)

16 bits {starting_address[9:2] / count[7:0]}

repeat (count)

8 bits data[31:24], 8 bits data[23:16], 8 bits data[15:8], 8 bits data[7:0]

Note: The starting address is a DWORD address. Appending two 0 bits will form the register address.

As an example, the following is a 3 burst sequence, with 1, 2 and 3 DWORDs starting at register addresses 40h, 80h and C0h respectively:

A5h, (Burst Sequence Valid Flag)

3h, (number_of_bursts)

16{10h, 1h}, (starting_address1 divided by 4 / count1)

11h, 12h, 13h, 14h, (4 x count1 of data)

16{20h, 2h}, (starting_address2 divided by 4 / count2)

21h, 22h, 23h, 24h, 25h, 26h, 27h, 28h, (4 x count2 of data)

16{30h, 3h}, (starting_address3 divided by 4 / count3)

31h, 32h, 33h, 34h, 35h, 36h, 37h, 38h, 39h, 3Ah, 3Bh, 3Ch (4 x count3 of data)

In order to avoid overwriting the Switch CSR interface, MAC CSR or MII Management interfaces, the EEPROM Loader waits until the following bits are cleared before performing any register write:

- [CSR Busy \(CSR_BUSY\)](#) bit of the [Switch Fabric CSR Interface Command Register \(SWITCH_CSR_CMD\)](#)
- [Host MAC CSR Busy \(HMAC_CSR_BUSY\)](#) bit of the [Host MAC CSR Interface Command Register \(MAC_CSR_CMD\)](#)
- [MII Busy \(MIIBZY\)](#) bit of the [Host MAC MII Access Register \(HMAC_MII_ACC\)](#)

The EEPROM Loader checks that the EEPROM address space is not exceeded. If so, it will stop and set the [EEPROM Loader Address Overflow \(LOADER_OVERFLOW\)](#) bit in the [EEPROM Command Register \(E2P_CMD\)](#). The address limit is based on the [eeprom_size_strap](#) which specifies a range of sizes. The address limit is set to the largest value of the specified range.

14.4.6 EEPROM LOADER FINISHED WAIT-STATE

Once finished with the last burst, the EEPROM Loader will go into a wait-state and the [EEPROM Controller Busy \(EPC_BUSY\)](#) bit of the [EEPROM Command Register \(E2P_CMD\)](#) will be cleared. This can optionally generate an interrupt.

14.5 I²C Master EEPROM Controller Registers

This section details the directly addressable I²C Master EEPROM Controller related System CSRs. These registers should only be used if an EEPROM has been connected to the device. For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map," on page 33](#).

TABLE 14-5: I²C MASTER EEPROM CONTROLLER REGISTERS

Address	Register Name (SYMBOL)
1B4h	EEPROM Command Register (E2P_CMD)
1B8h	EEPROM Data Register (E2P_DATA)

14.5.1 EEPROM COMMAND REGISTER (E2P_CMD)

Offset: 1B4h Size: 32 bits

This read/write register is used to control the read and write operations of the serial EEPROM.

Bits	Description	Type	Default
31	<p>EEPROM Controller Busy (EPC_BUSY) When a 1 is written into this bit, the operation specified in the EPC_COMMAND field of this register is performed at the specified EEPROM address. This bit will remain set until the selected operation is complete. In the case of a read, this indicates that the Host can read valid data from the EEPROM Data Register (E2P_DATA). The E2P_CMD and E2P_DATA registers should not be modified until this bit is cleared. In the case where a write is attempted and an EEPROM is not present, the EPC_BUSY bit remains set until the EEPROM Controller Timeout (EPC_TIMEOUT) bit is set. At this time the EPC_BUSY bit is cleared.</p> <p>Note: EPC_BUSY is set immediately following power-up, or pin reset, or DIGITAL_RST reset. This bit is also set following the setting of the Host MAC Reset (HMAC_RST) bit in the Reset Control Register (RESET_CTL). After the EEPROM Loader has finished loading, the EPC_BUSY bit is cleared. Refer to chapter Section 14.4, "EEPROM Loader," on page 465 for more information.</p>	R/W SC	0b

Bits	Description	Type	Default																																				
30:28	<p>EEPROM Controller Command (EPC_COMMAND) This field is used to issue commands to the EEPROM controller. The EEPROM controller will execute a command when the EPC_BUSY bit is set. A new command must not be issued until the previous command completes. The field is encoded as follows:</p> <table border="1"> <thead> <tr> <th>[30]</th><th>[29]</th><th>[28]</th><th>Operation</th></tr> </thead> <tbody> <tr> <td>0</td><td>0</td><td>0</td><td>READ</td></tr> <tr> <td>0</td><td>0</td><td>1</td><td>RESERVED</td></tr> <tr> <td>0</td><td>1</td><td>0</td><td>RESERVED</td></tr> <tr> <td>0</td><td>1</td><td>1</td><td>WRITE</td></tr> <tr> <td>1</td><td>0</td><td>0</td><td>RESERVED</td></tr> <tr> <td>1</td><td>0</td><td>1</td><td>RESERVED</td></tr> <tr> <td>1</td><td>1</td><td>0</td><td>RESERVED</td></tr> <tr> <td>1</td><td>1</td><td>1</td><td>RELOAD</td></tr> </tbody> </table> <p>Note: Only the READ, WRITE and RELOAD commands are valid for I²C mode. If an unsupported command is attempted, the EPC_BUSY bit will be cleared and EPC_TIMEOUT will be set.</p> <p>The EEPROM operations are defined as follows:</p> <p>READ (Read Location) This command will cause a read of the EEPROM location pointed to by the EPC_ADDRESS bit field. The result of the read is available in the EEPROM Data Register (E2P_DATA).</p> <p>WRITE (Write Location) If erase/write operations are enabled in the EEPROM, this command will cause the contents of the EEPROM Data Register (E2P_DATA) to be written to the EEPROM location selected by the EPC_ADDRESS field.</p> <p>RELOAD (EEPROM Loader Reload) Instructs the EEPROM Loader to reload the device from the EEPROM. If a value of A5h is not found in the first address of the EEPROM, the EEPROM is assumed to be un-programmed and the RELOAD operation will fail. The CFG_LOADED bit indicates a successful load. Following this command, the device will enter the not ready state. The Device Ready (READY) bit in the Hardware Configuration Register (HW_CFG) should be polled to determine then the RELOAD is complete.</p>	[30]	[29]	[28]	Operation	0	0	0	READ	0	0	1	RESERVED	0	1	0	RESERVED	0	1	1	WRITE	1	0	0	RESERVED	1	0	1	RESERVED	1	1	0	RESERVED	1	1	1	RELOAD	R/W	000b
[30]	[29]	[28]	Operation																																				
0	0	0	READ																																				
0	0	1	RESERVED																																				
0	1	0	RESERVED																																				
0	1	1	WRITE																																				
1	0	0	RESERVED																																				
1	0	1	RESERVED																																				
1	1	0	RESERVED																																				
1	1	1	RELOAD																																				
27:19	RESERVED	RO	-																																				
18	<p>EEPROM Loader Address Overflow (LOADER_OVERFLOW) This bit indicates that the EEPROM Loader tried to read past the end of the EEPROM address space. This indicates misconfigured EEPROM data.</p> <p>This bit is cleared when the EEPROM Loader is restarted with a RELOAD command, Host MAC Reset (HMAC_RESET), or a Digital Reset (DIGITAL_RST).</p>	RO	0b																																				

Bits	Description	Type	Default
17	EEPROM Controller Timeout (EPC_TIMEOUT) This bit is set when a timeout occurs, indicating the last operation was unsuccessful. If an EEPROM WRITE operation is performed and no response is received from the EEPROM within 30 ms, the EEPROM controller will timeout and return to its idle state. This bit is also set if the EEPROM fails to respond with the appropriate ACKs, if the EEPROM slave device holds the clock low for more than 30 ms, if the I2C bus is not acquired within 1.92 seconds, or if an unsupported EPC_COMMAND is attempted. This bit is cleared when written high.	R/WC	0b
16	Configuration Loaded (CFG_LOADED) When set, this bit indicates that a valid EEPROM was found and the EEPROM Loader completed normally. This bit is set upon a successful load. It is cleared on power-up, pin and DIGITAL_RST resets, Host MAC Reset (HMAC_RESET), or at the start of a RELOAD. This bit is cleared when written high.	R/WC	0b
15:0	EEPROM Controller Address (EPC_ADDRESS) This field is used by the EEPROM Controller to address a specific memory location in the serial EEPROM. This address must be byte aligned.	R/W	0000h

14.5.2 EEPROM DATA REGISTER (E2P_DATA)

Offset: 1B8h Size: 32 bits

This read/write register is used in conjunction with the [EEPROM Command Register \(E2P_CMD\)](#) to perform read and write operations with the serial EEPROM.

Bits	Description	Type	Default
31:8	RESERVED	RO	-
7:0	EEPROM Data (EEPROM_DATA) This field contains the data read from or written to the EEPROM.	R/W	00h

15.0 IEEE 1588

15.1 Functional Overview

The device provides hardware support for the IEEE 1588-2008 Precision Time Protocol (PTP), allowing clock synchronization with remote Ethernet devices, packet time stamping, and time driven event generation.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

Time stamping is supported on all ports, with an individual PTP Timestamp sub-module connected to each port. Any port may function as a master or a slave clock per the IEEE 1588-2008 specification, and the device as a whole may function as a transparent or boundary clock. Both end-to-end and peer-to-peer link delay mechanisms are supported as are one-step and two-step operations.

A 32-bit seconds and 30-bit nanoseconds tunable clock is provided that is used as the time source for all PTP timestamp related functions. A 1588 Clock Events sub-module provides 1588 Clock comparison based interrupt generation and timestamp related GPIO event generation. GPIO pins can be used to trigger a timestamp capture when configured as an input, or output a signal based on a 1588 Clock Target compare event.

All features of the IEEE 1588 unit can be monitored and configured via their respective configuration and status registers. A detailed description of all 1588 CSRs is included in [Section 15.8, "1588 Registers"](#).

15.1.1 IEEE 1588-2008

IEEE 1588-2008 specifies a Precision Time Protocol (PTP) used by master and slave clock devices to pass time information in order to achieve clock synchronization. Ten network message types are defined:

- Sync
- Follow_Up
- Delay_Req
- Delay_Resp
- PDelay_Req
- PDelay_Resp
- PDelay_Resp_Follow_Up
- Announce
- Signaling
- Management

The first seven message types are used for clock synchronization. Using these messages, the protocol software may calculate the offset and network delay between timestamps, adjusting the slave clock frequency as needed. Refer to the IEEE 1588-2008 protocol for message definitions and proper usage.

A PTP domain is segmented into PTP sub-domains, which are then segmented into PTP communication paths. Within each PTP communication path there is a maximum of one master clock, which is the source of time for each slave clock. The determination of which clock is the master and which clock(s) is(are) the slave(s) is not fixed, but determined by the IEEE 1588-2008 protocol. Similarly, each PTP sub-domain may have only one master clock, referred to as the Grand Master Clock.

PTP communication paths are conceptually equivalent to Ethernet collision domains and may contain devices which extend the network. However, unlike Ethernet collision domains, the PTP communication path does not stop at a network switch, bridge, or router. This leads to a loss of precision when the network switch/bridge/router introduces a variable delay. Boundary clocks are defined which conceptually bypass the switch/bridge/router (either physically or via device integration). Essentially, a boundary clock acts as a slave to an upstream master, and as a master to a downstream slave. A boundary clock may contain multiple ports, but a maximum of one slave port is permitted.

Although boundary clocks solve the issue of the variable delay influencing the synchronization accuracy, they add clock jitter as each boundary clock tracks the clock of its upstream master. Another approach that is supported is the concept of transparent clocks. These devices measure the delay they have added when forwarding a message (the residence time) and report this additional delay either in the forwarded message (one-step) or in a subsequent message (two-step).

The PTP relies on the knowledge of the path delays between the master and the slave. With this information, and the knowledge of when the master has sent the packet, a slave can calculate its clock offset from the master and make appropriate adjustments. There are two methods of obtaining the network path delay. Using the end-to-end method, packets are exchanged between the slave and the master. Any intermediate variable bridge or switch delays are com-

compensated by the transparent clock method described above. Using the round trip time and accounting for the residence time reported, the slave can calculate the mean delay from the master. Each slave sends and receives its own messages and calculates its own delay. While the end-to-end method is the simplest, it does add burden on the master since the master must process packets from each slave in the system. This is amplified when boundary clocks are replaced by transparent clocks. Also, the end-to-end delays must be recalculated if there is a change in the network topology. Using the peer-to-peer method, packets are exchanged only between adjacent master, slaves and transparent clocks. Each peer pair calculates the receive path delay. As time synchronization packets are forwarded between the master and the slave, the transparent clock adds the pre-measured receive path delay into the residence time. The final receiver adds its receive path delay. Using the peer-to-peer method, the full path delay is accounted for without the master having to service each slave. The peer-to-peer method better supports network topology changes since each path delay is kept up-to-date regardless of the port status.

The PTP implementation consists of the following major function blocks:

- [PTP Timestamp and Residence Time Correction](#)
This block provides time stamping and packet modification functions.
- [1588 Clock](#)
This block provides a tunable clock that is used as the time source for all PTP timestamp related functions.
- [1588 Clock Events](#)
This block provides clock comparison-based interrupt generation and timestamp related GPIO event generation.
- [1588 GPIOs](#)
This block provides for time stamping GPIO input events and for outputting clock comparison-based interrupt status.
- [1588 Interrupt](#)
This block provides interrupt generation, masking and status.
- [1588 Registers](#)
This block provides contains all configuration, control and status registers.

15.2 PTP Timestamp and Residence Time Correction

This sub-module handles all PTP packet tasks related to recording timestamps of packets, accounting for the frame forwarding delay through the switch and inserting timestamps into packets.

Modes supported are:

- Boundary Clock, master and slave, one-step and two-step, end-to-end or peer-to-peer delay
 - All 1588 packets are to and from the Host MAC (as directed by switch core)
 - Special Host VLAN tagging (via switch core) indicates ingress port and desired egress port
 - RX and TX timestamps saved in registers for S/W
 - RX timestamp stored in packet for ease of retrieval by S/W
 - Egress timestamp of Sync packet inserted on-the-fly for one-step
 - TX timestamp of Delay_Req packet stored in received Delay_Resp packet for ease of retrieval
 - Correction Field and ingress timestamp of Pdelay_Req packet saved in registers for one-step turnaround time
 - Correction Field of Pdelay_Resp packet automatically calculated and inserted on-the-fly for one-step
 - PTP checksums and Ethernet FCS updated on-the-fly
 - Ingress and egress timestamps corrected for latency
 - Asymmetry corrections
 - Peer delay correction on received Sync packets
- Transparent Clock with Ordinary Clock, master and slave, one-step and two-step, end-to-end or peer-to-peer delay
 - Peer-to-peer received 1588 packets forwarded to Host (peer-to-peer mode) or to other network port (end-to-end mode) (as directed by switch core)
 - All other received 1588 packets forwarded to Host and other network port (as directed by switch core)
 - Special Host VLAN tagging (via switch core) indicates ingress port and desired egress port
 - RX and TX timestamps saved in registers for S/W
 - RX timestamp stored in packet for ease of retrieval by S/W
 - Residence time correction on forwarded Sync, Delay_Req, Pdelay_Req and Pdelay_Resp packets
ingress timestamp subtracted from Correction Field on receive
egress timestamp added to Correction Field on-the-fly during transmit
 - Egress timestamp of Host Sync packet inserted on-the-fly for one-step (for master Ordinary Clock)
 - Correction Field and ingress timestamp of Pdelay_Req packet saved in registers for one-step turnaround time (peer-to-peer mode)
 - Correction Field of Host Pdelay_Resp packet automatically calculated and inserted on-the-fly for one-step (peer-to-peer mode)
 - PTP checksums and Ethernet FCS updated on-the-fly
 - Ingress and egress timestamps corrected for latency
 - Asymmetry corrections
 - Peer delay correction on received Sync packets

Functions include:

- Detecting a PTP packet
 - 802.3/SNAP or Ethernet II encoding
 - Skipping over VLAN tags
 - Ethernet, IPv4 or IPv6 message formats
 - Skipping over IP extension headers
 - Checking the MAC and / or the IP addresses
- Recording the timestamp of received packets into registers
 - Accounting for the ingress latency
- Recording the timestamp of received packets into the packet and updating the layer 3 checksum and layer 2 FCS fields
 - Accounting for the ingress latency

LAN9352

- Forwarding or filtering PTP packets as needed to support ordinary, boundary or transparent clock mode
- Recording the timestamp of transmitted packets into registers
 - Accounting for the egress latency
- Updating the correction field to account for the residence time in the switch and updating the layer 3 checksum and layer 2 FCS
 - Accounting for the peer delay and link asymmetry
- One-step on-the-fly timestamp insertion for Sync packets and updating the layer 3 checksum and layer 2 FCS
- One-step on-the-fly turnaround time insertion for Pdelay_Req packets and updating the layer 3 checksum and layer 2 FCS

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

Three instances of this sub-module are used, however the sub-module on the port that connects to the Host MAC typically would not be configured to operate.

15.2.1 RECEIVE FRAME PROCESSING

15.2.1.1 Ingress Time Snapshot

For each Ethernet frame, the receive frame processing detects the SFD field of the frame and temporarily saves the current 1588 Clock value.

INGRESS LATENCY

The ingress latency is the amount of time between the start of the frame's first symbol after the SFD on the network medium and the point when the 1588 clock value is internally captured. It is specified by the [RX Latency \(RX_LATENCY\[15:0\]\)](#) field in the [1588 Port x Latency Register \(1588_LATENCY_x\)](#) and is subtracted from the 1588 Clock value at the detection of the SFD. The setting is used to adjust the internally captured 1588 clock value such that the resultant timestamp more accurately corresponds to the start of the frame's first symbol after the SFD on the network medium.

The ingress latency consists of the receive latency of the PHY and the latency of the 1588 frame detection circuitry. The value depends on the port mode. Typical values are:

- 100BASE-TX: 285ns
- 100BASE-FX: 231ns plus the receive latency of the fiber transceiver
- 10BASE-T: 1674ns

15.2.1.2 1588 Receive Parsing

The 1588 Receive parsing block parses the incoming frame to identify 1588 PTP messages.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

The Receive parsing block may be programmed to detect PTP messages encoded in UDP/IPv4, UDP/IPv6 and Layer 2 Ethernet formats via the [RX IPv4 Enable \(RX_IPV4_EN\)](#), [RX IPv6 Enable \(RX_IPV6_EN\)](#) and [RX Layer 2 Enable \(RX_LAYER2_EN\)](#) bits in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#).

VLAN tagged and non-VLAN tagged frame formats are supported. Multiple VLAN tags are handled as long as they all use the standard type of 0x8100. Both Ethernet II (type field) and 802.3 (length field) w/ SNAP frame formats are supported.

The following tests are made to determine that the packet is a PTP message.

- MAC Destination Address checking is enabled via the [RX MAC Address Enable \(RX_MAC_ADDR_EN\)](#) in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#).

For the Layer 2 message format, the addresses of 01:1B:19:00:00:00 or 01:80:C2:00:00:0E may be enabled via the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). Either address is allowed for Peer delay and non-Peer delay messages.

For IPv4/UDP messages, any of the IANA assigned multicast IP destination addresses for IEEE 1588 (224.0.1.129 and 224.0.1.130 through .132), as well as the IP destination address for the Peer Delay Mechanism (224.0.0.107) may be enabled via the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). These IP addresses map to the 802.3 MAC addresses of 01:00:5e:00:01:81 through 01:00:5e:00:01:84 and 01:00:5e:00:00:6B. Any of these addresses are allowed for Peer delay and non-Peer delay messages.

For IPv6/UDP messages, any of the IANA assigned multicast IP destination addresses for IEEE 1588 (FF0X:0:0:0:0:0:181 and FF0X:0:0:0:0:0:182 through :184), as well as the IP destination address for the Peer Delay Mechanism (FF02:0:0:0:0:0:6B) may be enabled via the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). These IP addresses map to the 802.3 MAC addresses of 33:33:00:00:01:81 through 33:33:00:00:01:84 and 33:33:00:00:00:6B. Any of these addresses are allowed for Peer delay and non-Peer delay messages.

A user defined MAC address defined in the [1588 User MAC Address High-WORD Register \(1588_USER_MAC_HI\)](#) and the [1588 User MAC Address Low-DWORD Register \(1588_USER_MAC_LO\)](#) may also be individually enabled for the above formats.

- If the Type / Length Field indicates an EtherType then

For the Layer 2 message format, the EtherType must equal 0x88F7.

For IPv4/UDP messages, the EtherType must equal 0x0800.

For IPv6/UDP messages, the EtherType must equal 0x86DD.

- If the Type / Length Field indicates a Length and the next 3 bytes equal 0xAAAA03 (indicating that a SNAP header is present) and the SNAP header has a OUI equal to 0x000000 then

For the Layer 2 message format, the EtherType in the SNAP header must equal 0x88F7.

For IPv4/UDP messages, the EtherType in the SNAP header must equal 0x0800.

For IPv6/UDP messages, the EtherType in the SNAP header must equal 0x86DD.

- For IPv4/UDP messages, the Version field in the IPv4 header must equal 4, the IHL field must be 5 and the Protocol field must equal 17 (UDP) or 51 (AH). IPv4 options are not supported.
- For IPv6/UDP messages, the Version field in the IPv6 header must equal 6 and the Next Header field must equal 17 (UDP) or one of the IPv6 extension header values (0 - Hop-by-Hop Options, 60 - Destination Options, 43 - Routing, 44 - Fragment, 51 - Authentication Header (AH))
- For IPv4/UDP messages, Destination IP Address checking is enabled via the [RX IP Address Enable \(RX_IP_ADDR_EN\)](#) in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). Any of the IANA assigned multicast IP destination addresses for IEEE 1588 (224.0.1.129 and 224.0.1.130 through .132), as well as the IP destination address for the Peer Delay Mechanism (224.0.0.107) may be enabled via the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). Any of these addresses are allowed for Peer delay and non-Peer delay messages.
- For IPv6/UDP messages, Destination IP Address checking is enabled via the [RX IP Address Enable \(RX_IP_ADDR_EN\)](#) in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). Any of the IANA assigned multicast IP destination addresses for IEEE 1588 (FF0X:0:0:0:0:0:181 and FF0X:0:0:0:0:0:182 through :184), as well as the IP destination address for the Peer Delay Mechanism (FF02:0:0:0:0:0:6B) may be enabled via the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#). Any of these addresses are allowed for Peer delay and non-Peer delay messages.
- For IPv4/UDP if the Protocol field in the fixed header was 51 (AH), the Next Header field is checked for 17 (UDP) and the AH header is skipped.
- For IPv6/UDP if the Next Header field in the fixed header was one of the IPv6 extension header values, the Next Header field in the extension header is checked for 17 (UDP) or one of the IPv6 extension header values. If it is one of the IPv6 extension header values, the process repeats until either a value of 17 (UDP) or a value of 59 (No Next Header) are found or the packet ends.

15.2.1.3 Receive Message Ingress Time Recording

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked for ALL of the following.

- The messageType field of the PTP header is checked and only those messages enabled via the [RX PTP Message Type Enable \(RX_PTP_MESSAGE_EN\[15:0\]\)](#) bits in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#) will have their ingress times saved. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled.
- The versionPTP field of the PTP header is checked against the [RX PTP Version \(RX_PTP_VERSION\[3:0\]\)](#) field in

the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will have their ingress times saved. A setting of 0 allows any PTP version.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

- If enabled via the [RX PTP Domain Match Enable \(RX_PTP_DOMAIN_EN\)](#) bit in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#), the domainNumber field of the PTP header is checked against the [RX PTP Domain \(RX_PTP_DOMAIN\[7:0\]\)](#) value in the same register. Only those messages with a matching domain will have their ingress times saved.
- If enabled via the [RX PTP Alternate Master Enable \(RX_PTP_ALT_MASTER_EN\)](#) bit in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#), the alternateMasterFlag in the flagField of the PTP header is checked and only those messages with an alternateMasterFlag set to 0 will have their ingress times saved.

At the end of the frame, the frame's FCS and the UDP checksum (for IPv4 and IPv6 formats) are verified. FCS checking can be disabled using the [RX PTP FCS Check Disable \(RX_PTP_FCS_DIS\)](#) bit in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). UDP checksum checking can be disabled using the [RX PTP UDP Checksum Check Disable \(RX_PTP_UDP_CHKSUM_DIS\)](#) bit in the same register.

Note: A IPv4 UDP checksum value of 0x0000 indicates that the checksum is not included and is considered a pass. A IPv6 UDP checksum value of 0x0000 is invalid and is considered a fail.

Note: For IPv6, the UDP checksum calculation includes the IPv6 Pseudo header. Part of the IPv6 Pseudo header is the final IPv6 destination address.

If the IPv6 packet does not contain a Routing header, then the final IPv6 destination address is the destination address contained in the IPv6 header.

If the IPv6 packet does contain a Routing header, then the final IPv6 destination address is the address in the last element of the Routing header.

Note: The UDP checksum is calculated over the entire UDP payload as indicated by the UDP length field and not the assumed PTP packet length.

Note: The UDP checksum calculation does not include layer 2 pad bytes, if any.

If the FCS and checksum tests pass:

- The latency adjusted, 1588 Clock value, saved above at the start of the frame, is recorded into the [1588 Port x RX Ingress Time Seconds Register \(1588_RX_INGRESS_SEC_x\)](#) and [1588 Port x RX Ingress Time NanoSeconds Register \(1588_RX_INGRESS_NS_x\)](#).
- The messageType and sequenceId fields and 12-bit CRC of the portIdentity field of the PTP header are recorded into the [Message Type \(MSG_TYPE\)](#), [Sequence ID \(SEQ_ID\)](#) and [Source Port Identity CRC \(SRC_PRT_CRC\)](#) fields of the [1588 Port x RX Message Header Register \(1588_RX_MSG_HEADER_x\)](#).

The 12-bit CRC of the portIdentity field is created by using the polynomial of $X^{12} + X^{11} + X^3 + X^2 + X + 1$.

- The corresponding maskable [1588 RX Timestamp Interrupt \(1588_RX_TS_INT\[2:0\]\)](#) is set in the [1588 Interrupt Status Register \(1588_INT_STS\)](#).

Up to four receive events are saved per port with the count shown in the [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) field in the [1588 Port x Capture Information Register \(1588_CAP_INFO_x\)](#). Additional events are not recorded. When the appropriate [1588 RX Timestamp Interrupt \(1588_RX_TS_INT\[2:0\]\)](#) bit is written as a one to clear, [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) will decrement. If there are remaining events, the capture registers will update to the next event and the interrupt will set again.

PDELAY_REQ INGRESS TIME SAVING

One-step Pdelay_Resp messages sent by the Host, require their correctionField to be calculated on-the-fly to include the turnaround time between the ingress of the Pdelay_Req and the egress time of the Pdelay_Resp.

The [1588 Port x RX Pdelay_Req Ingress Time Seconds Register \(1588_RX_PDREQ_SEC_x\)](#) and the [1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register \(1588_RX_PDREQ_NS_x\)](#) hold the ingress time of the Pdelay_Req message.

The [1588 Port x RX Pdelay_Req Ingress Correction Field High Register \(1588_RX_PDREQ_CF_HI_x\)](#) and the [1588 Port x RX Pdelay_Req Ingress Correction Field Low Register \(1588_RX_PDREQ_CF_LOW_x\)](#) hold the correctionField of the Pdelay_Req message.

These registers can be set by S/W prior to sending the Pdelay_Resp message.

Alternatively, these registers can be updated by the H/W when the Pdelay_Req message is received. This function is enabled by the [Auto Update \(AUTO\)](#) bit in the [1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register \(1588_RX_PDREQ_NS_x\)](#) independent from the [RX PTP Message Type Enable \(RX_PTP_MESSAGE_EN\[15:0\]\)](#) bits.

As above (including all applicable notes):

- The versionPTP and domainNumber fields and alternateMasterFlag in the flagField of the PTP header are checked, if enabled.
- **Note:** Support for the IEEE 1588-2002 (v1) packet format is not provided.
- At the end of the frame, the frame's FCS and the UDP checksum (for IPv4 and IPv6 formats) are verified, if enabled.

If all tests pass, then the Pdelay_Req message information is updated.

