

VSC8584-11 Datasheet
Quad-Port 10/100/1000BASE-T PHY with Synchronous
Ethernet, VeriTime™, Intellisec™, and QSGMII/SGMII
MAC



a  MICROCHIP company



a  MICROCHIP company

Microsemi Headquarters

One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113

Outside the USA: +1 (949) 380-6100

Sales: +1 (949) 380-6136

Fax: +1 (949) 215-4996

Email: sales.support@microsemi.com

www.microsemi.com

©2018 Microsemi, a wholly owned subsidiary of Microchip Technology Inc. All rights reserved. Microsemi and the Microsemi logo are registered trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

About Microsemi

Microsemi, a wholly owned subsidiary of Microchip Technology Inc. (Nasdaq: MCHP), offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions, security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Learn more at www.microsemi.com.

Contents

1	Revision History	1
1.1	Revision 4.2	1
1.2	Revision 4.1	1
1.3	Revision 4.0	1
1.4	Revision 2.0	2
2	Overview	3
2.1	Key Features	4
2.1.1	Low Power	4
2.1.2	Advanced Carrier Ethernet Support	4
2.1.3	Wide Range of Support	4
2.1.4	Flexibility	4
2.1.5	IEEE 1588v2	5
2.1.6	MACsec Encryption	5
2.2	Block Diagram	5
3	Functional Descriptions	7
3.1	Operating Modes	7
3.1.1	QSGMII/SGMII MAC to 1000BASE-X Link Partner	7
3.1.2	QSGMII/SGMII MAC to 100BASE-FX Link Partner	8
3.1.3	QSGMII/SGMII MAC to AMS and 1000BASE-X Media SerDes	9
3.1.4	QSGMII/SGMII MAC to AMS and 100BASE-FX Media SerDes	9
3.1.5	QSGMII/SGMII MAC-to-AMS and Protocol Transfer Mode	10
3.1.6	QSGMII/SGMII MAC to Cat5 Link Partner	11
3.1.7	QSGMII/SGMII MAC-to-Protocol Transfer Mode	11
3.1.8	1000BASE-X MAC to Cat5 Link Partner	12
3.2	SerDes MAC Interface	12
3.2.1	1000BASE-X MAC	12
3.2.2	SGMII MAC	13
3.2.3	QSGMII MAC	13
3.3	SerDes Media Interface	14
3.3.1	QSGMII/SGMII to 1000BASE-X	14
3.3.2	QSGMII/SGMII to 100BASE-FX	14
3.3.3	QSGMII to SGMII Protocol Conversion	14
3.3.4	Unidirectional Transport for Fiber Media	15
3.4	PHY Addressing and Port Mapping	15
3.4.1	PHY Addressing	15
3.4.2	SerDes Port Mapping	15
3.5	Cat5 Twisted Pair Media Interface	15
3.5.1	Voltage Mode Line Driver	15
3.5.2	Cat5 Autonegotiation and Parallel Detection	16
3.5.3	Automatic Crossover and Polarity Detection	16
3.5.4	Manual MDI/MDIX Setting	17
3.5.5	Link Speed Downshift	17
3.5.6	Energy Efficient Ethernet	17
3.5.7	Ring Resiliency	18
3.6	MACsec Block Operation	18
3.6.1	MACsec Architecture	18
3.6.2	MACsec Target Applications	21
3.6.3	Formats, Transforms, and Classification	23
3.6.4	MACsec Integration in PHY	25

3.6.5	MACsec Pipeline Operation	26
3.6.6	Debug Fault Code in FCS	46
3.6.7	Capture FIFO	48
3.6.8	Flow Control Buffer	50
3.6.9	Media Access Control	53
3.7	Automatic Media Sense Interface Mode	56
3.8	Reference Clock	57
3.8.1	Configuring the Reference Clock	57
3.8.2	Single-Ended REFCLK Input	57
3.8.3	Differential REFCLK Input	58
3.9	IEEE 1588 Reference Clock	58
3.10	Ethernet Inline Powered Devices	58
3.11	IEEE 802.3af PoE Support	60
3.12	ActiPHY Power Management	60
3.12.1	Low Power State	61
3.12.2	Link Partner Wake-Up State	61
3.12.3	Normal Operating State	61
3.13	IEEE 1588 Block Operation	61
3.13.1	IEEE 1588 Block	62
3.13.2	IEEE 1588 One-Step E2E TC in Systems	64
3.13.3	IEEE 1588 TC and BC in Systems	64
3.13.4	Enhancing IEEE 1588 Accuracy for CE Switches and MACs	65
3.13.5	MACsec Support	65
3.13.6	Supporting One-Step Boundary Clock/Ordinary Clock	65
3.13.7	Supporting Two- Step Boundary/Ordinary Clock	67
3.13.8	Supporting One-Step End-to-End Transparent Clock	69
3.13.9	Supporting One-Step Peer-to-Peer Transparent Clock	72
3.13.10	Supporting Two-Step Transparent Clock	76
3.13.11	Calculating OAM Delay Measurements	78
3.13.12	Supporting Y.1731 One-Way Delay Measurements	78
3.13.13	Supporting Y.1731 Two-Way Delay Measurements	80
3.13.14	Device Synchronization for IEEE 1588 Support	82
3.13.15	Time Stamp Update Block	83
3.13.16	Analyzer	86
3.13.17	Time Stamp Processor	107
3.13.18	Time Stamp FIFO	108
3.13.19	Serial Time Stamp Output Interface	109
3.13.20	Rewriter	110
3.13.21	Local Time Counter	111
3.13.22	Serial Time of Day	113
3.13.23	Programmable Offset for LTC Load Register	115
3.13.24	Adjustment of LTC counter	115
3.13.25	Pulse per Second Output	116
3.13.26	Accuracy and Resolution	117
3.13.27	Loopbacks	117
3.13.28	IEEE 1588 Register Access using SMI (MDC/MDIO)	117
3.13.29	1588_DIFF_INPUT_CLK Configuration	118
3.14	Daisy-Chained SPI Time Stamping	118
3.15	SPI I/O Register Access	118
3.16	Media Recovered Clock Outputs	120
3.16.1	Clock Selection Settings	120
3.16.2	Clock Output Squelch	120
3.17	Serial Management Interface	120
3.17.1	SMI Frames	121
3.17.2	SMI Interrupt	122
3.18	LED Interface	122

3.18.1	LED Modes	122
3.18.2	Extended LED Modes	124
3.18.3	LED Behavior	124
3.18.4	Basic Serial LED Mode	125
3.18.5	Enhanced Serial LED Mode	126
3.18.6	LED Port Swapping	126
3.19	Fast Link Failure Indication	126
3.20	Integrated Two-Wire Serial Multiplexer	126
3.20.1	Read/Write Access Using the Two-Wire Serial MUX	127
3.21	GPIO Pins	128
3.22	Testing Features	129
3.22.1	Ethernet Packet Generator	129
3.22.2	CRC Counters	129
3.22.3	Loopbacks	130
3.22.4	VeriPHY Cable Diagnostics	134
3.22.5	JTAG Boundary Scan	135
3.22.6	JTAG Instruction Codes	135
3.22.7	Boundary Scan Register Cell Order	137
3.23	100BASE-FX Far-End Fault Indication (FEFI)	137
3.23.1	100BASE-FX Halt Code Transmission and Reception	137
3.24	Configuration	137
3.24.1	Initialization	138
4	Registers	139
4.1	Register and Bit Conventions	139
4.2	IEEE 802.3 and Main Registers	140
4.2.1	Mode Control	141
4.2.2	Mode Status	142
4.2.3	Device Identification	143
4.2.4	Autonegotiation Advertisement	143
4.2.5	Link Partner Autonegotiation Capability	144
4.2.6	Autonegotiation Expansion	144
4.2.7	Transmit Autonegotiation Next Page	145
4.2.8	Autonegotiation Link Partner Next Page Receive	145
4.2.9	1000BASE-T/X Control	145
4.2.10	1000BASE-T Status	146
4.2.11	MMD Access Control Register	147
4.2.12	MMD Address or Data Register	147
4.2.13	1000BASE-T Status Extension 1	147
4.2.14	100BASE-TX/FX Status Extension	147
4.2.15	1000BASE-T Status Extension 2	148
4.2.16	Bypass Control	149
4.2.17	Error Counter 1	150
4.2.18	Error Counter 2	150
4.2.19	Error Counter 3	150
4.2.20	Extended Control and Status	150
4.2.21	Extended PHY Control Set 1	151
4.2.22	Extended PHY Control Set 2	152
4.2.23	Interrupt Mask	153
4.2.24	Interrupt Status	154
4.2.25	Device Auxiliary Control and Status	154
4.2.26	LED Mode Select	155
4.2.27	LED Behavior	156
4.2.28	Extended Page Access	157
4.3	Extended Page 1 Registers	157
4.3.1	SerDes Media Control	158
4.3.2	Cu Media CRC Good Counter	159

4.3.3	Extended Mode Control	159
4.3.4	ActiPHY Control	160
4.3.5	PoE and Miscellaneous Functionality	161
4.3.6	Ethernet Packet Generator Control 1	161
4.3.7	Ethernet Packet Generator Control 2	162
4.4	Extended Page 2 Registers	162
4.4.1	Cu PMD Transmit Control	163
4.4.2	EEE Control	165
4.4.3	Extended Chip ID, Address 18E2 (0x12)	166
4.4.4	Entropy Data, Address 19E2 (0x13)	166
4.4.5	Extended Interrupt Mask, Address 28E2 (0x1C)	166
4.4.6	Extended Interrupt Status, Address 29E2 (0x1D)	167
4.4.7	Ring Resiliency Control (0x1E)	168
4.5	Extended Page 3 Registers	168
4.5.1	MAC SerDes PCS Control	169
4.5.2	MAC SerDes PCS Status	170
4.5.3	MAC SerDes Clause 37 Advertised Ability	171
4.5.4	MAC SerDes Clause 37 Link Partner Ability	171
4.5.5	MAC SerDes Status	171
4.5.6	Media/MAC SerDes Transmit Good Packet Counter	171
4.5.7	Media/MAC SerDes Transmit CRC Error Counter	172
4.5.8	Media SerDes PCS Control	172
4.5.9	Media SerDes PCS Status	173
4.5.10	Media SerDes Clause 37 Advertised Ability	174
4.5.11	Media SerDes Clause 37 Link Partner Ability	174
4.5.12	Media SerDes Status	174
4.5.13	Media/MAC SerDes Receive CRC Good Counter	175
4.5.14	Media/MAC SerDes Receive CRC Error Counter	175
4.6	Extended Page 4 Registers	176
4.6.1	CSR Access Controls and Status	176
4.6.2	1588_PPS_0/1 Mux Control	177
4.6.3	SPI Daisy-Chain Controls and Status	177
4.7	General Purpose Registers	178
4.7.1	Reserved General Purpose Address Space	179
4.7.2	LED/SIGDET/GPIO Control	179
4.7.3	GPIO Control 2	180
4.7.4	GPIO Input	181
4.7.5	GPIO Output	182
4.7.6	GPIO Pin Configuration	182
4.7.7	Microprocessor Command	183
4.7.8	MAC Configuration and Fast Link	184
4.7.9	Two-Wire Serial MUX Control 1	184
4.7.10	Two-Wire Serial MUX Control 2	185
4.7.11	Two-Wire Serial MUX Data Read/Write	185
4.7.12	Recovered Clock 1 Control	185
4.7.13	Recovered Clock 2 Control	186
4.7.14	Enhanced LED Control	187
4.7.15	Global Interrupt Status	188
4.8	Clause 45 Registers to Support Energy Efficient Ethernet and 802.3bf	189
4.8.1	PMA/PMD Status 1	190
4.8.2	PCS Status 1	190
4.8.3	EEE Capability	190
4.8.4	EEE Wake Error Counter	190
4.8.5	EEE Advertisement	191
4.8.6	EEE Link Partner Advertisement	191

5 Electrical Specifications 193

5.1	DC Characteristics	193
5.1.1	VDD25 and VDDMDIO (2.5 V)	193
5.1.2	VDDMDIO (1.2 V)	193
5.1.3	Supply Voltage	194
5.1.4	LED and GPIO	194
5.1.5	Internal Pull-Up or Pull-Down Resistors	194
5.1.6	Reference Clock	195
5.1.7	1588 Reference Clock	195
5.1.8	SerDes Interface (SGMII)	195
5.1.9	Enhanced SerDes Interface (QSGMII)	197
5.1.10	Current Consumption	198
5.1.11	Thermal Diode	199
5.2	AC Characteristics	200
5.2.1	Reference Clock	200
5.2.2	Recovered Clock	201
5.2.3	SerDes Outputs	202
5.2.4	SerDes Driver Jitter	203
5.2.5	SerDes Inputs	203
5.2.6	SerDes Receiver Jitter Tolerance	203
5.2.7	Enhanced SerDes Interface	203
5.2.8	Basic Serial LEDs	206
5.2.9	Enhanced Serial LEDs	207
5.2.10	Serial CPU Interface (SI) for Slave Mode	207
5.2.11	JTAG Interface	209
5.2.12	Serial Management Interface	210
5.2.13	Reset Timing	211
5.2.14	IEEE 1588 Timing Specifications	212
5.2.15	Serial Timestamp Interface	212
5.2.16	Local Time Counter Load/Save Timing	212
5.2.17	Daisy-Chain SPI Timestamping Inputs	213
5.3	Operating Conditions	214
5.4	Stress Ratings	214
6	Pin Descriptions	216
6.1	Pin Identifications	216
6.2	Pin Diagram	216
6.3	Pins by Function	218
6.3.1	1588 Support	218
6.3.2	1588 Support and GPIO	219
6.3.3	GPIO and Two-Wire Serial	219
6.3.4	JTAG	220
6.3.5	Miscellaneous	220
6.3.6	Power Supply and Ground	221
6.3.7	SerDes MAC Interface	222
6.3.8	SerDes Media Interface	222
6.3.9	Serial Management Interface	223
6.3.10	SIGDET/GPIO	223
6.3.11	SPI Interface	223
6.3.12	Twisted Pair Interface	224
7	Package Information	225
7.1	Package Drawing	225
7.2	Thermal Specifications	226
7.3	Moisture Sensitivity	227
8	Design Considerations	228
8.1	Clause 45 register address post-increment	228

