VSC8574

Quad-Port 10/100/1000BASE-T PHY with IEEE 1588, Synchronous Ethernet, and QSGMII/SGMII MAC

Datasheet

Advance Product Information
Subject to Change

VMDS-10324
Revision 2.1
February 2011

Confidential

Vitesse
Corporate Headquarters
741 Calle Plano
Camarillo, California 93012
United States

www.vitesse.com

Vitesse Semiconductor Corporation (“Vitesse”) retains the right to make changes to its products or specifications to improve performance, reliability or manufacturability. All information in this document, including descriptions of features, functions, performance, technical specifications and availability, is subject to change without notice at any time. While the information furnished herein is held to be accurate and reliable, no responsibility will be assumed by Vitesse for its use. Furthermore, the information contained herein does not convey to the purchaser of microelectronic devices any license under the patent right of any manufacturer.

Vitesse products are not intended for use in products or applications, including, but not limited to, medical devices (including life support and implantable medical devices), nuclear products, or other safety-critical uses where failure of a Vitesse product could reasonably be expected to result in personal injury or death. Anyone using a Vitesse product in such an application without express written consent of an officer of Vitesse does so at their own risk, and agrees to fully indemnify Vitesse for any damages that may result from such use or sale.

Safety of Laser Products, IEC 60825. While Vitesse products support IEC 60825, use of Vitesse products does not ensure compliance to IEC 60825. Buyers are responsible for ensuring compliance to IEC 60825. Buyers must fully indemnify Vitesse for any damages resulting from non-compliance to IEC 60825.

Vitesse Semiconductor Corporation is a registered trademark. All other products or service names used in this publication are for identification purposes only, and may be trademarks or registered trademarks of their respective companies. All other trademarks or registered trademarks mentioned herein are the property of their respective holders.

Contents

Figures

Tables

Revision History

This section describes the changes that were implemented in this document. The changes are listed by revision, starting with the most current publication.

Revision 2.1

Revision 2.1 of this datasheet was published in February 2011. The following is a summary of the changes implemented in the datasheet:

- Several pin descriptions were modified. For more information, see "Pin Descriptions," page 179.
- The XIC package was replaced by the XKS package. For more information about the package description, see "Package Information," page 194.
- VSC8574-03 is now available. It has an extended operating temperature range of –40 °C ambient to 125 °C junction.
- The maximum operating temperature for VSC8574 was increased from 90 °C junction to 125 °C junction.
- For carrier class applications, the maximum operating temperature for both the VSC8574 and VSC8574-03 devices is 110 °C junction.
- Thermal resistances were added. For more information, see "Thermal Specifications," page 196.

Revision 2.0

Revision 2.0 of this datasheet was published in December 2010. This was the first publication of the document.

1 Introduction

This document consists of descriptions and specifications for both functional and physical aspects of the VSC8574 quad-port 10/100/1000BASE-T PHY device with four dual media ports for the Ethernet market segment.

In addition to datasheets, the Vitesse Web site offers an extensive library of documentation, support files, and application materials specific to each device. The address of the Vitesse Web site is www.vitesse.com.

1.1 Register and Bit Conventions

This document refers to registers by their address and bit number in decimal notation. A range of bits is indicated with a colon. For example, a reference to address 26, bits 15 through 14 is shown as 26.15:14.

A register with an E and a number attached (example 27E1) means it is a register contained within extended register page number 1. A register with a G attached (example 13G) means it is a GPIO page register.

Bit numbering follows the IEEE standard with bit 15 being the most significant bit and bit 0 being the least significant bit.

2 Product Overview

The VSC8574 is a low-power, quad-port Gigabit Ethernet transceiver with four SerDes interfaces for quad-port dual media capability. It also includes an integrated quad port I2C multiplexer (MUX) to control SFPs or PoE modules. It has a low electromagnetic interference (EMI) line driver, and integrated line side termination resistors that conserve both power and printed circuit board (PCB) space.

The VSC8574 includes Vitesse's unique IEEE 1588 timestamping engine with dual encapsulation support. The VSC8574 also includes dual recovered clock outputs to support Synchronous Ethernet applications. Programmable clock squelch control is included to inhibit undesirable clocks from propagating and to help prevent timing loops. In addition, the VSC8574 also supports a ring resiliency feature that allows a 1000BASE-T connected PHY port to switch between master and slave timing without having to interrupt the 1000BASE-T link.

Using Vitesse's EcoEthernet v2.0 PHY technology, the VSC8574 supports energy efficiency features such as Energy Efficient Ethernet (EEE), ActiPHY link down power savings, and PerfectReach that can adjust power based on the cable length. It also supports fully optimized power consumption in all link speeds.

Vitesse's mixed signal and digital signal processing (DSP) architecture is a key operational feature of the VSC8574, assuring robust performance even under less-than-favorable environmental conditions. It supports both half-duplex and full-duplex 10BASE-T, 100BASE-TX, and 1000BASE-T communication speeds over Category 5 (Cat5) unshielded twisted pair (UTP) cable at distances greater than 100 m, displaying excellent tolerance to NEXT, FEXT, echo, and other types of ambient environmental and system electronic noise. The device also supports four dual media ports and can support up to four 100BASE-FX, 1000BASE-X, and/or triple-speed copper SFPs.

The following illustrations show a high-level, general view of typical VSC8574 applications.

Figure 1. Dual Media Application Diagram

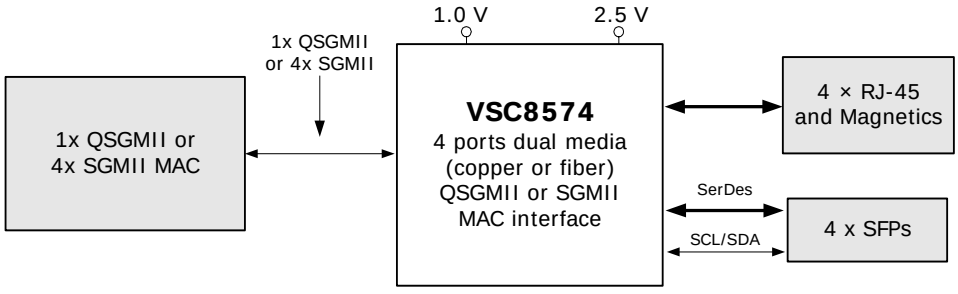

Figure 2. Copper Transceiver Application Diagram

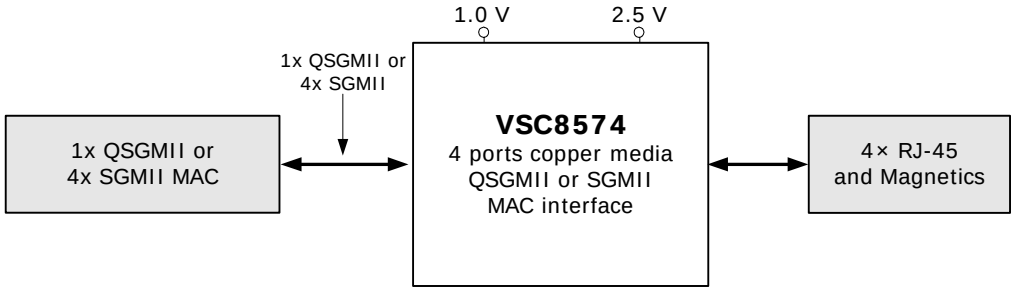

2.1 Key Features

This section lists the main features and benefits of the VSC8574 device.

Low Power

- Low power consumption of approximately 425 mW per port in 1000BASE-T mode, 200 mW per port in 100BASE-TX mode, 225 mW per port in 10BASE-T mode, and less than 115 mW per port in 100BASE-FX and 1000BASE-X modes
- ActiPHY™ link down power savings
- PerfectReach™ smart cable reach algorithm
- IEEE 802.3az Energy Efficient Ethernet idle power savings

Advanced Carrier Ethernet Support

- Support for 802.3ah unidirectional transport for 100BASE-FX and 1000BASE-X fiber media
- Recovered clock outputs with programmable clock squelch control and fast link failure indication (<1 ms) for G.8261 Synchronous Ethernet applications
- Support for IEEE 1588-2008 timestamping with dual encapsulation
- Flexible transmit and receive frequency timing per PHY port
- Ring resiliency for maintaining linkup integrity when switching between 1000BASE-T master and slave timing
- Integrated quad I2C mux to control SFP and PoE modules
- Supports IEEE 802.3bf Timing and Synchronization standard

Wide Range of Support

- Compliant with IEEE 802.3 (10BASE-T, 100BASE-TX, 1000BASE-T, 100BASE-FX, and 1000BASE-X) specifications
- Support for >16 kB jumbo frames in all speeds with programmable synchronization FIFOs

- Integrated on-die temperature monitoring feature with threshold settings accessible through register bits to allow a host to either poll through the MDC or MDIO pins, or be interrupted through the hardware pin
- Supports Cisco SGMII v1.9, Cisco QSGMII v1.3 and 1000BASE-X MACs, IEEE 1149.1 JTAG boundary scan, and IEEE 1149.6 AC-JTAG
- Support for applications that need to meet 2KV CDE, IEC 61000-4-2 at 8KV
- Available in a low-cost, 256-pin BGA package with a 17 mm × 17 mm body size

Flexibility

- VeriPHY® cable diagnostics suite provides extensive network cable information such as cable length, termination status, and open/short fault location
- Patented, low EMI line driver with integrated line side termination resistors
- Four programmable direct-drive LEDs per port with adjustable brightness levels using register controls; bi-color LED support using two LED pins
- Serial LED interface option
- Extensive test features including near end, far end, copper media connector, SerDes MAC/media loopback, and Ethernet packet generator with CRC error counter to decrease time-to-market

2.2 Block Diagram

The following illustration shows the primary functional blocks of the VSC8574 device.

Figure 3. Block Diagram

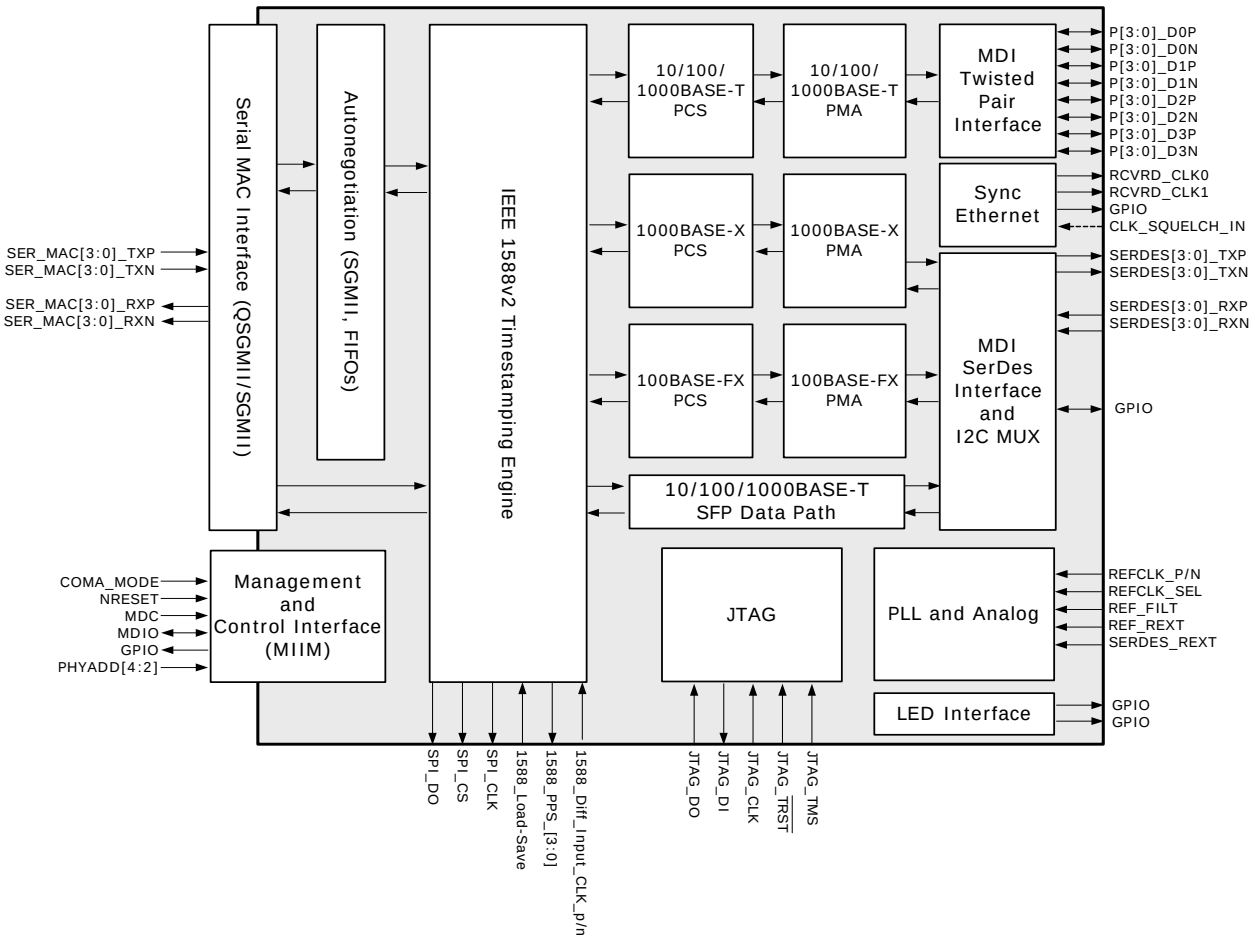

3 Functional Descriptions

This section provides detailed information about the functionality of the VSC8574 device, available configurations, operational features, and testing functionality. It includes descriptions of the various device interfaces and their configuration. With the information in this section, the device setup parameters can be determined for configuring the VSC8574 part for use in a particular application.

3.1 Operating Modes

The following table lists the operating modes of the VSC8574 device.

Table 1. Operating Modes

Operating Mode	Supported Media	Notes
SerDes MAC-to-Cat5 Link Partner	1000BASE-T only	See Figure 2, page 16.
QSGMII/SGMII MAC-to-Cat5 Link Partner	10/100/1000BASE-T	See Figure 2, page 16.
QSGMII/SGMII MAC-to-1000BASE-X Link Partner	1000BASE-X	See Figure 1, page 15.
QSGMII/SGMII MAC-to-100BASE-FX Link Partner	100BASE-FX	See Figure 1, page 15.
QSGMII/SGMII MAC-to-SGMII/1000BASE-X	SFP/Fiber (1000BASE-X, 10/100/1000BASE-T Cu SFP)	See Figure 1, page 15.
QSGMII/SGMII MAC-to-AMS and 1000BASE-X SerDes	1000BASE-X, 10/100/ 1000BASE-T	See Figure 1, page 15.
QSGMII/SGMII MAC-to-AMS and 100BASE-FX SerDes	100BASE-FX, 10/100/ 1000BASE-T	See Figure 1, page 15.
QSGMII/SGMII MAC-to-AMS and SGMII/ 1000BASE-X	SFP/Fiber (1000BASE-X, 10/100/1000BASE-T, 1000BASE-X, 10/100/ 1000BASE-T Cu SFP)	See Figure 1, page 15.

3.2 SerDes MAC Interface

The VSC8574 SerDes MAC interface performs data serialization and deserialization functions using an integrated SerDes block. The interface operates in three modes: 1000BASE-X compliant mode, SGMII mode, and QSGMII mode. Both the SerDes and the enhanced SerDes blocks have the termination resistor integrated into the device. The SerDes block also has the AC decoupling capacitors integrated in the receive path. Integrated AC decoupling is not supported in the enhanced SerDes block. Register 19G is a global register and only needs to be set once to configure the device. The other register bits are configured on a per-port basis and the operation either needs to be repeated for each port, or a broadcast write needs to be used by setting register 22, bit 0 to configure all the ports simultaneously.

3.2.1 SerDes MAC

When connected to a SerDes MAC compliant to 1000BASE-X, the VSC8574 device provides data throughput at a rate of 1000 Mbps only; 10 Mbps and 100 Mbps rates are not supported. To configure the device for SerDes MAC mode, set register 19G, bits 15:14 = 01, and register 23, bit 12 = 1. The device also supports 1000BASE-X Clause 37 MAC-side autonegotiation and is enabled through register 16E3, bit 7. To configure the rest of the device for 1000 Mbps operation, select 1000BASE-T only by disabling the 10BASE-T/100BASE-TX advertisements in register 4.

Figure 4. SerDes MAC Interface

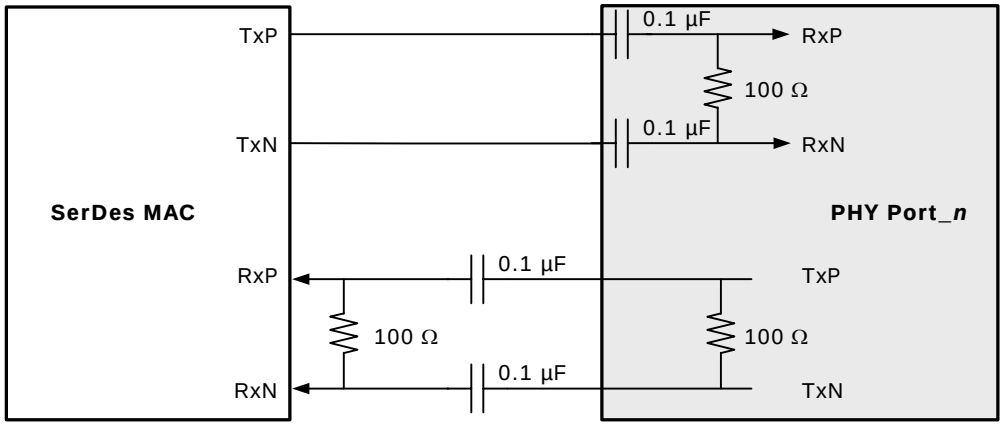

3.2.2 SGMII MAC

When configured to detect and switch between 10BASE-T, 100BASE-T, and 1000BASE-T data rates, the VSC8574 device can be connected to an SGMII-compatible MAC. To configure the device for SGMII MAC mode, set register 19G, bits 15:14 = 01 and register 23, bit 12 = 0. This device also supports SGMII MAC-side autonegotiation and is enabled through register 16E3, bit 7.

Figure 5. SGMII MAC Interface

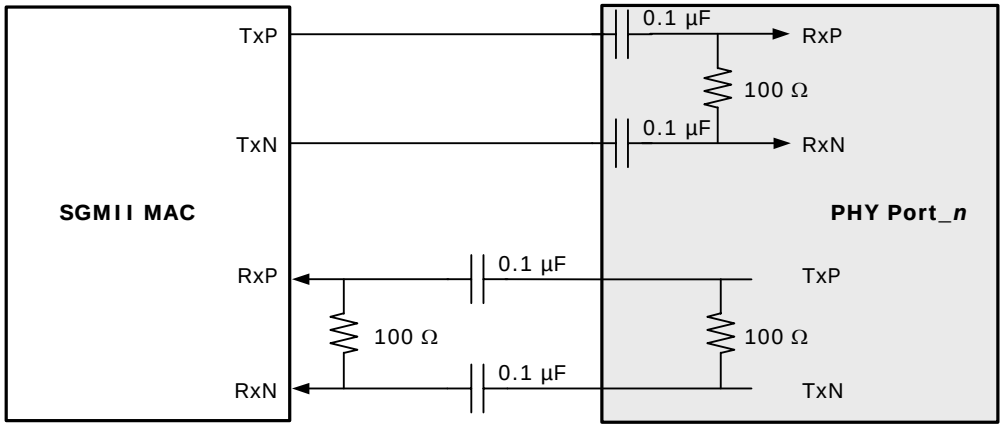

3.2.3 QSGMII MAC

The VSC8574 device supports a QSGMII MAC to convey four ports of network data and port speed between 10BASE-T, 100BASE-T, and 1000BASE-T data rates and operates in both half-duplex and full-duplex at all port speeds. The MAC interface protocol for each port within QSGMII can be either 1000BASE-X or SGMII, if the QSGMII MAC that the VSC8574 is connecting to supports this functionality. To configure the device for QSGMII MAC mode, set register 19G, bits 15:14 = 00 or 10. This device also supports SGMII MAC-side autonegotiation on each individual port and is enabled through register 16E3, bit 7, of that port.

Figure 6. QSGMII MAC Interface

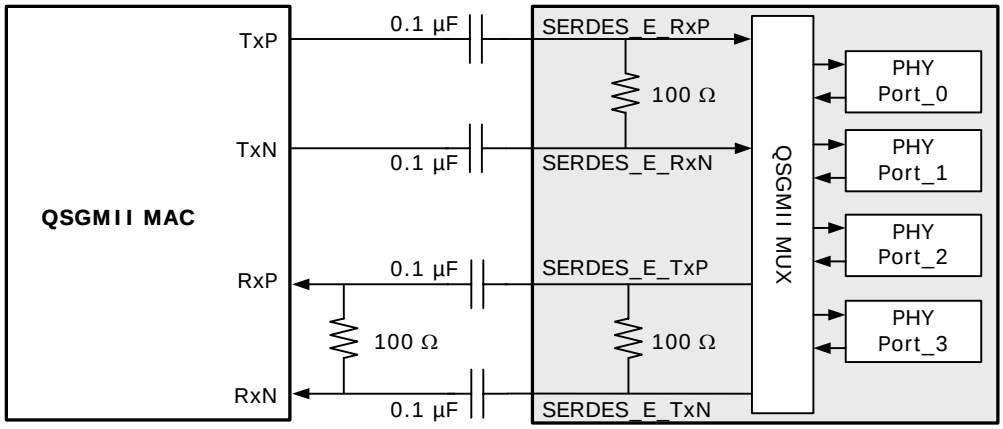

3.3 SerDes Media Interface

The VSC8574 device SerDes media interface performs data serialization and deserialization functions using an integrated SerDes block in the SerDes media interface. The interface operates at 1.25 Gbps speed, providing full-duplex and half-duplex for 10/100/1000 Mbps bandwidth that can connect directly to 100BASE-FX/1000BASE-X-compliant optical devices as well as to 10/100/1000BASE-T copper SFP devices. The interface also provides support for unidirectional transport as defined in IEEE 802.3-2008, Clause 66. The interface has three SerDes operating modes:

- QSGMII/SGMII to 1000BASE-X
- QSGMII/SGMII to 100BASE-FX
- QSGMII/SGMII to SGMII/1000BASE-X SGMII/1000BASE-X protocol transfer

The SerDes media block has the termination resistor integrated into the device. It also has the AC decoupling capacitors integrated in the receive path.

A software reset through register 0, bit 15 is required when changing operating modes between 100BASE-FX and 1000BASE-X.

3.3.1 QSGMII/SGMII to 1000BASE-X

The 1000BASE-X SerDes media in QSGMII mode supports IEEE 802.3 Clause 36 and Clause 37, which describe 1000BASE-X fiber autonegotiation. In this mode, control and status of the SerDes media is displayed in the VSC8574 device registers 0 through 15 in a manner similar to what is described in IEEE 802.3 Clause 28. In this mode, connected copper SFPs can only operate at 1000BASE-T speed. A link in this mode is established using autonegotiation (enabled or disabled) between the PHY and the link partner. To configure the PHY in this mode, set register 23, bits 10:8 = 010. To configure 1000BASE-X autonegotiation for this mode, set register 0, bit 12. Setting this mode and configurations can be performed individually on each of the four ports. Ethernet packet generator (EPG), cyclical redundancy check (CRC) counters, and loopback modes are supported in 1000BASE-X mode.

3.3.2 QSGMII/SGMII to 100BASE-FX

The VSC8574 supports 100BASE-FX communication speed for connecting to fiber modules such as GBICs and SFPs. This capability is facilitated by using the connections on the SerDes pins when connected to a MAC through QSGMII only. Ethernet packet generator (EPG), cyclical redundancy check (CRC) counters, and loopback modes are supported in the 100BASE-FX mode. Setting this mode and configurations can be performed individually on each of the four ports. To configure the PHY in this mode, set register 23, bits 10:8 = 011.

3.3.3 QSGMII/SGMII to SGMII Protocol Conversion

QSGMII to SGMII (protocol transfer) mode is a feature that links a fiber module or triple speed 10/100/1000-T copper SFP to the QSGMII MAC through the VSC8574 device. SGMII can be converted to QSGMII with protocol conversion using this mode.

To configure the PHY in this mode, set register 23, bits 10:8 = 001. To establish the link, assert the TBD pins.

All relevant LED modes are supported except for collision, duplex, and autonegotiation fault. The triple-speed copper SFP's link status and data type plugged into the port can be indicated by the PHY's LEDs. Setting this particular mode and configuration can be performed individually on each of the four ports within a QSGMII grouping.

3.3.4 Unidirectional Transport for Fiber Media

The VSC8574 device supports IEEE 802.3ah for unidirectional transport across its 1000BASE-X and 100BASE-FX fiber media. This feature enables transmission across fiber media, regardless of whether the PHY has determined that a valid link has been established (register 1, bit 2). The only valid operating modes for unidirectional fiber mode are 100BASE-FX or 1000BASE-X fiber media.

To enable this feature, set register 0, bit 5 to 1. For status of the unidirectional ability, read register 1, bit 7.

Note Automatic media sensing does not work with this feature. In addition, because unidirectional fiber media must have autonegotiation disabled, SGMII autonegotiation must also be disabled (register 16E3, bit 7 = 0).

3.4 PHY Addressing and Port Mapping

This section contains information about PHY addressing and port mapping.

3.4.1 PHY Addressing

The VSC8574 includes three external PHY address pins, PHYADD[4:2], to allow control of multiple PHY devices on a system board sharing a common management bus. These pins set the most significant bits of the PHY address port map. Based on the settings of these three address control pins, the internal PHYs in the VSC8574 device take on these address ranges.

3.4.2 SerDes Port Mapping

The VSC8574 includes three 1.25 GHz SerDes macros and one 5 GHz enhanced SerDes macros. Depending on the configuration, some of the SerDes macros are configured as either 1000BASE-X/100BASE-FX SerDes media interfaces or SGMII MAC interfaces. Also, based on configuration, the enhanced SerDes macros are configured as either QSGMII MAC interfaces or SGMII MAC interfaces. The following table shows the different operating modes based on the settings of register 19G, bits 15:14.

Table 2. Operating Modes

Mode	19G[15:14]	Operating Mode
Mode 0	00	QSGMII to CAT5 mode
Mode 1	01	SGMII to CAT5 mode
Mode 2	10	QSGMII to CAT5 and fiber media mode
Mode 3	11	SGMII to CAT5 and fiber media mode

The following table shows the SerDes port mapping in the modes of operation shown in the previous table.

Table 3. SerDes Port Mapping

Interface Pins	Mode 0 / 2	Mode 1 / 3
SERDES_E0_TXP, SERDES_E0_TXN	QSGMII	SGMII0
SERDES_E0_RXP, SERDES_E0_RXN	QSGMII	SGMII0
SERDES_E1_TXP, SERDES_E1_TXN		SGMII1
SERDES_E1_RXP, SERDES_E1_RXN		SGMII1
SERDES_E2_TXP, SERDES_E2_TXN		SGMII2
SERDES_E2_RXP, SERDES_E2_RXN		SGMII2
SERDES_E3_TXP, SERDES_E3_TXN		SGMII3
SERDES_E3_RXP, SERDES_E3_RXN		SGMII3

3.5 Cat5 Twisted Pair Media Interface

The VSC8574 twisted pair interface is compliant with IEEE 802.3-2008 and the IEEE 802.3az standard for energy efficient Ethernet.

3.5.1 Voltage Mode Line Driver

Unlike many other gigabit PHYs, the VSC8574 uses a patented voltage mode line driver that allows it to fully integrate the series termination resistors, which are required to connect the PHY's Cat5 interface to an external 1:1 transformer. Also, the interface does not require the user to place an external voltage on the center tap of the magnetic. The following illustration shows the connections.

Figure 7. Cat5 Media Interface

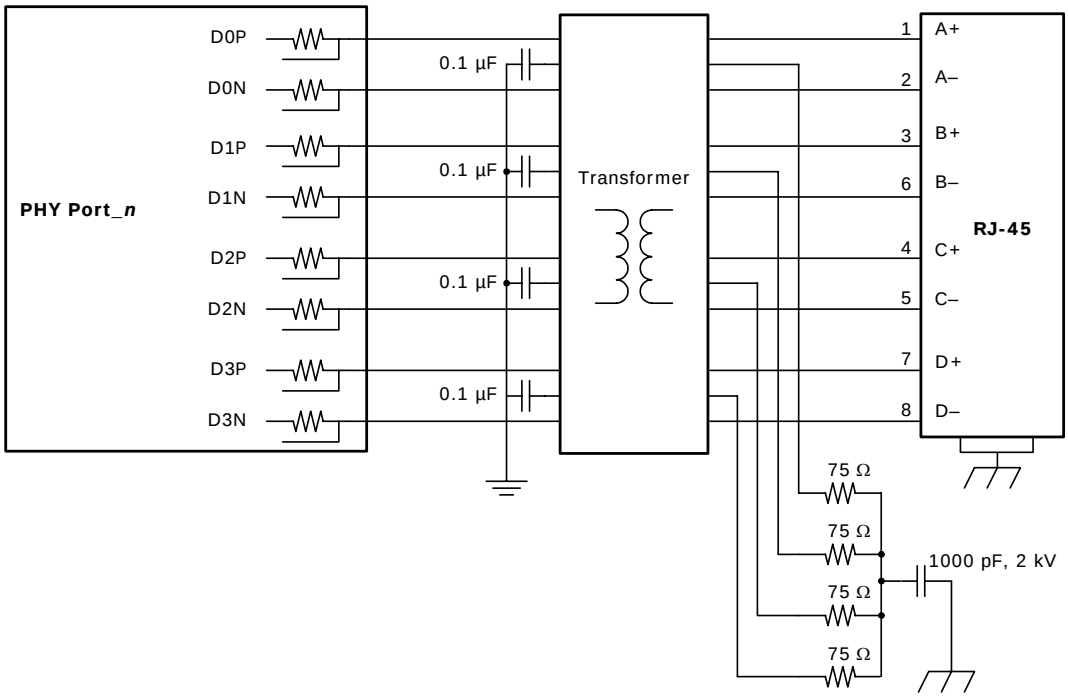

3.5.2 Cat5 Autonegotiation and Parallel Detection

The VSC8574 supports twisted pair autonegotiation, as defined by IEEE 802.3-2008 Clause 28 and IEEE 802.3az. The autonegotiation process evaluates the advertised capabilities of the local PHY and its link partner to determine the best possible operating mode. In particular, autonegotiation can determine speed, duplex configuration, and master or slave operating modes for 1000BASE-TX. Autonegotiation also enables a connected MAC to communicate with its link partner MAC through the VSC8574 using optional next pages, which set attributes that may not otherwise be defined by the IEEE standard.

If the Category 5 (Cat5) link partner does not support autonegotiation, the VSC8574 automatically uses parallel detection to select the appropriate link speed.

Autonegotiation is disabled by clearing register 0, bit 12. When autonegotiation is disabled, the state of register bits 0.6, 0.13, and 0.8 determine the device operating speed and duplex mode.

Note While 10BASE-T and 100BASE-T do not require autonegotiation, Clause 40 has defined 1000BASE-T to require autonegotiation.

3.5.3 1000BASE-T Forced Mode Support

VSC8574 provides support for a 1000BASE-T forced test mode. In this mode, the PHY can be forced into 1000BASE-T mode and does not require manual setting of master/slave at the two ends of the link. This mode is for test purposes only, and is not intended for use in normal operation. To configure a PHY in this mode, set register 17E2, bit 5 = 1 and register 0, bits 6 and 13 = 10.

3.5.4 Automatic Crossover and Polarity Detection

For trouble-free configuration and management of Ethernet links, the VSC8574 includes a robust automatic crossover detection feature for all three speeds on the twisted pair interface (10BASE-T, 100BASE-T, and 1000BASE T). Known as HP Auto-MDIX, the function is fully compliant with Clause 40 of IEEE 802.3-2008.

Additionally, the device detects and corrects polarity errors on all MDI pairs — a useful capability that exceeds the requirements of the standard.

Both HP Auto-MDIX detection and polarity correction are enabled in the device by default. Default settings can be changed using device register bits 18.5:4. Status bits for each of these functions are located in register 28.

Note The VSC8574 can be configured to perform HP Auto-MDIX, even when autonegotiation is disabled and the link is forced into 10/100 speeds. To enable this feature, set register 18.7 to 0. To disable the feature, set register 0.12 to 0.

The HP Auto-MDIX algorithm successfully detects, corrects, and operates with any of the MDI wiring pair combinations listed in the following table.

Table 4. Supported MDI Pair Combinations

1, 2	3, 6	4, 5	7, 8	Mode
A	B	C	D	Normal MDI
B	A	D	C	Normal MDI-X
A	B	D	C	Normal MDI with pair swap on C and D pair
B	A	C	D	Normal MDI-X with pair swap on C and D pair

3.5.5 Manual HP Auto-MDIX Setting

As an alternative to HP Auto-MDIX detection, the PHY can be forced to be MDI or MDI-X using register 19E1, bits 3:2. Setting these bits to 10 forces MDI and setting 11 forces MDI-X. Leaving the bits 00 enables the HP Auto-MDIX setting to be based on register 18, bits 7 and 5.

3.5.6 Link Speed Downshift

For operation in cabling environments that are incompatible with 1000BASE-T, the VSC8574 provides an automatic link speed downshift option. When enabled, the device automatically changes its 1000BASE-T autonegotiation advertisement to the next slower speed after a set number of failed attempts at 1000BASE-T. No reset is required to get out of this state when a subsequent link partner with 1000BASE-T support is connected. This feature is useful in setting up in networks using older cable installations that include only pairs A and B, and not pairs C and D.

To configure and monitor link speed downshifting, set register 20E1, bits 4:1. For more information, see Table 68, page 134.

3.5.7 Energy Efficient Ethernet

The VSC8574 supports the IEEE 802.3az Energy Efficient Ethernet standard that is currently in development. This new standard provides a method for reducing power consumption on an Ethernet link during times of low utilization. It uses low power idles (LPI) to achieve this objective.

Figure 8. Low Power Idle Operation

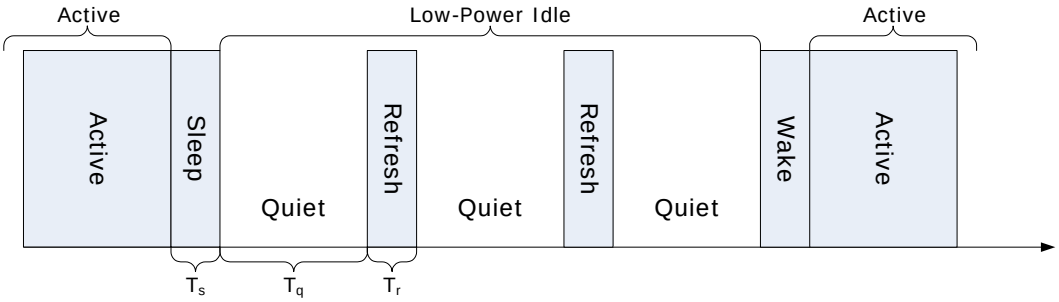

Using LPI, the usage model for the link is to transmit data as fast as possible and then return to a low power idle state. Energy is saved on the link by cycling between active and low power idle states. During LPI, power is reduced by turning off unused circuits and using this method, energy use scales with bandwidth utilization.

The VSC8574 uses LPI to optimize power dissipation in 100BASE-TX and 1000BASE-T modes of operation. In addition, the IEEE 802.3az standard defines a 10BASE-Te mode that reduces transmit signal amplitude from 5 Vp-p to ~3.3 Vp-p. This mode reduces power consumption in 10 Mbps link speed and fully interoperates with legacy 10BASE-T compliant PHYs over 100 m Cat5 cable or better.

To configure the VSC8574 in 10BASE-Te mode, set register 17E2.15 to 1 for each port. Additional energy efficient Ethernet features are controlled through Clause 45 registers. For more information, see "Clause 45 Registers to Support Energy Efficient Ethernet," page 157.

3.5.8 Ring Resiliency

Ring resiliency changes the timing reference between the master and slave PHYs without altering the master/slave configuration. The master PHY transmitter sends clock data based on the local clock and initiates timing recovery in the receiver. The slave PHY instructs node to switch the local timing reference to the recovered clock from other PHYs in the box, freezes timing recovery, and locks clock frequency for the transmitter. The master PHY makes a smooth transition to transmission from local clock to recovered clock after timing lock is achieved.

Ring resiliency can be used in synchronous Ethernet systems, because the local clocks in each node are synchronized to a grandmaster clock.

3.6 Automatic Media Sense Interface Mode

Automatic media sense (AMS) mode automatically sets the media interface to Cat5 mode or SerDes mode. The active media mode chosen is based on the automatic media sense preferences set in the device register 23, bit 11. The following illustration shows a block diagram of AMS functionality on ports 0 through 3 of the VSC8574 device.

Figure 9. Automatic Media Sense Block Diagram

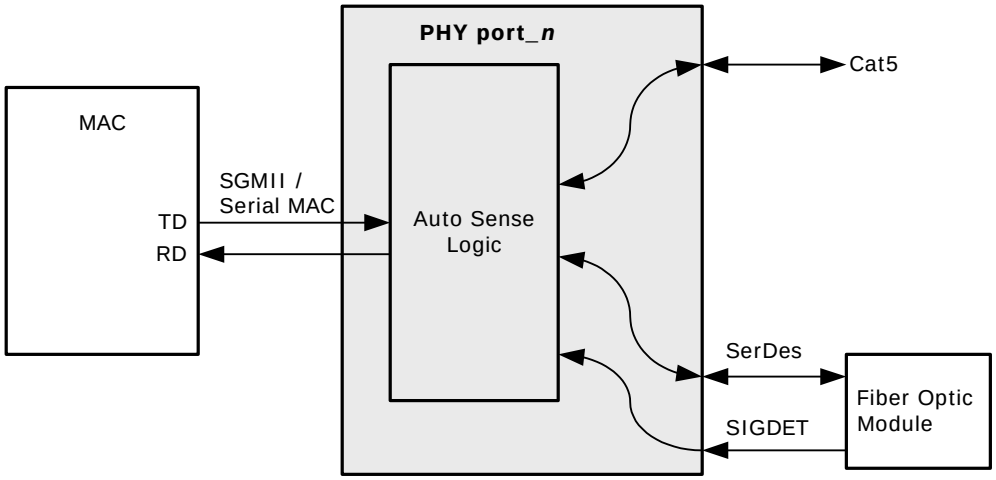

When both the SerDes and Cat5 media interfaces attempt to establish a link, the preferred media interface overrides a linkup of the nonpreferred media interface. For example, if the preference is set for SerDes mode and Cat5 media establishes a link, Cat5 becomes the active media interface. However, after the SerDes media interface establishes a link, the Cat5 interface drops its link because the preference was set for SerDes mode. In this scenario, the SerDes preference determines the active media source until the SerDes link is lost. Also, Cat5 media cannot link up unless there is no

SerDes media link established. The following table shows the possible link conditions based on preference settings.

Table 5. AMS Media Preferences

Preference Setting	Cat5 Linked, Fiber Not Linked	SerDes Linked, Cat5 Not Linked	Cat5 Linked, SerDes Attempts to Link	SerDes Linked, Cat5 Attempts to Link	Both Cat5 and SerDes Attempt to Link
SerDes	Cat5	SerDes	SerDes	SerDes	SerDes
Cat5	Cat5	SerDes	Cat5	Cat5	Cat5

The status of the media mode selected by the AMS can be read from device register 20E1, bits 7:6. It indicates whether copper media, SerDes media, or no media is selected. Each PHY has four automatic media sense modes. The difference between the modes is based on the SerDes media modes:

- SGMII or QSGMII MAC to AMS and 1000BASE-X SerDes
- SGMII or QSGMII MAC to AMS and 100BASE-FX SerDes
- SGMII or QSGMII MAC to AMS and SGMII (protocol transfer)

For more information about SerDes media mode functionality with AMS enabled, see "SerDes Media Interface," page 21.

3.7 Reference Clock

The device reference clock supports 25 MHz and 125 MHz clock signals. The reference clock can be either differential or single-ended. For more information, see "Reference Clock," page 167.

3.7.1 Configuring the REFCLK

The REFCLK_P and REFCLK_N pins configure the reference clock speed. The following table shows the functionality and associated reference clock frequency.

Table 6. REFCLK Frequency Selection

REFCLK_P/N	Frequency
0	125 MHz
1	25 MHz

3.7.2 Single-Ended REFCLK Input

To use a single-ended reference clock, an external resistor network is required. The purpose of the network is to limit the amplitude and to adjust the center of the swing. The configurations for a single-ended REFCLK are shown in the following illustrations.

Figure 10. 2.5 V CMOS Single-Ended REFCLK Input Resistor Network

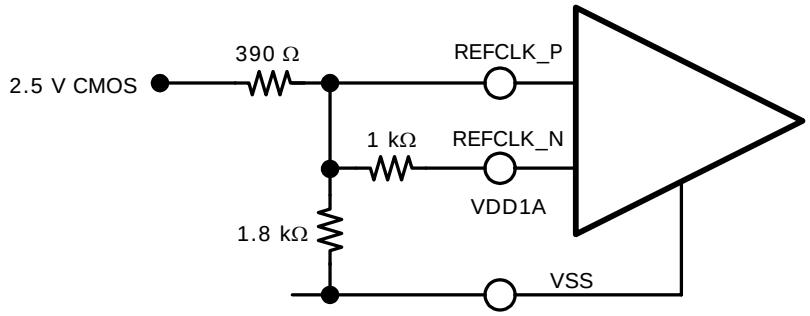

Figure 11. 3.3 V CMOS Single-Ended REFCLK Input Resistor Network

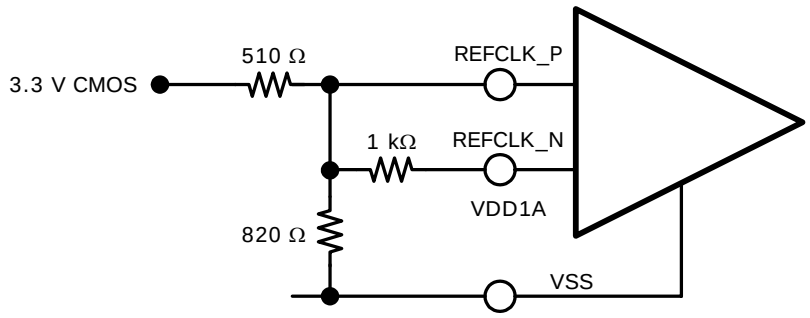

3.8 Ethernet Inline Powered Devices

The VSC8574 can detect legacy inline powered devices in Ethernet network applications. Inline powered detection capability is useful in systems that enable IP phones and other devices (such as wireless access points) to receive power directly from their Ethernet cable, similar to office digital phones receiving power from a private branch exchange (PBX) office switch over telephone cabling. This type of setup eliminates the need for an external power supply and enables the inline powered device to remain active during a power outage, assuming that the Ethernet switch is connected to an uninterrupted power supply, battery, back-up power generator, or other uninterruptable power source.

For more information about legacy inline powered device detection, visit the Cisco Web site at www.cisco.com. The following illustration shows an example of an inline powered Ethernet switch application.

Figure 12. Inline Powered Ethernet Switch Diagram

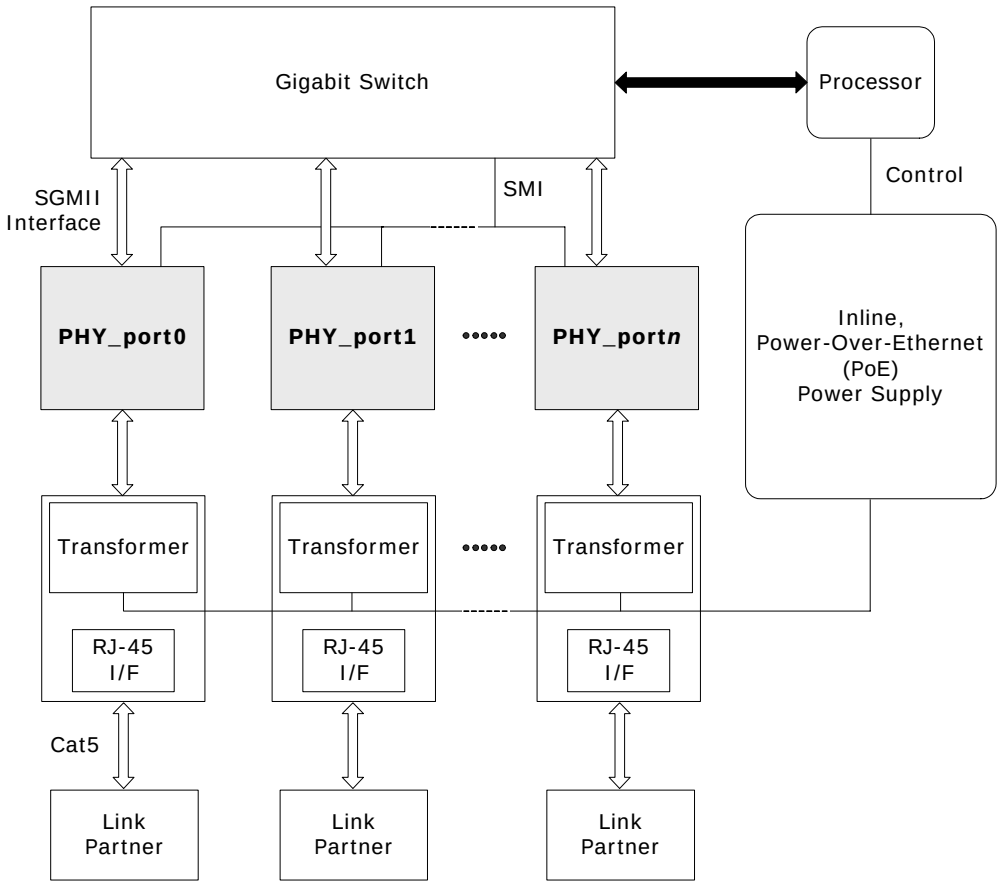

The following procedure describes the process that an Ethernet switch must perform to process inline power requests made by a link partner (LP) that is, in turn, capable of receiving inline power:

1. Enable the inline powered device detection mode on each VSC8574 PHY using its serial management interface. Set register bit 23E1.10 to 1.

2. Ensure that the VSC8574 autonegotiation enable bit (register 0.12) is also set to 1. In the application, the device sends a special fast link pulse (FLP) signal to the LP. Reading register bit 23E1.9:8 returns 00 during the search for devices that require power over Ethernet (PoE).

3. The VSC8574 PHY monitors its inputs for the FLP signal looped back by the LP. An LP capable of receiving PoE loops back the FLP pulses when the LP is in a powered down state. This is reported when VSC8574 register bit 23E1.9:8 reads back 01. It can also be verified as an inline power detection interrupt by reading VSC8574 register bit 26.9, which should be a 1, and which is subsequently cleared and the interrupt de-asserted after the read. When an LP device does not loop back the FLP after a specific time, VSC8574 register bit 23E1.9:8 automatically resets to 10.

4. If the VSC8574 PHY reports that the LP requires PoE, the Ethernet switch must enable inline power on this port, externally of the PHY.

5. The PHY automatically disables inline powered device detection when the VSC8574 register bits 23E1.9:8 automatically reset to 10, and then automatically changes to its normal autonegotiation process. A link is then autonegotiated and established when the link status bit is set (register bit 1.2 is set to 1).

6. In the event of a link failure (indicated when VSC8574 register bit 1.2 reads 0), it is recommended that the inline power be disabled to the inline powered device external to the PHY. The VSC8574 PHY disables its normal autonegotiation process and re-enables its inline powered device detection mode.

3.9 ActiPHY Power Management

In addition to the IEEE-specified power-down control bit (device register bit 0.11), the device also includes an ActiPHY power management mode for each PHY. This mode enables support for power-sensitive applications. It utilizes a signal-detect function that monitors the media interface for the presence of a link to determine when to automatically power-down the PHY. The PHY wakes up at a programmable interval and attempts to wake up the link partner PHY by sending a burst of FLP over copper media.

The ActiPHY power management mode in the VSC8574 is enabled on a per-port basis during normal operation at any time by setting register bit 28.6 to 1.

The following operating states are possible when ActiPHY mode is enabled:

- Low power state
- Link partner wake-up state
- Normal operating state (link-up state)

The VSC8574 switches between the low power state and LP wake-up state at a programmable rate (the default is two seconds) until signal energy has been detected on the media interface pins. When signal energy is detected, the PHY enters the normal operating state. If the PHY is in its normal operating state and the link fails, the PHY returns to the low power state after the expiration of the link status time-out timer. After reset, the PHY enters the low power state.

When autonegotiation is enabled in the PHY, the ActiPHY state machine operates as described. When autonegotiation is disabled and the link is forced to use 10BASE-T or 100BASE-TX modes while the PHY is in its low power state, the PHY continues to transition between the low power and LP wake-up states until signal energy is detected on the media pins. At that time, the PHY transitions to the normal operating state and stays in that state even when the link is dropped. When autonegotiation is disabled while the PHY is in the normal operation state, the PHY stays in that state when the link is dropped and does not transition back to the low power state.

The following illustration shows the relationship between ActiPHY states and timers.

Figure 13. ActiPHY State Diagram

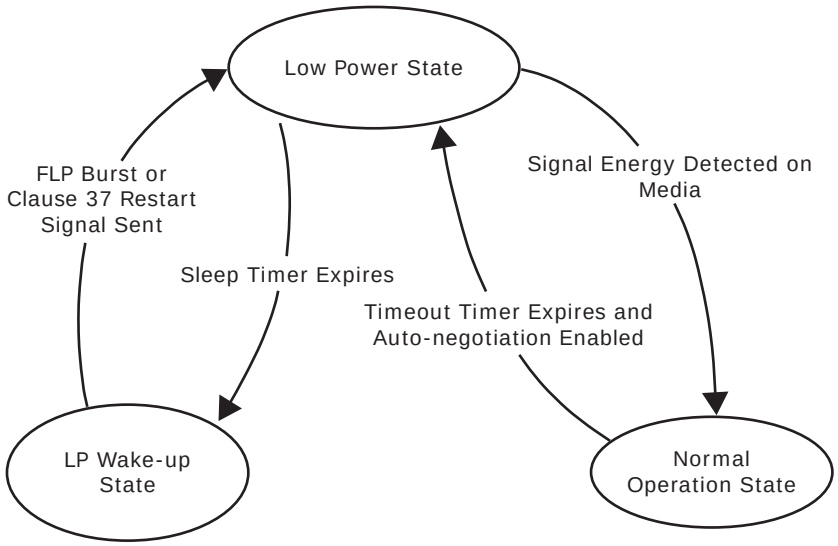

3.9.1 Low Power State

In the low power state, all major digital blocks are powered down. However, the SMI interface (MDC, MDIO, and TBD) functionality is provided.

In this state, the PHY monitors the media interface pins for signal energy. The PHY comes out of low power state and transitions to the normal operating state when signal energy is detected on the media. This happens when the PHY is connected to one of the following:

- Autonegotiation-capable link partner
- Another PHY in enhanced ActiPHY LP wake-up state

In the absence of signal energy on the media pins, the PHY periodically transitions from low-power state to LP wake-up state, based on the programmable sleep timer (register bits 20E1.14:13). The actual sleep time duration is randomized from –80 ms to 60 ms to avoid two linked PHYs in ActiPHY mode entering a lock-up state during operation.

3.9.2 Link Partner Wake-Up State

In the link partner wake-up state, the PHY attempts to wake up the link partner. Up to three complete FLP bursts are sent on alternating pairs A and B of the Cat5 media for a duration based on the wake-up timer, which is set using register bits 20E1.12:11.

In this state, SMI interface (MDC, MDIO, and TBD) functionality is provided.

After sending signal energy on the relevant media, the PHY returns to the low power state.

3.9.3 Normal Operating State

In the normal operating state, the PHY establishes a link with a link partner. When the media is unplugged or the link partner is powered down, the PHY waits for the duration of the programmable link status time-out timer, which is set using register bit 28.7 and bit 28.2. It then enters the low power state.

3.10 IEEE 1588 Timestamping Engine

This section provides information about the IEEE 1588 block for the VSC8574 device.

3.10.1 IEEE 1588 Block Operation

This section describes the basic operation of the architecture, when configured to work in each of the different IEEE 1588 operation modes. The following illustration shows a block diagram of the IEEE 1588 architecture in the VSC8574 device.