15.2.1.4 Ingress Packet Modifications

INGRESS TIME INSERTION INTO PACKETS

As an alternate to reading the receive time stamp from registers and matching it to the correct frame received in the Host MAC, the saved, latency adjusted, 1588 Clock value can be stored into the packet.

This function is enabled via the [RX PTP Insert Timestamp Enable \(RX_PTP_INSERT_TS_EN\)](#) and [RX PTP Insert Timestamp Seconds Enable \(RX_PTP_INSERT_TS_SEC_EN\)](#) bits in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

Note: Inserting the ingress time into the packet is an additional, separately enabled, feature verses the Ingress Time Recording described above. The capture registers are still updated as is the appropriate [1588 RX Timestamp Interrupt \(1588_RX_TS_INT\[2:0\]\)](#) bit and the [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) field.

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked for ALL of the following.

- The messageType field of the PTP header is checked and only those messages enabled via the [RX PTP Message Type Enable \(RX_PTP_MESSAGE_EN\[15:0\]\)](#) bits in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#) will be have their ingress times inserted. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled.
- The versionPTP field of the PTP header is checked against the [RX PTP Version \(RX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will be have their ingress times inserted. A setting of 0 allows any PTP version.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

Note: The domainNumber field and alternateMasterFlag in the flagField of the PTP header are not tested for purpose of ingress time insertion.

The packet is modified as follows:

- The four bytes of nanoseconds are stored at an offset from the start of the PTP header

The offset is specified in [RX PTP Insert Timestamp Offset \(RX_PTP_INSERT_TS_OFFSET\[5:0\]\)](#) field in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

The lowest two bits of the seconds are stored into the upper 2 bits of the nanoseconds.

- If also enabled, bits 3:0 of the seconds are stored into bits 3:0 of a reserved byte in the PTP header. Bits 7:4 are set to zero.

The offset of this reserved byte is specified by the [RX PTP Insert Timestamp Seconds Offset \(RX_PTP_INSERT_TS_SEC_OFFSET\[5:0\]\)](#) field in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

Note: For version 2 of IEEE 1588, the four reserved bytes starting at offset 16 should be used for the nanoseconds. The reserved byte at offset 5 should be used for the seconds.

DELAY REQUEST EGRESS TIME INSERTION INTO DELAY REPONSE PACKET

LAN9352

Normally, in ordinary clock operation, the egress times of transmitted Delay_Req packets are saved and read by the Host S/W. To avoid the need to read these timestamps via register access, the egress time of the last transmitted Delay_Req packet on the port can be inserted into Delay_Resp packets received on the port.

This function is enabled via the [RX PTP Insert Delay Request Egress in Delay Response Enable \(RX_PT_P_INSERT_DREQ_DRESP_EN\)](#) bit in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

As with any Ingress Time Insertion, Delay_Resp messages must be enable in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#) and the [RX PTP Insert Timestamp Enable \(RX_PT_P_INSERT_TS_EN\)](#) must be set.

Note: Inserting the delay request egress time into the packet is an additional, separately enabled, feature verses the Egress Time Recording described above.

As with [INGRESS TIME INSERTION INTO PACKETS](#), above:

- The versionPTP field of the PTP header is checked and the domainNumber field and alternateMasterFlag in the flagField of the PTP header are not checked.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

- The four bytes of nanoseconds / 2 bits of seconds are stored at the specified offset of the PTP header.
- Bits 3:0 of the seconds are stored at the specified offset in the PTP header, if enabled.

Effectively, this function is the same as the [INGRESS TIME INSERTION INTO PACKETS](#) except that the egress time of the Delay_Req is inserted instead of the ingress time of the Delay_Resp.

INGRESS CORRECTION FIELD RESIDENCE TIME ADJUSTMENT

In order to support one-step transparent clock operation, the residence time delay through the device is accounted for by adjusting the correctionField of certain packets.

This function is enabled per PTP message type via the [RX PTP Correction Field Message Type Enable \(RX_PT_P_CF_MSG_EN\[15:0\]\)](#) bits in the [1588 Port x RX Correction Field Modification Register \(1588_RX_CF_MOD_x\)](#). Typically the Sync message is enabled for both end-to-end and peer-to-peer transparent clocks, the Delay_Req, PDelay_Req and PDelay_Resp messages are enabled only for end-to-end transparent clocks.

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked.

- The versionPTP field of the PTP header is checked against the [RX PTP Version \(RX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will be have their correction field modified. A setting of 0 allows any PTP version.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

Note: The domainNumber field and alternateMasterFlag in the flagField of the PTP header are not tested for purpose of correction field modification.

The correctionField is modified as follows:

Note: If the original correctionField contains a value of 0x7FFFFFFFFFFFFFFF, it is not modified.

If adjustment to the correctionField would result in a value that is larger than 0x7FFFFFFFFFFFFFFF, that value is used instead.

- For **Sync** packets, the value of the [RX Peer Delay \(RX_PEER_DELAY\[15:0\]\)](#) field in the [1588 Port x Asymmetry and Peer Delay Register \(1588_ASYM_PEERDLY_x\)](#) (for the particular ingress port) is added to the correctionField.

This function is used for one-step peer-to-peer transparent clocks. If peer-to-peer transparent clock mode is not being used, the register should be set to zero. If one-step transparent clock mode is not being used, correction field modifications would not be enabled for Sync messages.

- For **Sync** and **Pdelay_Resp** packets, the value of the [Port Delay Asymmetry \(DELAY_ASYM\[15:0\]\)](#) field in the [1588 Port x Asymmetry and Peer Delay Register \(1588_ASYM_PEERDLY_x\)](#) (for the particular ingress port) is added to the correctionField.

This function is used for one-step transparent clocks. If one-step end-to-end transparent clock mode is not being

used, correction field modifications would not be enabled for PDelay_Resp messages. If one-step transparent clock mode is not being used, correction field modifications would not be enabled for Sync or PDelay_Resp messages.

- The nanoseconds portion of the ingress time are subtracted from the correctionField.
- Bits 3:0 of the seconds portion of the ingress time are inserted into bits 3:0 of a reserved byte in the PTP header. Bit 7 is set to a one as an indication to the transmitter that residence time correction is being done. Bits 6:4 are set to zero.

The offset of this reserved byte is specified by the [RX PTP Insert Timestamp Seconds Offset \(RX_PT-P_INSERT_TS_SEC_OFFSET\[5:0\]\)](#) field in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

Note: Proper operation of the transmitter portion of the Correction Field Residence Time Adjustment requires that the reserved byte resides after the versionPTP field and before the correctionField.

For version 2 of IEEE 1588, the reserved byte at offset 5 should be used for the seconds.

Note: Since the modification of the packet occurs on ingress, any packets that are forwarded to the Host software will also have the ingress time subtracted from the original correctionField. If necessary, the original correctionField can be reconstructed by adding the ingress time.

FRAME UPDATING

Frames are modified even if their original FCS or UDP checksum is invalid.

- For IPv4, the UDP checksum is set to 0.

If the original UDP checksum was invalid, a receive symbol error is forced and the [1588 Port x RX Checksum Dropped Count Register \(1588_RX_CHKSUM_DROPPED_CNT_x\)](#) incremented.

This can be disabled by the [RX PTP Bad UDP Checksum Force Error Disable \(RX_PTP_BAD_UDP_CHKSUM_FORCE_ERR_DIS\)](#) field in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

Note: An original UDP checksum value of 0x0000 indicates that the checksum is not included and is considered a pass.

Note: The original UDP checksum is calculated over the entire UDP payload as indicated by the UDP length field and not the assumed PTP packet length.

Note: The original UDP checksum calculation does not include layer 2 pad bytes, if any.

- For IPv6, the two bytes beyond the end of the PTP message are modified so that the original UDP checksum is correct for the modified payload. These bytes are updated by accumulating the differences between the original frame data and the substituted data using the mechanism defined in IETF RFC 1624.

If the original UDP checksum was invalid, a receive symbol error is forced and the [1588 Port x RX Checksum Dropped Count Register \(1588_RX_CHKSUM_DROPPED_CNT_x\)](#) incremented.

This can be disabled by the [RX PTP Bad UDP Checksum Force Error Disable \(RX_PTP_BAD_UDP_CHKSUM_FORCE_ERR_DIS\)](#) field in the [1588 Port x RX Timestamp Insertion Configuration Register \(1588_RX_TS_INSERT_CONFIG_x\)](#).

Note: Since the two bytes beyond the end of the PTP message are modified based on the differences between the original frame data and the substituted data, an invalid incoming checksum would always result in an outgoing checksum error.

Note: An original UDP checksum value of 0x0000 is invalid and is considered a fail.

Note: For IPv6, the UDP checksum calculation includes the IPv6 Pseudo header. Part of the IPv6 Pseudo header is the final IPv6 destination address.

If the IPv6 packet does not contain a Routing header, then the final IPv6 destination address is the destination address contained in the IPv6 header.

If the IPv6 packet does contain a Routing header, then the final IPv6 destination address is the address in the last element of the Routing header.

Note: The original UDP checksum is calculated over the entire UDP payload as indicated by the UDP length field and not the assumed PTP packet length.

Note: The original UDP checksum calculation does not include layer 2 pad bytes, if any.

Note: The two bytes beyond the end of the PTP message are located by using the messageLength field from the PTP header.

- The frame FCS is recomputed.
If the original FCS was invalid, a bad FCS is forced.
 - If the frame has a receive symbol error(s), a receive symbol error indication will be propagated at the same nibble location(s).
- Note:** FCS and UDP checksums are only updated if the frame was actually modified. If no modifications are done, the existing FCS and checksums are left unchanged.

15.2.1.5 Ingress Message Filtering

PTP messages can be filtered upon receive. Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked for ANY of the following.

- The messageType field of the PTP header is checked and those messages that have their [RX PTP Message Type Filter Enable \(RX_PTP_MSG_FLTR_EN\[15:0\]\)](#) bits in the [1588 Port x RX Filter Configuration Register \(1588_RX_FILTER_CONFIG_x\)](#) set will be filtered. Typically Delay_Req and Delay_Resp messages are filtered in peer-to-peer transparent clocks.
- The versionPTP field of the PTP header is checked against the [RX PTP Version \(RX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). If the [RX PTP Version Filter Enable \(RX_PTP_VERSION_FLTR_EN\)](#) bit in the [1588 Port x RX Filter Configuration Register \(1588_RX_FILTER_CONFIG_x\)](#) is set, messages with a non-matching version will be filtered. A version setting of 0 allows any PTP version and would not cause filtering.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

- If enabled via the [RX PTP Domain Filter Enable \(RX_PTP_DOMAIN_FLTR_EN\)](#) bit in the [1588 Port x RX Filter Configuration Register \(1588_RX_FILTER_CONFIG_x\)](#), messages whose domainNumber field in the PTP header does not match the [RX PTP Domain \(RX_PTP_DOMAIN\[7:0\]\)](#) value in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#) will be filtered.
- If enabled via the [RX PTP Alternate Master Filter Enable \(RX_PTP_ALT_MASTER_FLTR_EN\)](#) bit in the [1588 Port x RX Filter Configuration Register \(1588_RX_FILTER_CONFIG_x\)](#), messages whose alternateMasterFlag in the flagField of the PTP header is set will be filtered.

At the end of the frame, the frame's FCS and the UDP checksum (for IPv4 and IPv6 formats) are verified. FCS checking can be disabled using the [RX PTP FCS Check Disable \(RX_PTP_FCS_DIS\)](#) bit in the [1588 Port x RX Timestamp Configuration Register \(1588_RX_TIMESTAMP_CONFIG_x\)](#). UDP checksum checking can be disabled using the [RX PTP UDP Checksum Check Disable \(RX_PTP_UDP_CHKSUM_DIS\)](#) bit in the same register.

Note: A IPv4 UDP checksum value of 0x0000 indicates that the checksum is not included and is considered a pass. A IPv6 UDP checksum value of 0x0000 is invalid and is considered a fail.

Note: For IPv6, the UDP checksum calculation includes the IPv6 Pseudo header. Part of the IPv6 Pseudo header is the final IPv6 destination address.

If the IPv6 packet does not contain a Routing header, then the final IPv6 destination address is the destination address contained in the IPv6 header.

If the IPv6 packet does contain a Routing header, then the final IPv6 destination address is the address in the last element of the Routing header.

Note: The UDP checksum is calculated over the entire UDP payload as indicated by the UDP length field and not the assumed PTP packet length.

Note: The UDP checksum calculation does not include layer 2 pad bytes, if any.

If the FCS and checksum tests pass, the frame is filtered by inserting a receive symbol error and the [1588 Port x RX Filtered Count Register \(1588_RX_FILTERED_CNT_x\)](#) is incremented.

Note: The MAC will count this as an errored packet.

Note: Message filtering is an additional, separately enabled, feature versus any packet ingress time recording and packet modification. Although these functions typically would not be used together on the same message type.

15.2.2 TRANSMIT FRAME PROCESSING

15.2.2.1 Egress Time Snapshot

For each Ethernet frame, the transmit frame processing detects the SFD field of the frame and temporarily saves the current 1588 Clock value.

EGRESS LATENCY

The egress latency is the amount of time between the point when the 1588 clock value is internally captured and the start of the frame's first symbol after the SFD on the network medium. It is specified by the [TX Latency \(TX_LATENCY\[15:0\]\)](#) field in the [1588 Port x Latency Register \(1588_LATENCY_x\)](#) and is added to the 1588 Clock value at the detection of the SFD. The setting is used to adjust the internally captured 1588 clock value such that the resultant timestamp more accurately corresponds to the start of the frame's first symbol after the SFD on the network medium.

The egress latency consists of the transmit latency of the PHY and the latency of the 1588 frame detection circuitry. The value depends on the port mode. Typical values are:

- 100BASE-TX: 95ns
- 100BASE-FX: 68ns plus the transmit latency of the fiber transceiver
- 10BASE-T: 1139ns

15.2.2.2 1588 Transmit Parsing

The 1588 Transmit parsing block parses the outgoing frame to identify 1588 PTP messages.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

The Transmit parsing block may be programmed to detect PTP messages encoded in UDP/IPv4, UDP/IPv6 and Layer 2 Ethernet formats via the [TX IPv4 Enable \(TX_IPV4_EN\)](#), [TX IPv6 Enable \(TX_IPV6_EN\)](#) and [TX Layer 2 Enable \(TX_LAYER2_EN\)](#) bits in the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#).

VLAN tagged and non-VLAN tagged frame formats are supported. Multiple VLAN tags are handled as long as they all use the standard type of 0x8100. Both Ethernet II (type field) and 802.3 (length field) w/ SNAP frame formats are supported.

The following tests are made to determine that the packet is a PTP message.

- MAC Destination Address checking is enabled via the [TX MAC Address Enable \(TX_MAC_ADDR_EN\)](#) in the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#).

For the Layer 2 message format, the addresses of 01:1B:19:00:00:00 or 01:80:C2:00:00:0E may be enabled via the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). Either address is allowed for Peer delay and non-Peer delay messages.

For IPv4/UDP messages, any of the IANA assigned multicast IP destination addresses for IEEE 1588 (224.0.1.129 and 224.0.1.130 through .132), as well as the IP destination address for the Peer Delay Mechanism (224.0.0.107) may be enabled via the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). These IP addresses map to the 802.3 MAC addresses of 01:00:5e:00:01:81 through 01:00:5e:00:01:84 and 01:00:5e:00:00:6B. Any of these addresses are allowed for Peer delay and non-Peer delay messages.

For IPv6/UDP messages, any of the IANA assigned multicast IP destination addresses for IEEE 1588 (FF0X:0:0:0:0:0:181 and FF0X:0:0:0:0:0:182 through :184), as well as the IP destination address for the Peer Delay Mechanism (FF02:0:0:0:0:0:0:6B) may be enabled via the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). These IP addresses map to the 802.3 MAC addresses of 33:33:00:00:01:81 through 33:33:00:00:01:84 and 33:33:00:00:00:6B. Any of these addresses are allowed for Peer delay and non-Peer delay messages.

A user defined MAC address defined in the [1588 User MAC Address High-WORD Register \(1588_USER_MAC_HI\)](#) and the [1588 User MAC Address Low-DWORD Register \(1588_USER_MAC_LO\)](#) may also be individually enabled for the above formats.

- If the Type / Length Field indicates an EtherType then

For the Layer 2 message format, the EtherType must equal 0x88F7.

For IPv4/UDP messages, the EtherType must equal 0x0800.

For IPv6/UDP messages, the EtherType must equal 0x86DD.

- If the Type / Length Field indicates a Length and the next 3 bytes equal 0xAAAA03 (indicating that a SNAP header is present) and the SNAP header has a OUI equal to 0x000000 then

For the Layer 2 message format, the EtherType in the SNAP header must equal 0x88F7.

For IPv4/UDP messages, the EtherType in the SNAP header must equal 0x0800.

For IPv6/UDP messages, the EtherType in the SNAP header must equal 0x86DD.

- For IPv4/UDP messages, the Version field in the IPv4 header must equal 4, the IHL field must be 5 and the Protocol field must equal 17 (UDP) or 51 (AH). IPv4 options are not supported.
- For IPv6/UDP messages, the Version field in the IPv6 header must equal 6 and the Next Header field must equal 17 (UDP) or one of the IPv6 extension header values (0 - Hop-by-Hop Options, 60 - Destination Options, 43 - Routing, 44 - Fragment, 51 - Authentication Header (AH))
- For IPv4/UDP messages, Destination IP Address checking is enabled via the [TX IP Address Enable \(TX_IP_ADDR_EN\)](#) in the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). Any of the IANA assigned multicast IP destination addresses for IEEE 1588 (224.0.1.129 and 224.0.1.130 through .132), as well as the IP destination address for the Peer Delay Mechanism (224.0.0.107) may be enabled via the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). Any of these addresses are allowed for Peer delay and non-Peer delay messages.
- For IPv6/UDP messages, Destination IP Address checking is enabled via the [TX IP Address Enable \(TX_IP_ADDR_EN\)](#) in the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). Any of the IANA assigned multicast IP destination addresses for IEEE 1588 (FF0X:0:0:0:0:0:181 and FF0X:0:0:0:0:0:182 through :184), as well as the IP destination address for the Peer Delay Mechanism (FF02:0:0:0:0:0:6B) may be enabled via the [1588 Port x TX Parsing Configuration Register \(1588_TX_PARSE_CONFIG_x\)](#). Any of these addresses are allowed for Peer delay and non-Peer delay messages.
- For IPv4/UDP if the Protocol field in the fixed header was 51 (AH), the Next Header field is checked for 17 (UDP) and the AH header is skipped.
- For IPv6/UDP if the Next Header field in the fixed header was one of the IPv6 extension header values, the Next Header field in the extension header is checked for 17 (UDP) or one of the IPv6 extension header values. If it is one of the IPv6 extension header values, the process repeats until either a value of 17 (UDP) or a value of 59 (No Next Header) are found or the packet ends.

15.2.2.3 Transmit Message Egress Time Recording

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked for ALL of the following.

- The messageType field of the PTP header is checked and only those messages enabled via the [TX PTP Message Type Enable \(TX_PTP_MESSAGE_EN\[15:0\]\)](#) bits in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#) will be have their egress times saved. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled.
- The versionPTP field of the PTP header is checked against the [TX PTP Version \(TX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will be have their egress times saved. A setting of 0 allows any PTP version.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

- If enabled via the [TX PTP Domain Match Enable \(TX_PTP_DOMAIN_EN\)](#) bit in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#), the domainNumber field of the PTP header is checked against the [TX PTP Domain \(TX_PTP_DOMAIN\[7:0\]\)](#) value in the same register. Only those messages with a matching domain will be have their egress times saved.
- If enabled via the [TX PTP Alternate Master Enable \(TX_PTP_ALT_MASTER_EN\)](#) bit in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#), the alternateMasterFlag in the flagField of the PTP header is checked and only those messages with an alternateMasterFlag set to 0 will be have their egress times saved.

At the end of the frame, the frame's FCS and the UDP checksum (for IPv4 and IPv6 formats) are verified. FCS checking can be disabled using the [TX PTP FCS Check Disable \(TX_PTP_FCS_DIS\)](#) bit in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#). UDP checksum checking can be disabled using the [TX PTP UDP Checksum Check Disable \(TX_PTP_UDP_CHKSUM_DIS\)](#) bit in the same register.

Note: A IPv4 UDP checksum value of 0x0000 indicates that the checksum is not included and is considered a pass. A IPv6 UDP checksum value of 0x0000 is invalid and is considered a fail.

Note: For IPv6, the UDP checksum calculation includes the IPv6 Pseudo header. Part of the IPv6 Pseudo header is the final IPv6 destination address.

If the IPv6 packet does not contain a Routing header, then the final IPv6 destination address is the destination address contained in the IPv6 header.

If the IPv6 packet does contain a Routing header, then the final IPv6 destination address is the address in the last element of the Routing header.

Note: The UDP checksum is calculated over the entire UDP payload as indicated by the UDP length field and not the assumed PTP packet length.

Note: The UDP checksum calculation does not include layer 2 pad bytes, if any.

If the FCS and checksum tests pass:

- The latency adjusted, 1588 Clock value, saved above at the start of the frame, is recorded into the [1588 Port x TX Egress Time Seconds Register \(1588_TX_EGRESS_SEC_x\)](#) and [1588 Port x TX Egress Time NanoSeconds Register \(1588_TX_EGRESS_NS_x\)](#).
- The messageType and sequenceId fields and 12-bit CRC of the portIdentity field of the PTP header are recorded into the [Message Type \(MSG_TYPE\)](#), [Sequence ID \(SEQ_ID\)](#) and [Source Port Identity CRC \(SRC_PRT_CRC\)](#) fields of the [1588 Port x TX Message Header Register \(1588_TX_MSG_HEADER_x\)](#).

The 12-bit CRC of the portIdentity field is created by using the polynomial of $X^{12} + X^{11} + X^3 + X^2 + X + 1$.

- The corresponding maskable [1588 TX Timestamp Interrupt \(1588_TX_TS_INT\[2:0\]\)](#) is set in the [1588 Interrupt Status Register \(1588_INT_STS\)](#).

Up to four transmit events are saved per port with the count shown in the [1588 TX Timestamp Count \(1588_TX_TS_CNT\[2:0\]\)](#) field in the [1588 Port x Capture Information Register \(1588_CAP_INFO_x\)](#). Additional events are not recorded. When the appropriate [1588 TX Timestamp Interrupt \(1588_TX_TS_INT\[2:0\]\)](#) bit is written as a one to clear, [1588 TX Timestamp Count \(1588_TX_TS_CNT\[2:0\]\)](#) will decrement. If there are remaining events, the capture registers will update to the next event and the interrupt will set again.

TIME STAMPS FROM FORWARDED PACKETS

The transmitter will also save egress times for frames that are forwarded from another port. Typically, these are of no use to the Host S/W and would need to be discarded. Since these messages also typically have their correction field adjusted for residence time, they can be distinguished from messages from the Host.

If [EGRESS CORRECTION FIELD RESIDENCE TIME ADJUSTMENT](#), below, is performed on a message, egress times are not saved if the [TX PTP Suppress Timestamps when Correction Field Adjusted \(TX_PTP_SUPP_CF_TS\)](#) bit in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#) is set.

DELAY_REQ EGRESS TIME SAVING

Normally, in ordinary clock operation, the egress time of transmitted Delay_Req packets are saved and read by the Host S/W. To avoid the need to read these timestamps via register access, the egress time of the last transmitted Delay_Req packet on the port can be inserted into Delay_Resp packets received on the port.

The [1588 Port x TX Delay_Req Egress Time Seconds Register \(1588_TX_DREQ_SEC_x\)](#) and the [1588 Port x TX Delay_Req Egress Time NanoSeconds Register \(1588_TX_DREQ_NS_x\)](#) hold the egress time of the Delay_Req message.

These registers are updated by the H/W when the Delay_Req message is transmitted independent of the settings in the [TX PTP Message Type Enable \(TX_PTP_MESSAGE_EN\[15:0\]\)](#) bits.

As above (including all applicable notes):

- The versionPTP and domainNumber fields and alternateMasterFlag in the flagField of the PTP header are checked, if enabled.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

- At the end of the frame, the frame's FCS and the UDP checksum (for IPv4 and IPv6 formats) are verified, if enabled.

If all tests pass, then the Delay_Req message information is updated and available for the receive function.

15.2.2.4 Egress Packet Modifications

Modifications to frames on egress are divided into two categories, those to support one-step transparent clock residence time corrections and those to support one-step operations from the Host software.

Bit 7 of the PTP header's reserved byte (the byte which is also used to hold the ingress time seconds) is used to indicate packets that need to have their correction field adjusted for residence time. This bit is set on ingress when the correction field adjustment process is started.

When bit 7 of the PTP header's reserved byte is cleared, the alternate function, if any, for the message type is to be performed, if it is enabled. The Host S/W should normally have bit 7 cleared.

Note: The offset of the reserved byte is specified by the [TX PTP 1 Reserved Byte Offset \(TX_PTP_1_RSVD_OFFSET\[5:0\]\)](#) field in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#).

Proper operation of the transmitter requires that the reserved byte resides after the versionPTP field and before the correctionField.

For version 2 of IEEE 1588, the reserved byte at offset 5 should be used.

EGRESS CORRECTION FIELD RESIDENCE TIME ADJUSTMENT

In order to support one-step transparent clock operation, the residence time delay through the device is accounted for by adjusting the correctionField of certain packets.

This function is enabled per PTP message type via the [TX PTP Correction Field Message Type Enable \(TX_PT_P_CF_MSG_EN\[15:0\]\)](#) bits in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#).

Typically the Sync message is enabled for both end-to-end and peer-to-peer transparent clocks, the Delay_Req, PDelay_Req and PDelay_Resp messages are enabled only for end-to-end transparent clocks.

As described above, messages from the Host S/W would normally have bit 7 of the PTP header's reserved byte clear and are not modified in this manner. Typically, bit 7 is only set on ingress when the correction field adjustment process is started.

Note: The Host S/W should normally keep bit 7 of the PTP header's reserved byte clear for Sync, Delay_Req, PDelay_Req and PDelay_Resp messages so that residence time adjustment is not performed.

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked.

- The versionPTP field of the PTP header is checked against the [TX PTP Version \(TX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will have their correction field modified. A setting of 0 allows any PTP version.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

Note: The domainNumber field and alternateMasterFlag in the flagField of the PTP header are not tested for purpose of correction field modification.

The correctionField is modified as follows:

Note: If the original correctionField contains a value of 0x7FFFFFFFFFFFFFFF, it is not modified.

If adjustment to the correctionField would result in a value that is larger than 0x7FFFFFFFFFFFFFFF, that value is used instead.

- For **Delay_Req** and **Pdelay_Req** packets, the value of the [Port Delay Asymmetry \(DELAY_ASYM\[15:0\]\)](#) field in the [1588 Port x Asymmetry and Peer Delay Register \(1588_ASYM_PEERDLY_x\)](#) (for the particular egress port) is subtracted from the correctionField.

This function is used for one-step end-to-end transparent clocks. If one-step end-to-end transparent clock mode is not being used, correction field modifications would not be enabled for Delay_Req and PDelay_Req messages.

- The nanoseconds portion of the egress time are added to the correctionField.
- In order to detect and correct for a potential rollover of the nanoseconds portion the clock, egress seconds bits 3:0 minus ingress seconds bits 3:0 (without borrow) is added to the correctionField.

The ingress time is available in bits 3:0 of a reserved byte in the PTP header.

Note: The offset of the reserved byte is specified by the [TX PTP 1 Reserved Byte Offset \(TX_PTP_1_RSVD_OFFSET\[5:0\]\)](#) field in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#).

Proper operation of the transmitter requires that the reserved byte resides after the versionPTP field and before the correctionField.

For version 2 of IEEE 1588, the reserved byte at offset 5 should be used.

EGRESS TIME INSERTION - SYNC MESSAGE ALTERNATE FUNCTION

While functioning as an ordinary clock master, one-step transmission of Sync messages from the Host S/W requires the actual egress time to be inserted into the ten byte, originTimestamp field. The 32-bit nanoseconds portion and the lower 32 bits of the seconds portion come from the latency adjusted, 1588 Clock value, saved above at the start of the frame. The upper 16 bits of seconds are taken from the [1588 TX One-Step Sync Upper Seconds Register \(1588_TX_ONE_STEP_SYNC_SEC\)](#). The Host software is responsible for maintaining this register if required.

Note: Inserting the egress time into the packet is an additional, separately enabled, feature verses the Egress Time Recording described above.