8.2	Clause 45, register 7.60	228
8.3	10BASE-T signal amplitude	228
8.4	10BASE-T Half-Duplex linkup after initial reset from power up	228
8.5	Link performance in 100BASE-TX and 1000BASE-T modes	228
8.6	Clause 36 PCS incompatibilities in 1000BASE-X media mode	228
8.7	1000BASE-X parallel detect mode with Clause 37 auto-negotiation enabled	229
8.8	Near-end loopback non-functional in protocol transfer mode	229
8.9	LED Duplex/Collision function not working when in protocol transfer mode 10/100 Mbps	229
8.10	Fast Link Failover indication delay when using interrupts	229
8.11	Anomalous Fast Link Failure indication in 1000BT Energy Efficient Ethernet mode	229
8.12	Anomalous PCS error indications in Energy Efficient Ethernet mode	229
8.13	EEE allowed only for MACs supporting IEEE 802.3az-2010 in 1588 applications	229
8.14	Auto-Negotiation Management Register Test failures	230
8.15	1588 bypass switch may drop packets	230
8.16	1588 bypass shall be enabled during engine reconfiguration	230
8.17	Time stamp errors due to IEEE 1588 reference clock interruption	230
8.18	Soft power-down procedure	230
8.19	Egress time stamp variation in PHY Packet Mode of operation	230
8.20	Changing speed from 1 Gbps to 100 Mbps hangs MACsec engine	230
8.21	VTSS_MACSEC_uncontrolled_counters_get show incorrect counter values	231
8.22	Controlled port counter if_in_octets does not get set correctly	231
8.23	Far-end loopback requires disabling Tx preamble fix for 1000M traffic	231
8.24	Non-standard preamble frames discarded by MACsec engine	231
8.25	Pause control frames pause_timer may be exceeded by MACsec engine	231
8.26	MACsec engine may shrink the minimum 100M IPG during wire-speed transmission	231
8.27	Operate MACsec with flow control to handle bandwidth expansion	232
8.28	Frames with nibbles after the FCS may be discarded by MACsec engine	232
8.29	10BASE-T ingress time stamp errors when MACsec enabled	232
8.30	MACsec datapath switch may cause large time stamp errors in the 1588 engine	232
8.31	Packet mode datapath while processing 1588 traffic	232
9	Ordering Information	233

Figures

Figure 1	Dual Media Application Diagram	3
Figure 2	Copper Transceiver Application Diagram	4
Figure 3	Fiber Media Transceiver Application Diagram	4
Figure 4	Block Diagram	6
Figure 5	SGMII MAC to 100BASE-X Link Partner	7
Figure 6	QSGMII MAC to 100BASE-X Link Partner	8
Figure 7	QSGMII/SGMII MAC to 100BASE-FX Link Partner	8
Figure 8	QSGMII/SGMII MAC to AMS and 100BASE-X Media SerDes	9
Figure 9	QSGMII/SGMII MAC to AMS and 100BASE-FX Media SerDes	10
Figure 10	QSGMII/SGMII MAC-to-AMS and Protocol Transfer Mode	10
Figure 11	QSGMII/SGMII MAC to Cat5 Link Partner	11
Figure 12	QSGMII/SGMII MAC to Protocol Transfer Mode	11
Figure 13	100BASE-X MAC to Cat5 Link Partner	12
Figure 14	SerDes MAC Interface	13
Figure 15	SGMII MAC Interface	13
Figure 16	QSGMII MAC Interface	14
Figure 17	Cat5 Media Interface	16
Figure 18	Low Power Idle Operation	17
Figure 19	MACsec Architecture	19
Figure 20	Secure Enterprise Infrastructure and WAN	21
Figure 21	Secure Carrier Ethernet Connection	22
Figure 22	Secure Mobile Backhaul with IEEE 1588	22
Figure 23	Untagged Ethernet	23
Figure 24	Standard MACsec Transform of Untagged Ethernet	23
Figure 25	Single-Tagged Ethernet	23
Figure 26	Standard MACsec Transform of Single-Tagged Ethernet	23
Figure 27	Dual-Tagged Ethernet	24
Figure 28	Standard MACsec Transform of Dual-Tagged Ethernet	24
Figure 29	Single-Tagged Ethernet	24
Figure 30	MACsec Transform to Single Tag Bypass	24
Figure 31	Dual-Tagged Ethernet	24
Figure 32	MACsec Transform to Single and Dual Tag Bypass	25
Figure 33	EoMPLS with One Label	25
Figure 34	Standard and Advanced MACsec Transform	25
Figure 35	EoMPLS with Two Labels	25
Figure 36	Standard and Advanced MACsec Transform	25
Figure 37	MACsec in PHY	26
Figure 38	MACsec Egress Data Flow	28
Figure 39	MACsec Ingress Data Flow	28
Figure 40	VLAN Tag Bypass Format	35
Figure 41	EoMPLS Header Bypass Format	36
Figure 42	Capture FIFO Layout	49
Figure 43	Line Back-Pressure by Remote Link Partner	50
Figure 44	Host Back-Pressure by Remote Link Partner	51
Figure 45	Advanced Flow Control Handling	52
Figure 46	MAC Block Diagram	53
Figure 47	Automatic Media Sense Block Diagram	56
Figure 48	2.5 V CMOS Single-Ended REFCLK Input Resistor Network	57
Figure 49	3.3 V CMOS Single-Ended REFCLK Input Resistor Network	57
Figure 50	5 V CMOS Single-Ended REFCLK Input Resistor Network	58
Figure 51	AC Coupling for REFCLK Input	58
Figure 52	Inline Powered Ethernet Switch Diagram	59
Figure 53	ActiPHY State Diagram	60
Figure 54	IEEE 1588 Architecture	62

Figure 55	IEEE 1588 Block Diagram	63
Figure 56	TC and BC Linecard Application	64
Figure 57	One-Step E2E BC	66
Figure 58	Two-Step E2E BC	68
Figure 59	One-Step E2E TC Mode A	70
Figure 60	One-Step E2E TC Mode B	71
Figure 61	Delay Measurements	73
Figure 62	One-Step P2P TC Mode B	76
Figure 63	Two-Step E2E TC	77
Figure 64	Y.1731 1DM PDU Format	78
Figure 65	Y.1731 One-Way Delay	79
Figure 66	Y.1731 DMM PDU Format	80
Figure 67	Y.1731 Two-Way Delay	81
Figure 68	RFC6374 DMM/DMR OAM PDU Format	82
Figure 69	Draft-bhh DMM/DMR/1DM OAM PDU Formats	82
Figure 70	PTP Packet Encapsulations	84
Figure 71	OAM Packet Encapsulations	84
Figure 72	TSU Block Diagram	85
Figure 73	Analyzer Block Diagram	86
Figure 74	Type II Ethernet Basic Frame Format	89
Figure 75	Ethernet Frame with SNAP	89
Figure 76	Ethernet Frame with VLAN Tag and SNAP	89
Figure 77	Ethernet Frame with VLAN Tags and SNAP	89
Figure 78	PBB Ethernet Frame Format (No B-Tag)	89
Figure 79	PBB Ethernet Frame Format (1 B-Tag)	89
Figure 80	MPLS Label Format	92
Figure 81	MPLS Label Stack within an Ethernet Frame	92
Figure 82	MPLS Labels and Control Word	93
Figure 83	IPv4 with UDP	95
Figure 84	IPv6 with UDP	95
Figure 85	ACH Header Format	96
Figure 86	ACH Header with Protocol ID Field	96
Figure 87	IPSec Header Format	96
Figure 88	IPv6 with UDP and IPSec	97
Figure 89	PTP Frame Layout	100
Figure 90	OAM 1DM Frame Header Format	100
Figure 91	OAM DMM Frame Header Format	101
Figure 92	OAM DMR Frame Header Format	101
Figure 93	RFC6374 DMM/DMR OAM PDU Format	101
Figure 94	G8113.1/draft-bhh DMM/DMR/1DM OAM PDU Format	102
Figure 95	Serial Time Stamp/Frame Signature Output	110
Figure 96	Preamble Reduction in Rewriter	111
Figure 97	Local Time Counter Load/Save Timing	112
Figure 98	Standard PPS and 1PPS with TOD Timing Relationship	113
Figure 99	ToD Octet Waveform	114
Figure 100	SPI Time Stamping Format	118
Figure 101	SPI Write Cycles	119
Figure 102	SPI Read Cycle	119
Figure 103	SMI Read Frame	121
Figure 104	SMI Write Frame	121
Figure 105	MDINT Configured as an Open-Drain (Active-Low) Pin	122
Figure 106	Two-Wire Serial MUX with SFP Control and Status	127
Figure 107	Two-Wire Serial MUX Read and Write Register Flow	128
Figure 108	Far-End Loopback Diagram	130
Figure 109	Near-End Loopback Diagram	131
Figure 110	Connector Loopback Diagram	131
Figure 111	Data Loops of the SerDes Macro	132
Figure 112	Test Access Port and Boundary Scan Architecture	135
Figure 113	Register Space Diagram	139

Figure 114	SGMII DC Transmit Test Circuit	196
Figure 115	SGMII DC Definitions	196
Figure 116	SGMII DC Driver Output Impedance Test Circuit	196
Figure 117	SGMII DC Input Definitions	197
Figure 118	Thermal Diode	200
Figure 119	Test Circuit for Recovered Clock Output Signals	201
Figure 120	QSGMII Transient Parameters	204
Figure 121	Basic Serial LED Timing	206
Figure 122	Enhanced Serial LED Timing	207
Figure 123	SI Input Data Timing Diagram for Slave Mode	207
Figure 124	SI Output Data Timing Diagram for Slave Mode	208
Figure 125	Test Circuit for SI_DO Disable	209
Figure 126	JTAG Interface Timing Diagram	210
Figure 127	Test Circuit for TDO Disable Time	210
Figure 128	Serial Management Interface Timing	211
Figure 129	SPI Interface Timing	212
Figure 130	Local Time Counter Load/Save Timing Diagram	213
Figure 131	Daisy-Chained SPI Timestamping Input Timing Diagram	213
Figure 132	Top-Left Pin Diagram	217
Figure 133	Top-Right Pin Diagram	218
Figure 134	Package Drawing	226

Tables

Table 1	Operating Modes	7
Table 2	MAC Interface Mode Mapping	15
Table 3	Supported MDI Pair Combinations	17
Table 4	Standard MACsec Frame Combinations	23
Table 5	Advanced MACsec Frame Combinations	24
Table 6	MACsec Tag Parsing Checks	30
Table 7	Match Criteria and Maskable Bits	32
Table 8	Egress SA Flow Actions	33
Table 9	Ingress SA Flow Actions	34
Table 10	Transform Record Format (Non-XPN)	37
Table 11	Context Control Word Fields	38
Table 12	Transform Record Format (XPN)	38
Table 13	Egress SA Counters	42
Table 14	Egress Global Counters	42
Table 15	Ingress SA Counters	43
Table 16	Ingress Global Counters	43
Table 17	Egress Per-User Global Counters	43
Table 18	802.1AE Correlation	44
Table 19	Ingress Per-User Global Counters	44
Table 20	FCS Fault Codes	47
Table 21	Ingress Global Stat Event Vector Format	47
Table 22	Ingress SA Stat Event Vector Format	47
Table 23	Egress Global Stat Event Vector Format	48
Table 24	Egress SA Stat Event Vector Format	48
Table 25	AMS Media Preferences	56
Table 26	REFCLK Frequency Selection	57
Table 27	Flows Per Engine Type	87
Table 28	Ethernet Comparator: Next Protocol	88
Table 29	Comparator ID Codes	88
Table 30	Ethernet Comparator (Next Protocol)	90
Table 31	Ethernet Comparator (Flow)	90
Table 32	MPLS Comparator: Next Word	93
Table 33	MPLS Comparator: Per-Flow	93
Table 34	MPLS Range_Upper/Lower Label Map	93
Table 35	Next MPLS Comparator	94
Table 36	Next-Protocol Registers in OAM-Version of MPLS Block	94
Table 37	Comparator Field Summary	97
Table 38	IP/ACH Next-Protocol Comparison	98
Table 39	IP/ACH Comparator Flow Verification Registers	99
Table 40	PTP Comparison	102
Table 41	PTP Comparison: Common Controls	104
Table 42	PTP Comparison: Additions for OAM-Optimized Engine	104
Table 43	Frame Signature Byte Mapping	105
Table 44	Frame Signature Address Source	105
Table 45	LTC Time Load/Save Options	114
Table 46	Output Pulse Frequencies	116
Table 47	Daisy-Chain Parameters	118
Table 48	SI_ADDR Mapping	119
Table 49	LED Drive State	122
Table 50	LED Mode and Function Summary	123
Table 51	Extended LED Mode and Function Summary	124
Table 52	LED Serial Bitstream Order	125
Table 53	Register Bits for GPIO Control and Status	128
Table 54	SerDes Macro Address Map	134

Table 55	JTAG Instruction Codes	135
Table 56	IDCODE JTAG Device Identification Register Descriptions	136
Table 57	USERCODE JTAG Device Identification Register Descriptions	136
Table 58	JTAG Instruction Code IEEE Compliance	137
Table 59	IEEE 802.3 Registers	140
Table 60	Main Registers	140
Table 61	Mode Control, Address 0 (0x00)	141
Table 62	Mode Status, Address 1 (0x01)	142
Table 63	Identifier 1, Address 2 (0x02)	143
Table 64	Identifier 2, Address 3 (0x03)	143
Table 65	Device Autonegotiation Advertisement, Address 4 (0x04)	143
Table 66	Autonegotiation Link Partner Ability, Address 5 (0x05)	144
Table 67	Autonegotiation Expansion, Address 6 (0x06)	144
Table 68	Autonegotiation Next Page Transmit, Address 7 (0x07)	145
Table 69	Autonegotiation LP Next Page Receive, Address 8 (0x08)	145
Table 70	1000BASE-T/X Control, Address 9 (0x09)	145
Table 71	1000BASE-T Status, Address 10 (0x0A)	146
Table 72	MMD EEE Access, Address 13 (0x0D)	147
Table 73	MMD Address or Data Register, Address 14 (0x0E)	147
Table 74	1000BASE-T Status Extension 1, Address 15 (0x0F)	147
Table 75	100BASE-TX/FX Status Extension, Address 16 (0x10)	147
Table 76	1000BASE-T Status Extension 2, Address 17 (0x11)	148
Table 77	Bypass Control, Address 18 (0x12)	149
Table 78	Extended Control and Status, Address 19 (0x13)	150
Table 79	Extended Control and Status, Address 20 (0x14)	150
Table 80	Extended Control and Status, Address 21 (0x15)	150
Table 81	Extended Control and Status, Address 22 (0x16)	150
Table 82	Extended PHY Control 1, Address 23 (0x17)	151
Table 83	Extended PHY Control 2, Address 24 (0x18)	152
Table 84	Interrupt Mask, Address 25 (0x19)	153
Table 85	Interrupt Status, Address 26 (0x1A)	154
Table 86	Auxiliary Control and Status, Address 28 (0x1C)	154
Table 87	LED Mode Select, Address 29 (0x1D)	156
Table 88	LED Behavior, Address 30 (0x1E)	156
Table 89	Extended/GPIO Register Page Access, Address 31 (0x1F)	157
Table 90	Extended Registers Page 1 Space	158
Table 91	SerDes Media Control, Address 16E1 (0x10)	158
Table 92	Cu Media CRC Good Counter, Address 18E1 (0x12)	159
Table 93	Extended Mode Control, Address 19E1 (0x13)	159
Table 94	Extended PHY Control 3, Address 20E1 (0x14)	160
Table 95	Extended PHY Control 4, Address 23E1 (0x17)	161
Table 96	EPG Control Register 1, Address 29E1 (0x1D)	161
Table 97	EPG Control Register 2, Address 30E1 (0x1E)	162
Table 98	Extended Registers Page 2 Space	163
Table 99	Cu PMD Transmit Control, Address 16E2 (0x10)	163
Table 100	EEE Control, Address 17E2 (0x11)	165
Table 101	Extended Chip ID, Address 18E2 (0x12)	166
Table 102	Entropy Data, Address 19E2 (0x13)	166
Table 103	Extended Interrupt Mask, Address 28E2 (0x1C)	166
Table 104	Extended Interrupt Status, Address 29E2 (0x1D)	167
Table 105	Ring Resiliency, Address 30E2 (0x1E)	168
Table 106	Extended Registers Page 3 Space	168
Table 107	MAC SerDes PCS Control, Address 16E3 (0x10)	169
Table 108	MAC SerDes PCS Status, Address 17E3 (0x11)	170
Table 109	MAC SerDes CI37 Advertised Ability, Address 18E3 (0x12)	171
Table 110	MAC SerDes CI37 LP Ability, Address 19E3 (0x13)	171
Table 111	MAC SerDes Status, Address 20E3 (0x14)	171
Table 112	Media/MAC SerDes Tx Good Packet Counter, Address 21E3 (0x15)	171
Table 113	Media/MAC SerDes Tx CRC Error Counter, Address 22E3 (0x16)	172