Figure 14. IEEE 1588 Architecture

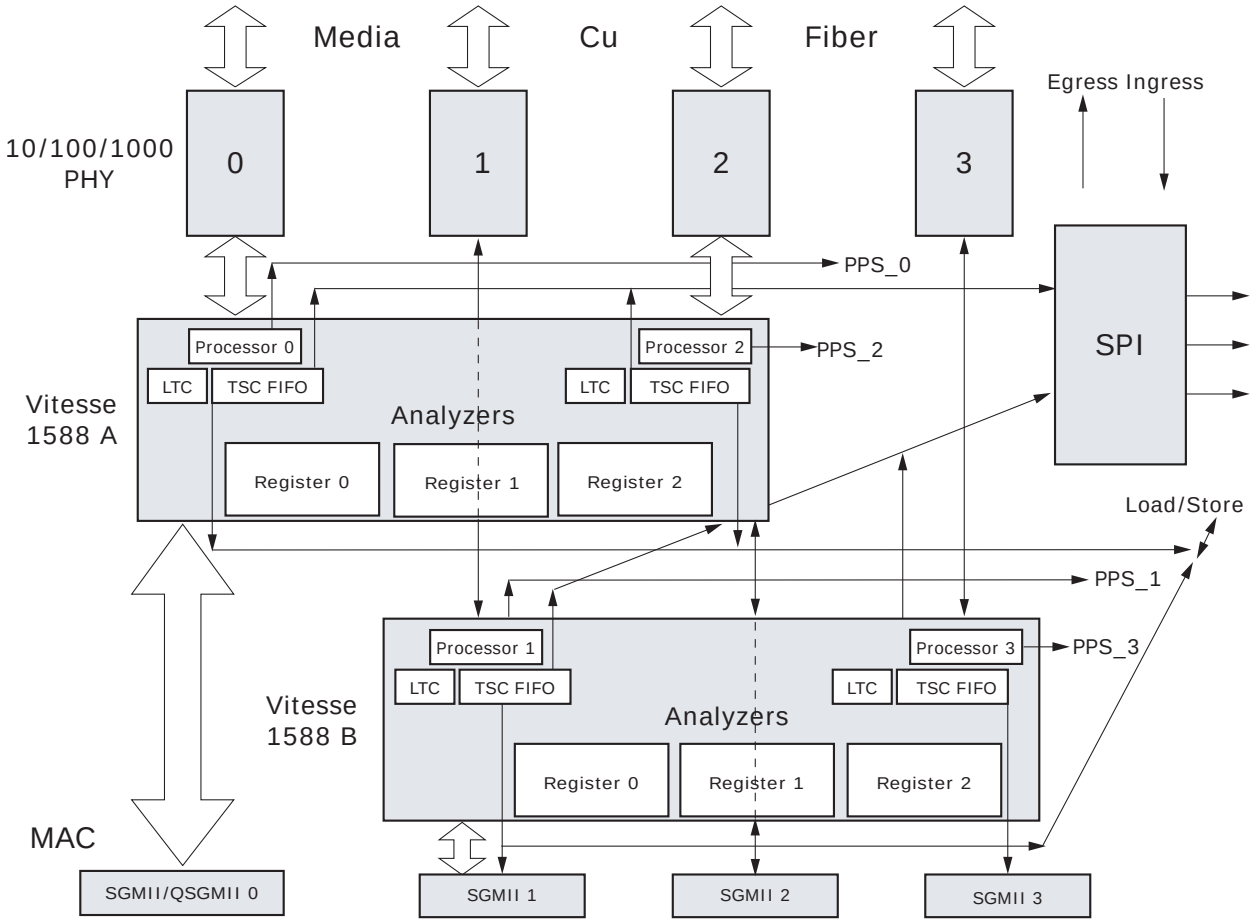

3.10.1.1 One-Step End-to-End Transparent Clock

End-to-end transparent clocks add the residence time (time it takes to traverse from the input to the output port(s) of the system) to all Sync and Delay_Req frames. It does not need to have any knowledge of the actual time, but if it is not locked to the frequency of the 1588 time, it will produce an error that is the ppm difference in frequency times the residence time.

When the TC is frequency-locked by means of 1588 or other methods (SyncE), the error is only caused by sampling inaccuracies.

When an E2E TC needs to recover the frequency using 1588, it must have a PHY with 1588 timestamping support before the 1588 engine, or another way of adding the local time to the correction field. The 1588 Engine is then able to receive Sync frames and adjust the local frequency to match the 1588 time. This can be done by adjusting the Time counter in each PHY or by adjusting the global Timetick clock.

Figure 15. One-Step E2E TC Mode A

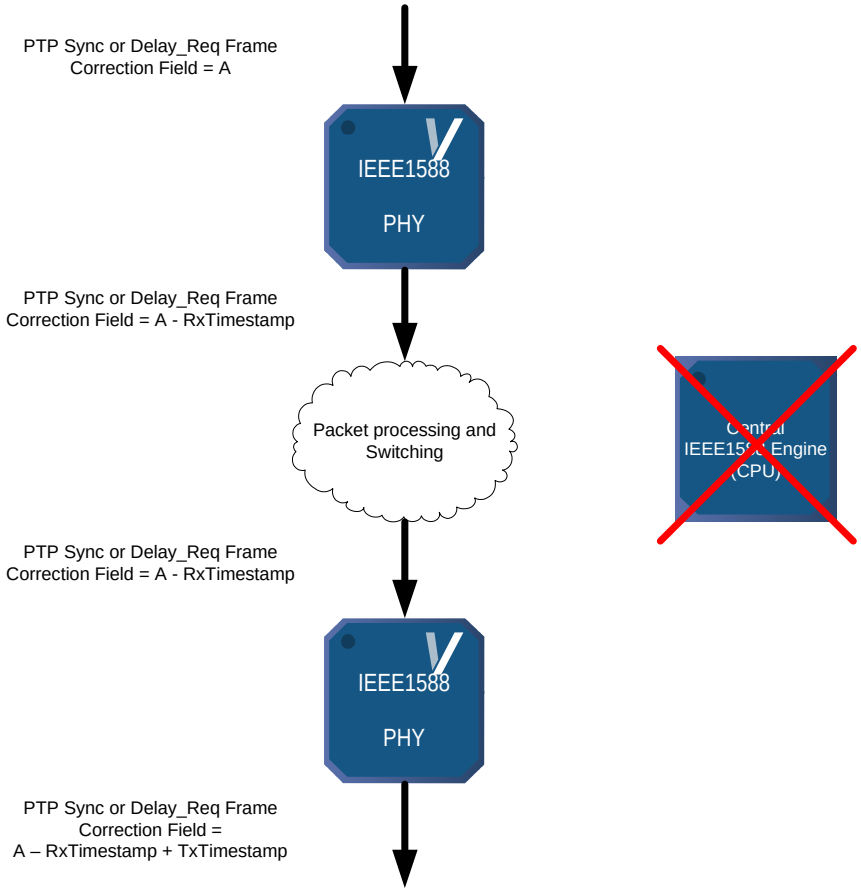

3.10.1.2 Ingress, Mode A

When the system works in one-step E2E TC mode, the system needs to forward Sync and Delay_Req frames through the system and add residence time = Egress timestamp – Ingress timestamp to the correction field in the frame before it leaves the system.

Each time the Timestamp block detects a rising edge on the Start_of_Frame_Indicator pulse (synchronized to the clock domain of the PHY core), it saves the value of the Local_Time that is receives from the Local Time Counter into a raw_timestamp register and converts this to raw_timestamp_ns. It then subtracts the value in the local_correction register from the raw_timestamp_ns value and stores the result in an active_timestamp_ns register. The local_correction register is programmed with the fixed latency from the measurement point to the place that the SFD is detected in the PCS/PMA logic. The timestamp block also contains a register that can be programmed with the known link asymmetry. This value is added or subtracted from the correction field, depending on the frame type.

When the frame leaves the PCS/PMA block it is loaded into a FIFO block that delays and stores the frame data for a few clock cycles to allow for later modifications of the frame. The data is also copied to the analyzer block that parses the incoming frame to detect whether it is a 1588 Sync or Delay_Req frame belonging to the PTP domain that the system is operating on.

If the analyzer detects that the frame is a 1588 Sync or Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Subtract), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is not matched, it signals to the Timestamp block and the Rewriter block to ignore the frame (NOP), which allows it to pass unmodified and flushes the saved timestamp in the Timestamp block.

If the Timestamp block gets the Subtract action, it subtracts the value in the active_timestamp_ns register from the value it receives from the analyzer (the original correction field from the frame) and outputs the value on the New_Field bus to the Rewriter block.

The Rewriter block continuously takes data out of the FIFO block and feeds it to the system side PCS/PMA block (or SGMII/QSGMII interface) and has a counter that keeps track of the byte positions of the frame. When the Rewriter block receives a signal from the Timestamp block (rising edge of NF) to rewrite a specific position in the frame (that information comes from the analyzer block), it overwrites the position with the data on the New_Field bus and replaces the FCS of the frame. The rewriter also checks the original FCS of the frame to ensure that a frame that is received with a bad FCS and then modified by the rewriter is also sent out with a bad FCS. This is achieved by inverting the new FCS.

The result is that the frame send towards the system now has a correction field containing the value: original correction field – RX timestamp (converted to ns).

The following full calculations are performed:

- Sync frames: Internal Correction field = Original Correction field – (Raw_Timestamp_ns – Local_correction) + Asymmetry

- Delay_req frames: Internal Correction field = Original Correction field – (Raw_Timestamp_ns – Local_correction)

3.10.1.3 Egress, Mode A

The egress side works that same way as ingress, but the analyzer is setup to add the active_timestamp to the correction field.

When a frame is received from the system side PCS/PMA (or SGMII/QSGMII) block it is loaded into a FIFO block that delays and stores the frame data for a few clock cycles to allow for later modifications of the frame. The data is also copied to the analyzer block that parses the incoming frame to detect whether it is a 1588 Sync or Delay_Req frame belonging to the PTP domain that the system is operating on.

If the analyzer detects that the frame is a 1588 Sync or Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is not matched, it signals the Timestamp block and the Rewriter block to ignore the frame (let it pass unmodified and flush the saved timestamp in the Timestamp block).

If the Timestamp block gets the Add action, it adds the current value of the active_timestamp_ns register with the value of the correction field it gets from the analyzer and outputs the value on the New_Field bus to the Rewriter block.

The Rewriter block continuously takes data out of the FIFO block and feeds it to the line side PCS/PMA block and has a counter that keeps track of the byte positions of the frame (detects SFD and counts until next SFD). When the Rewriter block receives a signal from the Analyzer block to rewrite a specific position in the frame, it overwrites the position with the data on the New_Field bus and replaces the FCS of the frame. The rewriter also checks the original FCS of the frame and ensures that a frame that is received with a bad FCS and then modified by the rewriter is also sent out with a bad FCS. This is achieved by inverting the new FCS.

The following full calculations are performed:

- Sync frames: Correction field = Internal Correction field + (Raw_Timestamp_ns + Local_correction)
- Delay_req frames: Correction field = Internal Correction field + (Raw_Timestamp_ns + Local_correction) – Asymmetry

Figure 16. One-Step E2E TC Mode B

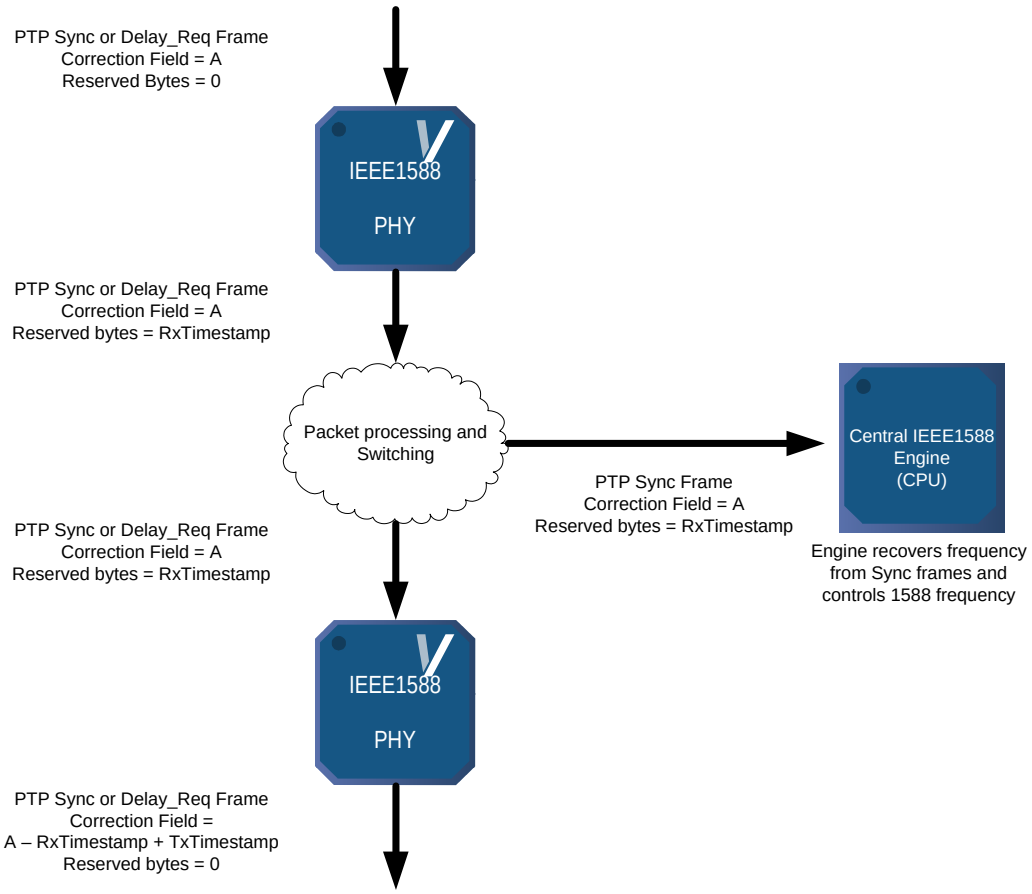

3.10.1.4 Ingress, Mode B

In Ingress Mode B, all calculations are performed at the Egress port. When the system works in one-step E2E TC mode, the system need to forward Sync and Delay_Req frames through the system and add residence time = Egress timestamp – Ingress timestamp to the correction field in the frame before it leaves the system.

On the ingress side, when the analyzer detects Sync or Delay_Req frames it adds the RX timestamp to the four reserved bytes in the PTP frame.

The following full calculations are performed:

- Sync frames: Reserved_bytes = Raw_Timestamp_ns – Local_correction + Asymmetry
- Delay_req frames: Reserved_bytes = Raw_Timestamp_ns – Local_correction

3.10.1.5 Egress, Mode B

All the calculations are done at the egress side. When the analyzer detects Sync or Delay_Req frames it performs the following calculation:

Correction field = Original correction field + TX timestamp – RX timestamp

The value of the RX timestamp is extracted from four reserved bytes in the PTP header.

The four reserved bytes are cleared back to 0 before transmission.

The result is that every Sync and Delay_Req frame that belongs to the PTP domain(s) and is configured as one-step E2E TC in the system will exit the system with a correction field that contains the following:

Correction field = Original correction field + TX timestamp – RX timestamp

All this is done without any interaction with a CPU system, other than the initial setup. There is no bandwidth expansion. Standard switching/routing tunneling can be done between the ingress and egress PHY, provided that the analyzers in the ingress PHY and egress PHY are set up to catch the sync and Delay_req on both. If the PTP sync and Delay_req frames are modified inside the system, the egress analyzer must be able to detect the egress sync and delay_req frames; otherwise, the egress sync and Delay_req frames will have an incorrect correction field.

The following full calculations are performed:

- Sync frames: Correction field = Original Correction field + (Raw_Timestamp_ns + Local_correction) – Reserved_bytes
- Delay_req frames: Correction field = Original Correction field + (Raw_Timestamp_ns + Local_correction) – Reserved_bytes – Asymmetry

3.10.2 Supporting IEEE 1588 Timestamping Applications

This section describes the integrated PTP block that supports IEEE 1588-2008 timestamping with dual encapsulation support.

Figure 17. Linecard E2E TC PHY application

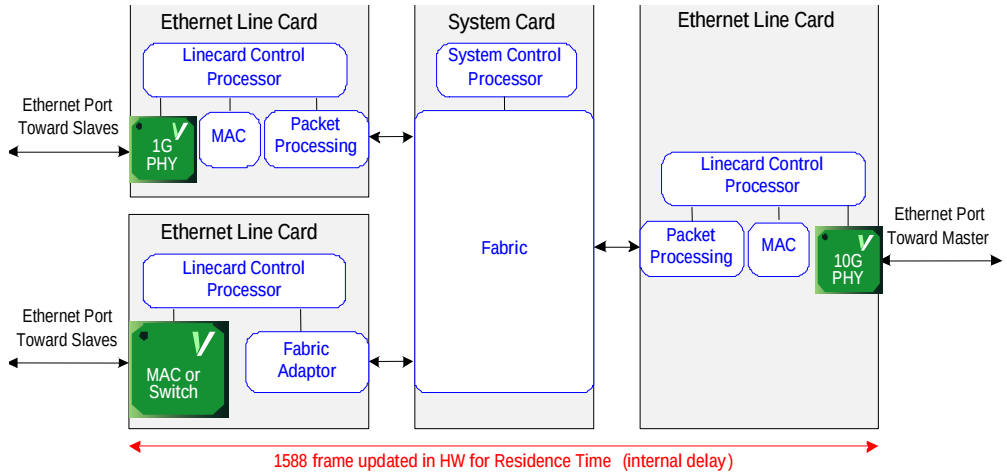

3.10.3 Application 1: IEEE 1588 One-Step E2E TC in Systems

Unique advantages for implementing IEEE 1588-2008:

- When several VSC8574 devices or Vitesse PHYs with integrated 1588 timestamping blocks are used on all ports within the system that support IEEE 1588 one-step E2E TC, the rest of the system does not need to be IEEE 1588 aware and there is no CPU maintenance needed once the system is set up
- As all the PHYs in a system can be configured the same way, it supports fail-over of IEEE 1588 masters without any CPU intervention
- VSC8574 and another Vitesse PHYs with integrated 1588 timestamping block also works for pizza box solutions, where the switch/router can be upgraded to support IEEE 1588 E2E TC
- The requirements to the rest of the system are:
- Delivery of a synchronous global timetick clock (or reference clock) to the PHYs
- Delivery of a global timetick load, that synchronizes the local time counters in each port.
- CPU access to each PHY to setup the required configuration. Can be MDC/MDIO or a dedicated CPU interface.

3.10.4 Application 2: IEEE 1588 TC and BC in Systems

This is the basically the same system as above, with the addition of a central IEEE 1588 engine (Boundary Clock). The 1588 engine is most likely a CPU system, possible together with hardware support functions to generate Sync frames (for BC and ordinary clock masters). The switch/fabric needs to have the ability to redirect (and copy) PTP frames to the 1588 engine for processing.

Figure 18. BC Linecard Application

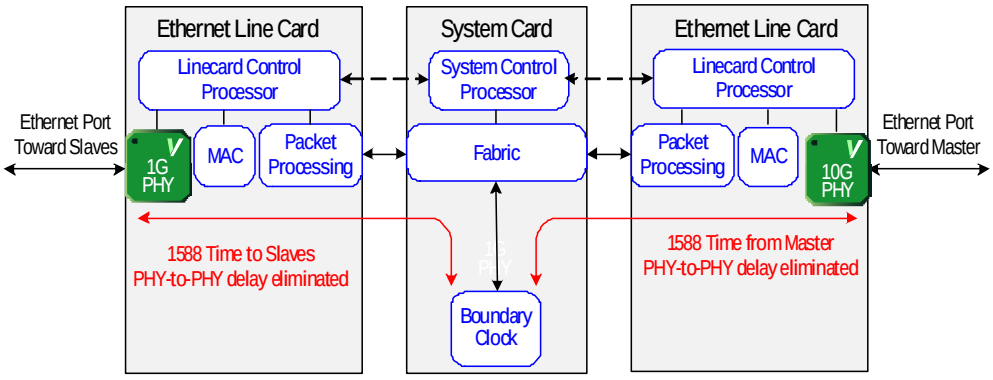

This solution also works for pizza boxes. To ensure that blade redundancy works, it the PHYs for the redundant blades must have the same 1588-in-the-PHY configuration.

The requirements to the rest of the system are:

- Delivery of a synchronous global timetick clock (or reference clock) to the PHYs
- Delivery of a global timetick load, that synchronizes the local time counters in each port
- CPU access to each PHY to setup the required configuration. For one-step support this can be MDC/MDIO. For two-step support, a higher speed CPU interface might be required (depending on the number of timestamps that are required to be read by the CPU). In blade systems it might be required to have a local CPU on the blade that collects the information and sends it to the 1588 engine by means of the control plane or the data plane. In advanced MAC/Switch devices this might be an internal CPU
- Fabric must be able to detect 1588 frames and redirect some of them to the central 1588 engine

The same solution can also be used to add Y.1731 delay measurement support. This does not require a local CPU on the blade, but the fabric must be able to redirect OAM frames to a local/central OAM processor

3.10.5 Application 3: Enhancing IEEE 1588 Accuracy for CE Switches and MACs

Connecting VSC8574 or other Vitesse PHYs have integrated 1588 time stamping in front of the CE Switches and MACs improves the accuracy of the 1588 timestamp calculation. This is due to the clock boundary for the XAUI/SGMII/QSGMII interface. It will also add support for one-step TC and BC on the Jaguar and Caracal family of devices.

3.10.6 Supporting One-Step Peer-to-Peer Transparent Clock

In P2P TC, the P2P TC device is actively sending out and receiving pDelay_Req and pDelay_Resp messages and calculating the path delays to each neighbor node in the PTP network. When a Sync frame traverses a P2P TC, the correction field is updated with both the residence time and the calculated path delay on the port that the Sync frame came in on.

Figure 19. Delay Measurements

To calculate the path delays on a link, the 1588 engine (located somewhere in the system) generates Pdelay_Req messages on all ports. When transmitted, the actual TX timestamp t3 is saved for the CPU to read.

When a P2P TC, BC, or OC receives a Pdelay_Req frame, it saves the Rx timestamp (t4) and generates a Pdelay_Resp frame, which adds t5 – t4 to the correction field copied from the received Pdelay_Req frame, where t5 is the time that the Pdelay_Resp leaves the port (t5).

When a P2P TC receives the Pdelay_Resp frame, it saves the RX timestamp (t6) and then calculates the path delay as (t6 – t3 – the correction field of the frame)/2. The timestamp corrections are combined into a single formula as follows:

Path delay = (t6 – (t3 + (t5 – t4))/2 = (t6 – t3 – t5 + t4)/2 = ((t4 – t3) + (t6 – t5))/2

The two path delays are divided by two, but in such a way as to cancel out any timing difference between the two devices.

A slight modification can be made to the algorithm to remove the hassle of the CPU reading the t3. To modify the algorithm, send the Pdelay_Req message with an origin timestamp and correction field of 0, and configure the Egress PHY to subtract the t3 timestamp from the correction field when exiting the Egress PHY. When Pdelay_resp package comes back on the link containing the value – t3 + (t5 – t4) in the correction field, the PHY is configured to add t6 to the correction field and the 1588 engine can read the resulting correction field from the frame and divide by two to get the link delay.

A P2P TC adds the calculated one-way path delay to the Ingress correction field, and this ensures that the timestamp + correction field in the egress Sync frames is accurate and a slave connected to the P2P TC only needs to take add the link delay from the TC to the slave into account.

The following section describes both the standard and modified methods for taking P2P measurements.

3.10.6.1 Ingress, Special

If the analyzer detects that the frame is a 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Subtract_p2p), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Subtract), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field in the PTP header, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Resp frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add). It also delivers the write offset and data size (location of the correction field, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Subtract_p2p action, it subtracts the value in the active_timestamp_ns_p2p register from the correction_field data and outputs the value on the New_Field bus to the Rewriter block.

If the Timestamp block gets the Subtract action, it subtracts the value in the active_timestamp_ns register from the correction field value and outputs the value on the New_field bus to the Rewriter block.

If the Timestamp block gets the Add action, it adds the correction field value to the value in the active_timestamp_ns register and outputs the value on the New_Field bus to the Rewriter block.

The following full calculations are performed:

- Sync frames: Internal Correction field = Original Correction field – (Raw_Timestamp_ns – Local_correction) + Path_delay + Asymmetry
- Pdelay_req frames: Internal Correction field = Original Correction field – (Raw_Timestamp_ns – Local_correction)
- Pdelay_resp frames: Internal Correction field = Original Correction field + (Raw_Timestamp_ns – Local_correction) + Asymmetry

3.10.6.2 Egress, Special

If the analyzer detects that the frame is a 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Subtract), along with the original correction field of the frame (will have the value of

0). It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Resp frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add), along with the original correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is not matched, it signals to the Timestamp block and the Rewriter block to ignore the frame (let it pass unmodified and flush the saved timestamp in the Timestamp block).

The following full calculations are performed:

- Sync frames: Correction field = Internal Correction field + (Raw_Timestamp_ns + Local_correction)
- Pdelay_req frames: Correction field = Internal Correction field – (Raw_Timestamp_ns + Local_correction) – Asymmetry
- Pdelay_resp frames: Correction field = Original Correction field + (Raw_Timestamp_ns + Local_correction)

Figure 20. One-Step P2P TC Standard

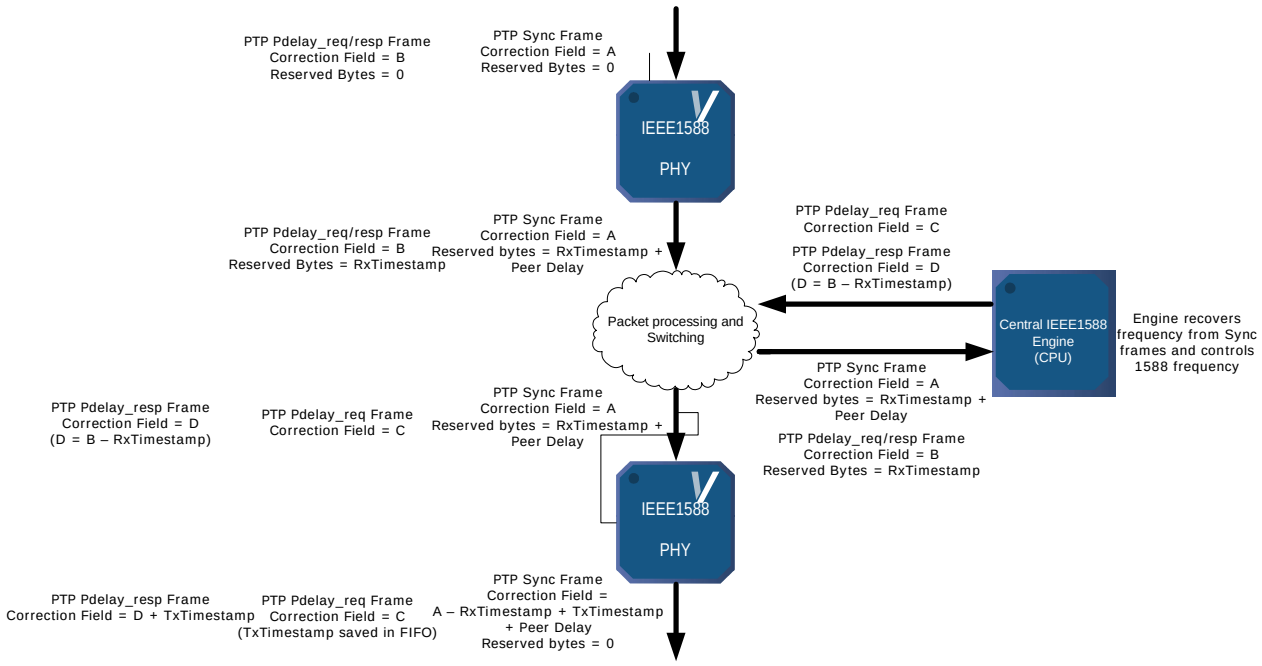

3.10.6.3 Ingress, Standard

If the analyzer detects that the frame is a 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (subtract_p2p), along with the correction field of the frame. It also delivers the write

offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the reserved 4 bytes in the PTP header we use to save the ns part of the RX timestamp, 4 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Resp frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the reserved 4 bytes in the PTP header we use to save the ns part of the RX timestamp, 4 bytes wide) to the rewriter.

If the Timestamp block gets the Subtract_p2p action, it subtracts the value in the active_timestamp_ns_p2p register from the correction_field data and outputs the value on the New_Field bus to the Rewriter block.

If the Timestamp block gets the Write action, it outputs the value of the active_timestamp_ns register on the New_field bus to the Rewriter block.

The following full calculations are performed:

- Sync frames: Internal Correction field = Original Correction field – (Raw_Timestamp_ns – Local_correction) + Path_delay + Asymmetry
- Pdelay_req frames: Reserved_bytes = Raw_Timestamp_ns – Local_correction
- Pdelay_resp frames: Reserved_bytes = Raw_Timestamp_ns – Local_correction + Asymmetry

3.10.6.4 Egress, Standard

If the analyzer detects that the frame is a 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add), along with the correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Pdelay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write, Save), along with the original correction field of the frame (will have the value of 0). It also delivers the write offset and data size (0 No data is actually written into the frame) to the rewriter. In addition it outputs the field that holds the frame identifier (sequenceId from the PTP header) to the timestamp FIFO, to save along with the TX timestamp.

If the analyzer detects that the frame is a 1588 Pdelay_Resp frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Add - this requires that the 1588 engine has subtracted the RX timestamp from the correction field), along with the original correction field of the frame. It also delivers the write offset and data size (location of the correction field inside the frame, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Write, Save action it outputs the value of the active_timestamp_ns register on the New_field bus to the Rewriter block and sets the save_timestamp bit.

If the Timestamp block gets the Add action, it adds the correction field value to the value in the active_timestamp_ns register and outputs the value on the New_Field bus to the Rewriter block.

The TX Timestamp FIFO block contains an (implementation specific) amount of buffer memory. It simply stores the TX timestamp values that it receives from the Timestamp block together with the frame identifier data it receives from the Analyzer block and has a CPU interface that allows the 1588 engine to read out the timestamp sets (Frame identifier + New TX timestamp).

The following full calculations are performed:

- Sync frames: Correction field = Internal Correction field + (Raw_Timestamp_ns + Local_correction)
- Pdelay_req frames: FIFO = Raw_Timestamp_ns + Local_correction – Asymmetry
- Pdelay_resp frames: Correction field = Internal Correction field + (Raw_Timestamp_ns + Local_correction)

3.10.7 Supporting One-Step Boundary Clock / Ordinary Clock

In one-step boundary clock, the BC device acts as an ordinary clock slave on one port and as master on all the other ports. On the master ports, Sync frames are transmitted from the 1588 engine that holds the Origin timestamp. These frames will have the correction field or the full TX timestamp updated on the way out though the PHY.

Master ports also receive Delay_req from the slaves and respond with Delay_resp messages. The Delay_req messages are time stamped on the way through the PHY and the 1588 engine receives the Delay_req frame and generates a Delay_resp message. The Delay_resp messages are not event messages and are passed though the PHY as any other frame.

The port that is configured as slave receives Sync frames from its master. The Sync frames get an RX timestamp added in the PHY and forwarded to the 1588 engine.

The 1588 engine also generates Delay_req frames that are send out on the port that is configures as slave port.

Boundary clocks and ordinary clocks must also reply to Pdelay_req messages just as P2P TC, but the procedure is the same as for the P2P TC, so this is not described here.

Figure 21. One-Step E2E BC

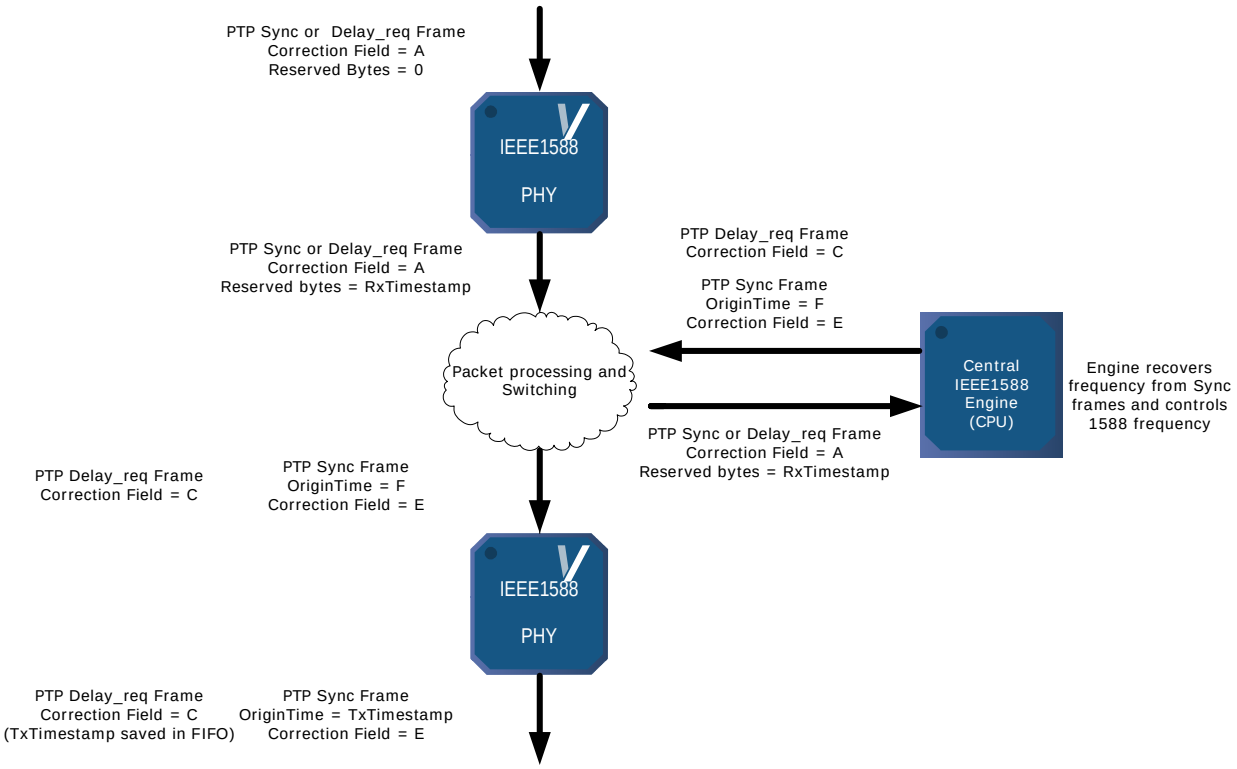

3.10.7.1 Ingress

If the analyzer detects that the frame is a 1588 Sync or Delay_req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the four reserved bytes in the PTP header, 4 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds this timestamp (ns part of it) to the four reserved bytes in the frame and recalculates FCS.

The following full calculations are performed:

- Sync frames: Reserved_bytes = (Raw_Timestamp_ns – Local_correction) + Asymmetry
- Delay_req frames: Reserved_bytes = (Raw_Timestamp_ns – Local_correction)

3.10.7.2 Egress

If the analyzer detects that the frame is a 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write), along with correction field of the frame (contains the offset from the system PTP domain to any other domain that the frame belong to - set by the 1588 engine that

generates the frame). It also delivers the write offset and data size (location of the TX timestamp inside the frame, 10 bytes wide) to the rewriter.

If the analyzer detects that the frame is a 1588 Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write, Save). It also delivers the write offset and data size (location of the TX timestamp inside the frame, 10 bytes wide to the rewriter. It also outputs 10 bytes of frame identifier to the TX Timestamp FIFO, to be saved along with the TX timestamp.

If the Timestamp block gets the Write, Save action it outputs the current value from the active_timestamp register on the New_field bus to the Rewriter Timestamp_fifo blocks and signals to the Timestamp FIFO block that it must save the New_field data along with the frame identifier data it received from the Analyzer block.

The TX Timestamp FIFO block contains an (implementation specific) amount of buffer memory. It simply stores the TX timestamp values that it receives from the Timestamp block together with the frame identifier data it receives from the Analyzer block and has a CPU interface that allows the 1588 engine to read out the timestamp sets (Frame identifier + New TX timestamp).

The following full calculations are performed:

- Sync frames: OriginTimestamp = (Raw_Timestamp + Local_correction)
- Delay_req frames: OriginTimestamp = (Raw_Timestamp + Local_correction) – Asymmetry

3.10.8 Supporting Two-Step Boundary / Ordinary Clock

Two-steps clocks are used in systems that cannot update the correction field on-the-fly and this requires more CPU power.

Every time a TX timestamp is sent in a frame, the 1588 engine needs to read the actual TX transmission time from the Timestamp FIFO and issue a follow-up message containing this timestamp.

Figure 22. Two-Step E2E BC

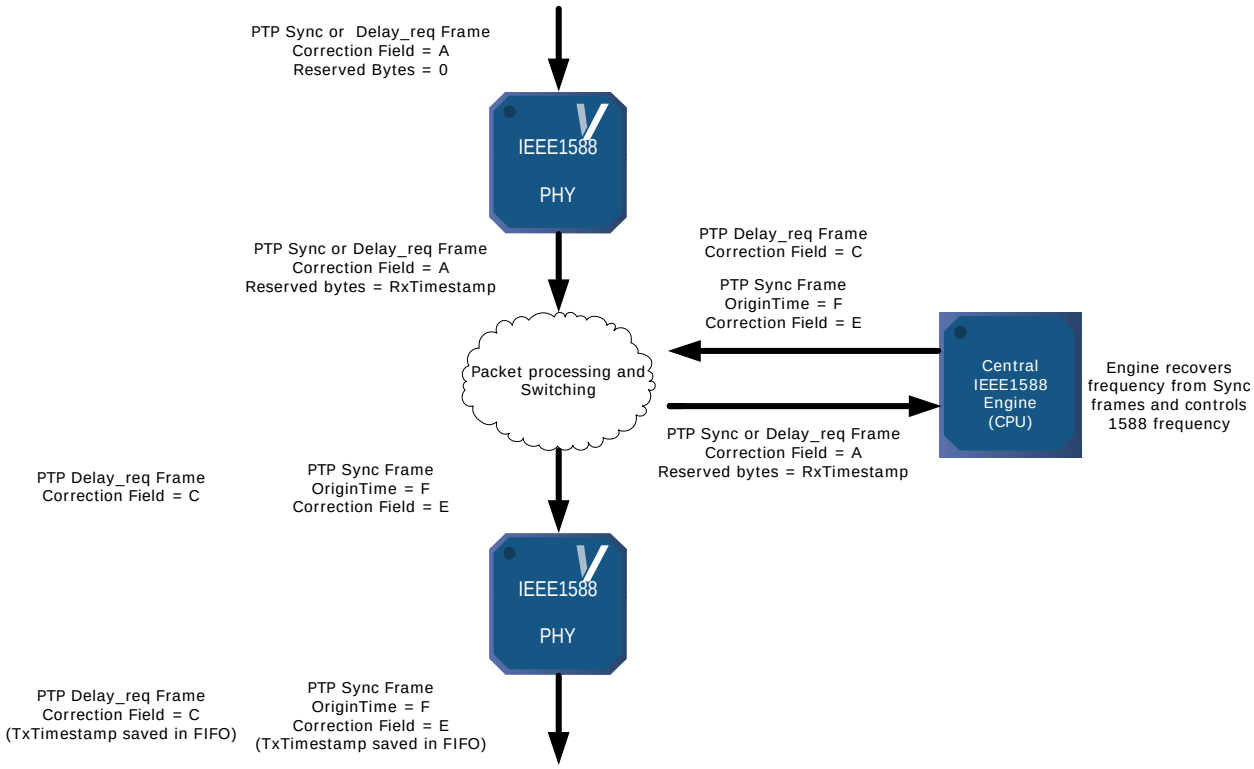

3.10.8.1 Ingress

If the analyzer detects that the frame is a 1588 Sync or Delay_req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the four reserved bytes in the PTP header, 4 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds this timestamp (ns part of it) to the four reserved bytes in the frame and recalculates FCS.

Note When secure timing delivery is required, the 1588 engine must revert the four reserved bytes back to 0 before performing integrity check.

The following full calculations are performed:

- Sync frames: Reserved_bytes = (Raw_Timestamp – Local_correction) + Asymmetry
- Delay_req frames: Reserved_bytes = (Raw_Timestamp – Local_correction)

3.10.8.2 Egress

If the analyzer detects that the frame is a 1588 Sync or Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform

(Write, Save). The analyzer also delivers the write offset and data size (but as nothing is to be overwritten the values will be 0) to the rewriter. The analyzer outputs up to 15 bytes of frame identifier to the TX Timestamp FIFO to be saved along with the TX timestamp. The frame identifier must include, at minimum, the sequenceId field so the CPU can match the timestamp with the follow-up frame.

If the Timestamp block gets the Write, Save action it outputs the current value from the active_timestamp register on the New_field bus to the Rewriter (and timestamp FIFO) and sets the save_timestamp bit. The Timestamp_fifo block saves the New_field data along with the frame identifier data it received from the Analyzer block.

The following full calculations are performed:

- Sync frames: FIFO = (Raw_Timestamp + Local_correction)
- Delay_req frames: FIFO = (Raw_Timestamp + Local_correction) – Asymmetry

3.10.9 Supporting Two-Step Transparent Clock

In two-step transparent clocks, the RX and TX timestamps are saved for the 1588 Engine to read and the follow-up message is redirected to the 1588 engine so that it can update the correction field with the residence time.

Even though two-step transparent clocks can be used with this architecture, it is also possible to process the frames in the same manner as a one-step TC, because the slaves are required to take both the corrections fields from the Sync frames and the follow-up frames into account. This significantly reduces the CPU load for the TC. The following illustration shows two-step transparent clock normal operation.

Figure 23. Two-Step E2E TC

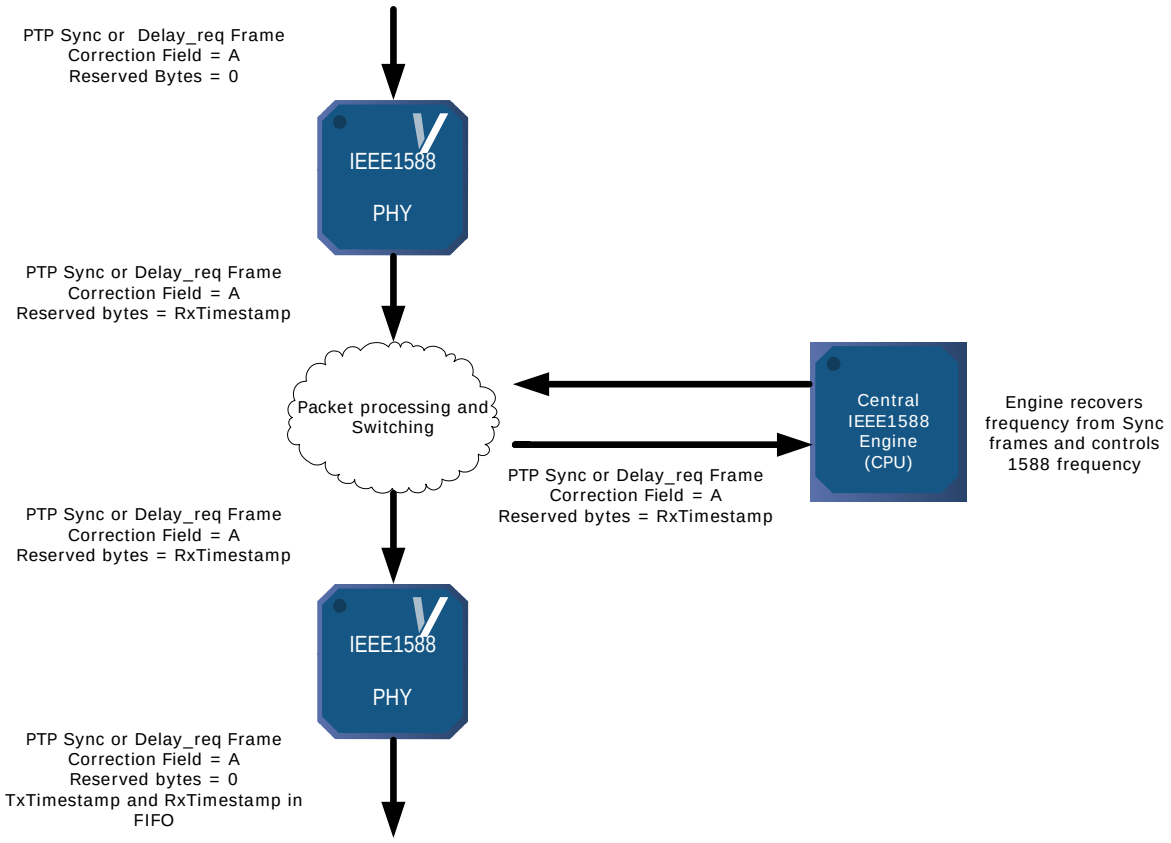

3.10.9.1 Ingress

If the analyzer detects that the frame is a 1588 Sync or Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write). The analyzer also delivers the write offset and data size to the rewriter (four reserved bytes in the PTP header, which will be passed out on the egress port of the system). A changed reserved value may be significant in security protection. This method allowed the frames to be copied to the 1588 engine, so that it can extract the RX timestamp and that it knows that it needs to read the TX timestamps to be ready for the follow up message. It is also possible to save the RX timestamp value along with the TX timestamp in the TX timestamp FIFO.

If the Timestamp block gets the Write action, it outputs the current value from the active_timestamp register on the New_field bus to the Rewriter and the rewriter writes the ns part of the timestamp into the reserved bytes and recalculates FCS.

The following full calculations are performed:

- Sync frames: Reserved_bytes = (Raw_Timestamp_ns – Local_correction) + Asymmetry
- Delay_req frames: Reserved_bytes = Raw_Timestamp_ns – Local_correction

3.10.9.2 Egress

If the analyzer detects that the frame is a 1588 Sync or Delay_Req frame belonging to the PTP domain(s) of system, it signals to the timestamp block which action to perform (Write, Save). The analyzer also delivers the write offset and data size (but as nothing is to be overwritten the values will be 0) to the rewriter. The analyzer outputs 10 bytes of frame identifier to the TX Timestamp FIFO to be saved along with the TX timestamp. The frame identifier must include, at minimum, the sequenceId field so the CPU can match the timestamp with the follow-up frame. The analyzer also outputs the offset for the reserved fields in the PTP header to the rewriter, so that the rewriter field is reset to 0 and the temporary RX timestamp value is cleared.

If the Timestamp block gets the Write, Save action it outputs the current value from the active_timestamp register on the New_field bus to the Rewriter (and timestamp FIFO) and sets the save_timestamp bit. The Timestamp_fifo block saves the New_field data along with the frame identifier data it received from the Analyzer block. The frame identifier data that is saved can contain the reserved field in the PTP header that was written with the RX timestamp, so that the CPU now can read the set of TX and RX timestamp from the TX timestamp FIFO.

The following full calculations are performed:

- Sync frames: FIFO = Raw_Timestamp_ns + Local_correction (reserved_bytes containing the RX timestamp saved together with TX timestamp)
- Delay_req frames: FIFO = Raw_Timestamp_ns + Local_correction – Asymmetry (reserved_bytes containing the RX timestamp saved together with TX timestamp)

3.10.10 Calculating Y.1731 OAM Delay Measurements

Frame delay measurements can be made as one-way and two-way delay measurements. Vitesse recommends that the delay measurement be measured before the packets enter the queues, if the purpose is to measure the delay for different priority traffic, but it can be used with timestamping in the PHY to measure the delay through the network devices placed in the path between the measurement points.

The function is mainly an on-demand OAM function, but it can run continuously.

3.10.11 One-Way Delay Measurements

One-way delay measurements require that the two peers are synchronized in time. When they are not synchronized, only frame delay variations can be measured.

The MEP periodically sends out 1DM OAM frames containing a TxTimeStampf value in IEEE 1588 format.

The receiver notes the time of reception of the 1DM frame and calculates the delay.

Figure 24. Y.1731 1DM PDU Format

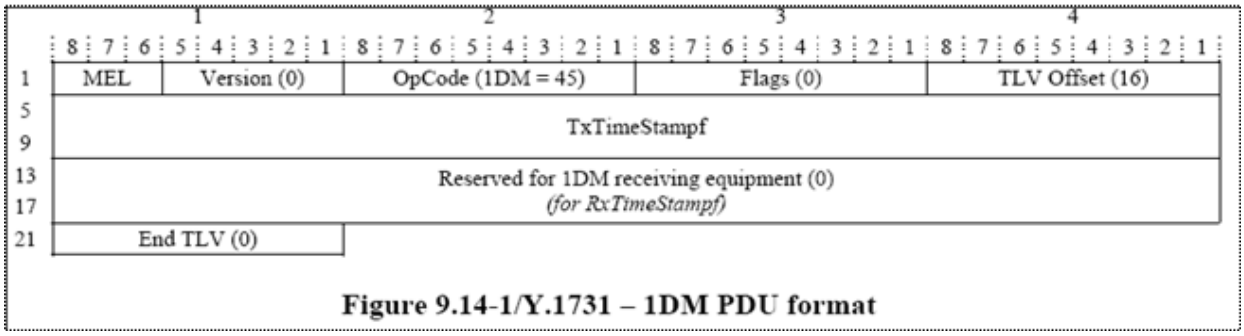

Figure 9.14-1/Y.1731 – 1DM PDU format

1. For one-way delay measurements, both MEPs must support IEEE 1588 and be in sync.
2. 1DM frame is generated by the CPU, but with an empty Tx timestamp.
3. The frame is transmitted by the initiating MEP.
4. The 1DM frame is classified as an outgoing 1DM frame by the Egress PHY and the PHY rewrites the frame with the time as TxFCf.
5. The receiving PHY classifies the incoming 1DM frame and writes the receive timestamp in reserved place (RxTimeStampf).
6. The frame is received by the peer MEP.
7. The frame is forwarded to the CPU that can calculate the delay.

Figure 25. Y.1731 One-Way Delay

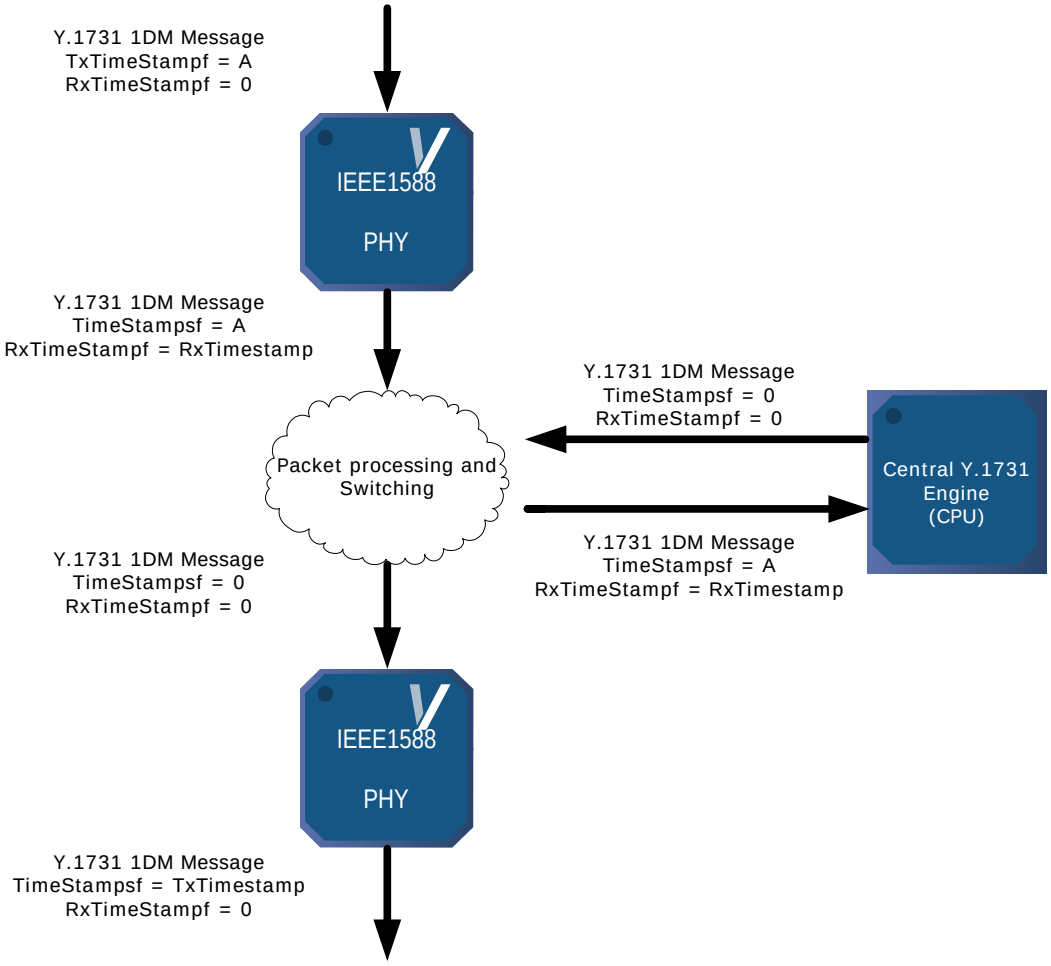

3.10.11.1 Ingress

If the analyzer detects that the frame is a Y.1731 1DM PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). The analyzer also delivers the write offset and data size (location of the RxTimeStampf location in the frame, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds this timestamp to the reserved bytes in the frame and recalculates FCS.