This function is enabled via the [TX PTP Sync Message Egress Time Insertion \(TX_PTP_SYNC_TS_INSERT\)](#) bit in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#) and is used only on frames which have bit 7 of the PTP header's reserved byte cleared.

Note: The offset of the reserved byte is specified by the [TX PTP 1 Reserved Byte Offset \(TX_PTP_1_RSVD_OFFSET\[5:0\]\)](#) field in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#).

Proper operation of the transmitter requires that the reserved byte resides after the versionPTP field and before the correctionField.

For version 2 of IEEE 1588, the reserved byte at offset 5 should be used.

Following the determination of packet format and qualification of the packet as a PTP message above, the PTP header is checked.

- The versionPTP field of the PTP header is checked against the [TX PTP Version \(TX_PTP_VERSION\[3:0\]\)](#) field in the [1588 Port x TX Timestamp Configuration Register \(1588_TX_TIMESTAMP_CONFIG_x\)](#). Only those messages with a matching version will have their egress time inserted. A setting of 0 allows any PTP version.

Note: The domainNumber field and alternateMasterFlag in the flagField of the PTP header are not tested.

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

EGRESS CORRECTION FIELD TURNAROUND TIME ADJUSTMENT - PDELAY_RESP MESSAGE ALTERNATE FUNCTION

One-step Pdelay_Resp messages sent by the Host, require their correctionField to be calculated on-the-fly to include the turnaround time between the ingress of the Pdelay_Req and the egress time of the Pdelay_Resp.

$Pdelay_Resp.CF = Pdelay_Req.CF + Pdelay_Resp.egress\ time - Pdelay_Req.ingress\ time.$

Note: Adjusting the Correction Field in the packet is an additional, separately enabled, feature verses the Egress Time Recording described above.

Note: If the original correctionField contains a value of 7FFFFFFFFFFFFFFF, it is not modified.

If adjustment to the correctionField would result in a value that is larger than 7FFFFFFFFFFFFFFF, that value is used instead.

The [1588 Port x RX Pdelay_Req Ingress Time Seconds Register \(1588_RX_PDREQ_SEC_x\)](#) and the [1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register \(1588_RX_PDREQ_NS_x\)](#) hold the ingress time of the Pdelay_Req message.

The [1588 Port x RX Pdelay_Req Ingress Correction Field High Register \(1588_RX_PDREQ_CF_HI_x\)](#) and the [1588 Port x RX Pdelay_Req Ingress Correction Field Low Register \(1588_RX_PDREQ_CF_LOW_x\)](#) hold the correctionField of the Pdelay_Req message.

These registers are set by S/W prior to sending the Pdelay_Resp message or by the automatic updating described above in [PDELAY_REQ INGRESS TIME SAVING](#).

LAN9352

The egress time is the latency adjusted, 1588 Clock value, saved above at the start of the Pdelay_Resp frame.

Note: Since only four bits worth of seconds of the Pdelay_Req ingress time are stored, the Host must send the Pdelay_Resp within 16 seconds.

This function is enabled via the [TX PTP Pdelay_Resp Message Turnaround Time Insertion \(TX_PTP_PDRE-SP_TA_INSERT\)](#) bit in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#) and is used only on frames which have bit 7 of the PTP header's reserved byte cleared.

As with Egress Time Insertion above:

- The versionPTP field of the PTP header is checked and the domainNumber field and alternateMasterFlag in the flagField of the PTP header are not checked

Note: Support for the IEEE 1588-2002 (v1) packet format is not provided.

CLEARING RESERVED FIELDS

If the frame is modified on egress for Correction Field Residence Time Adjustment:

- The reserved byte at the location specified by the [TX PTP 1 Reserved Byte Offset \(TX_PTP_1_RSVD_OFFSET\[5:0\]\)](#) is cleared.
- The four reserved bytes used for [INGRESS TIME INSERTION INTO PACKETS](#) are cleared if the [TX PTP Clear Four Byte Reserved Field \(TX_PTP_CLR_4_RSVD\)](#) bits in the [1588 Port x TX Modification Register \(1588_TX_MOD_x\)](#) is set.

Note: The offset of the four reserved bytes is specified in [TX PTP 4 Reserved Bytes Offset \(TX_PTP_4_RSVD_OFFSET\[5:0\]\)](#).

FRAME UPDATING

Frames are modified even if their original FCS or UDP checksum is invalid.

- For IPv4, the UDP checksum is set to 0 under the following conditions.

If the [TX PTP Clear UDP/IPv4 Checksum Enable \(TX_PTP_CLR_UDPV4_CHKSUM\)](#) bit in the [1588 Port x TX Modification Register 2 \(1588_TX_MOD2_x\)](#) is set, the UDP checksum is set to 0 for Sync messages if Sync Egress Time Insertion is enabled and for Pdelay_Resp messages if Pdelay_Resp Correction Field Turnaround Time Adjustment is enabled. The ptp_version field is also checked.

- When Residence Time Correction is performed, the UDP checksum is already set to 0 by the ingress port.
- For IPv6, the two bytes beyond the end of the PTP message are modified to correct for the UDP checksum. These bytes are updated by accumulating the differences between the original frame data and the substituted data using the mechanism defined in IETF RFC 1624.

The existing two bytes are included in the calculation and are updated.

It is assumed that the original UDP checksum is valid and is not checked.

Note: Since the two bytes beyond the end of the PTP message are modified based on the differences between the original frame data and the substituted data, an invalid incoming checksum would result in an outgoing checksum error.

Note: The two bytes beyond the end of the PTP message are located by using the messageLength field from the PTP header.

- The frame FCS is recomputed
It is assumed that the original FCS is valid and is not checked.
- If the frame has a transmit symbol error(s), a transmit symbol error indication will be propagated at the same nibble location(s)

Note: The FCS and IPv6/UDP checksum are updated and the reserved byte cleared only if the frame was actually modified.


The IPv4/UDP checksum is cleared as indicated above and could be the only modification in the message. If the IPv4/UDP checksum is cleared, the FCS is recomputed.

If no modifications are done, the existing FCS, checksums and reserved bytes are left unchanged.

15.3 1588 Clock

The tunable 1588 Clock is the time source for all PTP related functions of the device. The block diagram is shown in Figure 15-1.

FIGURE 15-1: 1588 CLOCK BLOCK DIAGRAM



The 1588 Clock consists of a 32-bit wide seconds portion and a 30-bit wide nanoseconds portion. Running at a nominal reference frequency of 100MHz, the nanoseconds portion is normally incremented by a value of 10 every reference clock period. Upon reaching or exceeding its maximum value of 10^9 , the nanoseconds portion rolls over to or past zero and the seconds portion is incremented.

The 1588 Clock can be read by setting the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#). This saves the current value of the 1588 Clock into the [1588 Clock Seconds Register \(1588_CLOCK_SEC\)](#), [1588 Clock NanoSeconds Register \(1588_CLOCK_NS\)](#) and [1588 Clock Sub-NanoSeconds Register \(1588_CLOCK_SUBNS\)](#) where it can be read.

Although the IEEE 1588-2008 specification calls for a 48-bit seconds counter, the hardware only supports 32 bits. For purposes of event timestamping, residence time correction or other comparisons, the 136 year rollover time of 32 bits is sufficient. Rollover can be detected and corrected by comparing the two values of interest. To support one-step operations, the device can insert the Egress Timestamp into the origin Timestamp field of Sync messages. However, the Host must maintain the [1588 TX One-Step Sync Upper Seconds Register \(1588_TX_ONE_STEP_SYNC_SEC\)](#). The Host should avoid sending a Sync message if there is a possibility that the 32-bit seconds counter will reach its rollover value before the message is transmitted.

A 32-bit sub-nanoseconds counter is used to precisely tune the rate of the 1588 Clock by accounting for the difference between the nominal 10ns and the actual rate of the master clock. Every reference clock period the sub-nanoseconds counter is incremented by the [Clock Rate Adjustment Value \(1588_CLOCK_RATE_ADJ_VALUE\)](#) in the [1588 Clock Rate Adjustment Register \(1588_CLOCK_RATE_ADJ\)](#), specified in 2^{-32} nanoseconds. When the sub-nanoseconds counter rolls over past zero, the nanoseconds portion of the 1588 Clock is incremented by 9 or 11 instead of the normal value of 10. The choice to speed up or slow down is determined by the [Clock Rate Adjustment Direction \(1588_CLOCK_RATE_ADJ_DIR\)](#) bit. The ability to adjust for 1 ns approximately every 43 seconds allows for a tuning precision of approximately 2.3^{-9} percent. The maximum adjustment is 1 ns every 4 clocks (40 ns) or 2.5 percent.

LAN9352

In addition to adjusting the frequency of the 1588 Clock, the Host may directly set the 1588 Clock, make a one-time step adjustment of the 1588 Clock or specify a temporary rate. The choice of method depends on needed adjustment. For initial adjustments, direct or one-time step adjustments may be best. For on-going minor adjustments, the temporary rate adjustment may be best. Ideally, the frequency will be matched and once the 1588 Clock is synchronized, no further adjustments would be needed.

In order to perform a direct writing of the 1588 Clock, the desired value is written into the [1588 Clock Seconds Register \(1588_CLOCK_SEC\)](#), [1588 Clock NanoSeconds Register \(1588_CLOCK_NS\)](#) and [1588 Clock Sub-NanoSeconds Register \(1588_CLOCK_SUBNS\)](#). The [Clock Load \(1588_CLOCK_LOAD\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is then set.

In order to perform a one-time positive or negative adjustment to the seconds portion of the 1588 Clock, the desired change and direction are written into the [1588 Clock Step Adjustment Register \(1588_CLOCK_STEP_ADJ\)](#). The [Clock Step Seconds \(1588_CLOCK_STEP_SECONDS\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is then set. The internal sub-nanoseconds counter and the nanoseconds portion of the 1588 Clock are not affected. If a nanoseconds portion rollover coincides with the 1588 Clock adjustment, the 1588 Clock adjustment is applied in addition to the seconds increment.

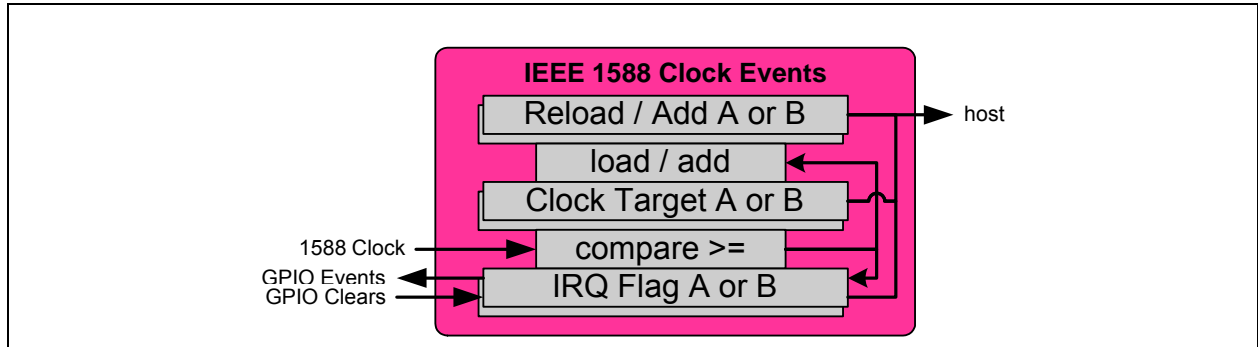
In order to perform a one-time positive adjustment to the nanoseconds portion of the 1588 Clock, the desired change is written into the [1588 Clock Step Adjustment Register \(1588_CLOCK_STEP_ADJ\)](#). The [Clock Step NanoSeconds \(1588_CLOCK_STEP_NANOSECONDS\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is then set. If the addition to the nanoseconds portion results in a rollover past zero, then the seconds portion of the 1588 Clock is incremented. The normal (9, 10 or 11 ns) increment to the nanoseconds portion is suppressed for one clock. This can be compensated for by specifying an addition value 10ns higher. A side benefit is that using an addition value of 0 effectively pauses the 1588 Clock for 10ns while a value less than 10 slows the clock down just briefly. The internal sub-nanoseconds counter of the 1588 Clock is not affected by the adjustment, however, if a sub-nanoseconds counter rollover coincides with the 1588 Clock adjustment it will be missed.

In order to perform a temporary rate adjustment of the 1588 Clock, the desired temporary rate and direction are written into the [1588 Clock Temporary Rate Adjustment Register \(1588_CLOCK_TEMP_RATE_ADJ\)](#) and the duration of the temporary rate, specified in reference clock cycles, is written into the [1588 Clock Temporary Rate Duration Register \(1588_CLOCK_TEMP_RATE_DURATION\)](#). The [Clock Temporary Rate \(1588_CLOCK_TEMP_RATE\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is then set. Once the temporary rate duration expires, the [Clock Temporary Rate \(1588_CLOCK_TEMP_RATE\)](#) bit will self-clear and the [1588 Clock Rate Adjustment Register \(1588_CLOCK_RATE_ADJ\)](#) will once again control the 1588 Clock rate. This method of adjusting the 1588 Clock may be preferred since it avoids large discrete changes in the 1588 Clock value. For a maximum setting in both the [1588 Clock Temporary Rate Adjustment Register \(1588_CLOCK_TEMP_RATE_ADJ\)](#) and [1588 Clock Temporary Rate Duration Register \(1588_CLOCK_TEMP_RATE_DURATION\)](#), the 1588 Clock can be adjusted by approximately 1 second.

15.4 1588 Clock Events

The 1588 Clock Events block is responsible for generating and controlling all 1588 Clock related events. Two clock event channels, A and B, are available. The block diagram is shown in Figure 15-2.

FIGURE 15-2: 1588 CLOCK EVENT BLOCK DIAGRAM



For each clock event channel, a comparator compares the 1588 Clock with a Clock Target loaded in the [1588 Clock Target x Seconds Register \(1588_CLOCK_TARGET_SEC_x\)](#) and [1588 Clock Target x NanoSeconds Register \(1588_CLOCK_TARGET_NS_x\)](#).

The Clock Target Register pair requires two 32-bit write cycles, one to each half, before the register pair is affected. The writes may be in any order. There is a register pair for each clock event channel (A and B).

The Clock Target can be read by setting the [Clock Target Read \(1588_CLOCK_TARGET_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#). This saves the current value of the both Clock Targets (A and B) into the [1588 Clock Target x Seconds Register \(1588_CLOCK_TARGET_SEC_x\)](#) and [1588 Clock Target x NanoSeconds Register \(1588_CLOCK_TARGET_NS_x\)](#) where they can be read.

When the 1588 Clock reaches or passes the Clock Target for a clock event channel, a clock event occurs which triggers the following:

- The maskable interrupt for that clock event channel ([1588 Timer Interrupt A \(1588_TIMER_INT_A\)](#) or [1588 Timer Interrupt B \(1588_TIMER_INT_B\)](#)) is set in the [1588 Interrupt Status Register \(1588_INT_STS\)](#).
- The [Reload/Add A \(RELOAD_ADD_A\)](#) or [Reload/Add B \(RELOAD_ADD_B\)](#) bit in the [1588 General Configuration Register \(1588_GENERAL_CONFIG\)](#) is checked to determine the new Clock Target behavior:

–RELOAD_ADD = 1:

The new Clock Target is loaded from the Reload / Add Registers ([1588 Clock Target x Reload / Add Seconds Register \(1588_CLOCK_TARGET_RELOAD_SEC_x\)](#) and [1588 Clock Target x Reload / Add NanoSeconds Register \(1588_CLOCK_TARGET_RELOAD_NS_x\)](#)).

–RELOAD_ADD = 0:

The Clock Target is incremented by the Reload / Add Registers ([1588 Clock Target x Reload / Add Seconds Register \(1588_CLOCK_TARGET_RELOAD_SEC_x\)](#) and [1588 Clock Target x Reload / Add NanoSeconds Register \(1588_CLOCK_TARGET_RELOAD_NS_x\)](#)). The Clock Target NanoSeconds rolls over at 10^9 and the carry is added to the Clock Target Seconds.

The Clock Target Reload / Add Register pair requires two 32-bit write cycles, one to each half, before the register pair is affected. The writes may be in any order. There is a register pair for each clock event channel (A and B).

Note: Writing the 1588 Clock may cause the interrupt event to occur if the new 1588 Clock value is set equal to or greater than the current Clock Target.

The Clock Target reload function (RELOAD_ADD = 1) allows the Host to pre-load the next trigger time in advance. The add function (RELOAD_ADD = 0), allows for a automatic repeatable event.

15.5 1588 GPIOs

In addition to time stamping PTP packets, the 1588 Clock value can be saved into a set of clock capture registers based on the GPIO inputs. The GPIO inputs can also be used to clear the 1588 Clock Target compare event interrupt. When configured as outputs, GPIOs can be used to output a signal based on an 1588 Clock Target compare events.

Note: The IEEE 1588 Unit supports up to 8 GPIO signals.

15.5.1 1588 GPIO INPUTS

15.5.1.1 GPIO Event Clock Capture

When the GPIO pins are configured as inputs, and enabled with the [GPIO Rising Edge Capture Enable 7-0 \(GPIO_RE_CAPTURE_ENABLE\[7:0\]\)](#) or [GPIO Falling Edge Capture Enable 7-0 \(GPIO_FE_CAPTURE_ENABLE\[7:0\]\)](#) bits in the [1588 GPIO Capture Configuration Register \(1588_GPIO_CAP_CONFIG\)](#), a rising or falling edge, respectively, will capture the 1588 Clock into the [1588 GPIO x Rising Edge Clock Seconds Capture Register \(1588_GPIO_RE_CLOCK_SEC_CAP_x\)](#) and the [1588 GPIO x Rising Edge Clock NanoSeconds Capture Register \(1588_GPIO_RE_CLOCK_NS_CAP_x\)](#) or [1588 GPIO x Falling Edge Clock Seconds Capture Register \(1588_GPIO_FE_CLOCK_SEC_CAP_x\)](#) and the [1588 GPIO x Falling Edge Clock NanoSeconds Capture Register \(1588_GPIO_FE_CLOCK_NS_CAP_x\)](#) where x equals the number of the active GPIO input.

GPIO inputs must be stable for greater than 40 ns to be recognized as capture events and are edge sensitive.

The GPIO inputs have a fixed capture latency of 65 ns that can be accounted for by the Host driver. The GPIO inputs have a capture latency uncertainty of +/-5 ns.

The corresponding, maskable, interrupt flags [1588 GPIO Rising Edge Interrupt \(1588_GPIO_RE_INT\[7:0\]\)](#) or [1588 GPIO Falling Edge Interrupt \(1588_GPIO_FE_INT\[7:0\]\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) will also be set. This is in addition to the interrupts available in the [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#).

A lock enable bit is provided for each timestamp enabled GPIO, [Lock Enable GPIO Rising Edge \(LOCK_GPIO_RE\)](#) and [Lock Enable GPIO Falling Edge \(LOCK_GPIO_FE\)](#) in the [1588 GPIO Capture Configuration Register \(1588_GPIO_CAP_CONFIG\)](#), which prevents the corresponding GPIO clock capture registers from being overwritten if the GPIO interrupt in [1588 Interrupt Status Register \(1588_INT_STS\)](#) is already set.

15.5.1.2 GPIO Timer Interrupt Clear

The GPIO inputs can also be configured to clear the [1588 Timer Interrupt A \(1588_TIMER_INT_A\)](#) or [1588 Timer Interrupt B \(1588_TIMER_INT_B\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) by setting the corresponding enable and select bits in the [1588 General Configuration Register \(1588_GENERAL_CONFIG\)](#).

The polarity of the GPIO input is determined by the [GPIO Interrupt/1588 Polarity 7-0 \(GPIO_POL\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#).

GPIO inputs must be active for greater than 40 ns to be recognized as interrupt clear events and are edge sensitive.

15.5.2 1588 GPIO OUTPUTS

Upon detection of a Clock Target A or B compare event, the corresponding clock event channel can be configured to output a 100 ns pulse, toggle its output, or reflect its 1588 Timer Interrupt bit. The selection is made using the [Clock Event Channel A Mode \(CLOCK_EVENT_A\)](#) and [Clock Event Channel B Mode \(CLOCK_EVENT_B\)](#) bits of the [1588 General Configuration Register \(1588_GENERAL_CONFIG\)](#).

A GPIO pin is configured as a 1588 event output by setting the corresponding [1588 GPIO Output Enable 7-0 \(1588_GPIO_OE\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#). These bits override the GPIO Direction bits of the [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#) and allow for GPIO output generation based on the 1588 Clock Target compare event. The choice of the event channel is controlled by the [1588 GPIO Channel Select 7-0 \(GPIO_CH_SEL\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#).

Note: The [1588 GPIO Output Enable 7-0 \(1588_GPIO_OE\[7:0\]\)](#) bits do not override the [GPIO Buffer Type 7-0 \(GPIOBUF\[7:0\]\)](#) in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#).

The clock event polarity, which determines whether the 1588 GPIO output is active high or active low, is controlled by the [GPIO Interrupt/1588 Polarity 7-0 \(GPIO_POL\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#).

The GPIO outputs have a latency of approximately 40 ns when using “100 ns pulse” or “Interrupt bit” modes and 30 ns when using “toggle” mode. On chip delays contribute an uncertainty of +/-4ns to these values.

15.6 Software Triggered Clock Capture

As an alternative, the GPIO Capture registers can be used by Host software to record software events by specifying the GPIO register set in the [1588 Manual Capture Select 3-0 \(1588_MANUAL_CAPTURE_SEL\[3:0\]\)](#) and setting the [1588 Manual Capture \(1588_MANUAL_CAPTURE\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#).

This also causes the corresponding bit in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) to set.

Note: The interrupts available in the [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#) are not set by the using this method.

Note: The [Lock Enable GPIO Rising Edge \(LOCK_GPIO_RE\)](#) and [Lock Enable GPIO Falling Edge \(LOCK_GPIO_FE\)](#) bits do not apply to manual clock capture.

The full set of GPIO Capture registers is always available regardless of the number of GPIOs supported by the device.

15.7 1588 Interrupt

The IEEE 1588 unit provides multiple interrupt conditions. These include timestamp indication on the transmitter and receiver side of each port, individual GPIO input timestamp interrupts, and a clock comparison event interrupts. All 1588 interrupts are located in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) and are fully maskable via their respective enable bits in the [1588 Interrupt Enable Register \(1588_INT_EN\)](#).

All 1588 interrupts are ANDed with their individual enables and then ORed, generating the [1588 Interrupt Event \(1588_EVNT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#).

When configured as inputs, the GPIOs have the added functionality of clearing the [1588 Timer Interrupt A \(1588_TIMER_INT_A\)](#) or [1588 Timer Interrupt B \(1588_TIMER_INT_B\)](#) bits of the [1588 Interrupt Status Register \(1588_INT_STS\)](#) as described in [Section 15.5.1.2](#).

Refer to [Section 8.0, "System Interrupts," on page 73](#) for additional information on the device interrupts.

15.8 1588 Registers

This section details the directly addressable PTP timestamp related registers.

Each port has a PTP timestamp block with related registers. These sets of registers are identical in functionality for each port, and thus their register descriptions have been consolidated. In these cases, the register names will be amended with a lowercase “x” in place of the port designation. The wildcard “x” should be replaced with “0”, “1” or “2” respectively.

For GPIO related registers, the wildcard “x” should be replaced with “0” through “7”.

Similarly, for Clock Compare events, the wildcard “x” should be replaced with “A” or “B”.

Port and GPIO registers share a common address space. Port vs. GPIO registers are selected by using the [Bank Select](#) (BANK_SEL[2:0]) in the [1588 Bank Port GPIO Select Register](#) (1588_BANK_PORT_GPIO_SEL). The port accessed (“x”) is set by the [Port Select](#) (PORT_SEL[1:0]) field. The GPIO accessed (“x”) is set by the [GPIO Select](#) (GPIO_SEL[2:0]) field.

Note: The IEEE 1588 Unit supports 8 GPIO signals.

For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map,"](#) on page 33.

TABLE 15-1: 1588 CONTROL AND STATUS REGISTERS

BANK SELECT	ADDRESS OFFSET	Register Name (Symbol)
na	100h	1588 Command and Control Register (1588_CMD_CTL)
na	104h	1588 General Configuration Register (1588_GENERAL_CONFIG)
na	108h	1588 Interrupt Status Register (1588_INT_STS)
na	10Ch	1588 Interrupt Enable Register (1588_INT_EN)
na	110h	1588 Clock Seconds Register (1588_CLOCK_SEC)
na	114h	1588 Clock NanoSeconds Register (1588_CLOCK_NS)
na	118h	1588 Clock Sub-NanoSeconds Register (1588_CLOCK_SUBNS)
na	11Ch	1588 Clock Rate Adjustment Register (1588_CLOCK_RATE_ADJ)
na	120h	1588 Clock Temporary Rate Adjustment Register (1588_CLOCK_TEMP_RATE_ADJ)
na	124h	1588 Clock Temporary Rate Duration Register (1588_CLOCK_TEMP_RATE_DURATION)
na	128h	1588 Clock Step Adjustment Register (1588_CLOCK_STEP_ADJ)
na	12Ch	1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) x=A
na	130h	1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) x=A
na	134h	1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) x=A
na	138h	1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=A
na	13Ch	1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) x=B
na	140h	1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) x=B
na	144h	1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) x=B

TABLE 15-1: 1588 CONTROL AND STATUS REGISTERS (CONTINUED)

BANK SELECT	ADDRESS OFFSET	Register Name (Symbol)
na	148h	1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=B
na	14Ch	1588 User MAC Address High-WORD Register (1588_USER_MAC_HI)
na	150h	1588 User MAC Address Low-DWORD Register (1588_USER_MAC_LO)
na	154h	1588 Bank Port GPIO Select Register (1588_BANK_PORT_GPIO_SEL)
0	158h	1588 Port x Latency Register (1588_LATENCY_x)
0	15Ch	1588 Port x Asymmetry and Peer Delay Register (1588_ASYM_PEERDLY_x)
0	160h	1588 Port x Capture Information Register (1588_CAP_INFO_x)
1	158h	1588 Port x RX Parsing Configuration Register (1588_RX_PARSE_CONFIG_x)
1	15Ch	1588 Port x RX Timestamp Configuration Register (1588_RX_TIMESTAMP_CONFIG_x)
1	160h	1588 Port x RX Timestamp Insertion Configuration Register (1588_RX_TS_INSERT_CONFIG_x)
1	164h	1588 Port x RX Correction Field Modification Register (1588_RX_CF_MOD_x)
1	168h	1588 Port x RX Filter Configuration Register (1588_RX_FILTER_CONFIG_x)
1	16Ch	1588 Port x RX Ingress Time Seconds Register (1588_RX_INGRESS_SEC_x)
1	170h	1588 Port x RX Ingress Time NanoSeconds Register (1588_RX_INGRESS_NS_x)
1	174h	1588 Port x RX Message Header Register (1588_RX_MSG_HEADER_x)
1	178h	1588 Port x RX Pdelay_Req Ingress Time Seconds Register (1588_RX_PDREQ_SEC_x)
1	17Ch	1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register (1588_RX_PDREQ_NS_x)
1	180h	1588 Port x RX Pdelay_Req Ingress Correction Field High Register (1588_RX_PDREQ_CF_HI_x)
1	184h	1588 Port x RX Pdelay_Req Ingress Correction Field Low Register (1588_RX_PDREQ_CF_LOW_x)
1	188h	1588 Port x RX Checksum Dropped Count Register (1588_RX_CHKSUM_DROPPED_CNT_x)
1	18Ch	1588 Port x RX Filtered Count Register (1588_RX_FILTERED_CNT_x)
2	158h	1588 Port x TX Parsing Configuration Register (1588_TX_PARSE_CONFIG_x)
2	15Ch	1588 Port x TX Timestamp Configuration Register (1588_TX_TIMESTAMP_CONFIG_x)
2	164h	1588 Port x TX Modification Register (1588_TX_MOD_x)
2	168h	1588 Port x TX Modification Register 2 (1588_TX_MOD2_x)
2	16Ch	1588 Port x TX Egress Time Seconds Register (1588_TX_EGRESS_SEC_x)

TABLE 15-1: 1588 CONTROL AND STATUS REGISTERS (CONTINUED)