Table 114	Media SerDes PCS Control, Address 23E3 (0x17)	172
Table 115	Media SerDes PCS Status, Address 24E3 (0x18)	173
Table 116	Media SerDes CI37 Advertised Ability, Address 25E3 (0x19)	174
Table 117	MAC SerDes CI37 LP Ability, Address 26E3 (0x1A)	174
Table 118	Media SerDes Status, Address 27E3 (0x1B)	174
Table 119	Media/MAC SerDes Receive CRC Good Counter, Address 28E3 (0x1C)	175
Table 120	Media/MAC SerDes Receive CRC Error Counter, Address 29E3 (0x1D)	175
Table 121	Extended Registers Page 4 Space	176
Table 122	CSR Access Control, Address 16E4	176
Table 123	CSR Buffer, Address 17E4	176
Table 124	1588_PPS_0 Mux Control, Address 21E4	177
Table 125	CSR Buffer, Address 18E4	177
Table 126	CSR Access Control, Address 19E4	177
Table 127	CSR Status, Address 20E4	177
Table 128	SPI Daisy-Chain Control, Address 26E4	178
Table 129	SPI Daisy-Chain Status, Address 27E4	178
Table 130	SPI Daisy-Chain Counter, Address 28E4	178
Table 131	1588 RefClk Input Buffer Control (LSW), Address 29E4	178
Table 132	1588 RefClk Input Buffer Control (MSW), Address 30E4	178
Table 133	General Purpose Registers Page Space	179
Table 134	LED/SIGDET/GPIO Control, Address 13G (0x0D)	179
Table 135	GPIO Control 2, Address 14G (0x0E)	180
Table 136	GPIO Input, Address 15G (0x0F)	181
Table 137	GPIO Output, Address 16G (0x10)	182
Table 138	GPIO Input/Output Configuration, Address 17G (0x11)	182
Table 139	Microprocessor Command Register, Address 18G	183
Table 140	MAC Configuration and Fast Link Register, Address 19G (0x13)	184
Table 141	Two-Wire Serial MUX Control 1, Address 20G (0x14)	184
Table 142	Two-Wire Serial MUX Interface Status and Control, Address 21G (0x15)	185
Table 143	Two-Wire Serial MUX Data Read/Write, Address 22G (0x16)	185
Table 144	Recovered Clock 1 Control, Address 23G (0x17)	185
Table 145	Recovered Clock 2 Control, Address 24G (0x18)	186
Table 146	Enhanced LED Control, Address 25G (0x19)	187
Table 147	Global Interrupt Status, Address 29G (0x1D)	188
Table 148	Clause 45 Registers Page Space	189
Table 149	PMA/PMD Status 1	190
Table 150	PCS Status 1, Address 3.1	190
Table 151	EEE Capability, Address 3.20	190
Table 152	EEE Wake Error Counter, Address 3.22	191
Table 153	EEE Advertisement, Address 7.60	191
Table 154	EEE Advertisement, Address 7.61	191
Table 155	802.3bf Registers	192
Table 156	VDD25 and VDDMDIO	193
Table 157	VDDMDIO	193
Table 158	Supply Voltage Specifications	194
Table 159	LED and GPIO	194
Table 160	Internal Pull-Up or Pull-Down Resistors	194
Table 161	Reference Clock DC Characteristics	195
Table 162	1588 Reference Clock DC Characteristics	195
Table 163	SerDes Driver DC Specifications	195
Table 164	SerDes Receiver DC Specifications	197
Table 165	Enhanced SerDes Driver DC Specifications	197
Table 166	Enhanced SerDes Receiver DC Specifications	198
Table 167	Current Consumption (1588 and MACsec Disabled)	198
Table 168	1588 Current Consumption	199
Table 169	MACsec Current Consumption	199
Table 170	Thermal Diode Parameters	200
Table 171	Reference Clock AC Characteristics for QSGMII 125 MHz Differential Clock	200
Table 172	Recovered Clock AC Characteristics	201

Table 173	SerDes Outputs AC Specifications	202
Table 174	SerDes Driver Jitter Characteristics	203
Table 175	SerDes Input AC Specifications	203
Table 176	SerDes Receiver Jitter Tolerance	203
Table 177	Enhanced SerDes Outputs AC Specifications, SGMII Mode	204
Table 178	Enhanced SerDes Outputs AC Specifications, QSGMII Mode	204
Table 179	Enhanced SerDes Input AC Specifications, SGMII Mode	205
Table 180	Enhanced SerDes Inputs AC Specifications, QSGMII Mode	205
Table 181	Enhanced SerDes Receiver Jitter Tolerance	206
Table 182	Basic Serial LEDs AC Characteristics	206
Table 183	Enhanced Serial LEDs AC Characteristics	207
Table 184	SI Timing Specifications for Slave Mode	208
Table 185	JTAG Interface AC Specifications	209
Table 186	Serial Management Interface AC Characteristics	210
Table 187	Reset Timing Specifications	211
Table 188	IEEE 1588 Timing Specifications AC Characteristics	212
Table 189	SPI Timing	212
Table 190	Local Time Counter Load/Save Timing Specifications	213
Table 191	Daisy-Chained SPI Timestamping Input Timing Specifications	213
Table 192	PHY Latency in IEEE 1588 Timing Bypass Mode	213
Table 193	Recommended Operating Conditions	214
Table 194	Stress Ratings	214
Table 195	Pin Type Symbol Definitions	216
Table 196	1588 Support Pins	218
Table 197	1588 Support and GPIO Pins	219
Table 198	GPIO and Two-Wire Serial Pins	219
Table 199	JTAG Pins	220
Table 200	Miscellaneous Pins	220
Table 201	Power Supply and Ground Pins	221
Table 202	SerDes MAC Interface Pins	222
Table 203	SerDes Media Interface Pins	222
Table 204	SMI Pins	223
Table 205	SIGDET/GPIO Pins	223
Table 206	SPI Interface Pins	223
Table 207	Twisted Pair Interface Pins	224
Table 208	Thermal Resistances	227
Table 209	Ordering Information	233

1 Revision History

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

1.1 Revision 4.2

Revision 4.2 was published in April 2019. The following is a summary of the changes in this document.

- The VeriPHY information was updated in the Flexibility section. For more information, see [Flexibility](#), page 4.
- VeriPHY™ Cable Diagnostics section was updated. For more information, see [VeriPHY Cable Diagnostics](#), page 134.
- Removed the Mean Square Error Noise section.
- The references to the VeriPHY were removed from the Extended Registers Page 1 Space table. For more information, see [Table 90](#), page 158.
- The VeriPHY Control 1/2/3 sections were removed.

1.2 Revision 4.1

Revision 4.1 was published in August 2017. The following is a summary of the changes in this document.

- All references to LVDS were clarified to reflect LVDS compatibility.
- All references to CLK1588P/N were clarified as 1588_DIFF_INPUT_CLK_P/N.
- All references to serial parallel interface were corrected to serial peripheral interface.
- Pin E16 name was corrected to remove GPIO functionality.
- Operating modes were updated to correctly reflect available functionality. For more information, see [Operating Modes](#), page 7.
- The Analyzer block diagram was updated. For more information, see [Figure 73](#), page 86.
- A note was added about the use of recovered clock outputs and fast link failure indication in EEE mode. For more information, see [Media Recovered Clock Outputs](#), page 120 and [Fast Link Failure Indication](#), page 126.
- The equipment loop description was updated to correctly reflect available functionality. For more information, see [Equipment Loop](#), page 133.
- EEE Control register descriptions were updated to indicate sticky bits. For more information, see [Table 100](#), page 165.
- Media/MAC SerDes transmit CRC error counter register descriptions were updated. For more information, see [Table 113](#), page 172.
- Specifications for the IEEE 1588 timing, timestamp interface, and local time counter were updated. For more information, see [Table 188](#), page 212, [Serial Timestamp Interface](#), page 212, and [Local Time Counter Load/Save Timing](#), page 212.
- Information for daisy-chained SPI timestamping was added. For more information, see [Daisy-Chained SPI Timestamping Inputs](#), page 213.
- Design considerations were updated. For more information, see [Design Considerations](#), page 228.
- Temperature specifications were added to the part ordering information. For more information, see [Table 209](#), page 233.

1.3 Revision 4.0

Revision 4.0 was published in November 2015. The following is a summary of the changes in this document.

- Current consumption tables were updated to better reflect device performance.
- The PHY Latency values in IEEE 1588 Timing Bypass Mode were updated.
- The following design considerations were added.
 - LED Duplex/Collision function not working when in protocol transfer mode 10/100 Mbps.
 - Anomalous Fast Link Failure indication in 1000BT Energy Efficient Ethernet mode.
 - Energy Efficient Ethernet allowed only for MACs supporting IEEE 802.3az-2010 in 1588 applications.

- Auto-Negotiation Management Register Test failures.
- Egress time stamp variation in PHY Packet Mode of operation.
- Changing speed from 1 Gbps to 100 Mbps hangs MACsec engine.
- VTSS_MACSEC_uncontrolled_counters_get show incorrect counter values.
- Controlled port counter if_in_octets does not get set correctly.
- Far-end loopback non-functional for bypass mode of MACsec.
- Link performance in 100BASE-TX and 1000BASE-T modes design consideration was updated.

1.4 Revision 2.0

Revision 2.0 was published in August 2015. It was the first publication of this document.

2 Overview

VSC8584-11 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 two-wire serial 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 VSC8584-11 device includes Intellisec™, Microsemi’s implementation of IEEE 802.1AE 128/256-bit MACsec protocols to meet the security requirements for protecting data traversing Ethernet LANs. It does input classification, frame encryption/decryption, performance, and latency monitoring.

The VSC8584-11 includes VeriTime™, Microsemi’s patent-pending timing technology that delivers the industry’s most accurate IEEE 1588v2 timing implementation. IEEE 1588v2 timing integrated in the VSC8584-11 device is the quickest, lowest cost method of implementing the timing accuracy to maintain existing timing-critical capabilities during the migration from TDM to packet-based architectures.

The VSC8584-11 device offers a seamless integration between IEEE 1588v2 and the MACsec engine with no loss of precision.

The VSC8584-11 device supports 1-step and 2-step PTP implementations to provide accuracies below 8 ns that greatly minimize internal system delays and variabilities for ordinary clock, boundary clock, and transparent clock applications. Complete Y.1731 OAM performance monitoring capabilities, master, slave, boundary, and transparent clock configurations, and sophisticated classifications (including UDP, IPv4, IPv6 packets and VLAN, and MPLS-TP encapsulations) are also supported.

The VSC8584-11 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 Microsemi’s EcoEthernet v2.0 PHY technology, the VSC8584-11 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.

Microsemi’s mixed signal and digital signal processing (DSP) architecture is a key operational feature of the VSC8584-11, 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 that can support up to four 100BASE-FX, 1000BASE-X fiber, and/or triple-speed copper SFPs.

The following illustrations show a high-level, general view of typical VSC8584-11 applications.

Figure 1 • Dual Media Application Diagram

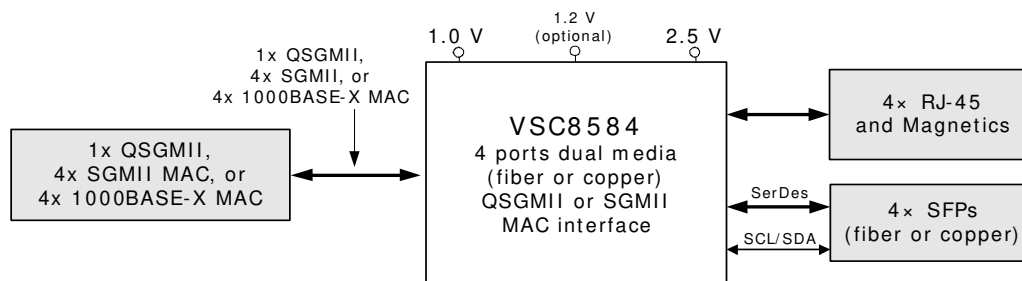




Figure 2 • Copper Transceiver Application Diagram

Figure 3 • Fiber Media Transceiver Application Diagram


2.1 Key Features

This section lists the main features and benefits of the VSC8584-11 device.

2.1.1 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 (MACsec adds 250 mW per port to the power consumption)
- ActiPHY™ link down power savings
- PerfectReach™ smart cable reach algorithm
- IEEE 802.3az-2010 Energy Efficient Ethernet idle power savings, even for legacy non-EEE systems

2.1.2 Advanced Carrier Ethernet Support

- Recovered clock outputs with programmable clock squelch control and fast link failure indication (typical <1 ms; worst-case <3 ms) for G.8261 Synchronous Ethernet applications
- Ring resiliency for maintaining linkup integrity when switching between 1000BASE-T master and slave timing
- Supports IEEE 802.3bf timing and synchronization standard
- Integrated quad two-wire serial MUX to control SFP and PoE modules
- Support for 802.3ah unidirectional transport for 100BASE-FX and 1000BASE-X fiber media

2.1.3 Wide Range of Support

- Compliant with IEEE 802.3 (10BASE-T, 10BASE-Te, 100BASE-TX, 1000BASE-T, 100BASE-FX, and 1000BASE-X) specifications
- Support for >16 kB jumbo frames in all speeds with programmable synchronization FIFOs
- Supports Cisco QSGMII v1.3, Cisco SGMII v1.9, 1000BASE-X MACs, and IEEE 1149.1 JTAG boundary scan
- Available in a low-cost, 256-pin BGA package with a 17 mm × 17 mm body size

2.1.4 Flexibility

- VeriPHY® cable diagnostics suite provides extensive network cable operating conditions and status
- 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

Note: All MAC interfaces must be the same — all QSGMII or SGMII.