The following calculation is performed for 1DM frames:

RxTimeStampf = (Raw_Timestamp – Local_correction)

3.10.11.2 Egress

If the analyzer detects that the frame is a Y.1731 1DM PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the TxTimeStampf location in the frame, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds this timestamp to the reserved bytes in the frame and recalculates FCS.

The following calculation is performed for 1DM frames:

TxTimeStampf = (Raw_Timestamp + Local_correction)

3.10.12 Two-Way Delay Measurements

When performing two-way delay measurements, the initiating MEP transmits DMM frames containing a TxTimeStampf value. The receiving MEP replies with a DMR frame that is the same as the DMM frame, but with destination and source MAC address swapped and with a different OAMPDU opcode.

When the DMR frame is received back at the initiating MEP, the time of reception is noted and the total delay is calculated.

As an option, it is allowed to include two additional timestamps in the DMR frame: RxTimeStampf and TxTimeStampb. These contain the time that the DMM page is received for processing and the time the responding DMR reply is sent back, both in IEEE 1588 format.

Including these timestamps allow for exclusion of the processing time in the peer MEP, but it does not require that the two MEPs are synchronized.

Figure 26. Y.1731 DMM PDU Format

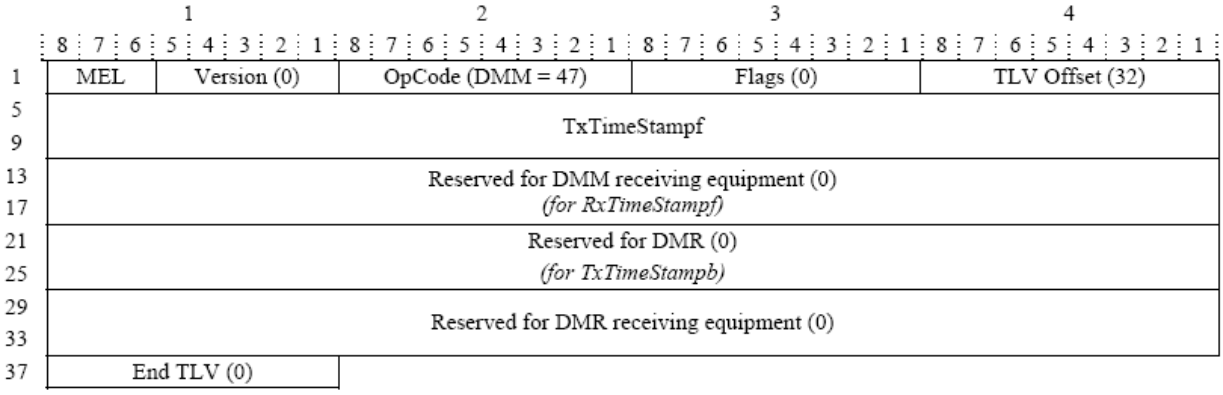

Figure 9.15-1/Y.1731 – DMM PDU format

In that case, the following frame flow is needed (two-way delay measurement):

1. DMM frame is generated by the CPU (initiating MEP), but with an empty Tx timestamp.
2. In the egress PHY the DMM frame is classified as an outgoing DMM frame from the MEP and the PHY rewrites the frame with the time as TxTimeStampf.
3. In the ingress PHY the frame is classified as an incoming DMM belonging to the MEP and the RxTimeStampf in the frame is written (the frame has a reserved space for this).
4. The DMM frame is forwarded to the MEP (CPU).
5. The CPU processes the frame (swaps SA/DA MAC addresses, modifies the opcode to DMT) and sends out a DMT frame.
6. The outgoing DMT frame is detected in the egress PHY and the TxTimeStampb is written into the frame.
7. In the ingress PHY the frame is classified as an incoming DMT belonging to the MEP and the RxTimeStampb in the frame in written (the frame has a reserved space for this).
8. The frame is forwarded to the CPU that can calculate the delays.

Figure 27. Y.1731 Two-Way Delay

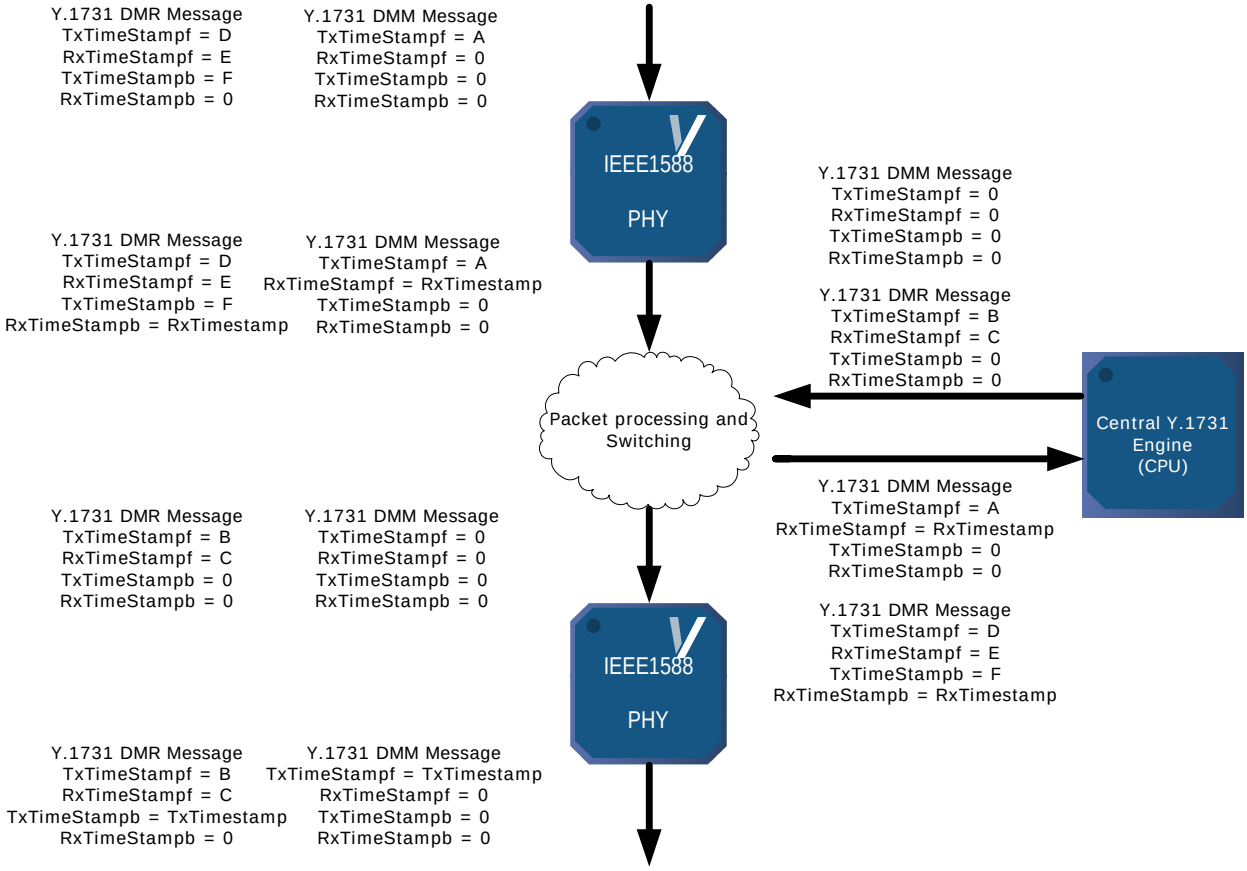

3.10.12.1 Ingress

If the analyzer detects that the frame is a Y.1731 DMM PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the RxTimeStampf location in the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a Y.1731 DMT PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the RxTimeStampb location in the frame, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds this timestamp to the reserved bytes in the frame and recalculates FCS.

The following calculations are performed:

- DMM frames: RxTimeStampf = (Raw_Timestamp – Local_correction)
- DMT frames: RxTimeStampb = (Raw_Timestamp – Local_correction)

3.10.12.2 Egress

If the analyzer detects that the frame is a Y.1731 DMM PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the TxTimeStampf location in the frame, 8 bytes wide) to the rewriter.

If the analyzer detects that the frame is a Y.1731 DMT PDU frame belonging to the MEP, it signals to the timestamp block which action to perform (Write). It also delivers the write offset and data size (location of the TxTimeStampb location in the frame, 8 bytes wide) to the rewriter.

If the Timestamp block gets the Write action, it puts the value of the active_timestamp register out on the New_field bus to the Rewriter block and the rewriter block adds the timestamp to the reserved bytes in the frame and recalculates FCS as follows:

- DMM frames: TxTimeStampf = (Raw_Timestamp + Local_correction)
- DMT frames: TxTimeStampb = (Raw_Timestamp + Local_correction)

3.10.13 IEEE 1588 Device Synchronization

It is important to keep all the Local Clock blocks synchronized to the accurate time over a complete system. To maintain ns accuracy, the signal routing and internal signal delays must be taken into account when configuring a system.

The architecture described in this document assumes that there is a global synchronous clock available in the system. If the system is a telecom system where the system is locked to a PRC, the system clock can be adjusted to match the PRC, meaning that once locked, the frequency of the system clock ensures that the local clocks are progressing (counting) with the accurate frequency. This system clock can be locked to the PRC using 1588, SyncE, SDH, or by other means.

A global timing signal must also be distributed to all the devices. This could be a 1 pps pulse or another slow synchronization pulse, like a 4 kHz synchronization frequency. It can also just be a one-shot pulse. The system CPU can load each local counter with the time value that happens next time the synchronization pulse goes high (+ the known delay of the synchronization pulse traces). It can also just load the same approximate time value into all the local clock blocks (again + the known delay of the synchronization pulse traces) and load them in parallel. Then the local time can be adjusted to match the actual time by adjusting the local clock blocks using the ±1 ns function.

If the Save signal is triggered synchronously on all PHYs of the system, software can read the saved timestamp in each PHY and correct the time accordingly. On a blade with multiple PHYs, it is possible to connect the 1588_PPS_1 pin on one PHY to the 1588_LOAD_SAVE pin on the next PHY. If the routing delay (both internal chip delay and trace delay) is known, Vitesse recommends that the value saved in the next PHYs correspond to this delay.

If the global system clock is not synchronous, the PPM offset between system clock and the 1588 time progress can be calculated. This PPM offset can be used to calculate how many local-time-clocks is takes to reach a time offset of 1 ns and this value can be programmed into each local time block. The CPU still need to keep track of the smaller PPM offset and adjust the local time blocks with ± writes when necessary.

By measuring the skew between the 1 pps test output from each PHY, it is possible to measure the nominal correction values for the time counters in a system. These can be incorporated into the software of the system. Variations from system to system and temperature variations should be minimized by design.

3.10.14 Time Stamp Update Block

The IEEE1588 block is also called the Time Stamp Update block (TSU) and supports the implementation of IEEE 1588v2 and ITU-T Y.1731 in PHY hardware by providing a mechanism for time-stamp update (PTP) and time-stamping (OAM).

The TSU block works with other blocks to identify PTP/OAM messages, process these messages, and insert accurate timestamp updates/timestamps where necessary. For 1588 timing distribution the VSC8574 device supports ordinary clocks, boundary clocks, end-to-end transparent clocks, and peer-to-peer transparent clocks in a chassis based 1588 capable system. One-step and two-step processing is also supported. For details on the timing protocol, refer to IEEE 1588v2. For OAM details refer to ITU-T Y.1731. The TSU block implements part of the functionality required for full 1588 compliance.

The 1588 protocol has four different types of messages that require action by the TSU: Sync, Delay_Req, Pdelay_Req, and Pdelay_Resp. These frames may be encapsulated in other protocols, several layers deep. The processor is able to detect PTP messages within these other protocols. The supported encapsulations are as follows:

- Ethernet
- UDP over IPv4
- UDP over IPv6
- MPLS

- Pseudo-wires
- PBB and PBB-TE tunnels
- IP/IP tunnel

ITU-T Y.1731 OAM frames for delay measurement (1DM, DMM, and DMR) with the following supported encapsulations:

- Ethernet (Y.1731 Ethernet OAM)
- Ethernet in MPLS pseudo-wires (Y.1731 Ethernet OAM)
- MPLS (MPLS-TP Y.1731-like OAM via ACH as described in draft-bhh-mpls-tp-oam-y1731-04.txt)

The following illustration shows an overview of the supported PTP encapsulations. Note that the implementation is flexible such that encapsulations not defined here may also be covered.

Figure 28. PTP Packet Encapsulations

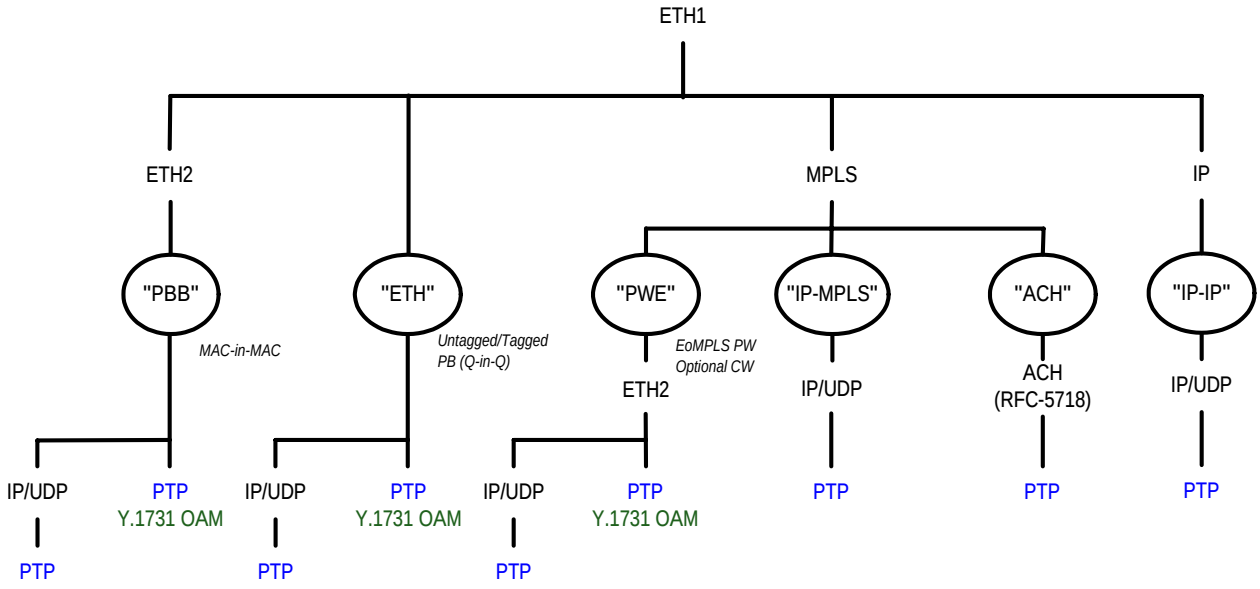

The following illustration shows the same overview of the supported encapsulations with the focus on OAM.

Figure 29. OAM Packet Encapsulations

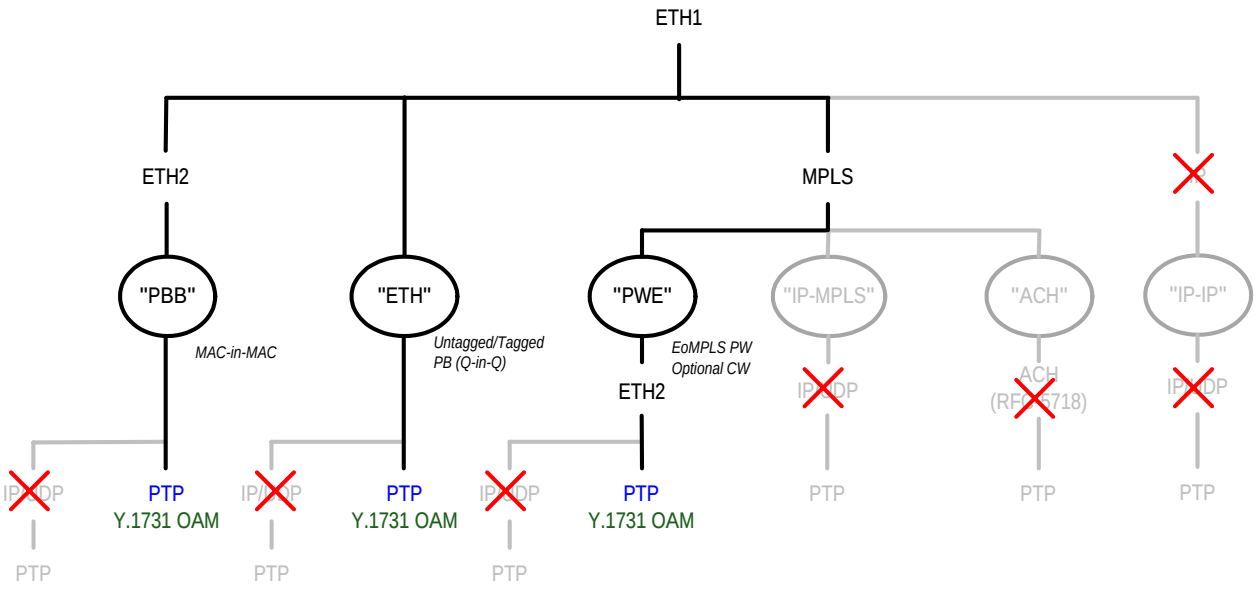

There is one TSU per channel in the VSC8574 device. The TSU detects and updates up to two different encapsulations of PTP/OAM. Non-matching frames are transferred transparently. This includes IFG, preamble, and SFD. For all frames there is no bandwidth expansion/shrink.

Once these frames are detected in the receive path, they are stamped with the ingress time and forwarded for further PTP/OAM processing. In the transmit path, the correction field of the appropriate PTP message (or the Rx and Tx fields of the OAM frame) is updated with the correct timestamp. A local time counter is maintained to provide the timestamps. Full implementation of the 1588 protocol requires interaction with the TSU block over the CPU interface and external processing.

The system has an ingress processor, egress processor, and a local time counter. The ingress and egress processing logic blocks are identical except that the timestamp FIFO is only required in the egress direction because the CPU needs to know the actual timestamps of the transmitted PTP frames. The CPU reads the timestamps and any associated frame information out of the timestamp FIFO. The FIFO saves the generated timestamps along with information that uniquely identifies the frame to be read out by the CPU.

The ingress and egress processing blocks run on the same clock as the data paths for the corresponding directions. The local time counter is the primary reference clock for the system and it maintains the local reference time used by the TSU logic. It should be synchronized by an external entity. The block provides a method to load and view its value when the 1588_LOAD_SAVE pin is asserted. The block also provides a one pulse-per-second output signal with a programmable duty cycle. The local time counter runs at several clock frequencies.

The following illustration shows the block diagram of the TSU.

Figure 30. TSU Block Diagram

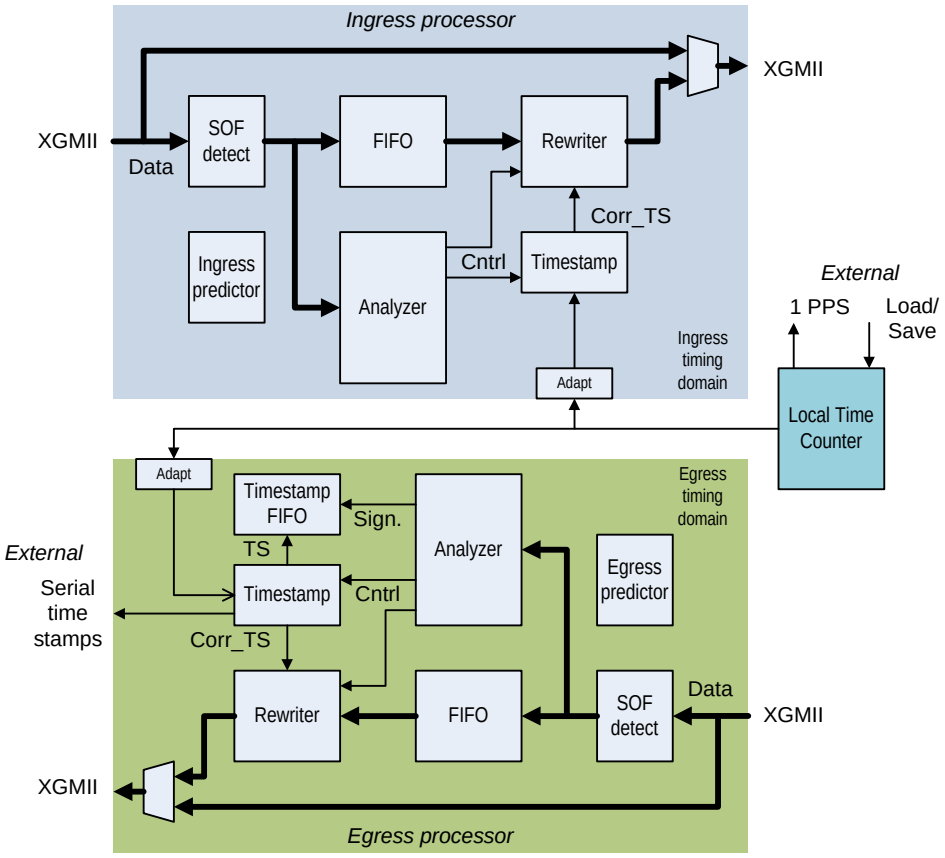

In both directions, the input data from the PHY layer is first fed to an SOF detect block. Data is then fed to both the programmable time-delay FIFO and the analyzer. The FIFO delays the data by the time needed to complete the operations necessary to update the PTP frame. That is, the data is delayed to the input of the rewriter so that the rewriter operations are known when the frame arrives. This includes the analyzer and timestamp processor block's functions.

The analyzer block checks the data stream and searches for PTP/OAM frames. When one is detected, it determines the appropriate operations to be performed based on the operating mode and the type of frame detected.

Note The analyzer blocks of different channels share configuration registers and have identical setups.

The timestamp block waits for an SOF to be detected, captures a timestamp from the local time counter, and builds the new timestamp that is to be written into the PTP/OAM frame. Captured timestamps can be read by the CPU and via a serial interface (GPIO).

The rewriter block handles the actual writing of the new timestamp into the PTP/OAM frame. It is also able to clear parts of the frame such as the UDP checksum, if required. The block also calculates the new FCS to be written to the PTP frame after updating the fields with the new timestamp.

The VSC8574 device has variable latency in the PCS and WIS blocks. These variations are predicted and used to compensate/maximize the accuracy of the 1588 timestamp logic. A predictor is required in the ingress and egress paths to calculate the variable delay and add it to the timestamp.

If the time stamp update function is not to be used the block can be bypassed. This bypass is intelligent in that it only applies the asserted register bit when it detects an IDLE column ||I|| in the bypass path registers. This allows a seamless transition from bypass to active data path at a safe time during startup. When the TSU is bypassed, the block can be configured and then enabled and taken out of bypass mode. The change in bypass mode takes effect only when an IDLE is in the bypass register. This allows the TSU block to be switched on without corrupting data. The data path blocks in the TSU have IDLE columns in them after reset.

The TSU has the following known limitation.

- Pause frames can pass unmodified, but the latency may cause a violation of the allowed pause flow-control latency limits per IEEE 802.3.

3.10.15 Analyzer

The packet analyzer parses incoming packets looking for PTP/OAM frames. It determines the offset of the correction field within the packet for all PTP frames/for the time stamp in Y.1731 OAM frames. The analyzer has the following characteristics:

- Can compare against two different filter sets plus one optimized for OAM
- Each filter targets PTP or OAM frames
- Flexible comparator sequence with fixed start (Ethernet/SNAP) and end (PTP/OAM) comparator. Configurable intermediate comparators (Ethernet/SNAP, 2x IP/UDP/ ACH, and MPLS)

The following illustration shows a block diagram of the analyzer.

Figure 31. Analyzer Block Diagram

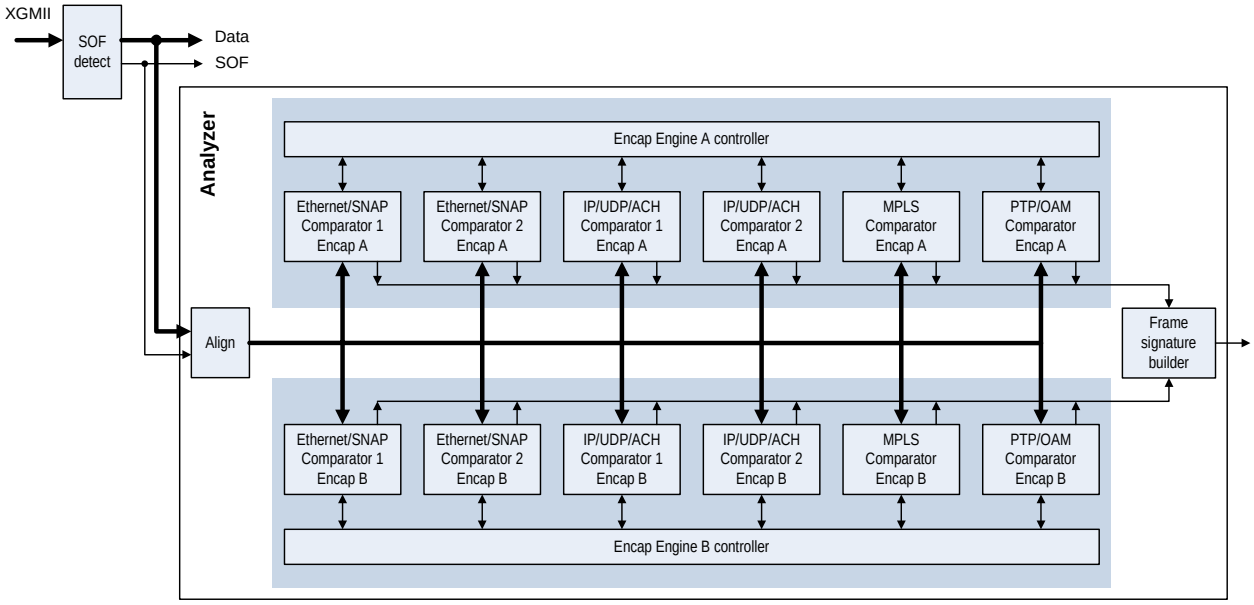

The analyzer process is divided into engines and stages. Each engine represents a particular encapsulation stack that must be matched. There are up to six stages in each engine. Each stage uses a comparator block that looks for a particular protocol. The comparison is performed stage-by-stage until the entire frame header has been parsed.

Each engine has its own master enable, so that it can be shut down for major reconfiguration, such as changes in encapsulation order, without stopping traffic. Other enabled engines are not affected.

The SOF detect block searches for the SFD in the preamble and uses that to indicate the SOF position. This information is carried along in the pipeline and also passed to the analyzer.

The first stage of the analyzer is a data path aligner that aligns the first byte of the packet (without the preamble & SFD) to byte 0 of the analyzer data path.

The encapsulation engine handles numerous types of encapsulation stacks. These can be broken down to their individual protocols, and a comparator is defined for each type. The order in which these are applied is configurable. Each comparator outputs a pattern/flow match bit and an offset to the start of the next protocol. The cumulative offset points to the time stamp field.

The sequence in which the protocol comparators are applied is determined by configuration registers associated with each comparator and the transfer of parameters between comparators is controlled by the encapsulation engine controller.

It receives the pattern match and offset information from one comparator stage and feeds the start-of-protocol position to the next comparator. This continues until the entire encapsulation stack has been parsed and always ends with the PTP/OAM stage or until a particular comparator stage cannot find a match in any of its flows. If at any point along the way no valid match is found in a particular stage, the analyzer sends the NOP communication to the timestamp block indicating that this frame does not need modification and that it should discard its timestamp.

There are two types of engines in the analyzer, one optimized for PTP frames and the other optimized for OAM frames. The two engine types are mostly identical except that the IP comparators are removed from the OAM engines. The following table shows the comparator layout per engine type and the number of flows in each comparator. There are two PTP engines and one OAM engine in each analyzer. Additional differences in the Ethernet and MPLS blocks are defined in their respective sections. For more information, see "Ethernet/SNAP/LLC Comparator," page 64 and "MPLS Comparator," page 68.

Table 7. Flows Per Engine Type

	Number of Flows	
Comparator	**PTP Engine**	**OAM Engine**
Ethernet 1	8	8
Ethernet 2	8	8
MPLS	8	8
IP/ACH 1	8	0
IP/ACH 2	8	0
PTP/OAM	6	6

Each comparator stage has an offset register that points to the beginning of the next protocol relative to the start of the current one. The offset is in bytes, and the first byte of the current protocol counts as byte 0. As an example, the offset register for a stage would be programmed to 10 when the header to match is 10 byte long. With the exception of the MPLS stage (offsets are automatically calculated in that stage), it is the responsibility of the programmer to determine the value to put in these registers. This value must be calculated based upon the expected length of the header and is not expected to change from frame-to-frame when matching a given flow.

Table 8. Ethernet Comparator: Next Protocol

Parameter	Width	Description
Encap_Engine_ENA	1 bit	For each encapsulation engine and enable bit that turns the engine on or off. The engine enables and disables either during IDLE (all 8 bytes must be IDLE) or at the end of a frame. If the enable bit is changed during the middle of a frame, the engine will wait until it sees either of those conditions before turning on or off.
Encap_Flow_Mode	1 bit	There is a separate bit for each engine. For each encapsulation engine: 1 = Strict flow matching, a valid frame must use the same flow IDs in all comparators in the engine except the PTP and MPLS comparators. 0 = A valid frame may match any enabled flow in all comparators If more than one encapsulation produces a match, the analyzer sends NOP to the rewriter and sets a sticky bit.

The following table shows the ID codes comparators use in the sequencing registers. The PTP packet target encapsulations require only up to five comparators.

Table 9. Comparator ID Codes

ID	Name	Sequence
0	Ethernet Comparator 1	Must be the first
1	Ethernet Comparator 2	Intermediate
2	IP/UDP/ACH Comparator 1	Intermediate
3	IP/UDP/ACH Comparator 2	Intermediate
4	MPLS Comparator	Intermediate
5	PTP/Y.1731 OAM Comparator	Must be the last

The following sections describe the comparators. The frame format of each comparator type is described first, followed by match/mask parameter definition. All upper and lower bound ranges are inclusive and all match/mask registers work the same way. If the corresponding mask bit is 1, then the match bit is compared to the incoming frame. If a mask bit is 0, then the corresponding match bit is ignored (a wildcard).

3.10.15.1 Ethernet/SNAP/LLC Comparator

There are two such comparators in each engine. The first stage of each engine is always an Ethernet/SNAP/LLC comparator. The other comparator can be configured to be at any point in the chain.

Ethernet frames can have multiple formats. Frames that have an actual length value in the ether-type field (Ethernet type I) can have one of three formats: Ethernet with an EtherType (Ethernet type II), Ethernet with LLC, or Ethernet with LLC & SNAP. Each of these formats can be compounded by having one or two VLAN tags.

TYPE II ETHERNET

Type II Ethernet is the most common and basic type of Ethernet frame. The Length/EtherType field contains an EtherType value and either 0, 1, or 2 VLAN tags. Both VLAN can be of type S/C (with EtherType 0x8a88/0x8100). The payload would be the start of the next protocol.

Figure 32. Type II Ethernet Basic Frame Format

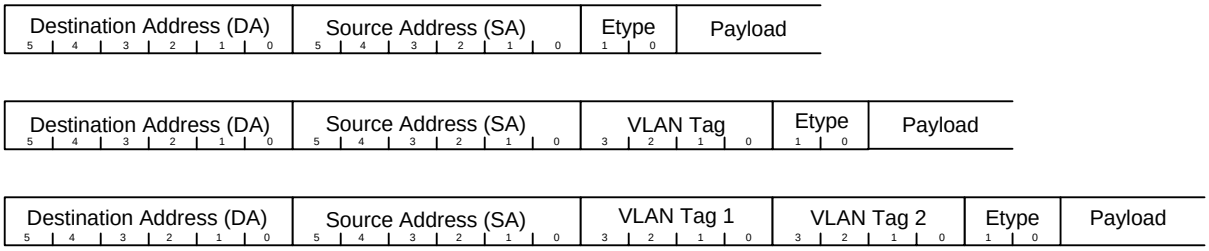

Ethernet WITH LLC AND SNAP

If an Ethernet frame with LLC contains a SNAP header, it always follows a three-octet LLC header. The LLC values for DSAP & SSAP are either 0xAA or 0xAB and the control field contains 0x03. The SNAP header is five octets long and consists of two fields, the

3-octet OUI value and the 2-octet EtherType. As with the other types of Ethernet frames, this format can have 0, 1, or 2 VLAN tags. The OUI portion of the SNAP header is hard configured to be 0 or 0xf8.

The following illustration shows an Ethernet frame with a length in the Length/ EtherType field, an LLC header, and a SNAP header.

Figure 33. Ethernet Frame with SNAP

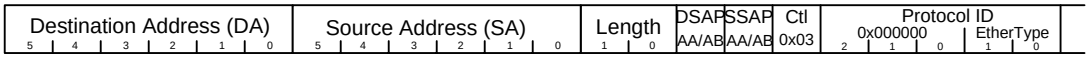

The following illustration shows an Ethernet frame with an LLC/SNAP header and a VLAN tag in the SNAP header. The Ethertype in the SNAP header is the VLAN identifier and tag immediately follows the SNAP header.

Figure 34. Ethernet Frame with VLAN Tag and SNAP

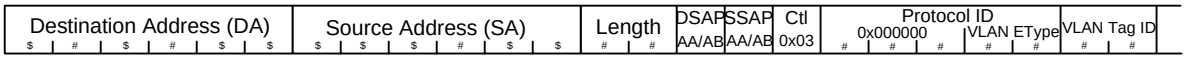

The following illustration shows the longest form of the Ethernet frame header that needs to be supported: two VLAN tags, an LLC header, and a SNAP header.

Figure 35. Ethernet Frame with VLAN Tags and SNAP

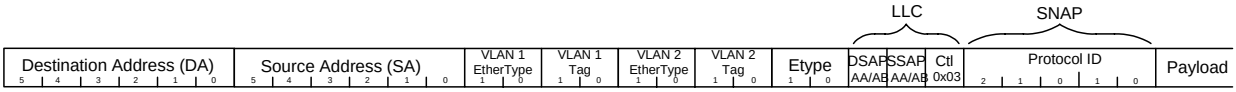

PBB (PROVIDER BACKED BRIDGES) SUPPORT

The provider backed bridges protocol is supported using two Ethernet comparator blocks back-to-back. The first portion of the frame has a type II Ethernet frame with either 0 or 1 VLAN tags followed by an I-tag. The following illustrations show two examples of the PBB Ethernet frame format.

Figure 36. PBB Ethernet Frame Format (No B-Tag)

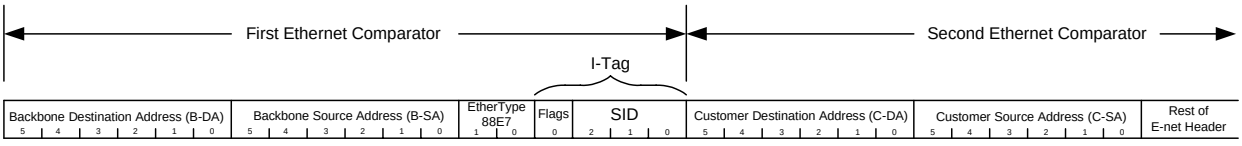

Figure 37. PBB Ethernet Frame Format (1 B-Tag)

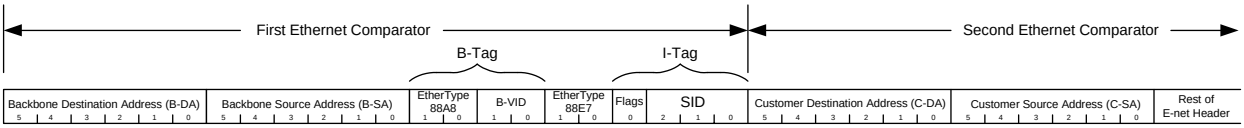

ETHERNET COMPARISON

The Ethernet comparator block has two forms of comparison, as follows:

- Next protocol comparison is common for all flows in the comparator. It is the single set of registers and is used to verify what the next protocol in the encapsulated stack will be.
- Flow comparison is used to match any of the possible flows within the comparator.

ETHERNET NEXT PROTOCOL COMPARISON

The next protocol comparison field looks at the last EtherType field in the header (there can be multiple in the header) to verify the next protocol. It may also look at VLAN tags and the EtherType field when it is used as a length. Each has a pattern match/mask or range, and an offset.

The following table lists the next protocol parameters for the Ethernet comparator.

Table 10. Ethernet Comparator (Next Protocol)

Parameter	Width	Description
Eth_Nxt_Comparator	3 bit	Pointer to the next comparator.
Eth_Frame_Sig_Offset	5 bit	Points to the start of the field used to build the frame signature.
Eth_VLAN-TPID_CFG	16 bit	Globally defines the value of the TPID for an S-tag, B-tag, or any other tag type other than a C-tag or I-tag.
Eth_PBB_ENA	1 bit	Configures if the packet carries PBB or not. This configuration bit is only present in the first Ethernet comparator block. PBB is disabled in Ethernet comparator block 2.
Eth_Etype_Match_Enable	1 bit	Configures if the Ethertype field match register is used or not. Only valid when the packet is a type II Ethernet packet.
Eth_Etype_Match	16 bit	If the packet is a type II Ethernet packet and Eth_Etype_Match_Enable is a 1, the Ethertype field in the packet is compared against this value.

ETHERNET FLOW COMPARISON

The Ethernet flow is determined by looking at VLAN tags and either the source address (SA) or the destination address (DA). There are a configurable number of these matched sets. The following table lists the flow parameters for the Ethernet comparator.

Table 11. Ethernet Comparator (Flow)

Parameter	Width	Description
Eth_Flow_Enable	1 bit/ flow	0 = Flow disabled 1 = Flow enabled
Eth_Channel_Mask	1 bit/ channel/ flow	0 = Do not use this flow match group for this channel 1 = Use this flow match group for this channel
Eth_VLAN_Tags	2 bit	Configures the number of VLAN tags in the frame (0, 1, or 2)

Table 11. Ethernet Comparator (Flow)

Parameter	Width	Description
Eth_VLAN_Tag1_Type	1 bit	Configures the VLAN tag type for VLAN tag 1 If PBB is not enabled: 0 = C-tag, value of 0x8100 1 = S-tag, match to the value in CONF_VLAN_TPID (global for all ports/directions) If PBB enabled: 0 = S-tag (or B-tag), to the value in CONF_VLAN_TPID (global for all ports/directions) There must be 2 VLAN tags, 1 S-tag and one I-tag 1 = I-tag
Eth_VLAN_Tag2_Type	1 bit	Configures the VLAN tag type for VLAN tag 2 If PBB is not enabled: 0 = C-tag, value of 0x8100 1 = S-tag, match to the value in CONF_VLAN_TPID (global for all ports/directions) If PBB enabled: The second tag is always an I-tag and this register control bit is not used. The second tag in PBB is always an I-tag.
Eth_Ethertype_Mode	1 bit	0 = Only type 2 Ethernet frames supported, no SNAP/LLC expected 1 = Type 1 & 2 Ethernet packets supported. Logic looks at the Ethertype/length field to determine the packet type. If the field is a length (less than 0x0600), then the packet is a type 1 packet and MUST include a SNAP & 3-byte LLC header. If the field is not a length, it is assumed to be an Ethertype and SNAP/LLC must not be present
Eth_VLAN_Verify_Ena	1 bit	0 = Parse for presence of VLAN tags but do not check the values. For PBB mode, the I-tag is still always checked. 1 = Verify the VLAN tag configuration including number and value of the tags.
Eth_VLAN_Tag_Mode	2 bit	0 = No range checking on either VLAN tag 1 = Range checking on VLAN tag 1 2 = Range checking on VLAN tag 2
Eth_Addr_Match	48 bit	Matches an address field selected by Eth_Addr_Match_Mode
Eth_Addr_Match_Select	2 bit	Selects the address to match 0 = Match the destination address 1 = Match the source address 2 = Match either the source or destination address 3 = Reserved, do not use
Eth_Addr_Match_Mode	3 bits per flow	Selects the address match mode. One or multiple bits can be set in this mode register allowing any combination of match types. For unicast or multicast modes, only the MSB of the address field is checked (0 = unicast; 1 = multicast). See section 3.2.3.1 of 802.3 for more details. 0 = Match the full 48-bit address 1 = Match any unicast address 2 = Match any multicast address
Eth_VLAN_Tag1_Match	12 bit	Match field for the first VLAN tag (if configured to be present).

Table 11. Ethernet Comparator (Flow)

Parameter	Width	Description
Eth_VLAN_Tag1_Mask	12 bit	Mask for the first VLAN tag. If a match set is not used, set this register to all 0s.
Eth_VLAN_Tag2_Match	12 bit	Match field for the update VLAN tag (if configured to be present).
Eth_VLAN_Tag2_Mask	12 bit	Mask for the second VLAN tag. If a match set is not used, set this register to all 0s.
Eth_VLAN_Tag_Range_Upper	12 bit	Upper limit of the range for one of the VLAN fields selected by ETH_VLAN_TAG_MODE register. If PBB mode is enabled, this register is not used for range checking but rather is the upper 12 bit of the I-tag.
Eth_VLAN_Tag_Range_Lower	12 bit	Lower limit of the range for one of the VLAN fields selected by ETH_VLAN_TAG_MODE register. If PBB mode is enabled, this register is not used for range checking but rather is the lower 12 bit of the I-tag SID.
Eth_Nxt_Prot_Grp_Sel	1 bit	Per flow, maps a particular flow to a next-protocol group register set. This register only appears in the Ethernet block in the OAM-optimized engine.

If the Ethernet block is part of the OAM optimized engine, there are two sets of next-protocol configuration registers. Both sets are identical except one has an _A suffix and the other has a _B suffix. In the per-flow registers an additional register, ETH_NXT_PROT_SEL, is included to map a particular flow with a set of next protocol register set. This function allows the Ethernet block within the OAM-optimized engine to act like two separate engines with a configurable number of flows assignable to each with a total maximum number of eight flows. It effectively allows two separate protocol encapsulation stacks to be handled within the engine.

3.10.15.2 MPLS Comparator

The MPLS comparator block counts MPLS labels to find the start of the next protocol. The MPLS header can have anywhere from 1 to 4 labels. Each label is 32 bit long and has the format shown in the following illustration.

Figure 38. MPLS Label Format

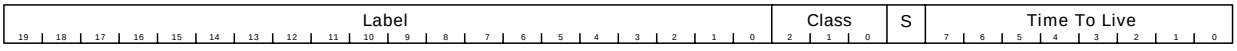

The S bit is used to indicate the last label in the stack, as follows: If S = 0, then there is another label. If S = 1, then this is the last label in the stack.

Also, the MPLS stack can optionally be followed by a control word (CW). This is configurable per flow.

The following illustration shows a simple Ethernet packet with either one label or three labels and no control word.

Figure 39. MPLS Label Stack within an Ethernet Frame

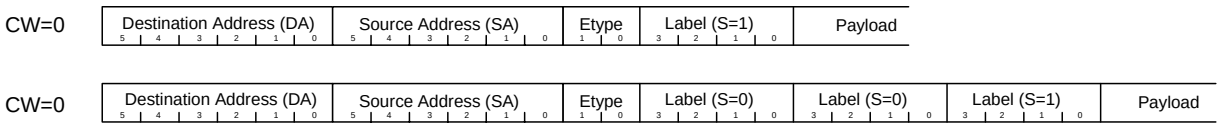

The following illustration shows an Ethernet frame with four labels and a control word. Keep in mind that this comparator is used to compare the MPLS labels and control words; the Ethernet portion is checked in the first stage.

Figure 40. MPLS Labels and Control Word

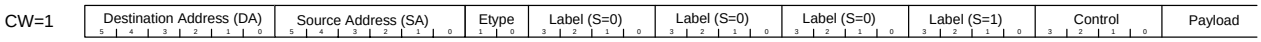

There could be VLAN tags between the SA and the Etype fields and, potentially, an LLC and SNAP header before the MPLS stack, but these would be handled in the Ethernet/ LLC/SNAP comparator.

The only configuration registers that apply to all flows within the comparator are the match_mode register and the nxt_comparator register. The match mode register determines how the match filters are used and there is one per stage. Each flow has it own complete set of match registers.

Table 12. MPLS Comparator: Next Word

Parameter	Width	Description
MPLS_Nxt_Comparator	3 bit	Pointer to the next comparator

Table 13. MPLS Comparator: Per-Flow

Parameter	Width	Description
MPLS_Flow_Enable	1 bit per flow	0 = Flow disabled 1 = Flow enabled
MPLS_Channel_Mask	1 bit per channel per flow	0 = Do not use this flow match group for this channel 1 = Use this flow match group for this channel
MPLS_Ctl_Word	1 bit	Indicates if there is a 32-bit control word after the last label. This should only be set if the control word is not expected to be an ACH header. ACH headers are checked in the IP block. If the control word is a non-ACH control word, only the upper 4 bits of the control are checked and are expected to be 0. 0 = There is no control word after the last label 1 = There is expected to be a control word after the last label
MPLS_REF_PNT	1 bit	The MPLS comparator implements a searching algorithm to properly parse the MPLS header. The search can be performed from either the top of the stack or the end of the stack. 0 = All searching is performed starting from the top of the stack 1 = All searching if performed from the end of the stack

Table 13. MPLS Comparator: Per-Flow

Parameter	Width	Description	
MPLS_STACK_DEPTH	4 bit	Each bit represents a possible stack depth, as shown in the following list.	
		MPLS_STACK_DEPTH Bit	Allowed Stack Depth
		0	1
		1	2
		2	3
		3	4

Table 14. MPLS Range_Upper/Lower Label Map

Parameter	MPLS_REF_PNT = 0, top-of-stack referenced	MPLS_REF_PNT=1, end-of-stack referenced
MPLS_Range_Upper/Lower_0	Top label	Third label before the end label
MPLS_Range_Upper/Lower_1	First label after the top label	Second label before the end label
MPLS_Range_Upper/Lower_2	Second label after the top label	First label before the end label
MPLS_Range_Upper/Lower_3	Third label after the top label	End label

The offset to the next protocol is calculated automatically. It is based upon the number of labels found and whether a control word is configured to be present. It points to the first octet after the last label or after the control word, if present.

Table 15. Next MPLS Comparator

Parameter	Width	Description
MPLS_Range_Lower	20 bit × 4 labels	Lower value of the label range when range checking is enabled
MPLS_Range_Upper	20 bit × 4 labels	Upper value of the label range when range checking is enabled

If an exact label match is desired, set the upper and lower range values to the same value. If a label value is a don't care, then set the upper value to the maximum value and the lower value to 0.

The MPLS comparator block used in the OAM-optimized engine differs from the one used in the PTP-optimized engine.

Just like the Ethernet comparator block, there are two sets of next protocol blocks along with a next protocol association configuration field per-flow. This allows two different encapsulations to occur in a single engine.

Table 16. Next-Protocol Registers in OAM-Version of MPLS Block

Parameter	Width	Description
MPLS_Nxt_Prot_Grp_Sel	1 bit per flow	Maps each flow to next-protocol-register set A or B

3.10.15.3 IP/UDP/ACH Comparator

The IP/UDP/ACH comparator is used to verify one of three possible formats, IPv4, IPV6, and ACH. Additionally, IPv4 and IPv6 can also have a UDP header after the IP header.

There are two of these comparators and they can operate at stages 2, 3, or 4 of the analyzer pipeline. Note that if there is an IP-in-IP encapsulation, a UDP header will only exist with the inner encapsulation.

3.10.15.4 IPv4 Header Format

The following illustration shows an IPv4 frame header followed immediately by a UDP header. IPv4 does not always have the UDP header, but the comparator is designed to work with or without it. The Header Length field is used to verify the offset to the next protocol. It is a count of 32-bit words and does not include the UDP header. If the IPv4 frame contains a UDP header, the Source and Destination ports are also checked. These values are the same for all flows within the comparator. Note that IPv4 options, extended headers, and UDP fragments are not supported.

Figure 41. IPv4 with UDP

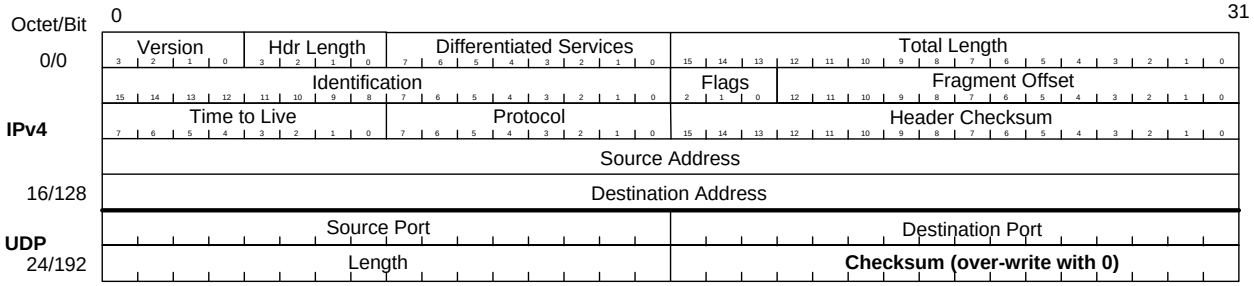

Per flow validation is performed on the Source or Destination Address in the IPv4 header. The comparator can be configured to indicate a match in the flow if the source, destination, or either the source or destination fields match.

3.10.15.5 IPv6 Header Format

The following illustration shows an IPv6 frame header followed immediately by a UDP header. IPv6 does not always have the UDP header, but the comparator is designed to work with or without it. The Next Header field is used to verify the offset to the next protocol. It is a count of 32-bit words and does not include the UDP header. If the IPv6 frame contains a UDP header, the Source and Destination ports are also checked. These values are the same for all flows within the comparator.

Figure 42. IPv6 with UDP

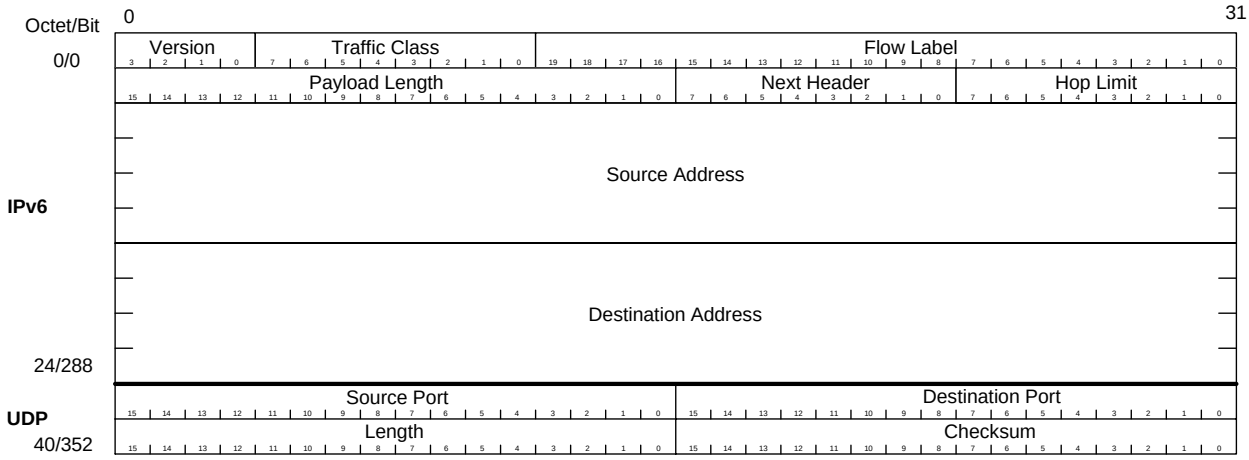

Per flow validation is performed on the Source or Destination Address in the IPv6 header. The comparator can be configured to indicate a match in the flow if the source, destination, or either the source or destination fields match.