BANK SELECT	ADDRESS OFFSET	Register Name (Symbol)
2	170h	1588 Port x TX Egress Time NanoSeconds Register (1588_TX_EGRESS_NS_x)
2	174h	1588 Port x TX Message Header Register (1588_TX_MSG_HEADER_x)
2	178h	1588 Port x TX Delay_Req Egress Time Seconds Register (1588_TX-_DREQ_SEC_x)
2	17Ch	1588 Port x TX Delay_Req Egress Time NanoSeconds Register (1588_TX-_DREQ_NS_x)
2	180h	1588 TX One-Step Sync Upper Seconds Register (1588_TX_ONE_STEP_SYNC_SEC)
3	15Ch	1588 GPIO Capture Configuration Register (1588_GPIO_CAP_CONFIG)
3	16Ch	1588 GPIO x Rising Edge Clock Seconds Capture Register (1588_GPIO_RE_-CLOCK_SEC_CAP_x)
3	170h	1588 GPIO x Rising Edge Clock NanoSeconds Capture Register (1588_GPIO_RE_CLOCK_NS_CAP_x)
3	178h	1588 GPIO x Falling Edge Clock Seconds Capture Register (1588_GPIO_FE_-CLOCK_SEC_CAP_x)
3	17Ch	1588 GPIO x Falling Edge Clock NanoSeconds Capture Register (1588_GPIO_FE_CLOCK_NS_CAP_x)

15.8.1 1588 COMMAND AND CONTROL REGISTER (1588_CMD_CTL)

Offset: 100h Size: 32 bits
Bank: na

Bits	Description	Type	Default
31:14	RESERVED	RO	-
13	Clock Target Read (1588_CLOCK_TARGET_READ) Writing a one to this bit causes the current values of both of the 1588 clock targets (A and B) to be saved into the 1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) and the 1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) so they can be read. Writing a zero to this bit has no affect.	WO SC	0b
12:9	1588 Manual Capture Select 3-0 (1588_MANUAL_CAPTURE_SEL[3:0]) These bits specify which GPIO 1588 Clock Capture Registers are used during a manual capture. Bit 3 selects the rising edge (0) or falling edge (1) registers. Bits 2-0 select the GPIO number. Note: All 8 GPIO register sets are available.	R/W	0000b
8	1588 Manual Capture (1588_MANUAL_CAPTURE) Writing a one to this bit causes the current value of the 1588 clock to be saved into the GPIO 1588 Clock Capture Registers specified above. The corresponding bit in the 1588 Interrupt Status Register (1588_INT_STS) is also set. Writing a zero to this bit has no affect.	WO SC	0b
7	Clock Temporary Rate (1588_CLOCK_TEMP_RATE) Writing a one to this bit enables the use of the temporary clock rate adjustment specified in the 1588 Clock Temporary Rate Adjustment Register (1588_CLOCK_TEMP_RATE_ADJ) for the duration specified in the 1588 Clock Temporary Rate Duration Register (1588_CLOCK_TEMP_RATE_DURATION) . Writing a zero to this bit has no affect.	WO SC	0b
6	Clock Step NanoSeconds (1588_CLOCK_STEP_NANOSECONDS) Writing a one to this bit adds the value of the Clock Step Adjustment Value (1588_CLOCK_STEP_ADJ_VALUE) field in the 1588 Clock Step Adjustment Register (1588_CLOCK_STEP_ADJ) to the nanoseconds portion of the 1588 Clock. Writing a zero to this bit has no affect.	WO SC	0b

Bits	Description	Type	Default
5	Clock Step Seconds (1588_CLOCK_STEP_SECONDS) Writing a one to this bit adds or subtracts the value of the Clock Step Adjustment Value (1588_CLOCK_STEP_ADJ_VALUE) field in the 1588 Clock Step Adjustment Register (1588_CLOCK_STEP_ADJ) to or from the seconds portion of the 1588 Clock. The choice of adding or subtracting is set using the Clock Step Adjustment Direction (1588_CLOCK_STEP_ADJ_DIR) bit. Writing a zero to this bit has no affect.	WO SC	0b
4	Clock Load (1588_CLOCK_LOAD) Writing a one to this bit writes the value of the 1588 Clock Seconds Register (1588_CLOCK_SEC) , the 1588 Clock NanoSeconds Register (1588_CLOCK_NS) and the 1588 Clock Sub-NanoSeconds Register (1588_CLOCK_SUBNS) into the 1588 Clock. Writing a zero to this bit has no affect.	WO SC	0b
3	Clock Read (1588_CLOCK_READ) Writing a one to this bit causes the current value of the 1588 clock to be saved into the 1588 Clock Seconds Register (1588_CLOCK_SEC) , the 1588 Clock NanoSeconds Register (1588_CLOCK_NS) and the 1588 Clock Sub-NanoSeconds Register (1588_CLOCK_SUBNS) so it can be read. Writing a zero to this bit has no affect.	WO SC	0b
2	1588 Enable (1588_ENABLE) Writing a one to this bit will enable the 1588 unit. Reading this bit will return the current enabled value. Writing a zero to this bit has no affect. Note: Ports are individually enabled with the Time-Stamp Unit 2-0 Enable bits in the 1588 General Configuration Register (1588_GENERAL_CONFIG) .	R/W SC	Note 1:
1	1588 Disable (1588_DISABLE) Writing a one to this bit will cause the 1588 Enable (1588_ENABLE) to clear once all current frame processing is completed. No new frame processing will be started if this bit is set. Writing a zero to this bit has no affect.	WO SC	0b
0	1588 Reset (1588_RESET) Writing a one to this bit resets the 1588 H/W, state machines and registers and disables the 1588 unit. Any frame modifications in progress are halted at the risk of causing frame data or FCS errors. 1588_Reset should only be used once the 1588 unit is disabled as indicated by the 1588 Enable (1588_ENABLE) bit. Note: Writing a zero to this bit has no affect.	WO SC	0b

Note 1: The default value of this field is determined by the configuration strap [1588_enable_strap](#).

15.8.2 1588 GENERAL CONFIGURATION REGISTER (1588_GENERAL_CONFIG)

Offset: 104h Size: 32 bits
Bank: na

Bits	Description	Type	Default
31:19	RESERVED	RO	-
18	Time-Stamp Unit 2 Enable (TSU_ENABLE_2) This bit enables the receive and transmit functions of time-stamp unit 2. The 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) bit must also be set. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
17	Time-Stamp Unit 1 Enable (TSU_ENABLE_1) This bit enables the receive and transmit functions of time-stamp unit 1. The 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) bit must also be set. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
16	Time-Stamp Unit 0 Enable (TSU_ENABLE_0) This bit enables the receive and transmit functions of time-stamp unit 0. The 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) bit must also be set. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
15	GPIO 1588 Timer Interrupt B Clear Enable (GPIO_1588_TIMER_INT_B_CLEAR_EN) This bit enables the selected GPIO to clear the 1588_TIMER_INT_B bit of the 1588 Interrupt Status Register (1588_INT_STS) . The GPIO input is selected using the GPIO 1588 Timer Interrupt B Clear Select (GPIO_1588_TIMER_INT_B_CLEAR_SEL[2:0]) bits in this register. The polarity of the GPIO input is determined by GPIO Interrupt/1588 Polarity 7-0 (GPIO_POL[7:0]) in the General Purpose I/O Configuration Register (GPIO_CFG) . Note: The GPIO must be configured as an input for this function to operate. For the clear function, GPIO inputs are edge sensitive and must be active for greater than 40 ns to be recognized.	R/W	0b
14:12	GPIO 1588 Timer Interrupt B Clear Select (GPIO_1588_TIMER_INT_B_CLEAR_SEL[2:0]) These bits determine which GPIO is used to clear the 1588 Timer Interrupt B (1588_TIMER_INT_B) bit of the 1588 Interrupt Status Register (1588_INT_STS) . Note: The IEEE 1588 Unit supports 8 GPIO signals.	R/W	000b

Bits	Description	Type	Default
11	<p>GPIO 1588 Timer Interrupt A Clear Enable (GPIO_1588_TIMER_INT_A_CLEAR_EN) This bit enables the selected GPIO to clear the 1588 Timer Interrupt A (1588_TIMER_INT_A) bit of the 1588 Interrupt Status Register (1588_INT_STS).</p> <p>The GPIO input is selected using the GPIO 1588 Timer Interrupt A Clear Select (GPIO_1588_TIMER_INT_A_CLEAR_SEL[2:0]) bits in this register.</p> <p>The polarity of the GPIO input is determined by GPIO Interrupt/1588 Polarity 7-0 (GPIO_POL[7:0]) in the General Purpose I/O Configuration Register (GPIO_CFG).</p> <p>Note: The GPIO must be configured as an input for this function to operate. For the clear function, GPIO inputs are edge sensitive and must be active for greater than 40 ns to be recognized.</p>	R/W	0b
10:8	<p>GPIO 1588 Timer Interrupt A Clear Select (GPIO_1588_TIMER_INT_A_CLEAR_SEL[2:0]) These bits determine which GPIO is used to clear the 1588_TIMER_INT_A bit of the 1588 Interrupt Status Register (1588_INT_STS).</p> <p>Note: The IEEE 1588 Unit supports 8 GPIO signals.</p>	R/W	000b
7:6	RESERVED	RO	-
5:4	<p>Clock Event Channel B Mode (CLOCK_EVENT_B) These bits determine the output on Clock Event Channel B when a Clock Target compare event occurs.</p> <p>00: 100ns pulse output 01: Toggle output 10: 1588_TIMER_INT_B bit value in 1588_INT_STS_EN register output 11: RESERVED</p> <p>Note: The General Purpose I/O Configuration Register (GPIO_CFG) is used to enable the clock event onto the GPIO pins as well as to set the polarity and output buffer type.</p>	R/W	00b
3:2	<p>Clock Event Channel A Mode (CLOCK_EVENT_A) These bits determine the output on Clock Event Channel A when a Clock Target compare event occurs.</p> <p>00: 100ns pulse output 01: Toggle output 10: 1588_TIMER_INT_A bit value in 1588_INT_STS_EN register output 11: RESERVED</p> <p>Note: The General Purpose I/O Configuration Register (GPIO_CFG) is used to enable the clock event onto the GPIO pins as well as to set the polarity and output buffer type.</p>	R/W	00b

Bits	Description	Type	Default
1	<p>Reload/Add B (RELOAD_ADD_B) This bit determines the course of action when a Clock Target compare event for Clock Event Channel B occurs.</p> <p>When set, the 1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) and 1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) are loaded from the 1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) and 1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=B.</p> <p>When low, the Clock Target Registers are incremented by the Clock Target Reload Registers.</p> <p>0: Increment upon a clock target compare event 1: Reload upon a clock target compare event</p>	R/W	0b
0	<p>Reload/Add A (RELOAD_ADD_A) This bit determines the course of action when a Clock Target compare event for Clock Event Channel A occurs.</p> <p>When set, the 1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) and 1588 Clock Target x NanoSeconds Register (1588_CLOCK_TARGET_NS_x) are loaded from the 1588 Clock Target x Reload / Add Seconds Register (1588_CLOCK_TARGET_RELOAD_SEC_x) and 1588 Clock Target x Reload / Add NanoSeconds Register (1588_CLOCK_TARGET_RELOAD_NS_x) x=A.</p> <p>When low, the Clock Target Registers are incremented by the Clock Target Reload Registers.</p> <p>0: Increment upon a clock target compare event 1: Reload upon a clock target compare event</p>	R/W	0b

15.8.3 1588 INTERRUPT STATUS REGISTER (1588_INT_STS)

Offset: 108h Size: 32 bits
Bank: na

This read/write register contains the 1588 interrupt status bits.

Writing a 1 to a interrupt status bits acknowledges and clears the individual interrupt. If enabled in the [1588 Interrupt Enable Register \(1588_INT_EN\)](#), these interrupt bits are cascaded into the [1588 Interrupt Event \(1588_EVNT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#). Status bits will still reflect the status of the interrupt source regardless of whether the source is enabled as an interrupt. The [1588 Interrupt Event Enable \(1588_EVNT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must also be set in order for an actual system level interrupt to occur. Refer to [Section 8.0, "System Interrupts," on page 73](#) for additional information.

Bits	Description	Type	Default
31:24	1588 GPIO Falling Edge Interrupt (1588_GPIO_FE_INT[7:0]) This interrupt indicates that a falling event occurred and the 1588 Clock was captured. Note: As 1588 capture inputs, GPIO inputs are edge sensitive and must be low for greater than 40 ns to be recognized as interrupt inputs. These bits can also be set due to a manual capture via 1588 Manual Capture (1588_MANUAL_CAPTURE) .	R/WC	00h
23:16	1588 GPIO Rising Edge Interrupt (1588_GPIO_RE_INT[7:0]) This interrupt indicates that a rising event occurred and the 1588 Clock was captured. Note: As 1588 capture inputs, GPIO inputs are edge sensitive and must be high for greater than 40 ns to be recognized as interrupt inputs. These bits can also be set due to a manual capture via 1588 Manual Capture (1588_MANUAL_CAPTURE) .	R/WC	00h
15	RESERVED	RO	-
14:12	1588 TX Timestamp Interrupt (1588_TX_TS_INT[2:0]) This interrupt (one bit per port) indicates that a PTP packet was transmitted and its egress time stored. Up to four events, as indicated by the 1588 TX Timestamp Count (1588_TX_TS_CNT[2:0]) field in the 1588 Port x Capture Information Register (1588_CAP_INFO_x) , are buffered per port.	R/WC	000b
11	RESERVED	RO	-
10:8	1588 RX Timestamp Interrupt (1588_RX_TS_INT[2:0]) This interrupt (one bit per port) indicates that a PTP packet was received and its ingress time and associated data stored. Up to four events, as indicated by the 1588 RX Timestamp Count (1588_RX_TS_CNT[2:0]) field in the 1588 Port x Capture Information Register (1588_CAP_INFO_x) , are buffered per port.	R/WC	000b
7:2	RESERVED	RO	-

Bits	Description	Type	Default
1	1588 Timer Interrupt B (1588_TIMER_INT_B) This interrupt indicates that the 1588 clock equaled or passed the Clock Event Channel B Clock Target value in the 1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) and 1588 Clock Target x Nano-Seconds Register (1588_CLOCK_TARGET_NS_x) x=B. Note: This bit is also cleared by an active edge on a GPIO if enabled. For the clear function, GPIO inputs are edge sensitive and must be active for greater than 40 ns to be recognized as a clear input.	R/WC	0b
0	1588 Timer Interrupt A (1588_TIMER_INT_A) This interrupt indicates that the 1588 clock equaled or passed the Clock Event Channel A Clock Target value in the 1588 Clock Target x Seconds Register (1588_CLOCK_TARGET_SEC_x) and 1588 Clock Target x Nano-Seconds Register (1588_CLOCK_TARGET_NS_x) x=A. Note: This bit is also cleared by an active edge on a GPIO if enabled. For the clear function, GPIO inputs are edge sensitive and must be active for greater than 40 ns to be recognized as a clear input.	R/WC	0b

15.8.4 1588 INTERRUPT ENABLE REGISTER (1588_INT_EN)

Offset: 10Ch Size: 32 bits
Bank: na

This read/write register contains the 1588 interrupt enable bits.

If enabled, these interrupt bits are cascaded into the [1588 Interrupt Event \(1588_EVNT\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#). Writing a 1 to an interrupt enable bits will enable the corresponding interrupt as a source. Status bits will still reflect the status of the interrupt source regardless of whether the source is enabled as an interrupt in this register. The [1588 Interrupt Event Enable \(1588_EVNT_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must also be set in order for an actual system level interrupt to occur. Refer to [Section 8.0, "System Interrupts," on page 73](#) for additional information.

Bits	Description	Type	Default
31:24	1588 GPIO Falling Edge Interrupt Enable (1588_GPIO_FE_EN[7:0])	R/W	00h
23:16	1588 GPIO Rising Edge Interrupt Enable (1588_GPIO_RE_EN[7:0])	R/W	00h
15	RESERVED	RO	-
14:12	1588 TX Timestamp Enable (1588_TX_TS_EN[2:0])	R/W	000b
11	RESERVED	RO	-
10:8	1588 RX Timestamp Enable (1588_RX_TS_EN[2:0])	R/W	000b
7:2	RESERVED	RO	-
1	1588 Timer B Interrupt Enable (1588_TIMER_EN_B)	R/W	0b
0	1588 Timer A Interrupt Enable (1588_TIMER_EN_A)	R/W	0b

15.8.5 1588 CLOCK SECONDS REGISTER (1588_CLOCK_SEC)

Offset: 110h Size: 32 bits
Bank: na

This register contains the seconds portion of the 1588 Clock. It is used to read the 1588 Clock following the setting of the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) and to directly change the 1588 Clock when the [Clock Load \(1588_CLOCK_LOAD\)](#) bit is set.

Bits	Description	Type	Default
31:0	Clock Seconds (1588_CLOCK_SEC) This field contains the seconds portion of the 1588 Clock.	R/W	00000000h

Note: The value read is the saved value of the 1588 Clock when the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is set.

LAN9352

15.8.6 1588 CLOCK NANOSECONDS REGISTER (1588_CLOCK_NS)

Offset: 114h Size: 32 bits
Bank: na

This register contains the nanoseconds portion of the 1588 Clock. It is used to read the 1588 Clock following the setting of the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) and to directly change the 1588 Clock when the [Clock Load \(1588_CLOCK_LOAD\)](#) bit is set.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Clock NanoSeconds (1588_CLOCK_NS) This field contains the nanoseconds portion of the 1588 Clock.	R/W	00000000h

Note: The value read is the saved value of the 1588 Clock when the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is set.

15.8.7 1588 CLOCK SUB-NANOSECONDS REGISTER (1588_CLOCK_SUBNS)

Offset: 118h Size: 32 bits
Bank: na

This register contains the sub-nanoseconds portion of the 1588 Clock. It is used to read the 1588 Clock following the setting of the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) and to directly change the 1588 Clock when the [Clock Load \(1588_CLOCK_LOAD\)](#) bit is set.

Bits	Description	Type	Default
31:0	Clock Sub-NanoSeconds (1588_CLOCK_SUBNS) This field contains the sub-nanoseconds portion of the 1588 Clock.	R/W	00000000h

Note: The value read is the saved value of the 1588 Clock when the [Clock Read \(1588_CLOCK_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is set.

15.8.8 1588 CLOCK RATE ADJUSTMENT REGISTER (1588_CLOCK_RATE_ADJ)

Offset: 11Ch Size: 32 bits
Bank: na

This register is used to adjust the rate of the 1588 Clock. Every 10 ns, 1588 Clock is normally incremented by 10 ns. This register is used to occasionally change that increment to 9 or 11 ns.

Bits	Description	Type	Default
31	Clock Rate Adjustment Direction (1588_CLOCK_RATE_ADJ_DIR) This field specifies if the 1588 Rate Adjustment causes the 1588 Clock to be faster or slower than the reference clock. 0 = slower (1588 Clock increments by 9 ns) 1 = faster (1588 Clock increments by 11 ns)	R/W	0b
30	RESERVED	RO	-
29:0	Clock Rate Adjustment Value (1588_CLOCK_RATE_ADJ_VALUE) This field indicates an adjustment to the reference clock period of the 1588 Clock in units of 2^{-32} ns. On each 10 ns reference clock cycle, this value is added to the 32-bit sub-nanoseconds portion of the 1588 Clock. When the sub-nanoseconds portion wraps around to zero, the 1588 Clock will be adjusted by 1ns.	R/W	00000000h

15.8.9 1588 CLOCK TEMPORARY RATE ADJUSTMENT REGISTER (1588_CLOCK_TEMP_RATE_ADJ)

Offset: 120h Size: 32 bits
Bank: na

This register is used to temporarily adjust the rate of the 1588 Clock. Every 10 ns, 1588 Clock is normally incremented by 10 ns. This register is used to occasionally change that increment to 9 or 11 ns.

Bits	Description	Type	Default
31	Clock Temporary Rate Adjustment Direction (1588_CLOCK_TEMP_RATE_ADJ_DIR) This field specifies if the 1588 Temporary Rate Adjustment causes the 1588 Clock to be faster or slower than the reference clock. 0 = slower (1588 Clock increments by 9 ns) 1 = faster (1588 Clock increments by 11 ns)	R/W	0b
30	RESERVED	RO	-
29:0	Clock Temporary Rate Adjustment Value (1588_CLOCK_TEMP_RATE_ADJ_VALUE) This field indicates a temporary adjustment to the reference clock period of the 1588 Clock in units of 2^{-32} ns. On each 10ns reference clock cycle, this value is added to the 32-bit sub-nanoseconds portion of the 1588 Clock. When the sub-nanoseconds portion wraps around to zero, the 1588 Clock will be adjusted by 1ns (a 9 or 11 ns increment instead of the normal 10ns).	R/W	00000000h

15.8.10 1588 CLOCK TEMPORARY RATE DURATION REGISTER
(1588_CLOCK_TEMP_RATE_DURATION)

Offset: 124h Size: 32 bits
Bank: na

This register specifies the active duration of the temporary clock rate adjustment.

Bits	Description	Type	Default
31:0	Clock Temporary Rate Duration (1588_CLOCK_TEMP_RATE_DURATION) This field specifies the duration of the temporary rate adjustment in reference clock cycles.	R/W	00000000h

15.8.11 1588 CLOCK STEP ADJUSTMENT REGISTER (1588_CLOCK_STEP_ADJ)

Offset: 128h Size: 32 bits
Bank: na

This register is used to perform a one-time adjustment to either the seconds portion or the nanoseconds portion of the 1588 Clock. The amount and direction can be specified.

Bits	Description	Type	Default
31	Clock Step Adjustment Direction (1588_CLOCK_STEP_ADJ_DIR) This field specifies if the Clock Step Adjustment Value (1588_CLOCK_STEP_ADJ_VALUE) is added to or subtracted from the 1588 Clock. 0 = subtracted 1 = added Note: Only addition is supported for the nanoseconds portion of the 1588 Clock	R/W	0b
30	RESERVED	RO	-
29:0	Clock Step Adjustment Value (1588_CLOCK_STEP_ADJ_VALUE) When the nanoseconds portion of the 1588 Clock is being adjusted, this field specifies the amount to add. This is in lieu of the normal 9, 10 or 11 ns increment. When the seconds portion of the 1588 Clock is being adjusted, the lower 4 bits of this field specify the amount to add to or subtract.	R/W	00000000h

15.8.12 1588 CLOCK TARGET X SECONDS REGISTER (1588_CLOCK_TARGET_SEC_X)

Offset: Channel A: [12Ch](#) Size: 32 bits
 Channel B: [13Ch](#)
 Bank: Channel A: [na](#)
 Channel B: [na](#)

This read/write register combined with [1588 Clock Target x NanoSeconds Register \(1588_CLOCK_TARGET_NS_x\)](#) form the 1588 Clock Target value. The 1588 Clock Target value is compared to the current 1588 Clock value and can be used to trigger an interrupt upon at match. Refer to [Section 15.4, "1588 Clock Events"](#) for additional information.

Bits	Description	Type	Default
31:0	Clock Target Seconds (CLOCK_TARGET_SEC) This field contains the seconds portion of the 1588 Clock Compare value.	R/W	00000000h

Note: Both this register and the [1588 Clock Target x NanoSeconds Register \(1588_CLOCK_TARGET_NS_x\)](#) must be written for either to be affected.

Note: The value read is the saved value of the 1588 Clock Target when the [Clock Target Read \(1588_CLOCK_TARGET_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is set.

Note: When the [Clock Target Read \(1588_CLOCK_TARGET_READ\)](#) bit is set, the previous value written to this register is overwritten. Normally, a read command should not be requested in between writing this register and the [1588 Clock Target x NanoSeconds Register \(1588_CLOCK_TARGET_NS_x\)](#).

15.8.13 1588 CLOCK TARGET X NANOSECONDS REGISTER (1588_CLOCK_TARGET_NS_X)

Offset: Channel A: [130h](#) Size: 32 bits
 Channel B: [140h](#)
 Bank: Channel A: [na](#)
 Channel B: [na](#)

This read/write register combined with [1588 Clock Target x Seconds Register \(1588_CLOCK_TARGET_SEC_x\)](#) form the 1588 Clock Target value. The 1588 Clock Target value is compared to the current 1588 Clock value and can be used to trigger an interrupt upon at match. Refer to [Section 15.4, "1588 Clock Events"](#) for additional information.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Clock Target NanoSeconds (CLOCK_TARGET_NS) This field contains the nanoseconds portion of the 1588 Clock Compare value.	R/W	00000000h

Note: Both this register and the [1588 Clock Target x Seconds Register \(1588_CLOCK_TARGET_SEC_x\)](#) must be written for either to be affected.

Note: The value read is the saved value of the 1588 Clock Target when the [Clock Target Read \(1588_CLOCK_TARGET_READ\)](#) bit in the [1588 Command and Control Register \(1588_CMD_CTL\)](#) is set.

Note: When the [Clock Target Read \(1588_CLOCK_TARGET_READ\)](#) bit is set, the previous value written to this register is overwritten. Normally, a read command should not be requested in between writing this register and the [1588 Clock Target x Seconds Register \(1588_CLOCK_TARGET_SEC_x\)](#).

15.8.14 1588 CLOCK TARGET X RELOAD / ADD SECONDS REGISTER (1588_CLOCK_TARGET_RELOAD_SEC_X)

Offset: Channel A: [134h](#) Size: 32 bits
Channel B: [144h](#)
Bank: Channel A: [na](#)
Channel B: [na](#)

This read/write register combined with [1588 Clock Target x Reload / Add NanoSeconds Register \(1588_CLOCK_TARGET_RELOAD_NS_x\)](#) form the 1588 Clock Target Reload value. The 1588 Clock Target Reload is the value that is reloaded or added to the 1588 Clock Compare value when a clock compare event occurs. Refer to [Section 15.4, "1588 Clock Events"](#) for additional information.

Bits	Description	Type	Default
31:0	Clock Target Reload Seconds (CLOCK_TARGET_RELOAD_SEC) This field contains the seconds portion of the 1588 Clock Target Reload value that is reloaded to the 1588 Clock Compare value.	R/W	00000000h

Note: Both this register and the [1588 Clock Target x Reload / Add NanoSeconds Register \(1588_CLOCK_TARGET_RELOAD_NS_x\)](#) must be written for either to be affected.

15.8.15 1588 CLOCK TARGET X RELOAD / ADD NANOSECONDS REGISTER (1588_CLOCK_TARGET_RELOAD_NS_X)

Offset: Channel A: [138h](#) Size: 32 bits
 Channel B: [148h](#)
 Bank: Channel A: [na](#)
 Channel B: [na](#)

This read/write register combined with [1588 Clock Target x Reload / Add Seconds Register \(1588_CLOCK_TARGET_RELOAD_SEC_x\)](#) form the 1588 Clock Target Reload value. The 1588 Clock Target Reload is the value that is reloaded or added to the 1588 Clock Compare value when a clock compare event occurs. Refer to [Section 15.4, "1588 Clock Events"](#) for additional information.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Clock Target Reload NanoSeconds (CLOCK_TARGET_RELOAD_NS) This field contains the nanoseconds portion of the 1588 Clock Target Reload value that is reloaded to the 1588 Clock Compare value.	R/W	00000000h

Note: Both this register and the [1588 Clock Target x Reload / Add Seconds Register \(1588_CLOCK_TARGET_RELOAD_SEC_x\)](#) must be written for either to be affected.

LAN9352

15.8.16 1588 USER MAC ADDRESS HIGH-WORD REGISTER (1588_USER_MAC_HI)

Offset: 14Ch Size: 32 bits
Bank: na

This read/write register combined with the [1588 User MAC Address Low-DWORD Register \(1588_USER_MAC_LO\)](#) forms the 48-bit user defined MAC address. The Auxiliary MAC address can be enabled for each protocol via their respective User Defined MAC Address Enable bit in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#).

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	User MAC Address High (USER_MAC_HI) This field contains the high 16 bits of the user defined MAC address used for PTP packet detection. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0000h

15.8.17 1588 USER MAC ADDRESS LOW-DWORD REGISTER (1588_USER_MAC_LO)

Offset: 150h Size: 32 bits
Bank: na

This read/write register combined with the [1588 User MAC Address High-WORD Register \(1588_USER_MAC_HI\)](#) forms the 48-bit user defined MAC address. The Auxiliary MAC address can be enabled for each protocol via their respective User Defined MAC Address Enable bit in the [1588 Port x RX Parsing Configuration Register \(1588_RX_PARSE_CONFIG_x\)](#).