2.1.5 IEEE 1588v2

- Support for IEEE 1588-2008 time stamping with encapsulation support
- Separate bypass bits for egress and ingress paths
- General PBB support for four encapsulations across three encapsulation engines
- MPLS-TP OAM support in third encapsulation engine
- ETH1 comparator bypass for time-stamping all packets
- Full 48-bit arithmetic to time-stamp
- Improved time-precision of local time counter (LTC) load with programmable offset register
- Auto-clear Load/Save signal for more deterministic software writes to LTC
- LTC block clock output based on LTC timer
- Programmable duty cycle to enable use of the synthesis pulse to synchronize out of sync devices
- Ability to load/read ToD information serially
- Improved time-precision for PPS output including external cable delay measurement for PPS
- Ability to issue interrupt when data in reserved field
- IP frame signature offset width increased to support IPv6
- IEEE-1588 v2 with MACsec support
- Ability to store frame signature from all the engines

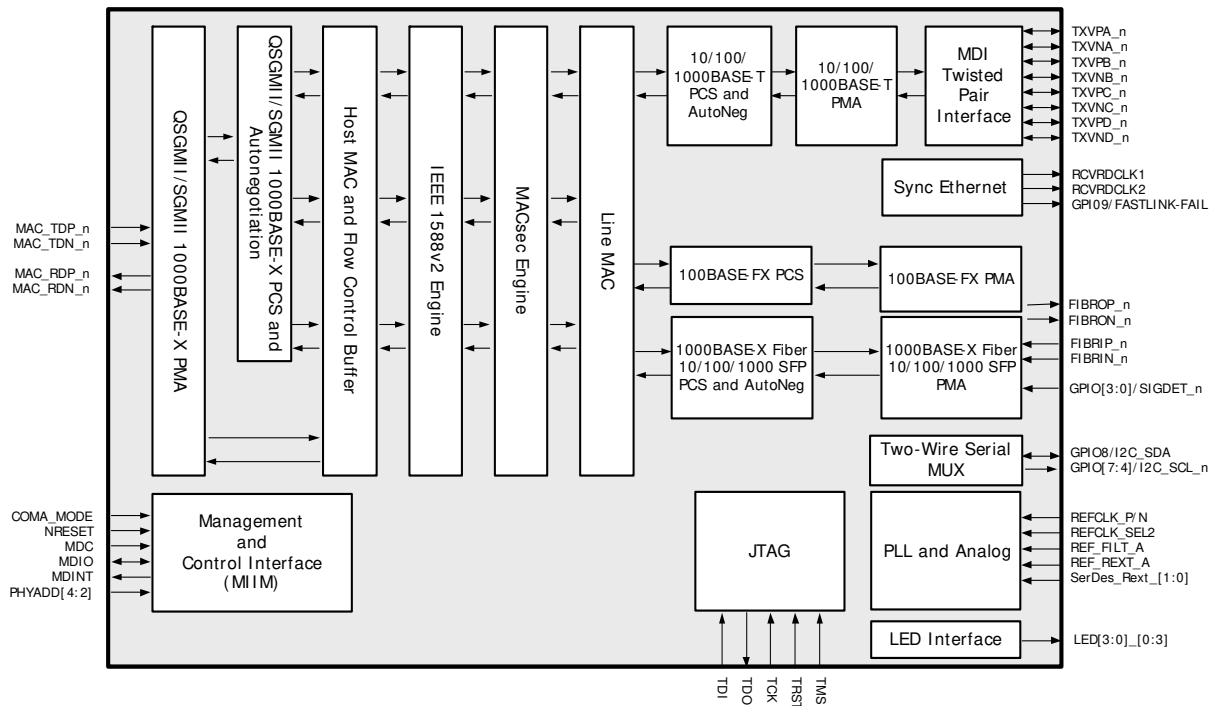
2.1.6 MACsec Encryption

- Fully IEEE 802.1AE-2006 compliant supporting GCM-AES-128, and fully IEEE 802.1AEbn-2011 compliant supporting GCM-AES-256
- Fixed latency for accurate 1588 timestamp support
- Supports full-duplex operation at all speeds
- Patent-pending architecture for minimal and predictable delays when used in conjunction with IEEE 1588
- VLAN and MPLS header bypassing
- 16 secure associations (SA) per port

2.2 Block Diagram

The following illustration shows the primary functional blocks of the VSC8584-11 device.

Figure 4 • Block Diagram



Note: All MAC interfaces must be the same—all QSGMII, SGMII, or 1000BASE-X.

3 Functional Descriptions

This section describes the functional aspects of the VSC8584-11 device, including available configurations, operational features, and testing functionality. It also defines the device setup parameters that configure the device for a particular application.

3.1 Operating Modes

The following table lists the operating modes of the VSC8584-11 device.

Table 1 • Operating Modes

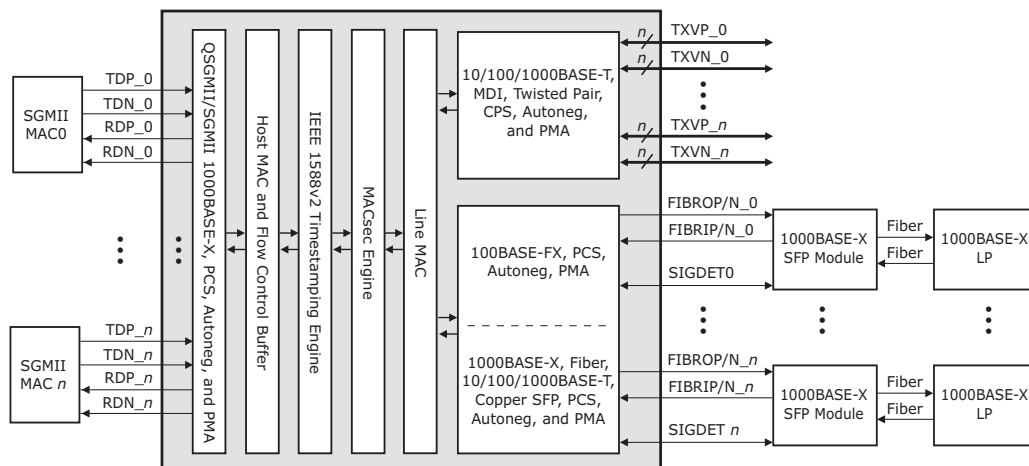
Operating Mode	Supported Media	Notes
QSGMII/SGMII MAC-to-1000BASE-X Link Partner	1000BASE-X	See Figure 5, page 7.
QSGMII/SGMII MAC-to-100BASE-FX Link Partner	100BASE-FX	See Figure 7, page 8.
QSGMII/SGMII MAC-to-AMS and 1000BASE-X SerDes	1000BASE-X, 10/100/1000BASE-T	See Figure 8, page 9.
QSGMII/SGMII MAC-to-AMS and 100BASE-FX SerDes	100BASE-FX, 10/100/1000BASE-T	See Figure 9, page 10.
QSGMII/SGMII MAC-to-AMS and Protocol Transfer Mode	SFP/Fiber protocol transfer mode (10/100/1000BASE-T Cu SFP), 10/100/1000BASE-T	See Figure 10, page 10.
QSGMII/SGMII MAC-to-Cat5 Link Partner	10/100/1000BASE-T	See Figure 11, page 11.
QSGMII/SGMII MAC-to-Protocol Transfer Mode	SFP/Fiber protocol transfer mode (10/100/1000BASE-T Cu SFP)	See Figure 12, page 11.
1000BASE-X MAC to Cat5 Link Partner	1000BASE-T only	See Figure 13, page 12.

Note: All MAC interfaces must be the same—all QSGMII or SGMII.

3.1.1 QSGMII/SGMII MAC to 1000BASE-X Link Partner

The following illustrations show the register settings used to configure a QSGMII/SGMII MAC to 1000BASE-X link partner.

Figure 5 • SGMII MAC to 1000BASE-X Link Partner

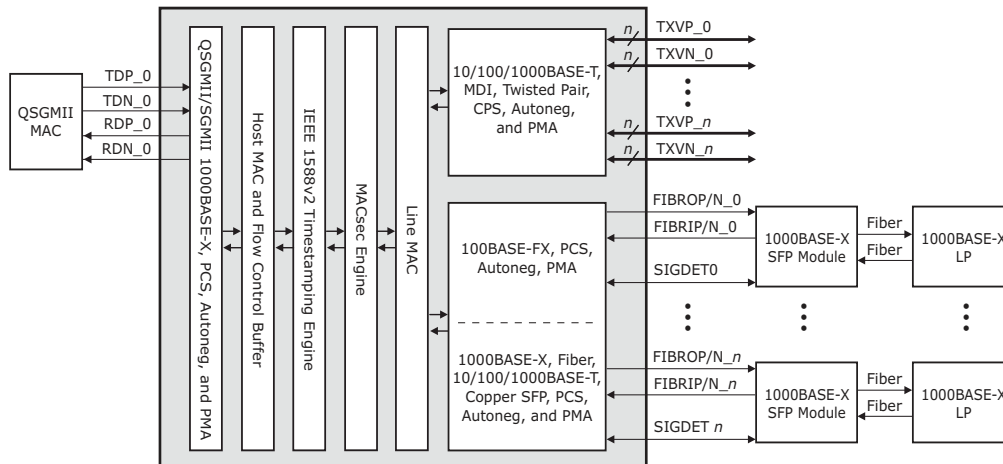


3.1.1.1 MAC Interface SGMII

- Set register 19G bits 15:14 = 00

- Set Register 23 (main register) bit 12 = 0
- Set Register 18G = 0x80F0. For more information, see Table 139, page 183.

Figure 6 • QSGMII MAC to 1000BASE-X Link Partner



3.1.1.2 MAC Interface QSGMII

- Set register 19G bits 15:14 = 01
- Set Register 23 (main register) bit 12 = 0
- Set Register 18G = 0x80E0. For more information, see Table 139, page 183.

3.1.1.3 Media Interface 1000BASE-X SFP Fiber (1000BASE-X Link Partner)

- Set register 23 bits 10:8 = 010.
- Set register 0 bit 12 = 1 (enable autonegotiation)
- Set Register 18G = 0x8FC1. For more information, see Table 139, page 183.


The F in 0x8FC1 identifies the port. To exclude a port from the configuration, set its bit to 0. For example, the configuration of port 0 and port 1 to 1000BASE-X is 0011 or 3, making the bit setting 0x83C1.

Note: Whenever there is a mode change in register 23, a software reset (register 0 bit 15) is required to make the mode change effective. This register will read the currently active mode and not what was just written.

3.1.2 QSGMII/SGMII MAC to 100BASE-FX Link Partner

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC to 100BASE-FX link partner.

Figure 7 • QSGMII/SGMII MAC to 100BASE-FX Link Partner



3.1.2.1 Media Interface 100BASE-FX SFP Fiber (100BASE-FX Link Partner)

- Set register 23 bits 10:8 = 011
- Set register 0 bit 12 = 0 (autonegotiation not present in 100BASE-FX PHY)
- Set Register 18G = 0x8FD1. For more information, Table 139, page 183.

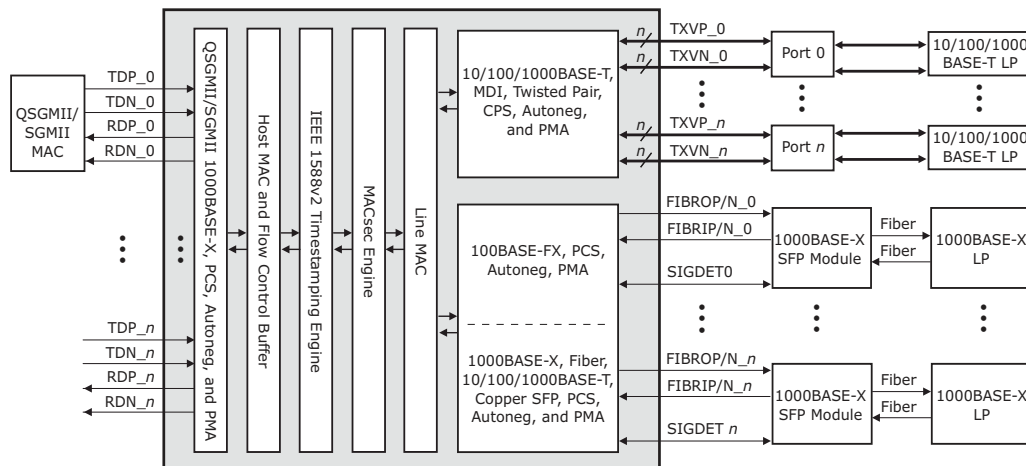
For QSGMII only port 0 is used.

Note: Whenever there is a mode change a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.3 QSGMII/SGMII MAC to AMS and 1000BASE-X Media SerDes

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC to AMS and 1000BASE-X media SerDes.

Figure 8 • QSGMII/SGMII MAC to AMS and 1000BASE-X Media SerDes



3.1.3.1 Media Interface 1000BASE-X SFP Fiber (1000BASE-X Link Partner)

- Set register 23 bits 10:8 = 010
- Set register 0 bit 12 = 1 (enable autonegotiation)

3.1.3.2 AMS Preference Setup

- Set register 23 bit 10 = 1 (enable AMS)
- Set register 23 bit 11 to the port preferences

The media selected by AMS can be read from register 20E1 bits 7:6. For more information, see Table 25, page 56.


For QSGMII only port 0 is used.

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.4 QSGMII/SGMII MAC to AMS and 100BASE-FX Media SerDes

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC to AMS and 100BASE-FX media SerDes.

Figure 9 • QSGMII/SGMII MAC to AMS and 100BASE-FX Media SerDes



3.1.4.1 Media Interface 100BASE-FX SFP Fiber (100BASE-FX Link Partner)

- Set register 23 bits 10:8 = 011
- Set register 0 bit 12 = 1 (enable autonegotiation)

3.1.4.2 AMS Preference Setup

- Set register 23 bit 10 = 1 (enable AMS)
- Set register 23 bit 11 to the port preferences

The media selected by AMS can be read from register 20E1 bits 7:6. For more information, see [Table 25](#), page 56.

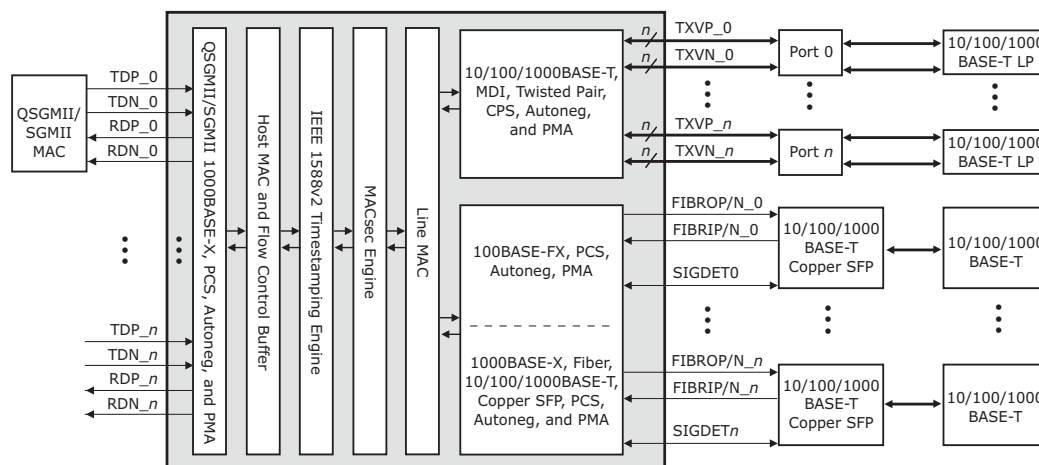
For QSGMII, only port 0 is used.