If the IPv6 frame is the inner most IP protocol, then the checksum field must be valid. This is accomplished using a pair of pad bytes after the PTP frame. The checksum is computed using one's compliment of the one's compliment sum of the IPv6 header, UDP header, and payload including the pad bytes. If any of the fields in the frame are updated, the pad byte field must be updated so that the checksum field does not have to be modified.

Note IPv6 extension headers are not supported.

3.10.15.6 ACH Header Format

The following illustrations show ACH headers. They can appear after a MPLS label stack in place of the control word. ACH is verified as a protocol only. There are no flows within the protocol for ACH. The ACH header can optionally have a Protocol ID field. The protocol is verified using the Version, Channel type, and optional Protocol ID field.

Figure 43. ACH Header Format

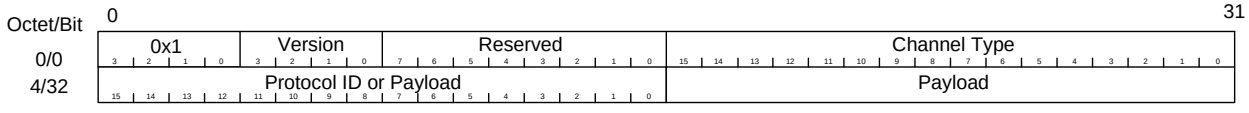

Figure 44. ACH Header with Protocol ID Field

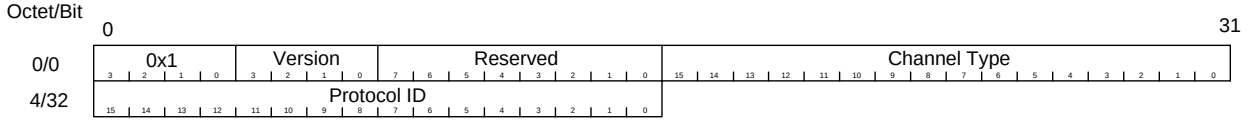

3.10.15.7 IPSec

IPSec adds security to the IP frame using an Integrity Check Value (ICV), a variable-length checksum that is encoded with a special key. The key value is known by the sender and the receiver, but not any of the devices in between. A frame must have a correct ICV to be valid. The sequence number field is a continuously incrementing value that is used to prevent replay attacks (resending a known good frame).

Little can be done with frames when IPSec is used because the 1588 block cannot recalculate the ICV and the frame cannot be modified on egress. Therefore, one-step processing cannot be performed, only two-step processing can be done. The only task here is to verify the presence of the protocol header. Stored timestamps in the TS FIFO are used to create follow-up messages. On ingress, the timestamp can be added to the PTP frame by writing it into the reserved bytes or by overwriting the CRC with it and appending a new CRC. The CPU must know how to handle these cases correctly.

The following illustration shows the format of the IPSec frame. It normally appears between the IP header (IPv4 or IPv6) and the UDP header or at the start of the payload.

Figure 45. IPSec Header Format

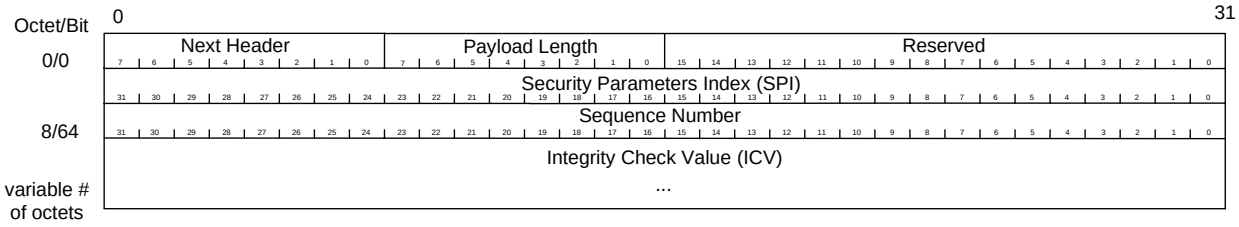

There is only one set of match/mask registers associated with IPSec and they are used to verify the presence of the IPSec header. The following illustration shows the largest possible IP frame header with IPv6, IPSec, and UDP.

Figure 46. IPv6 with UDP and IPSec

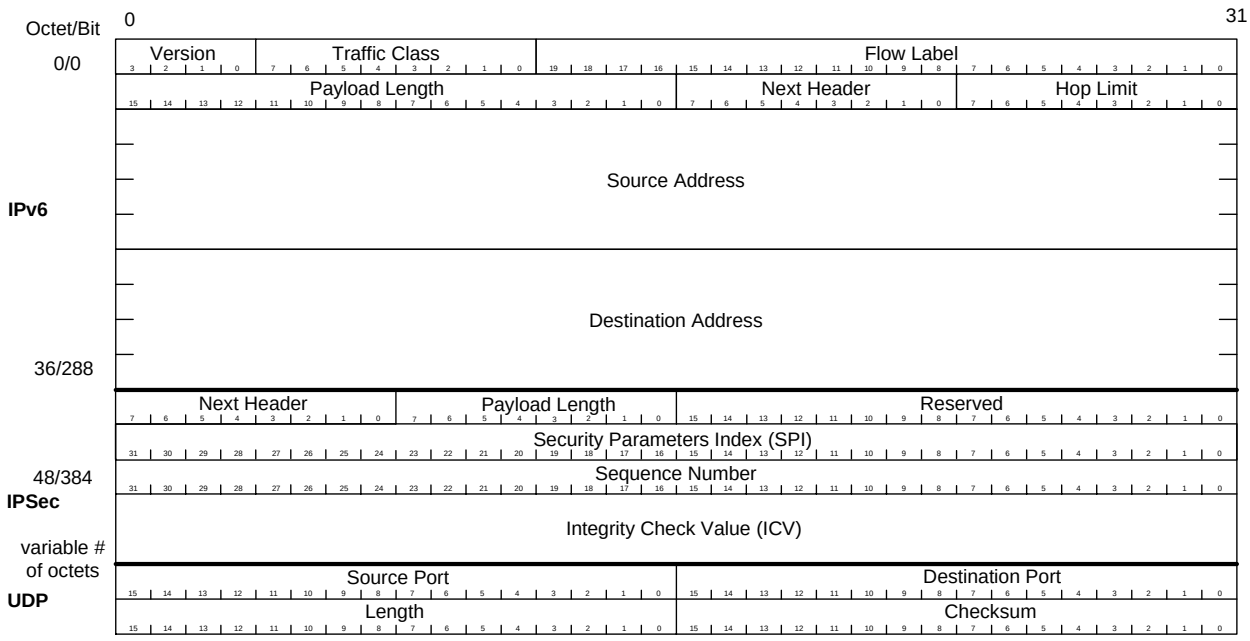

3.10.15.8 Comparator Field Summary

The following table shows a summary of the fields and widths to verify IPv4, IPv6, and ACH protocols.

Table 17. Comparator Field Summary

Protocol	Next Protocol Fields	NPF Bit Widths	Flow Fields	Flow Bit Widths
IPv4	Header length	One 4-bit field	Source/ Destination Address	One 32-bit field
	UDP Source/Destination Port	One 32-bit field		
IPv6	Next header	One 8-bit field	Source/ Destination Address	One 128-bit field
	UDP Source/Destination Port	One 32-bit field		
ACH	Entire ACH header	One 64-bit field		
IPSec	Next Header/Payload Length/ SPI	One 64-bit field		

IP/ACH COMPARATOR NEXT PROTOCOL

The following table shows the registers used to verify the current header protocol and the next protocol. They are universal and cover IPv4, IPv6, and ACH. They can also be used to verify other future protocols.

Table 18. IP/ACH Next-Protocol Comparison

Parameter	Width	Description
IP_Mode	2 bit	Specifies the mode of the comparator. If IPv4 or IPv6 is selected, the version field is automatically checked to be either 4 or 6 respectively. If another protocol mode is selected, then the version field is not automatically checked. In IPv4, the fragment offset field must be 0, and the MF flag bit (LSB of the flag field) must be 0. 0 = IPv4 1 = IPv6 2 = Other protocol, 32-bit address match 3 = Other protocol, 128-bit address match
IP_Prot_Match_1	8 bit	Match bit for Protocol field in IPv4 or next header field in IPv6
IP_Prot_Mask_1	8 bit	Mask bits for IP_Prot_Match_1. For each bit, if it is a 1, the corresponding match bit is valid. If it is 0, the corresponding match bit is ignored. Disable this match/mask set by setting the mask register to all 0's.
IP_Prot_Offset_1	5 bit	Indicates the starting position relative to the beginning of the IP frame header to start matching for the match/mask 1 register pair.
IP_Prot_Match_2	64 bit	Match bits for the IPSec header or any other desired field. For ACH, this register should be used to match the ACH header.
IP_Prot_Mask_2	64 bit	Mask bits for IP_Prot_Match_2. For each bit, if it is a 1, the corresponding match bit is valid. If it is 0, the corresponding match bit is ignored. Disable this match/mask set by setting the mask register to all 0's.
IP_Prot_Offset_2	7 bit	Indicates the starting position relative to the beginning of the IP frame header to start matching for the match/mask two-register pair.
IP_Nxt_Protocol	8 bit	Points to the start of the next protocol relative to the beginning of this header. It is the responsibility of the programmer to determine this offset, it is not calculated automatically. Each flow within an encapsulation engine must have the same encapsulation order and each header must be the same length. This field is current protocol header length in bytes.
IP_Nxt_Comparator	3 bit	Pointer to the next comparator. 0 = Reserved 1 = Ethernet comparator 2 2 = IP/UDP/ACH comparator 1 3 = IP/UDP/ACH comparator 2 4 = Reserved 5 = PTP/OAM comparator 6,7 = Reserved
IP_Flow_Offset	5 bit	Indicates the starting position relative to the beginning of the IP frame header to start matching for the flow match/mask register pair. When used with IPv4 or 6, this will point to the first byte of the source address. When used with a protocol other that IPv4 or 6, this register points to the beginning of the field that will be used for flow matching.
IP_UDP_Checksum_Clear_Ena	1 bit	If set, the 2-byte UDP checksum should be cleared (written with zeroes). This would only be used for UDP in IPv4.

Table 18. IP/ACH Next-Protocol Comparison

Parameter	Width	Description
IP_UDP_Checksum_Update_Ena	1 bit	If set, the last two bytes in the UDP frame must be updated to reflect changes in the PTP or OAM frame. This is necessary to preserve the validity of the IPv6 UDP checksum. Note that IP_UDP_Checksum_Clear_Ena & IP_UDP_Checksum_Update_Ena should never be set at the same time.
IP_UDP_Checksum_Offset	8 bit	This configuration field is only used if the protocol is IPv4. This register points to the location of the UDP checksum relative to the start of this header. This info is used later by the PTP/Y.1731 block to inform the rewriter of the location of the checksum in a UDP frame. This is normally right after the Log Message Interval field.
IP_UDP_Checksum_Width	2 bit	Specifies the length of the UDP checksum in bytes (normally 2 bytes)

The IP/ACH Comparator Flow Verification registers are used to verify the current frame against a particular flow within the engine. When this engine is used to verify IPv4 or IPv6 protocol, the flow is verified using either the source or destination address in the frame.

If the protocol is something other than IPv4 or IPv6, then the flow match can be used to match either a 32 or 128 bit field pointed to by the IP_Flow_Offset register. Mask bits can be used to shorten the length of the match, but there is no concept of source or destination address in this mode.

Table 19. IP/ACH Comparator Flow Verification Registers

Parameter	Width	Description
IP_Flow_Ena	1 bit per flow	0 = Flow disabled 1 = Flow enabled
IP_Flow_Match_Mode	2 bit per flow	This register is only valid when the comparator block is configured to match on IPv4 or IPv6. It allows the match to be performed on the source address, destination address, or either address. 0 = Match on the source address 1 = Match on the destination address 2 = Match on either the source or the destination address
IP_Flow_Match	128 bit	Match bits for source & destination address in IPv4 & 6. Also used as the flow match for protocols other than IPv4 or 6. When used with IPv4, only the upper 32 bits are used and the remaining bits are not used.
IP_Flow_Mask	128 bit	Mask bits for IP_Flow_Match. For each bit, if it is a 1, the corresponding match bit is valid. If it is 0, the corresponding match bit is ignored.
IP_Channel_Mask	1 bit per channel per flow	Enable for this match set for this channel
IP_Frame_Sig_Offset	5 bit	Points to the start of the field that will be used to build the frame signature. This register is only present in comparators where frame signature is supported. In other words, if there is no frame signature FIFO in a particular direction, this register will be removed.

3.10.15.9 PT/Y.1731 OAM Comparator

The PTP/OAM comparator is always the last stage in the analyzer for each encapsulation engine. It can validate IEEE 1588 PTP frames or Y.1731 OAM frames.

3.10.15.10 PTP Frame Header

The following illustration shows the header of a PTP frame.

Figure 47. PTP Frame Layout

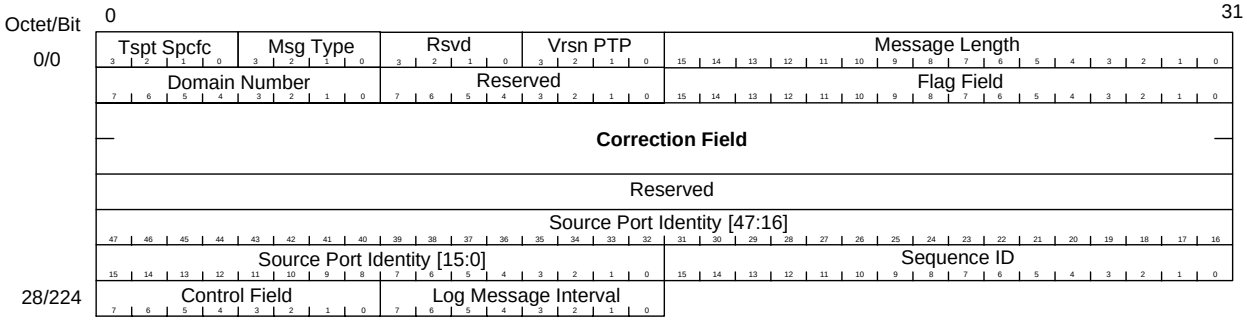

Unlike most of the other stages, there is no protocol validation for PTP frames; only interpretation of the header to determine what action to take. The first eight bytes of the header are used to determine the action to be taken. These match fields in the flow comparison registers with a corresponding set of command registers for each flow.

3.10.15.11 Y.1731 OAM Frame Header

1DM, DMM, and DMR are the three supported Y.1731 frame headers. The following illustration shows the header part of a 1DM Y.1731 OAM frame.

Figure 48. OAM 1DM Frame Header Format

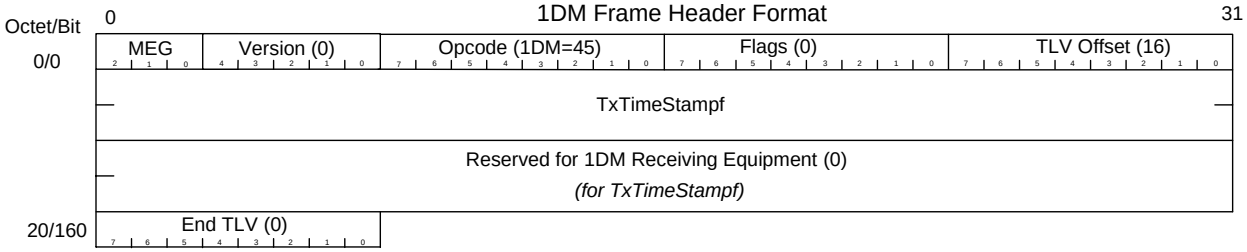

The following illustration shows a DMM frame header.

Figure 49. OAM DMM Frame Header Format

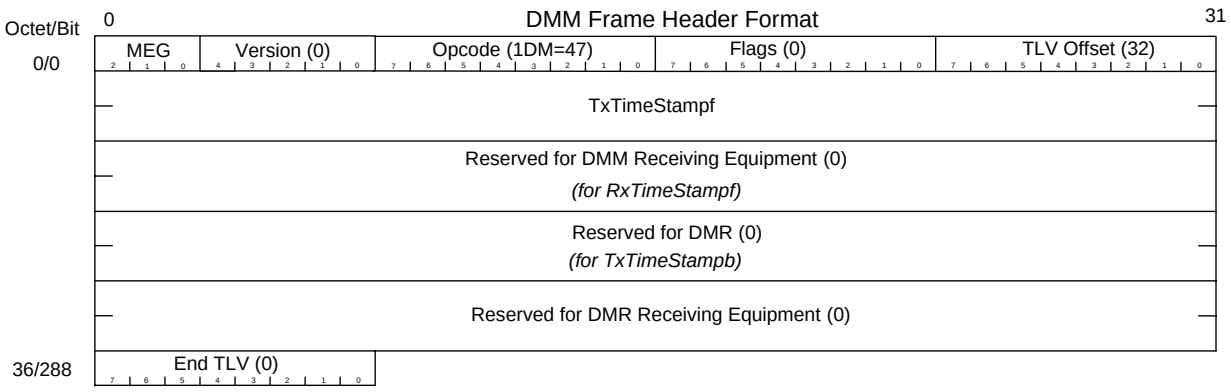

The following illustration shows a DMR frame header.

Figure 50. OAM DMR Frame Header Format

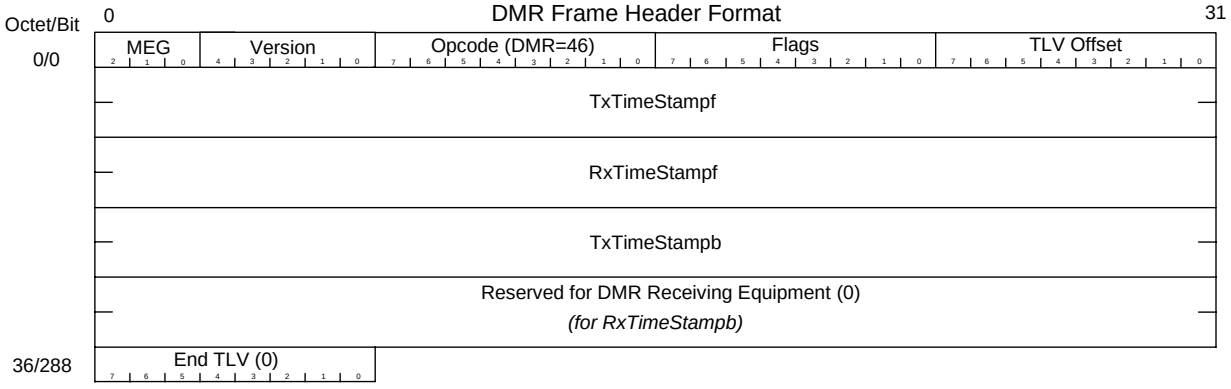

As with PTP, there is no protocol validation for Y.1731 frames; only interpretation of the header to determine what action to take. The first four bytes of the header are used to determine the action to be taken.

3.10.15.12 PTP Comparator Action Control Registers

The following registers perform matching on the frame header and define what action is to be taken based upon the match. There is one mask register for all flows, and the rest of the registers are unique for each flow. The number of flows supported in this comparator is separately configurable from the other comparator blocks.

Table 20. PTP Comparison

Parameter	Width	Description
PTP_Flow_Match	64 bit	Matches bits in the PTP/Y.1731 frame starting at the beginning of the protocol header
PTP_Flow_Mask	64 bit	Mask bits for PTP_Flow_Match
PTP_Domain_Range_Lower	8 bit	Lower range of the domain field to match

Table 20. PTP Comparison

Parameter	Width	Description		
PTP_Domain_Range_Upper	8 bit	Upper range of the domain field to match		
PTP_Domain_Range_ Enable	1 bit	Enable for range checking		
PTP_Domain_Offset	5 bit	Pointer to the domain field, or whatever field is to be used for range checking		
PTP_Action_Command	3 bit	**Command Value**	**Mnemonic**	**Action**
		0	NOP	Do nothing
		1	SUB	New correction field = Current correction field – Captured local time (not supported)
		2	SUB_P2P	New correction field = Current correction field – Local latency + path_delay (not supported)
		3	ADD	New correction field = Current correction field + Captured local time (not supported)
		4	SUB_ADD	New correction field = Current correction field + (Captured local time + Local latency – Time storage field)
		5	WRITE_1588	Write captured local time to time storage field
		6	WRITE_P2P	Active_timestamp_ns = captured local time and path_delay written to time storage field and correction field (deprecated command)
		7	WRITE_NS	Write local time in nanoseconds to the new field
		8	WRITE_NS_ P2P	Write local time in nanoseconds + p2p_delay to the new field and correction field
PTP_Save_Local_Time	1 bit	When set, saves the local time to the Timestamp FIFO		
PTP_Correction_Field_Offset	5 bit	Points to the location of the correction field for updating the timestamp. Location is relative to the first byte of the PTP/OAM header.		
PTP_Time_Storage_Field_ Offset	6 bit	Points to a location in a PTP frame where a time value can be stored or read.		
PTP_Add_Delay_Asymmetry_ Enable	1 bit	When enabled, the add_delay_asymmetry signal to the timestamp block is set		
PTP_Subtract_Delay_ Asymmetry_Enable	1 bit	When enabled, the subtract_delay_asymmetry signal to the timestamp block is set		

Table 20. PTP Comparison

Parameter	Width	Description
PTP_Zero_Field_Offset	6 bit	Points to a location in the PTP/OAM frame to be zeroed if this function is enabled
PTP_Zero_Field_Byte_Count	4 bit	The number of bytes to be zeroed. If this field is 0, then this function is not enabled.
PTP_Modified_Frame_Byte_ Offset	3 bit	Indicates the position relative to the start of the PTP frame in bytes where the Modified_Frame_Status bit resides. This value is also used to calculate the offset from the beginning of the Ethernet packet to this field for use by the Rewriter.
PTP_Modified_Frame_Status _Update	1 bit	If set, tells the rewriter to update the value of this bit. Configuration registers inside the rewriter indicate if the bit will be set to 0 or 1.
PTP_Rewrite_Bytes	4 bits	Number of bytes in the PTP or OAM frame that must be modified by the Rewriter for the timestamp
PTP_Rewrite_Offset	8 bits	Points to where in the frame relative to the SFD that the timestamp should be updated
PTP_New_CF_Loc	8 bits	Location of the new correction field relative to the PTP header start
PTP_Channel_Mask	1 bit per channel per flow	Enable for this match set for this channel
PTP_Flow_Enable	1 bit	When set, the fields associated with this flow are all valid

The following table shows controls that are common to all flows.

Table 21. PTP Comparison: Common Controls

Parameter	Width	Description
PTP_IP_CHKSUM_Sel	1 bit	0 = Use IP checksum controls from comparator 1 1 = Use IP checksum controls from comparator 2
FSB_Adr_Sel	2 bits	Selects the source of the address for use in the frame signature builder

The following table shows the one addition, per-flow, register.

Table 22. PTP Comparison: Additions for OAM-Optimized Engine

Parameter	Width	Description
PTP_NXT_Prot_Group_Mask	2 bits	There are two bits for each flow. Each bit indicates if the flow can be associated with next-protocol group A or B. One or both bits may be set. If a bit is 1 for a particular next-protocol group, then a flow match is valid if the prior comparator stages also produced matches with the same next-protocol group.

3.10.15.13 Future Protocol Compatibility

Except for MPSL, the comparators are not hardwired to their intended protocols. They can be used as generic field and range comparators because all of the offsets or

pointers to the beginning of the fields are configurable. The IP comparator is the most generic and would probably be the first choice for validating a new protocol.

Additionally, if there are not enough comparison resources in a single comparator block to handle a new protocol, two comparators back-to-back can used by splitting up the comparison work. One portion can be validated in one comparator and then handed off to another. The only restriction is that there must be at least one 64-bit word of separation between the start of the protocol and where the second starts to operate.

3.10.15.14 Reconfiguration

There are two ways to perform reconfiguration:

1. Disable an entire encapsulation engine.
   Once an engine has been disabled, any of the configuration registers associated with it may be modified in any order. If other encapsulation engines are still active, they will still operate normally.

2. Disable a flow in an active engine.
   Each stage in the engine has an enable bit for each flow. If a flow is disabled in a stage, its registers may be modified. Once reconfiguration for a flow in a stage is complete, it can be enabled.

The disabling of engines and flows is always done in a clean manner so that partial matches do not occur. Flows and engines are always enabled or disabled during inter-packet gaps or at the end of a packet. This guarantees that when a new packet is received that it will be analyzed cleanly.

If strict flow matching is enabled and a flow is disabled in one of the stages, then the entire flow is automatically disabled.

If any register in a stage that applies to all flows needs to be modified, then the entire encapsulation engine must be disabled.

3.10.15.15 Frame Signature Builder

Along with timestamp and CRC updates, the analyzer outputs a frame signature that can be stored in the timestamp FIFO to help match frames with other info in the FIFO. This information is used by the CPU so that it can match timestamps in the timestamp FIFO with actual frames. The frame signature is up to 16 bytes long and contains information from the Ethernet header (SA or DA), IP header (SA or DA), and from the PTP or OAM frame. The frame signature is only used in the egress direction.

The PTP block contains a set of mapping registers to configure which bytes are mapped into the frame signature. The following tables show the mapping for each byte.

Table 23. Frame Signature Byte Mapping

Select	Source Byte
0-23	PTP header byte number = (31-select)
24	PTP header byte number 6
25	PTP header byte number 4
26	PTP header byte number 0
27	Reserved

Table 23. Frame Signature Byte Mapping

Select	Source Byte
28-35	Selected address byte (select-28)

Table 24. Frame Signature Address Source

Parameter	Width	Description	
FSB_Map_Reg_0-15	6 bits	For each byte of the frame signature, use Table 23, page 81 to select which available byte is used. Frame signature byte 0 is the LSB. If not all 16 bytes are needed, the frame signature should be packed towards the LSB and the upper unused byte configuration values do not need to be programmed.	
FSB_Adr_Sel	2 bits	Selects the source of the address for use in the frame signature builder according to the following list	
		Select Value	**Address Source**
		0	Ethernet block 1
		1	Ethernet block 2
		2	IP block 1
		3	IP block 2

Configuration registers in each comparator block supply an address to select if it is the source address or the destination address.

3.10.15.16 Configuration Sharing

The analyzer services both channels. Each flow within each comparator has a channel-mask register that indicates which channels the flow is valid for. Each flow can be valid for channel A, channel B, or both channels.

The total of eight flows can be allocated the two channels if the analyzer configuration cannot be shared. They can each have four distinct flows (or three for the one, and five for the other, etc.).

3.10.15.17 OAM-Optimized Engine

In addition to the descriptions of the Ethernet and MPLS blocks in the OAM optimized engine, there is the notion of protocol-A/protocol-B. When a match occurs in the Ethernet 1 block the status of the protocol set that produced the match is indicated. There are two bits, one for protocol A and another for protocol B. If both sets produce a match (this is actually very likely), then both bits are set.

These bits are then carried to the next comparison block and only allow flow matches for the protocol sets that produced matches in the prior block. This block also produces a set of protocol match bits that are also carried forward.

This feature is provided to prevent a match with protocol set A in the first block and protocol set B in the second block.

3.10.16 Timestamp Processor

The primary function of the timestamp processor block is to generate a new Timestamp_field or new Correction_field (Transparent clocks) for the rewriter block.

The timestamp block generates an output that is either a snapshot of the corrected Local Time (struct Timestamp) or a signed (two's complement) 64 bit Correction_field.

In the ingress direction the timestamp block calculates a new timestamp for the rewriter that indicates the earlier time when the corresponding PTP event frame entered the chip (crossed the reference plane referred to in the IEEE1588 standard).

In the egress direction the timestamp block calculates a new timestamp for the rewriter in time for the PCS block to transmit the new timestamp field in the frame. In this case the Timestamp field indicates when the corresponding PTP event frame will exit the chip.

Transparent clocks correct PTP event messages for the time resided in the transparent clock. Peer-to-Peer transparent clocks additionally correct for the propagation time on the inbound link (Path_delay). The Path_delay [ns] input to the timestamp block is software programmed based upon 1588 path delay measurements.

In general, the IEEE 1588 standard allows for a transparent clock to update the Correction_Field for both PTP event messages as well as the associated follow up message. However, the TSP only updates PTP event messages. Also, the 1588 standard allows that end-to-end transparent clocks correct and forward all PTP-timing messages while Peer-to-Peer transparent clocks only correct and forward Sync and Follow_Up messages. Again, the TSP only updates PTP event messages (not Follow_Up messages).

Internally the timestamp block generates an Active_timestamp from the captured/ timestamped Local time (Raw_timestamp). The Active_time stamp is the Raw_timestamp corrected for the both fixed (programmed) local chip, and variable chip latencies relative to where the Start_of_Frame_Indicator captures the local time. The timestamp block operates on the Active_timestamp based on the Command code.

The Active_timestamp is calculated differently in the Ingress and Egress directions and the equations are given below.

In the ingress direction:

```
Active_timestamp = Raw_timestamp - Local_latency - Variable_latency
```

In the egress direction:

```
Active_timestamp = Raw_timestamp + Local_latency + Variable_latency
```

In addition, the following values are also calculated for use by the commands:

```
Active_timestamp_ns = Active_timestamp converted to nanoseconds

Active_timestamp_p2p_ns = active_timestamp_ns + path delay
```

The Local_latency is a programmed fixed value while the Variable_latency is predicted from the PCS logic based upon the current state of the ingress or egress data pipeline.

For the option of Peer-to-Peer transparent clocks, the ingress Active_timestamp calculation includes an additional Path_delay component. The path delay is always added for a transparent clock per the standard. The path delay is always added to the correction field.

The signed 32-bit two's complement Delay Asymmetry register (bits 31–0) can be programmed by the user. Bit 31 is the sign bit. Bits 15–0 are scaled nanoseconds just

like for the CorrectionField format. The DelayAsymmetry register (whether it be positive or negative) will be sign extended and added to the 64-bit correction field (signed add) if the Add_Delay_Asymmetry bit is set. The DelayAsymmetry register (whether it be positive or negative) will be sign extended and subtracted from the 64-bit correction field (signed Subtract) if the Subtract_Delay_Asymmetry bit is set.

The timestamp block keeps a shadow copy of the programmed CPU latency values (Local_latency, Path_delay, and Delay_Asymmetry) to protect against CPU updates.

3.10.17 Timestamp FIFO

The Timestamp FIFO stores timestamps along with frame signature information. This information can be read out by a CPU or pushed out on a dedicated Serial Timestamp Output Interface and used in 2-step processing mode to create follow-up messages.

The timestamp FIFO takes a frame signature from the analyzer and the updated correction field, and the full data set for that timestamp is saved to the FIFO. This creates an interrupt to the CPU. If the FIFO ever overflows this is indicated with an interrupt.

The stored frame signature can be of varying sizes controlled by [REGISTER]. Only the indicated number of signature bytes is saved with each timestamp. This value is register configurable. The saved values are packed so that reducing the number of signature bytes allows more timestamps to be saved. The minimum size of the FIFO is 200 bytes or 10 20-byte values. This is the minimum number of timestamps that must be available.

The packing of the timestamp data is done by logic before the write occurs to the FIFO. When no compression is used, each timestamp may contain 208 bits of information, 128 bits of frame signature and 80 bits of timestamp data. Therefore a full sized timestamp is 26 bytes long. Compressing the frame signature can reduce this to as little as 10 bytes (or 4 bytes if [REGISTER] = 1) if no signature information is saved ([REGISTER] = 0). The value to store is built up in a 51-byte register and when it contains 26 valid bytes, the value is written to the FIFO and the register shifted down by 26 bytes. Data in the FIFO is packed end-to-end and it is up to the reader of the data to unpack the data.

The timestamps in the FIFO are visible and accessible for the CPU as a set of 32-bit registers. Seven (7) register reads from seven separate locations are required to read a full timestamp if all bits are used. The MSB of the first read address contains the current FIFO empty flag, which can be used by the CPU to determine if the current timestamps are available for reading. If the bit is set, the FIFO is empty and no timestamps are available. The value that was read can be discarded because it does not contain any valid timestamp data. If the bit is 0 (deasserted), the value contains 16 valid data bits of a timestamp. The remaining bits should be read from the other registers in the other locations and properly packed to recreate the timestamp. Care should be taken to read the timestamps one at a time as each read of the last (7th) address will trigger a pop of the FIFO.

Timestamps are packed into seven registers in [REGISTER]. If the TS FIFO registers are read to the point that the FIFO goes empty and there are remaining valid bytes in the 51-byte packing register, then the packing register is written to the FIFO. In this case the registers may not be fully packed with timestamps, and flag bits are used to indicate where the valid data ends within the set of seven registers. The three flag bits are encoded in register 0 (together with the empty flag) and are encoded as follows:

000 = Only a partial timestamp is valid in the seven register set

001 = One timestamp begins in the current seven register set

010 = Two timestamps begin in the current seven register set.

011 = Three timestamps begin in the current seven register set (ts_4byte mode)

100 = Four timestamps begin in the current seven register set (ts_4byte mode)

101 = Five timestamps begin in the current seven register set (ts_4byte mode)

110 = Six timestamps begin in the current seven register set (ts_4byte mode)

111 = The current seven register set is fully packed with valid timestamp data

The FIFO empty bit is also visible in the status [REGISTER] so the CPU can poll this bit to know when timestamps are available although this is not necessary. It can also wait for a maskable interrupt which will assert whenever the timestamp FIFO level reaches the threshold given in [REGISTER]. The FIFO level is also visible in a status register. If the timestamp FIFO overflows writes to the FIFO are inhibited. The data in the FIFO is still available for reading but new timestamps are dropped.

3.10.17.1 FIFO Sizing

The required size of the timestamp FIFO is calculated using the following assumptions:

- The worst case is end-to-end transparent clocks with two-step processing
- 1588 processor is a node with two masters and Ns slaves with Ns > 200
- 128 Sync frames/sec are sent unicast to each slave from each master
- Each master has a limited PTP message rate and does not burst PTP frames
- The maximum number of back-to-back PTP messages is (masters × 2) as each master will send spaced out messages and only back-to-back masters will cause back-to-back PTP messages
- The timestamp FIFO read data is a full 26-bytes wide (seven 32-bit reads)

Therefore, the FIFO size is only required to be equal to the number of Masters, which is 2. To be conservative and allow extra storage, a FIFO size of 8 26-byte timestamps is used. This is 208 bytes of storage. This also meets the requirement of 20 minimum size timestamps.

The rate at which timestamps can be read out of the FIFO relative to the rate at which they arrive is an issue in two-step processing modes where the CPU needs to read the timestamps out of the timestamp FIFO to create follow-up messages in a timely manner. Assuming the following to be the worst case:

- End-to-end transparent clocks with two-step processing
- 1588 processor is a node with two masters and Ns slaves with Ns > 200
- 128 Sync frames/sec are sent unicast to each slave
- Each slave responds with 64 delay_req frames/sec

- The timestamp FIFO read data is a full 26-bytes wide (seven 32-bit reads)
- The time taken for each 32-bit read is Tr
- The max bandwidth utilization is U where 0 < U <1

Each slave has

((128 × 2)+64) = 320 TS/sec

Number of timestamps read out is

1/(7 × Tr) TS/sec

Therefore:

320 × Ns < U/(7 × Tr) or Ns < U/(320 × 7 × Tr)

This relation can be used to calculate the maximum number of slaves supported with two masters, given the read access time and the maximum bandwidth utilization.

3.10.18 Serial Timestamp Output Interface

The timestamp FIFO serial interface block writes, or pushes, timestamp/frame signature pairs that have been enqueued and packed into timestamp FIFOs to the external chip interface consisting of three output pins: 1588_SPI_DO, 1588_SPI_CLK, and 1588_SPI_CS. There is one interface for all channels.

When the SI interface is enabled, the timestamp/frame signature pairs are dequeued from timestamp FIFO(s) and unpacked. Unpacked timestamp/frame signature pairs are then serialized and sent one at a time to the external interface. Unpacking shifts the timestamp/frame signature into alignment considering the configured size of the timestamps and frame signatures (a single SI write may require multiple reads from a timestamp FIFO). The timestamp FIFO serial interface is an alternative to the register interface described in the TS FIFO section. When the SI interface is enabled (via [REGISTER]), the TS_FIFO_SI prevents register access (register reads return empty with invalid data and flags). Also, when the SI interface is enabled the output enable drivers for pins are turned on (by default the interface is disabled and tri-stated).

Timestamp/Frame signature pairs from two egress timestamp FIFOs are serialized one at a time and transmitted to the interface pins. The TS_FIFO_SI arbitrates in a round-robin fashion amongst the ports that have non-empty timestamp FIFOs. The port associated with each transmitted timestamp/frame signature pair is indicated in a serial address that precedes the data phase of the serial transmission. Because the timestamp FIFOs are instantiated in the per port clock domains, a small single entry asynchronous SI FIFO (per port) ensures that the timestamp/frame signature pairs are synchronized, staged, and ready for serial transmission. When an SI FIFO is empty, the SI FIFO control fetches and/or unpacks a single timestamp/frame signature performing any timestamp FIFO dequeues necessary. The SI FIFO goes empty following the completion of the last data bit of the serial transmission. Enabled ports ([REGISTER]) participate in the round-robin selection.

Counters accumulate the number of timestamp/frame signature pairs transmitted from the SI interface for each channel. Counters per channel also accumulate the number of timestamp/frame signature pairs that have been dropped due to a TS FIFO overflow.

As shown in the following illustration, the SPI compatible interface asserts a chip select (SI_EN) for each write followed by a write command data bit equal to 1, followed by a don't care bit (0), followed by an address phase, followed by a data phase, followed by a deselect where SI_EN is negated. Each write command corresponds to a single timestamp/frame signature pair. The length of the data phase depends upon the sum of the configured lengths of the timestamp and signature, respectively. The address phase is fixed at five bits. The 1588_SPI_CLK pin is toggled to transfer each 1588_SPI_DO bit (as well as the command and address bits). The Timestamp and Frame Identifier Data from the following illustration are sent MSB first down to LSB (bit 0) in the same format as stored in the seven registers of TS FIFO CSRs. For more information, see "Timestamp FIFO," page 84.

The frequency of the generated output 1588_SPI_CLK can be flexibly programmed from 10 MHz up to 62.5 MHz using [REGISTER] to set the number of CSR clocks that the 1588_SPI_CLK is both high and low. For example, to generate a 1588_SPI_CLK that is a divide-by-6 of the csr_clk, the CSR register would be set such that both SPI_CLKs_lo_cycs_i and SPI_CLKs_hi_cycs_i equal 3. Also, the number of CSR clocks after SI_EN asserts before the first 1588_SPI_CLK is programmable (si_en_on_cycs_i), as is the number of clocks before SI_EN negates after the last 1588_SPI_CLK (si_en_off_cycs_i). The number of clocks during which SI_EN is negated between writes is also programmable (si_en_des_cycs_i). The 1588_SPI_CLK may also be configured to be inverted (SPI_CLK_pol_i).

Without considering de-selection between writes, if the PTP 16-bit SequenceID (frame signature) is used as frame identifier each 32 bit timestamp write take 2 + 5 + 32 + 16 = 55 clocks (at 40 MHz) ~1400 ns. This corresponds to a timestamp bandwidth of > 0.7 M timestamp/second.

Figure 51. Serial Timestamp/Frame Signature Output

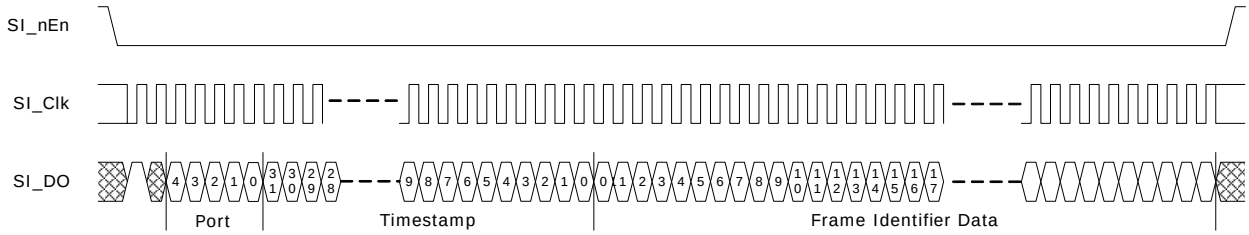

3.10.19 Rewriter

When the rewriter block gets a valid indication it overwrites the input data starting at the offset specified in Rewrite_offset and replaces N bytes of the input data with updated N bytes. Frames are modified by the rewriter as indicated by the analyzer-only PTP/OAM frames are modified by the rewriter.

The output of the rewriter block is the frame data stream that includes both unmodified frames and modified PTP frames. The block also outputs a count of the number of modified PTP frames in INGR_RW_MODFRM_CNT. This counter accumulates the number of PTP frames to which a write was performed and includes errored frames.

3.10.19.1 Rewriter Ethernet FCS Calculation

The rewriter block has to recalculate the Ethernet CRC for the PTP message to modify the contents by writing a new timestamp or clear bytes. Two versions of the Ethernet CRC are calculated in accordance with IEEE 802.3 Clause 3.2.9: one on the unmodified input data stream and one on the modified output data stream. The input frame FCS is checked against the input calculated FCS and if the values match, the frame is good. If they do not, then the frame is considered a bad or errored frame. The new calculated output FCS is used to update the FCS value in the output data frame. If the frame was good, then the FCS is used directly. If the frame was bad, the calculated output FCS is inverted before writing to the frame. Each version of the FCS is calculated in parallel by a separate FCS engine.

A count of the number of PTP/OAM frames that are in error is kept in the counter register (TBD).

3.10.19.2 Rewriter UDP Checksum Calculation

For IPv6/UDP, the rewriter also calculates the value to write into the dummy blocks to correct the UDP checksum. The checksum correction is calculated by taking the original frame's checksum, the value in the dummy bytes, and the new data to be written; and using them to modify the existing value in the dummy byte location. The new dummy byte value is then written to the frame to ensure a valid checksum. The location of the dummy bytes is given by the analyzer. The UDP checksum correction is only performed if Fix_UDP_i is asserted.

3.10.20 Local Time Counter

The local time counter keeps the local time for the device and the time is monitored and synchronized to an external reference by the CPU. The source clock for the counter is selected externally to be a 250 MHz, 200 MHz, 156.25 MHz, 125 MHz, or some other frequency. The clock may be a line clock or a dedicated LTC clock. An external clock mux may be required to select the desired clock source.

To support other frequencies, a flexible counter system is used that can convert almost any frequency in the 125 MHz to 250 MHz range into a usable source clock. The frequency is programmed in terms of the clock period. The seq_a_i input is the nearest whole number of nanoseconds to add to the local time for each ltc_clk cycle. The seq_e_i is the number of femtoseconds of error, and the seq_add_sub_e_i input indicates the direction of the error. When an internal counter adds up the amount of error to a whole nanosecond, then an extra nanosecond is either added or subtracted from the local time. For example to program a frequency of 156.25 MHz (6.4 ns period):

```
Seq_a_i = 6 (6 ns)
Seq_e_i = 400000 (0.4 ns)
Seq_add_sub_e_i = 1 (add an extra nanosecond, in this case add 7 ns)
```

Likewise, to program a 5.9 ns period:

```
Seq_a_i = 6 (6 ns)
Seq_e_i = 100000 (0.1 ns)
Seq_add_sub_e_i = 0 (subtract an extra nanosecond, in this case add 5 ns)
```

To support automatic PPM adjustments, an internal counter runs on the same clock as the local time counter and increments using the same sequence to count nanoseconds. The maximum (rollover) value of the counter in nanoseconds is given in adj_cnt_max_i. At rollover, the next increment of the local time counter is increased or decreased by 1 ns which is indicated by the adj_add_sub_i input. If the input is set to 0x1, addition occurs. When it is set to 0x2, subtraction occurs. When it is 0x0 or 0x3, no adjustment is performed.

To support ppm adjustments on an as-needed basis, the CPU writes to a one-shot input that adds or subtracts 1 from the local time counter once per write. The input is edge detected and the correction applied. External circuitry should reset the input as only rising edges are acted upon. The oneshot_add_i input adds 1 ns and the oneshot_sub_i subtracts 1 ns when pulsed.

When the load_ena_i input is asserted and the ltc_load_save_i input is pulsed, the value on the local_time_i input is latched into the local time counter. This is used to load the counter with the correct local time as set by an external source. If the save_ena_i input is asserted when ltc_load_save_i is asserted, the value in the local time counter is registered and presented on the saved_time_o output for reading by the CPU.

The block also provides a one pulse-per-second output signal with a programmable pulse width which available on an external pin. The pulse width of the 1 pps output is given in the configuration input one_pps_pw_i.

The one pulse-per-second output has an alternate mode that increases the frequency of the pulses. This mode may be used for applications such as locking an external DPLL to the IEEE1588 frequency.

The alternate mode is enabled by setting LTC_CTRL.LTC_ALT_MODE. In the alternate mode the one pulse-per-second output signal is driven with a single bit of the nanosecond field counter of the local time. The alternate mode bit is selected with LTC_CTRL.LTC_ALT_MODE_PPS_BIT[4:0]. The pulse width is not programmable in the alternate mode. The output frequencies that result are 1 divided by powers of 2 nanoseconds (bit 4 = 1/32 ns, bit 5 = 1/64 ns, bit 6 = 1/128 ns, …). The output pulses may jitter by the amount of the programmed nanoseconds of the adder to the local nanoseconds counter, and any automatic or one-shot adjustments.

The following table shows the possible output pulse frequencies (including the range of 4 kHz to 10 MHz) usable for external applications.

Table 25. Output Pulse Frequencies

Nanosecond Counter Bit	Output Pulse Frequency
4	31.25 MHz
5	15.625 MHz
6	7.8125 MHz
7	3.90625 MHz
8	1.953125 MHz
9	976.5625 kHz
10	488.28125 kHz
11	244.140625 kHz
12	122.0703125 kHz
13	61.03515625 kHz

Table 25. Output Pulse Frequencies *(continued)*

Nanosecond Counter Bit	Output Pulse Frequency
14	30.51757813 kHz
15	15.25878906 kHz
16	7.629394531 kHz
17	3.814697266 kHz

The following is an example of an automatic adjustment calculation:

If the 250 MHz clock is off by 100 PPM (.01%) then the 4 ns period is off by 0.0004 ns (0.4 ps) every cycle. So a 1 ns needs to be adjusted every 1 ns/(0.0004 ns) cycles which is 2500. So the adj_cnt_max_i would be set to 2500 x 4 ns which is 10,000 ns.

3.11 Media Recovered Clock Outputs

For Synchronous Ethernet applications, the VSC8574 includes two recovered clock output pins, RCVRDCLK1 and RCVRDCLK2, controlled by registers 23G and 24G, respectively. The recovered clock pins are synchronized to the clock of the active media link.

To enable recovered clock output, set register 23G or 24G, bit 15, to 1. By default, the recovered clock output pins are disabled and held low, including when NRESET is asserted. Registers 23G and 24G also control the PHY port for clock output, the clock source, the clock frequency (either 25 MHz or 125 MHz), and squelch conditions.

3.11.1 Clock Selection Settings

On each pin, the recovered clock supports the following sources, as set by registers 23G or 24G, bits 2:0:

- Fiber SerDes media
- Copper media
- Copper transmitter TCLK output (RCVRDCLK1 only)

Note When using the automatic media sense feature, the recovered clock output cannot automatically change between each active media. Changing the media source must be managed through the recovered clock register settings.

The 10BASE-T mode and 1000BASE-T master mode are not effective for Synchronous Ethernet clock recovery. For 10BASE-T mode, the receiver does not produce a reliable continuous clock source. For 1000BASE-T master mode, the clock is based on the VSC8574 REFCLK input, which is a local clock.

3.11.2 Clock Output Squelch

Under certain conditions, the PHY outputs a clock based on the REFCLK_P and REFCLK_N pins, such as when there is no link present or during autonegotiation. To prevent an undesirable clock from appearing on the recovered clock pins, the VSC8574 squelches, or inhibits, the clock output based on any of the following criteria:

No link is detected (the link status register 1, bit 2 = 0).

The link is found to be unstable using the fast link failure detection feature. The GPIO9/FASTLINK-FAIL pin is asserted high when enabled.

The active link is in 10BASE-T or in 1000BASE-T master mode. These modes produce unreliable recovered clock sources.

CLK_SQUELCH_IN is enabled to squelch the clock.

Use registers 23G or 24G, bits 5:4 to configure the clock squelch criteria. These registers can also disable the squelch feature. The CLK_SQUELCH_IN pin functionality is controlled by register bit 14G.14. When this bit is set to 1, the CLK_SQUELCH_IN pin also controls the squelching of the clock. When enabled, both RCVRDCLK1 and RCVRDCLK2 are squelched when the CLK_SQUELCH_IN pin is high.

3.12 Serial Management Interface

The VSC8574 device includes an IEEE 802.3-compliant serial management interface (SMI) that is affected by use of its MDC and MDIO pins. The SMI provides access to device control and status registers. The register set that controls the SMI consists of 32 16-bit registers, including all required IEEE-specified registers. Also, there are additional pages of registers accessible using device register 31.

Energy efficient Ethernet control registers are available through the SMI using Clause 45 registers and Clause 22 register access in registers 13 through 14. For more information, see Table 46, page 121 and Table 111, page 157.

The SMI is a synchronous serial interface with input data to the VSC8574 on the MDIO pin that is clocked on the rising edge of the MDC pin. The output data is sent on the MDIO pin on the rising edge of the MDC signal. The interface can be clocked at a rate from 0 MHz to 12.5 MHz, depending on the total load on MDIO. An external 2-kΩ pull-up resistor is required on the MDIO pin.

3.12.1 SMI Frames

Data is transferred over the SMI using 32-bit frames with an optional, arbitrary-length preamble. Before the first frame can be sent, at least two clock pulses on MDC must be provided with the MDIO signal at logic one to initialize the SMI state machine. The following illustrations show the SMI frame format for read and write operations.

Figure 52. SMI Read Frame

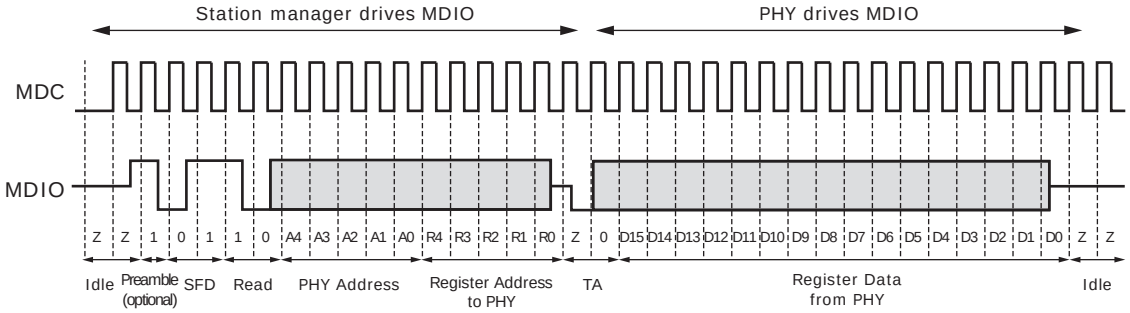

Figure 53. SMI Write Frame

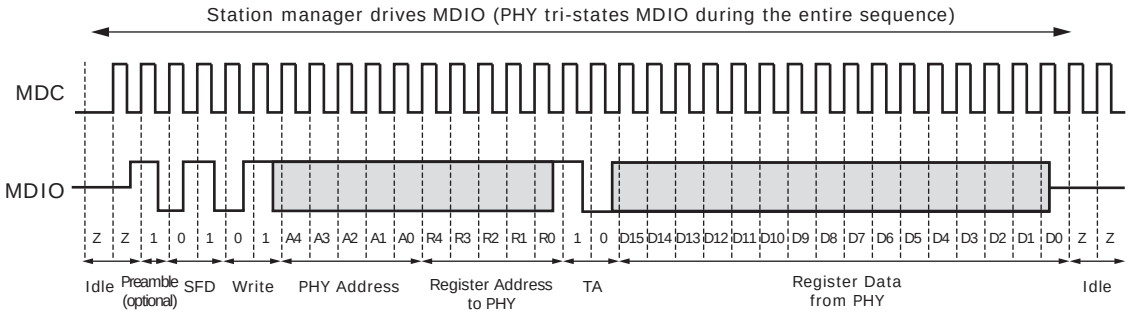

The following list provides additional information about the terms used in the SMI read and write timing diagrams.