Bits	Description	Type	Default
31:0	User MAC Address Low (USER_MAC_LO) This field contains the low 32 bits of the user defined MAC address used for PTP packet detection. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	00000000h

LAN9352

15.8.18 1588 BANK PORT GPIO SELECT REGISTER (1588_BANK_PORT_GPIO_SEL)

Offset: 154h Size: 32 bits
Bank: na

Bits	Description	Type	Default
31:11	RESERVED	RO	-
10:8	GPIO Select (GPIO_SEL[2:0]) This field specifies which GPIO the various GPIO x registers will access.	R/W	000b
7:6	RESERVED	RO	-
5:4	Port Select (PORT_SEL[1:0]) This field specifies which port the various Port x registers will access.	R/W	00b
3	RESERVED	RO	-
2:0	Bank Select (BANK_SEL[2:0]) This field specifies which bank of registers is accessed. 000: Ports General 001: Ports RX 010: Ports TX 011: GPIOs 1xx: Reserved	R/W	000b

15.8.19 1588 PORT X LATENCY REGISTER (1588_LATENCY_X)

Offset: 158h Size: 32 bits
Bank: 0

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:16	<p>TX Latency (TX_LATENCY[15:0]) This field specifies the egress delay in nanoseconds between the PTP timestamp point and the network medium. The setting is used to adjust the internally captured 1588 clock value such that the resultant timestamp more accurately corresponds to the start of the frame's first symbol after the SFD on the network medium.</p> <p>The value depends on the port mode. Typical values are:</p> <ul style="list-style-type: none"> 100BASE-TX: 95ns 100BASE-FX: 68ns plus the receive latency of the fiber transceiver 10BASE-T: 1139ns <p>Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.</p>	R/W	<p>20 for Port 0 Note 2</p> <p>95 for Port 1 Note 3</p> <p>95 for Port 2 Note 4</p>
15:0	<p>RX Latency (RX_LATENCY[15:0]) This field specifies the ingress delay in nanoseconds between the network medium and the PTP timestamp point. The setting is used to adjust the internally captured 1588 clock value such that the resultant timestamp more accurately corresponds to the start of the frame's first symbol after the SFD on the network medium.</p> <p>The value depends on the port mode. Typical values are:</p> <ul style="list-style-type: none"> 100BASE-TX: 285ns 100BASE-FX: 231ns plus the receive latency of the fiber transceiver 10BASE-T: 1674ns <p>Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.</p>	R/W	<p>20 for Port 0 Note 2</p> <p>285 for Port 1 Note 3</p> <p>285 for Port 2 Note 4</p>

Note 2: 1588 functions on Port 0 are not normally used and the default is irrelevant.

Note 3: The default value is appropriate for 100BASE-TX mode. For other modes (100BASE-FX or 10BASE-T) the proper value needs to be set, via S/W or EEPROM.

Note 4: The default value is appropriate for 100BASE-TX mode. For other modes (100BASE-FX or 10BASE-T) the proper value needs to be set, via S/W or EEPROM.

15.8.20 1588 PORT X ASYMMETRY AND PEER DELAY REGISTER (1588_ASYM_PEERDLY_X)

Offset: 15Ch Size: 32 bits
Bank: 0

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:16	Port Delay Asymmetry (DELAY_ASYM[15:0]) This field specifies the previously known delay asymmetry in nanoseconds. This is a signed 2's complement number. Positive values occur when the master-to-slave or responder-to-requestor propagation time is longer than the slave-to-master or requestor-to-responder propagation time. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0000h
15:0	RX Peer Delay (RX_PEER_DELAY[15:0]) This field specifies the measured peer delay in nanoseconds used during peer-to-peer mode.	R/W	0000h

15.8.21 1588 PORT X CAPTURE INFORMATION REGISTER (1588_CAP_INFO_X)

Offset: 160h Size: 32 bits
Bank: 0

This read only register provides information about the receive and transmit capture buffers.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:7	RESERVED	RO	-
6:4	1588 TX Timestamp Count (1588_TX_TS_CNT[2:0]) This field indicates how many transmit timestamps are available to be read. It is incremented when a PTP packet is transmitted and decremented when the appropriate 1588 TX Timestamp Interrupt (1588_TX_TS_INT[2:0]) bit is written with a 1.	RO	000b
3	RESERVED	RO	-
2:0	1588 RX Timestamp Count (1588_RX_TS_CNT[2:0]) This field indicates how many receive timestamps are available to be read. It is incremented when a PTP packet is received and decremented when the appropriate 1588 RX Timestamp Interrupt (1588_RX_TS_INT[2:0]) bit is written with a 1.	RO	000b

15.8.22 1588 PORT X RX PARSING CONFIGURATION REGISTER (1588_RX_PARSE_CONFIG_X)

Offset: 158h Size: 32 bits
Bank: 1

This register is used to configure the PTP receive message detection.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:15	RESERVED	RO	-
14	RX Layer 2 Address 1 Enable (RX_LAYER2_ADD1_EN) This bit enables the Layer 2 MAC address of 01:80:C2:00:00:0E for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
13	RX Layer 2 Address 2 Enable (RX_LAYER2_ADD2_EN) This bit enables the Layer 2 MAC address of 01:1B:19:00:00:00 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
12	RX Address 1 Enable (RX_ADD1_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:81 and IPv4 destination address of 224.0.1.129 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:81 and IPv6 destination address of FF0X:0:0:0:0:0:181 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
11	RX Address 2 Enable (RX_ADD2_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:82 and IPv4 destination address of 224.0.1.130 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:82 and IPv6 destination address of FF0X:0:0:0:0:0:182 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

Bits	Description	Type	Default
10	RX Address 3 Enable (RX_ADD3_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:83 and IPv4 destination address of 224.0.1.131 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:83 and IPv6 destination address of FF0X:0:0:0:0:0:183 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
9	RX Address 4 Enable (RX_ADD4_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:84 and IPv4 destination address of 224.0.1.132 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:84 and IPv6 destination address of FF0X:0:0:0:0:0:184 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
8	RX Address 5 Enable (RX_ADD5_EN) This bit enables the IPv4 MAC address of 01:00:5e:00:00:6B and IPv4 destination address of 224.0.0.107 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:00:6B and IPv6 destination address of FF02:0:0:0:0:0:6B for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
7	RX User Defined Layer 2 MAC Address Enable (RX_LAYER2_USER_MAC_EN) This bit enables a user defined Layer 2 MAC address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
6	RX User Defined IPv6 MAC Address Enable (RX_IPV6_USER_MAC_EN) This bit enables a user defined IPv6 MAC address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
5	RX User Defined IPv4 MAC Address Enable (RX_IPV4_USER_MAC_EN) This bit enables the user defined IPv4 MAC address in PTP messages. The address is defined via the 1588 User MAC Address High-WORD Register (1588_USER_MAC_HI) and the 1588 User MAC Address Low-DWORD Register (1588_USER_MAC_LO) . Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

LAN9352

Bits	Description	Type	Default
4	RX IP Address Enable (RX_IP_ADDR_EN) This bit enables the checking of the IP destination address in PTP messages for both IPv4 and IPv6 formats. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
3	RX MAC Address Enable (RX_MAC_ADDR_EN) This bit enables the checking of the MAC destination address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
2	RX Layer 2 Enable (RX_LAYER2_EN) This bit enables the detection of the layer 2 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
1	RX IPv6 Enable (RX_IPV6_EN) This bit enables the detection of the UDP/IPv6 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
0	RX IPv4 Enable (RX_IPV4_EN) This bit enables the detection of the UDP/IPv4 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b

15.8.23 1588 PORT X RX TIMESTAMP CONFIGURATION REGISTER (1588_RX_TIMESTAMP_CONFIG_X)

Offset: 15Ch Size: 32 bits
Bank: 1

This register is used to configure PTP receive message timestamping.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:24	RX PTP Domain (RX_PTP_DOMAIN[7:0]) This field specifies the PTP domain in use. If RX PTP Domain Match Enable (RX_PTP_DOMAIN_EN) is set, the domainNumber in the PTP message must matches the value in this field in order to recorded the ingress time. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	00h
23	RX PTP Domain Match Enable (RX_PTP_DOMAIN_EN) When this bit is set, the domainNumber in the PTP message is checked against the value in RX PTP Domain (RX_PTP_DOMAIN[7:0]) . Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
22	RX PTP Alternate Master Enable (RX_PTP_ALT_MASTER_EN) When this bit is set, the alternateMasterFlag in the PTP message is checked for a zero value. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
21	RX PTP UDP Checksum Check Disable (RX_PTP_UDP_CHKSUM_DIS) When this bit is cleared, ingress times are not saved and ingress messages are not filtered if the frame has an invalid UDP checksum. When this bit is set, the UDP checksum check is bypassed and the ingress time is saved and ingress messages are filtered regardless. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
20	RX PTP FCS Check Disable (RX_PTP_FCS_DIS) When this bit is cleared, ingress times are not saved and ingress messages are not filtered if the frame has an invalid FCS. When this bit is set, the FCS check is bypassed. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

Bits	Description	Type	Default
19:16	RX PTP Version (RX_PTP_VERSION[3:0]) This field specifies the PTP version in use. A setting of 0 allows any PTP version. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	2h
15:0	RX PTP Message Type Enable (RX_PTP_MESSAGE_EN[15:0]) These bits individually enable timestamping of their respective message types. Bit 0 of this field corresponds to a message type value of 0 (Sync), bit 1 to message type value 1 (Delay_Req), etc. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled.	R/W	0000h

15.8.24 1588 PORT X RX TIMESTAMP INSERTION CONFIGURATION REGISTER (1588_RX_TS_INSERT_CONFIG_X)

Offset: 160h Size: 32 bits
Bank: 1

This register is used to configure PTP message timestamp insertion.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:18	RESERVED	RO	-
17	RX PTP Insert Delay Request Egress in Delay Response Enable (RX_PTP_INSERT_DREQ_DRESP_EN) When this bit is set, the egress time of the last Delay_Req packet is inserted into received Delay_Resp packets. This bit has no affect if RX_PTP_INSERT_TS_EN is a low or if detection of the Delay_Resp message type is not enabled.	R/W	0b
16	RX PTP Bad UDP Checksum Force Error Disable (RX_PTP_BAD_UDP_CHKSUM_FORCE_ERR_DIS) When this bit is cleared, ingress packets that have an invalid UDP checksum will have a receive symbol error forced if the packet is modified for timestamp or correction field reasons. When this bit is set, the UDP checksum check is bypassed. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
15	RX PTP Insert Timestamp Seconds Enable (RX_PTP_INSERT_TS_SEC_EN) When RX_PTP_INSERT_TS_EN is set, this bit enables bits 3:0 of the seconds portion of the receive ingress time to be inserted into the PTP message. This bit has no affect if RX_PTP_INSERT_TS_EN is a low.	R/W	0b
14	RESERVED	RO	-
13:8	RX PTP Insert Timestamp Seconds Offset (RX_PTP_INSERT_TS_SEC_OFFSET[5:0]) This field specifies the offset into the PTP header where the seconds portion of the receive ingress time is inserted. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	000101b
7	RX PTP Insert Timestamp Enable (RX_PTP_INSERT_TS_EN) When set, receive ingress times are inserted into the PTP message.	R/W	0b
6	RESERVED	RO	-

LAN9352

Bits	Description	Type	Default
5:0	RX PTP Insert Timestamp Offset (RX_PTP_INSERT_TS_OFFSET[5:0]) This field specifies the offset into the PTP header where the receive ingress time is inserted. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	010000b

15.8.25 1588 PORT X RX CORRECTION FIELD MODIFICATION REGISTER (1588_RX_CF_MOD_X)

Offset: 164h Size: 32 bits
Bank: 1

This register is used to configure RX PTP message correction field modifications.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	RX PTP Correction Field Message Type Enable (RX_PTP_CF_MSG_EN[15:0]) These bits individually enable correction field modification of their respective message types. Bit 0 of this field corresponds to a message type value of 0 (Sync), bit 1 to message type value 1 (Delay_Req), etc. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled	R/W	0000h

LAN9352

15.8.26 1588 PORT X RX FILTER CONFIGURATION REGISTER (1588_RX_FILTER_CONFIG_X)

Offset: 168h Size: 32 bits
Bank: 1

This register is used to configure PTP message filtering.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:19	RESERVED	RO	-
18	RX PTP Alternate Master Filter Enable (RX_PTP_ALT_MASTER_FLTR_EN) This bit enables message filtering based on the alternateMasterFlag flagField bit. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
17	RX PTP Domain Filter Enable (RX_PTP_DOMAIN_FLTR_EN) This bit enables message filtering based on the PTP domain. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
16	RX PTP Version Filter Enable (RX_PTP_VERSION_FLTR_EN) This bit enables message filtering based on the PTP version. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
15:0	RX PTP Message Type Filter Enable (RX_PTP_MSG_FLTR_EN[15:0]) These bits enable individual message filtering. Bit 0 of this field corresponds to a message type value of 0 (Sync), bit 1 to message type value 1 (Delay_Req), etc. Typically Delay_Req and Delay_Resp messages are filtered for peer-to-peer transparent clocks.	R/W	0000h

15.8.27 1588 PORT X RX INGRESS TIME SECONDS REGISTER (1588_RX_INGRESS_SEC_X)

Offset: 16Ch Size: 32 bits
Bank: 1

This read only register combined with the [1588 Port x RX Ingress Time NanoSeconds Register \(1588_RX_INGRESS_NS_x\)](#) contains the RX timestamp captures. Up to four captures are buffered.

Note: Values are only valid if the [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) field indicates that at least one timestamp is available.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the receive ingress time.	RO	00000000h

15.8.28 1588 PORT X RX INGRESS TIME NANOSECONDS REGISTER (1588_RX_INGRESS_NS_X)

Offset: 170h Size: 32 bits
Bank: 1

This read only register combined with the [1588 Port x RX Ingress Time Seconds Register \(1588_RX_INGRESS_SEC_x\)](#) contains the RX timestamp capture. Up to four captures are buffered.

Note: Values are only valid if the [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) field indicates that at least one timestamp is available.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the receive ingress time.	RO	00000000h

15.8.29 1588 PORT X RX MESSAGE HEADER REGISTER (1588_RX_MSG_HEADER_X)

Offset: 174h Size: 32 bits
Bank: 1

This read only register contains the RX message header. Up to four captures are buffered.

Note: Values are only valid if the [1588 RX Timestamp Count \(1588_RX_TS_CNT\[2:0\]\)](#) field indicates that at least one timestamp is available.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:20	Source Port Identity CRC (SRC_PRT_CRC) This field contains the 12-bit CRC of the sourcePortIdentity field of the received PTP packet.	RO	000h
19:16	Message Type (MSG_TYPE) This field contains the messageType field of the received PTP packet.	RO	0h
15:0	Sequence ID (SEQ_ID) This field contains the sequenceId field of the received PTP packet.	RO	0000h

15.8.30 1588 PORT X RX PDELAY_REQ INGRESS TIME SECONDS REGISTER
(1588_RX_PDREQ_SEC_X)

Offset: 178h Size: 32 bits
Bank: 1

This register combined with the 1588 Port x RX Pdelay_Req Ingress Time NanoSeconds Register (1588_RX_PDREQ_NS_x) contains the ingress time of the last Pdelay_Req message. This register is automatically updated if the Auto Update (AUTO) bit is set.

Note: Port and GPIO registers share a common address space. Port registers are selected by the Bank Select (BANK_SEL[2:0] in the 1588 Bank Port GPIO Select Register (1588_BANK_PORT_GPIO_SEL). The port accessed ("x") is set by the Port Select (PORT_SEL[1:0]) field.

Bits	Description	Type	Default
31:4	RESERVED	RO	-
3:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the receive ingress time.	R/W	0h

15.8.31 1588 PORT X RX PDELAY_REQ INGRESS TIME NANOSECONDS REGISTER (1588_RX_PDREQ_NS_X)

Offset: 17Ch Size: 32 bits
Bank: 1

This register combined with the [1588 Port x RX Pdelay_Req Ingress Time Seconds Register \(1588_RX_PDREQ_SEC_x\)](#) contains the ingress time of the last Pdelay_Req message. This register is automatically updated if the [Auto Update \(AUTO\)](#) bit is set.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31	Auto Update (AUTO) If this bit is set, the TS_NS field in this register, the TS_SEC field in 1588_RX_PDREQ_SEC_x and the CF field in 1588_RX_PDREQ_CF_HI_x / 1588_RX_PDREQ_CF_LO_x are updated when a PDelay_Req message is received. When cleared, S/W is responsible to maintain those fields. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the receive ingress time.	R/W	00000000h

15.8.32 1588 PORT X RX PDELAY_REQ INGRESS CORRECTION FIELD HIGH REGISTER
(1588_RX_PDREQ_CF_HI_X)

Offset: 180h Size: 32 bits
Bank: 1

This register combined with the 1588 Port x RX Pdelay_Req Ingress Correction Field Low Register (1588_RX_PDREQ_CF_LOW_x) contains the correction field from the last Pdelay_Req message. Only the nanoseconds portion is used.

This register is automatically updated if the Auto Update (AUTO) bit is set.

Note: Port and GPIO registers share a common address space. Port registers are selected by the Bank Select (BANK_SEL[2:0] in the 1588 Bank Port GPIO Select Register (1588_BANK_PORT_GPIO_SEL). The port accessed ("x") is set by the Port Select (PORT_SEL[1:0]) field.

Bits	Description	Type	Default
31:0	Correction Field (CF[63:32]) This field contains the upper 32 bits of the correction field.	R/W	00000000h

15.8.33 1588 PORT X RX PDELAY_REQ INGRESS CORRECTION FIELD LOW REGISTER (1588_RX_PDREQ_CF_LOW_X)

Offset: 184h Size: 32 bits
Bank: 1

This register combined with the [1588 Port x RX Pdelay_Req Ingress Correction Field High Register \(1588_RX_PDREQ_CF_HI_X\)](#) contains the correction field from the last Pdelay_Req message. Only the nanoseconds portion is used.

This register is automatically updated if the [Auto Update \(AUTO\)](#) bit is set.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:16	Correction Field (CF[31:16]) This field contains the low middle 16 bits of the correction field.	R/W	0000h
15:0	RESERVED	RO	-

15.8.34 1588 PORT X RX CHECKSUM DROPPED COUNT REGISTER (1588_RX_CHKSUM_DROPPED_CNT_X)

Offset: 188h Size: 32 bits
Bank: 1

This register counts the number of packets dropped at ingress due to a bad UDP checksum. The packet will also be counted as an error by the receiving MAC.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:0	Bad Checksum Dropped Count (BAD_CHKSUM_DROPPED_CNT[31:0]) This field is a count of packets dropped at ingress due to a bad UDP checksum. It can be cleared by writing a zero value at the risk of losing any previous count. Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.	R/W	00000000h

15.8.35 1588 PORT X RX FILTERED COUNT REGISTER (1588_RX_FILTERED_CNT_X)

Offset: 18Ch Size: 32 bits
Bank: 1

This register counts the number of packets filtered at ingress due to [Ingress Message Filtering](#). The packet will also be counted as an error by the receiving MAC.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:0	<p>Filtered Count (FILTERED_CNT[31:0]) This field is a count of packets dropped at ingress due to Ingress Message Filtering. It can be cleared by writing a zero value at the risk of losing any previous count.</p> <p>Note: This counter will stop at its maximum value of FFFF_FFFFh. Minimum rollover time at 100 Mbps is approximately 481 hours.</p>	R/W	00000000h

LAN9352

15.8.36 1588 PORT X TX PARSING CONFIGURATION REGISTER (1588_TX_PARSE_CONFIG_X)

Offset: 158h Size: 32 bits
Bank: 2

This register is used to configure the PTP transmit message detection.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:15	RESERVED	RO	-
14	TX Layer 2 Address 1 Enable (TX_LAYER2_ADD1_EN) This bit enables the Layer 2 MAC address of 01:80:C2:00:00:0E for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
13	TX Layer 2 Address 2 Enable (TX_LAYER2_ADD2_EN) This bit enables the Layer 2 MAC address of 01:1B:19:00:00:00 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
12	TX Address 1 Enable (TX_ADD1_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:81 and IPv4 destination address of 224.0.1.129 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:81 and IPv6 destination address of FF0X:0:0:0:0:0:0:181 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
11	TX Address 2 Enable (TX_ADD2_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:82 and IPv4 destination address of 224.0.1.130 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:82 and IPv6 destination address of FF0X:0:0:0:0:0:0:182 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

Bits	Description	Type	Default
10	TX Address 3 Enable (TX_ADD3_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:83 and IPv4 destination address of 224.0.1.131 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:83 and IPv6 destination address of FF0X:0:0:0:0:0:183 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
9	TX Address 4 Enable (TX_ADD4_EN) This bit enables the IPv4 MAC address of 01:00:5E:00:01:84 and IPv4 destination address of 224.0.1.132 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:01:84 and IPv6 destination address of FF0X:0:0:0:0:0:184 for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
8	TX Address 5 Enable (TX_ADD5_EN) This bit enables the IPv4 MAC address of 01:00:5e:00:00:6B and IPv4 destination address of 224.0.0.107 for PTP packets. This bit enables the IPv6 MAC address of 33:33:00:00:00:6B and IPv6 destination address of FF02:0:0:0:0:0:6B for PTP packets. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
7	TX User Defined Layer 2 MAC Address Enable (TX_LAYER2_USER_MAC_EN) This bit enables a user defined Layer 2 MAC address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
6	TX User Defined IPv6 MAC Address Enable (TX_IPV6_USER_MAC_EN) This bit enables a user defined IPv6 MAC address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
5	TX User Defined IPv4 MAC Address Enable (TX_IPV4_USER_MAC_EN) This bit enables the user defined IPv4 MAC address in PTP messages. The address is defined via the 1588 User MAC Address High-WORD Register (1588_USER_MAC_HI) and the 1588 User MAC Address Low-DWORD Register (1588_USER_MAC_LO) . Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

LAN9352

Bits	Description	Type	Default
4	TX IP Address Enable (TX_IP_ADDR_EN) This bit enables the checking of the IP destination address in PTP messages for both IPv4 and IPv6 formats. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
3	TX MAC Address Enable (TX_MAC_ADDR_EN) This bit enables the checking of the MAC destination address in PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
2	TX Layer 2 Enable (TX_LAYER2_EN) This bit enables the detection of the layer 2 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
1	TX IPv6 Enable (TX_IPV6_EN) This bit enables the detection of the UDP/IPv6 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b
0	TX IPv4 Enable (TX_IPV4_EN) This bit enables the detection of the UDP/IPv4 formatted PTP messages. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	1b

15.8.37 1588 PORT X TX TIMESTAMP CONFIGURATION REGISTER (1588_TX_TIMESTAMP_CONFIG_X)

Offset: 15Ch Size: 32 bits
Bank: 2

This register is used to configure PTP transmit message timestamping.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:24	TX PTP Domain (TX_PTP_DOMAIN[7:0]) This field specifies the PTP domain in use. If TX PTP Domain Match Enable (TX_PTP_DOMAIN_EN) is set, the domainNumber in the PTP message must match the value in this field in order to recorded the egress time. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	00h
23	TX PTP Domain Match Enable (TX_PTP_DOMAIN_EN) When this bit is set, the domainNumber in the PTP message is checked against the value in TX PTP Domain (TX_PTP_DOMAIN[7:0]) . Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
22	TX PTP Alternate Master Enable (TX_PTP_ALT_MASTER_EN) When this bit is set, the alternateMasterFlag in the PTP message is checked for a zero value. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
21	TX PTP UDP Checksum Check Disable (TX_PTP_UDP_CHKSUM_DIS) When this bit is cleared, egress times are not saved if the frame has an invalid UDP checksum. When this bit is set, the UDP checksum check is bypassed and the egress time is saved regardless. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
20	TX PTP FCS Check Disable (TX_PTP_FCS_DIS) When this bit is cleared, egress times are not saved if the frame has an invalid FCS. When this bit is set, the FCS check is bypassed. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

LAN9352

Bits	Description	Type	Default
19:16	TX PTP Version (TX_PTP_VERSION[3:0]) This field specifies the PTP version in use. A setting of 0 allows any PTP version. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	2h
15:0	TX PTP Message Type Enable (TX_PTP_MESSAGE_EN[15:0]) These bits individually enable timestamping of their respective message types. Bit 0 of this field corresponds to a message type value of 0 (Sync), bit 1 to message type value 1 (Delay_Req), etc. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled	R/W	0000h

15.8.38 1588 PORT X TX MODIFICATION REGISTER (1588_TX_MOD_X)

Offset: 164h Size: 32 bits
Bank: 2

This register is used to configure TX PTP message modifications.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31	TX PTP Clear Four Byte Reserved Field (TX_PTP_CLR_4_RSVD) This bit enables the clearing of the four byte reserved field if the frame was modified on transmission.	R/W	0b
30	TX PTP Suppress Timestamps when Correction Field Adjusted (TX_PTP_SUPP_CF_TS) This bit prevents egress times from being saved if correction field modification is done. This is used to suppress timestamps from frames forwarded across the switch. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b
29	TX PTP Pdelay_Resp Message Turnaround Time Insertion (TX_PTP_PDRESP_TA_INSERT) Note: This bit enables the turnaround time between the received Pdelay_Req and the transmitted Pdelay_Resp to be inserted into the correction field of Pdelay_Resp messages sent by the Host.	R/W	0b
28	TX PTP Sync Message Egress Time Insertion (TX_PTP_SYNC_TS_INSERT) This bit enables the egress time to be inserted into the originTimestamp field of Sync messages sent by the Host.	R/W	0b
27:22	TX PTP 4 Reserved Bytes Offset (TX_PTP_4_RSVD_OFFSET[5:0]) This field specifies the offset into the PTP header of the four reserved bytes which the transmitter would clear if enabled. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	010000b
21:16	TX PTP 1 Reserved Byte Offset (TX_PTP_1_RSVD_OFFSET[5:0]) This field specifies the offset into the PTP header where the transmitter can retrieve the seconds portion of the ingress time. Note: The host S/W must not change this field while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	000101b

LAN9352

Bits	Description	Type	Default
15:0	TX PTP Correction Field Message Type Enable (TX_PTP_CF_MSG_EN[15:0]) These bits individually enable correction field modification of their respective message types. Bit 0 of this field corresponds to a message type value of 0 (Sync), bit 1 to message type value 1 (Delay_Req), etc. Typically Sync, Delay_Req, PDelay_Req and PDelay_Resp messages are enabled	R/W	0000h

15.8.39 1588 PORT X TX MODIFICATION REGISTER 2 (1588_TX_MOD2_X)

Offset: 168h Size: 32 bits
Bank: 2

This register is used to configure TX PTP message modifications.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:1	RESERVED	RO	-
0	TX PTP Clear UDP/IPv4 Checksum Enable (TX_PTP_CLR_UDPV4_CHKSUM) This bit enables the clearing of the UDP/IPv4 checksum when Pdelay_Resp Message Turnaround Time Insertion or Sync Message Egress Time Insertion is enabled. Note: The host S/W must not change this bit while the 1588 Enable (1588_ENABLE) bit in 1588 Command and Control Register (1588_CMD_CTL) is set.	R/W	0b

15.8.40 1588 PORT X TX EGRESS TIME SECONDS REGISTER (1588_TX_EGRESS_SEC_X)

Offset: 16Ch Size: 32 bits
Bank: 2

This read only register combined with the 1588 Port x TX Egress Time NanoSeconds Register (1588_TX_EGRESS_NS_x) contains the TX timestamp captures. Up to four captures are buffered.

- Note:** Values are only valid if the 1588 TX Timestamp Count (1588_TX_TS_CNT[2:0]) field indicates that at least one timestamp is available.
- Note:** Port and GPIO registers share a common address space. Port registers are selected by the Bank Select (BANK_SEL[2:0] in the 1588 Bank Port GPIO Select Register (1588_BANK_PORT_GPIO_SEL). The port accessed ("x") is set by the Port Select (PORT_SEL[1:0]) field.

Bits	Description	Type	Default
31:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the transmit egress time.	RO	00000000h

15.8.41 1588 PORT X TX EGRESS TIME NANOSECONDS REGISTER (1588_TX_EGRESS_NS_X)

Offset: 170h Size: 32 bits
Bank: 2

This read only register combined with the [1588 Port x TX Egress Time Seconds Register \(1588_TX_EGRESS_SEC_x\)](#) contains the TX timestamp capture. Up to four captures are buffered.

Note: Values are only valid if the [1588 TX Timestamp Count \(1588_TX_TS_CNT\[2:0\]\)](#) field indicates that at least one timestamp is available.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the transmit egress time.	RO	00000000h

15.8.42 1588 PORT X TX MESSAGE HEADER REGISTER (1588_TX_MSG_HEADER_X)

Offset: 174h Size: 32 bits
Bank: 2

This read only register contains the TX message header. Up to four captures are buffered.