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.5 QSGMII/SGMII MAC-to-AMS and Protocol Transfer Mode

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC-to-AMS and Protocol Transfer mode.

Figure 10 • QSGMII/SGMII MAC-to-AMS and Protocol Transfer Mode



3.1.5.1 Media Interface 10/100/1000BASE-T Cu-SFP

- Set register 23 bits 10:8 = 001
- Set register 23 bit 12 = 0
- Set register 0 bit 12 = 1 (enable autonegotiation)

3.1.5.2 AMS Preference Setup

- Set register 23 bit 10 = 1 (enable AMS)
- Set register 23 bit 11 to the port preferences

The media selected by AMS can be read from register 20E1 bits 7:6. For more information, see Table 25, page 56.


For QSGMII, only port 0 is used.

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.6 QSGMII/SGMII MAC to Cat5 Link Partner

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC to Cat5 link partner.

Figure 11 • QSGMII/SGMII MAC to Cat5 Link Partner




For QSGMII, only port 0 is used.

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.7 QSGMII/SGMII MAC-to-Protocol Transfer Mode

The following illustration shows the register settings used to configure a QSGMII/SGMII MAC to Protocol Transfer Mode.

Figure 12 • QSGMII/SGMII MAC to Protocol Transfer Mode



3.1.7.1 Media Interface 10/100/1000BASE-T Cu SFP

- Set register 23 bits 10:8 = 001
- Set register 23 bit 11 = 0
- Set register 0 bit 12 = 1 (enable autonegotiation)

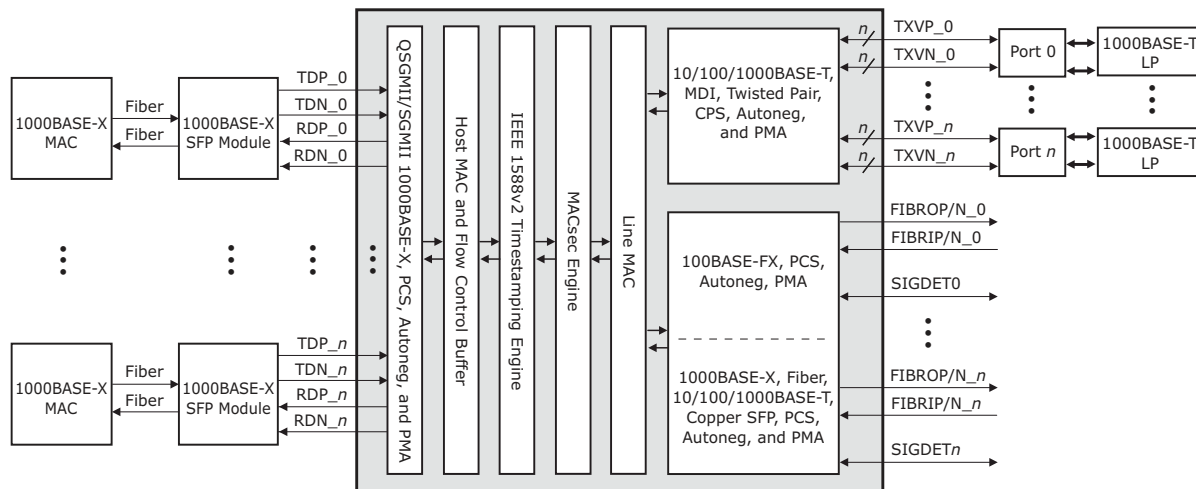
For QSGMII, only port 0 is used.

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.1.8 1000BASE-X MAC to Cat5 Link Partner

The following illustration shows the register settings used to configure a 1000BASE-X MAC to Cat5 link partner.

Figure 13 • 1000BASE-X MAC to Cat5 Link Partner



In this mode, the device provides data throughput of 1000 Mbps only.

3.1.8.1 MAC Interface

- Set Register 18G = 0x80F0 to configure the 1000BASE-X MAC SerDes. For more information, see [Table 139](#), page 183.
- Set register 19G bits 15:14 = 00
- Set Register 23 (main register) bit 12 = 1

3.1.8.2 Clause 37 MAC Autonegotiation

- Set Register 16E3 bit 7 = 1

Note: Whenever there is a mode change, a software reset (register 0 bit 15) is required to make the mode change effective. This register is cleared when read.

3.2 SerDes MAC Interface

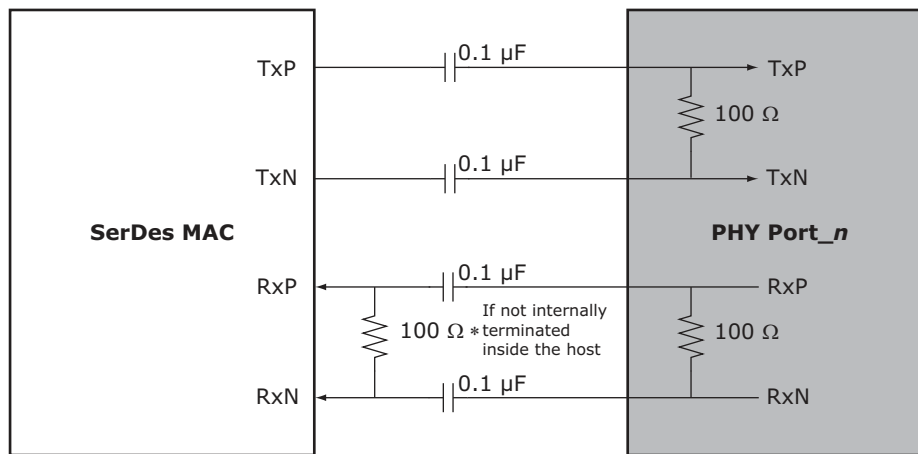
The VSC8584-11 SerDes MAC interface performs data serialization and deserialization functions using an integrated SerDes block. The interface operates in 1000BASE-X compliant mode, QSGMII mode, or SGMII mode. Register 19G is a global register and needs to be set once to configure the device to the desired mode. 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. The SerDes and enhanced SerDes blocks have the termination resistor integrated into the device.

3.2.1 1000BASE-X MAC

When connected to a SerDes MAC compliant to 1000BASE-X, the VSC8584-11 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 = 0, 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.

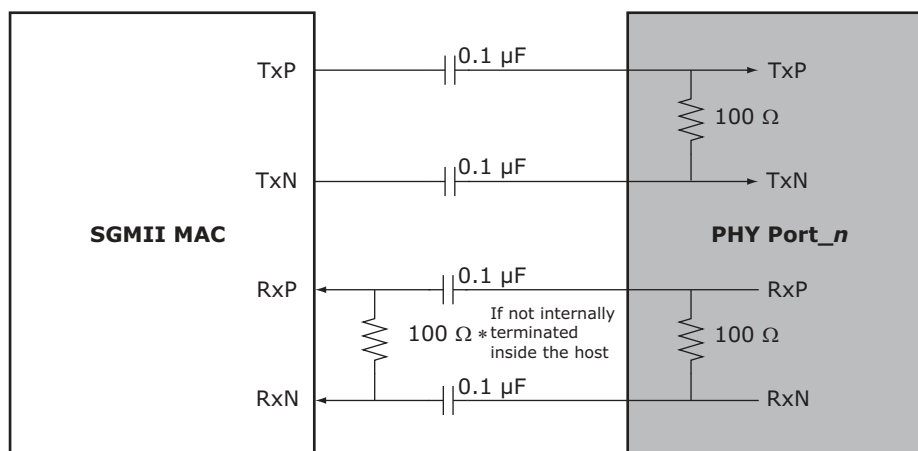
Figure 14 • SerDes MAC Interface



3.2.2 SGMII MAC

When configured to detect and switch between 10BASE-T, 100BASE-T, and 1000BASE-T data rates, the VSC8584-11 device can be connected to an SGMII-compatible MAC. To configure the device for SGMII MAC mode, set register 19G, bits 15:14 = 00 and register 23, bit 12 = 0. In addition, set register 18G as desired. This device also supports SGMII MAC-side autonegotiation and is enabled through register 16E3, bit 7.


Figure 15 • SGMII MAC Interface



3.2.3 QSGMII MAC

The VSC8584-11 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 VSC8584-11 is connecting to supports this functionality. To configure the device for QSGMII MAC mode, set register 19G, bits 15:14 = 01. In addition, set register 18G as desired. The device also supports SGMII MAC-side autonegotiation on each individual port and is enabled through register 16E3, bit 7, of that port.

Figure 16 • QSGMII MAC Interface



3.3 SerDes Media Interface

The VSC8584-11 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 SerDes interface has the following operating modes:

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

The SerDes media block has the termination resistor integrated into the device. 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/SGMII 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 VSC8584-11 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 VSC8584-11 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/SGMII. 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 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 VSC8584-11 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 relevant signal detect 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 VSC8584-11 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 or not 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 VSC8584-11 includes four external PHY address pins, PHYADD[4:1], to allow control of multiple PHY devices on a system board sharing a common management bus. These pins, with the physical address of the PHY, form the PHY address port map. The equation $[(PHYADD[4:1], 1'b0) + \text{physical address of the port (0 to 3)}]$, and the setting of PHY address reversal bit in register 20E1, bit 9, determine the PHY address.

3.4.2 SerDes Port Mapping

The VSC8584-11 includes seven 1.25 GHz SerDes macros and one 5 GHz enhanced SerDes macro. Three of the seven SerDes macros are configured as SGMII MAC interfaces and the remaining four are configured as 1000BASE-X/100BASE-FX SerDes media interfaces. The enhanced SerDes macro can be configured as either a QSGMII MAC interface or the fourth SGMII MAC interface. The following table shows the different operating modes based on the settings of register 19G, bits 15:14.

Table 2 • MAC Interface Mode Mapping

19G[15:14]	Operating Mode
00	SGMII
01	QSGMII
10	Reserved
11	Reserved

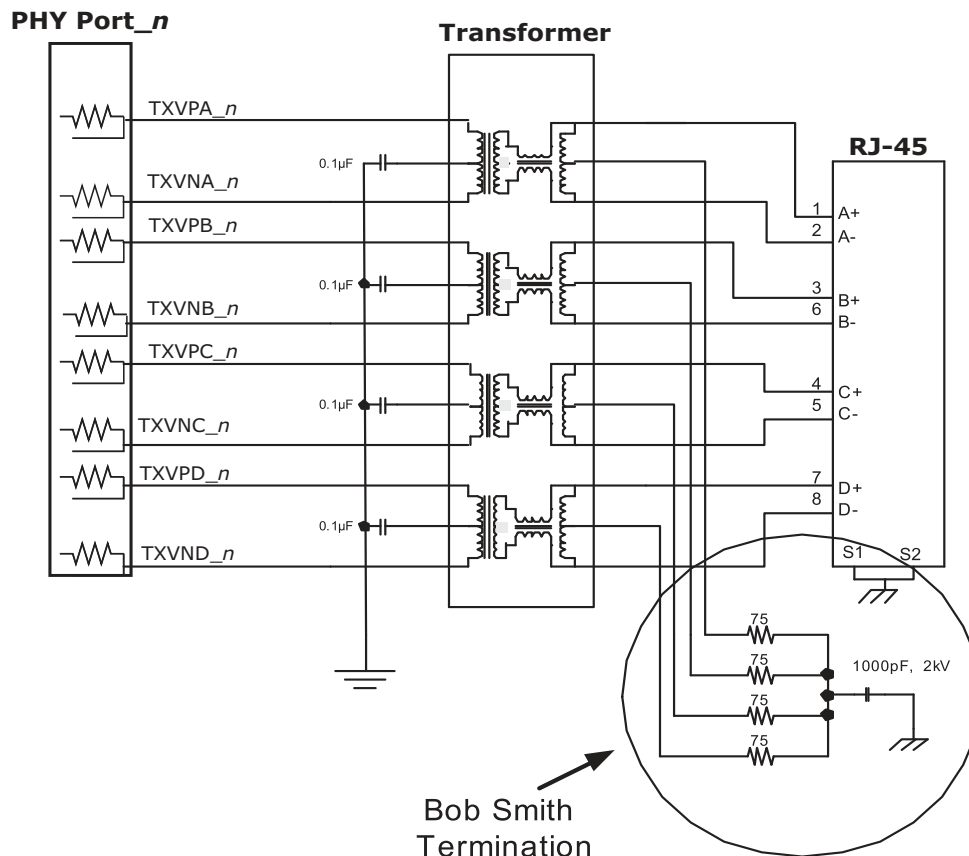
3.5 Cat5 Twisted Pair Media Interface

The VSC8584-11 twisted pair interface is compliant with IEEE 802.3-2008 and the IEEE 802.3az-2010 standard for energy efficient Ethernet.

3.5.1 Voltage Mode Line Driver

Unlike many other gigabit PHYs, the VSC8584-11 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 17 • Cat5 Media Interface



3.5.2 Cat5 Autonegotiation and Parallel Detection

The VSC8584-11 supports twisted pair autonegotiation, as defined by IEEE 802.3-2008 Clause 28 and IEEE 802.3az-2010. 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 VSC8584-11 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 VSC8584-11 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-TX do not require autonegotiation, Clause 40 has defined 1000BASE-T to require autonegotiation.

3.5.3 Automatic Crossover and Polarity Detection

For trouble-free configuration and management of Ethernet links, the VSC8584-11 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 VSC8584-11 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 use the feature, also 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, which shows that twisted pair A (of four twisted pairs A, B, C, and D) is connected to the RJ45 connector 1,2 in normal MDI mode.

Table 3 • Supported MDI Pair Combinations

RJ45 Connections				
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

3.5.4 Manual MDI/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.

3.5.5 Link Speed Downshift


For operation in cabling environments that are incompatible with 1000BASE-T, the VSC8584-11 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 94, page 160.

3.5.6 Energy Efficient Ethernet

The VSC8584-11 supports the IEEE 802.3az-2010 Energy Efficient Ethernet standard. This 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 18 • 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 VSC8584-11 uses LPI to optimize power dissipation in 100BASE-TX and 1000BASE-T modes of operation. In addition, the IEEE 802.3az-2010 standard defines a 10BASE-Te mode that reduces transmit signal amplitude from 5 V peak-to-peak to approximately 3.3 V peak-to-peak. 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 VSC8584-11 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 and 802.3bf](#), page 189.

3.5.7 Ring Resiliency

Ring resiliency changes the timing reference between the master and slave PHYs without altering the master/slave configuration in 1000BASE-T mode. The master PHY transmitter sends data based on the local clock and initiates timing recovery in the receiver. The slave PHY instructs the 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.

Note: For ring resiliency to successfully exchange master/slave timing over 1000BASE-T, the link partner must also support ring resiliency.