- **Idle** During idle, the MDIO node goes to a high-impedance state. This allows an external pull-up resistor to pull the MDIO node up to a logical 1 state. Because the idle mode does not contain any transitions on MDIO, the number of bits is undefined during idle.
- **Preamble** By default, preambles are not expected or required. The preamble is a string of ones. If it exists, the preamble must be at least 1 bit; otherwise, it can be of an arbitrary length.
- **Start of Frame (SFD)** A pattern of 01 indicates the start of frame. If the pattern is not 01, all following bits are ignored until the next preamble pattern is detected.
- **Read or Write Opcode** A pattern of 10 indicates a read. A 01 pattern indicates a write. If the bits are not either 01 or 10, all following bits are ignored until the next preamble pattern is detected.
- **PHY Address** The particular VSC8574 responds to a message frame only when the received PHY address matches its physical address. The physical address is 5 bits long (4:0).
- **Register Address** The next five bits are the register address.
- **Turnaround** The two bits used to avoid signal contention when a read operation is performed on the MDIO are called the turnaround (TA) bits. During read operations, the VSC8574 drives the second TA bit, a logical 0.

- **Data** The 16-bits read from or written to the device are considered the data or data stream. When data is read from a PHY, it is valid at the output from one rising edge of MDC to the next rising edge of MDC. When data is written to the PHY, it must be valid around the rising edge of MDC.
- **Idle** The sequence is repeated.

3.12.2 SMI Interrupt

The SMI includes an output interrupt signal, MDINT, for signaling the station manager when certain events occur in the VSC8574.

3.13 LED Interface

The VSC8574 outputs two LED signals per port, LED0 and LED1, through direct-drive signal outputs, or four LED signals per port (LED0, LED1, LED2, and LED3) through an enhanced serial LED mode. For more information, see "Enhanced Serial LED Mode," page 97. The polarity of the LED outputs is programmable and can be changed through register 17E2, bits 13:10. The default polarity is active low.

Basic serial LED mode is also supported, in which all possible signals that can be displayed on LEDs are sent out on the serial LED stream for further processing by an external programmable device. For more information, see "Basic Serial LED Mode," page 96. Four of the GPIO pins, TBD, are multipurpose and can be configured to serve as LED input pins. For more information, see Table 95, page 148.

3.13.1 LED Modes

Each LED pin can be configured to display different status information that can be selected by setting the LED mode in register 29. The modes listed in the following table are equivalent to the setting used in register 29 to configure each LED pin. The default LED state is active low and can be changed by modifying the value in register 17E2, bits 13:10. The blink/pulse-stretch is dependent on the LED behavior setting in register 30.

The following table provides a summary of the LED modes and functions.

Table 26. LED Mode and Function Summary

Mode	Function Name	LED State and Description
0	Link/Activity	1: No link in any speed on any media interface. 0: Valid link at any speed on any media interface. Blink or pulse-stretch = Valid link at any speed on any media interface with activity present.
1	Link1000/Activity	1: No link in 1000BASE-T or 1000BASE-X. 0: Valid 1000BASE-T or 1000BASE-X. Blink or pulse-stretch = Valid 1000BASE-T or 1000BASE-X link with activity present.
2	Link100/Activity	1: No link in 100BASE-TX or 100BASE-FX. 0: Valid 100BASE-TX or 100BASE-FX. Blink or pulse-stretch = Valid 100BASE-TX or 100BASE-FX link with activity present.

Table 26. LED Mode and Function Summary *(continued)*

Mode	Function Name	LED State and Description
3	Link10/Activity	1: No link in 10BASE-T. 0: Valid 10BASE-T link. Blink or pulse-stretch = Valid 10BASE-T link with activity present.
4	Link100/1000/Activity	1: No link in 100BASE-TX, 100BASE-FX, 1000BASE-X, or 1000BASE-T. 0: Valid 100BASE-TX, 100BASE-FX, 1000BASE-X, or 1000BASE-T link. Blink or pulse-stretch = Valid 100BASE-TX, 100BASE-FX, 1000BASE-X, or 1000BASE-T link with activity present.
5	Link10/1000/Activity	1: No link in 10BASE-T, 1000BASE-X, or 1000BASE-T. 0: Valid 10BASE-T, 1000BASE-X, or 1000BASET-T link. Blink or pulse-stretch = Valid 10BASE-T, 1000BASE-X, or 1000BASE-T link with activity present.
6	Link10/100/Activity	1: No link in 10BASE-T, 100BASE-FX, or 100BASE-TX. 0: Valid 10BASE-T 100BASE-FX, or 100BASE-TX, link. Blink or pulse-stretch = Valid 10BASE-T, 100BASE-FX, or 100BASE-TX link with activity present.
7	Link100BASE-FX/ 1000BASE-X/Activity	1: No link in 100BASE-FX or 1000BASE-X. 0: Valid 100BASE-FX or 1000BASE-X link. Blink or pulse-stretch = Valid 100BASE-FX or 1000BASE-X link with activity present.
8	Duplex/Collision	1: Link established in half-duplex mode, or no link established. 0: Link established in full-duplex mode. Blink or pulse-stretch = Link established in half-duplex mode but collisions are present.
9	Collision	1: No collision detected. Blink or pulse-stretch = Collision detected.
10	Activity	1: No activity present. Blink or pulse-stretch = Activity present (becomes TX activity present when register bit 30.14 is set to 1).
11	100BASE-FX/1000BASE-X Fiber Activity	1: No 100BASE-FX or 1000BASE-X activity present. Blink or pulse-stretch = 100BASE-FX or 1000BASE-X activity present (becomes RX activity present when register bit 30.14 is set to 1).
12	Autonegotiation Fault	1: No autonegotiation fault present. 0: Autonegotiation fault occurred.
13	Serial Mode	Serial stream. See "Basic Serial LED Mode," page 96. Only relevant on PHY port 0 and reserved in others.
14	Force LED Off	1: De-asserts the LED(1).
15	Force LED On	0: Asserts the LED(1).

1. Setting this mode suppresses LED blinking after reset.

3.13.2 Extended LED Modes

In addition to the LED modes in register 29, there are also additional LED modes that are enabled on the LED0_3:0 pins whenever the corresponding register 19E1, bits 15 to 12 are set to 1. Each of these bits enables extended modes on a specific LED pin and

these extended modes are shown in the following table. For example, LED0 = mode 17 means that register 19E1 bit 12 = 1 and register 29 bits 3 to 0 = 0001.

The following table provides a summary of the extended LED modes and functions.

Table 27. Extended LED Mode and Function Summary

Mode	Function Name	LED State and Description
16	Link1000BASE-X Activity	1: No link in 1000BASE-X. 0: Valid 1000BASE-X link.
17	Link100BASE-FX Activity	1: No link in 100BASE-FX. 0: Valid 100BASE-FX link.
18	1000BASE-X Activity	1: No 1000BASE-X activity present. Blink or pulse-stretch = 1000BASE-X activity present.
19	100BASE-FX Activity	1: No 100BASE-FX activity present. Blink or pulse-stretch = 100BASE-FX activity present.
20	Force LED Off	1: De-asserts the LED.
21	Force LED On	0: Asserts the LED. LED pulsing is disabled in this mode.
22	Fast Link Fail	1: Enable fast link fail on the LED0_0 pin 0: Disable

3.13.3 LED Behavior

Several LED behaviors can be programmed into the VSC8574. Use the settings in register 30 and 19E1 to program LED behavior, which includes the following.

LED Combine Enables an LED to display the status for a combination of primary and secondary modes. This can be enabled or disabled for each LED pin. For example, a copper link running in 1000BASE-T mode and activity present can be displayed with one LED by configuring an LED pin to Link1000/Activity mode. The LED asserts when linked to a 1000BASE-T partner and also blinks or performs pulse-stretch when activity is either transmitted by the PHY or received by the Link Partner. When disabled, the combine feature only provides status of the selected primary function. In this example, only Link1000 asserts the LED, and the secondary mode, activity, does not display when the combine feature is disabled.

LED Blink or Pulse-Stretch This behavior is used for activity and collision indication. This can be uniquely configured for each LED pin. Activity and collision events can occur randomly and intermittently throughout the link-up period. Blink is a 50% duty cycle oscillation of asserting and de-asserting an LED pin. Pulse-stretch guarantees that an LED is asserted and de-asserted for a specific period of time when activity is either present or not present. These rates can also be configured using a register setting.

Rate of LED Blink or Pulse-Stretch This behavior controls the LED blink rate or pulse-stretch length when blink/pulse-stretch is enabled on an LED pin. The blink rate, which alternates between a high and low voltage level at a 50% duty cycle, can be set to 2.5 Hz, 5 Hz, 10 Hz, or 20 Hz. For pulse-stretch, the rate can be set to 50 ms, 100 ms, 200 ms, or 400 ms. The blink rate selection for PHY0 globally sets the rate used for all LED pins on all PHY ports.

LED Pulsing Enable To provide additional power savings, the LEDs (when asserted) can be pulsed at 5 kHz, 20% duty cycle.

LED Blink After Reset The LEDs will blink for one second after power-up and after any time all resets have been de-asserted. This can be disabled through register 19E1, bit 11 = 0.

Fiber LED Disable This bit controls whether the LEDs indicate the fiber and copper status (default) or the copper status only.

Pulse Programmable Control These bits add the ability to width and frequency of LED pulses. This feature facilitates power reduction options.

Fast Link Failure For more information about this feature, see "Fast Link Failure Indication," page 97.

3.13.4 Basic Serial LED Mode

Optionally, the VSC8574 can be configured so that access to all its LED signals is available through two pins. This option is enabled by setting LED0 on PHY0 to serial LED mode. When serial LED mode is enabled on PHY0, the LED0_0 pin becomes the serial data pin, and the LED0_1 pin becomes the serial clock pin. All other LED pins can still be configured normally. The serial LED mode clocks the 48 LED status bits on the rising edge of the serial clock.

The LED behavior settings can also be used in serial LED mode. The controls are used on a per-PHY basis, where the LED combine and LED blink or pulse-stretch setting of LED0_n for each PHY is used to control the behavior of each bit of the serial LED stream for each corresponding PHY. To configure LED behavior, set device register 30.

The serial bitstream outputs, 1 through 48, of each LED signal are shown in the following table beginning with PHY port 0 and ending with PHY port 3. The individual signals can be clocked in the order shown.

Table 28. LED Serial Stream Order

PHY0	PHY1	PHY*n*
Bit 1. Link/Activity	Bit 13. Link/Activity	Bit *n*. Link/Activity
Bit 2. Link1000/Activity	Bit 14. Link1000/Activity	Bit *n*. Link1000/Activity
Bit 3. Link100/Activity	Bit 15. Link100/Activity	Bit *n*. Link100/Activity
Bit 4. Link10/Activity	Bit 16. Link10/Activity	Bit *n*. Link10/Activity
Bit 5. Fiber Link/Activity	Bit 17. Fiber Link/Activity	Bit *n*. Fiber Link/Activity
Bit 6. Duplex/Collision	Bit 18. Duplex/Collision	Bit *n*. Duplex/Collision
Bit 7. Collision	Bit 19. Collision	Bit *n*. Collision
Bit 8. Activity	Bit 20. Activity	Bit *n*. Activity
Bit 9. Fiber Activity	Bit 21. Fiber Activity	Bit *n*. Fiber Activity
Bit 10. TX Activity	Bit 22. TX Activity	Bit *n*. TX Activity
Bit 11. RX Activity	Bit 23. RX Activity	Bit *n*. RX Activity
Bit 12. Autonegotiation Fault	Bit 24. Autonegotiation Fault	Bit *n*. Autonegotiation Fault

3.13.5 Enhanced Serial LED Mode

VSC8574 can be configured to output up to four LED signals per port on a serial stream that can be de-serialized externally to drive LEDs on the system board. This functionality is controlled by setting register 25G, bits 7:1. In this mode, the serial LED_DATA is shifted out on the falling edge of LED_CLK and is latched in the external serial-to-parallel converter on the rising edge of LED_CLK. The falling edge of LED_LD signal can be used to shift the data from the shift register in the converter to the parallel output drive register. When a separate parallel output drive register is not used in the external serial-to-parallel converter, the LEDs will blink at a high frequency as the data bits are being shifted through, which may be undesirable. The LED_PULSE signal provides a 5 kHz pulse stream whose duty cycle can be modulated to turn on/off LEDs at a high rate. This signal can be tied to the output enable signal of the serial-to-parallel converter to provide the LED dimming functionality to save energy.

3.13.6 LED Port Swapping

For additional hardware configurations, the VSC8574 can have its LED port order swapped. This is a useful feature to help simplify PCB layout design. Register 25G bit 0 controls the LED port swapping mode.

3.14 Fast Link Failure Indication

To aid Synchronous Ethernet applications, the VSC8574 can indicate the onset of a link failure in less than 1 ms. By comparison, the IEEE 802.3 standard establishes a delay of up to 750 ms before indicating that a 1000BASE-T link is no longer present. A fast link failure indication is critical to support ports used in a synchronization timing link application. The fast link failure indication works for all copper and fiber media speeds. Fast link failure is supported for each PHY port through the GPIO9/FASTLINK-FAIL pin.

Note For all links except 1000BASE-T, the fast link failure indication matches the link status register (address 1, bit 2). For 1000BASE-T links, the link failure is based on a circuit that analyzes the integrity of the link, and at the indication of failure, will assert.

3.15 Integrated I2C Multiplexer

The VSC8574 includes an integrated quad I2C multiplexer (MUX), eliminating the need for an external two-wire serial device for the control and status of SFP or PoE modules. There are five I2C controller pins: four clocks and one shared data pin. Each SFP or PoE connects to the multipurpose GPIO[7:4]_I2C_SCL_[3:0] and GPIO8/I2C_SDA device pins, which must be configured to the corresponding I2C function. For more information about configuring the pins, see “I2C MUX Control 1,” page 150. For SFP modules, VSC8574 can also provide control for the MODULE_DETECT and TX_DIS module pins using the multipurpose LED and GPIO pins.

Figure 54. I2C MUX with SFP Control and Status

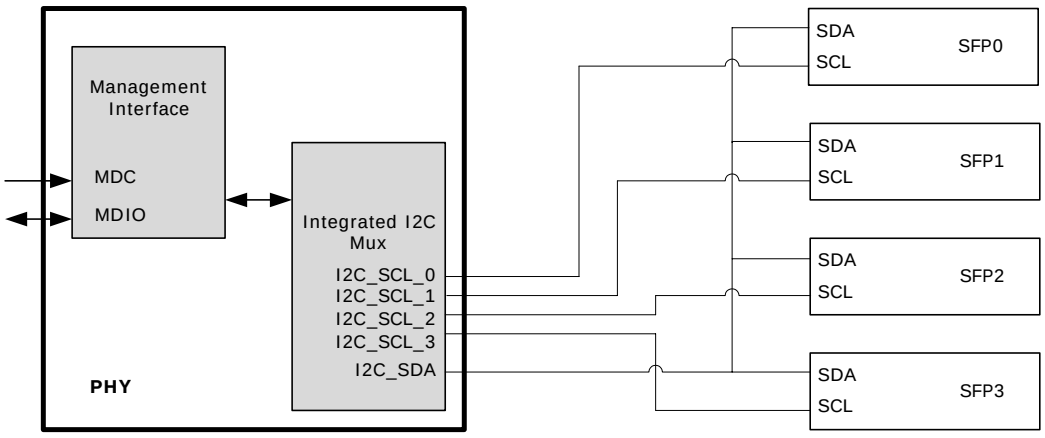

3.15.1 Read/Write Access Using the I2C MUX

Using the integrated I2C MUX, the VSC8574 device can read and write to an SFP or PoE module through the SCL and SDA pins. If the ability is required to write to the slave I2C device, refer to the device's specific datasheet for more information.

Note The VSC8574 device does not automatically increment the I2C address. Each desired address must be intentionally set.

Main control of the integrated I2C MUX is available through register 20G. The I2C MUX pins are enabled or disabled using register 20G bits 3:0. Register 20G bits 15:9 set the I2C device address (the default is 0xA0). Using register 20G bits 5:4, the two-wire serial frequency can be changed from 100 kHz to other speeds, such as 50 kHz, 100 kHz (the default), 400 kHz, and 2 MHz. Registers 21G and 22G provide status and control of the read/write process. The following illustration shows the read and write register flow.

Figure 55. I2C MUX Read and Write Register Flow

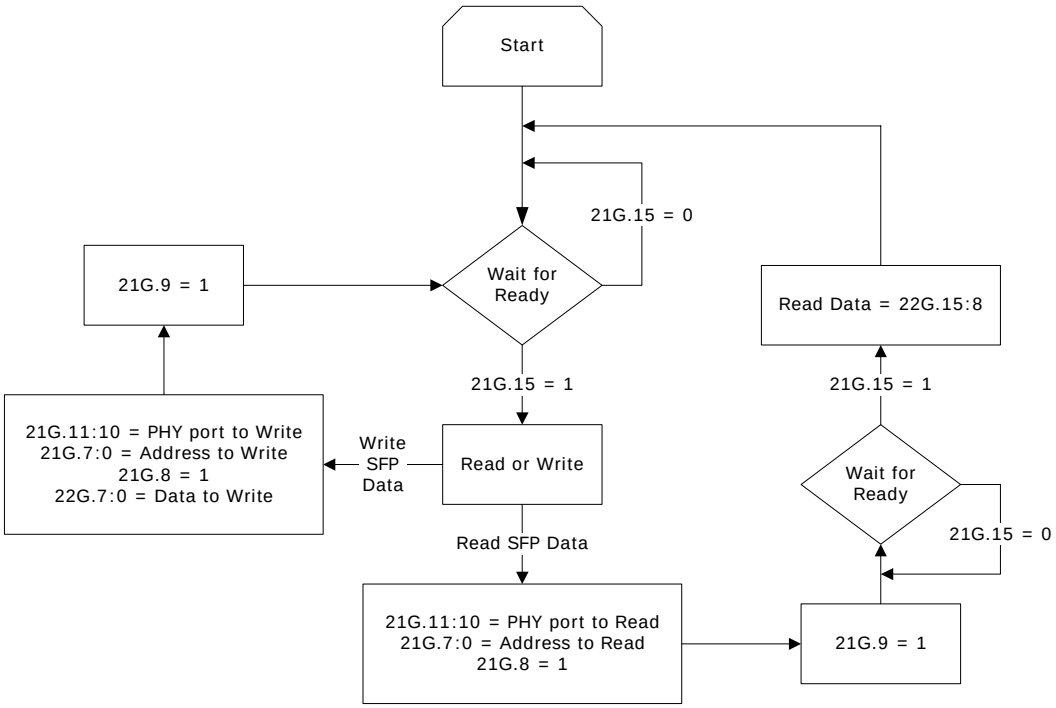

To read a value from a specific address of the I2C slave device:

1. Read the VSC8574 device register 21G bit 15, and ensure that it is set.
2. Write the PHY port address to be read to register 21G bits 11:10.
3. Write the I2C address to be read to register 21G bits 7:0.
4. Set both register 21G bits 8 and 9 to 1.
5. When register 21G bit 15 changes to 1, read the 8-bit data value found at register 22G bits 15:8. This is the contents of the address just read by the PHY.

To write a value to a specific address of the I2C slave device:

1. Read the VSC8574 device register 21G bit 15 and ensure that it is set.
2. Write the PHY port address to be written to register 21G bits 11:10.
3. Write the address to be written to register 21G bits 7:0.
4. Set register 21 bit 8 to 0.
5. Set register 22G bits 7:0 with the 8-bit value to be written to the slave device.
6. Set register 21G bit 9 to 1.

To avoid collisions during read and write transactions on the I2C serial bus, always wait until register 21G bit 15 changes to 1 before performing another I2C read or write operation.

3.16 GPIO Pins

The VSC8574 provides 10 multiplexed general purpose input/output (GPIO) pins. The GPIO_TBD pins can be reconfigured as LED pins. All device GPIO pins and their behavior are controlled using registers. For more information, see "General Purpose Registers," page 147.

3.17 Testing Features

The VSC8574 device includes several testing features designed to facilitate performing system-level debugging and in-system production testing. This section describes the available features.

3.17.1 Digital Temperature, Core Voltage, and I/O Voltage Monitor

The VSC8574 device contains an accurate temperature monitoring circuit for the purpose of measuring core temperature. The temperature is converted to a digital representation and can be read out using software through the SMI bus. The temperature value that is read out is accurate to within ±TBD °C.

In addition to providing the core temperature, the monitoring circuit also provides a readout of the I/O and core supply voltages. The voltage value that is read out is accurate to within ±TBD V.

For more information about the temperature monitor control registers, see Table 107, page 154, Table 108, page 156, and Table 109, page 156. For more information about the temperature monitor interrupt source control bit 12, see Table 110, page 157.

3.17.2 Ethernet Packet Generator

The Ethernet packet generator (EPG) can be used at each of the 10/100/1000BASE-T speed settings for copper Cat5 media and fiber media to isolate problems between the MAC and the VSC8574, or between a locally connected PHY and its remote link partner. Enabling the EPG feature effectively disables all MAC interface transmit pins and selects the EPG as the source for all data transmitted onto the twisted pair interface. This feature is not used when the SerDes media is set to pass-through mode.

Important The EPG is intended for use with laboratory or in-system testing equipment only. Do not use the EPG testing feature when the VSC8574 is connected to a live network.

To enable the VSC8574 EPG feature, set the device register bit 29E1.15 to 1.

When the EPG is enabled, packet loss occurs during transmission of packets from the MAC to the PHY. However, the PHY receive output pins to the MAC are still active when the EPG is enabled. When it is necessary to disable the MAC receive pins as well, set the register bit 0.10 to 1.

When the device register bit 29E1.14 is set to 1, the PHY begins transmitting Ethernet packets based on the settings in registers 29E1 and 30E1. These registers set:

- Source and destination addresses for each packet
- Packet size
- Interpacket gap
- FCS state
- Transmit duration
- Payload pattern

When register bit 29E1.13 is set to 0, register bit 29E1.14 is cleared automatically after 30,000,000 packets are transmitted.

3.17.3 CRC Counters

Two sets of cyclical redundancy check (CRC) counters are available in all PHYs in VSC8574. One set monitors traffic on the copper interface and the other set monitors traffic on the SerDes interface.

The device CRC counters operate in the 100BASE-FX/1000BASE-X over SerDes mode as well as in the 10/100/1000BASE-T mode as follows:

After receiving a packet on the media interface, register bit 15 in register 18E1 or register 28E3 is set and cleared after being read.

The packet then is counted by either the good CRC counter or the bad CRC counter.

Both CRC counters are also automatically cleared when read.

The good CRC counter's highest value is 9,999 packets. After this value is reached, the counter clears on the 10,000th packet and continues to count additional packets beyond that value. The bad CRC counter stops counting when it reaches its maximum counter limit of 255 packets.

Copper Interface CRC Counters Two separate CRC counters are available and reside between the copper interface PCSs and SerDes MAC interface. There is a 14-bit good CRC counter available through register bits 18E1.13:0 and a separate 8-bit bad CRC counter available in register bits 23E1.7:0.

SerDes Interface CRC Counters Two separate CRC counters are available and reside between the SerDes media interface PCSs and SerDes MAC interface. There is a 14-bit good CRC counter available through register bits 28E3.13:0 and a separate 8-bit bad CRC counter available in register bits 29E3.7:0.

3.17.4 Far-End Loopback

The far-end loopback testing feature is enabled by setting register bit 23.3 to 1. When enabled, it forces incoming data from a link partner on the current media interface, into the MAC interface of the PHY, to be retransmitted back to the link partner on the media interface as shown in the following illustration. In addition, the incoming data also appears on the receive data pins of the MAC interface. Data present on the transmit data pins of the MAC interface is ignored when using this testing feature.

Figure 56. Far-End Loopback Diagram

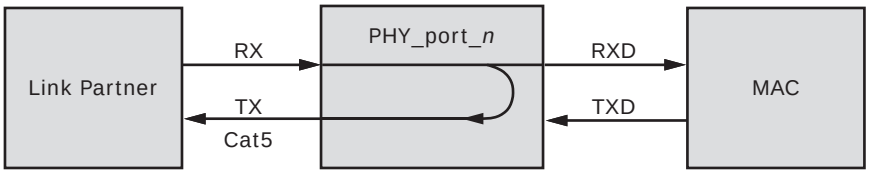

3.17.5 Near-End Loopback

When the near-end loopback testing feature is enabled, transmitted data (TXD) is looped back in the PCS block onto the receive data signals (RXD), as shown in the following illustration. When using this testing feature, no data is transmitted over the network. To enable near-end loopback, set the device register bit 0.14 to 1.

Figure 57. Near-End Loopback Diagram

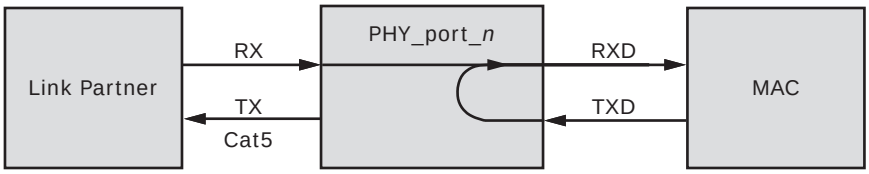

3.17.6 Connector Loopback

The connector loopback testing feature allows the twisted pair interface to be looped back externally. When using this feature, the PHY must be connected to a loopback connector or a loopback cable. Connect pair A to pair B, and pair C to pair D, as shown in the following illustration. The connector loopback feature functions at all available interface speeds.

Figure 58. Connector Loopback Diagram

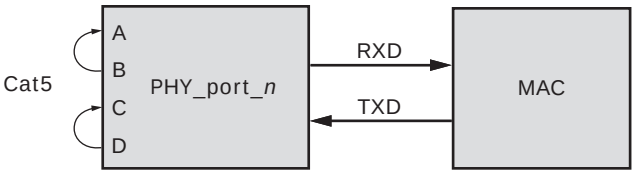

When using the connector loopback testing feature, the device autonegotiation, speed, and duplex configuration is set using device registers 0, 4, and 9. For 1000BASE-T connector loopback, the following additional writes are required. Execute the additional writes in the following order:

1. Enable the 1000BASE-T connector loopback. Set register bit 24.0 to 1.
2. Disable pair swap correction. Set register bit 18.5 to 1.

3.17.7 SerDes Loopbacks

For test purposes, the SerDes and enhanced SerDes macro interfaces provide several data loops. The following illustration shows the SerDes loopbacks.

Figure 59. Data Loops of the SerDes Macro

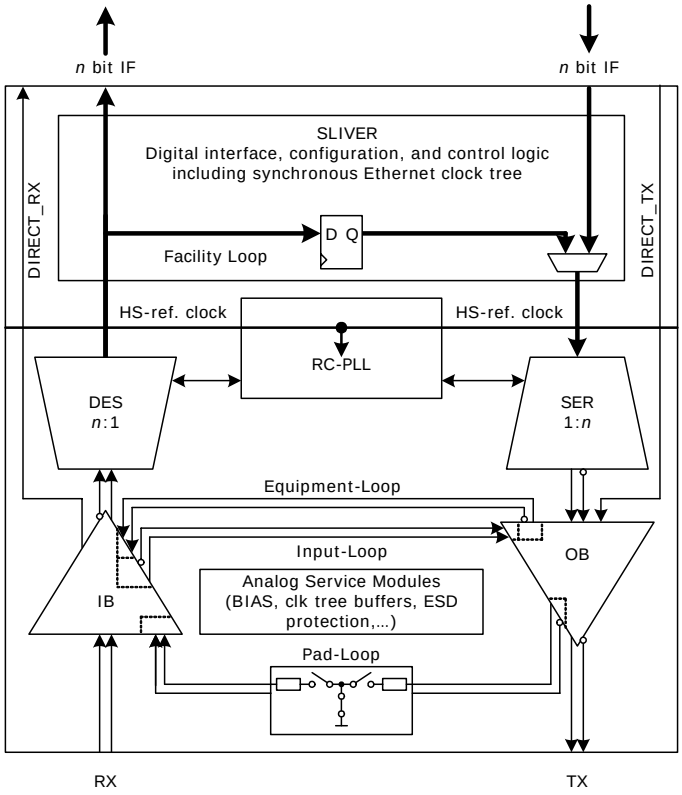

SGMII Mode When the MAC interface is configured in SGMII mode, write the following 16-bit value to register 18G:

Bits 15:12 0x9

Bits 11:8: Port address (0x0 to 0xb)

Bits 7:4: Loopback type

Bits 3:0: 0x2

where loopback type is:

0x0: No loopback

0x1: Pad loopback

0x2: Input loopback

0x4: Facility loopback

0x8: Equipment loopback

QSGMII Mode When the MAC interface is configured in QSGMII mode, the configuration command is the same as SGMII mode, except that the port addresses (bits 11:8) are 0xC for the Enhanced SerDes macro for ports 0–3.

Note Loopback configuration affects all four ports associated with a QSGMII. Individual port loopback within a QSGMII is not possible.

Fiber Media Port Mode When the SerDes is configured as a fiber media port, write the following 16-bit value to register 18G:

Bits 15:12: 0x8

Bits 11:8: Port address

Bits 7:4: Loopback type

Bits 3:0: 0x2

where port address is:

0x1: Fiber8 port

0x2: Fiber9 port

0x4: Fiber10 port

0x8: Fiber11 port

Port addresses for fiber media SerDes can be OR'ed together to address multiple ports using a single command. bit 18G.15 will be cleared when the internal configuration is complete.

Facility Loop The recovered and de-multiplexer deserializer data output is looped back to the serializer data input and replaces the data delivered by the digital core. This test loop provides the possibility to test the complete analog macro data path from outside including input buffer, clock and data recovery, serialization and output buffer. The data received by the input buffer must be transmitted by the output buffer after some delay.

Equipment Loop The 1-bit data stream at the serializer output is looped back to the deserializer and replaces the received data stream from the input buffer. This test loop provides the possibility to verify the digital data path internally. The transmit data goes through the serialization, the clock and data recovery and deserialization before the data is fed back to the digital core.

Input Loop The received 1-bit data stream of the input buffer is looped back asynchronously to the output buffer. This test loop provides the possibility to the test only the analog parts of the SGMII interface because only the input and output buffer are part of this loop.

Pad Loop The 1-bit data stream at the output buffer output is looped back to the input buffer input and added to the differential pad signal. Therefore, the input pad must not be driven when the output loop is activated. The test loop provides a means to test the complete SGMII macro data path, including the input and output buffers.

3.17.8 VeriPHY Cable Diagnostics

The VSC8574 includes a comprehensive suite of cable diagnostic functions that are available using SMI reads and writes. These functions enable a variety of cable operating conditions and status to be accessed and checked. The VeriPHY suite has the ability to identify the cable length and operating conditions and to isolate a variety of common faults that can occur on the Cat5 twisted pair cabling.

Note When a link is established on the twisted pair interface in the 1000BASE-T mode, VeriPHY can run without disrupting the link or disrupting any data transfer. However, when a link is established in 100BASE-TX or 10BASE-T modes, VeriPHY causes the link to drop while the diagnostics are running. After diagnostics are finished, the link is re-established.

The following diagnostic functions are part of the VeriPHY suite:

- Detecting coupling between cable pairs
- Detecting cable pair termination
- Determining cable length

Coupling Between Cable Pairs Shorted wires, improper termination, or high crosstalk resulting from an incorrect wire map can cause error conditions, such as anomalous coupling between cable pairs. These conditions can prevent the device from establishing a link in any speed.

Cable Pair Termination Proper termination of Cat5 cable requires a 100 Ω differential impedance between the positive and negative cable terminals. IEEE 802.3 allows for a termination of 115 Ω maximum and 85 Ω minimum. If the termination falls outside of this range, it is reported by the VeriPHY diagnostics as an anomalous termination. The diagnostics can also determine the presence of an open or shorted cable pair.

Cable Length When the Cat5 cable in an installation is properly terminated, VeriPHY reports the approximate cable length in meters.

Mean Square Error Noise The average absolute error can be read out when either a 100BASE-TX or 1000BASE-T link is established. In the case of 1000BASE-T link, there are four average absolute error terms, one for each twisted pair over which signal is received. Use the following script to read average absolute error for 100BASE-TX:

```
PhyWrite( <phy>, 31, 0x52b5 );

PhyWrite( <phy>, 16, 0xa3c0 );

PhyRead( <phy>, 16 );

tmp17 = PhyRead( <phy>, 17 );

tmp18 = PhyRead( <phy>, 18 );

mse = (tmp18 << 4) | (tmp17 >> 12);

PhyWrite( <phy>, 31, 0 );
```

The returned average absolute error is in units of 1/2,048 and can be found in the mse variable.

```
PhyWrite( <phy>, 31, 0x52b5 );

PhyWrite( <phy>, 16, 0xa3c0 );

PhyRead( <phy>, 16 );

tmp17 = PhyRead( <phy>, 17 );

tmp18 = PhyRead( <phy>, 18 );

mseA = (tmp18 << 4) | (tmp17 >> 12);

mseB = tmp17 & 0x0fff;

PhyWrite( <phy>, 16, 0xa3c2 );

PhyRead( <phy>, 16 );

tmp17 = PhyRead( <phy>, 17 );

tmp18 = PhyRead( <phy>, 18 );

mseC = (tmp18 << 4) | (tmp17 >> 12);

mseD = tmp17 & 0x0fff;

PhyWrite( <phy>, 31, 0 );
```

The returned average absolute error is in units of 1/2,048 and can be found in the mseA, mseB, mseC, and mseD variables for each twisted pair.

3.17.9 JTAG Boundary Scan

The VSC8574 supports the test access port (TAP) and boundary scan architecture described in IEEE 1149.1. The device includes an IEEE 1149.1-compliant test interface, referred to as a JTAG TAP interface.

The JTAG boundary scan logic on the VSC8574, accessed using its TAP interface, consists of a boundary scan register and other logic control blocks. The TAP controller includes all IEEE-required signals (TMS, TCK, TDI, and TDO), in addition to the optional asynchronous reset signal TRST. The following illustration shows the TAP and boundary scan architecture.

Important When JTAG is not in use, the TRST pin must be tied to ground with a pull-down resistor for normal operation.

Figure 60. Test Access Port and Boundary Scan Architecture

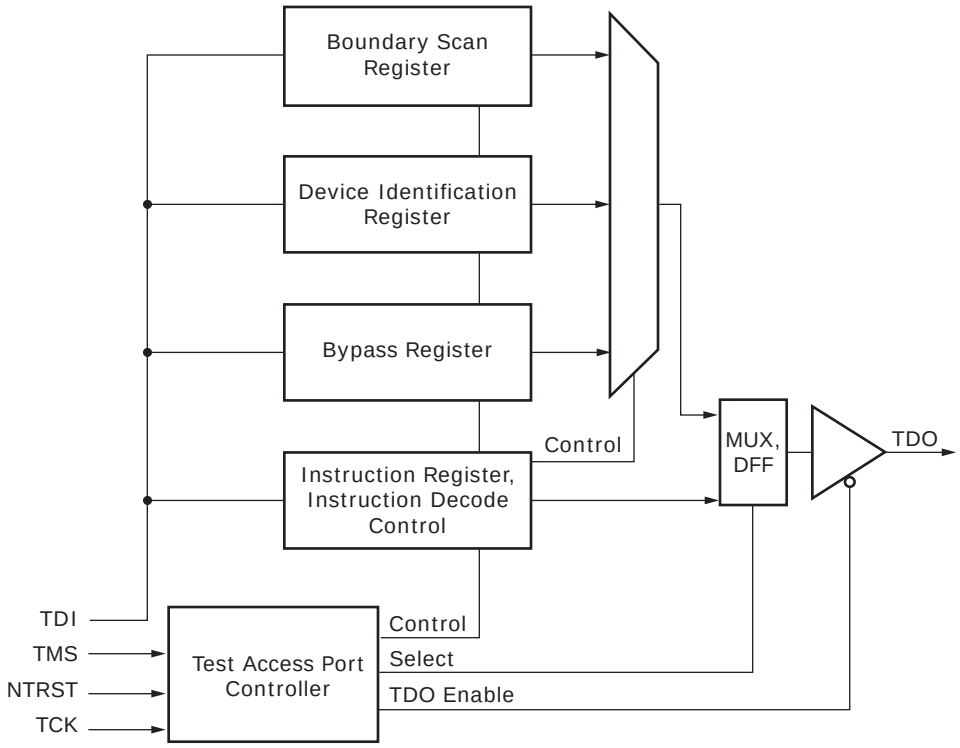

After a TAP reset, the device identification register is serially connected between TDI and TDO by default. The TAP instruction register is loaded either from a shift register when a new instruction is shifted in, or, if there is no new instruction in the shift register, a default value of 6'b100100 (IDCODE) is loaded. Using this method, there is always a valid code in the instruction register, and the problem of toggling instruction bits during a shift is avoided. Unused codes are mapped to the BYPASS instruction.

3.17.10 JTAG Instruction Codes

The VSC8574 supports the following instruction codes:

Table 29. JTAG Instruction Codes

Instruction Code	Description
BYPASS	The bypass register contains a single shift-register stage and is used to provide a minimum-length serial path (one TCK clock period) between TDI and TDO to bypass the device when no test operation is required.
CLAMP	Allows the state of the signals driven from the component pins to be determined from the boundary scan register while the bypass register is selected as the serial path between TDI and TDO. While the CLAMP instruction is selected, the signals driven from the component pins do not change.

Table 29. JTAG Instruction Codes *(continued)*

Instruction Code	Description
EXTEST	Allows tests of the off-chip circuitry and board-level interconnections by sampling input pins and loading data onto output pins. Outputs are driven by the contents of the boundary scan cells, which have to be updated with valid values, with the PRELOAD instruction, prior to the EXTEST instruction.
HIGHZ	Places the component in a state in which all of its system logic outputs are placed in a high-impedance state. In this state, an in-circuit test system can drive signals onto the connections normally driven by a component output without incurring a risk of damage to the component. This makes it possible to use a board where not all of the components are compatible with the IEEE 1149.1 standard.
IDCODE	Provides the version number (bits 31:28), device family ID (bits 27:12), and the manufacturer identity (bits 11:1) to be serially read from the device.
SAMPLE/PRELOAD	Allows a snapshot of inputs and outputs during normal system operation to be taken and examined. It also allows data values to be loaded into the boundary scan cells prior to the selection of other boundary scan test instructions.
USERCODE	Provides the version number (bits 31:28), part number (bits 27:12), and the manufacturer identity (bits 11:1) to be serially read from the device.

The following tables provide information about the IDCODE and USERCODE binary values stored in the device JTAG registers.

Table 30. IDCODE JTAG Device Identification Register Descriptions

Description	Device Version	Family ID	Manufacturing Identity	LSB
Bit field	31–28	27–12	11–1	0
Binary value	0000	1000 0101 0111 0100	000 0111 0100	1

Table 31. USERCODE JTAG Device Identification Register Descriptions

Description	Device Version	Model Number	Manufacturing Identity	LSB
Bit field	31–28	27–12	11–1	0
Binary value	0000	1011 0000 0000 0001	000 0111 0100	1

The following table provides information about the location and IEEE compliance of the JTAG instruction codes used in the VSC8574. Instructions not explicitly listed in the table are reserved. For more information about these IEEE specifications, visit the IEEE Web site at www.IEEE.org.

Table 32. JTAG Instruction Code IEEE Compliance

Instruction	Code	Selected Register	Register Width	IEEE 1149.1	IEEE 1149.6
Width	IEEE 1149.1	IEEE 1149.6			
EXTEST	6'b000000	Boundary Scan	161	Mandatory	
SAMPLE/ PRELOAD	6'b000001	Boundary Scan	161	Mandatory	

Table 32. JTAG Instruction Code IEEE Compliance *(continued)*

Instruction	Code	Selected Register	Register Width	IEEE 1149.1	IEEE 1149.6
IDCODE	6'b100100	Device Identification	32	Optional	
USERCODE	6'b100101	Device Identification	32	Optional	
CLAMP	6'b000010	Bypass Register	1	Optional	
HIGHZ	6'b000101	Bypass Register	1	Optional	
BYPASS	6'b111111	Bypass Register	1	Mandatory	
EXTEST_PULSE	6'b000011	Boundary Scan	161		Mandatory
EXTEST_TRAIN	6'b000100	Boundary Scan	161		Mandatory

3.17.11 Boundary Scan Register Cell Order

All inputs and outputs are observed in the boundary scan register cells. All outputs are additionally driven by the contents of boundary scan register cells. Bidirectional pins have all three related boundary scan register cells: input, output, and control.

The complete boundary scan cell order is available as a BSDL file format on the Vitesse Web site at www.vitesse.com.

3.17.12 JTAG Boundary Scan Interface

The IEEE 1149.6 AC-JTAG solution integrated on all SerDes ports of the VSC8574 extends the capability of IEEE 1149.1 boundary scan for robust board-level testing. This interface is backward-compatible to the IEEE 1149.1 standard.

3.18 100FX Halt Code Transmission and Reception

The VSC8574 device supports transmission and reception of halt code words in 100BASE-FX mode. There are three separate scripts provided to initiate transmission of halt code words, stop transmission of halt code words and detect reception of halt code words. Use the following scripts to implement each of these functions:

Sending the HALT codeword:

```
PhyWrite( <phy>, 31, 0x52b5 );

PhyWrite( <phy>, 16, 0xac82 );

reg18 = PhyRead( <phy>, 18 );

reg18 = (reg18 & 0xf0) | 0x0c;

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, 0xe739 );

PhyWrite( <phy>, 16, 0x8c82 );
```

```
PhyWrite< <phy>, 16, 0xbe80 );

reg17 = PhyRead( <phy>, 17 );

reg18 = PhyRead( <phy>, 18 );

reg17 = reg17 | 0x0040;

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, reg17 );

PhyWrite( <phy>, 16, 0x9e80 );

PhyWrite( <phy>, 31, 0 );
```

Stop sending the HALT codeword:

```
PhyWrite( <phy>, 31, 0x52b5 );

PhyWrite< <phy>, 16, 0xbe80 );

reg17 = PhyRead( <phy>, 17 );

reg18 = PhyRead( <phy>, 18 );

reg17 = reg17 & ~0x0040;

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, reg17 );

PhyWrite( <phy>, 16, 0x9e80 );

PhyWrite( <phy>, 31, 0 );
```

Detecting whether the HALT codeword is being sent by the link partner:

```
long patternset[5] = {

     0xce739,

     0xe739c,

     0x739ce,

     0x39ce7,

     0x9ce73

};
```

Turning on the pattern checker:

```
PhyWrite( <phy>, 31, 0x52b5 );

PhyWrite( <phy>, 16, 0xbe80 );
```

```
reg18 = PhyRead( <phy>, 18 );

reg17 = PhyRead( <phy>, 17 );

reg17 = reg17 | 4;

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, reg17 );

PhyWrite( <phy>, 16, 0x9e80 );
```

Sweeping through all five pattern shifts checking for a match:

```
for (i = 0, matchfailed = 1; i < 5 && matchfailed; ++i) {

PhyWrite( <phy>, 16, 0xac84 );

reg18 = PhyRead( <phy>, 18 );

reg18 = (reg18 & 0xf0) | (patternset[i] >> 16)

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, patternset[i] & 0xffff );

PhyWrite( <phy>, 16, 0x8c84 );


PhyWrite( <phy>, 16, 0xbe84 ); // Dummy read to clear latched mismatch

PhyWrite( <phy>, 16, 0xbe84 ); // Read pattern check failure status

matchfailed = PhyRead( <phy>, 17 ) & 1; // Extract pattern check failure
status

}
```

Turning off the pattern checker:

```
PhyWrite( <phy>, 16, 0xbe80 );

reg18 = PhyRead( <phy>, 18 );

reg17 = PhyRead( <phy>, 17 );

reg17 = reg17 & ~4;

PhyWrite( <phy>, 18, reg18 );

PhyWrite( <phy>, 17, reg17 );

PhyWrite( <phy>, 16, 0x9e80 );


PhyWrite( <phy>, 31, 0 );
```

```
HALT_codeword_detected = !matchfailed;
```

3.19 Configuration

The VSC8574 can be configured by setting internal memory registers using the management interface.

4 Registers

This section provides information about how to configure the VSC8574 using its internal memory registers and the management interface.

The VSC8574 uses several different types of registers:

- IEEE Clause 22 device registers with addresses from 0 to 31
- Three pages of extended registers with addresses from 16E1–30E1, 16E2–30E2, and 16E3–30E3
- General-purpose registers with addresses from 0G to 30G
- IEEE Clause 45 devices registers accessible through the Clause 22 registers 13 and 14 to support IEEE 802.3az energy efficient Ethernet registers

The following illustration shows the relationship between the device registers and their address spaces.

Figure 61. Register Space Diagram

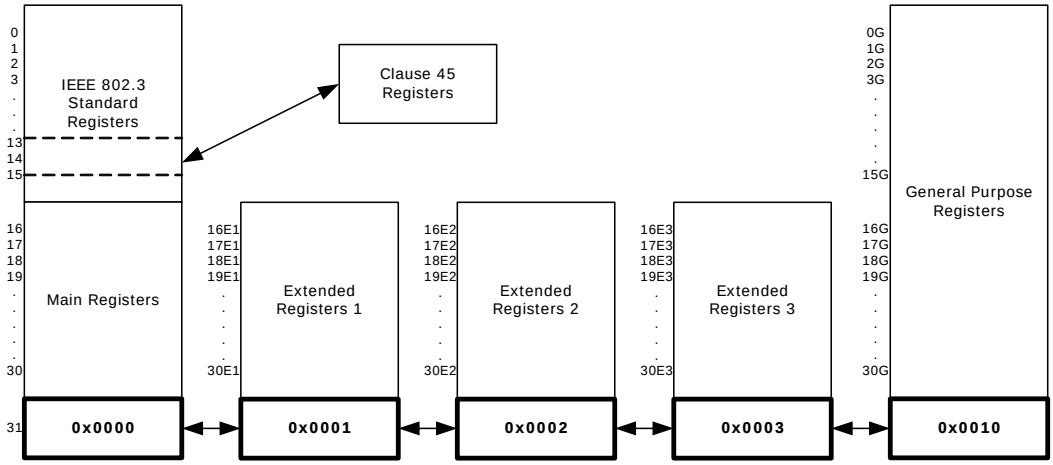

Reserved Registers For main registers 16–31, extended registers 16E1–30E1, 16E2–30E2, 16E3–30E3, and general purpose registers 0G–30G, any bits marked as Reserved should be processed as read-only and their states as undefined.

Reserved Bits In writing to registers with reserved bits, use a read-modify-then-write technique, where the entire register is read but only the intended bits to be changed are modified. Reserved bits cannot be changed and their read state cannot be considered static or unchanging.

4.1 IEEE 802.3 and Main Registers

In the VSC8574, the page space of the standard registers consists of the IEEE 802.3 standard registers and the Vitesse standard registers. The following table lists the names of the registers associated with the addresses as specified by IEEE 802.3.

Table 33. IEEE 802.3 Registers

Address	Name
0	Mode Control
1	Mode Status
2	PHY Identifier 1
3	PHY Identifier 2
4	Autonegotiation Advertisement
5	Autonegotiation Link Partner Ability
6	Autonegotiation Expansion
7	Autonegotiation Next-Page Transmit
8	Autonegotiation Link Partner Next-Page Receive
9	1000BASE-T Control
10	1000BASE-T Status
11–12	Reserved
13	Clause 45 Access Registers from IEEE 802.3 Table 22-6 and 22.24.3.11-12 and Annex 22D
14	Clause 45 Access Registers from IEEE 802.3 Table 22-6 and 22.24.3.11-12 and Annex 22D
15	1000BASE-T Status Extension 1

The following table lists the names of the registers in the main page space of the device. These registers are accessible only when register address 31 is set to 0x0000.

Table 34. Main Registers

Address	Name
16	100BASE-TX status extension
17	1000BASE-T status extension 2
18	Bypass control
19	Error Counter 1
20	Error Counter 2
21	Error Counter 3
22	Extended control and status
23	Extended PHY control 1
24	Extended PHY control 2
25	Interrupt mask
26	Interrupt status
27	Reserved
28	Auxiliary control and status
29	LED mode select

Table 34. Main Registers *(continued)*

Address	Name
30	LED behavior
31	Extended register page access

4.1.1 Mode Control

The device register at memory address 0 controls several aspects of VSC8574 functionality. The following table shows the available bit settings in this register and what they control.

Table 35. Mode Control, Address 0 (0x00)

Bit	Name	Access	Description	Default
15	Software reset	R/W	Self-clearing. Restores all serial management interface (SMI) registers to default state, except for sticky and super-sticky bits. 1: Reset asserted. 0: Reset de-asserted. Wait [X] after setting this bit to initiate another SMI register access.	0
14	Loopback	R/W	1: Loopback enabled. 0: Loopback disabled. When loop back is enabled, the device functions at the current speed setting and with the current duplex mode setting (bits 6, 8, and 13 of this register).	0
13	Forced speed selection LSB	R/W	Least significant bit. MSB is bit 6. 00: 10 Mbps. 01: 100 Mbps. 10: 1000 Mbps. 11: Reserved.	0
12	Autonegotiation enable	R/W	1: Autonegotiation enabled. 0: Autonegotiation disabled.	1
11	Power-down	R/W	1: Power-down enabled.	0
10	Isolate	R/W	1: Disable MAC interface outputs and ignore MAC interface inputs.	0
9	Restart autonegotiation	R/W	Self-clearing bit. 1: Restart autonegotiation on media interface.	0
8	Duplex	R/W	1: Full-duplex. 0: Half-duplex.	0
7	Collision test enable	R/W	1: Collision test enabled.	0
6	Forced speed selection MSB	R/W	Most significant bit. LSB is bit 13. 00: 10 Mbps. 01: 100 Mbps. 10: 1000 Mbps. 11: Reserved.	10

Table 35. Mode Control, Address 0 (0x00) *(continued)*

Bit	Name	Access	Description	Default
5	Unidirectional enable	R/W	When bit 0.12 = 1 or bit 0.8 = 0, this bit is ignored. When bit 0.12 = 0 and bit 0.8 = 1, the behavior is as follows: 1: Enable transmit from media independent interface regardless of whether the PHY has determined that a valid link has been established. 0: Enable transmit from media independent interface only when the PHY has determined that a valid link has been established **Note** This bit is only applicable in 100BASE-FX and 1000BASE-X fiber media modes for ports 8–11	0
4:0	Reserved		Reserved.	00000

4.1.2 Mode Status

The register at address 1 in the device main registers space allows you to read the currently enabled mode setting. The following table shows possible readouts of this register.

Table 36. Mode Status, Address 1 (0x01)

Bit	Name	Access	Description	Default
15	100BASE-T4 capability	RO	1: 100BASE-T4 capable.	0
14	100BASE-TX FDX capability	RO	1: 100BASE-TX FDX capable.	1
13	100BASE-TX HDX capability	RO	1: 100BASE-TX HDX capable.	1
12	10BASE-T FDX capability	RO	1: 10BASE-T FDX capable.	1
11	10BASE-T HDX capability	RO	1: 10BASE-T HDX capable.	1
10	100BASE-T2 FDX capability	RO	1: 100BASE-T2 FDX capable.	0
9	100BASE-T2 HDX capability	RO	1: 100BASE-T2 HDX capable.	0
8	Extended status enable	RO	1: Extended status information present in register 15.	1
7	Unidirectional ability	RO	1: PHY able to transmit from media independent interface regardless of whether the PHY has determined that a valid link has been established. 0: PHY able to transmit from media independent interface only when the PHY has determined that a valid link has been established. **Note** This bit is only applicable to 100BASE-FX and 1000BASE-X fiber media modes.	1
6	Preamble suppression capability	RO	1: MF preamble can be suppressed. 0: MF required.	1
5	Autonegotiation complete	RO	1: Autonegotiation complete.	0
4	Remote fault	RO	Latches high. 1: Far-end fault detected.	0

Table 36. Mode Status, Address 1 (0x01) *(continued)*

Bit	Name	Access	Description	Default
3	Autonegotiation capability	RO	1: Autonegotiation capable.	1
2	Link status	RO	Latches low. 1: Link is up.	0
1	Jabber detect	RO	Latches high. 1: Jabber condition detected.	0
0	Extended capability	RO	1: Extended register capable.	1

4.1.3 Device Identification

All 16 bits in both register 2 and register 3 in the VSC8574 are used to provide information associated with aspects of the device identification. The following tables list the expected readouts.