- Note:** Values are only valid if the [1588 TX Timestamp Count \(1588_TX_TS_CNT\[2:0\]\)](#) field indicates that at least one timestamp is available.
- Note:** Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed (“x”) is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:20	Source Port Identity CRC (SRC_PRT_CRC) This field contains the 12-bit CRC of the sourcePortIdentity field of the transmitted PTP packet.	RO	000h
19:16	Message Type (MSG_TYPE) This field contains the messageType field of the transmitted PTP packet.	RO	0h
15:0	Sequence ID (SEQ_ID) This field contains the sequenceId field of the transmitted PTP packet.	RO	0000h

15.8.43 1588 PORT X TX DELAY_REQ EGRESS TIME SECONDS REGISTER (1588_TX_DREQ_SEC_X)

Offset: 178h Size: 32 bits
Bank: 2

This register combined with the [1588 Port x TX Delay_Req Egress Time NanoSeconds Register \(1588_TX_DREQ_NS_x\)](#) contains the egress time of the last Delay_Req message. The contents of this field are normally only used to insert the egress time into received Delay_Resp messages.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

BITS	DESCRIPTION	TYPE	DEFAULT
31:4	RESERVED	RO	-
3:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the transmit egress time.	RO	0h

15.8.44 1588 PORT X TX DELAY_REQ EGRESS TIME NANoseconds REGISTER (1588_TX_DREQ_NS_X)

Offset: 17Ch Size: 32 bits
Bank: 2

This register combined with the [1588 Port x TX Delay_Req Egress Time Seconds Register \(1588_TX_DREQ_SEC_x\)](#) contains the egress time of the last Delay_Req message. The contents of this field are normally only used to insert the egress time into received Delay_Resp messages.

Note: Port and GPIO registers share a common address space. Port registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The port accessed ("x") is set by the [Port Select \(PORT_SEL\[1:0\]\)](#) field.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the transmit egress time.	RO	00000000h

15.8.45 1588 TX ONE-STEP SYNC UPPER SECONDS REGISTER (1588_TX_ONE_STEP_SYNC_SEC)

Offset: 180h Size: 32 bits
Bank: 2

This register contains the highest 16 bits of the originTimestamp which is inserted into Sync messages when one-step timestamp insertion is enabled.

Note: This is a static field that is maintained by the Host. It is not incremented when the lower 32 bits of the 1588 Clock rollover.

Note: This register applies to all ports.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	Clock Seconds High (1588_CLOCK_SEC_HI) This field contains the highest 16 bits of seconds of the 1588 Clock.	R/W	0000h

15.8.46 1588 GPIO CAPTURE CONFIGURATION REGISTER (1588_GPIO_CAP_CONFIG)

Offset: 15Ch Size: 32 bits
Bank: 3

Note: Port and GPIO registers share a common address space. GPIO registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#).

Note: The IEEE 1588 Unit supports 8 GPIO signals.

Bits	Description	Type	Default
31:24	Lock Enable GPIO Falling Edge (LOCK_GPIO_FE) These bits enable/disable the GPIO falling edge lock. This lock prevents a 1588 capture from overwriting the Clock value if the GPIO interrupt in the 1588 Interrupt Status Register (1588_INT_STS) is already set due to a previous capture. 0: Disables GPIO falling edge lock 1: Enables GPIO falling edge lock	R/W	FFh
23:16	Lock Enable GPIO Rising Edge (LOCK_GPIO_RE) These bits enable/disable the GPIO rising edge lock. This lock prevents a 1588 capture from overwriting the Clock value if the GPIO interrupt in the 1588 Interrupt Status Register (1588_INT_STS) is already set due to a previous capture. 0: Disables GPIO rising edge lock 1: Enables GPIO rising edge lock	R/W	FFh
15:8	GPIO Falling Edge Capture Enable 7-0 (GPIO_FE_CAPTURE_ENABLE[7:0]) These bits enable the falling edge of the respective GPIO input to capture the 1588 clock value and to set the respective 1588_GPIO interrupt in the 1588 Interrupt Status Register (1588_INT_STS) . 0: Disables GPIO Capture 1: Enables GPIO Capture Note: The GPIO must be configured as an input for this function to operate. GPIO inputs are edge sensitive and must be low for greater than 40 ns to be recognized.	R/W	00h
7:0	GPIO Rising Edge Capture Enable 7-0 (GPIO_RE_CAPTURE_ENABLE[7:0]) These bits enable the rising edge of the respective GPIO input to capture the 1588 clock value and to set the respective 1588_GPIO interrupt in the 1588 Interrupt Status Register (1588_INT_STS) . 0: Disables GPIO Capture 1: Enables GPIO Capture Note: The GPIO must be configured as an input for this function to operate. GPIO inputs are edge sensitive and must be high for greater than 40 ns to be recognized.	R/W	00h

15.8.47 1588 GPIO X RISING EDGE CLOCK SECONDS CAPTURE REGISTER (1588_GPIO_RE_CLOCK_SEC_CAP_X)

Offset: 16Ch Size: 32 bits
Bank: 3

This read only register combined with the [1588 GPIO x Rising Edge Clock NanoSeconds Capture Register \(1588_GPIO_RE_CLOCK_NS_CAP_x\)](#) forms the GPIO rising edge timestamp capture.

Note: Values are only valid if the appropriate [1588 GPIO Rising Edge Interrupt \(1588_GPIO_RE_INT\[7:0\]\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) indicates that a timestamp is available.

Note: Unless the corresponding [Lock Enable GPIO Rising Edge \(LOCK_GPIO_RE\)](#) bit is set, a new capture may occur between reads of this register and the [1588 GPIO x Rising Edge Clock NanoSeconds Capture Register \(1588_GPIO_RE_CLOCK_NS_CAP_x\)](#). Software techniques are required to avoid reading intermediate values.

Note: Port and GPIO registers share a common address space. GPIO registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The GPIO accessed ("x") is set by the [GPIO Select \(GPIO_SEL\[2:0\]\)](#) field.

Note: All 8 GPIO register sets are available.

Bits	Description	Type	Default
31:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the timestamp upon the rising edge of a GPIO or upon a software commanded manual capture.	RO	00000000h

LAN9352

15.8.48 1588 GPIO X RISING EDGE CLOCK NANOSECONDS CAPTURE REGISTER (1588_GPIO_RE_CLOCK_NS_CAP_X)

Offset: 170h Size: 32 bits
Bank: 3

This read only register combined with the [1588 GPIO x Rising Edge Clock Seconds Capture Register \(1588_GPIO_RE_CLOCK_SEC_CAP_x\)](#) forms the GPIO rising edge timestamp capture.

Note: Values are only valid if the appropriate [1588 GPIO Rising Edge Interrupt \(1588_GPIO_RE_INT\[7:0\]\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) indicates that a timestamp is available.

Note: Unless the corresponding [Lock Enable GPIO Rising Edge \(LOCK_GPIO_RE\)](#) bit is set, a new capture may occur between reads of this register and the [1588 GPIO x Rising Edge Clock Seconds Capture Register \(1588_GPIO_RE_CLOCK_SEC_CAP_x\)](#). Software techniques are required to avoid reading intermediate values.

Note: Port and GPIO registers share a common address space. GPIO registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The GPIO accessed ("x") is set by the [GPIO Select \(GPIO_SEL\[2:0\]\)](#) field.

Note: All 8 GPIO register sets are available.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the timestamp upon the rising edge of a GPIO or upon a software commanded manual capture.	RO	00000000h

15.8.49 1588 GPIO X FALLING EDGE CLOCK SECONDS CAPTURE REGISTER (1588_GPIO_FE_CLOCK_SEC_CAP_X)

Offset: 178h Size: 32 bits
Bank: 3

This read only register combined with the [1588 GPIO x Falling Edge Clock NanoSeconds Capture Register \(1588_GPIO_FE_CLOCK_NS_CAP_x\)](#) forms the GPIO falling edge timestamp capture.

Note: Values are only valid if the appropriate [1588 GPIO Falling Edge Interrupt \(1588_GPIO_FE_INT\[7:0\]\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) indicates that a timestamp is available.

Note: Unless the corresponding [Lock Enable GPIO Falling Edge \(LOCK_GPIO_FE\)](#) bit is set, a new capture may occur between reads of this register and the [1588 GPIO x Falling Edge Clock NanoSeconds Capture Register \(1588_GPIO_FE_CLOCK_NS_CAP_x\)](#). Software techniques are required to avoid reading intermediate values.

Note: Port and GPIO registers share a common address space. GPIO registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The GPIO accessed ("x") is set by the [GPIO Select \(GPIO_SEL\[2:0\]\)](#) field.

Note: All 8 GPIO register sets are available.

Bits	Description	Type	Default
31:0	Timestamp Seconds (TS_SEC) This field contains the seconds portion of the timestamp upon the falling edge of a GPIO or upon a software commanded manual capture.	RO	00000000h

LAN9352

15.8.50 1588 GPIO X FALLING EDGE CLOCK NANOSECONDS CAPTURE REGISTER (1588_GPIO_FE_CLOCK_NS_CAP_X)

Offset: 17Ch Size: 32 bits
Bank: 3

This read only register combined with the [1588 GPIO x Falling Edge Clock Seconds Capture Register \(1588_GPIO_FE_CLOCK_SEC_CAP_x\)](#) forms the GPIO falling edge timestamp capture.

Note: Values are only valid if the appropriate [1588 GPIO Falling Edge Interrupt \(1588_GPIO_FE_INT\[7:0\]\)](#) in the [1588 Interrupt Status Register \(1588_INT_STS\)](#) indicates that a timestamp is available.

Note: Unless the corresponding [Lock Enable GPIO Falling Edge \(LOCK_GPIO_FE\)](#) bit is set, a new capture may occur between reads of this register and the [1588 GPIO x Falling Edge Clock Seconds Capture Register \(1588_GPIO_FE_CLOCK_SEC_CAP_x\)](#). Software techniques are required to avoid reading intermediate values.

Note: Port and GPIO registers share a common address space. GPIO registers are selected by the [Bank Select \(BANK_SEL\[2:0\]\)](#) in the [1588 Bank Port GPIO Select Register \(1588_BANK_PORT_GPIO_SEL\)](#). The GPIO accessed ("x") is set by the [GPIO Select \(GPIO_SEL\[2:0\]\)](#) field.

Note: All 8 GPIO register sets are available.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29:0	Timestamp NanoSeconds (TS_NS) This field contains the nanoseconds portion of the timestamp upon the falling edge of a GPIO or upon a software commanded manual capture.	RO	00000000h

16.0 GENERAL PURPOSE TIMER & FREE-RUNNING CLOCK

This chapter details the General Purpose Timer (GPT) and the Free-Running Clock.

16.1 General Purpose Timer

The device provides a 16-bit programmable General Purpose Timer that can be used to generate periodic system interrupts. The resolution of this timer is 100 μ s.

The GPT loads the [General Purpose Timer Count Register \(GPT_CNT\)](#) with the value in the [General Purpose Timer Pre-Load \(GPT_LOAD\)](#) field of the [General Purpose Timer Configuration Register \(GPT_CFG\)](#) when the [General Purpose Timer Enable \(TIMER_EN\)](#) bit of the [General Purpose Timer Configuration Register \(GPT_CFG\)](#) is asserted (1). On a chip-level reset or when the [General Purpose Timer Enable \(TIMER_EN\)](#) bit changes from asserted (1) to de-asserted (0), the [General Purpose Timer Pre-Load \(GPT_LOAD\)](#) field is initialized to FFFFh. The [General Purpose Timer Count Register \(GPT_CNT\)](#) is also initialized to FFFFh on reset.

Once enabled, the GPT counts down until it reaches 0000h. At 0000h, the counter wraps around to FFFFh, asserts the [GP Timer \(GPT_INT\)](#) interrupt status bit in the [Interrupt Status Register \(INT_STS\)](#), asserts the IRQ interrupt (if [GP Timer Interrupt Enable \(GPT_INT_EN\)](#) is set in the [Interrupt Enable Register \(INT_EN\)](#)) and continues counting. [GP Timer \(GPT_INT\)](#) is a sticky bit. Once this bit is asserted, it can only be cleared by writing a 1 to the bit. Refer to [Section 8.2.7, "General Purpose Timer Interrupt," on page 77](#) for additional information on the GPT interrupt.

Software can write a pre-load value into the [General Purpose Timer Pre-Load \(GPT_LOAD\)](#) field at any time (e.g., before or after the [General Purpose Timer Enable \(TIMER_EN\)](#) bit is asserted). The [General Purpose Timer Count Register \(GPT_CNT\)](#) will immediately be set to the new value and continue to count down (if enabled) from that value.

16.2 Free-Running Clock

The Free-Running Clock (FRC) is a simple 32-bit up-counter that operates from a fixed 25 MHz clock. The current FRC value can be read via the [Free Running 25MHz Counter Register \(FREE_RUN\)](#). On assertion of a chip-level reset, this counter is cleared to zero. On de-assertion of a reset, the counter is incremented once for every 25 MHz clock cycle. When the maximum count has been reached, the counter rolls over to zeros. The FRC does not generate interrupts.

Note: The free running counter can take up to 160 ns to clear after a reset event.

16.3 General Purpose Timer and Free-Running Clock Registers

This section details the directly addressable general purpose timer and free-running clock related System CSRs. For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map," on page 33](#).

TABLE 16-1: MISCELLANEOUS REGISTERS

ADDRESS	Register Name (SYMBOL)
08Ch	General Purpose Timer Configuration Register (GPT_CFG)
090h	General Purpose Timer Count Register (GPT_CNT)
09Ch	Free Running 25MHz Counter Register (FREE_RUN)

16.3.1 GENERAL PURPOSE TIMER CONFIGURATION REGISTER (GPT_CFG)

Offset: 08Ch Size: 32 bits

This read/write register configures the device's General Purpose Timer (GPT). The GPT can be configured to generate host interrupts at the interval defined in this register. The current value of the GPT can be monitored via the [General Purpose Timer Count Register \(GPT_CNT\)](#). Refer to [Section 16.1, "General Purpose Timer," on page 559](#) for additional information.

Bits	Description	Type	Default
31:30	RESERVED	RO	-
29	General Purpose Timer Enable (TIMER_EN) This bit enables the GPT. When set, the GPT enters the run state. When cleared, the GPT is halted. On the 1 to 0 transition of this bit, the GPT_LOAD field of this register will be preset to FFFFh. 0: GPT Disabled 1: GPT Enabled	R/W	0b
28:16	RESERVED	RO	-
15:0	General Purpose Timer Pre-Load (GPT_LOAD) This value is pre-loaded into the GPT. This is the starting value of the GPT. The timer will begin decrementing from this value when enabled.	R/W	FFFFh

16.3.2 GENERAL PURPOSE TIMER COUNT REGISTER (GPT_CNT)

Offset: 090h Size: 32 bits

This read-only register reflects the current general purpose timer (GPT) value. The register should be used in conjunction with the [General Purpose Timer Configuration Register \(GPT_CFG\)](#) to configure and monitor the GPT. Refer to [Section 16.1, "General Purpose Timer," on page 559](#) for additional information.

Bits	Description	Type	Default
31:16	RESERVED	RO	-
15:0	General Purpose Timer Current Count (GPT_CNT) This 16-bit field represents the current value of the GPT.	RO	FFFFh

16.3.3 FREE RUNNING 25MHZ COUNTER REGISTER (FREE_RUN)

Offset: 09Ch Size: 32 bits

This read-only register reflects the current value of the free-running 25MHz counter. Refer to [Section 16.2, "Free-Running Clock," on page 559](#) for additional information.

Bits	Description	Type	Default
31:0	Free Running Counter (FR_CNT) This field reflects the current value of the free-running 32-bit counter. At reset, the counter starts at zero and is incremented by one every 25 MHz cycle. When the maximum count has been reached, the counter will rollover to zero and continue counting. Note: The free running counter can take up to 160nS to clear after a reset event.	RO	00000000h

17.0 GPIO/LED CONTROLLER

17.1 Functional Overview

The GPIO/LED Controller provides 8 configurable general purpose input/output pins, **GPIO[7:0]**. These pins can be individually configured to function as inputs, push-pull outputs or open drain outputs and each is capable of interrupt generation with configurable polarity. Alternatively, 6 GPIO pins can be configured as LED outputs, enabling these pins to drive Ethernet status LEDs for external indication of various attributes of the ports. All GPIOs also provide extended 1588 functionality. Refer to [Section 15.5, "1588 GPIOs," on page 492](#) for additional details.

GPIO and LED functionality is configured via the GPIO/LED System Control and Status Registers (CSRs). These registers are defined in [Section 17.4, "GPIO/LED Registers," on page 566](#).

17.2 GPIO Operation

The GPIO controller is comprised of 8 programmable input/output pins. These pins are individually configurable via the GPIO CSRs. On application of a chip-level reset:

- All GPIOs are set as inputs ([GPIO Direction 7-0 \(GPIO_DIR\[7:0\]\)](#) cleared in [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#))
- All GPIO interrupts are disabled ([GPIO Interrupt Enable\[7:0\] \(GPIO\[7:0\]_INT_EN\)](#) cleared in [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#))
- All GPIO interrupts are configured to low logic level triggering ([GPIO Interrupt/1588 Polarity 7-0 \(GPIO_POL\[7:0\]\)](#) cleared in [General Purpose I/O Configuration Register \(GPIO_CFG\)](#))

Note: [GPIO\[5:0\]](#) may be configured as LED outputs by default, dependent on the [LED_en_strap\[5:0\]](#) configuration straps. Refer to [Section 17.3, "LED Operation"](#) for additional information.

The direction and buffer type of all GPIOs are configured via the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#) and [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#). The direction of each GPIO, input or output, should be configured first via its respective [GPIO Direction 7-0 \(GPIO_DIR\[7:0\]\)](#) bit in the [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#). When configured as an output, the output buffer type for each GPIO is selected by the [GPIO Buffer Type 7-0 \(GPIOBUF\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#). Push/pull and open-drain output buffers are supported for each GPIO. When functioning as an open-drain driver, the GPIO output pin is driven low when the corresponding [GPIO Data 7-0 \(GPIO\[7:0\]\)](#) bit in the [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#) is cleared to 0 and is not driven when set to 1.

When a GPIO is enabled as a push/pull output, the value output to the GPIO pin is set via the corresponding [GPIO Data 7-0 \(GPIO\[7:0\]\)](#) bit in the [General Purpose I/O Data & Direction Register \(GPIO_DATA_DIR\)](#). For GPIOs configured as inputs, the corresponding [GPIO Data 7-0 \(GPIO\[7:0\]\)](#) bit reflects the current state of the GPIO input.

In GPIO mode, the input buffers are disabled when the pin is set to an output and the pull-ups are normally enabled.

Note: Upon reset, GPIOs that were outputs may generate an active interrupt status as the system settles - typically when a low GPIO pin slowly rises due to the internal pull-up. The interrupt status bits within the [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#) should be cleared as part of the device initialization software routine.

17.2.1 GPIO INTERRUPTS

Each GPIO provides the ability to trigger a unique GPIO interrupt in the [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#). Reading the [GPIO Interrupt\[7:0\] \(GPIO\[7:0\]_INT\)](#) bits of this register provides the current status of the corresponding interrupt and each interrupt is enabled by setting the corresponding [GPIO Interrupt Enable\[7:0\] \(GPIO\[7:0\]_INT_EN\)](#) bit. The GPIO/LED Controller aggregates the enabled interrupt values into an internal signal that is sent to the System Interrupt Controller and is reflected via the [Interrupt Status Register \(INT_STS\) GPIO Interrupt Event \(GPIO\)](#) bit. For more information on interrupts, refer to [Section 8.0, "System Interrupts," on page 73](#).

As interrupts, GPIO inputs are level sensitive and must be active for greater than 40 ns to be recognized.

17.2.1.1 GPIO Interrupt Polarity

The interrupt polarity can be set for each individual GPIO via the [GPIO Interrupt/1588 Polarity 7-0 \(GPIO_POL\[7:0\]\)](#) bits in the [General Purpose I/O Configuration Register \(GPIO_CFG\)](#). When set, a high logic level on the GPIO pin will set the corresponding interrupt bit in the [General Purpose I/O Interrupt Status and Enable Register \(GPIO_INT_STS_EN\)](#). When cleared, a low logic level on the GPIO pin will set the corresponding interrupt bit.

17.3 LED Operation

GPIO[5:0] can be individually selected to function as a LED. These pins are configured as LED outputs by setting the corresponding [LED Enable 5-0 \(LED_EN\[5:0\]\)](#) bit in the [LED Configuration Register \(LED_CFG\)](#). When configured as an LED, the pin is either a push-pull or open-drain / open-source output and the GPIO related input buffer and pull-up are disabled. The default configuration, including polarity, is determined by input straps or EEPROM entries. Refer to [Section 7.0, "Configuration Straps," on page 60](#) for additional information.

The functions associated with each LED pin are configurable via the [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) bits of the [LED Configuration Register \(LED_CFG\)](#). These bits allow the configuration of each LED pin to indicate various port related functions. The behaviors of each LED for each [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) configuration are described in the following tables. Detailed definitions for each LED indication type are provided in [Section 17.3.1](#) and [Section 17.3.2](#).

The default values of the [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) and [LED Enable 5-0 \(LED_EN\[5:0\]\)](#) bits of the [LED Configuration Register \(LED_CFG\)](#) are determined by the [LED_fun_strap\[2:0\]](#) and [LED_en_strap\[5:0\]](#) configuration straps. For more information on the [LED Configuration Register \(LED_CFG\)](#) and its related straps, refer to [Section 17.4.1, "LED Configuration Register \(LED_CFG\)," on page 567](#).

All LED outputs may be disabled by setting the [LED Disable \(LED_DIS\)](#) bit in the [Power Management Control Register \(PMT_CTRL\)](#). Open-drain / open-source LEDs are un-driven. Push-pull LEDs are still driven but are set to their inactive state.

TABLE 17-1: LED OPERATION AS A FUNCTION OF LED_FUN[2:0] = 000B - 011B

	000b	001b	010b	011b
LED5 (GPIO5)	Link / Activity Port 2	100Link / Activity Port 2	TX Port 0	Activity Port 2
LED4 (GPIO4)	Full-duplex / Collision Port 2	Full-duplex / Collision Port 2	Link / Activity Port 2	Link Port 2
LED3 (GPIO3)	Speed Port 2	10Link / Activity Port 2	Speed Port 2	Speed Port 2
LED2 (GPIO2)	Link / Activity Port 1	100Link / Activity Port 1	RX Port 0	Activity Port 1
LED1 (GPIO1)	Full-duplex / Collision Port 1	Full-duplex / Collision Port 1	Link / Activity Port 1	Link Port 1
LED0 (GPIO0)	Speed Port 1	10Link / Activity Port 1	Speed Port 1	Speed Port 1

TABLE 17-2: LED OPERATION AS A FUNCTION OF LED_FUN[2:0] = 100B - 111B

	100b	101b	110b	111b
LED5 (GPIO5)	Activity Port 2	Activity Port 2	Reserved	TX_EN Port 0
LED4 (GPIO4)	Link Port 2	10Link Port 2		TX_EN Port 2
LED3 (GPIO3)	Full-duplex / Collision Port 2	100Link Port 2		RX_DV Port 2
LED2 (GPIO2)	Activity Port 1	Activity Port 1		RX_DV Port 0
LED1 (GPIO1)	Link Port 1	10Link Port 1		TX_EN Port 1
LED0 (GPIO0)	Full-duplex / Collision Port 1	100Link Port 1		RX_DV Port 1

The various LED indication functions listed in the previous tables are described in the following sections.

17.3.1 LED FUNCTION DEFINITIONS WHEN LED_FUN[2:0] = 000B - 101B

The following LED rules apply when [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) is 000b through 101b:

- “Active” is defined as the pin being driven to the opposite value latched at reset on the related hard-straps. The LED polarity cannot be modified via soft-straps.
- “Inactive” is defined as the pin not being driven.
- The input buffers and pull-ups are disabled on the shared GPIO/LED pins.

The following LED function definitions apply when [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) is 000b through 101b:

- **TX** - The signal is pulsed active for 80 ms to indicate activity from the Switch Fabric to the external MII pins. This signal is then made inactive for a minimum of 80 ms, after which the process will repeat if TX activity is again detected.

Note: Link indication does not affect this function.

- **RX** - The signal is pulsed active for 80 ms to indicate activity from the external MII pins to the Switch Fabric. This signal is then made inactive for a minimum of 80 ms, after which the process will repeat if RX activity is again detected.

Note: Link indication does not affect this function.

- **Activity** - The signal is pulsed active for 80ms to indicate transmit or receive activity on the port. The signal is then made inactive for a minimum of 80mS, after which the process will repeat if RX or TX activity is again detected.

Note: The idle condition is *inactive* in contrast to that of the Link / Activity function.

Note: The signal will be held inactive if the PHY does not have a valid link.

- **Link** - A steady active output indicates that the port has a valid link (10Mbps or 100Mbps), while a steady inactive output indicates no link on the port.
- **Link / Activity** - A steady active output indicates that the port has a valid link, while a steady inactive output indicates no link on the port. When the port has a valid link, the signal is pulsed inactive for 80 ms to indicate transmit or receive activity on the port. The signal is then made active for a minimum of 80 ms, after which the process will repeat if RX or TX activity is again detected.
- **100Link** - A steady active output indicates the port has a valid link and the speed is 100 Mbps. The signal will be held inactive if the port does not have a valid link or the speed is not 100 Mbps.

- **100Link / Activity** - A steady active output indicates the port has a valid link and the speed is 100 Mbps. The signal is pulsed inactive for 80 ms to indicate TX or RX activity on the port. The signal is then driven active for a minimum of 80 ms, after which the process will repeat if RX or TX activity is again detected. The signal will be held inactive if the port does not have a valid link or the speed is not 100 Mbps.
- **10Link** - A steady active output indicates the port has a valid link and the speed is 10 Mbps. This signal will be held inactive if the port does not have a valid link or the speed is not 10 Mbps.
- **10Link / Activity** - A steady active output indicates the port has a valid link and the speed is 10 Mbps. The signal is pulsed inactive for 80 ms to indicate transmit or receive activity on the port. The signal is then driven active for a minimum of 80 ms, after which the process will repeat if RX or TX activity is again detected. This signal will be held inactive if the port does not have a valid link or the speed is not 10 Mbps.
- **Full-duplex / Collision** - A steady active output indicates the port is in full-duplex mode. In half-duplex mode, the signal is pulsed active for 80 ms to indicate a network collision. The signal is then made inactive for a minimum of 80 ms, after which the process will repeat if another collision is detected. The signal will be held inactive if the port does not have a valid link.
- **Speed** - A steady active output indicates a valid link with a speed of 100 Mbps. A steady inactive output indicates a speed of 10 Mbps. The signal will be held inactive if the port does not have a valid link.

17.3.2 LED FUNCTION DEFINITIONS WHEN LED_FUN[2:0] = 111B

When [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) is 111b, the following LED rules apply:

- The LED pins are push-pull drivers.
- The LED pin is driven high when the function signal is high and is driven low when the function signal is low.
- The input buffers and pull-ups are disabled on the shared GPIO/LED pins.

When [LED Function 2-0 \(LED_FUN\[2:0\]\)](#) is 111b, the following LED function definitions apply:

- **TX_EN** - Non-stretched TX_EN signal from the Switch Fabric.
Note: Link indication does not affect this function.
- **RX_DV** - Non-stretched RX_DV signal to the Switch Fabric.
Note: Link indication does not affect this function.

17.4 GPIO/LED Registers

This section details the directly addressable General Purpose I/O (GPIO) and LED related System CSRs. For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map," on page 33](#).

TABLE 17-3: GPIO/LED REGISTERS

ADDRESS	REGISTER NAME (SYMBOL)
1BCh	LED Configuration Register (LED_CFG)
1E0h	General Purpose I/O Configuration Register (GPIO_CFG)
1E4h	General Purpose I/O Data & Direction Register (GPIO_DATA_DIR)
1E8h	General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN)

17.4.1 LED CONFIGURATION REGISTER (LED_CFG)

Offset: 1BCh Size: 32 bits

This read/write register configures the **GPIO[5:0]** pins as LED pins and sets their functionality.