3.6 MACsec Block Operation

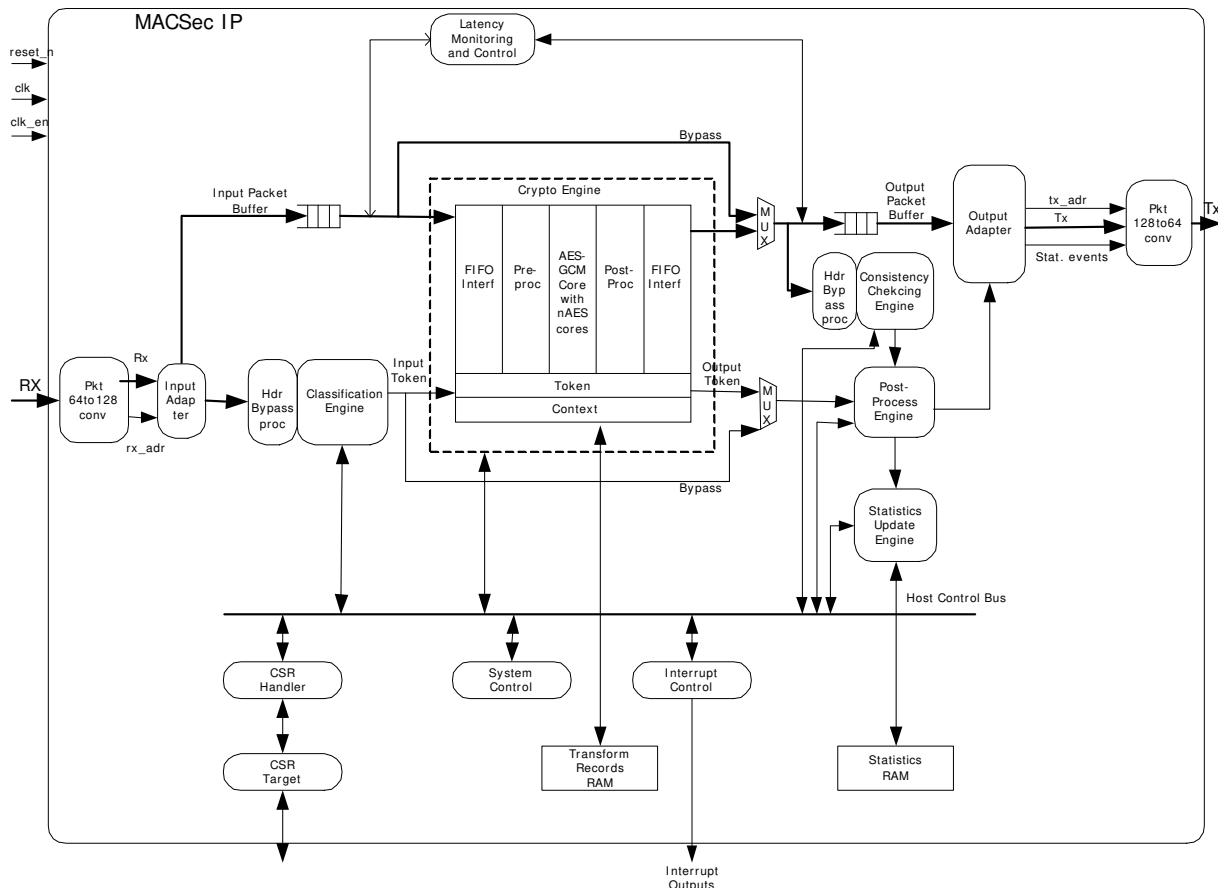
The VSC8584-11 device includes a high-performance streaming MACsec frame processing engine that provides hardware acceleration for the complete MACsec frame transform along with frame classification and statistics counter updates. The following list includes some of the major features of the MACsec engine.

- Fully IEEE 802.1AE-2006, IEEE 802.1AEbn, and IEEE 802.1AEbw-2013 compliant.
- 16 secure associations (SA) per direction and 16 ingress consistency check rules.
- MACsec cipher suite GCM-AES-128 support.
- MACsec cipher suite GCM-AES-256 support.
- MACsec cipher suite GCM-AES-XPB-128/256 support.
- VLAN and Q-in-Q tag detection.
- MACsec tag detection and sub-classification (Untagged, Tagged, BadTag, KaY).
- Programmable “control” packet classification.
- 8-entry programmable non-match flow operation selection (drop, bypass), depending on MACsec tag sub-classification and control packet classification.
- Programmable confidentiality offset (0 B – 127 B).
- SecTAG insertion and removal.
- Integrity Check Value (ICV) checking/removal and calculation/insertion.
- Packet number generation and checking.
- IEEE 802.1AE MACsec statistics counter support.
- Ingress path consistency checking (ICC)–64/16 entry programmable matching table with separate drop/transfer decisions.
- MTU checking and oversize dropping dependent on VLAN User priority for VLAN frames and at global level for non-VLAN frames.
- Advanced MACsec transformations–VLAN tag bypass and EoMPLS header bypass.
- Hardware offload for the nextPN and lowestPN update from the host(KaY).
- Support for AES-ECB, AES-CTR, and AES-GCM/GMAC transformation for FIPS certification of the crypto core.
- Patent-pending architecture to enable use with IEEE 1588v2 with minimal and predictable delays.

3.6.1 MACsec Architecture

The MACsec block operates as a frame processing pipeline whose main function is the implementation of the MACsec transform on Ethernet frames. The following illustration shows the MACsec data flow in one direction.

Figure 19 • MACsec Architecture



The following sections describe the blocks in the MACsec data flow.

3.6.1.1 PKT64to128 Block

The Packet 64to128 block is the Rx interface of the MACsec IP with the other blocks. It converts the 64-bit packet interface to the 128-bit packet interface with which the MACsec IP works. It also presents the port information associated with the current frame. In the egress configuration, the PKT64to128 block has a FIFO to temporarily handle back-pressure from the MACsec IP due to frame expansion. Based on packet expansion within the MACsec IP, the PKT64to128 block provides flow control feedback to the flow control buffer, which manages all data build up that occurs as a result of MACsec frame expansion.

3.6.1.2 Input Adapter Block

The Input Adapter manages the Input Packet interface to ensure interface protocol compliance.

3.6.1.3 Input Classification Engine Block

The Input Classification engine inspects the received frame data and performs the following functions:

- **Control Frame Classification.** A total of 29 programmable rules to classify the frame as a control frame.
- **VLAN Tag Detection.** Programmable functionality to detect VLAN tags and extract information before further classification.
- **MACsec Tag Detection.** Programmable functionality to detect MACsec tags and check whether they are valid (also detects special KaY packet tags).
- **Default Frame Handling.** Classifies packets into eight classes based on the outputs of the control frame classification and MACsec tag detection modules, with control registers to define what to do with a packet (drop or bypass) for each class.

- **Flow Lookup Frame Classification and Frame Handling.** Classifies frames based on frame header field contents and outputs of the control frame classification, VLAN tag detection, and MACsec tag detection modules. Flow control registers define what to do with a frame (drop, bypass, or MACsec process) when matching entries. A programmable per-rule priority level resolves any overlap between these rules.
- **Flow Lookup/Default Classification Multiplexer.** Gives priority to the decision from the flow lookup frame classification. The default frame handling is used for a frame only if none of the flow lookup entries match.

3.6.1.4 Latency Monitoring and Control

The Latency Monitoring and Control module monitors the latency that the first word of each frame incurs going through the pipeline, and optionally stalls the output side until this latency matches a programmable value. This ensures each frame incurs the same latency through the pipeline, irrespective of any processing time differences.

3.6.1.5 MACsec Crypto Engine

The MACsec Crypto engine performs the standard MACsec encapsulation/decapsulation processing. This engine is able to perform a MACsec transform on a frame using GCM-AES-128 according to the IEEE 802.1AE-2006 MACsec specification and its amendment, IEEE 802.1AEbn-2011, which adds the GCM-AES-256 cipher suite. The crypto engine also transforms a frame using GCM-AES-128/256 according to IEEE 802.1aebw-2013. This includes modifications to the Ethernet frame header, insertion/removal of the MACsec header (SecTAG), encryption/decryption, authentication, and authentication result insertion/verification. It does not perform MACsec header parsing, but relies on external logic to provide a processing token that tells it how to process the incoming frame.

In addition to the MACsec specifications 0-byte, 30-byte, and 50-byte confidentiality offset, the MACsec crypto engine supports byte-grained confidentiality offsets from 1 to 127 bytes. The MACsec crypto engine supports one or two VLAN tag bypass operation wherein VLAN tags that bypass MACsec processing are fully excluded from the encryption and authentication, such that the receiver side must be able to remove the bypassed VLAN tags without breaking the MACsec packet. It also supports MPLS header bypass wherein the MPLS link header is excluded from encryption and authentication and the client Ethernet frame is subjected to MACsec transformations.

3.6.1.6 Consistency Checking Engine Block

The Consistency Checking engine checks the contents of a frame at the output of the MACsec Crypto engine (after any MACsec decryption) against a set of 16/64 programmable rules (depending on the configuration) for consistency. A programmable per-rule priority level resolves any overlap between these rules. This engine is not present in the egress configuration.

3.6.1.7 Output Post-Processing Engine Block

The Output Post-Processing engine checks the classification and MACsec Crypto engine processing results against a fixed set of MACsec compliance rules, resulting in a drop decision if the rules are violated. Additionally, it performs programmable MTU checking on the MACsec Crypto engine output frame, with individual global and per-VLAN-user-priority MTU settings.

It combines these internal decisions with decisions made by the Classification and Consistency Checking engines into a final pass/drop decision to the output adapter.

Furthermore, based on all the information from the MACsec Crypto engine and the consistency checking engine available to it, the Output Post-Processing engine decides which statistics counters to increment.

3.6.1.8 Statistics Update Engine Block

The Statistics Update engine updates the statistics counter in the statistics RAM, as instructed by the Output Post-process engine. This allows the updating to be scheduled with external statistics access and to occur in parallel with the post-processing of the next frame. This engine also can be configured to skip certain statistics counters.

3.6.1.9 Output Adapter Block

The Output Adapter block manages the output packet interface and ensures interface protocol compliance by isolating the MACsec IP from this interface.

3.6.1.10 PKT128to64 Block

The Packet 128to64 block is the interface of the MACsec IP with the other blocks. It converts the 128-bit packet interface of MACsec IP to the 64-bit packet interface used to communicate with other blocks. It also prepares the security fail debug code to be put into FCS field for packets failing security check.

3.6.2 MACsec Target Applications

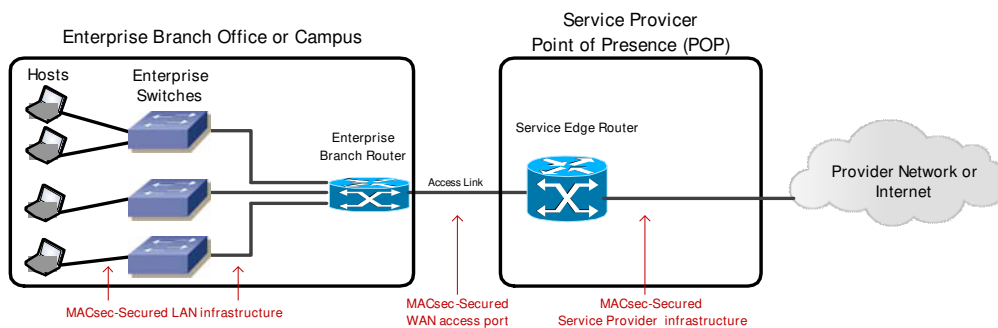
The MACsec engine targets the following applications.

- Secure enterprise infrastructure and WAN ports
- Secure end-to-end Carrier Ethernet connections
- Secure Carrier Ethernet Mobile Backhaul, including high precision IEEE 1588v2 timing

3.6.2.1 MACsec Secured Enterprise Infrastructure and WAN Port

The following illustration shows an enterprise branch office or campus where a Local Area Network (LAN) connected to a Wide Area Network (WAN) operated by a service provider is protected using MACsec.

Figure 20 • Secure Enterprise Infrastructure and WAN




Each host has a dedicated physical link to an Enterprise Ethernet switch, and the switches are connected to an enterprise branch router that also provides WAN access. In smaller configurations, hosts can also connect directly to the branch router. All internal branch office Ethernet ports are secured using MACsec.

The branch router connects across an access link to a service provider's service edge router, and this access link is secured using MACsec. MACsec may also be used to secure the service provider's network.

The 802.1X security protocols can be used for authentication and to automate the distribution and management of MACsec encryption keys. The VSC8584-11 device supports 128-bit and 256-bit encryption.

3.6.2.2 MACsec Secured Carrier Ethernet Connection

The following illustration shows a Carrier Ethernet network providing end-to-end MACsec secured WAN connectivity for an enterprise.

Figure 21 • Secure Carrier Ethernet Connection

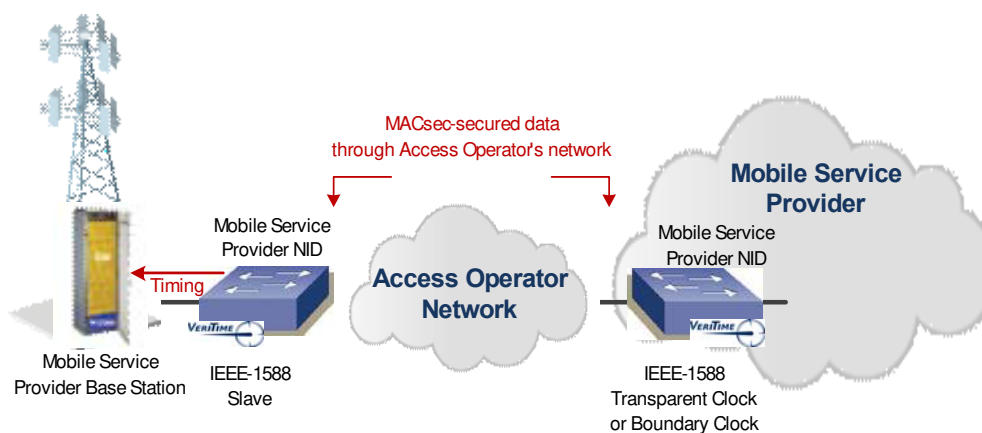
With traditional MACsec, VLAN tags or MPLS labels are fully encrypted and hidden from the Carrier Ethernet network thereby limiting the enterprise to only the simplest point-to-point private line connectivity services.

The VSC8584-11 device supports leaving the VLAN tags or MPLS labels unencrypted for use by the Carrier Ethernet network while fully securing the enterprise's Ethernet data inside these encapsulations. This approach uses standard, non-proprietary encapsulation formats with 128-bit and 256-bit encryption.

By enabling these features, the enterprise is able to take advantage of the latest Layer-2 (L2) VPN services available from a Carrier Ethernet network. These L2 VPN services can be point-to-point or multipoint, and can use standardized Metro Ethernet Forum (MEF) Carrier Ethernet and Internet Engineering Task Force (IETF) MPLS service offerings including multiple Virtual Private Lines per WAN port.

3.6.2.3 MACsec Secured Mobile Backhaul with IEEE 1588

The following illustration shows a typical mobile backhaul application where multiple network operators collaborate to deliver mobile service. In this application, a mobile service provider uses MACsec to secure the backhaul connections end-to-end through the network.

Figure 22 • Secure Mobile Backhaul with IEEE 1588

The mobile service provider may choose to leave VLAN tags or MPLS labels unencrypted so that the access operator can map the virtual private line services.