Table 37. Identifier 1, Address 2 (0x02)

Bit	Name	Access	Description	Default
15:0	Organizationally unique identifier (OUI)	RO	OUI most significant bits (3:18)	0×0007

Table 38. Identifier 2, Address 3 (0x03)

Bit	Name	Access	Description	Default
15:10	OUI	RO	OUI least significant bits (19:24)	0x0001
9:4	Vitesse model number	RO	VSC8574 (0x2e)	101110
3:0	Device revision number	RO		0000

4.1.4 Autonegotiation Advertisement

The bits in address 4 in the main registers space control the VSC8574 ability to notify other devices of the status of its autonegotiation feature. The following table shows the available settings and readouts.

Table 39. Device Autonegotiation Advertisement, Address 4 (0x04)

Bit	Name	Access	Description	Default
15	Next page transmission request	R/W	1: Request enabled	0
14	Reserved	RO	Reserved	0
13	Transmit remote fault	R/W	1: Enabled	0
12	Reserved	R/W	Reserved	0
11	Advertise asymmetric pause	R/W	1: Advertises asymmetric pause	0
10	Advertise symmetric pause	R/W	1: Advertises symmetric pause	0
9	Advertise100BASE-T4	R/W	1: Advertises 100BASE-T4	0
8	Advertise100BASE-TX FDX	R/W	1: Advertise 100BASE-TX FDX	1
7	Advertise100BASE-TX HDX	R/W	1: Advertises 100BASE-TX HDX	1
6	Advertise10BASE-T FDX	R/W	1: Advertises 10BASE-T FDX	1

Table 39. Device Autonegotiation Advertisement, Address 4 (0x04) *(continued)*

Bit	Name	Access	Description	Default
5	Advertise10BASE-T HDX	R/W	1: Advertises 10BASE-T HDX	1
4:0	Advertise selector	R/W		00001

4.1.5 Link Partner Autonegotiation Capability

The bits in main register 5 can be used to determine if the Cat5 link partner (LP) used with the VSC8574 is compatible with the autonegotiation functionality.

Table 40. Autonegotiation Link Partner Ability, Address 5 (0x05)

Bit	Name	Access	Description	Default
15	LP next page transmission request	RO	1: Requested	0
14	LP acknowledge	RO	1: Acknowledge	0
13	LP remote fault	RO	1: Remote fault	0
12	Reserved	RO	Reserved	0
11	LP advertise asymmetric pause	RO	1: Capable of asymmetric pause	0
10	LP advertise symmetric pause	RO	1: Capable of symmetric pause	0
9	LP advertise 100BASE-T4	RO	1: Capable of 100BASE-T4	0
8	LP advertise 100BASE-TX FDX	RO	1: Capable of 100BASE-TX FDX	0
7	LP advertise 100BASE-TX HDX	RO	1: Capable of 100BASE-TX HDX	0
6	LP advertise 10BASE-T FDX	RO	1: Capable of 10BASE-T FDX	0
5	LP advertise 10BASE-T HDX	RO	1: Capable of 10BASE-T HDX	0
4:0	LP advertise selector	RO		00000

4.1.6 Autonegotiation Expansion

The bits in main register 6 work together with those in register 5 to indicate the status of the LP autonegotiation functioning. The following table shows the available settings and readouts.

Table 41. Autonegotiation Expansion, Address 6 (0x06)

Bit	Name	Access	Description	Default
15:5	Reserved	RO	Reserved.	All zeros
4	Parallel detection fault	RO	This bit latches high. 1: Parallel detection fault.	0
3	LP next page capable	RO	1: LP is next page capable.	0
2	Local PHY next page capable	RO	1: Local PHY is next page capable.	1
1	Page received	RO	This bit latches low. 1: New page is received.	0
0	LP is autonegotiation capable	RO	1: LP is capable of autonegotiation.	0

4.1.7 Transmit Autonegotiation Next Page

The settings in register 7 in the main registers space provide information about the number of pages in an autonegotiation sequence. The following table shows the settings available.

Table 42. Autonegotiation Next Page Transmit, Address 7 (0x07)

Bit	Name	Access	Description	Default
15	Next page	R/W	1: More pages follow	0
14	Reserved	RO	Reserved	0
13	Message page	R/W	1: Message page 0: Unformatted page	1
12	Acknowledge 2	R/W	1: Complies with request 0: Cannot comply with request	0
11	Toggle	RO	1: Previous transmitted LCW = 0 0: Previous transmitted LCW = 1	0
10:0	Message/unformatted code	R/W		00000000001

4.1.8 Autonegotiation Link Partner Next Page Receive

The bits in register 8 of the main register space work together with register 7 to determine certain aspects of the LP autonegotiation. The following table shows the possible readouts.

Table 43. Autonegotiation LP Next Page Receive, Address 8 (0x08)

Bit	Name	Access	Description	Default
15	LP next page	RO	1: More pages follow	0
14	Acknowledge	RO	1: LP acknowledge	0
13	LP message page	RO	1: Message page 0: Unformatted page	0
12	LP acknowledge 2	RO	1: LP complies with request	0
11	LP toggle	RO	1: Previous transmitted LCW = 0 0: Previous transmitted LCW = 1	0
10:0	LP message/unformatted code	RO		All zeros

4.1.9 1000BASE-T Control

The VSC8574's 1000BASE-T functionality is controlled by the bits in register 9 of the main register space. The following table shows the settings and readouts available.

Table 44. 1000BASE-T Control, Address 9 (0x09)

Bit	Name	Access	Description	Default
15:13	Transmitter test mode	R/W	000: Normal 001: Mode 1: Transmit waveform test 010: Mode 2: Transmit jitter test as master 011: Mode 3: Transmit jitter test as slave 100: Mode 4: Transmitter distortion test 101–111: Reserved	000
12	Master/slave manual configuration	R/W	1: Master/slave manual configuration enabled	0
11	Master/slave value	R/W	This register is only valid when bit 9.12 is set to 1. 1: Configure PHY as master during negotiation 0: Configure PHY as slave during negotiation	0
10	Port type	R/W	1: Multi-port device 0: Single-port device	1
9	1000BASE-T FDX capability	R/W	1: PHY is 1000BASE-T FDX capable	1
8	1000BASE-T HDX capability	R/W	1: PHY is 1000BASE-T HDX capable	1
7:0	Reserved	R/W	Reserved	0x00

Note Transmitter test mode (bits 15:13) operates in the manner described in IEEE 802.3 section 40.6.1.1.2. When using any of the transmitter test modes, the automatic media sense feature must be disabled. For more information, see "Extended PHY Control Set 2," page 127.

4.1.10 1000BASE-T Status

The bits in register 10 of the main register space can be read to obtain the status of the 1000BASE-T communications enabled in the device. The following table shows the readouts.

Table 45. 1000BASE-T Status, Address 10 (0x0A)

Bit	Name	Access	Description	Default
15	Master/slave configuration fault	RO	This bit latches high. 1: Master/slave configuration fault detected 0: No master/slave configuration fault detected	0
14	Master/slave configuration resolution	RO	1: Local PHY configuration resolved to master 0: Local PHY configuration resolved to slave	1
13	Local receiver status	RO	1: Local receiver is operating normally	0
12	Remote receiver status	RO	1: Remote receiver OK	0

Table 45. 1000BASE-T Status, Address 10 (0x0A) *(continued)*

Bit	Name	Access	Description	Default
11	LP 1000BASE-T FDX capability	RO	1: LP 1000BASE-T FDX capable	0
10	LP 1000BASE-T HDX capability	RO	1: LP 1000BASE-T HDX capable	0
9:8	Reserved	RO	Reserved	00
7:0	Idle error count	RO	Self-clearing register	0x00

4.1.11 MMD Access Control Register

The bits in register 13 of the main register space are a window to the EEE registers as defined in IEEE 802.3az Clause 45.

Table 46. MMD EEE Access, Address 13 (0x0D)

Bit	Name	Access	Description
15:14	Function	R/W	00: Address 01: Data, no post increment 10: Data, post increment for read and write 11: Data, post increment for write only
13:5	Reserved	R/W	Reserved
4:0	DVAD	R/W	Device address as defined in IEEE 802.3az table 45–1

4.1.12 MMD Address or Data Register

The bits in register 14 of the main register space are a window to the EEE registers as defined in IEEE 802.3az Clause 45.

Table 47. MMD Address or Data Register, Address 14 (0x0E)

Bit	Name	Access	Description
15:0	Register Address/Data	R/W	When register 13.15:14 = 2'b00, address of register of the device that is specified by 13.4:0. Otherwise, the data to be written to or read from the register.

4.1.13 1000BASE-T Status Extension 1

Register 15 provides additional information about the operation of the device 1000BASE-T communications. The following table shows the readouts available.

Table 48. 1000BASE-T Status Extension 1, Address 15 (0x0F)

Bit	Name	Access	Description	Default
15	1000BASE-X FDX capability	RO	1: PHY is 1000BASE-X FDX capable	0
14	1000BASE-X HDX capability	RO	1: PHY is 1000BASE-X HDX capable	0
13	1000BASE-T FDX capability	RO	1: PHY is 1000BASE-T FDX capable	1
12	1000BASE-T HDX capability	RO	1: PHY is 1000BASE-T HDX capable	1

Table 48. 1000BASE-T Status Extension 1, Address 15 (0x0F) *(continued)*

Bit	Name	Access	Description	Default
11:0	Reserved	RO	Reserved	0x000

4.1.14 100BASE-TX Status Extension

Register 16 in the main registers page space of the VSC8574 provides additional information about the status of the device's 100BASE-TX operation.

Table 49. 100BASE-TX Status Extension, Address 16 (0x10)

Bit	Name	Access	Description	Default
15	100BASE-TX Descrambler	RO	1: Descrambler locked	0
14	100BASE-TX lock error	RO	Self-clearing bit. 1: Lock error detected	0
13	100BASE-TX disconnect state	RO	Self-clearing bit. 1: PHY 100BASE-TX link disconnect detected	0
12	100BASE-TX current link status	RO	1: PHY 100BASE-TX link active	0
11	100BASE-TX receive error	RO	Self-clearing bit. 1: Receive error detected	0
10	100BASE-TX transmit error	RO	Self-clearing bit. 1: Transmit error detected	0
9	100BASE-TX SSD error	RO	Self-clearing bit. 1: Start-of-stream delimiter error detected	0
8	100BASE-TX ESD error	RO	Self-clearing bit. 1: End-of-stream delimiter error detected	0
7:0	Reserved	RO	Reserved	

4.1.15 1000BASE-T Status Extension 2

The second status extension register is at address 17 in the device main registers space. It provides information about another set of parameters associated with 1000BASE-T communications. For information about the first status extension register, see Table 48, page 121.

Table 50. 1000BASE-T Status Extension 2, Address 17 (0x11)

Bit	Name	Access	Description	Default
15	1000BASE-T descrambler	RO	1: Descrambler locked.	0
14	1000BASE-T lock error	RO	Self-clearing bit. 1: Lock error detected	0
13	1000BASE-T disconnect state	RO	Self-clearing bit. 1: PHY 1000BASE-T link disconnect detected	0
12	1000BASE-T current link status	RO	1: PHY 1000BASE-T link active	0

Table 50. 1000BASE-T Status Extension 2, Address 17 (0x11) *(continued)*

Bit	Name	Access	Description	Default
11	1000BASE-T receive error	RO	Self-clearing bit. 1: Receive error detected	0
10	1000BASE-T transmit error	RO	Self-clearing bit. 1: Transmit error detected	0
9	1000BASE-T SSD error	RO	Self-clearing bit. 1: Start-of-stream delimiter error detected	0
8	1000BASE-T ESD error	RO	Self-clearing bit. 1: End-of-stream delimiter error detected	0
7	1000BASE-T carrier extension error	RO	Self-clearing bit. 1: Carrier extension error detected	0
6	Non-compliant BCM5400 detected	RO	1: Non-compliant BCM5400 link partner detected	0
5	MDI crossover error	RO	1: MDI crossover error was detected	0
4:0	Reserved	RO	Reserved	

4.1.16 Bypass Control

The bits in this register control aspects of functionality in effect when the device is disabled for the purpose of traffic bypass. The following table shows the settings available.

Table 51. Bypass Control, Address 18 (0x12)

Bit	Name	Access	Description	Default
15	Transmit disable	R/W	1: PHY transmitter disabled	0
14	4B5B encoder/decoder	R/W	1: Bypass 4B/5B encoder/decoder	0
13	Scrambler	R/W	1: Bypass scrambler	0
12	Descrambler	R/W	1: Bypass descrambler	0
11	PCS receive	R/W	1: Bypass PCS receiver	0
10	PCS transmit	R/W	1: Bypass PCS transmit	0
9	LFI timer	R/W	1: Bypass Link Fail Inhibit (LFI) timer	0
8	Reserved	RO	Reserved	
7	HP Auto-MDIX at forced 10/100	R/W	Sticky bit. 1: Disable HP Auto-MDIX at forced 10/100 speeds	1
6	Non-compliant BCM5400 detect disable	R/W	Sticky bit. 1: Disable non-compliant BCM5400 detection	0
5	Disable pair swap correction (HP Auto-MDIX when autonegotiation enabled)	R/W	Sticky bit. 1: Disable the automatic pair swap correction	0
4	Disable polarity correction	R/W	Sticky bit. 1: Disable polarity inversion correction on each subchannel	0

Table 51. Bypass Control, Address 18 (0x12) *(continued)*

Bit	Name	Access	Description	Default
3	Parallel detect control	R/W	Sticky bit. 1: Do not ignore advertised ability 0: Ignore advertised ability	1
2	Pulse shaping filter	R/W	1: Disable pulse shaping filter	0
1	Disable automatic 1000BASE-T next page exchange	R/W	Sticky bit. 1: Disable automatic 1000BASE T next page exchanges	0
0	Reserved	RO	Reserved	

Note If bit 18.1 is set to 1 in this register, automatic exchange of next pages is disabled, and control is returned to the user through the SMI after the base page is exchanged. The user then must send the correct sequence of next pages to the link partner, determine the common capabilities, and force the device into the correct configuration following the successful exchange of pages.

4.1.17 Error Counter 1

The bits in register 19 provide an error counter. The following table shows the settings available.

Table 52. Extended Control and Status, Address 19 (0x13)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved.	
7:0	100/1000BASE-TX receive error counter	RO	8-bit counter that saturates when it reaches 255. These bits are self-clearing when read.	0x00

4.1.18 Error Counter 2

The bits in register 20 provide an error counter. The following table shows the settings available.

Table 53. Extended Control and Status, Address 20 (0x14)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved.	
7:0	100/1000BASE-TX false carrier counter	RO	8-bit counter that saturates when it reaches 255. These bits are self-clearing when read.	0x00

4.1.19 Error Counter 3

The bits in register 21 provide an error counter. The following table shows the settings available.

Table 54. Extended Control and Status, Address 21 (0x15)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved.	
7:0	Copper media link disconnect counter	RO	8-bit counter that saturates when it reaches 255. These bits are self-clearing when read.	0x00

4.1.20 Extended Control and Status

The bits in register 22 provide additional device control and readouts. The following table shows the settings available.

Table 55. Extended Control and Status, Address 22 (0x16)

Bit	Name	Access	Description	Default
15	Force 10BASE-T link high	R/W	Sticky bit. 1: Bypass link integrity test 0: Enable link integrity test	0
14	Jabber detect disable	R/W	Sticky bit. 1: Disable jabber detect	0
13	Disable 10BASE-T echo	R/W	Sticky bit. 1: Disable 10BASE-T echo	1
12	Disable SQE mode	R/W	Sticky bit. 1: Disable SQE mode	1
11:10	10BASE-T squelch control	R/W	Sticky bit. 00: Normal squelch 01: Low squelch 10: High squelch 11: Reserved	00
9	Sticky reset enable	R/W	Super-sticky bit. 1: Enabled	1
8	EOF Error	RO	This bit is self-clearing. 1: EOF error detected	0
7	10BASE-T disconnect state	RO	This bit is self-clearing. 1: 10BASE-T link disconnect detected	0
6	10BASE-T link status	RO	1: 10BASE-T link active	0
5:1	Reserved	RO	Reserved	
0	SMI broadcast write	R/W	Sticky bit. 1: Enabled	0

The following information applies to the extended control and status bits:

- When bit 22.15 is set, the link integrity state machine is bypassed and the PHY is forced into a link pass status.
- When bits 22.11:10 are set to 00, the squelch threshold levels are based on the IEEE standard for 10BASE-T. When set to 01, the squelch level is decreased, which

can improve the bit error rate performance on long loops. When set to 10, the squelch level is increased and can improve the bit error rate in high-noise environments.

- When bit 22.9 is set, all sticky register bits retain their values during a software reset. Clearing this bit causes all sticky register bits to change to their default values upon software reset. Super-sticky bits retain their values upon software reset regardless of the setting of bit 22.9.
- When bit 22.0 is set, if a write to any PHY register (registers 0–31, including extended registers), the same write is broadcast to all PHYs. For example, if bit 22.0 is set to 1 and a write to PHY0 is executed (register 0 is set to 0x1040), all PHYs' register 0s are set to 0x1040. Disabling this bit restores normal PHY write operation. Reads are still possible when this bit is set, but the value that is read corresponds only to the particular PHY being addressed.

4.2 Extended PHY Control Set 1

The following table shows the settings available.

Table 56. Extended PHY Control 1, Address 23 (0x17)

Bit	Name	Access	Description	Default
15:13	Reserved	RO	Reserved.	
12	MAC interface mode	R/W	Super-sticky bit. 0: SGMII 1: 1000BASE-X. **Note** Register 19G.15:14 must be = 01 for this selection to be valid.	0
11	AMS preference	R/W	Super-sticky bit. 1: Cat5 copper preferred. 0: SerDes fiber/SFP preferred. **Note** Register 19G.15:14 must be = 10 for this selection to be valid.	0
10:8	Media operating mode	R/W	Super-sticky bits. 000: Cat5 copper only. 001: SerDes fiber/SFP protocol transfer mode only. 010: 1000BASE-X fiber/SFP media only with autonegotiation performed by the PHY. 011: 100BASE-FX fiber/SFP on the fiber media pins only. 101: Automatic media sense (AMS) with Cat5 media or SerDes fiber/SFP protocol transfer mode. 110: AMS with Cat5 media or 1000BASE-X fiber/SFP media with autonegotiation performed by PHY. 111: AMS with Cat5 media or 100BASE-FX fiber/SFP media. 100: AMS. **Note** Register 19G.15:14 must be = 10 for the any of the fiber media selections to be valid.	000

Table 56. Extended PHY Control 1, Address 23 (0x17) *(continued)*

Bit	Name	Access	Description	Default
7:6	Force AMS override	R/W	00: Normal AMS selection 01: Force AMS to select SerDes media only 10: Force AMS to select copper media only 11: Reserved **Note** Register 19G.15:14 must be = 10 for this selection to be valid.	00
5:4	Reserved	RO	Reserved.	
3	Far-end loopback mode	R/W	1: Enabled.	0
2:0	Reserved	RO	Reserved.	

Note After configuring bits 13:8 of the extended PHY control register set 1, a software reset (register 0, bit 15) must be written to change the device operating mode. On read, these bits only indicate the actual operating mode and not the pending operating mode setting before a software reset has taken place.

4.2.1 Extended PHY Control Set 2

The second set of extended controls is located in register 24 in the main register space for the device. The following table shows the settings and readouts available.

Table 57. Extended PHY Control 2, Address 24 (0x18)

Bit	Name	Access	Description	Default
15:13	100BASE-TX edge rate control	R/W	Sticky bit. 011: +5 edge rate (slowest) 010: +4 edge rate 001: +3 edge rate 000: +2 edge rate 111: +1 edge rate 110: Default edge rate 101: –1 edge rate 100: –2 edge rate (fastest)	110
12	PICMG 2.16 reduced power mode	R/W	Sticky bit. 1: Enabled	0
11:6	Reserved	RO	Reserved	
5:4	Jumbo packet mode	R/W	Sticky bit. 00: Normal IEEE 1.5 kB packet length 01: 9 kB jumbo packet length (12 kB with 60 ppm or better reference clock) 10: 12 kB jumbo packet length (16 kB with 70 ppm or better reference clock) 11: Reserved	00
3:1	Reserved	RO	Reserved	
0	1000BASE-T connector loopback	R/W	1: Enabled	0

Note When bits 5:4 are set to jumbo packet mode, the default maximum packet values are based on 100 ppm driven reference clock to the device. Controlling the ppm offset between the MAC and the PHY as specified in the bit description results in a higher jumbo packet length.

4.2.2 Interrupt Mask

These bits control the device interrupt mask. The following table shows the settings available.

Table 58. Interrupt Mask, Address 25 (0x19)

Bit	Name	Access	Description	Default
15	MDINT interrupt status enable	R/W	Sticky bit. 1: Enabled.	0
14	Speed state change mask	R/W	Sticky bit. 1: Enabled.	0
13	Link state change mask	R/W	Sticky bit. 1: Enabled.	0
12	FDX state change mask	R/W	Sticky bit. 1: Enabled.	0
11	Autonegotiation error mask	R/W	Sticky bit. 1: Enabled.	0
10	Autonegotiation complete mask	R/W	Sticky bit. 1: Enabled.	0
9	Inline powered device (PoE) detect mask	R/W	Sticky bit. 1: Enabled.	0
8	Symbol error interrupt mask	R/W	Sticky bit. 1: Enabled.	0
7	Fast link failure interrupt mask	R/W	Sticky bit. 1: Enabled.	0
6	TX FIFO over/underflow interrupt mask	R/W	Sticky bit. 1: Enabled.	0
5	RX FIFO over/underflow interrupt mask	R/W	Sticky bit. 1: Enabled.	0
4	AMS media changed mask	R/W	Sticky bit. 1: Enabled.	0
3	False carrier interrupt mask	R/W	Sticky bit. 1: Enabled.	0
2	Link speed downshift detect mask	R/W	Sticky bit. 1: Enabled.	0
1	Master/Slave resolution error mask	R/W	Sticky bit. 1: Enabled.	0
0	RX_ER interrupt mask	R/W	Sticky bit. 1: Enabled.	0

Note When bit 25.15 is set, the MDINT pin is enabled. When enabled, the state of this pin reflects the state of bit 26.15. Clearing this bit only inhibits the MDINT pin from being asserted. Also, before enabling this bit, read register 26 to clear any previously inactive interrupts pending that will cause bit 25.15 to be set.

4.2.3 Interrupt Status

The status of interrupts already written to the device is available for reading from register 26 in the main registers space. The following table shows the expected readouts.

Table 59. Interrupt Status, Address 26 (0x1A)

Bit	Name	Access	Description	Default
15	Interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
14	Speed state change status	RO	Self-clearing bit. 1: Interrupt pending.	0
13	Link state change status	RO	Self-clearing bit. 1: Interrupt pending.	0
12	FDX state change status	RO	Self-clearing bit. 1: Interrupt pending.	0
11	Autonegotiation error status	RO	Self-clearing bit. 1: Interrupt pending.	0
10	Autonegotiation complete status	RO	Self-clearing bit. 1: Interrupt pending.	0
9	Inline powered device detect status	RO	Self-clearing bit. 1: Interrupt pending.	0

Table 59. Interrupt Status, Address 26 (0x1A) *(continued)*

Bit	Name	Access	Description	Default
8	Symbol error status	RO	Self-clearing bit. 1: Interrupt pending.	0
7	Fast link failure detect status	RO	Self-clearing bit. 1: Interrupt pending.	0
6	TX FIFO over/underflow detect status	RO	Self-clearing bit. 1: Interrupt pending.	0
5	RX FIFO over/underflow detect status	RO	Self-clearing bit. 1: Interrupt pending.	0
4	AMS media changed mask	RO	Self-clearing bit. 1: Interrupt pending.	0
3	False carrier interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
2	Link speed downshift detect status	RO	Self-clearing bit. 1: Interrupt pending.	0
1	Master/Slave resolution error status	RO	Self-clearing bit. 1: Interrupt pending.	0
0	RX_ER interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0

The following information applies to the interrupt status bits:

- All set bits in this register are cleared after being read (self-clearing). If bit 26.15 is set, the cause of the interrupt can be read by reading bits 26.14:0.
- For bits 26.14 and 26.12, bit 0.12 must be set for this interrupt to assert.
- For bit 26.2, bits 4.8:5 must be set for this interrupt to assert.
- For bit 26.0, this interrupt will not occur when RX_ER is used for carrier-extension decoding of a link partner's data transmission.

4.2.4 Device Auxiliary Control and Status

Register 28 provides control and status information for several device functions not controlled or monitored by other device registers. The following table shows the settings available and the expected readouts.

Table 60. Auxiliary Control and Status, Address 28 (0x1C)

Bit	Name	Access	Description	Default
15	Autonegotiation complete	RO	Duplicate of bit 1.5	0
14	Autonegotiation disabled	RO	Inverted duplicate of bit 0.12	0
13	HP Auto-MDIX crossover indication	RO	1: HP Auto-MDIX crossover performed internally	0
12	CD pair swap	RO	1: CD pairs are swapped	0
11	A polarity inversion	RO	1: Polarity swap on pair A	0
10	B polarity inversion	RO	1: Polarity swap on pair B	0
9	C polarity inversion	RO	1: Polarity swap on pair C	0
8	D polarity inversion	RO	1: Polarity swap on pair D	0

Table 60. Auxiliary Control and Status, Address 28 (0x1C) *(continued)*

Bit	Name	Access	Description	Default
7	ActiPHY link status time-out control [1]	R/W	Sticky bit. Bits 7 and 2 are part of the ActiPHY Link Status time-out control. Bit 7 is the MSB. 00: 1 second 01: 2 seconds 10: 3 seconds 11: 4 seconds	0
6	ActiPHY mode enable	R/W	Sticky bit. 1: Enabled	0
5	FDX status	RO	1: Full-duplex 0: Half-duplex	00
4:3	Speed status	RO	00: Speed is 10BASE-T 01: Speed is 100BASE-TX or 100BASE-FX 10: Speed is 1000BASE-T or 1000BASE-X 11: Reserved	0
2	ActiPHY link status time-out control [0]	R/W	Sticky bit. Bits 7 and 2 are part of the ActiPHY Link Status time-out control. Bit 7 is the MSB. 00: 1 second 01: 2 seconds 10: 3 seconds 11: 4 seconds	1
1:0	Reserved	RO	Reserved	

4.2.5 LED Mode Select

The device LED outputs are controlled using the bits in register 29 of the main register space. The following table shows the information needed to access the functionality of each of the outputs. For more information about LED modes, see Table 26, page 93. For information about enabling the extended LED mode bits in Register 19E1 bits 13 to 12, see Table 27, page 95.

Table 61. LED Mode Select, Address 29 (0x1D)

Bit	Name	Access	Description	Default
15:12	LED3 mode select	R/W	Sticky bit. Select from LED modes 0–15.	1000
11:8	LED2 mode select	R/W	Sticky bit. Select from LED modes 0–15.	0000
7:4	LED1 mode select	R/W	Sticky bit. Select from LED modes 0–15.	0010
3:0	LED0 mode select	R/W	Sticky bit. Select from LED modes 0–15.	0001

4.2.6 LED Behavior

The bits in register 30 control and enable you to read the status of the pulse or blink rate of the device LEDs. The following table shows the settings you can write to the register or read from the register.

Table 62. LED Behavior, Address 30 (0x1E)

Bit	Name	Access	Description	Default
15:13	Reserved	RO	Reserved	
12	LED pulsing enable	R/W	Sticky bit. 0: Normal operation 1: LEDs pulse with a 5 kHz, programmable duty cycle when active	0
11:10	LED blink/pulse-stretch rate	R/W	Sticky bit. 00: 2.5 Hz blink rate/400 ms pulse-stretch 01: 5 Hz blink rate/200 ms pulse-stretch 10: 10 Hz blink rate/100 ms pulse-stretch 11: 20 Hz blink rate/50 ms pulse-stretch The blink rate selection for PHY0 globally sets the rate used for all LED pins on all PHY ports.	01
9:7	Reserved	RO	Reserved	
6	LED1 pulse-stretch/ blink select	R/W	Sticky bit. 1: Pulse-stretch 0: Blink	0
5	LED0 pulse-stretch/ blink select	R/W	Sticky bit. 1: Pulse-stretch 0: Blink	0
4:2	Reserved	RO	Reserved	
1	LED1 combine feature disable	R/W	Sticky bit. 0: Combine enabled (link/activity, duplex/ collision) 1: Disable combination (link only, duplex only)	0
0	LED0 combine feature disable	R/W	Sticky bit. 0: Combine enabled (link/activity, duplex/ collision) 1: Disable combination (link only, duplex only)	0

Note Bits 30.11:10 are active only in port 0 and affect the behavior of LEDs for all the ports.

4.2.7 Extended Page Access

To provide functionality beyond the IEEE 802.3-specified registers and main device registers, the VSC8574 includes an extended set of registers that provide an additional 15 register spaces.

The register at address 31 controls the access to the extended registers for the VSC8574. Accessing the GPIO page register space is similar to accessing the extended page registers. The following table shows the settings available.

Table 63. Extended/GPIO Register Page Access, Address 31 (0x1F)

Bit	Name	Access	Description	Default
15:0	Extended/GPIO page register access	R/W	0x0000: Register 16–30 accesses main register space. Writing 0x0000 to register 31 restores the main register access. 0x0001: Registers 16–30 access extended register space 1 0x0002: Registers 16–30 access extended register space 2 0x0003: Registers 16–30 access extended register space 3 0x0010: Registers 0–30 access GPIO register space	0x0000

4.3 Extended Page 1 Registers

To access the extended page 1 registers (16E1–30E1), enable extended register access by writing 0x0001 to register 31. Writing 0x0000 to register 31 restores the main register access.

When extended page 1 register access is enabled, reads and writes to registers 16–30 affect the extended registers 16E1–30E1 instead of those same registers in the IEEE-specified register space. Registers 0–15 are not affected by the state of the extended page register access.

Table 64. Extended Registers Page 1 Space

Address	Name
16E1	SerDes Media Control
17E1	Reserved
18E1	Cu Media CRC good counter
19E1	SerDes loopback and SIGDET control
20E1	Extended PHY control 3 (ActiPHY)
21E1–22E1	Reserved
23E1	Extended PHY control 4 (PoE and CRC error counter)
24E1	VeriPHY 1
25E1	VeriPHY 2
26E1	VeriPHY 3
27E1–28E1	Reserved
29E1	Ethernet packet generator (EPG) 1
30E1	EPG 2

4.3.1 SerDes Media Control

Register 16E1 controls some functions of the SerDes media interface on ports 0–3. These settings are only valid for those ports. The following table shows the setting available in this register.

Table 65. SerDes Media Control, Address 16E1 (0x10)

Bit	Name	Access	Description	Default
15:14	Transmit remote fault	R/W	Remote fault indication sent to link partner (LP)	00
13:12	Link partner (LP) remote fault	RO	Remote fault bits sent by LP during autonegotiation	00
11:10	Reserved	RO	Reserved	
9	Allow 1000BASE-X link-up	R/W	Sticky bit. 1: Allow 1000BASE-X fiber media link-up capability 0: Suppress 1000BASE-X fiber media link-up capability	1
8	Allow 100BASE-FX link-up	R/W	Sticky bit. 1: Allow 100BASE-FX fiber media link-up capability 0: Suppress 100BASE-FX fiber media link-up capability	1
7	Reserved	RO	Reserved	
6	Far end fault detected in 100BASE-FX	RO	Self-clearing bit. 1: Far end fault in 100BASE-FX detected	0
5:0	Reserved	RO	Reserved	

4.3.2 Cu Media CRC Good Counter

Register 18E1 makes it possible to read the contents of the CRC good counter for packets that are received on the Cu media interface; the number of CRC routines that have executed successfully. The following table shows the expected readouts.

Table 66. Cu Media CRC Good Counter, Address 18E1 (0x12)

Bit	Name	Access	Description	Default
15	Packet since last read	RO	Self-clearing bit. 1: Packet received since last read.	0
14	Reserved	RO	Reserved.	
13:0	Cu Media CRC good counter contents	RO	Self-clearing bit. Counter containing the number of packets received on the Cu media interface with valid CRCs. This counter does not saturate and will roll over.	0x000

4.3.3 Extended Mode Control

Register 19E1 controls the extended LED and other chip modes. The following table shows the settings available.

Table 67. Extended Mode Control, Address 19E1 (0x13)

Bit	Name	Access	Description	Default
15	LED3 Extended Mode	R/W	1: See "Extended LED Modes," page 94	0
14	LED2 Extended Mode	R/W	1: See "Extended LED Modes," page 94	0
13	LED1 Extended Mode	R/W	1: See "Extended LED Modes," page 94	0
12	LED0 Extended Mode	R/W	1: See "Extended LED Modes," page 94	0
11	LED Reset Blink Suppress	R/W	1: Blink LEDs after COMA_MODE is de-asserted 0: Suppress LED blink after COMA_MODE is de-asserted	0
10:4	Reserved	RO	Reserved	0
3:2	Force MDI crossover	R/W	00: Normal HP Auto-MDIX operation 01: Reserved 10: Copper media forced to MDI 11: Copper media forced MDI-X	00
1	Reserved	RO	Reserved	
0	GPIO[3:0]/SIGDET[3:0] pin polarity	R/W	1: Active low 0: Active high	0

4.3.4 ActiPHY Control

Register 20E1 controls the device ActiPHY sleep timer, its wake-up timer, and its link speed downshifting feature. The following table shows the settings available.

Table 68. Extended PHY Control 3, Address 20E1 (0x14)

Bit	Name	Access	Description	Default
15	Disable carrier extension	R/W	1: Disable carrier extension in SGMII-1000BASE-T copper links	0
14:13	ActiPHY sleep timer	R/W	Sticky bit. 00: 1 second 01: 2 seconds 10: 3 seconds 11: 4 seconds	01
12:11	ActiPHY wake-up timer	R/W	Sticky bit. 00: 160 ms 01: 400 ms 10: 800 ms 11: 2 seconds	11
10	Reserved	RO	Reserved	
9	PHY address reversal	R/W	1: Enabled	Address
8	Reserved	RO	Valid only on PHY0	

Table 68. Extended PHY Control 3, Address 20E1 (0x14) *(continued)*

Bit	Name	Access	Description	Default
7:6	Media mode status	RO	00: No media selected 01: Copper media selected 10: SerDes media selected 11: Reserved	00
5	Enable 10BASE-T no preamble mode	R/W	Sticky bit. 1: 10BASE-T will assert RX_DV indication when data is presented to the receiver even without a preamble preceding it	0
4	Enable link speed autodownshift feature	R/W	Sticky bit. 1: Enable auto link speed downshift from 1000BASE-T	0
3:2	Link speed auto downshift control	R/W	Sticky bit. 00: Downshift after 2 failed 1000BASE-T autonegotiation attempts 01: Downshift after 3 failed 1000BASE-T autonegotiation attempts 10: Downshift after 4 failed 1000BASE-T autonegotiation attempts 11: Downshift after 5 failed 1000BASE-T autonegotiation attempts	01
1	Link speed auto downshift status	RO	0: No downshift 1: Downshift is required or has occurred	0
0	Reserved	RO	Reserved	

4.3.5 PoE and Miscellaneous Functionality

The register at address 23E1 controls various aspects of inline powering and the CRC error counter in the VSC8574.

Table 69. Extended PHY Control 4, Address 23E1 (0x17)

Bit	Name	Access	Description	Default
15:11	PHY address	RO	PHY address; latched on reset	
10	Inline powered device detection	R/W	Sticky bit. 1: Enabled	0
9:8	Inline powered device detection status	RO	Only valid when bit 10 is set. 00: Searching for devices 01: Device found; requires inline power 10: Device found; does not require inline power 11: Reserved	00
7:0	Cu Media CRC error counter	RO	Self-clearing bit	

RC error counter for packets received on the Cu media interface. The value saturates at 0xFF and subsequently clears when read and restarts count.0x00

4.3.6 VeriPHY Control 1

Register 24E1 in the extended register space provides control over the device VeriPHY diagnostics features. There are three separate VeriPHY control registers. The following table shows the settings available and describes the expected readouts.

Table 70. VeriPHY Control Register 1, Address 24E1 (0x18)

Bit	Name	Access	Description	Default
15	VeriPHY trigger	R/W	Self-clearing bit. 1: Triggers the VeriPHY algorithm and clears when VeriPHY has completed. Settings in registers 24E–26E become valid after this bit clears.	0
14	VeriPHY valid	RO	1: VeriPHY results in registers 24E–26E are valid.	0
13:8	Pair A (1, 2) distance	RO	Loop length or distance to anomaly for pair A (1, 2).	0x00
7:6	Reserved	RO	Reserved.	
5:0	Pair B (3, 6) distance	RO	Loop length or distance to anomaly for pair B (3, 6).	0x00

Note The resolution of the 6-bit length field is 3 meters.

4.3.7 VeriPHY Control 2

The register at address 25E1 consists of the second of the three device registers that provide control over VeriPHY diagnostics features. The following table shows the expected readouts.

Table 71. VeriPHY Control Register 2, Address 25E1 (0x19)

Bit	Name	Access	Description	Default
15:14	Reserved	RO	Reserved	
13:8	Pair C (4, 5) distance	RO	Loop length or distance to anomaly for pair C (4, 5)	0x00
7:6	Reserved	RO	Reserved	
5:0	Pair D (7, 8) distance	RO	Loop length or distance to anomaly for pair D (7, 8)	0x00

Note The resolution of the 6-bit length field is 3 meters.

4.3.8 VeriPHY Control 3

The register at address 26E1 consists of the third of the three device registers that provide control over VeriPHY diagnostics features. Specifically, this register provides

information about the termination status (fault condition) for all four link partner pairs. The following table shows the expected readouts.

Table 72. VeriPHY Control Register 3, Address 26E1 (0x1A)

Bit	Name	Access	Description	Default
15:12	Pair A (1, 2) termination status	RO	Termination fault for pair A (1, 2)	0x00
11:8	Pair B (3, 6) termination status	RO	Termination fault for pair B (3, 4)	0x00
7:4	Pair C (4, 5) termination status	RO	Termination fault for pair C (4, 5)	0x00
3:0	Pair D (7, 8) termination status	RO	Termination fault for pair D (7, 8)	0x00

The following table shows the meanings for the various fault codes.

Table 73. VeriPHY Control Register 3 Fault Codes

Code	Denotes
0000	Correctly terminated pair
0001	Open pair
0010	Shorted pair
0100	Abnormal termination
1000	Cross-pair short to pair A
1001	Cross-pair short to pair B
1010	Cross-pair short to pair C
1011	Cross-pair short to pair D
1100	Abnormal cross-pair coupling with pair A
1101	Abnormal cross-pair coupling with pair B
1110	Abnormal cross-pair coupling with pair C
1111	Abnormal cross-pair coupling with pair D

4.3.9 Ethernet Packet Generator Control 1

The EPG control register provides access to and control of various aspects of the EPG testing feature. There are two separate EPG control registers. The following table shows the settings available in the first register.

Table 74. EPG Control Register 1, Address 29E1 (0x1D)

Bit	Name	Access	Description	Default
15	EPG enable	R/W	1: Enable EPG	0
14	EPG run or stop	R/W	1: Run EPG	0
13	Transmission duration	R/W	1: Continuous (sends in 10,000-packet increments) 0: Send 30,000,000 packets and stop	0
12:11	Packet length	R/W	00: 125 bytes 01: 64 bytes 10: 1518 bytes 11: 10,000 bytes (jumbo packet)	0
10	Interpacket gap	R/W	1: 8,192 ns 0: 96 ns	0

Table 74. EPG Control Register 1, Address 29E1 (0x1D) *(continued)*

Bit	Name	Access	Description	Default
9:6	Destination address	R/W	Lowest nibble of the 6-byte destination address	0001
5:2	Source address	R/W	Lowest nibble of the 6-byte destination address	0000
1	Payload type	R/W	1: Randomly generated payload pattern 0: Fixed based on payload pattern	0
0	Bad frame check sequence (FCS) generation	R/W	1: Generate packets with bad FCS 0: Generate packets with good FCS	0

The following information applies to the EPG control number 1:

- Do not run the EPG when the VSC8574 is connected to a live network.
- bit 29E1.13 (continuous EPG mode control): When enabled, this mode causes the device to send continuous packets. When disabled, the device continues to send packets only until it reaches the next 10,000-packet increment mark. It then ceases to send packets.
- The 6-byte destination address in bits 9:6 is assigned one of 16 addresses in the range of 0xFF FF FF FF FF F0 through 0xFF FF FF FF FF FF.
- The 6-byte source address in bits 5:2 is assigned one of 16 addresses in the range of 0xFF FF FF FF FF F0 through 0xFF FF FF FF FF FF.
- If any of bits 13:0 are changed while the EPG is running (bit 14 is set to 1), bit 14 must be cleared and then set back to 1 for the change to take effect and to restart the EPG.

4.3.10 Ethernet Packet Generator Control 2

Register 30E1 consists of the second set of bits that provide access to and control over the various aspects of the EPG testing feature. The following table shows the settings available.

Table 75. EPG Control Register 2, Address 30E1 (0x1E)

Bit	Name	Access	Description	Default
15:0	EPG packet payload	R/W	Data pattern repeated in the payload of packets generated by the EPG	0x00

Note If any of bits 15:0 in this register are changed while the EPG is running (bit 14 of register 29E1 is set to 1), that bit (29E1.14) must first be cleared and then set back to 1 for the change to take effect and to restart the EPG.

4.4 Extended Page 2 Registers

To access the extended page 2 registers (16E2–30E2), enable extended register access by writing 0x0002 to register 31. For more information, see Table 63, page 132.

When extended page 2 register access is enabled, reads and writes to registers 16–30 affect the extended registers 16E2–30E2 instead of those same registers in the IEEE-specified register space. Registers 0–15 are not affected by the state of the extended page register access.

Writing 0x0000 to register 31 restores the main register access.

The following table lists the addresses and register names in the extended register page 2 space. These registers are accessible only when the device register 31 is set to 0x0002.

Table 76. Extended Registers Page 2 Space

Address	Name
16E2	Cu PMD Transmit Control
17E2	EEE Control
18E2–30E2	Reserved

4.4.1 Cu PMD Transmit Control

The register at address 16E2 consists of the bits that provide control over the amplitude settings for the transmit side Cu PMD interface. These bits provide the ability to make small adjustments in the signal amplitude to compensate for minor variations in the magnetics from different vendors. Extreme caution must be exercised when changing these settings from the default values as they have a direct impact on the signal quality. Changing these settings also affects the linearity and harmonic distortion of the transmitted signals. Please contact the Vitesse Applications Support team for further help with changing these values.

Table 77. Cu PMD Transmit Control, Address 16E2 (0x10)

Bit	Name	Access	Description	Default
15:12	1000BASE-T signal amplitude trim	R/W	Change 1000BASE-T signal amplitude	0000
11:8	100BASE-TX signal amplitude trim	R/W	Change 100BASE-TX signal amplitude	0000
7:4	10BASE-T signal amplitude trim	R/W	Change 10BASE-T signal amplitude	0101
3:0	10BASE-Te signal amplitude trim	R/W	Change 10BASE-Te signal amplitude	0000

4.4.2 EEE Control

The register at address 17E2 consists of the bits that provide additional control over the chip behavior in energy efficient Ethernet (IEEE 802.3az) mode for debug and to allow interoperation with legacy MACs that do not support IEEE 802.3az.

Table 78. EEE Control, Address 17E2 (0x11)

Bit	Name	Access	Description	Default
15	Enable 10BASE-Te	R/W	Enable energy efficient (IEEE 802.3az) 10BASE-Te operating mode.	0
14	Enable LED in fiber unidirectional mode	R/W	1: Enable LED functions in fiber unidirectional mode.	0
13:10	Invert LED polarity	R/W	Invert polarity of LED[3:0]_[3:0] signals. Default is to drive an active low signal on the LED pins. This also applies to enhanced serial LED mode. For more information, see "Enhanced Serial LED Mode," page 97.	0000
9:6	Reserved	R/O	Reserved.	
5	Enable 1000BASE-T force mode	R/W	1: Enable 1000BASE-T force mode to allow PHY to link up in 1000BASE-T mode without forcing master/slave when register 0, bits 6 and 13 are set to 2'b10.	0
4	Force transmit LPI	R/W	1: Enable the EPG to transmit LPI on the MDI instead of normal idles when receiving normal idles from the MAC. 0: Transmit idles being received from the MAC.	0
3	Inhibit 100BASE-TX transmit EEE LPI	R/W	1: Disable transmission of EEE LPI on transmit path MDI in 100BASE-TX mode when receiving LPI from MAC.	0
2	Inhibit 100BASE-TX receive EEE LPI	R/W	1: Disable transmission of EEE LPI on receive path MAC interface in 100BASE-TX mode when receiving LPI from the MDI.	0
1	Inhibit 1000BASE-T transmit EEE LPI	R/W	1: Disable transmission of EEE LPI on transmit path MDI in 1000BASE-T mode when receiving LPI from MAC.	0
0	Inhibit 1000BASE-T receive EEE LPI	R/W	1: Disable transmission of EEE LPI on receive path MAC interface in 1000BASE-T mode when receiving LPI from the MDI.	0

4.5 Extended Page 3 Registers

To access the extended page 3 registers (16E3–30E3), enable extended register access by writing 0x0003 to register 31. For more information, see Table 63, page 132.

When extended page 3 register access is enabled, reads and writes to registers 16–30 affect the extended registers 16E3–30E3 instead of those same registers in the IEEE-specified register space. Registers 0–15 are not affected by the state of the extended page register access.

Writing 0x0000 to register 31 restores the main register access.

The following table lists the addresses and register names in the extended register page 3 space. These registers are accessible only when the device register 31 is set to 0x0003.

Table 79. Extended Registers Page 3 Space

Address	Name
16E3	MAC SerDes PCS Control
17E3	MAC SerDes PCS Status
18E3	MAC SerDes Clause 37 Advertised Ability
19E3	MAC SerDes Clause 37 Link Partner Ability
20E3	MAC SerDes Status
21E3	MAC SerDes Transmit Good Packet Counter
22E3	MAC SerDes Transmit CRC Error Counter
23E3	Media SerDes PCS Control
24E3	Media SerDes PCS Status
25E3	Media SerDes Clause 37 Advertised Ability
26E3	Media SerDes Clause 37 Link Partner Ability
27E3	Media SerDes status
28E3	Fiber Media CRC Good Counter
29E3	Fiber Media CRC Error Counter
30E3	Reserved

4.5.1 MAC SerDes PCS Control

The register at address 16E3 consists of the bits that provide access to and control over MAC SerDes PCS block. The following table shows the settings available.

Table 80. MAC SerDes PCS Control, Address 16E3 (0x10)

Bit	Name	Access	Description	Default
15	MAC interface disable	R/W	Sticky bit. 1: 1000BASE-X MAC interface disable when media link down.	0
14	MAC interface restart	R/W	Sticky bit. 1: 1000BASE-X MAC interface restart on media link change.	0
13	MAC interface PD enable	R/W	Sticky bit. 1: MAC interface autonegotiation parallel detect enable.	0
12	MAC interface autonegotiation restart	R/W	Self-clearing bit. 1: Restart MAC interface autonegotiation.	0
11	Force advertised ability	R/W	1: Force 16-bit advertised ability from register 18E3.	0

Table 80. MAC SerDes PCS Control, Address 16E3 (0x10) *(continued)*

Bit	Name	Access	Description	Default
10:8	SGMII preamble control	R/W	000: No effect on the start of packet. 001: If both the first two nibbles of the 10/100 packet are not 0x5, a byte of 0x55 must be prefixed to the output, otherwise there will be no effect on the start of packet. 010: If both the first two nibbles of the 10/100 packet are not 0x5, a byte of 0x55 must be prefixed to the output. An additional byte of 0x55 must be prefixed to the output if the next two nibbles are also not 0x5. 011–111: Reserved.	001
7	MAC SerDes autonegotiation enable	R/W	1: MAC SerDes ANEG enable.	0
6	SerDes polarity at input of MAC	R/W	1: Invert polarity of signal received at input of MAC.	0
5	SerDes polarity at output of MAC	R/W	1: Invert polarity of signal at output of MAC.	
4	Fast link status enable	R/W	1: Use fast link fail indication as link status indication to MAC SerDes 0: Use normal link status indication to MAC SerDes	0
3	Unidirectional enable	R/W	1: Enable transmit on MAC interface regardless of whether the PHY has determined that a valid link has been established. 0: Enable transmit on MAC interface only when the PHY has determined that a valid link has been established.	0
2:0	Reserved	RO	Reserved.	

4.5.2 MAC SerDes PCS Status

The register at address 17E3 consists of the bits that provide status from the MAC SerDes PCS block. The following table shows the settings available.

Table 81. MAC SerDes PCS Status, Address 17E3 (0x11)

Bit	Name	Access	Description
15:13	Reserved	RO	Reserved
12	SGMII alignment error	RO	1: SGMII alignment error occurred
11	MAC interface LP autonegotiation restart	RO	1: MAC interface link partner autonegotiation restart request occurred
10	Reserved	RO	Reserved
9:8	MAC remote fault	RO	01, 10, and 11: Remote fault detected from MAC 00: No remote fault detected from MAC
7	Asymmetric pause advertisement	RO	1: Asymmetric pause advertised by MAC
6	Symmetric pause advertisement	RO	1: Symmetric pause advertised by MAC

Table 81. MAC SerDes PCS Status, Address 17E3 (0x11) *(continued)*

Bit	Name	Access	Description
5	Full duplex advertisement	RO	1: Full duplex advertised by MAC
4	Half duplex advertisement	RO	1: Half duplex advertised by MAC
3	MAC interface LP autonegotiation capable	RO	1: MAC interface link partner autonegotiation capable
2	MAC interface link status	RO	1: MAC interface link status connected
1	MAC interface autonegotiation complete	RO	1: MAC interface autonegotiation complete
0	MAC interface PCS signal detect	RO	1: MAC interface PCS signal detect present

4.5.3 MAC SerDes Clause 37 Advertised Ability

The register at address 18E3 consists of the bits that provide access to and control over MAC SerDes Clause 37 advertised ability. The following table shows the settings available.

Table 82. MAC SerDes Cl37 Advertised Ability, Address 18E3 (0x12)

Bit	Name	Access	Description	Default
15:0	MAC SerDes advertised ability	R/W	Current configuration code word being advertised (this register is read/write if 16E3.11 = 1)	0x0000

4.5.4 MAC SerDes Clause 37 Link Partner Ability

The register at address 19E3 consists of the bits that provide status of the MAC SerDes link partner's Clause 37 advertised ability. The following table shows the settings available.

Table 83. MAC SerDes Cl37 LP Ability, Address 19E3 (0x13)

Bit	Name	Access	Description
15:0	MAC SerDes LP ability	RO	Last configuration code word received from link partner

4.5.5 MAC SerDes Status

The register at address 20E3 consists of the bits that provide access to MAC SerDes status. The following table shows the settings available.

Table 84. MAC SerDes Status, Address 20E3 (0x14)

Bit	Name	Access	Description
15	K28.5 comma realignment	RO	Self-clearing bit. 1: K28.5 comma re-alignment has occurred
14	SerDes signal detect	RO	Self-clearing bit. Sticky bit. 1: SerDes signal detection occurred
13:0	Reserved	RO	Reserved

4.5.6 MAC SerDes Transmit Good Packet Counter

The register at address 21E3 consists of the bits that provide status of the MAC SerDes transmit good packet counter. The following table shows the settings available.

Table 85. MAC SerDes Tx Good Packet Counter, Address 21E3 (0x15)

Bit	Name	Access	Description
15	Tx good packet counter active	RO	1: Transmit good packet counter active
14	Reserved	RO	Reserved
13:0	Tx good packet count	RO	Transmit good packet count modulo 10000

4.5.7 MAC SerDes Transmit CRC Error Counter

The register at address 22E3 consists of the bits that provide status of the MAC SerDes transmit packet count that had a CRC error. The following table shows the settings available.

Table 86. MAC SerDes Tx CRC Error Counter, Address 22E3 (0x16)

Bit	Name	Access	Description
15:8	Reserved	RO	Reserved
7:0	Tx CRC packet count	RO	Transmit CRC packet count (saturates at 255)

4.5.8 Media SerDes PCS Control

The register at address 23E3 consists of the bits that provide access to and control over Media SerDes PCS control. The following table shows the settings available.