Bits	Description	Type	Default
31:11	RESERVED	RO	-
10:8	LED Function 2-0 (LED_FUN[2:0]) These bits control the function associated with each LED pin as shown in Section 17.3, "LED Operation," on page 564 . Note: In order for these assignments to be valid, the particular pin must be enabled as an LED output pin via the LED_EN bits of this register.	R/W	Note 1
7:6	RESERVED	RO	-
5:0	LED Enable 5-0 (LED_EN[5:0]) This field toggles the functionality of the GPIO[5:0] pins between GPIO and LED. 0: Enables the associated pin as a GPIO signal 1: Enables the associated pin as a LED output When configured as LED outputs, the pins are either push-pull or open-drain/open-source outputs and the pull-ups and input buffers are disabled. Push-pull is selected when LED_FUN[2:0] = 111b, otherwise, they are open-drain/open-source. When open-drain/open-source, the polarity of the pins depends upon the strap value sampled at reset. If a high is sampled at reset, then this signal is active low. Note: The polarity is determined by the strap value sampled on reset (a hard-strap) and not the soft-strap value (of the shared strap) set via EEPROM. When configured as a GPIO output, the pins are configured per the General Purpose I/O Configuration Register (GPIO_CFG) and the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) . The polarity of the pins does not depend upon the strap value sampled at reset.	R/W	Note 2

Note 1: The default value of this field is determined by the configuration strap [LED_fun_strap\[2:0\]](#).

Note 2: The default value of this field is determined by the configuration strap [LED_en_strap\[5:0\]](#).

17.4.2 GENERAL PURPOSE I/O CONFIGURATION REGISTER (GPIO_CFG)

Offset: 1E0h Size: 32 bits

This read/write register configures the GPIO input and output pins. The polarity of the GPIO pins is configured here as well as the IEEE 1588 timestamping and clock compare event output properties. Refer to [Section 15.5, "1588 GPIOs," on page 492](#) for additional 1588 information.

Bits	Description	Type	Default
31:24	1588 GPIO Channel Select 7-0 (GPIO_CH_SEL[7:0]) These bits select the 1588 channel to be output on the corresponding GPIO[7:0]. Refer to Section 15.5, "1588 GPIOs," on page 492 for additional information. 0: Sets 1588 channel A as the output for the corresponding GPIO pin 1: Sets 1588 channel B as the output for the corresponding GPIO pin	R/W	00h
23:16	GPIO Interrupt/1588 Polarity 7-0 (GPIO_POL[7:0]) These bits set the interrupt input polarity and 1588 clock event output polarity of the 8 GPIO pins. The configured level (high/low) will set the corresponding GPIO_INT bit in the General Purpose I/O Interrupt Status and Enable Register (GPIO_INT_STS_EN) . 1588 clock events will be output active at the configured level (high/low). These bits also determine the polarity of the GPIO 1588 Timer Interrupt Clear inputs. Refer to Section 15.5, "1588 GPIOs," on page 492 for additional information. 0: Sets low logic level trigger on corresponding GPIO pin 1: Sets high logic level trigger on corresponding GPIO pin	R/W	00h
15:8	1588 GPIO Output Enable 7-0 (1588_GPIO_OE[7:0]) These bits configure the 8 GPIO pins to output 1588 clock compare events. 0: Disables the output of 1588 clock compare events 1: Enables the output of 1588 clock compare events Note: These bits override the direction bits in the General Purpose I/O Data & Direction Register (GPIO_DATA_DIR) register. However, the GPIO Buffer Type 7-0 (GPIOBUF[7:0]) in the General Purpose I/O Configuration Register (GPIO_CFG) is not overridden.	R/W	00h

Bits	Description	Type	Default															
7:0	<p>GPIO Buffer Type 7-0 (GPIOBUF[7:0]) This field sets the buffer types of the 8 GPIO pins.</p> <p>0: Corresponding GPIO pin configured as an open-drain driver 1: Corresponding GPIO pin configured as a push/pull driver</p> <p>As an open-drain driver, the output pin is driven low when the corresponding data register is cleared, and is not driven when the corresponding data register is set.</p> <p>As an open-drain driver used for 1588 Clock Events, the corresponding GPIO_POL_x bit determines when the corresponding pin is driven per the following table:</p> <table><tr><th>GPIOx Clock Event Polarity</th><th>1588 Clock Event</th><th>Pin State</th></tr><tr><td>0</td><td>no</td><td>not driven</td></tr><tr><td>0</td><td>yes</td><td>driven low</td></tr><tr><td>1</td><td>no</td><td>driven low</td></tr><tr><td>1</td><td>yes</td><td>not driven</td></tr></table>	GPIOx Clock Event Polarity	1588 Clock Event	Pin State	0	no	not driven	0	yes	driven low	1	no	driven low	1	yes	not driven	R/W	00h
GPIOx Clock Event Polarity	1588 Clock Event	Pin State																
0	no	not driven																
0	yes	driven low																
1	no	driven low																
1	yes	not driven																

17.4.3 GENERAL PURPOSE I/O DATA & DIRECTION REGISTER (GPIO_DATA_DIR)

Offset: 1E4h Size: 32 bits

This read/write register configures the direction of the GPIO pins and contains the GPIO input and output data bits.

Bits	Description	Type	Default
31:24	RESERVED	RO	-
23:16	GPIO Direction 7-0 (GPIODIR[7:0]) These bits set the input/output direction of the 8 GPIO pins. 0: GPIO pin is configured as an input 1: GPIO pin is configured as an output	R/W	00h
15:8	RESERVED	RO	-
7:0	GPIO Data 7-0 (GPIOD[7:0]) When a GPIO pin is enabled as an output, the value written to this field is output on the corresponding GPIO pin. Upon a read, the value returned depends on the current direction of the pin. If the pin is an input, the data reflects the current state of the corresponding GPIO pin. If the pin is an output, the data is the value that was last written into this register. The pin direction is determined by the GPIODIR bits of this register and the 1588_GPIO_OE bits in the General Purpose I/O Configuration Register (GPIO_CFG) .	R/W	00h

17.4.4 GENERAL PURPOSE I/O INTERRUPT STATUS AND ENABLE REGISTER (GPIO_INT_STS_EN)

Offset: 1E8h Size: 32 bits

This read/write register contains the GPIO interrupt status bits.

Writing a 1 to any of the interrupt status bits acknowledges and clears the interrupt. If enabled, these interrupt bits are cascaded into the [GPIO Interrupt Event \(GPIO\)](#) bit of the [Interrupt Status Register \(INT_STS\)](#). Writing a 1 to any of the interrupt enable bits will enable the corresponding interrupt as a source. Status bits will still reflect the status of the interrupt source regardless of whether the source is enabled as an interrupt in this register. The [GPIO Interrupt Event Enable \(GPIO_EN\)](#) bit of the [Interrupt Enable Register \(INT_EN\)](#) must also be set in order for an actual system level interrupt to occur. Refer to [Section 8.0, "System Interrupts,"](#) on page 73 for additional information.

BITS	DESCRIPTION	TYPE	DEFAULT
31:24	RESERVED	RO	-
23:16	GPIO Interrupt Enable[7:0] (GPIO[7:0]_INT_EN) When set, these bits enable the corresponding GPIO interrupt. Note: The GPIO interrupts must also be enabled via the GPIO Interrupt Event Enable (GPIO_EN) bit of the Interrupt Enable Register (INT_EN) in order to cause the interrupt pin (IRQ) to be asserted.	R/W	00h
15:8	RESERVED	RO	-
7:0	GPIO Interrupt[7:0] (GPIO[7:0]_INT) These signals reflect the interrupt status as generated by the GPIOs. These interrupts are configured through the General Purpose I/O Configuration Register (GPIO_CFG) . Note: As GPIO interrupts, GPIO inputs are level sensitive and must be active greater than 40 ns to be recognized as interrupt inputs.	R/WC	00h

18.0 MISCELLANEOUS

This chapter describes miscellaneous functions and registers that are present in the device.

18.1 Miscellaneous System Configuration & Status Registers

This section details the remainder of the directly addressable System CSRs. These registers allow for monitoring and configuration of various device functions such as the Chip ID/revision, byte order testing, and hardware configuration.

For an overview of the entire directly addressable register map, refer to [Section 5.0, "Register Map," on page 33](#).

TABLE 18-1: MISCELLANEOUS REGISTERS

ADDRESS	Register Name (SYMBOL)
050h	Chip ID and Revision (ID_REV)
064h	Byte Order Test Register (BYTE_TEST)
074h	Hardware Configuration Register (HW_CFG)

18.1.1 CHIP ID AND REVISION (ID_REV)

Offset: 050h Size: 32 bits

This read-only register contains the ID and Revision fields for the device.

Bits	Description	Type	Default
31:16	Chip ID This field indicates the chip ID.	RO	9352
15:0	Chip Revision This field indicates the design revision.	RO	Note 1

Note 1: Default value is dependent on device revision.

18.1.2 BYTE ORDER TEST REGISTER (BYTE_TEST)

Offset: 064h Size: 32 bits

This read-only register can be used to determine the byte ordering of the current configuration. Byte ordering is a function of the host data bus width and endianness. Refer to [Section 9.0, "Host Bus Interface," on page 86](#) for additional information on byte ordering.

The BYTE_TEST register can optionally be used as a dummy read register when assuring minimum write-to-read or read-to-read timing. Refer to [Section 9.0, "Host Bus Interface," on page 86](#) for additional information.

For host interfaces that are disabled during the reset state, the BYTE_TEST register can be used to determine when the device has exited the reset state.

Note: This register can be read while the device is in the reset or not ready / power savings states without leaving the host interface in an intermediate state. If the host interface is in a reset state, returned data may be invalid. However, during reset, the returned data will not match the normal valid data pattern.

Note: It is not necessary to read all four bytes of this register. DWORD access rules do not apply to this register.

Bits	Description	Type	Default
31:0	Byte Test (BYTE_TEST) This field reflects the current byte ordering	RO	87654321h

18.1.3 HARDWARE CONFIGURATION REGISTER (HW_CFG)

Offset: 074h Size: 32 bits

This register allows the configuration of various hardware features including TX/RX FIFO sizes and Host MAC transmit threshold properties. A detailed explanation of the allowable settings for FIFO memory allocation can be found in [Section 11.10.3, "FIFO Memory Allocation Configuration," on page 171](#).

Note: This register can be read while the device is in the reset or not ready / power savings states without leaving the host interface in an intermediate state. If the host interface is in a reset state, returned data may be invalid.

Note: It is not necessary to read all four BYTES of this register. DWORD access rules do not apply to this register.

Bits	Description	Type	Default
31:28	RESERVED	RO	-
27	Device Ready (READY) When set, this bit indicates that the device is ready to be accessed. Upon power-up, RST# reset, return from power savings states, Host MAC module level reset or digital reset, the host processor may interrogate this field as an indication that the device has stabilized and is fully active. This rising edge of this bit will assert the Device Ready (READY) bit in the Interrupt Status Register (INT_STS) and can cause an interrupt if enabled. Note: With the exception of the HW_CFG, PMT_CTRL, BYTE_TEST, and RESET_CTL registers, read access to any internal resources is forbidden while the READY bit is cleared. Writes to any address are invalid until this bit is set. Note: This bit is identical to bit 0 of the Power Management Control Register (PMT_CTRL) .	RO	0b
26	AMDIX_EN Strap State Port B This bit reflects the state of the auto_mdix_strap_2 strap that connects to the PHY. The strap value is loaded with the level of the auto_mdix_strap_2 during reset and can be re-written by the EEPROM Loader. The strap value can be overridden by bit 15 and 13 of the Port B PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) PHY Special Control/Status Indication Register.	RO	Note 2
25	AMDIX_EN Strap State Port A This bit reflects the state of the auto_mdix_strap_1 strap that connects to the PHY. The strap value is loaded with the level of the auto_mdix_strap_1 during reset and can be re-written by the EEPROM Loader. The strap value can be overridden by bit 15 and 13 of the Port A PHY x Special Control/Status Indication Register (PHY_SPECIAL_CONTROL_STAT_IND_x) .	RO	Note 3
24:22	RESERVED	RO	-
21	RESERVED - This bit must be written with 0b for proper operation.	R/W	0b
20	Must Be One (MBO). This bit must be set to '1' for normal device operation.	R/W	0b

Bits	Description	Type	Default
19:16	TX FIFO Size (TX_FIF_SZ) This field sets the size of the TX FIFOs in 1KB values to a maximum of 14KB. The TX Status FIFO consumes 512 bytes of the space allocated by TX_FIF_SZ, and the TX Data FIFO consumes the remaining space specified by TX_FIF_SZ. The minimum size of the TX FIFOs is 2KB (TX Data FIFO and Status FIFO combined). The TX Data FIFO is used for both TX data and TX commands. The RX Status and Data FIFOs consume the remaining space, which is equal to 16KB minus TX_FIF_SZ. See section Section 11.10.3, "FIFO Memory Allocation Configuration," on page 171 for more information.	R/W	5h
15:14	RESERVED	RO	-
13:12	RESERVED - This field must be written with 00b for proper operation.	R/W	00b
11:0	RESERVED	RO	-

Note 2: The default value of this field is determined by the configuration strap [auto_mdix_strap_2](#). See [Section 6.3, "Power Management," on page 49](#) for more information.

Note 3: The default value of this field is determined by the configuration strap [auto_mdix_strap_1](#). See [Section 6.3, "Power Management," on page 49](#) for more information.

19.0 JTAG

19.1 JTAG

A IEEE 1149.1 compliant TAP Controller supports boundary scan and various test modes.

The device includes an integrated JTAG boundary-scan test port for board-level testing. The interface consists of four pins (**TDO**, **TDI**, **TCK** and **TMS**) and includes a state machine, data register array, and an instruction register. The JTAG pins are described in [Table 3-11, “JTAG Pin Descriptions,” on page 29](#). The JTAG interface conforms to the IEEE Standard 1149.1 - 2001 *Standard Test Access Port (TAP) and Boundary-Scan Architecture*.

All input and output data is synchronous to the **TCK** test clock input. TAP input signals **TMS** and **TDI** are clocked into the test logic on the rising edge of **TCK**, while the output signal **TDO** is clocked on the falling edge.

JTAG pins are multiplexed with the GPIO/LED and EEPROM pins. The JTAG functionality is selected when the **TEST-MODE** pin is asserted.

The implemented IEEE 1149.1 instructions and their op codes are shown in [Table 19-1](#).

TABLE 19-1: IEEE 1149.1 OP CODES

INSTRUCTION	OP CODE	COMMENT
BYPASS 0	16'h0000	Mandatory Instruction
BYPASS 1	16'hFFFF	Mandatory Instruction
SAMPLE/PRELOAD	16'hFFF8	Mandatory Instruction
EXTEST	16'hFFE8	Mandatory Instruction
CLAMP	16'hFFE8	Optional Instruction
ID_CODE	16'hFFFE	Optional Instruction
HIGHZ	16'hFFCF	Optional Instruction
INT_DR_SEL	16'hFFFD	Private Instruction

Note: The JTAG device ID is 00111445h

Note: All digital I/O pins support IEEE 1149.1 operation. Analog pins and the **OSCI** / **OSCO** pins do not support IEEE 1149.1 operation.

LAN9352

19.1.1 JTAG TIMING REQUIREMENTS

This section specifies the JTAG timing of the device.

FIGURE 19-1: JTAG TIMING

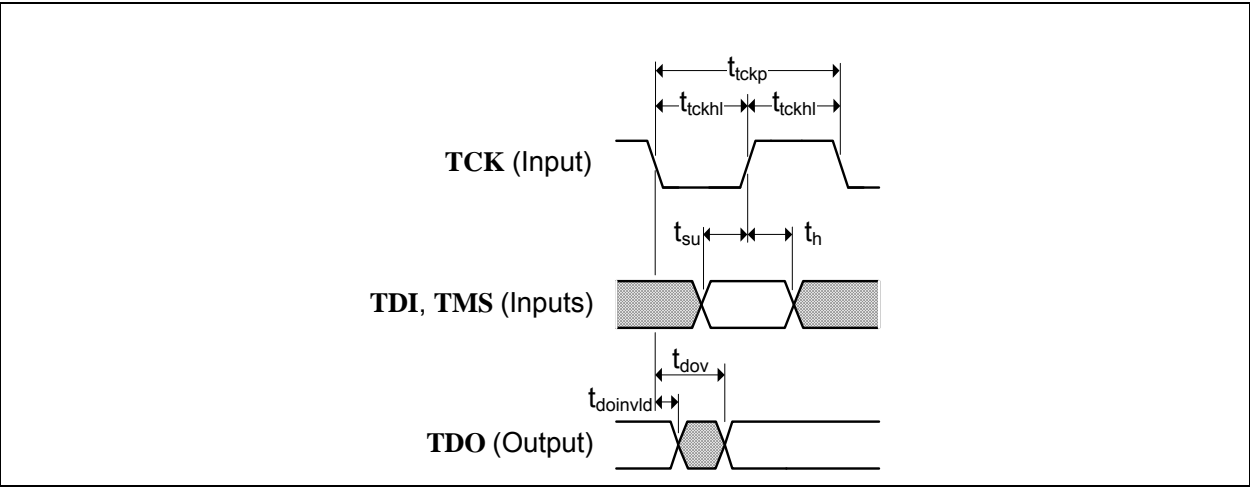


TABLE 19-2: JTAG TIMING VALUES

Symbol	Description	Min	Max	Units	Notes
t_{tckp}	TCK clock period	40		ns	
t_{tckhl}	TCK clock high/low time	$t_{tckp} \cdot 0.4$	$t_{tckp} \cdot 0.6$	ns	
t_{su}	TDI, TMS setup to TCK rising edge	5		ns	
t_h	TDI, TMS hold from TCK rising edge	5		ns	
t_{dov}	TDO output valid from TCK falling edge		15	ns	
$t_{doinvld}$	TDO output invalid from TCK falling edge	0		ns	

Note: Timing values are with respect to an equivalent test load of 25 pF.

20.0 OPERATIONAL CHARACTERISTICS

20.1 Absolute Maximum Ratings*

Supply Voltage (VDD12TX1 , VDD12TX2 , OSCVDD12 , VDDCR) (Note 1)	0 V to +1.5 V
Supply Voltage (VDD33TXRX1 , VDD33TXRX2 , VDD33BIAS , VDD33 , VDDIO) (Note 1)	0 V to +3.6 V
Ethernet Magnetics Supply Voltage	-0.5 V to +3.6 V
Positive voltage on input signal pins, with respect to ground (Note 2)	VDDIO + 2.0 V
Negative voltage on input signal pins, with respect to ground (Note 3)	-0.5 V
Positive voltage on OSCI , with respect to ground	+3.6 V
Storage Temperature	-55°C to +150°C
Junction Temperature	+150°C
Lead Temperature Range	Refer to JEDEC Spec. J-STD-020
HBM ESD Performance	JEDEC Class 3A

Note 1: When powering this device from laboratory or system power supplies, it is important that the absolute maximum ratings not be exceeded or device failure can result. Some power supplies exhibit voltage spikes on their outputs when AC power is switched on or off. In addition, voltage transients on the AC power line may appear on the DC output. If this possibility exists, it is suggested to use a clamp circuit.

Note 2: This rating does not apply to the following pins: [OSCI](#), [RBIAS](#)

Note 3: This rating does not apply to the following pins: [RBIAS](#)

*Stresses exceeding those listed in this section could cause permanent damage to the device. This is a stress rating only. Exposure to absolute maximum rating conditions for extended periods may affect device reliability. Functional operation of the device at any condition exceeding those indicated in [Section 20.2, "Operating Conditions"](#), [Section 20.5, "DC Specifications"](#), or any other applicable section of this specification is not implied. Note, device signals are NOT 5 volt tolerant.

20.2 Operating Conditions**

Supply Voltage (VDD12TX1 , VDD12TX2 , OSCVDD12 , VDDCR)	+1.14 V to +1.26 V
Analog Port Supply Voltage (VDD33TXRX1 , VDD33TXRX2 , VDD33BIAS , VDD33)	+3.0 V to +3.6 V
I/O Supply Voltage (VDDIO) (Note 1)	+1.62 V to +3.6 V
Ethernet Magnetics Supply Voltage	+2.25 V to +3.6 V
Ambient Operating Temperature in Still Air (T _A)	Note 4

Note 4: 0°C to +70°C for commercial version, -40°C to +85°C for industrial version.

Proper operation of the device is guaranteed only within the ranges specified in this section. After the device has completed power-up, **VDDIO and the magnetics power supply must maintain their voltage level with ±10%. Varying the voltage greater than ±10% after the device has completed power-up can cause errors in device operation.

Note: Do not drive input signals without power supplied to the device.

20.3 Package Thermal Specifications

TABLE 20-1: 72-PIN QFN PACKAGE THERMAL PARAMETERS

Parameter	Symbol	Value	Units	Comments
Thermal Resistance Junction to Ambient	Θ_{JA}	20.5	°C/W	Measured in still air
Thermal Resistance Junction to Bottom of Case	Ψ_{JT}	0.1	°C/W	Measured in still air
Thermal Resistance Junction to Top of Case	Θ_{JC}	1.6	°C/W	Airflow 1 m/s

TABLE 20-2: 80-PIN TQFP-EP PACKAGE THERMAL PARAMETERS

Parameter	Symbol	Value	Units	Comments
Thermal Resistance Junction to Ambient	Θ_{JA}	29.0	°C/W	Measured in still air
Thermal Resistance Junction to Bottom of Case	Ψ_{JT}	0.3	°C/W	Measured in still air
Thermal Resistance Junction to Top of Case	Θ_{JC}	12.3	°C/W	Airflow 1 m/s

Note: Thermal parameters are measured or estimated for devices in a multi-layer 2S2P PCB per JESDN51.

TABLE 20-3: MAXIMUM POWER DISSIPATION

Mode	Maximum Power (mW)
Internal Regulator Disabled, 2.5 V Ethernet Magnetics	832
Internal Regulator Disabled, 3.3 V Ethernet Magnetics	1012
Internal Regulator Enabled, 2.5 V Ethernet Magnetics	1068
Internal Regulator Enabled, 3.3 V Ethernet Magnetics	1255

20.4 Current Consumption and Power Consumption

This section details the device's typical supply current consumption and power dissipation for 10BASE-T, 100BASE-TX and power management modes of operation with the internal regulator enabled and disabled.

Note: Each mode in the current consumption and power dissipation tables assumes all PHYs are in the corresponding mode of operation.

20.4.1 INTERNAL REGULATOR DISABLED

TABLE 20-4: CURRENT CONSUMPTION AND POWER DISSIPATION (REGS. DISABLED)

		3.3 V Device Current (mA) (A) Note 5, Note 7	1.2 V Device Current (mA) (B) Note 6, Note 7	TX Magnetics Current (mA) (C) Note 8	Device Power with 2.5 V Magnetics (mW) Note 9, Note 10	Device Power with 3.3 V Magnetics (mW) Note 9, Note 11
Reset (RST#)	Typ.	32.6	42.7	0.0	159	159
D0, 100BASE-TX with Traffic (No EEE)	Typ.	70.7	96.2	82.0	554	620
D0, 100BASE-TX Idle (w/o EEE)	Typ.	69.2	94.9	82.0	548	613
D0, 100BASE-TX Idle (with EEE)	Typ.	67.7	80.2	0.0	320	320
D0, 10BASE-T with Traffic	Typ.	30.7	76.4	204.0	703	867
D0, 10BASE-T Idle	Typ.	30.7	74.7	202.0	696	858
D0, PHY Energy Detect Power Down	Typ.	14.7	70.6	0.0	134	134
D0, PHY General Power Down	Typ.	7.7	70.4	0.0	110	110
D1, 100BASE-TX Idle (w/o EEE)	Typ.	69.3	45.5	82.0	489	554
D1, 100BASE-TX Idle (with EEE)	Typ.	67.7	32.7	0.0	263	263
D1, 10BASE-T Idle	Typ.	30.3	24.8	202.0	635	797
D1, PHY Energy Detect Power Down	Typ.	15.1	19.7	0.0	74	74
D1, PHY General Power Down	Typ.	8.2	19.4	0.0	51	51
D2, 100BASE-TX Idle (w/o EEE)	Typ.	69.4	45.5	82.0	489	555
D2, 100BASE-TX Idle (with EEE)	Typ.	67.7	32.8	0.0	263	263
D2, 10BASE-T Idle	Typ.	30.0	24.8	202.0	634	796

TABLE 20-4: CURRENT CONSUMPTION AND POWER DISSIPATION (REGS. DISABLED)

D2, PHY Energy Detect Power Down	Typ.	15.0	7.6	0.0	59	59
D2, PHY General Power Down	Typ.	8.2	7.3	0.0	36	36
D3, PHY General Power Down	Typ.	8.1	3.9	0.0	32	32

Note 5: VDD33TXRX1, VDD33TXRX2, VDD33BIAS, VDD33, VDDIO

Note 6: VDD12TX1, VDD12TX2, OSCVDD12, VDDCR

Note 7: Current measurements do not include power applied to the magnetics or the optional external LEDs.

Note 8: The Ethernet component current is independent of the supply rail voltage (2.5V or 3.3V) of the transformer. Two copper TP operation is assumed. Current is half if one PHY is using 100BASE-FX mode. Current is zero if both PHYs are using 100BASE-FX mode.

Note 9: This includes the power dissipated by the transmitter by way of the current through the transformer.

Note 10: $3.3*(A) + 1.2*(B) + (2.5)*(C)$ @ Typ

Note 11: $3.3*(A) + 1.2*(B) + (3.3)*(C)$ @ Typ

20.4.2 INTERNAL REGULATOR ENABLED

TABLE 20-5: CURRENT CONSUMPTION AND POWER DISSIPATION (REGS. ENABLED)

		3.3 V Device Current (mA) (A) Note 12, Note 13, Note 14	TX Magnetics Current (mA) (C) Note 15	Device Power with 2.5 V Magnetics (mW) Note 16, Note 17	Device Power with 3.3 V Magnetics (mW) Note 16, Note 18
Reset (RST#)	Typ.	71.5	0.0	236	236
D0, 100BASE-TX with Traffic (No EEE)	Typ.	160.8	82.0	736	802
D0, 100BASE-TX Idle (w/o EEE)	Typ.	157.0	84.0	729	796
D0, 100BASE-TX Idle (with EEE)	Typ.	142.7	0.0	471	471
D0, 10BASE-T with Traffic	Typ.	100.6	206.0	847	1012
D0, 10BASE-T Idle	Typ.	98.5	213.0	858	1028
D0, PHY Energy Detect Power Down	Typ.	80.1	0.0	265	265

TABLE 20-5: CURRENT CONSUMPTION AND POWER DISSIPATION (REGS. ENABLED)

D0, PHY General Power Down	Typ.	73.3	0.0	242	242
D1, 100BASE-TX Idle (w/o EEE)	Typ.	111.1	83.0	575	641
D1, 100BASE-TX Idle (with EEE)	Typ.	98.0	0.0	324	324
D1, 10BASE-T Idle	Typ.	51.6	210.0	696	864
D1, PHY Energy Detect Power Down	Typ.	32.0	0.0	106	106
D1, PHY General Power Down	Typ.	25.6	0.0	85	85
D2, 100BASE-TX Idle (w/o EEE)	Typ.	110.5	83.0	573	639
D2, 100BASE-TX Idle (with EEE)	Typ.	97.9	0.0	324	324
D2, 10BASE-T Idle	Typ.	51.4	211.0	698	866
D2, PHY Energy Detect Power Down	Typ.	20.6	0.0	68	68
D2, PHY General Power Down	Typ.	14.0	0.0	47	47
D3, PHY General Power Down	Typ.	9.9	0.0	33	33

Note 12: [VDD33TXRX1](#), [VDD33TXRX2](#), [VDD33BIAS](#), [VDD33](#), [VDDIO](#)

Note 13: [VDD12TX1](#) and [VDD12TX2](#), are driven by the internal regulator via the PCB. The current is accounted for via [VDD33](#).

Note 14: Current measurements do not include power applied to the magnetics or the optional external LEDs.

Note 15: The Ethernet component current is independent of the supply rail voltage (2.5V or 3.3V) of the transformer. Two copper TP operation is assumed. Current is half if one PHY is using 100BASE-FX mode. Current is zero if both PHYs are using 100BASE-FX mode.

Note 16: This includes the power dissipated by the transmitter by way of the current through the transformer.