In addition to backhauling data, IEEE 1588 packet-based timing technology delivers high-precision frequency and phase synchronization to the base stations. IEEE 1588 packets may be encrypted along with backhaul data and tunneled through the access operator network, or delivered as an unencrypted synchronization service directly from the access operator network. To meet 4G/LTE specifications, nanosecond-accurate time stamping of IEEE 1588 packets is required. However, such tight tolerances

cannot be achieved using traditional MACsec, even if the IEEE 1588 packets themselves are unencrypted.

The Microsemi IEEE 1588 time stamping engine in the VSC8584-11 device works in conjunction with the MACsec engine to deliver 4G/LTE timing quality over Carrier Ethernet connections, while using MACsec for end-to-end security across the access operator network.

3.6.3 Formats, Transforms, and Classification

This section shows the frame formats before and after MACsec transformation with an overview of the classifiable fields that can be used for SA classification for different MACsec applications. Classification fields are selectable per SA. In depicting which fields may be used for pre-decrypt classification, it is assumed that the confidentiality offset field is not used (all fields after SecTAG are encrypted).

3.6.3.1 Standard MACsec Formats

The following table summarizes the MACsec frame combinations in the standard MACsec mode.

Table 4 • Standard MACsec Frame Combinations

Unencrypted Format	Pre-Encryption (Tx) Classification Fields	Pre-Decryption (Rx) Classification Fields
Untagged Ethernet	DA, SA, Etype	DA, SA, SecTAG
Single-tagged Ethernet	DA, SA, TPID, VID, Etype	DA, SA, SecTAG
Dual-tagged Ethernet	DA, SA, TPID1, VID1, TPID2, VID2, Etype	DA, SA, SecTAG

The following illustrations show each frame format before and after standard MACsec transformation.

Figure 23 • Untagged Ethernet



Figure 24 • Standard MACsec Transform of Untagged Ethernet




Figure 25 • Single-Tagged Ethernet




Figure 26 • Standard MACsec Transform of Single-Tagged Ethernet




Figure 27 • Dual-Tagged Ethernet





Figure 28 • Standard MACsec Transform of Dual-Tagged Ethernet



3.6.3.2 Advanced MACsec Formats

The following table summarizes the MACsec frame combinations in the advanced MACsec mode.

Table 5 • Advanced MACsec Frame Combinations

Unencrypted Format	Encrypted Format	Pre-Encryption (Tx) Classification Fields	Pre-Decryption (Rx) Classification Fields
Single-tagged Ethernet	MACsec plus single tag bypass	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTAG
Dual-tagged Ethernet	MACsec plus single tag bypass	DA, SA, TPID1, VID1, TPID2, VID2, Etype	DA, SA, TPID1, VID1, SecTAG
Dual-tagged Ethernet	MACsec plus dual tag bypass	DA, SA, TPID1, VID1, TPID2, VID2, Etype	DA, SA, TPID1, VID1, TPID2, VID2, SecTAG
EoMPLS with one Label	MACsec plus EoMPLS header bypass	C-DA, C-SA, MPLS Etype, 32-bit Label	C-DA, C-SA, MPLS Etype, 32-bit label, SecTAG
EoMPLS with two Labels	MACsec plus EoMPLS header bypass	C-DA, C-SA, MPLS Etype, 32-bit Label1, 32-bit Label2	C-DA, C-SA, MPLS Etype, 32-bit label1, 32-bit label2, SecTAG

The following illustrations show each frame format before and after advanced MACsec transformation.

Figure 29 • Single-Tagged Ethernet




Figure 30 • MACsec Transform to Single Tag Bypass




Figure 31 • Dual-Tagged Ethernet



Figure 32 • MACsec Transform to Single and Dual Tag Bypass




Figure 33 • EoMPLS with One Label




Figure 34 • Standard and Advanced MACsec Transform

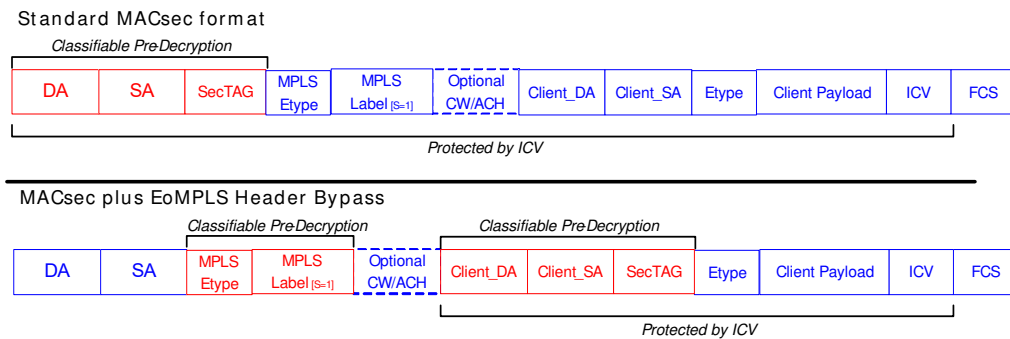


Figure 35 • EoMPLS with Two Labels


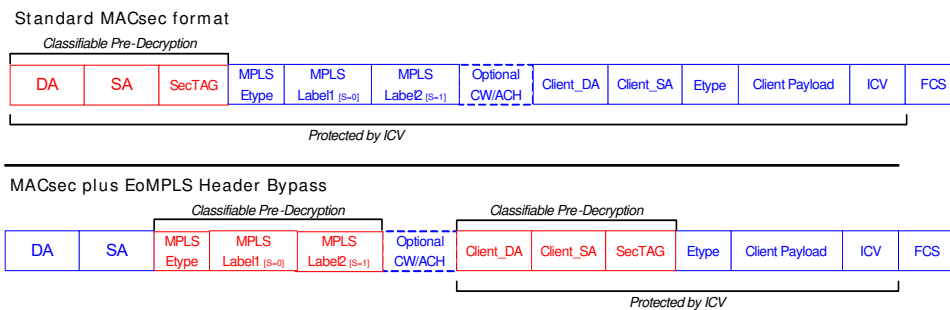


Figure 36 • Standard and Advanced MACsec Transform



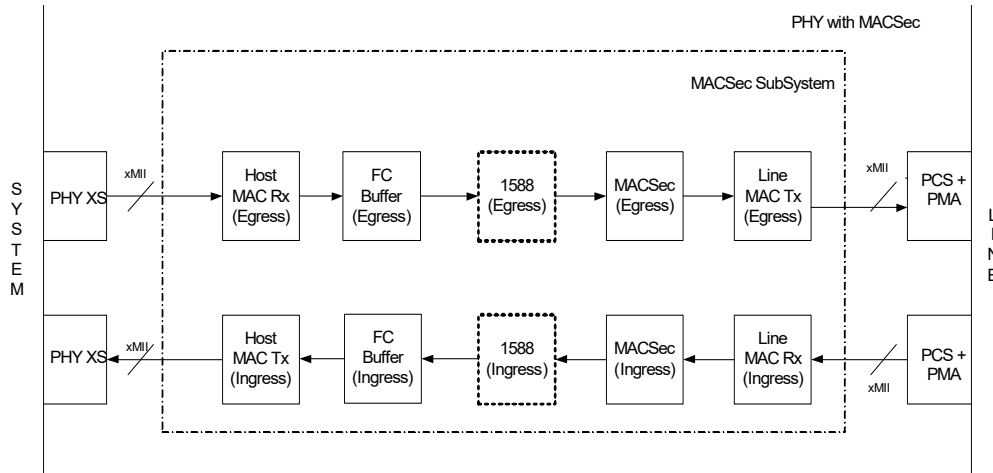
3.6.4 MACsec Integration in PHY

The MACsec block is designed to be integrated with a host MAC and a line MAC to form a plug-in MACsec solution between an existing Ethernet MAC (system-side) and an existing Ethernet PHY (line-side). MACsec adds bandwidth in egress. This increase in bandwidth is handled adding IEEE 802.3 pause flow control toward the system. The FC buffer block provides packet buffering and controls the

IEEE 802.3 pause flow control generation to handle MACsec frame expansion. The IEEE 1588 block is on the host side of MACsec, and the IEEE 1588 PTP frames are also subjected to MACsec transformations.

The following illustration shows the integration of MACsec in the PHY.

Figure 37 • MACsec in PHY



3.6.5 MACsec Pipeline Operation

MACsec ingress and egress pipeline operations are identical except for a few situations mentioned in the following sections. The MACsec block always operates in cut-through mode. The length of the frame is calculated on the fly and does not need to be known before the start of processing. This means that MACsec egress processing encrypts (protects) all bytes of the frame fed into the MACsec core. If the frame contains Ethernet padding, this padding is encrypted/protected by MACsec and the ICV is appended after it. For ingress processing, the MACsec block accepts frames with Ethernet padding and it strips Ethernet padding from short MACsec frames.

Ethernet frames are submitted to the MACsec egress/ingress block with their Ethernet header (destination address, source address, Ethertype) but without the leading preamble and start-of-frame bytes and trailing 4-byte CRC (FCS). It is the responsibility of the host/line MAC to strip and check the CRC of each incoming frame.

In the case of large frames, the first output data word of a frame may leave the MACsec pipeline before the last input data word of a frame enters, and errors such as ICV check verification or MTU checking may only be detected after the last byte of frame data has been processed. As a consequence, dropping a frame is accomplished by setting the frame abort signal and not by preventing the frame from appearing on the output. In other words, the system/line MAC transmits a frame with bad CRC. The engine can be programmed to drop frames completely (internal drop), but only if the decision to drop has been made by the flow lookup stage. The pipeline outputs the (processed) frames in the same order they are input, unless the frame is dropped internally. The MACsec block can also be bypassed completely to improve latency.

The SL field in MACsec indicates the end of the MACsec frame, which is needed to locate the ICV in case Ethernet padding follows the ICV. For such frames, the MACsec block uses the information from the SecTAG of the frame to calculate the actual MACsec frame length and uses this length during ingress processing. All data that follows the ICV is removed from the data stream by the MACsec block. This action is the de-padding action, using the MACsec protocol header. The ICV is assumed to be at the location as indicated by the SecTAG, otherwise the frame does not pass the MACsec integrity check.

If the SL field in the MACsec frame indicates a longer frame than the packet actually received by the MACsec block (if the frame does not pass MACsec PDU check), the MACsec block flags this situation as an integrity check failure or packet length error, depending on the difference in length.

Note: The de-padding action is applicable only for MACsec frames that are going to be decrypted/validated by the MACsec flow, and will not change the regular MACsec processing latency. No de-padding action is performed on bypass/drop frames.

After ingress MACsec processing, it is possible for the frame to become smaller than 64 bytes. Such frames are then padded by the host MAC (if enabled) and the packet processing switch/system receives 64 byte frames after Ethernet padding.

In the egress direction, the MACsec core calculates and updates the SL field in the MACsec header, and authenticates and encrypts (optionally) the frame if a frame's size (including the MACsec header and ICV) is less than 64-bytes.

Note: This short length field indicates frame data from after the first byte of the MACsec header to byte immediately before ICV.

For this feature to work, host MAC receiver should be configured to allow undersized frames and line MAC transmitter should be configured to pad frames.

Host and line MACs do not accept less than 64 byte frames (without Ethernet padding) from system/line interfaces. Also they do not remove the Ethernet padding from the frames.

Each frame at input is accompanied by the following signals:

- **Port Number.** Two-bit signal that indicates the source port (common, reserved, controlled, or uncontrolled) of the packet as defined in the IEEE 802.1AE standard.
- **Bad CRC/Package Error.** Bits that indicate that the packet has a bad CRC/package error.

Frames with a bad CRC or other packet errors are forwarded to the output with the same errors, unless their classification leads to a decision to drop them. Because error signals appear at the end of a frame and processing must start before the end of a frame is received, classification and processing is performed, but statistics are not updated.

The source port for MAC data/control frames is configurable. Typically, egress MAC data frames are put on the controlled port and MAC control frames are put on the uncontrolled port. All ingress frames are put on the common port. This configuration is controlled using `MAC_DATA_FRAMES_SRC_PORT` and `MAC_CTRL_FRAMES_SRC_PORT` in the `MACSEC_CTL_CFG` register. Control packet classification determines the frames that are assumed to have come from controlled/uncontrolled ports in egress and the frames that should go to controlled/uncontrolled ports in ingress.

LPI and fault signals that appear on the Ethernet interface can be detected by the MAC and converted into internal status frames (single-byte frames containing the state of the signals). The MACsec block can recognize these status frames on the input and propagate them to its output.

Status frames travel through the pipeline along with normal Ethernet frames, so they appear at the output after the preceding Ethernet frame and before any frames that appear after the status change. However, status frames do not take part in any operations of the pipeline. They are invisible to static classification, flow lookup, MACsec processing, and consistency checking.

The following illustrations show the egress and ingress MACsec data flows.

Figure 38 • MACsec Egress Data Flow

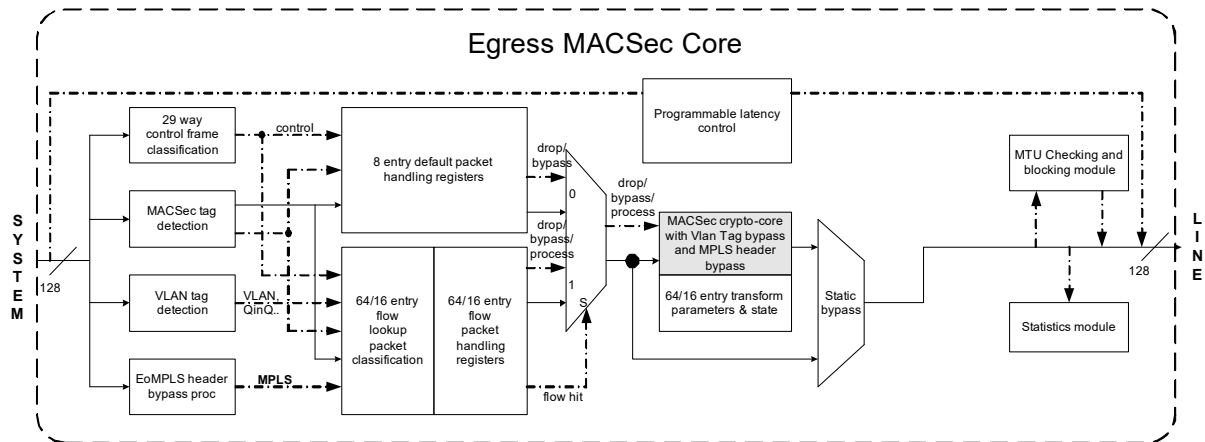
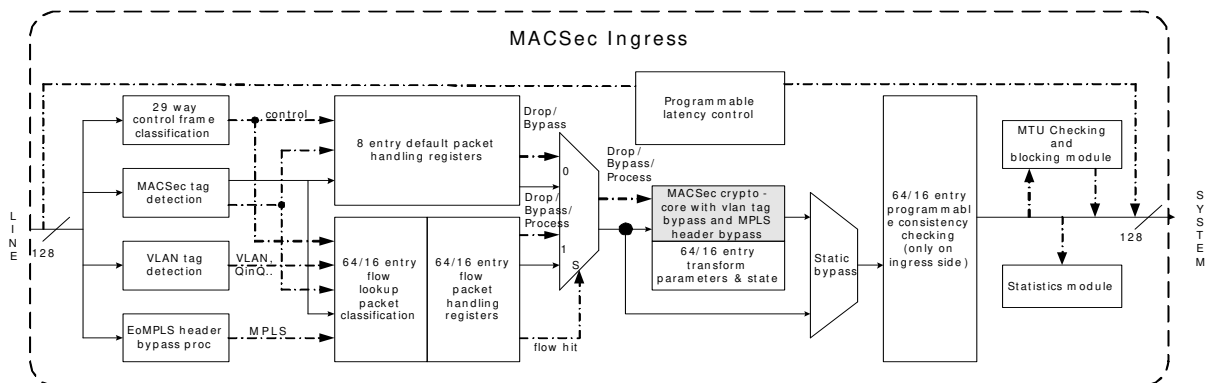


Figure 39 • MACsec Ingress Data Flow



The following sections describe the pipeline stages. Of these pipeline stages, the MACsec transform stage is the only one that can modify the frame data or drop a frame completely (no frame will appear at the output of the pipeline in that case). Other stages can only perform the following actions for frame data:

- Inspect the frame data, such as performing a classification based on fields in a header.
 - Drop the frame (which is already streaming out) by setting the frame abort signal along with the last word of data.
- Static Classification

This is the first stage of packet classification. Control packet classification, MACsec tag parsing, and VLAN tag parsing are carried out in parallel.