Table 87. Media SerDes PCS Control, Address 23E3 (0x17)

Bit	Name	Access	Description	Default
15:14	Reserved	RO	Reserved	
13	Media interface autonegotiation parallel-detection	R/W	Sticky bit. 1: SerDes media autonegotiation parallel detect enabled	0
12	Reserved	RO	Reserved	
11	Force advertised ability	R/W	1: Force 16-bit advertised ability from register 25E3.15:0	0
10:0	Reserved	RO	Reserved	

4.5.9 Media SerDes PCS Status

The register at address 24E3 consists of the bits that provide status of the Media SerDes PCS block. The following table shows the settings available.

Table 88. Media SerDes PCS Status, Address 24E3 (0x18)

Bit	Name	Access	Description
15:12	Reserved	RO	Reserved
11	Media interface link partner autonegotiation restart	RO	1: Media interface link partner autonegotiation restart request occurred
10	Reserved	RO	Reserved
9:8	Remote fault detected	RO	01, 10, 11: Remote fault detected from link partner
7	Link partner asymmetric pause	RO	1: Asymmetric pause advertised by link partner
6	Link partner symmetric pause	RO	1: Symmetric pause advertised by link partner
5	Link partner full duplex advertisement	RO	1: Full duplex advertised by link partner
4	Link partner half duplex advertisement	RO	1: Half duplex advertised by link partner
3	Link partner autonegotiation capable	RO	1: Media interface link partner autonegotiation capable
2	Media interface link status	RO	1: Media interface link status
1	Media interface autonegotiation complete	RO	1: Media interface autonegotiation complete
0	Media interface signal detect	RO	1: Media interface signal detect

4.5.10 Media SerDes Clause 37 Advertised Ability

The register at address 25E3 consists of the bits that provide access to and control over Media SerDes Clause 37 advertised ability. The following table shows the settings available.

Table 89. Media SerDes Cl37 Advertised Ability, Address 25E3 (0x19)

Bit	Name	Access	Description	Default
15:0	Media SerDes advertised ability	R/W	Current configuration code word being advertised. This register is read/write when 23E3.11 = 1.	0x0000

4.5.11 Media SerDes Clause 37 Link Partner Ability

The register at address 26E3 consists of the bits that provide status of the media SerDes link partner's Clause 37 advertised ability. The following table shows the settings available.

Table 90. MAC SerDes Cl37 LP Ability, Address 26E3 (0x1A)

Bit	Name	Access	Description
15:0	Media SerDes LP ability	RO	Last configuration code word received from link partner

4.5.12 Media SerDes Status

The register at address 27E3 consists of the bits that provide access to Media SerDes status. The following table shows the settings available.

Table 91. Media SerDes Status, Address 27E3 (0x1B)

Bit	Name	Access	Description
15	K28.5 comma realignment	RO	Self-clearing bit. 1: K28.5 comma re-alignment has occurred
14	Signal detect	RO	Self-clearing bit. Sticky bit. 1: SerDes media signal detect
13:0	Reserved	RO	Reserved

4.5.13 Fiber Media CRC Good Counter

Register 28E3 makes it possible to read the contents of the CRC good counter for packets that are received on the Fiber media interface; the number of CRC routines that have executed successfully. The following table shows the expected readouts.

Table 92. Fiber Media CRC Good Counter, Address 28E3 (0x1C)

Bit	Name	Access	Description	Default
15	Packet since last read	RO	Self-clearing bit. 1: Packet received since last read.	0
14	Reserved	RO	Reserved.	
13:0	Fiber media CRC good counter contents	RO	Self-clearing bit. Counter containing the number of packets with valid CRCs. This counter does not saturate and will roll over.	0x000

4.5.14 Fiber Media CRC Error Counter

Register 29E3 makes it possible to read the contents of the CRC error counter for packets that are received on the Fiber media interface. The following table shows the expected readouts.

Table 93. Fiber Media CRC Error Counter, Address 29E3 (0x1D)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved.	
7:0	Fiber Media CRC error counter	RO	Self-clearing bit. CRC error counter for packets received on the Fiber media interface. The value saturates at 0xFF and subsequently clears when read and restarts count.	0x00

4.6 General Purpose Registers

Accessing the General Purpose register space is similar to accessing the extended page registers. Set register 31 to 0x0010. This sets all 32 registers to the general purpose register space.

To restore main register page access, write 0x0000 to register 31.

The following table lists the addresses and register names in the general purpose register page space. These registers are accessible only when the device register 31 is set to 0x0010.

Table 94. General Purpose Registers Page Space

Address	Name
0G–12G	Reserved
13G	LED/SIGDET/GPIO Control
14G	GPIO Control 2
15G	GPIO Input
16G	GPIO Output
17G	GPIO Output Enable
18G	Reserved
19G	GPIO Control 3
20G	I2C MUX Control 1
21G	I2C MUX Control 2
22G	I2C MUX Data Read/Write
23G	Recovered Clock 0 Control
24G	Recovered Clock 1 Control
25G	Enhanced LED Control
26G	Temperature Monitor Control 1
27G	Temperature Monitor Control 2
28G	Temperature Monitor Control 3
29G	Global Interrupt Status

Table 94. General Purpose Registers Page Space *(continued)*

Address	Name
30G	Reserved

4.6.1 Reserved General Purpose Address Space

The bits in registers 0G to 12G, 18G, and 30G of the general purpose register space are reserved.

4.6.2 LED/SIGDET/GPIO Control

The LED control bits configure the LED[3:0]_[3:0] pins to function as either LED control pins for each PHY, or as general purpose I/O pins. The SIGDET control bits configure the GPIO[3:0]/SIGDET[3:0] pins to function either as signal detect pins for each fiber media port, or as GPIOs. The following table shows the values that can be written.

Table 95. LED/SIGDET/GPIO Control, Address 13G (0x0D)

Bit	Name	Access	Description	Default
15:14	LED3_[3:0]	R/W	00: LED operation for PHY3 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
13:12	LED2_[3:0]	R/W	00: LED operation for PHY2 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
11:10	LED1_[3:0]	R/W	00: LED operation for PHY1 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
9:8	LED0_[3:0]	R/W	00: LED operation for PHY0 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
7:6	GPIO3/SIGDET3 control	R/W	00: SIGDET operation 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
5:4	GPIO2/SIGDET2 control	R/W	00: SIGDET operation 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
3:2	GPIO5/I2C_SCL_1 control	R/W	00: SIGDET operation 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00
1:0	GPIO4/I2C_SCL_0 control	R/W	00: SIGDET operation 01, 10: Reserved 11: Controlled by MII registers 15G to 17G	00

4.6.3 GPIO Control 2

The GPIO control 2 register configures the functionality of the COMA_MODE and CLOCK_SQUELCH input pins.

Table 96. GPIO Control 2, Address 14G (0x0E)

Bit	Name	Access	Description	Default
15	Reserved	RO	Reserved	
14	Clock squelch input enable	R/W	Enables clock squelch control on the CLK_SQUELCH_IN pin.	0
13	COMA_MODE output enable (active low)	R/W	1: COMA_MODE pin is an input 0: COMA_MODE pin is an output	1
12	COMA_MODE output data	R/W	Value to output on the COMA_MODE pin when it is configured as an output	0
11	COMA_MODE input data	RO	Data read from the COMA_MODE pin	
10	Tri-state enable for I2C bus	R/W	1: Tri-state I2C bus output signals instead of driving them high. This allows those signals to be pulled above VDD25 using an external pull-up resistor. 0: Drive I2C bus output signals to high and low values as appropriate	0
9	Tri-state enable for LEDs	R/W	1: Tri-state LED output signals instead of driving them high. This allows those signals to be pulled above VDD25 using an external pull-up resistor. 0: Drive LED bus output signals to high and low values as appropriate.	0

4.6.4 GPIO Input

The input register contains information about the input to the device GPIO pins. Read from this register to access the data on the device GPIO pins. The following table shows the readout you can expect.

Table 97. GPIO Input, Address 15G (0x0F)

Bit	Name	Access	Description
15:8	Reserved	RO	Reserved
7:0	GPIO input	RO	Data read from the GPIO_TBD pins

4.6.5 GPIO Output

The output register allows you to access and control the output from the device GPIO pins. The following table shows the values you can write.

Table 98. GPIO Output, Address 16G (0x10)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved	
7:0	GPIO output	R/W	Data written to the GPIO_TBD pins	0x00

4.6.6 GPIO Pin Configuration

Register 17G in the GPIO register space controls whether a particular GPIO pin functions as an input or an output. The following table shows the settings available.

Table 99. GPIO Input/Output Configuration, Address 17G (0x11)

Bit	Name	Access	Description	Default
15:8	Reserved	RO	Reserved	
7:0	GPIO[3:0]/SIGDET[3:0], GPIO[7:4]/I2C_SCL_[3:0], and GPIO9/FASTLINK-FAIL input or output enable	R/W	1: Pin is configured as an output 0: Pin is configured as an input	0x00

4.6.7 GPIO Control 3

Register 19G in the GPIO register space controls the MAC interface mode and the selection of the source PHY for the fast link failure indication. The following table shows the settings available for the GPIO9/FASTLINK-FAIL pin.

Table 100. GPIO Control 3, Address 19G (0x13)

Bit	Name	Access	Description	Default
15:14	MAC interface mode select	R/W	Select MAC interface mode 00: QSGMII to CAT5 mode 01: SGMII to CAT5 mode 10: QSGMII to CAT5 and Fiber mode 11: Reserved	00
13:4	Reserved	RO	Reserved	
3:0	Fast link failure port setting	R/W	0000: PHY0 0001: PHY1 0010: PHY2 0011: PHY3 1100–1111: Output disabled	0xF

4.6.8 I2C MUX Control 1

The following table shows the settings available to control the integrated I2C MUX.

Table 101. I2C MUX Control 1, Address 20G (0x14)

Bit	Name	Access	Description	Default
15:9	I2C device address	R/W	Top 7 bits of the 8-bit address sent out on the two wire serial stream. The bottom bit is the read/write signal, which is controlled by register 21G, bit 8. SFPs use 0xA0.	0xA0
8:6	Reserved	RO	Reserved.	
5:4	I2C SCL clock frequency	R/W	00: 50 kHz 01: 100 kHz 10: 400 kHz 11: 2 MHz	01

Table 101. I2C MUX Control 1, Address 20G (0x14) *(continued)*

Bit	Name	Access	Description	Default
3	I2C MUX port 3 enable	R/W	1: Enabled. 0: I2C disabled. Becomes GPIO pin.	0
2	I2C MUX port 2 enable	R/W	1: Enabled. 0: I2C disabled. Becomes GPIO pin.	0
1	I2C MUX port 1 enable	R/W	1: Enabled. 0: I2C disabled. Becomes GPIO pin.	0
0	I2C MUX port 0 enable	R/W	1: Enabled. 0: I2C disabled. I2C MUX port 0 becomes GPIO pin if serial LED function is enabled, regardless of the settings of this bit.	0

4.6.9 I2C MUX Control 2

Register 21G is used to control the I2C MUX for status and control of I2C slave devices.

Table 102. I2C MUX Interface Status and Control, Address 21G (0x15)

Bit	Name	Access	Description	Default
15	I2C MUX ready	RO	1: I2C MUX is ready for read or write	
14:12	Reserved	RO	Reserved	
11:10	PHY port Address	R/W	Specific VSC8574 PHY port being addressed.	00
9	Enable I2C MUX access	R/W	Self-clearing bit. 1: Execute read or write through the I2C MUX based on the settings of register bit 21G.8	0
8	I2C MUX read or write	R/W	1: Read from I2C MUX 0: Write to I2C MUX	1
7:0	I2C MUX address	R/W	Sets the address of the I2C MUX used to direct read or write operations.	0x00

4.6.10 I2C MUX Data Read/Write

Register 22G in the extended register space enables access to the I2C MUX.

Table 103. I2C MUX Data Read/Write, Address 22G (0x16)

Bit	Name	Access	Description	Default
15:8	I2C MUX read data	RO	Eight-bit data read from I2C MUX; requires setting both register 21G.9 and 21G.8 to 1.	
7:0	I2C MUX write data	R/W	Eight-bit data to be written to I2C MUX.	0x00

4.6.11 Recovered Clock 0 Control

Register 23G in the extended register space controls the functionality of the recovered clock 0 output signal.

Table 104. Recovered Clock 0 Control, Address 23G (0x17)

Bit	Name	Access	Description	Default
15	Enable RCVRDCLK1	R/W	1: Enable recovered clock 1 output 0: Disable recovered clock 1 output	0
14:11	Clock source select	R/W	Select bits for source PHY for recovered clock: 0000: PHY0 0001: PHY1 0010: PHY2 0011: PHY3 0100–1111: Reserved	0000
10:8	Clock frequency select	R/W	Select output clock frequency: 000: 25 MHz output clock 001: 125 MHz output clock 010: 31.25 MHz output clock 011–111: Reserved	000
7:6	Reserved	RO	Reserved.	
5:4	Clock squelch level	R/W	Select clock squelch level 00: Automatically squelch clock to low when the link is not up, is unstable, is up in a mode that does not support the generation of a recovered clock (1000BASE-T master or 10BASE-T), or is up in EEE mode (100BASE-TX or 1000BASE-T slave). 01: Same as 00 except that the clock is also generated in 1000BASE-T master and 10BASE-T link-up modes. This mode also generates a recovered clock output in EEE mode during reception of LP_IDLE. 10: Squelch only when the link is not up. 11: Disable clock squelch. **Note** A clock from the SerDes or Cu PHY will be output on the recovered clock output in this mode when the link is down.	
3	Reserved	RO	Reserved.	
2:0	Clock selection for specified PHY	R/W	000: Serial media recovered clock (only valid for PHYs 8, 9, 10, and 11 which support dual media functionality) 001: Copper PHY recovered clock 010: Copper PHY transmitter TCLK 011–111: Reserved	000

4.6.12 Recovered Clock 1 Control

Register 24G in the extended register space controls the functionality of the recovered clock 1 output signal.

Table 105. Recovered Clock 1 Control, Address 24G (0x18)

Bit	Name	Access	Description	Default
15	Enable RCVRDCLK2	R/W	Enable recovered clock 2 output	0
14:11	Clock source select	R/W	Select bits for source PHY for recovered clock: 0000: PHY0 0001: PHY1 0010: PHY2 0011: PHY3 0100–1111: Reserved	0000
10:8	Clock frequency select	R/W	Select output clock frequency: 000: 25 MHz output clock 001: 125 MHz output clock 010: 31.25 MHz output clock 011–111: Reserved	000
7:6	Reserved	RO	Reserved	
5:4	Clock squelch level	R/W	Select clock squelch level: 00: Automatically squelch clock to low when the link is not up, is unstable, is up in a mode that does not support the generation of a recovered clock (1000BASE-T master or 10BASE-T), or is up in EEE mode (100BASE-TX or 1000BASE-T slave). 01: Same as 00 except that the clock is also generated in 1000BASE-T master and 10BASE-T link-up modes. This mode also generates a recovered clock output in EEE mode during reception of LP_IDLE 10: Squelch only when the link is not up 11: Disable clock squelch. Note A clock from the SerDes or Cu PHY will be output on the recovered clock output in this mode when the link is down.	
3	Reserved	RO	Reserved	
2:0	Clock selection for specified PHY	R/W	000: Serial media recovered clock (only valid for PHYs 8, 9, 10, and 11 which support dual media functionality). 001: Copper PHY recovered clock 010: Copper PHY transmitter TCLK 011–111: Reserved	000

4.6.13 Enhanced LED Control

The following table contains the bits to control advanced functionality of the parallel and serial LED signals.

Table 106. Enhanced LED Control, Address 25G (0x19)

Bit	Name	Access	Description	Default
15:8	LED pulsing duty cycle control	R/W	Programmable control for LED pulsing duty cycle when bit 30.12 is set to 1. Valid settings are between 0 and 198. A setting of 0 corresponds to a 0.5% duty cycle and 198 corresponds to a 99.5% duty cycle. Intermediate values change the duty cycle in 0.5% increments	00
7	Serial LED output enable	R/W	Enable the serial LED output functionality for LED0_1, LED0_2, and LED0_3 pins 1: Pins function as serial LED outputs 0: Pins retain their normal function	0
6	Reserved	R/W	Reserved	0
5:3	Serial LED frame rate selection	R/W	Select frame rate of serial LED stream 000: 2500 Hz frame rate 001: 1000 Hz frame rate 010: 500 Hz frame rate 011: 250 Hz frame rate 100: 200 Hz frame rate 101: 125 Hz frame rate 110: 40 Hz frame rate 111: Output basic serial LED stream See Table 28, page 96.	
2:1	Serial LED select	R/W	Select which LEDs from each PHY to enable on the serial stream 00: Enable all four LEDs of each PHY 01: Enable LEDs 2, 1 and 0 of each PHY 10: Enable LEDs 1 and 0 of each PHY 11: Enable LED 0 of each PHY	00
0	LED port swapping	R/W	See "LED Port Swapping," page 97.	

4.6.14 Temperature Monitor Control 1

The following table contains the bits to control the on-chip temperature monitor function.

Table 107. Temperature Monitor Control 1, Address 26G (0x1A)

Bit	Name	Access	Description	Default
15:11	Reserved	RO	Reserved	
10	Power up TMON2	R/W	Power-up control for TMON2 1: Power-up TMON2 0: Power-down TMON2	0
9	Power up TMON1	R/W	Power-up control for TMON1 1: Power-up TMON1 0: Power-down TMON1	0

Table 107. Temperature Monitor Control 1, Address 26G (0x1A) *(continued)*

Bit	Name	Access	Description	Default
8	Power up TMON0	R/W	Power-up control for TMON 01: Power-up TMON 00: Power-down TMON0	0
7	Reset TMONs	R/W	Reset all TMON blocks (active low) 1: De-assert reset 0: Assert reset	0
6	Initiate temperature measurement in selected TMON	R/W	Initiate temperature measurement cycle in the TMON selected in register 28G bits 15:14	0
5	Enable background temperature monitoring	R/W	Enable background temperature monitoring in the TMON selected through bits 10:8. When enabled, this mode initiates a temperature read cycle at a 40 Hz rate. The averaging selected in bits 1:0 determine the polling rate	0
4	Enable alarm threshold interrupt	R/W	Enable interrupt in 28G.11 on alarm threshold temperature condition being met 1: Enable interrupt 0: Disable interrupt	0
3	Enable unalarm threshold interrupt	R/W	Enable interrupt in 28G.10 on unalarm threshold temperature condition being met 1: Enable interrupt 0: Disable interrupt	0
2	Reverse alarm/unalarm conditions	R/W	Reverse alarm/unalarm conditions so that the interrupt is generated on a temperature reading less than the alarm threshold and greater than the unalarm threshold. This allows for monitoring a negative slope signal with the alarm mechanism. 1: Reverse alarm/unalarm conditions 0: Normal alarm/unalarm conditions	0
1:0	Temperature averaging rate	R/W	Select temperature averaging and polling rates 00: No averaging with polling at a 40 Hz rate 01: Average 2 successive readings with polling at a 20 Hz rate 10: Average 4 successive readings with polling at a 10 Hz rate 11: Average 8 successive readings with polling at a 5 Hz rate	

4.6.15 Temperature Monitor Control 2

The following table contains the bits to control the on-chip temperature monitor function.

Table 108. Temperature Monitor Control 2, Address 27G (0x1B)

Bit	Name	Access	Description	Default
15:8	Alarm threshold	R/W	Alarm threshold for the temperature monitor interrupt generation	0x00
7:0	Unalarm threshold	R/W	Unalarm threshold for the temperature monitor interrupt generation	0x00

4.6.16 Temperature Monitor Control 3

The following table contains the bits to control the on-chip temperature monitor function.

Table 109. Temperature Monitor Control 3, Address 28G (0x1C)

Bit	Name	Access	Description	Default
15:14	Select TMON	R/W	Select TMON to monitor temperature 00: Select TMON0 01: Select TMON1 10: Select TMON2 11: Reserved	00
13:12	Reserved	RO	Reserved	
11	Alarm condition interrupt	RO/SC	Interrupt asserted on alarm condition. The bit clears when it is read. 1: Alarm condition is triggered	
10	Unalarm condition interrupt	RO/SC	Interrupt asserted on unalarm condition. The bit clears when it is read 1: Unalarm condition is triggered	
9	Alarm state	RO	Alarm state where alarm goes on when alarm condition is met and goes off when unalarm condition is met. 1: Alarm on 0: Alarm off	
8	Temperature reading refreshed	RO/SC	Indication that temperature reading has been refreshed since last read. This bit is cleared when read.	
7:0	Temperature reading	RO	Most recent temperature reading. If polling is enabled, then this is the most recent averaged temperature. If polling is disabled, then this is the most recent TMON reading.	

4.6.17 Global Interrupt Status

The following table contains the interrupt status from the various sources to indicate which one caused that last interrupt on the pin.

Table 110. Global Interrupt Status, Address 29G (0x1D)

Bit	Name	Access	Description
15:13	Reserved	RO	Reserved
12	Temperature monitor interrupt source	RO	Indicates that the temperature monitor is the source of the interrupt. This bit is cleared when register 28G is read
11:4	Reserved	RO	Reserved
3	PHY3 interrupt source	RO	Indicates that PHY3 is the source of the interrupt. This bit is cleared when register 26 in PHY3 is read
2	PHY2 interrupt source	RO	Indicates that PHY2 is the source of the interrupt. This bit is cleared when register 26 in PHY2 is read
1	PHY1 interrupt source	RO	Indicates that PHY1 is the source of the interrupt. This bit is cleared when register 26 in PHY1 is read
0	PHY0 interrupt source	RO	Indicates that PHY0 is the source of the interrupt. This bit is cleared when register 26 in PHY0 is read

4.7 Clause 45 Registers to Support Energy Efficient Ethernet

This section describes the Clause 45 registers that are required to support energy efficient Ethernet. Access to these registers is through the IEEE standard registers 13 and 14 (MMD access control and MMD data or address registers) as described in section 4.2.11 and 4.2.12.

The following table lists the addresses and register names in the Clause 45 register page space.

Table 111. Clause 45 Registers Page Space

Address	Name
3.1	PCS Status 1
3.20	EEE Capability
3.22	EEE Wake Error Counter
7.60	EEE Advertisement
7.61	EEE Link Partner Advertisement

4.7.1 PCS Status 1

The bits in the PCS Status 1 register provide a status of the EEE operation from the PCS for the link that is currently active.

Table 112. PCS Status 1, Address 3.1

Bit	Name	Access	Description
15:12	Reserved	RO	Reserved
11	Tx LPI received	RO/LH	1: Tx PCS has received LPI 0: LPI not received
10	Rx LPI received	RO/LH	1: Rx PCS has received LPI 0: LPI not received
9	Tx LPI indication	RO	1: Tx PCS is currently receiving LPI 0: PCS is not currently receiving LPI
8	Rx LPI indication	RO	1: Rx PCS is currently receiving LPI 0: PCS is not currently receiving LPI
7:3	Reserved	RO	Reserved
2	PCS receive link status	RO	1: PCS receive link up 0: PCS receive link down
1:0	Reserved	RO	Reserved

4.7.2 EEE Capability

This register is used to indicate the capability of the PCS to support EEE functions for each PHY type. The following table shows the bit assignments for the EEE capability register.

Table 113. EEE Capability, Address 3.20

Bit	Name	Access	Description
15:3	Reserved	RO	Reserved
2	1000BASE-T EEE	RO	1: EEE is supported for 1000BASE-T 0: EEE is not supported for 1000BASE-T
1	100BASE-TX EEE	RO	1: EEE is supported for 100BASE-TX 0: EEE is not supported for 100BASE-TX
0	Reserved	RO	Reserved

4.7.3 EEE Wake Error Counter

This register is used by PHY types that support EEE to count wake time faults where the PHY fails to complete its normal wake sequence within the time required for the specific PHY type. The definition of the fault event to be counted is defined for each PHY and can occur during a refresh or a wakeup as defined by the PHY. This 16-bit counter is

reset to all zeros when the EEE wake error counter is read or when the PHY undergoes hardware or software reset.

Table 114. EEE Wake Error Counter, Address 3.22

Bit	Name	Access	Description
15:0	Wake error counter	RO	Count of wake time faults for a PHY

4.7.4 EEE Advertisement

This register defines the EEE advertisement that is sent in the unformatted next page following a EEE technology message code. The following table shows the bit assignments for the EEE advertisement register.

Table 115. EEE Advertisement, Address 7.60

Bit	Name	Access	Description	Default
15:3	Reserved	RO	Reserved	
2	1000BASE-T EEE	R/W	1: Advertise that the 1000BASE-T has EEE capability 0: Do not advertise that the 1000BASE-T has EEE capability	0
1	100BASE-TX EEE	R/W	1: Advertise that the 100BASE-TX has EEE capability 0: Do not advertise that the 100BASE-TX has EEE capability	0
0	Reserved	RO	Reserved	

4.7.5 EEE Link Partner Advertisement

All the bits in the EEE LP Advertisement register are read only. A write to the EEE LP advertisement register has no effect. When the AN process has been completed, this register will reflect the contents of the link partner's EEE advertisement register. The following table shows the bit assignments for the EEE advertisement register.

Table 116. EEE Advertisement, Address 7.61

Bit	Name	Access	Description
15:3	Reserved	RO	Reserved
2	1000BASE-T EEE	RO	1: Link partner is advertising EEE capability for 1000BASE-T 0: Link partner is not advertising EEE capability for 1000BASE-T
1	100BASE-TX EEE	RO	1: Link partner is advertising EEE capability for 100BASE-TX 0: Link partner is not advertising EEE capability for 100BASE-TX
0	Reserved	RO	Reserved

5 Electrical Specifications

This section provides the DC characteristics, AC characteristics, recommended operating conditions, and stress ratings for the VSC8574 device.

5.1 DC Characteristics

This section contains the DC specifications for the VSC8574 device.

5.1.1 VDD25

The following table shows the DC specifications for the pins referenced to VDD25. The specifications listed in the following table are valid only when VDD = 1.0 V, VDD1 = 1.0 V, VDD1A = 1.0 V, or VDD25A = 2.5 V.

Table 117. VDD25 DC Characteristics

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage	V_{OH}	2.0	2.8	V	IOH = –1.0 mA
Output low voltage	V_{OL}	–0.3	0.4	V	IOL = 1.0 mA
Input high voltage	V_{IH}	1.7	3.0	V	
Input low voltage	V_{IL}	–0.3	0.7	V	
Input leakage current	I_{ILEAK}	–32	32	µA	Internal resistor included
Output leakage current	I_{OLEAK}	–32	32	µA	Internal resistor included
Output low current drive strength	I_{OL}		6	mA	
Output high current drive strength	I_{OH}	–6		mA	

5.1.2 Internal Pull-Up or Pull-Down Resistors

Internal pull-up or pull-down resistors are specified in the following table. For more information about signals with internal pull-up or pull-down resistors, see "Pins by Function," page 181.

All internal pull-up resistors are connected to their respective I/O supply.

Table 118. Internal Pull-Up or Pull-Down Resistors

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Internal pull-up resistor	R_{PU}	30	50	90	kΩ
Internal pull-down resistor	R_{PD}	35	60	115	kΩ

5.1.3 Reference Clock

The following table shows the DC specifications for a differential reference clock input signal.

Table 119. Reference Clock DC Characteristics

Parameter	Symbol	Minimum	Maximum	Unit
Input voltage range	V_{IP},V_{IN}	–25	1260	mV
Input differential voltage	V_{ID}	150	1000	mV
Input common-mode voltage	V_{ICM}	0	1200(1)	mV

1. The maximum common-mode voltage is provided without a differential signal, because the input is internally AC-coupled. The common-mode voltage is only limited by the maximum and minimum input voltage range and the input signal's differential amplitude.

5.1.4 SerDes Interface (SGMII)

The SerDes output drivers are designed to operate in SGMII/LVDS mode and high-drive/PECL mode (SFP and 1000BASE-KX modes). The SGMII/LVDS mode meets or exceeds the DC requirements of Serial-GMII Specification Revision 1.9 (ENG-46158), unless otherwise noted. The following table lists the DC specifications for the SGMII driver. The values are valid for all configurations, unless stated otherwise.

Table 120. SerDes Driver DC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, VOA or VOB	V_{OH}		1250	mV	RL = 100 Ω ±1%
Output low voltage, VOA or VOB	V_{OL}	0		mV	RL = 100 Ω ±1%
Output differential peak voltage	$\|V_{OD}\|$	150	400	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%
Output differential peak voltage	$\|V_{OD}\|$	150	600	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%
Output differential peak voltage, SGMII mode	$\|V_{OD}\|$	150	400	mVp	
Output differential peak voltage, 1000BASE-KX mode	$\|V_{OD}\|$	400	600	mVp	VDD1 = 1.0 V
Output differential peak voltage, SFP mode	$\|V_{OD}\|$	300	600	mVp	
Output offset voltage(1)	V_{OS}	420	580	mV	VDD1 = 1.0 V RL = 100 Ω ±1%
Output offset voltage(1)	V_{OS}	445	605	mV	VDD1 = 1.0 V RL = 100 Ω ±1%
DC output impedance, single-ended	R_O	40	140	Ω	SGMII mode VC = 1.0 V and 1.2 V See Figure 64, page 163

Table 120. SerDes Driver DC Specifications *(continued)*

Parameter	Symbol	Minimum	Maximum	Unit	Condition
RO mismatch between A and B(2)	ΔR_O		10	%	SGMII mode VC = 1.0 V and 1.2 V See Figure 64, page 163
Change in |VOD| between 0 and 1	$\Delta\|V_{OD}\|$		25	mV	SGMII mode RL = 100 Ω ±1%
Change in VOS between 0 and 1	ΔV_{OS}		25	mV	SGMII mode RL = 100 Ω ±1%
Output current	$\|I_{OSA}\|$, $\|I_{OSB}\|$		40	mA	SGMII mode, driver shorted to GND
Output current	$\|I_{OSAB}\|$		12	mA	SGMII mode, drivers shorted together

1. Requires AC-coupling for SGMII compliance.
2. Matching of reflection coefficients. For more information about test methods, see IEEE 1596.3-1996.

Figure 62. SGMII DC Transmit Test Circuit

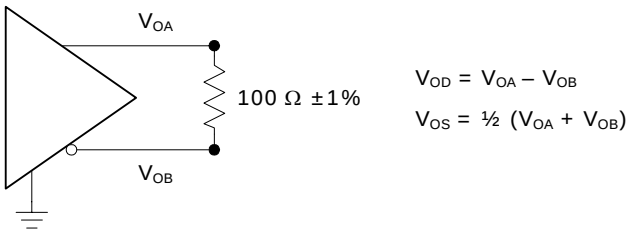

Figure 63. SGMII DC Definitions

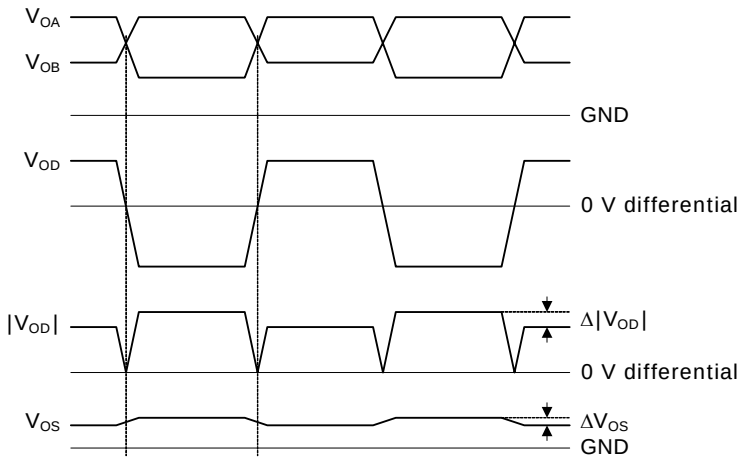

$\Delta|V_{OD}| = | \, |V_{OAH} - V_{OBL}| - |V_{OBH} - V_{OAL}| \, |$

$\Delta V_{OS} = | \, ½(V_{OAH} + V_{OBL}) - ½(V_{OAL} + V_{OBH}) \, |$

Figure 64. SGMII DC Driver Output Impedance Test Circuit

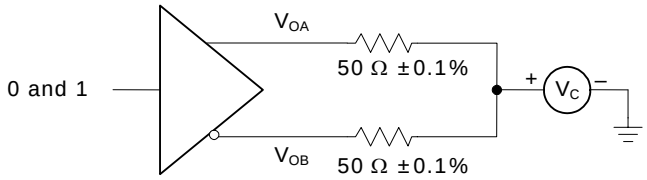

The following table lists the DC specifications for the SGMII receivers.

Table 121. SerDes Receiver DC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input voltage range, V_{IA} or V_{IB}	V_I	–25	1250	mV	
Input differential voltage	$\|V_{ID}\|$	50	2000	mVp-p	
Input common-mode voltage(1)	V_{ICM}	0	1200(2)	mV	Without any differential signal (internally AC-coupled)
Receiver differential input impedance	R_I	80	120	Ω	
Input differential hysteresis	V_{HYST}	25		mV	SGMII mode

1. SGMII compliancy requires external AC-coupling. When interfacing with specific Vitesse devices, DC-coupling is possible. For more information, contact your local Vitesse sales representative.
2. The maximum common-mode voltage is provided without a differential signal, because the input is internally AC-coupled. The common-mode voltage is only limited by the maximum and minimum input voltage range and the input signal's differential amplitude.

Figure 65. SGMII DC Input Definitions

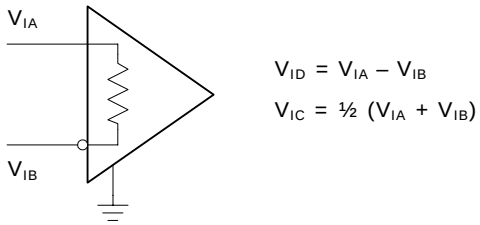

5.1.5 Enhanced SerDes Interface (QSGMII)

All DC specifications for the enhanced SerDes interface are compliant with QSGMII Specification Revision 1.3 and meet or exceed the requirements in the standard. They are also compliant with OIF-CEI-02.0 requirements where applicable.

The enhanced SerDes interface supports four major operating modes: SGMII, QSGMII, 2.5G, and SFP. The values in the following table apply to the modes specified in the condition column.

Table 122. Enhanced SerDes Driver DC Specifications

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Output differential peak voltage, SFP, 2.5G, and QSGMII modes	$\|V_{ODp}\|$	300		360	mV	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output differential peak voltage, SGMII mode	$\|V_{ODp}\|$	150		400	mVp	VDD1 = 1.0 V
Output differential peak voltage	$\|V_{ODp}\|$	400		600	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output differential peak voltage, SGMII mode	$\|V_{ODp}\|$	150		400	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output differential peak voltage, SFP mode	$\|V_{ODp}\|$	300		600	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output differential peak voltage, 2.5G mode	$\|V_{ODp}\|$	400		600	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output differential peak voltage, QSGMII mode	$\|V_{ODp}\|$	200		375	mVp	VDD1 = 1.0 V RL = 100 Ω ±1%, maximum drive
Output common-mode voltage	V_{OCM}	0		VDD1	mV	All modes RL = 100 Ω ±1%
DC output impedance, single-ended	R_O	40		140	Ω	SGMII mode VC = 1.0 V and 1.2 V See Figure 64, page 163
RO mismatch between A and B[1]	Δ_{RO}			10	%	SGMII mode VC = 1.0 V and 1.2 V See Figure 64, page 163
Change in |VOD| between 0 and 1	$\Delta\|V_{OD}\|$			25	mV	SGMII mode RL = 100 Ω ±1%
Change in VOS between 0 and 1	ΔV_{OS}			25	mV	SGMII mode RL = 100 Ω ±1%
Output current, SGMII and QSGMII modes	$\|I_{OSA}\|$, $\|I_{OSB}\|$			40	mA	Drivers shorted to ground
Output current, SGMII and QSGMII modes	$\|I_{OSAB}\|$			12	mA	Drivers shorted together

1. Matching of reflection coefficients. For more information about test methods, see IEEE 1596.3-1996.

The following table lists the DC specifications for the enhanced SerDes receiver.

Table 123. Enhanced SerDes Receiver DC Specifications

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Input voltage range, V_{IA} or V_{IB}[1]	V_I	–0.25		1.2	V
Input differential voltage	$\|V_{ID}\|$	100		1600	mV
Input common-mode voltage	V_{ICM}	0		1200	mV
Receiver differential input impedance	R_I	80	100	120	Ω

1. QSGMII DC input sensitivity is less than 400 mV.

5.1.6 Current Consumption

Current consumption is determined under the following operating conditions:

- QSGMII to 1000BASE-T mode
- QSGMII to 1000BASE-X dual media mode
- QSGMII to 100BASE-FX dual media mode
- SGMII to 1000BASE-T mode

The typical current consumption values in QSGMII to 1000BASE-T mode are based on nominal voltages with the MAC interface operating in QSGMII mode, and all media side ports operating in 1000BASE-T with full-duplex enabled. Data traffic is a 64-bit random data pattern at 100% utilization.

Table 124. QSGMII to 1000BASE-T Current Consumption

Parameter	Symbol	Typical	Maximum	Unit
Current with VDD at 1.0 V	I_{VDD}	TBD		mA
Current with VDD1 at 1.0 V	I_{VDD1}	TBD		mA
Current with VDD1A at 1.0 V	I_{VDD1A}	TBD		mA
Current with VDD25 at 2.5 V	I_{VDD25}	TBD		mA
Current with VDD25A at 2.5 V	I_{VDD25A}	TBD		mA
Worst-case power consumption	P_D		TBD	W

The typical current consumption values in QSGMII to 1000BASE-X dual media mode are based on nominal voltages with the MAC interface operating in QSGMII mode, 8 media side ports operating in 1000BASE-T with full-duplex enabled and 4 dual media ports operating in 1000BASE-X mode. Data traffic is a 64-bit random data pattern at 100% utilization.

Table 125. QSGMII to 1000BASE-X Current Consumption

Parameter	Symbol	Typical	Maximum	Unit
Current with VDD at 1.0 V	I_{VDD}	TBD		mA
Current with VDD1 at 1.0 V	I_{VDD1}	TBD		mA
Current with VDD1A at 1.0 V	I_{VDD1A}	TBD		mA
Current with VDD25 at 2.5 V	I_{VDD25}	TBD		mA

Table 125. QSGMII to 1000BASE-X Current Consumption *(continued)*

Parameter	Symbol	Typical	Maximum	Unit
Current with VDD25A at 2.5 V	I_{VDD25A}	TBD		mA
Worst-case power consumption	P_D		TBD	W

The typical current consumption values in QSGMII to 100BASE-FX dual media mode are based on nominal voltages with the MAC interface operating in QSGMII mode, 8 media side ports operating in 1000BASE-T with full-duplex enabled, and 4 dual media ports operating in 100BASE-FX mode. Data traffic is a 64-bit random data pattern at 100% utilization.

Table 126. QSGMII to 100BASE-FX Current Consumption

Parameter	Symbol	Typical	Maximum	Unit
Current with VDD at 1.0 V	I_{VDD}	TBD		mA
Current with VDD1 at 1.0 V	I_{VDD1}	TBD		mA
Current with VDD1A at 1.0 V	I_{VDD1A}	TBD		mA
Current with VDD25 at 2.5 V	I_{VDD25}	TBD		mA
Current with VDD25A at 2.5 V	I_{VDD25A}	TBD		mA
Worst-case power consumption	P_D		TBD	W

The typical current consumption values in SGMII to 1000BASE-T mode are based on nominal voltages with the MAC interface operating in SGMII mode and all media side ports operating in 1000BASE-T with full-duplex enabled. Data traffic is a 64-bit random data pattern at 100% utilization.

Table 127. SGMII to 1000BASE-T Current Consumption

Parameter	Symbol	Typical	Maximum	Unit
Current with VDD at 1.0 V	I_{VDD}	TBD		mA
Current with VDD1 at 1.0 V	I_{VDD1}	TBD		mA
Current with VDD1A at 1.0 V	I_{VDD1A}	TBD		mA
Current with VDD25 at 2.5 V	I_{VDD25}	TBD		mA
Current with VDD25A at 2.5 V	I_{VDD25A}	TBD		mA
Worst-case power consumption	P_D		TBD	W

5.2 AC Characteristics

This section provides the AC specifications for the VSC8574 device.

5.2.1 Reference Clock

The following table shows the AC specifications for a differential reference clock input. Performance is guaranteed for 125 MHz differential clocks only, however 25 MHz and single-ended clocks are also supported.

Table 128. Reference Clock AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Reference clock frequency, REFCLK_SEL2 = 1	*f*		125.00		MHz	±100 ppm Jitter < 1 ps RMS
Reference clock frequency, REFCLK_SEL2 = 0	*f*		25.00		MHz	±100 ppm
Duty cycle	DC	40	50	60	%	
Rise time and fall time	t_r, t_f			1.5	ns	20% to 80% threshold
Peak-to-peak jitter, iitter bandwidth between TBD kHz and TBD MHz	JPP_{REFCLK}			TBD	ps	
Peak-to-peak jitter, iitter bandwidth between TBD kHz and TBD MHz	JPP_{REFCLK}			TBD	ps	
Jitter transfer, REFCLK to SGMII output, jitter bandwidth between TBD kHz and TBD MHz	JBW_{REFCLK}			TBD	dB	
Jitter transfer, REFCLK to SGMII output, jitter bandwidth between TBD kHz and TBD MHz	JBW_{REFCLK}			TBD	dB	

5.2.2 Recovered Clock Outputs

The following table shows the AC specifications for the RCVRDCLK1 and RCVRDCLK2 outputs.

Table 129. Recovered Clock AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Recovered clock frequency	*f*		125.00		MHz	
Recovered clock frequency	*f*		31.25		MHz	
Recovered clock frequency	*f*		25.00		MHz	
Recovered clock cycle time	t_{RCYC}		8.0		ns	

Table 129. Recovered Clock AC Characteristics *(continued)*

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Recovered clock cycle time	t_{RCYC}		32.0		ns	
Recovered clock cycle time	t_{RCYC}		40.0		ns	
Frequency stability	$f_{STABILITY}$			50	ppm	
Duty cycle	DC	45	50	55	%	
Clock rise time and fall time	t_R, t_F			600	ps	20% to 80%
Peak-to-peak jitter, copper media interface	JPP_{CLK_Cu}			TBD	ps	Jitter bandwidth between TBD kHz and TBD MHz
Peak-to-peak jitter, fiber media interface	JPP_{CLK_Fi}			TBD	ps	Jitter bandwidth between TBD kHz and TBD MHz

5.2.3 SerDes Outputs

The values listed in the following table are valid for all configurations, unless otherwise noted.

Table 130. SerDes Output AC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
VOD ringing compared to VS	V_{RING}		±10	%	SGMII mode RL = 100 Ω ±1%
VOD rise time and fall time	t_R, t_F	100	200	ps	20% to 80% of V_S SGMII mode RL = 100 Ω ±1%
Differential output voltage	V_{OD}		30	mVp-p	Tx disabled
Differential output return loss, 50 MHz to 625 MHz	R_{LO_DIFF}	≥10		dB	1000BASE-KX mode RL = 100 Ω ±1%
Differential output return loss, 625 MHz to 1250 MHz	R_{LO_DIFF}	10–10 × log (*f*/625 MHz)		dB	1000BASE-KX mode RL = 100 Ω ±1%
Common-mode return loss	RL_{OCM}	6		dB	1000BASE-KX mode 50 MHz to 625 MHz
Interpair skew	t_{SKEW}		20	ps	SGMII mode

5.2.4 SerDes Driver Jitter

The following table lists the jitter characteristics for the SerDes output driver.

Table 131. SerDes Driver Jitter Characteristics

Parameter	Symbol	Maximum	Unit	Condition
Total jitter	TJ_O	192	ps	Measured according to IEEE 802.3.38.5.

Table 131. SerDes Driver Jitter Characteristics *(continued)*

Parameter	Symbol	Maximum	Unit	Condition
Deterministic jitter	DJ_O	80	ps	Measured according to IEEE 802.3.38.5.

5.2.5 SerDes Inputs

The following table lists the AC specifications for the SerDes inputs.

Table 132. SerDes Input AC Specifications

Parameter	Maximum	Unit	Condition
Differential input return loss, 50 MHz to 625 MHz	≥10	dB	1000BASE-KX mode RL = 100 Ω ±1%
Differential input return loss, 625 MHz to 1250 MHz	10–10 × log (*f*/625 MHz)	dB	1000BASE-KX mode RL = 100 Ω ±1%

5.2.6 SerDes Receiver Jitter Tolerance

The following table lists jitter tolerances for the SerDes receiver.

Table 133. SerDes Receiver Jitter Tolerance

Parameter	Symbol	Minimum	Unit	Condition
Total jitter tolerance	TJT_I	600	ps	1000BASE-KX and SFP mode Above 637 kHz Measured according to IEEE 802.3 38.6.8
Deterministic jitter tolerance	DJT_I	370	ps	1000BASE-KX and SFP mode Above 637 kHz Measured according to IEEE 802.3 38.6.8
Cycle distortion jitter tolerance	JT_{CD}	1.4	ns	100BASE-FX mode Measured according to ISO/IEC 9314-3:1990
Data-dependent jitter tolerance	DDJ	2.2	ns	100BASE-FX mode Measured according to ISO/IEC 9314-3:1990
Random jitter tolerance, peak-to-peak	RJT	2.27	ns	100BASE-FX mode Measured according to ISO/IEC 9314-3:1990

5.2.7 Enhanced SerDes Interface

All AC specifications for the enhanced SerDes interface are compliant with QSGMII Specification Revision 1.3 and meet or exceed the requirements in the standard. They are also compliant with the OIF-CEI-02.0 requirements where applicable.

The enhanced SerDes interface supports four major modes: SGMII, QSGMII, 2.5G, and SFP. The values in the tables in the following sections apply to the modes listed in the condition column and are based on the test circuit shown in Figure 62, page 162. The

transmit and receive eye specifications relate to the eye diagrams shown in the following illustration, with the compliance load as defined in the test circuit.

Figure 66. QSGMII Transient Parameters

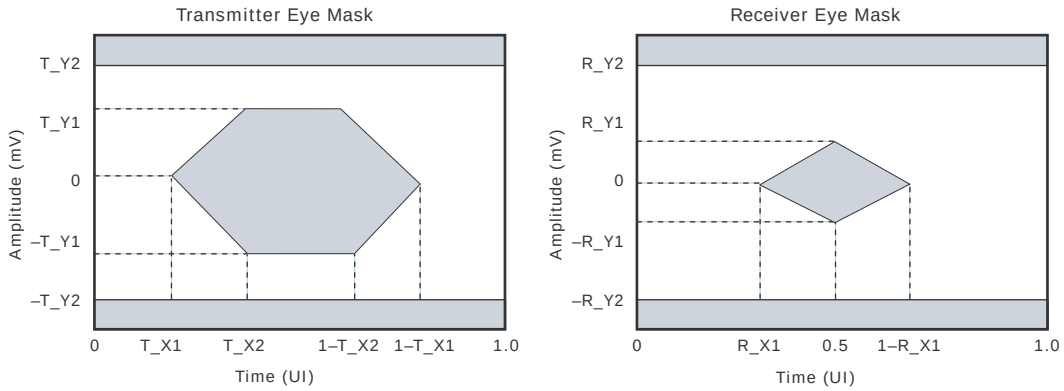

5.2.8 Enhanced SerDes Outputs

The following table provides the AC specifications for the enhanced SerDes outputs in SGMII mode.

Table 134. Enhanced SerDes Outputs AC Specifications in SGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 1.25G mode	UI				800 ps
Duty cycle	DC	48	52	%	
VOD ringing compared to VS	V_{RING}		±10	%	SGMII mode RL = 100 Ω ±1%
VOD rise time and fall time	t_R, t_F	100	200	ps	20% to 80% of VS SGMII mode RL = 100 Ω ±1%
Differential output voltage	V_{OD}		30	mVp-p	Tx disabled
Differential output return loss, 50 MHz to 625 MHz	R_{LO_DIFF}	≥10		dB	1000BASE-KX mode RL = 100 Ω ±1%
Differential output return loss, 625 MHz to 1250 MHz	R_{LO_DIFF}	10–10 × log (*f*/625 MHz)		dB	1000BASE-KX mode RL = 100 Ω ±1%
Common-mode return loss	RL_{OCM}	6		dB	1000BASE-KX mode 50 MHz to 625 MHz
Interpair skew	t_{SKEW}		20	ps	SGMII mode

The following table provides the AC specifications for the enhanced SerDes outputs in QSGMII mode.

Table 135. Enhanced SerDes Outputs AC Specifications, QSGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 5G	UI				200 ps
VOD rise time and fall time	t_R, t_F	60	130	ps	20% to 80% of VS RL = 100 Ω ±1%
Differential output voltage	V_{OD}		30	mVp-p	Tx disabled
Differential output return loss, 100 MHz to 2.5 GHz	RL_{O_DIFF}	8		dB	RL = 100 Ω ±1%
Differential output return loss, 2.5 GHz to 5 GHz	RL_{O_DIFF}	8 dB – 16.6 log (*f*/2.5 GHz)		dB	1000BASE-KX mode RL = 100 Ω ±1%
Interpair skew	t_{SKEW}		20	ps	
Eye mask X1	T_X1		0.15	UI	
Eye mask X2	T_X2		0.4	UI	
Eye mask Y1	T_Y1	200		mV	
Eye mask Y2	T_Y2		450	mV	

The following table provides the AC specifications for the enhanced SerDes outputs in 2.5G mode.

Table 136. Enhanced SerDes Outputs AC Specifications, 2.5G Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 2.5G	UI				320 ps
VOD rise time and fall time	t_R, t_F	60	130	ps	20% to 80% of V_S SGMII mode RL = 100 Ω ±1%
Differential output voltage	V_{OD}		30	mVp-p	Tx disabled
Differential output return loss, 100 MHz to 625 MHz	RL_{TX_DIFF}	10		dB	RL = 100 Ω ±1%
Differential output return loss, 625 MHz to 3.125 GHz	RL_{TX_DIFF}	10–10 × log (*f*/625 MHz)		dB	RL = 100 Ω ±1%
Interpair skew	t_{SKEW}		20	ps	
Eye mask X1	T_X1		0.175	UI	
Eye mask X2	T_X2		0.390	UI	
Eye mask Y1	T_Y1	200		mV	
Eye mask Y2	T_Y2		400	mV	

5.2.9 Enhanced Serial LEDs

This section contains the AC specifications for the enhanced serial LEDs. The duty cycle of the LED_PULSE signal is programmable and can be varied between 0.5% and 99.5%.

Table 137. Enhanced Serial LEDs AC Characteristics

Parameter	Symbol	Minimum	Maximum	Unit
LED_CLK cycle time	t_{CYC}	TBD	TBD	ns
Pause between LED_DATA bit sequences	t_{PAUSE}	TBD	TBD	µs
LED_CLK to LED_DATA	t_{CO}	TBD	TBD	ns
LED_CLK to LED_LD	t_{CL}	TBD	TBD	ns
LED_LD pulse width	t_{LW}	TBD	TBD	ns
LED_PULSE cycle time	t_{PULSE}	199	201	µs

Figure 67. Enhanced Serial LED Timing

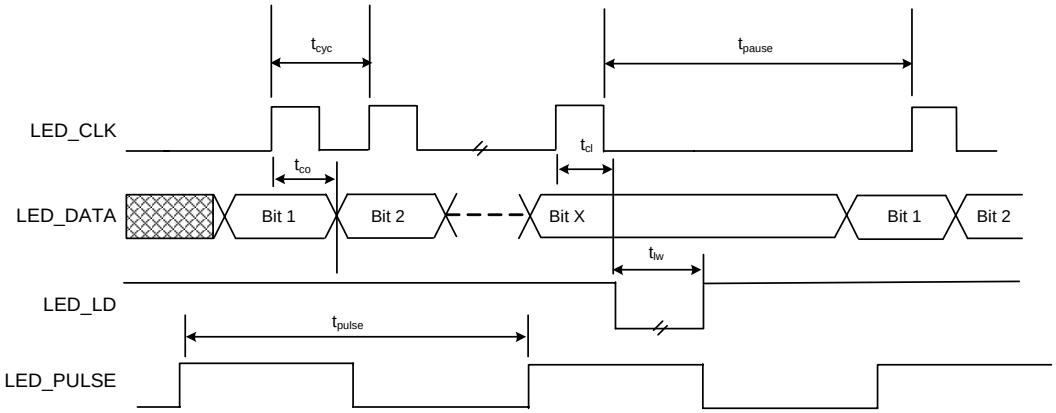

5.2.10 Enhanced SerDes Driver Jitter

The following table lists the jitter characteristics for the enhanced SerDes driver in QSGMII mode. For information about jitter characteristics for the enhanced SerDes driver in SGMII mode, see Table 131, page 168.