Note 17: $3.3*(A) + (2.5)*(C)$ @ Typ

Note 18: $3.3*(A) + (3.3)*(C)$ @ Typ

20.5 DC Specifications

TABLE 20-6: NON-VARIABLE I/O DC ELECTRICAL CHARACTERISTICS

Parameter	Symbol	Min	Typ	Max	Units	Notes
IS Type Input Buffer						
Low Input Level	V_{ILI}	-0.3		0.8	V	
High Input Level	V_{IHI}	2.0		3.6	V	
Schmitt Trigger Hysteresis ($V_{IHT} - V_{ILT}$)	V_{HYS}	121		151	mV	
Input Leakage ($V_{IN} = VSS$ or $VDD33$)	I_{IH}	-10		10	μA	Note 19
Input Capacitance	C_{IN}			3	pF	
Pull-Up Impedance ($V_{IN} = VSS$)	R_{DPU}	6		8.9	K Ω	
Pull-Down Impedance ($V_{IN} = VDD33$)	R_{DPD}	52		79	K Ω	
AI Type Input Buffer (FXSDENA/FXSDENB)						
Low Input Level	V_{IL}	-0.3		0.8	V	
High Input Level	V_{IH}	1.2		$VDD33+0.3$	V	
AI Type Input Buffer (RXPA/RXNA/RXPB/RXNB)						
Differential Input Level	$V_{IN-DIFF}$	0.1		$VDD33TXRXx$	V	
Common Mode Voltage	V_{CM}	1.0	$VDD33TXRXx-1.3$		V	
Input Capacitance	C_{IN}			5	pF	
AI Type Input Buffer (FXLOSEN Input)						
State A Threshold	V_{THA}	-0.3		0.8	V	
State B Threshold	V_{THB}	1.2		1.7	V	
State C Threshold	V_{THC}	2.3		$VDD33+0.3$	V	
ICLK Type Input Buffer (OSCI Input)						
Low Input Level	V_{ILI}	-0.3		0.35	V	Note 20
High Input Level	V_{IHI}	$OSCVDD12-0.35$		3.6	V	
Input Leakage	I_{ILCK}	-10		10	μA	

TABLE 20-6: NON-VARIABLE I/O DC ELECTRICAL CHARACTERISTICS (CONTINUED)

Parameter	Symbol	Min	Typ	Max	Units	Notes
ILVPECL Input Buffer						
Low Input Level	$V_{IL}-VDD33TXRXx$	$VDD33TXRXx+0.3$		-1.48	V	Note 21
High Input Level	$V_{IH}-VDD33TXRXx$	-1.14		0.3	V	Note 21
OLVPECL Output Buffer						
Low Output Level	V_{OL}			$VDD33TXRXx-1.62$	V	
High Output Level	V_{OH}	$VDD33TXRXx-1.025$			V	
Peak-to-Peak Differential (SFF mode)	$V_{DIFF-SFF}$	1.2	1.6	2.0	V	
Peak-to-Peak Differential (SFP mode)	$V_{DIFF-SFP}$	0.6	0.8	1.0	V	
Common Mode Voltage	V_{CM}	1.0	$VDD33TXRXx-1.3$		V	
Offset Voltage	V_{OFFSET}		40		mV	Note 22
Load Capacitance	C_{LOAD}			10	pF	

Note 19: This specification applies to all inputs and tri-stated bi-directional pins. Internal pull-down and pull-up resistors add +/- 50 μ A per-pin (typical).

Note 20: **OSCI** can optionally be driven from a 25 MHz singled-ended clock oscillator.

Note 21: LVPECL compatible.

Note 22: V_{OFFSET} is a function of the external resistor network configuration. The listed value is recommended to prevent issues due to crosstalk.

TABLE 20-7: VARIABLE I/O DC ELECTRICAL CHARACTERISTICS

Parameter	Symbol	Min	1.8 V Typ	3.3 V Typ	Max	Units	Notes
VIS Type Input Buffer							
Low Input Level	V_{ILI}	-0.3				V	
High Input Level	V_{IHI}				3.6	V	
Negative-Going Threshold	V_{ILT}	0.64	0.83	1.41	1.76	V	Schmitt trigger
Positive-Going Threshold	V_{IHT}	0.81	0.99	1.65	1.90	V	Schmitt trigger
Schmitt Trigger Hysteresis ($V_{IHT} - V_{ILT}$)	V_{HYS}	102	158	138	288	mV	
Input Leakage ($V_{IN} = V_{SS}$ or V_{DDIO})	I_{IH}	-10			10	μA	Note 23
Input Capacitance	C_{IN}				2	pF	
Pull-Up Impedance ($V_{IN} = V_{SS}$)	R_{DPU}	54	68	82		K Ω	
Pull-Up Current ($V_{IN} = V_{SS}$)	I_{DPU}	20	27	67		μA	
Pull-Down Impedance ($V_{IN} = V_{DD33}$)	R_{DPD}	54	68	85		K Ω	
Pull-Down Current ($V_{IN} = V_{DD33}$)	I_{DPD}	19	26	66		μA	
VO8 Type Buffers							
Low Output Level	V_{OL}				0.4	V	$I_{OL} = 8 \text{ mA}$
High Output Level	V_{OH}	$V_{DDIO} - 0.4$				V	$I_{OH} = -8 \text{ mA}$
VOD8 Type Buffer							
Low Output Level	V_{OL}				0.4	V	$I_{OL} = 8 \text{ mA}$
VO12 Type Buffers							
Low Output Level	V_{OL}				0.4	V	$I_{OL} = 12 \text{ mA}$
High Output Level	V_{OH}	$V_{DDIO} - 0.4$				V	$I_{OH} = -12 \text{ mA}$
VOD12 Type Buffer							
Low Output Level	V_{OL}				0.4	V	$I_{OL} = 12 \text{ mA}$
VOS12 Type Buffers							
High Output Level	V_{OH}	$V_{DDIO} - 0.4$				V	$I_{OH} = -12 \text{ mA}$
VO16 Type Buffers							
Low Output Level	V_{OL}				0.4	V	$I_{OL} = 16 \text{ mA}$
High Output Level	V_{OH}	$V_{DDIO} - 0.4$				V	$I_{OH} = -16 \text{ mA}$

Note 23: This specification applies to all inputs and tri-stated bi-directional pins. Internal pull-down and pull-up resistors add $\pm 50 \mu\text{A}$ per-pin (typical).

TABLE 20-8: 100BASE-TX TRANSCEIVER CHARACTERISTICS

Parameter	Symbol	Min	Typ	Max	Units	Notes
Peak Differential Output Voltage High	V_{PPH}	950	-	1050	mVpk	Note 24
Peak Differential Output Voltage Low	V_{PPL}	-950	-	-1050	mVpk	Note 24
Signal Amplitude Symmetry	V_{SS}	98	-	102	%	Note 24
Signal Rise and Fall Time	T_{RF}	3.0	-	5.0	ns	Note 24
Rise and Fall Symmetry	T_{RFS}	-	-	0.5	ns	Note 24
Duty Cycle Distortion	D_{CD}	35	50	65	%	Note 25
Overshoot and Undershoot	V_{OS}	-	-	5	%	
Jitter	-	-	-	1.4	ns	Note 26

Note 24: Measured at line side of transformer, line replaced by 100Ω (+/- 1%) resistor.

Note 25: Offset from 16 ns pulse width at 50% of pulse peak.

Note 26: Measured differentially.

TABLE 20-9: 10BASE-T TRANSCEIVER CHARACTERISTICS

Parameter	Symbol	Min	Typ	Max	Units	Notes
Transmitter Peak Differential Output Voltage	V_{OUT}	2.2	2.5	2.8	V	Note 27
Receiver Differential Squelch Threshold	V_{DS}	300	420	585	mV	

Note 27: Min/max voltages guaranteed as measured with 100Ω resistive load.

20.6 AC Specifications

This section details the various AC timing specifications of the device.

Note: The I²C timing adheres to the NXP *I²C-Bus Specification*. Refer to the NXP *I²C-Bus Specification* for detailed I²C timing information.


Note: The MII/SMI timing adheres to the *IEEE 802.3 Specification*.

Note: The RMII timing adheres to the RMII Consortium *RMII Specification R1.2*.

20.6.1 EQUIVALENT TEST LOAD

Output timing specifications assume the 25 pF equivalent test load, unless otherwise noted, as illustrated in [Figure 20-1](#).

FIGURE 20-1: OUTPUT EQUIVALENT TEST LOAD



20.6.2 POWER SEQUENCING TIMING

These diagrams illustrates the device power sequencing requirements. The **VDDIO**, **VDD33**, **VDD33TXRX1**, **VDD33TXRX2**, **VDD33BIAS** and magnetics power supplies must all reach operational levels within the specified time period t_{pon} . When operating with the internal regulators disabled, **VDDCR**, **OSCVDD12**, **VDD12TX1** and **VDD12TX2** are also included into this requirement.

In addition, once the **VDDIO** power supply reaches 1.0 V, it must reach 80% of its operating voltage level (1.44 V when operating at 1.8 V, 2.0 V when operating at 2.5 V, 2.64 V when operating at 3.3 V) within an additional 15ms. This requirement can be safely ignored if using an external reset as shown in [Section 20.6.3, "Reset and Configuration Strap Timing"](#).

Device power supplies can turn off in any order provided they all reach 0 volts within the specified time period t_{poff} .

FIGURE 20-2: POWER SEQUENCE TIMING - INTERNAL REGULATORS




FIGURE 20-3: POWER SEQUENCE TIMING - EXTERNAL REGULATORS




TABLE 20-10: POWER SEQUENCING TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{pon}	Power supply turn on time	-	-	50	ms
t_{poff}	Power supply turn off time	-	-	500	ms

20.6.3 RESET AND CONFIGURATION STRAP TIMING

This diagram illustrates the RST# pin timing requirements and its relation to the configuration strap pins and output drive. Assertion of RST# is not a requirement. However, if used, it must be asserted for the minimum period specified. The RST# pin can be asserted at any time, but must not be deasserted until t_{purstd} after all external power supplies have reached operational levels. Refer to [Section 6.2, "Resets,"](#) on page 42 for additional information.

FIGURE 20-4: RST# PIN CONFIGURATION STRAP LATCHING TIMING




TABLE 20-11: RST# PIN CONFIGURATION STRAP LATCHING TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{purstd}	External power supplies at operational level to RST# deassertion	25			ms
t_{rstia}	RST# input assertion time	200	-	-	μ s
t_{css}	Configuration strap pins setup to RST# deassertion	200	-	-	ns
t_{csh}	Configuration strap pins hold after RST# deassertion	10	-	-	ns
t_{odad}	Output drive after deassertion	3	-	-	us

Note: The clock input must be stable prior to RST# deassertion.

Note: Device configuration straps are latched as a result of RST# assertion. Refer to [Section 6.2.1, "Chip-Level Resets,"](#) on page 43 for details.

Note: Configuration strap latching and output drive timings shown assume that the Power-On reset has finished first otherwise the timings in [Section 20.6.4, "Power-On and Configuration Strap Timing"](#) apply.

20.6.4 POWER-ON AND CONFIGURATION STRAP TIMING

This diagram illustrates the configuration strap valid timing requirements in relation to power-on. In order for valid configuration strap values to be read at power-on, the following timing requirements must be met.

FIGURE 20-5: POWER-ON CONFIGURATION STRAP LATCHING TIMING




TABLE 20-12: POWER-ON CONFIGURATION STRAP LATCHING TIMING VALUES

Symbol	Description	Min	Typ	Max	Units
t_{cfg}	Configuration strap valid time	-	-	15	ms

Note: Configuration straps must only be pulled high or low. Configuration straps must not be driven as inputs.

Device configuration straps are also latched as a result of RST# assertion. Refer to [Section 20.6.3, "Reset and Configuration Strap Timing"](#) and [Section 6.2.1, "Chip-Level Resets,"](#) on page 43 for additional details.

20.6.5 HOST BUS INTERFACE I/O TIMING

Timing specifications for the Host Bus Interface are given in [Section 9.4.5, "Multiplexed Addressing Mode Timing Requirements," on page 102](#) and [Section 9.5.7, "Indexed Addressing Mode Timing Requirements," on page 128](#).

20.6.6 SPI/SQI SLAVE INTERFACE I/O TIMING

Timing specifications for the SPI/SQI Slave Bus Interface are given in [Section 10.4, "SPI/SQI Timing Requirements," on page 151](#).

20.6.7 I²C EEPROM I/O TIMING

Timing specifications for I²C EEPROM access are given in [Section 14.3, "I2C Master EEPROM Controller," on page 458](#).

20.6.8 JTAG TIMING

Timing specifications for the JTAG interface are given in [Table 19.1.1, "JTAG Timing Requirements," on page 578](#).

20.7 Clock Circuit

The device can accept either a 25 MHz crystal or a 25 MHz single-ended clock oscillator (± 50 ppm) input. If the single-ended clock oscillator method is implemented, **OSCO** should be left unconnected and **OSCI** should be driven with a clock signal that adheres to the specifications outlined throughout [Section 20.0, "Operational Characteristics"](#). See [Table 20-13](#) for the recommended crystal specifications.

TABLE 20-13: CRYSTAL SPECIFICATIONS

PARAMETER	SYMBOL	MIN	NOM	MAX	UNITS	NOTES
Crystal Cut	AT, typ					
Crystal Oscillation Mode	Fundamental Mode					
Crystal Calibration Mode	Parallel Resonant Mode					
Frequency	F_{fund}	-	25.000	-	MHz	
802.3 Frequency Tolerance at 25°C	F_{tol}	-	-	± 40	ppm	Note 28
802.3 Frequency Stability Over Temp	F_{temp}	-	-	± 40	ppm	Note 28
802.3 Frequency Deviation Over Time	F_{age}	-	± 3 to 5	-	ppm	Note 29
802.3 Total Allowable PPM Budget		-	-	± 50	ppm	Note 30
Shunt Capacitance	C_O	-	-	7	pF	
Load Capacitance	C_L	-	-	18	pF	
Drive Level	P_W	300 Note 31	-	-	μW	
Equivalent Series Resistance	R_1	-	-	100	Ω	
Operating Temperature Range		Note 32	-	Note 33	°C	
OSCI Pin Capacitance		-	3 typ	-	pF	Note 34
OSCO Pin Capacitance		-	3 typ	-	pF	Note 34

Note 28: The maximum allowable values for frequency tolerance and frequency stability are application dependent. Since any particular application must meet the IEEE ± 50 ppm Total PPM Budget, the combination of these two values must be approximately ± 45 ppm (allowing for aging).

Note 29: Frequency Deviation Over Time is also referred to as Aging.

Note 30: The total deviation for 100BASE-TX is ± 50 ppm.

Note 31: The minimum drive level requirement P_W is reduced to 100 μW with the addition of a 500 Ω series resistor, if $C_O \leq 5$ pF, $C_L \leq 12$ pF and $R_1 \leq 80$ Ω

Note 32: 0 °C for commercial version, -40 °C for industrial version

Note 33: +70 °C for commercial version, +85 °C for industrial version

Note 34: This number includes the pad, the bond wire and the lead frame. PCB capacitance is not included in this value. The **OSCI** pin, **OSCO** pin and PCB capacitance values are required to accurately calculate the value of the two external load capacitors. The total load capacitance must be equivalent to what the crystal expects to see in the circuit so that the crystal oscillator will operate at 25.000 MHz.

21.0 PACKAGE OUTLINES

21.1 72-QFN

FIGURE 21-1: 72-QFN PACKAGE





FIGURE 21-2: 72-QFN PACKAGE DIMENSIONS




Units		MILLIMETERS		
Dimension Limits		MIN	NOM	MAX
Number of Pins	N	72		
Pitch	e	0.50 BSC		
Overall Height	A	0.80	0.90	1.00
Standoff	A1	0.00	0.02	0.05
Contact Thickness	A3	0.20 REF		
Overall Width	E	10.00 BSC		
Exposed Pad Width	E2	5.90	6.00	6.10
Overall Length	D	10.00 BSC		
Exposed Pad Length	D2	5.90	6.00	6.10
Contact Width	b	0.18	0.25	0.30
Contact Length	L	0.30	0.40	0.50
Contact-to-Exposed Pad	K	1.50	1.60	-

Notes:


1. Pin 1 visual index feature may vary, but must be located within the hatched area.
2. Package is saw singulated.
3. Dimensioning and tolerancing per ASME Y14.5M.
 BSC: Basic Dimension. Theoretically exact value shown without tolerances.
 REF: Reference Dimension, usually without tolerance, for information purposes only.

21.2 80-TQFP-EP

FIGURE 21-3: 80-TQFP-EP PACKAGE



Dimension	Units	Limits	MILLIMETERS		
			MIN	NOM	MAX
Number of Leads	N		80		
Lead Pitch	e		0.50 BSC		
Overall Height	A		—	—	1.20
Molded Package Thickness	A2		0.95	1.00	1.05
Standoff	A1		0.05	—	0.15
Foot Length	L		0.45	0.60	0.75
Footprint	L1		1.00 REF		
Foot Angle	ϕ		0°	3.5°	7°
Overall PKG Length/Width	D/E		14.00 BSC		
PKG Body Length/Width	D1/E1		12.00 BSC		
Exposed Pad Length/Width	D2/E2		5.90	6.00	6.10
Lead Thickness	c		0.09	—	0.20
Lead Width	b		0.17	0.22	0.27
Mold Draft Angle Top	α		11°	12°	13°
Mold Draft Angle Bottom	β		11°	12°	13°



Lead Position Tolerance (dim "b")

ϕ	0.08	M	C
--------	------	---	---

Lead Coplanarity Tolerance

\bigcap	0.08	C
-----------	------	---

Notes:

- Pin 1 visual index feature may vary, but must be located within the hatched area.
- Chamfers at corners are optional; size may vary.
- Dimensions D1 and E1 do not include mold flash or protrusions. Mold flash or protrusions shall not exceed 0.25mm per side.
- Dimensioning and tolerancing per ASME Y14.5M

BSC: Basic Dimension. Theoretically exact value shown without tolerances.

REF: Reference Dimension, usually without tolerance, for information purposes only.

22.0 REVISION HISTORY

TABLE 22-1: REVISION HISTORY

Revision Level	Section/Figure/Entry	Correction
DS00001923A (06-30-15)		Initial Release

THE MICROCHIP WEB SITE

Microchip provides online support via our WWW site at www.microchip.com. This web site is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the web site contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, Microchip consultant program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

CUSTOMER CHANGE NOTIFICATION SERVICE

Microchip's customer notification service helps keep customers current on Microchip products. Subscribers will receive e-mail notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, access the Microchip web site at www.microchip.com. Under "Support", click on "Customer Change Notification" and follow the registration instructions.

CUSTOMER SUPPORT

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or field application engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the web site at: <http://microchip.com/support>

PRODUCT IDENTIFICATION SYSTEM

To order or obtain information, e.g., on pricing or delivery, refer to the factory or the listed sales office.

<u>PART NO.</u>	[X]	[X]	/	XX
Device	Tape and Reel Option	Temperature Range		Package
Device:	LAN9352			
Tape and Reel Option:	Blank = Standard packaging (tray) T = Tape and Reel (Note 1)			
Temperature Range:	Blank = 0°C to +70°C (Commercial) I = -40°C to +85°C (Industrial)			
Package:	ML = 72-pin QFN PT = 80-pin TQFP-EP			

Examples:

a) LAN9352/ML
Standard Packaging (Tray),
Commercial Temperature,
72-pin QFN

b) LAN9352TI/PT
Tape and Reel
Industrial Temperature,
80-pin TQFP-EP

Note 1: Tape and Reel identifier only appears in the catalog part number description. This identifier is used for ordering purposes and is not printed on the device package. Check with your Microchip Sales Office for package availability with the Tape and Reel option.

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specification contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is one of the most secure families of its kind on the market today, when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods used to breach the code protection feature. All of these methods, to our knowledge, require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Most likely, the person doing so is engaged in theft of intellectual property.
- Microchip is willing to work with the customer who is concerned about the integrity of their code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of their code. Code protection does not mean that we are guaranteeing the product as “unbreakable.”

Information contained in this publication regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications. MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION, INCLUDING BUT NOT LIMITED TO ITS CONDITION, QUALITY, PERFORMANCE, MERCHANTABILITY OR FITNESS FOR PURPOSE. Microchip disclaims all liability arising from this information and its use. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, dsPIC, FlashFlex, flexPWR, JukeBlox, KEELOQ, KEELOQ logo, Kleer, LANCheck, MediaLB, MOST, MOST logo, MPLAB, OptoLyzer, PIC, PICSTART, PIC³² logo, RightTouch, SpyNIC, SST, SST Logo, SuperFlash and UNI/O are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

The Embedded Control Solutions Company and mTouch are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Analog-for-the-Digital Age, BodyCom, chipKIT, chipKIT logo, CodeGuard, dsPICDEM, dsPICDEM.net, ECAN, In-Circuit Serial Programming, ICSP, Inter-Chip Connectivity, KleerNet, KleerNet logo, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, RightTouch logo, REAL ICE, SQI, Serial Quad I/O, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

Silicon Storage Technology is a registered trademark of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2015, Microchip Technology Incorporated, Printed in the U.S.A., All Rights Reserved.

ISBN: 9781632773340

QUALITY MANAGEMENT SYSTEM
CERTIFIED BY DNV
= ISO/TS 16949 =

Microchip received ISO/TS-16949:2009 certification for its worldwide headquarters, design and wafer fabrication facilities in Chandler and Tempe, Arizona; Gresham, Oregon and design centers in California and India. The Company's quality system processes and procedures are for its PIC® MCUs and dsPIC® DSCs, KEELOQ® code hopping devices, Serial EEPROMs, microperipherals, nonvolatile memory and analog products. In addition, Microchip's quality system for the design and manufacture of development systems is ISO 9001:2000 certified.



Worldwide Sales and Service

AMERICAS

Corporate Office

2355 West Chandler Blvd.
Chandler, AZ 85224-6199

Tel: 480-792-7200

Fax: 480-792-7277

Technical Support:

[http://www.microchip.com/
support](http://www.microchip.com/support)

Web Address:

www.microchip.com

Atlanta

Duluth, GA

Tel: 678-957-9614

Fax: 678-957-1455

Austin, TX

Tel: 512-257-3370

Boston

Westborough, MA

Tel: 774-760-0087

Fax: 774-760-0088

Chicago

Itasca, IL

Tel: 630-285-0071

Fax: 630-285-0075

Cleveland

Independence, OH

Tel: 216-447-0464

Fax: 216-447-0643

Dallas

Addison, TX

Tel: 972-818-7423

Fax: 972-818-2924

Detroit

Novi, MI

Tel: 248-848-4000

Houston, TX

Tel: 281-894-5983

Indianapolis

Noblesville, IN

Tel: 317-773-8323

Fax: 317-773-5453

Los Angeles

Mission Viejo, CA

Tel: 949-462-9523

Fax: 949-462-9608

New York, NY

Tel: 631-435-6000

San Jose, CA

Tel: 408-735-9110

Canada - Toronto

Tel: 905-673-0699

Fax: 905-673-6509

ASIA/PACIFIC

Asia Pacific Office

Suites 3707-14, 37th Floor
Tower 6, The Gateway
Harbour City, Kowloon

Hong Kong

Tel: 852-2943-5100

Fax: 852-2401-3431

Australia - Sydney

Tel: 61-2-9868-6733

Fax: 61-2-9868-6755

China - Beijing

Tel: 86-10-8569-7000

Fax: 86-10-8528-2104

China - Chengdu

Tel: 86-28-8665-5511

Fax: 86-28-8665-7889

China - Chongqing

Tel: 86-23-8980-9588

Fax: 86-23-8980-9500

China - Dongguan

Tel: 86-769-8702-9880

China - Hangzhou

Tel: 86-571-8792-8115

Fax: 86-571-8792-8116

China - Hong Kong SAR

Tel: 852-2943-5100

Fax: 852-2401-3431

China - Nanjing

Tel: 86-25-8473-2460

Fax: 86-25-8473-2470

China - Qingdao

Tel: 86-532-8502-7355

Fax: 86-532-8502-7205

China - Shanghai

Tel: 86-21-5407-5533

Fax: 86-21-5407-5066

China - Shenyang

Tel: 86-24-2334-2829

Fax: 86-24-2334-2393

China - Shenzhen

Tel: 86-755-8864-2200

Fax: 86-755-8203-1760

China - Wuhan

Tel: 86-27-5980-5300

Fax: 86-27-5980-5118

China - Xian

Tel: 86-29-8833-7252

Fax: 86-29-8833-7256

ASIA/PACIFIC

China - Xiamen

Tel: 86-592-2388138

Fax: 86-592-2388130

China - Zhuhai

Tel: 86-756-3210040

Fax: 86-756-3210049

India - Bangalore

Tel: 91-80-3090-4444

Fax: 91-80-3090-4123

India - New Delhi

Tel: 91-11-4160-8631

Fax: 91-11-4160-8632

India - Pune

Tel: 91-20-3019-1500

Japan - Osaka

Tel: 81-6-6152-7160

Fax: 81-6-6152-9310

Japan - Tokyo

Tel: 81-3-6880-3770

Fax: 81-3-6880-3771

Korea - Daegu

Tel: 82-53-744-4301

Fax: 82-53-744-4302

Korea - Seoul

Tel: 82-2-554-7200

Fax: 82-2-558-5932 or

82-2-558-5934

Malaysia - Kuala Lumpur

Tel: 60-3-6201-9857

Fax: 60-3-6201-9859

Malaysia - Penang

Tel: 60-4-227-8870

Fax: 60-4-227-4068

Philippines - Manila

Tel: 63-2-634-9065

Fax: 63-2-634-9069

Singapore

Tel: 65-6334-8870

Fax: 65-6334-8850

Taiwan - Hsin Chu

Tel: 886-3-5778-366

Fax: 886-3-5770-955

Taiwan - Kaohsiung

Tel: 886-7-213-7828

Taiwan - Taipei

Tel: 886-2-2508-8600

Fax: 886-2-2508-0102

Thailand - Bangkok

Tel: 66-2-694-1351

Fax: 66-2-694-1350

EUROPE

Austria - Wels

Tel: 43-7242-2244-39

Fax: 43-7242-2244-393

Denmark - Copenhagen

Tel: 45-4450-2828

Fax: 45-4485-2829

France - Paris

Tel: 33-1-69-53-63-20

Fax: 33-1-69-30-90-79

Germany - Dusseldorf

Tel: 49-2129-3766400

Germany - Munich

Tel: 49-89-627-144-0

Fax: 49-89-627-144-44

Germany - Pforzheim

Tel: 49-7231-424750

Italy - Milan

Tel: 39-0331-742611

Fax: 39-0331-466781

Italy - Venice

Tel: 39-049-7625286

Netherlands - Drunen

Tel: 31-416-690399

Fax: 31-416-690340

Poland - Warsaw

Tel: 48-22-3325737

Spain - Madrid

Tel: 34-91-708-08-90

Fax: 34-91-708-08-91

Sweden - Stockholm

Tel: 46-8-5090-4654

UK - Wokingham

Tel: 44-118-921-5800

Fax: 44-118-921-5820

01/27/15

Компания «Life Electronics» занимается поставками электронных компонентов импортного и отечественного производства от производителей и со складов крупных дистрибьюторов Европы, Америки и Азии.

С конца 2013 года компания активно расширяет линейку поставок компонентов по направлению коаксиальный кабель, кварцевые генераторы и конденсаторы (керамические, пленочные, электролитические), за счёт заключения дистрибьюторских договоров

Мы предлагаем:

- Конкурентоспособные цены и скидки постоянным клиентам.
- Специальные условия для постоянных клиентов.
- Подбор аналогов.
- Поставку компонентов в любых объемах, удовлетворяющих вашим потребностям.
- Приемлемые сроки поставки, возможна ускоренная поставка.
- Доставку товара в любую точку России и стран СНГ.
- Комплексную поставку.
- Работу по проектам и поставку образцов.
- Формирование склада под заказчика.
- Сертификаты соответствия на поставляемую продукцию (по желанию клиента).
- Тестирование поставляемой продукции.
- Поставку компонентов, требующих военную и космическую приемку.
- Входной контроль качества.
- Наличие сертификата ISO.

В составе нашей компании организован Конструкторский отдел, призванный помогать разработчикам, и инженерам.

Конструкторский отдел помогает осуществить:

- Регистрацию проекта у производителя компонентов.
- Техническую поддержку проекта.
- Защиту от снятия компонента с производства.
- Оценку стоимости проекта по компонентам.
- Изготовление тестовой платы монтаж и пусконаладочные работы.



Тел: +7 (812) 336 43 04 (многоканальный)

Email: org@lifeelectronics.ru

www.lifeelectronics.ru