Table 138. Enhanced SerDes Driver Jitter Characteristics, QSGMII Mode

Parameter	Symbol	Maximum	Unit	Condition
Total output jitter	TJO	60	ps	Measured according to IEEE 802.3.38.5.
Deterministic output jitter	DJO	10	ps	Measured according to IEEE 802.3.38.5.

5.2.11 Enhanced SerDes Inputs

The following table lists the AC specifications for the enhanced SerDes inputs in SGMII mode.

Table 139. Enhanced SerDes Input AC Specifications, SGMII Mode

Parameter	Symbol	Minimum	Unit	Condition
Unit interval, 1.25G	UI		ps	800 ps
Differential input return loss	RLI_DIFF	10	dB	50 MHz to 625 MHz RL = 100 Ω ±1%
Common-mode input return loss	RLICM	6	dB	50 MHz to 625 MHz

The following table lists the AC specifications for the enhanced SerDes inputs in QSGMII mode.

Table 140. Enhanced SerDes Inputs AC Specifications, QSGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 5G	UI				200 ps
Differential input return loss, 100 MHz to 2.5 GHz	RLI_DIFF	8		dB	RL = 100 Ω ±1%
Differential input return loss, 2.5 GHz to 5 GHz	RLI_DIFF	8 dB – 16.6 log (f/2.5 GHz)		dB	RL = 100 Ω ±1%
Common-mode input return loss	RLICM	6		dB	100 MHz to 2.5 GHz
Eye mask X1	R_X1		0.3	UI	
Eye mask Y1	R_Y1		50	mV	
Eye mask Y2	R_Y2		450	mV	

The following table lists the AC specifications for the enhanced SerDes inputs in 2.5G mode.

Table 141. Enhanced SerDes Inputs AC Specifications, 2.5G Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 2.5G	UI				320 ps
Differential output return loss	RLI_DIFF	10		dB	100 MHz to 2.5 GHz RL = 100 Ω ±1%
Common-mode input return loss	RLICM	6		dB	100 MHz to 2.5 GHz
Eye mask X1	R_X1		0.275	UI	
Eye mask X2	R_X2		0.5	UI	
Eye mask Y1	R_Y1	100		mV	
Eye mask Y2	R_Y2		800	mV	

5.2.12 Enhanced SerDes Receiver Jitter Tolerance

The following table lists the jitter tolerance for the enhanced SerDes receiver in QSGMII mode. For information about jitter tolerance for the enhanced SerDes receiver in SGMII mode, see Table 133, page 169.

Table 142. Enhanced SerDes Receiver Jitter Tolerance, QSGMII Mode

Parameter	Symbol	Maximum	Unit	Condition
Bounded high-probability jitter(1)	BHPJ	90	ps	
Sinusoidal jitter, maximum	SJ_{MAX}	1000	ps	
Sinusoidal jitter, high frequency	SJ_{HF}	10	ps	
Total jitter tolerance	TJT_I	120	ps	Does not include sinusoidal jitter.

1. This is the sum of uncorrelated bounded high probability jitter (0.15 UI), and correlated bounded high probability jitter (0.30 UI). Uncorrelated bounded high probability jitter is distribution where the value of the jitter shows no correlation to any signal level being transmitted, formally defined as deterministic jitter (DJ). Correlated bounded high probability jitter is jitter distribution where the value of the jitter shows a strong correlation to the signal level being transmitted. This jitter can be considered as being equalizable, due to its correlation to the signal level.

5.2.13 JTAG Interface

This section provides the AC specifications for the JTAG interface. The specifications meet or exceed the requirements of IEEE 1149.1-2001. The JTAG receive signal requirements are requested at the pin of the device. The JTAG_TRST signal is asynchronous to the clock, and does not have a setup or hold time requirement.

Table 143. JTAG Interface AC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
TCK frequency	f		10	MHz	
TCK cycle time	t_C	100		ns	
TCK high time	$t_{W(CH)}$	40		ns	
TCK low time	$t_{W(CL)}$	40		ns	
Setup time to TCK rising	t_{SU}	10		ns	
Hold time from TCK rising	t_H	10		ns	
TDO valid after TCK falling	$t_{V(C)}$		28	ns	CL = 10 pF
TDO hold time from TCK falling	$t_{H(TDO)}$	0		ns	CL = 0 pF
TDO disable time(1)	t_{DIS}		30	ns	See Figure 69, page 175.
TRST time low	$t_{W(TL)}$	30		ns	

1. The pin begins to float when a 300 mV change from the actual VOH/VOL level occurs.

Figure 68. JTAG Interface Timing Diagram

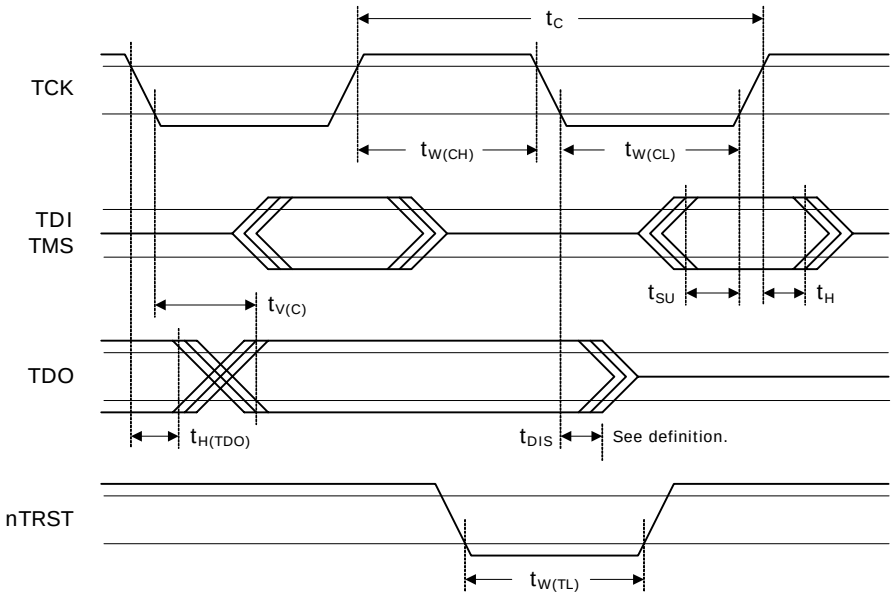

Figure 69. Test Circuit for TDO Disable Time

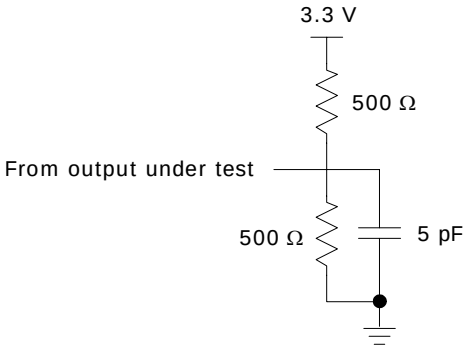

5.2.14 Serial Management Interface

This section contains the AC specifications for the serial management interface (SMI).

Table 144. SMI Interface AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit
MDC frequency	f_{CLK}	TBD	TBD	TBD	MHz
MDC cycle time	t_{CYC}	TBD	TBD	TBD	ns
MDC time high	t_{WH}	TBD	TBD	TBD	ns
MDC time low	t_{WL}	TBD	TBD	TBD	ns
Setup to MDC rising	t_{SU}	TBD	TBD	TBD	ns
Hold from MDC rising	t_H	TBD	TBD	TBD	ns

Table 144. SMI Interface AC Characteristics *(continued)*

Parameter	Symbol	Minimum	Typical	Maximum	Unit
MDC rise time	t_R	TBD	TBD	TBD	ns
MDC fall time	t_F	TBD	TBD	TBD	ns
MDC to MDIO valid	t_{CO}	TBD	TBD	TBD	ns

Figure 70. SMI Interface Timing

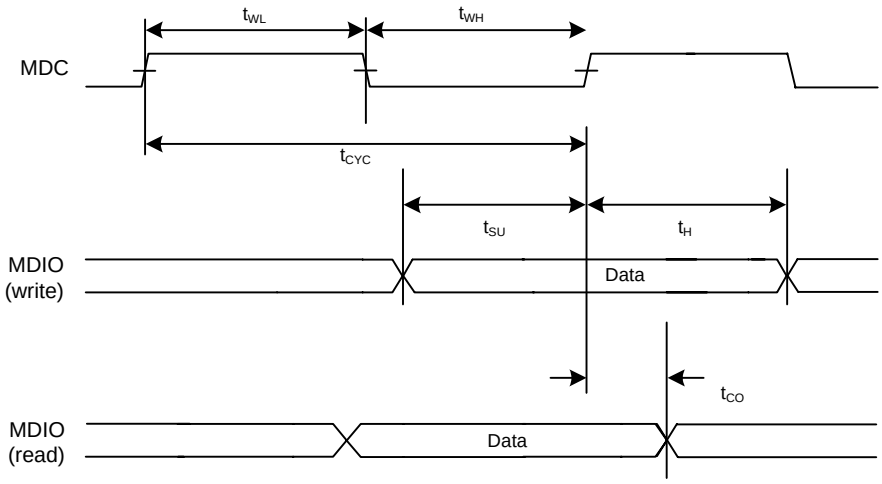

5.2.15 Reset Timing

This section contains the AC specifications that apply to device reset functionality. The signal applied to the NRESET input must comply with the specifications listed in the following table.

Table 145. NRESET Timing Specifications

Parameter	Symbol	Minimum	Maximum	Unit
NRESET assertion time after power supplies and clock stabilize	t_W	2		ms
Recovery time from reset inactive to device fully active	t_{REC}		5	ms
NRESET pulse width	$t_{W(RL)}$	100		ns
Wait time between NRESET de-assert and access of the SMI interface	t_{WAIT}	TBD	TBD	ms

Figure 71. Reset Timing

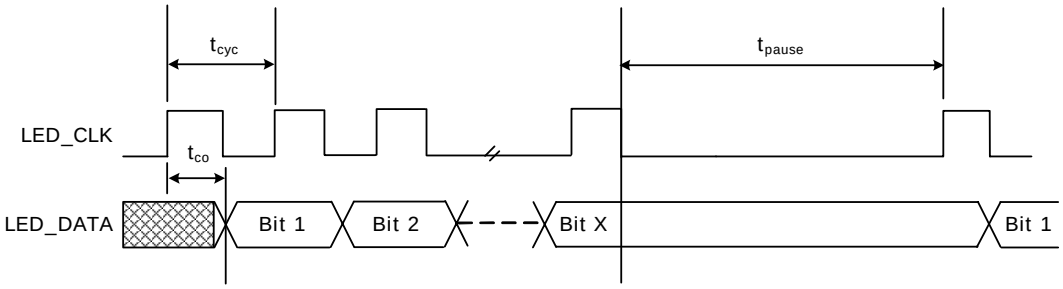

5.2.16 Serial LEDs

This section contains the AC specifications for the serial LEDs.

Table 146. Serial LEDs AC Characteristics

Parameter	Symbol	Minimum	Maximum	Unit
LED_CLK cycle time	t_{CYC}	TBD	TBD	ms
Pause between LED bit sequences	t_{PAUSE}	TBD	TBD	ms
LED_CLK to LED_DATA	t_{CO}	TBD	TBD	ns

Figure 72. Serial LED Timing

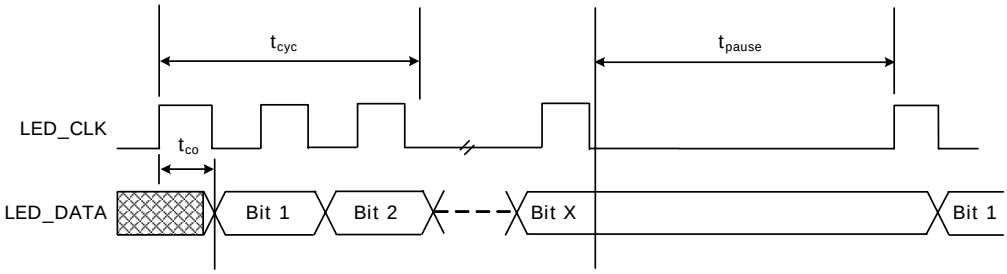

5.2.17 1588 Specifications

This section contains the AC specifications for the 1588 pins.

Table 147. 1588 Specifications Clock AC Characteristics

Parameter	Minimum	Typical	Maximum	Unit	Condition
1588 differential input clock	10		100	ps	
Reference clock frequency		125.00		MHz	±100 ppm Jitter < 10 ps RMS

5.3 Operating Conditions

The following table shows the recommended operating conditions for the VSC8574 device.

Table 148. Recommended Operating Conditions

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Power supply voltage for VDD	V_{DD}	0.95	1.00	1.05	V
Power supply voltage for VDD1	V_{DD1}	0.95	1.00	1.05	V
Power supply voltage for VDD1A	V_{DD1A}	0.95	1.00	1.05	V
Power supply voltage for VDD25	V_{DD25}	2.38	2.50	2.62	V
Power supply voltage for VDD25A	V_{DD25A}	2.38	2.50	2.62	V
VSC8574 operating temperature(1)	T	0		125	°C
VSC8574-03 operating temperature(1)	T	–40		125	°C

1. Minimum specification is ambient temperature, and the maximum is junction temperature. For carrier class applications, the maximum operating temperature is 110 °C junction.

5.4 Stress Ratings

This section contains the stress ratings for the VSC8574 device.

Warning Stresses listed in the following table may be applied to devices one at a time without causing permanent damage. Functionality at or exceeding the values listed is not implied. Exposure to these values for extended periods may affect device reliability.

Table 149. Stress Ratings

Parameter	Symbol	Minimum	Maximum	Unit
DC input voltage on any non-supply pin	I_{xx}	–0.5	3.6	V
DC input voltage on VDD supply pin	I_{DD}	–0.5	1.7	V
DC input voltage on VDD1 supply pin	I_{DD_VS}	–0.5	1.7	V
DC input voltage on VDD1A supply pin	I_{DD_AL}	–0.5	1.7	V
DC input voltage on VDD25 supply pin	I_{DD_IO}	–0.5	3.2	V
DC input voltage on VDD25A supply pin	I_{DD_AH}	–0.5	3.2	V
Storage temperature	T_S	–55	150	°C

Warning This device can be damaged by electrostatic discharge (ESD) voltage. Vitesse recommends that all integrated circuits be handled with appropriate precautions. Failure to observe proper handling and installation procedures may adversely affect reliability of the device.

6 Pin Descriptions

The VSC8574 device has 256 pins, which are described in this section.

The pin information is also provided as an attached Microsoft Excel file so that you can copy it electronically. In Acrobat, double-click the attachment icon.

6.1 Pin Identifications

This section contains the pin descriptions for the VSC8574 device. The following table provides notations for definitions of the various pin types.

Table 150. Pin Type Symbols

Symbol	Pin Type	Description
A_{BIAS}	Analog bias	Analog bias pin.
A_{DIFF}	Analog differential	Analog differential signal pair for twisted pair interface.
I	Input	Input with no on-chip pull-up or pull-down resistor.
I_A	Analog input	Analog input for sensing variable voltage levels.
I_{DIFF}	Input differential	Input differential signal pair
I_{PD}	Input with pull-down	Input with on-chip pull-down resistor to ground.
I_{PU}	Input with pull-up	Input with on-chip pull-up resistor to VDD25.
I_{PU}/O	Bidirectional with pull-up	Input and output signal with on-chip pull-up resistor VDD25.
I_{PU5V}	Input with pull-up	Input with on-chip pull-up resistor to VDD25. These pins are 5 V tolerant.
O	Output	Output signal.
O_{CRYST}	Crystal output	Crystal clock output pin. If not used, leave unconnected.
OD	Open drain	Open drain output.
O_{DIFF}	Output differential	Output differential signal pair.
O_{PU}	Output with pull-up	Output with on-chip pull-up resistor to VDD25.
OS	Open source	Open source output.

6.2 Pin Diagram

The following illustrations show the pin diagram for the VSC8574 device. For clarity, the device is shown in two halves, the top left and top right.

Figure 73. Pin Diagram, Top Left

	1	2	3	4	5	6	7	8
A	NC_1	TXVPA_3	TXVPB_3	TXVPC_3	TXVPD_3	TXVPA_2	TXVPB_2	TXVPC_2
B	VSS_1	TXVNA_3	TXVNB_3	TXVNC_3	TXVND_3	TXVNA_2	TXVNB_2	TXVNC_2
C	REFCLK_N	VDD25A_1	THERMDA	VDD25A_2	VSS_3	VDD25A_3	VDD1A_1	VDD1A_2
D	REFCLK_P	THERMDC_VSS	REF_FILT_A	REF_REXT_A	VSS_6	VSS_7	VSS_8	VSS_9
E	REFCLK_SEL2	TMS	TRST	VDD25A_6	VDD1_1	VSS_14	VSS_15	VSS_16
F	TDO	TDI	TCK	VSS_20	VDD1_3	VSS_21	VSS_22	VSS_23
G	LED0_0	LED1_0	LED2_0	LED3_0	VDD1_5	VSS_27	VSS_28	VSS_29
H	LED0_1	LED1_1	LED2_1	LED3_1	VDD1_7	VSS_33	VSS_34	VSS_35
J	LED0_2	LED1_2	LED2_2	LED3_2	VDD1_9	VSS_39	VSS_40	VSS_41
K	LED0_3	LED1_3	LED2_3	LED3_3	VDD1_11	VSS_45	VSS_46	VSS_47
L	RESERVED_5	1588_PPS_3	COMA_MODE	RESERVED_3	VDD1_13	VSS_51	VSS_52	VSS_53
M	RESERVED_6	MDINT	NRESET	VDD25_2	VDD1_15	VSS_57	VSS_58	VSS_59
N	RESERVED_7	MDIO	1588_PPS_1	1588_PPS_2	VDD1_17	VSS_63	VSS_64	VSS_65
P	RESERVED_8	MDC	VDD25_4	RESERVED_4	VDD25A_8	VDD1A_5	VDD1A_6	VDD1A_7
R	VSS_69	FIBROP_3	FIBRIP_3	RDP_3	TDP_3	FIBROP_2	FIBRIP_2	RDP_2
T	NC_3	FIBRON_3	FIBRIN_3	RDN_3	TDN_3	FIBRON_2	FIBRIN_2	RDN_2

Figure 74. Pin Diagram, Top Right

9	10	11	12	13	14	15	16	
TXVPD_2	TXVPA_1	TXVPB_1	TXVPC_1	TXVPD_1	TXVPA_0	TXVPB_0	NC_2	A
TXVND_2	TXVNA_1	TXVNB_1	TXVNC_1	TXVND_1	TXVNA_0	TXVNB_0	VSS_2	B
VDD1A_3	RESERVED_1	VDD25A_4	VSS_4	VDD1A_4	VDD25A_5	TXVNC_0	TXVPC_0	C
VSS_10	VSS_11	VSS_12	VSS_13	RESERVED_2	VSS_71	TXVND_0	TXVPD_0	D
VSS_17	VSS_18	VSS_19	VDD1_2	VDD25A_7	VSS_72	CLK_SQUELCH_IN	1588_SPI_CLK	E
VSS_24	VSS_25	VSS_26	VDD1_4	VSS_73	PHYADD4	VSS_74	RCVRDCLK1	F
VSS_30	VSS_31	VSS_32	VDD1_6	PHYADD2	PHYADD3	VSS_75	RCVRDCLK2	G
VSS_36	VSS_37	VSS_38	VDD1_8	VDD25_1	1588_SPI_DO	VSS_76	VSS_77	H
VSS_42	VSS_43	VSS_44	VDD1_10	VSS_78	1588_SPI_CS	1588_DIFF_INPUT_CLK_P	1588_DIFF_INPUT_CLK_N	J
VSS_48	VSS_49	VSS_50	VDD1_12	GPIO8/I2C_SDA	GPIO9/FASTLINK-FAIL	1588_LOAD_SAVE	1588_PPS_0	K
VSS_54	VSS_55	VSS_56	VDD1_14	GPIO4/I2C_SCL_0	GPIO5/I2C_SCL_1	GPIO6/I2C_SCL_2	GPIO7/I2C_SCL_3	L
VSS_60	VSS_61	VSS_62	VDD1_16	VDD25_3	GPIO1/SIGDET1	GPIO2/SIGDET2	GPIO3/SIGDET3	M
VSS_66	VSS_67	VSS_68	VDD1_18	SerDes_Rext_1	GPIO0/SIGDET0	TDP_0	TDN_0	N
VDD1A_8	VDD1A_9	VDD1A_10	VDD25A_9	VDD25A_10	SerDes_Rext_0	RDP_0	RDN_0	P
TDP_2	FIBROP_1	FIBRIP_1	RDP_1	TDP_1	FIBROP_0	FIBRIP_0	VSS_70	R
TDN_2	FIBRON_1	FIBRIN_1	RDN_1	TDN_1	FIBRON_0	FIBRIN_0	NC_4	T

6.3 Pins by Function

This section contains the functional pin descriptions for the VSC8574 device.

6.3.1 1588 Support

The following table lists the 1588 support pins.

Table 151. 1588 Support Pins

Name	Pin	Type	Description
1588_DIFF_INPUT_CLK_N 1588_DIFF_INPUT_CLK_P	J16 J15	I	125 MHz differential reference clock input pair
1588_LOAD_SAVE	K15	I	Sync signal to load the time to the 1588 engine
1588_PPS_0	K16	O	1588 local timer 0 pps configurable to PHY0
1588_PPS_1	N3	O	1588 local timer 1 pps configurable to PHY1
1588_PPS_2	N4	O	1588 local timer 2 pps configurable to PHY2
1588_PPS_3	L2	O	1588 local timer 1 pps configurable to PHY3
1588_SPI_CLK	E16	O	1588 SPI clock
1588_SPI_CS	J14	O	1588 SPI chip select
1588_SPI_DO	H14	O	1588 SPI data output

6.3.2 GPIO and SIGDET

The following table lists the GPIO and SIGDET pins.

Table 152. GPIO and SIGDET Pins

Name	Pin	Type	Description
GPIO0/SIGDET0 GPIO1/SIGDET1 GPIO2/SIGDET2 GPIO3/SIGDET3 GPIO4/I2C_SCL_0 GPIO5/I2C_SCL_1 GPIO6/I2C_SCL_2 GPIO7/I2C_SCL_3 GPIO8/I2C_SDA GPIO9/FASTLINK-FAIL	N14 M14 M15 M16 L13 L14 L15 L16 K13 K14	I_{PU}/O	General purpose input/output (GPIO). Six dedicated GPIO pins are provided. The fast link fail pin, four SIGDET pins, and five I2C controller pins can be configured to serve as GPIOs.

6.3.3 JTAG

The following table lists the JTAG test pins.

Table 153. JTAG Pins

Name	Pin	Type	Description
TCK	F3	I_{PU5V}	JTAG test clock input.
TDI	F2	I_{PU5V}	JTAG test serial data input.
TDO	F1	O_{PU}	JTAG test serial data output.
TMS	E2	I_{PU5V}	JTAG test mode select.

Table 153. JTAG Pins *(continued)*

Name	Pin	Type	Description
TRST	E3	I_{PU5V}	JTAG reset. **Important** When JTAG is not in use, this pin must be tied to ground with a pull-down resistor for normal operation.

6.3.4 Miscellaneous

The following table lists the miscellaneous pins.

Table 154. Miscellaneous Pins

Name	Pin	Type	Description
CLK_SQUELCH_IN	E15	I_A	Clock squelch input pin.
COMA_MODE	L3	I_{PU}	When this pin is asserted high, all PHYs are held in a powered down state. When de-asserted low, all PHYs are powered up and resume normal operation. This signal is also used to synchronize the operation of multiple chips on the same PCB to provide visual synchronization for LEDs driven by separate chips.
LED0_[0:3] LED1_[0:3] LED2_[0:3] LED3_[0:3]	G1, H1, J1, K1 G2, H2, J2, K2 G3, H3, J3, K3 G4, H4, J4, K4	O	LED direct-drive outputs. All LEDs pins are active-low. A serial LED stream can also be implemented. See "LED Mode Select," page 130. **Note** LEDbit_port, where port = PHY port number and bit = the particular LED for the port.
NC_1 NC_2 NC_3 NC_4	A1 A16 T1 T16	NC	No connect.
PHYADD2 PHYADD3 PHYADD4	G13 G14 F14	I_{PU}	Device SMI address bits 4:2.
RCVRDCLK1 RCVRDCLK2	F16 G16	O	Clock output can be enabled or disabled and also output a clock frequency of 125 MHz or 25 MHz based on the selected active recovered media programmed for this pin. This pin is not active when NRESET is asserted. When disabled, the pin is held low.
REF_FILT_A	D3	A_{BIAS}	Reference filter connects to an external 1 µF capacitor to analog ground.
REF_REXT_A	D4	A_{BIAS}	Reference external connects to an external 2 kΩ (1%) resistor to analog ground.
REFCLK_N REFCLK_P	C1 D1	I	125 MHz or 25 MHz reference clock input pair.
REFCLK_SEL2	E1	I_{PD}	Selects the reference clock speed: 0: 25 MHz (VSS) 1: 125 MHz (2.5 V) Use 125 MHz for typical applications.

Table 154. Miscellaneous Pins *(continued)*

Name	Pin	Type	Description
RESERVED_1 RESERVED_2 RESERVED_3 RESERVED_4 RESERVED_5 RESERVED_6 RESERVED_7 RESERVED_8	C10 D13 L4 P4 L1 M1 N1 P1	NC	Leave these pins unconnected (floating).
THERMDA	C3	A	Thermal diode anode.
THERMDC_VSS	D2	A	Thermal diode cathode.

6.3.5 Power Supply

The following table lists the power supply pins and associated functional pins. All power supply pins must be connected to their respective voltage input, even if certain functions are not used for a specific application.

Table 155. Power Supply Pins

Name	Pin	Type	Description
VDD1_[1:18]	E5, E12, F5, F12, G5, G12, H5, H12, J5, J12, K5, K12, L5, L12, M5, M12, N5, N12	Power	1.0 V internal digital logic.
VDD1A_[1:10]	C7, C8, C9, C13, P6, P7, P8, P9, P10, P11	Power	1.0 V analog power requiring additional PCB power supply filtering. Associated with the RDP/N_n and TDP/N_n pins.
VDD25_[1:4]	H13, M4, M13, P3	Power	2.5 V general digital power supply. Associated with the following pins: LED[3:0]_n, GPIO[9:0], JTAG pins, TXVP[A:D]_n, TXVN[A:D]_n, REF_FILT_A, REF_REXT_A, and RCVRDCLK[2:1].
VDD25A_[1:10]	C2, C4, C6, C11, C14, E4, E13, P5, P12, P13	Power	2.5 V general analog power supply.
VSS_[1:4] VSS_[6:78]	B1, B16, C5, C12 D5, D6, D7, D8, D9, D10, D11, D12, E6, E7, E8, E9, E10, E11, F4, F6, F7, F8, F9, F10, F11, G6, G7, G8, G9, G10, G11, H6, H7, H8, H9, H10, H11, J6, J7, J8, J9, J10, J11, K6, K7, K8, K9, K10, K11, L6, L7, L8, L9, L10, L11, M6, M7, M8, M9, M10, M11, N6, N7, N8, N9, N10, N11, R1, R16, D14, E14, F13, F15, G15, H15, H16, J13	Power	General device ground.

6.3.6 SerDes MAC Interface

The following table lists the SerDes MAC interface pins.

Table 156. SerDes MAC Interface Pins

Name	Pin	Type	Description
RDN_0 RDP_0	P16 P15	A_{DIFF}	PHY0 SGMII or QSGMII MAC receiver output pair.
RDN_1 RDN_2 RDN_3 RDP_1 RDP_2 RDP_3	T12 T8 T4 R12 R8 R4	O_{DIFF}	SerDes MAC receiver output pair.
SerDes_Rext_0	P14	Abias	SerDes bias pins. Connect to SerDes_Rext_1 with a 620 Ω 1% resistor.
SerDes_Rext_1	N13	Abias	SerDes bias pins. Connect to SerDes_Rext_0 with a 620 Ω 1% resistor.
TDN_0 TDP_0	N16 N15	A_{DIFF}	PHY0 SGMII or QSGMII MAC receiver output pair.
TDN_1 TDN_2 TDN_3 TDP_1 TDP_2 TDP_3	T13 T9 T5 R13 R9 R5	O_{DIFF}	SerDes MAC receiver output pair.

6.3.7 SerDes Media Interface

The following table lists the SerDes media interface pins.

Table 157. SerDes Media Interface Pins

Name	Pin	Type	Description
FIBRIN_0 FIBRIN_1 FIBRIN_2 FIBRIN_3 FIBRIP_0 FIBRIP_1 FIBRIP_2 FIBRIP_3	T15 T11 T7 T3 R15 R11 R7 R3	I_{DIFF}	SerDes media receiver input pair.
FIBRON_0 FIBRON_1 FIBRON_2 FIBRON_3 FIBROP_0 FIBROP_1 FIBROP_2 FIBROP_3	T14 T10 T6 T2 R14 R10 R6 R2	O_{DIFF}	SerDes media transmitter output pair.

6.3.8 Serial Management Interface

The following table lists the serial management interface (SMI) pins. The SMI pins are referenced to VDD25 and can be set to a 2.5 V power supply.

Table 158. SMI Pins

Name	Pin	Type	Description
MDC	P2	I_{PU}	Management data clock. A 0 MHz to 12.5 MHz reference input is used to clock serial MDIO data into and out of the PHY.
MDINT	M2	OS/OD	Management interrupt signal. Upon reset the device will configure these pins as active-low (open drain) or active-high (open source) based on the polarity of an external 10 kΩ resistor connection. These pins can be tied together in a wired-OR configuration with only a single pull-up or pull-down resistor.
MDIO	N2	OD	Management data input/output pin. Serial data is written or read from this pin bidirectionally between the PHY and Station Manager, synchronously on the positive edge of MDC. One external pull-up resistor is required at the Station Manager, and its value depends on the MDC clock frequency and the total sum of the capacitive loads from the MDIO pins.
NRESET	M3	I_{PU}	Device reset. Active low input that powers down the device and sets all register bits to their default state.

6.4 Twisted Pair Interface

The following table lists the twisted pair interface pins.

Table 159. Twisted Pair Interface Pins

Name	Pin	Type	Description
TXVNA_0 TXVNA_1 TXVNA_2 TXVNA_3	B14 B10 B6 B2	A_{DIFF}	TX/RX channel A negative signal
TXVNB_0 TXVNB_1 TXVNB_2 TXVNB_3	B15 B11 B7 B3	A_{DIFF}	TX/RX channel B negative signal
TXVNC_0 TXVNC_1 TXVNC_2 TXVNC_3	C15 B12 B8 B4	A_{DIFF}	TX/RX channel C negative signal
TXVND_0 TXVND_1 TXVND_2 TXVND_3	D15 B13 B9 B5	A_{DIFF}	TX/RX channel D negative signal
TXVPA_0 TXVPA_1 TXVPA_2 TXVPA_3	A14 A10 A6 A2	A_{DIFF}	TX/RX channel A positive signal

Table 159. Twisted Pair Interface Pins *(continued)*

Name	Pin	Type	Description
TXVPB_0	A15	A_{DIFF}	TX/RX channel B positive signal
TXVPB_1	A11		
TXVPB_2	A7		
TXVPB_3	A3		
TXVPC_0	C16	A_{DIFF}	TX/RX channel C positive signal
TXVPC_1	A12		
TXVPC_2	A8		
TXVPC_3	A4		
TXVPD_0	D16	A_{DIFF}	TX/RX channel D positive signal
TXVPD_1	A13		
TXVPD_2	A9		
TXVPD_3	A5		

6.5 Pins by Number

This section provides a numeric list of the VSC8574 pins.

A1	NC_1	C7	VDD1A_1	E13	VDD25A_7
A2	TXVPA_3	C8	VDD1A_2	E14	VSS_72
A3	TXVPB_3	C9	VDD1A_3	E15	CLK_SQUELCH_IN
A4	TXVPC_3	C10	RESERVED_1	E16	1588_SPI_CLK
A5	TXVPD_3	C11	VDD25A_4	F1	TDO
A6	TXVPA_2	C12	VSS_4	F2	TDI
A7	TXVPB_2	C13	VDD1A_4	F3	TCK
A8	TXVPC_2	C14	VDD25A_5	F4	VSS_20
A9	TXVPD_2	C15	TXVNC_0	F5	VDD1_3
A10	TXVPA_1	C16	TXVPC_0	F6	VSS_21
A11	TXVPB_1	D1	REFCLK_P	F7	VSS_22
A12	TXVPC_1	D2	THERMDC_VSS	F8	VSS_23
A13	TXVPD_1	D3	REF_FILT_A	F9	VSS_24
A14	TXVPA_0	D4	REF_REXT_A	F10	VSS_25
A15	TXVPB_0	D5	VSS_6	F11	VSS_26
A16	NC_2	D6	VSS_7	F12	VDD1_4
B1	VSS_1	D7	VSS_8	F13	VSS_73
B2	TXVNA_3	D8	VSS_9	F14	PHYADD4
B3	TXVNB_3	D9	VSS_10	F15	VSS_74
B4	TXVNC_3	D10	VSS_11	F16	RCVRDCLK1
B5	TXVND_3	D11	VSS_12	G1	LED0_0
B6	TXVNA_2	D12	VSS_13	G2	LED1_0
B7	TXVNB_2	D13	RESERVED_2	G3	LED2_0
B8	TXVNC_2	D14	VSS_71	G4	LED3_0
B9	TXVND_2	D15	TXVND_0	G5	VDD1_5
B10	TXVNA_1	D16	TXVPD_0	G6	VSS_27
B11	TXVNB_1	E1	REFCLK_SEL2	G7	VSS_28
B12	TXVNC_1	E2	TMS	G8	VSS_29
B13	TXVND_1	E3	TRST	G9	VSS_30
B14	TXVNA_0	E4	VDD25A_6	G10	VSS_31
B15	TXVNB_0	E5	VDD1_1	G11	VSS_32
B16	VSS_2	E6	VSS_14	G12	VDD1_6
C1	REFCLK_N	E7	VSS_15	G13	PHYADD2
C2	VDD25A_1	E8	VSS_16	G14	PHYADD3
C3	THERMDA	E9	VSS_17	G15	VSS_75
C4	VDD25A_2	E10	VSS_18	G16	RCVRDCLK2
C5	VSS_3	E11	VSS_19	H1	LED0_1
C6	VDD25A_3	E12	VDD1_2	H2	LED1_1

Pins by number *(continued)*

H3	LED2_1	K12	VDD1_12	N5	VDD1_17
H4	LED3_1	K13	GPIO8/I2C_SDA	N6	VSS_63
H5	VDD1_7	K14	GPIO9/FASTLINK-FAIL	N7	VSS_64
H6	VSS_33	K15	1588_LOAD_SAVE	N8	VSS_65
H7	VSS_34	K16	1588_PPS_0	N9	VSS_66
H8	VSS_35	L1	RESERVED_5	N10	VSS_67
H9	VSS_36	L2	1588_PPS_3	N11	VSS_68
H10	VSS_37	L3	COMA_MODE	N12	VDD1_18
H11	VSS_38	L4	RESERVED_3	N13	SerDes_Rext_1
H12	VDD1_8	L5	VDD1_13	N14	GPIO0/SIGDET0
H13	VDD25_1	L6	VSS_51	N15	TDP_0
H14	1588_SPI_DO	L7	VSS_52	N16	TDN_0
H15	VSS_76	L8	VSS_53	P1	RESERVED_8
H16	VSS_77	L9	VSS_54	P2	MDC
J1	LED0_2	L10	VSS_55	P3	VDD25_4
J2	LED1_2	L11	VSS_56	P4	RESERVED_4
J3	LED2_2	L12	VDD1_14	P5	VDD25A_8
J4	LED3_2	L13	GPIO4/I2C_SCL_0	P6	VDD1A_5
J5	VDD1_9	L14	GPIO5/I2C_SCL_1	P7	VDD1A_6
J6	VSS_39	L15	GPIO6/I2C_SCL_2	P8	VDD1A_7
J7	VSS_40	L16	GPIO7/I2C_SCL_3	P9	VDD1A_8
J8	VSS_41	M1	RESERVED_6	P10	VDD1A_9
J9	VSS_42	M2	MDINT	P11	VDD1A_10
J10	VSS_43	M3	NRESET	P12	VDD25A_9
J11	VSS_44	M4	VDD25_2	P13	VDD25A_10
J12	VDD1_10	M5	VDD1_15	P14	SerDes_Rext_0
J13	VSS_78	M6	VSS_57	P15	RDP_0
J14	1588_SPI_CS	M7	VSS_58	P16	RDN_0
J15	1588_DIFF_INPUT_CLK_P	M8	VSS_59	R1	VSS_69
J16	1588_DIFF_INPUT_CLK_N	M9	VSS_60	R2	FIBROP_3
K1	LED0_3	M10	VSS_61	R3	FIBRIP_3
K2	LED1_3	M11	VSS_62	R4	RDP_3
K3	LED2_3	M12	VDD1_16	R5	TDP_3
K4	LED3_3	M13	VDD25_3	R6	FIBROP_2
K5	VDD1_11	M14	GPIO1/SIGDET1	R7	FIBRIP_2
K6	VSS_45	M15	GPIO2/SIGDET2	R8	RDP_2
K7	VSS_46	M16	GPIO3/SIGDET3	R9	TDP_2
K8	VSS_47	N1	RESERVED_7	R10	FIBROP_1
K9	VSS_48	N2	MDIO	R11	FIBRIP_1
K10	VSS_49	N3	1588_PPS_1	R12	RDP_1
K11	VSS_50	N4	1588_PPS_2	R13	TDP_1

Pins by number *(continued)*

R14	FIBROP_0
R15	FIBRIP_0
R16	VSS_70
T1	NC_3
T2	FIBRON_3
T3	FIBRIN_3
T4	RDN_3
T5	TDN_3
T6	FIBRON_2
T7	FIBRIN_2
T8	RDN_2
T9	TDN_2
T10	FIBRON_1
T11	FIBRIN_1
T12	RDN_1
T13	TDN_1
T14	FIBRON_0
T15	FIBRIN_0
T16	NC_4

6.6 Pins by Name

This section provides an alphabetical list of the VSC8574 pins.

1588_DIFF_INPUT_CLK_N	J16	LED0_0	G1	REF_REXT_A	D4
1588_DIFF_INPUT_CLK_P	J15	LED0_1	H1	REFCLK_N	C1
1588_LOAD_SAVE	K15	LED0_2	J1	REFCLK_P	D1
1588_PPS_0	K16	LED0_3	K1	REFCLK_SEL2	E1
1588_PPS_1	N3	LED1_0	G2	RESERVED_1	C10
1588_PPS_2	N4	LED1_1	H2	RESERVED_2	D13
1588_PPS_3	L2	LED1_2	J2	RESERVED_3	L4
1588_SPI_CLK	E16	LED1_3	K2	RESERVED_4	P4
1588_SPI_CS	J14	LED2_0	G3	RESERVED_5	L1
1588_SPI_DO	H14	LED2_1	H3	RESERVED_6	M1
CLK_SQUELCH_IN	E15	LED2_2	J3	RESERVED_7	N1
COMA_MODE	L3	LED2_3	K3	RESERVED_8	P1
FIBRIN_0	T15	LED3_0	G4	SerDes_Rext_0	P14
FIBRIN_1	T11	LED3_1	H4	SerDes_Rext_1	N13
FIBRIN_2	T7	LED3_2	J4	TCK	F3
FIBRIN_3	T3	LED3_3	K4	TDI	F2
FIBRIP_0	R15	MDC	P2	TDN_0	N16
FIBRIP_1	R11	MDINT	M2	TDN_1	T13
FIBRIP_2	R7	MDIO	N2	TDN_2	T9
FIBRIP_3	R3	NC_1	A1	TDN_3	T5
FIBRON_0	T14	NC_2	A16	TDO	F1
FIBRON_1	T10	NC_3	T1	TDP_0	N15
FIBRON_2	T6	NC_4	T16	TDP_1	R13
FIBRON_3	T2	NRESET	M3	TDP_2	R9
FIBROP_0	R14	PHYADD2	G13	TDP_3	R5
FIBROP_1	R10	PHYADD3	G14	THERMDA	C3
FIBROP_2	R6	PHYADD4	F14	THERMDC_VSS	D2
FIBROP_3	R2	RCVRDCLK1	F16	TMS	E2
GPIO0/SIGDET0	N14	RCVRDCLK2	G16	TRST	E3
GPIO1/SIGDET1	M14	RDN_0	P16	TXVNA_0	B14
GPIO2/SIGDET2	M15	RDN_1	T12	TXVNA_1	B10
GPIO3/SIGDET3	M16	RDN_2	T8	TXVNA_2	B6
GPIO4/I2C_SCL_0	L13	RDN_3	T4	TXVNA_3	B2
GPIO5/I2C_SCL_1	L14	RDP_0	P15	TXVNB_0	B15
GPIO6/I2C_SCL_2	L15	RDP_1	R12	TXVNB_1	B11
GPIO7/I2C_SCL_3	L16	RDP_2	R8	TXVNB_2	B7
GPIO8/I2C_SDA	K13	RDP_3	R4	TXVNB_3	B3
GPIO9/FASTLINK-FAIL	K14	REF_FILT_A	D3	TXVNC_0	C15

Pins by name *(continued)*

TXVNC_1	B12	VDD1A_1	C7	VSS_19	E11
TXVNC_2	B8	VDD1A_2	C8	VSS_20	F4
TXVNC_3	B4	VDD1A_3	C9	VSS_21	F6
TXVND_0	D15	VDD1A_4	C13	VSS_22	F7
TXVND_1	B13	VDD1A_5	P6	VSS_23	F8
TXVND_2	B9	VDD1A_6	P7	VSS_24	F9
TXVND_3	B5	VDD1A_7	P8	VSS_25	F10
TXVPA_0	A14	VDD1A_8	P9	VSS_26	F11
TXVPA_1	A10	VDD1A_9	P10	VSS_27	G6
TXVPA_2	A6	VDD1A_10	P11	VSS_28	G7
TXVPA_3	A2	VDD25_1	H13	VSS_29	G8
TXVPB_0	A15	VDD25_2	M4	VSS_30	G9
TXVPB_1	A11	VDD25_3	M13	VSS_31	G10
TXVPB_2	A7	VDD25_4	P3	VSS_32	G11
TXVPB_3	A3	VDD25A_1	C2	VSS_33	H6
TXVPC_0	C16	VDD25A_2	C4	VSS_34	H7
TXVPC_1	A12	VDD25A_3	C6	VSS_35	H8
TXVPC_2	A8	VDD25A_4	C11	VSS_36	H9
TXVPC_3	A4	VDD25A_5	C14	VSS_37	H10
TXVPD_0	D16	VDD25A_6	E4	VSS_38	H11
TXVPD_1	A13	VDD25A_7	E13	VSS_39	J6
TXVPD_2	A9	VDD25A_8	P5	VSS_40	J7
TXVPD_3	A5	VDD25A_9	P12	VSS_41	J8
VDD1_1	E5	VDD25A_10	P13	VSS_42	J9
VDD1_2	E12	VSS_1	B1	VSS_43	J10
VDD1_3	F5	VSS_2	B16	VSS_44	J11
VDD1_4	F12	VSS_3	C5	VSS_45	K6
VDD1_5	G5	VSS_4	C12	VSS_46	K7
VDD1_6	G12	VSS_6	D5	VSS_47	K8
VDD1_7	H5	VSS_7	D6	VSS_48	K9
VDD1_8	H12	VSS_8	D7	VSS_49	K10
VDD1_9	J5	VSS_9	D8	VSS_50	K11
VDD1_10	J12	VSS_10	D9	VSS_51	L6
VDD1_11	K5	VSS_11	D10	VSS_52	L7
VDD1_12	K12	VSS_12	D11	VSS_53	L8
VDD1_13	L5	VSS_13	D12	VSS_54	L9
VDD1_14	L12	VSS_14	E6	VSS_55	L10
VDD1_15	M5	VSS_15	E7	VSS_56	L11
VDD1_16	M12	VSS_16	E8	VSS_57	M6
VDD1_17	N5	VSS_17	E9	VSS_58	M7
VDD1_18	N12	VSS_18	E10	VSS_59	M8

Pins by name *(continued)*

VSS_60	M9
VSS_61	M10
VSS_62	M11
VSS_63	N6
VSS_64	N7
VSS_65	N8
VSS_66	N9
VSS_67	N10
VSS_68	N11
VSS_69	R1
VSS_70	R16
VSS_71	D14
VSS_72	E14
VSS_73	F13
VSS_74	F15
VSS_75	G15
VSS_76	H15
VSS_77	H16
VSS_78	J13

7 Package Information

The VSC8574 device is offered with two operating temperature ranges. The range for VSC8574 is 0 °C ambient to 125 °C junction, and the range for VSC8574-03 is –40 °C ambient to 125 °C junction.

The VSC8574XKS and VSC8574XKS-03 are packaged in a lead(Pb)-free, 256-pin, plastic ball grid array (BGA) with a 17 mm × 17 mm body size, 1 mm pin pitch, and 1.8 mm maximum height.

Lead(Pb)-free products from Vitesse comply with the temperatures and profiles defined in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

This section provides the package drawing, thermal specifications, and moisture sensitivity rating for the VSC8574 device.

7.1 Package Drawing

The following illustration shows the package drawing for the VSC8574 device. The drawing contains the top view, bottom view, side view, dimensions, tolerances, and notes.

Figure 75. Package Drawing

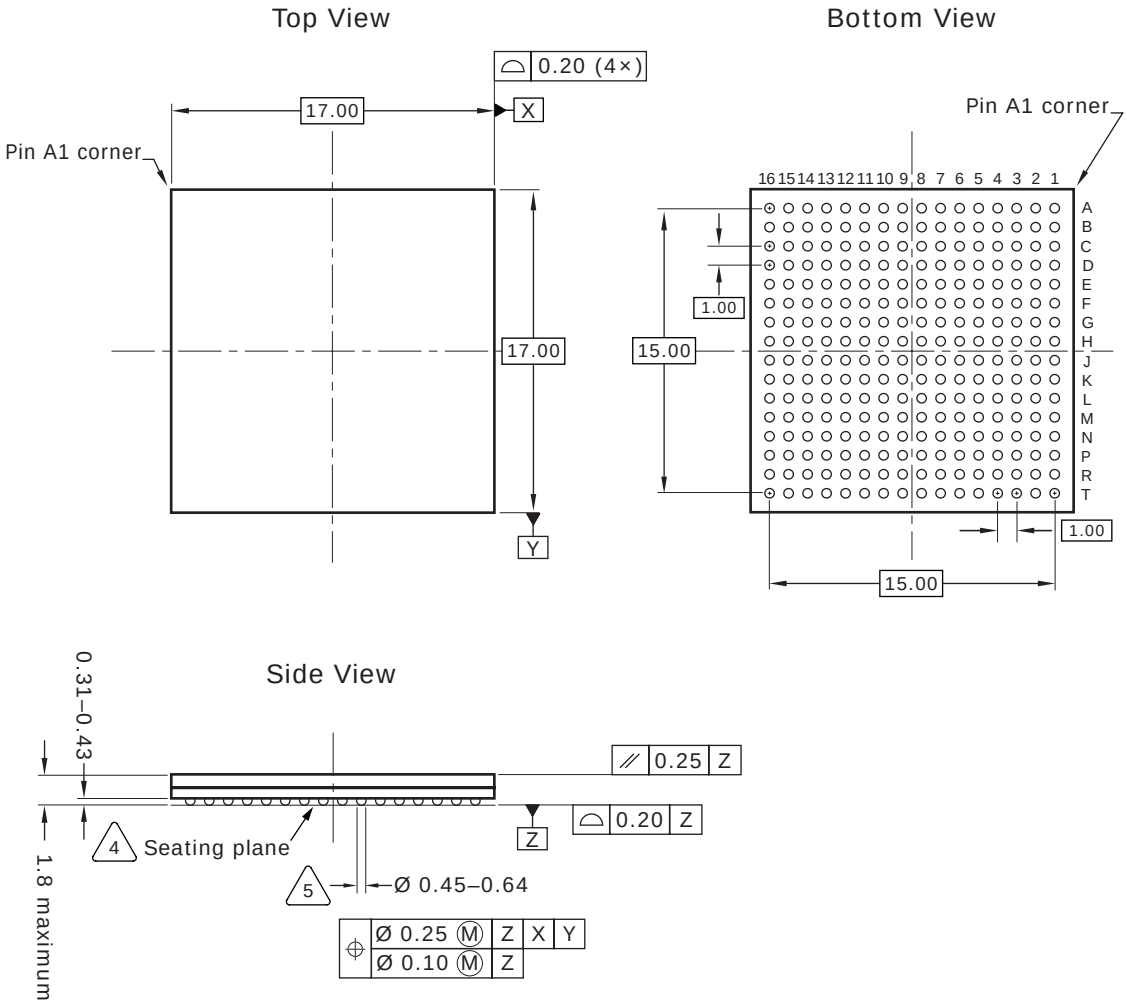

Notes

1. All dimensions and tolerances are in millimeters (mm).
2. Ball diameter is 0.50 mm.
3. Radial true position is represented by typical values.
4. Primary datum Z and seating plane are defined by the spherical crowns of the solder balls.
5. Dimension is measured at the maximum solder ball diameter, parallel to primary datum Z.

7.2 Thermal Specifications

Thermal specifications for this device are based on the JEDEC JESD51 family of documents. These documents are available on the JEDEC Web site at www.jedec.org. The thermal specifications are modeled using a four-layer test board with two signal layers, a power plane, and a ground plane (2s2p PCB). For more information about the thermal measurement method used for this device, see the JESD51-1 standard.

Table 160. Thermal Resistances

Symbol	°C/W	Parameter
θ_{JCtop}	5.9	Die junction to package case top
θ_{JB}	12.7	Die junction to printed circuit board
θ_{JA}	22	Die junction to ambient
θ_{JMA} at 1 m/s	18.5	Die junction to moving air measured at an air speed of 1 m/s
θ_{JMA} at 2 m/s	16.3	Die junction to moving air measured at an air speed of 2 m/s

To achieve results similar to the modeled thermal measurements, the guidelines for board design described in the JESD51 family of publications must be applied. For information about applications using BGA packages, see the following:

- JESD51-2A, *Integrated Circuits Thermal Test Method Environmental Conditions, Natural Convection (Still Air)*
- JESD51-6, *Integrated Circuit Thermal Test Method Environmental Conditions, Forced Convection (Moving Air)*
- JESD51-8, *Integrated Circuit Thermal Test Method Environmental Conditions, Junction-to-Board*
- JESD51-9, *Test Boards for Area Array Surface Mount Package Thermal Measurements*

7.3 Moisture Sensitivity

Moisture sensitivity level ratings for Vitesse products comply with the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. All Vitesse products are rated moisture sensitivity level 4 or better unless specified otherwise. For more information, see the IPC and JEDEC standard.

8 Ordering Information

The VSC8574 device is offered with two operating temperature ranges. The range for VSC8574 is 0 °C ambient to 125 °C junction, and the range for VSC8574-03 is –40 °C ambient to 125 °C junction.

The VSC8574XKS and VSC8574XKS-03 are packaged in a lead(Pb)-free, 256-pin, plastic ball grid array (BGA) with a 17 mm × 17 mm body size, 1 mm pin pitch, and 1.8 mm maximum height.

Lead(Pb)-free products from Vitesse comply with the temperatures and profiles defined in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

The following table lists the ordering information for the VSC8574 device.

Table 161. Ordering Information

Part Order Number	Description
VSC8574XKS	Lead(Pb)-free, 256-pin, plastic BGA with a 17 mm × 17 mm body size, 1 mm pin pitch, and 1.8 mm maximum height. The operating temperature is 0 °C ambient to 125 °C junction.
VSC8574XKS-03	Lead(Pb)-free, 256-pin, plastic BGA with a 17 mm × 17 mm body size, 1 mm pin pitch, and 1.8 mm maximum height. The operating temperature is –40 °C ambient to 125 °C junction.