Flow Lookup

Each table of 16 SA flows can match on a number of criteria. An action and a MACsec context is associated with each flow. If the packet does not match any of these 16 flows, one of eight default actions is selected, depending on the results of MACsec tag parsing and control packet classification.

MACsec Transform

This stage carries out the actual MACsec encryption and authentication. It uses the MACsec context associated with the flow that was matched in the previous stage. A MACsec context is a data structure containing all information (such as key and sequence numbers) needed to carry out a MACsec transform. This stage can also bypass or drop certain packets.

The MACsec transform stage can be bypassed by setting MACSEC_BYPASS_ENA = 1 in the MACSEC_ENA_CFG register. Setting the MACSEC_BYPASS_ENA = 0 and MACSEC_ENA = 1 results in traffic passing through the MACsec transform block. Setting the MACSEC_BYPASS_ENA and MACSEC_ENA bits to 0 results in all traffic being dropped at the input interface of MACsec.

Ingress Consistency Checking

This stage is not present in the egress-only version of the MACsec block. It extracts information from the decrypted packet and checks it against a table of 64/16 rules. Rules can either reject or pass certain packets. Separate default actions can be configured for control and non-control packets (in case no match is found in the table).

Output Postprocessor

This stage checks the results of the MACsec transform operation. It also checks that the length of a packet does not exceed the MTU (incrementing a counter if the MTU is exceeded, and optionally tagging the packet for deletion). Each of the eight VLAN user priorities and non-VLAN packets can have a different MTU. This stage implements the MACsec-compliant post-processing decision tree and updates all MACsec statistics.

3.6.5.1 Static Classification

Control packet classification, MACsec header parsing, and VLAN tag parsing are the three static classification operations performed in parallel to produce the following results:

- **Control.** Single bit that is set if the packet is classified as a control packet.
- **MACsec Tag Status.** One of four values: untagged, tagged, bad tag and KaY, where tagged means the packet has a valid non-KaY MACsec SecTAG.
- **VLAN Related Status Signals.** VLAN valid, VLAN ID, Inner VLAN ID, VLAN User Priority, Inner VLAN User Priority, QTAG valid, STAG valid, and Q-in-Q valid.
- **Parsed Ethertype.** First non-VLAN Ethertype found in the frame.

The following sections describe the static classification operations.

3.6.5.1.1 Control Packet Classification

Control packet classification is used to identify frames from uncontrolled ports and exclude them from MACsec processing. Frames such as the MAC control frames and MKA/EAPOL frames are forwarded without MACsec processing because they use uncontrolled ports for transmission. MKA/EAPOL frames are used for Key exchange and have Ethertype 0x888E.

The control packet classification logic classifies a packet as a control packet based on its destination address and/or its Ethertype. It yields a single-bit output (control) classifying the packet either as a control packet or not.

The control packet classification logic can match a packet based on 29 individually enabled criterion. If the packet matches one or more of the enabled criterion, the packet is classified as a control packet and the control output is set to 1. If no enabled criterion is matched, the packet is not classified as a control packet. The CTL_PACKET_CLASS_PARAMS and CTL_PACKET_CLASS_PARAMS2 registers configure control packet classification. The match criterion are as follows:

- The fixed Ethernet destination address 01_00_0C_CC_CC_CC. The corresponding register CP_MAC_DA_48* has this address as a reset value, but this value can be changed if needed.
- The fixed Ethernet destination address range 01_80_C2_00_00_0? (the first 44 bits must match; the trailing 4 bits are don't care). The corresponding register CP_MAC_DA_44* has this address range as a reset value, but this value can be changed if needed. It is always a range with 44 matching bits and 4 don't care bits.
- One free to program Ethernet destination address range specified by the CP_MAC_DA_START and CP_MAC_DA_END registers. Ethernet addresses are treated as unsigned 48-bit integers, as shown in the following examples.

If CP_MAC_DA_START = 00_80_C2_00_00_00 and CP_MAC_DA_END = 00_80_C2_00_00_0F, the matched range is identical to the range normally matched by MAC_DA_44.

If CP_MAC_DA_START = 00_00_00_00_00_00 and CP_MAC_DA_END = FF_FF_FF_FF_FF_FF, all destination addresses are matched (every packet is classified as a control packet).

If CP_MAC_DA_START = CP_MAC_DA_END, then only a single address will be matched.

- Eight individual Ethernet destination addresses: CP_MAC_DA_MATCH_0 through CP_MAC_DA_MATCH_7.

- Sixteen individual Ethertypes: CP_MAC_ET_MATCH_0 through CP_MAC_ET_MATCH_7 where each Ethertype compare value field shares a register with two destination address compare value bytes and CP_MAC_ET_MATCH_10 through MAC_ET_MATCH_17 registers.
- Two combinations of destination address and Ethertype: CP_MAC_DA_ET_MATCH_8 and CP_MAC_DA_ET_MATCH_9. A packet matches only if both the destination address and the Ethertype match.
- Even though the registers for destination addresses and Ethertypes 8 and 9 have the same format as those for destination addresses and Ethertypes 0 to 7 and 10 to 17, they have different semantics. Destination address 8 can only be enabled in combination with Ethertype 8, and only packets with both a matching destination address and a matching Ethertype will match this criterion. The same applies to destination address and Ethertype 9. On the other hand, destination addresses 0 to 7 can be enabled independent of Ethertypes 0 to 7. When both destination address 0 and Ethertype 0 are enabled, packets that have either a matching destination address or a matching Ethertype (or both) will be classified as control packets.
- After reset, control packet matching criteria are disabled. The registers for a matching criterion must be programmed to enable it.
- Either the first Ethertype after the DA/SA fields or the parsed Ethertype, determined by the VLAN parsing algorithm, is the Ethertype value (number 0 to 17, including the combined numbers 8, 9) from the packet that can be used to compare. This selection is done using the CP_MATCH_MODE register.
- Rules are enabled using the CP_MATCH_ENABLE register.

3.6.5.1.2 MACsec Tag Parsing

The MACsec tag parsing logic inspects MACsec tags. MACsec tags must follow the source address, without any intervening VLAN tags. (They may follow VLAN tags only in VLAN tag bypass mode.) MACsec tag parsing classifies each packet into one of four categories:

- **Untagged.** No MACsec tag (Ethertype differs from 0x88E5).
- **Bad Tag.** Invalid MACsec tag, as determined by the tag detection logic.
- **KaY Tag.** These packets are generated and/or handled by software and no MACsec processing is performed on them by the hardware except for straight bypass.
- **Tagged.** Valid MACsec tag that is not KaY.

The following table shows the IEEE 802.1AE checks that determine the status of the MACsec tag parsing.

Table 6 • MACsec Tag Parsing Checks

MACsec Tag (SecTAG) Check	Result
Ethertype is not MACsec type	Untagged
V bit = 1	Bad tag
C bit = 0 and E bit = 1	KaY
C bit = 1 and E bit = 0	Bad tag
SC bit = 1 and ES bit = 1	Bad tag
SC bit = 1 and SCB bit = 1	Bad tag
SL \geq 48	Bad tag
PN = 0	Bad tag
All other	Tagged

MACsec tag parsing checks are controlled by configuring the SAM_NM_PARSING register.

3.6.5.1.3 VLAN Tag Parsing

The VLAN tag parsing logic recognizes VLAN tags that immediately follow the source address. Both 802.1Q and 802.1s tags can be recognized. Packets with two VLAN tags can also be recognized.

The VLAN tag parsing logic generates the following signals that can be used by flow lookup and other processing stages.

- **VLAN Valid.** Single bit that is set when a VLAN tag (of any type) is successfully parsed.
- **Stag Valid.** Single bit that is set if the first valid VLAN tag is an 802.1S tag.
- **Qtag Valid.** Single bit that is set if the first valid VLAN tag is an 802.1Q tag.
- **Q-in-Q Found.** Single bit that is set if two valid VLAN tags were found.
- **VLAN User Priority.** Three-bit field derived from the first VLAN tag. For non-VLAN tag packets the default user priority is returned. User priority processing can be disabled to also return the default user priority.
- **VLAN ID.** Twelve-bit field taken from the first VLAN tag. Undefined for non-VLAN packets.
- **Inner VLAN User Priority.** Three-bit field derived from the second (inner) VLAN tag. This value is always passed through the re-mapping table (the SAM_CP_TAG2 register) and the result value is used in classification. Undefined for non-VLAN packets or VLAN packets without a second VLAN tag.
- **Inner VLAN ID.** Twelve-bit field that is taken from the second (inner) VLAN tag. Undefined for non-VLAN packets or VLAN packets without a second VLAN tag.
- **Ethertype.** Ethertype extracted from the packet after zero, one, or two VLAN tags.

VLAN parsing is controlled by configuring the SAM_CP_TAG, SAM_PP_TAGS, SAM_PP_TAGS2, and SAM_CP_TAG2 registers.

The parsed VLAN fields (including UP) are used in SA flow classification lookup. The MACsec block also maintains VLAN statistics on a per user priority basis. This includes dropped and oversize packets on a user priority basis.

3.6.5.2 Flow Lookup/SA Flow Classification

The flow lookup logic associates each packet with one of the two following flows:

- A table of SA matching flows, each of which can match a packet based on a set of match criterion. If a packet matches multiple (enabled) SA flows, the SA flow with the highest user-defined priority value is selected. The flow specifies which action must be performed (drop the packet, pass it unchanged, or perform a MACsec transform). Each SA flow for which a MACsec operation is specified corresponds to exactly one MACsec context (and hence to a single MACsec SA, either ingress or egress). In other words, all packets that are to be processed using a single MACsec SA have to be matched by a single SA flow.
- A table of eight non-matching flows. If no enabled SA flow matches a packet, a non-matching flow is selected based on the MACsec tag parsing result and the control bit (from the control packet classification). For these non-matching flows the only possible actions are bypass and drop (MACsec operations cannot be selected here).

The output of the flow lookup is as follows:

- **SA Hit.** Single bit signal that is set if the packet matched an enabled SA flow.
- **SA Index.** Index of the SA flow being matched. If no SA flow was matched, this field is composed from the control packet classification and MACsec tag parsing results, which identifies the non-matching flow used.

3.6.5.2.1 SA Match Criteria

Each SA flow has a set of registers that specify the match criteria using one of two following categories:

- The four MACsec tag match bits (untagged, tagged, bad_tag, and kay_tag in the SAM_MISC_MATCH registers). If the corresponding bit is set in the SAM_MISC_MATCH register, packets from that category (as classified by the MACsec classification logic) can be matched if the other criteria are also satisfied. If the corresponding bit is clear, packets from that category can not be matched.
- The mask-able match criteria. Each of these criteria can be masked by a mask bit in the SAM_MASK registers. If the corresponding mask bit is clear, the matching criterion is not tested and packets may be matched regardless of actual value in the packet. If the corresponding mask bit is set, the matching criterion is tested; if the packet has a different value from that specified in the flow, the packet will not be matched.

The following table shows the match criteria and maskable bits.

Table 7 • Match Criteria and Maskable Bits

Egress/Ingress SA Match Classifiers	Data Bits	Mask Bits
MAC SA and MAC DA (mask bit per byte)	96	12
MAC Ethertype (parsed Ethertype)	16	1
VLAN class / parsing result (vlan_valid, qtag_valid, stag_valid, qinq_found)	4	4
VLAN UP (parsed User Priority)	3	1
VLAN ID	12	1
Inner VLAN UP (inner User Priority when Q-in-Q is detected)	3	1
Inner VLAN ID (inner VLAN ID when Q-in-Q is detected)	12	1
Source port (controlled/uncontrolled/common/reserved)	2	1
Control packet	1	1
MACsec tag classifier output (untagged/tagged/bad tag/KaY)	0	4
MACsec SCI (compared only for MACsec tagged frames, available only in ingress)	64	1
MACsec TCI.AN (compared only for MACsec tagged frames, available only in ingress, individually masked)	8	8
Field_2B_16B (used in MPLS header bypass mode)	64	64
SA match priority	4	0
Entry enable	1	0

If all four MACsec tag match bits are set and none of the mask bits are set, the flow matches all possible packets. If none of the MACsec tag match bits are set, the flow does not match any packets.

If an exact match of the MAC source address is desired, all six mac_sa_mask bits must be set. If an exact match of the MAC destination address is desired, all six ma_da_mask bits must be set.

The SCI and TCI.AN fields are used in only in the ingress SA flow classification. The TCI.AN field match can be masked per bit. If an exact match of the TCI.AN field is desired, all eight tci_an_mask bits must be set. If a match on the SCI field is desired, make sure that the SCI field is expected in the packet and match on the SC bit in the TCI field (SC bit must be set). For packets without an SCI field, the TCI field in combination with the MAC source address determines the match criterion (as defined in the 802.1AE standard).

The VLAN ID output can be undefined for non-VLAN packets. When matching packets on VLAN ID, also match on vlan_valid = 1.

A packet is matched on the parsed Ethertype from the VLAN classification logic. This differs from the Ethertype used by the control packet classification logic.

Each flow can be enabled or disabled individually. Only enabled flows are selected when they match a frame. When multiple enabled flows match a frame, the one with the highest match_priority field (a number from 0 to 15) will be selected; among equal priority flows the one with the lowest index will be selected.

The match_priority field is always 4-bit wide (16 priority levels) regardless of the number of SAs supported in the given configuration. The SA_MATCH_PARAMS registers control the SA match criteria.

3.6.5.2.2 Enabling and Disabling Flows

SA_MATCH_CTL_PARAMS registers control the enabling and disabling of matching table entries in the main SA matching module. It is also possible to set, clear, and toggle enable bits with a single write action.

Note: To write the match registers of an SA flow or the MACsec context, the flow must be disabled first to ensure that all flow parameters are loaded into the engine when the flow is enabled again.

If the block supports more than 32 SAs, setting, clearing and toggling of enable bits for SA entries beyond 32 requires two write operations. The upper flags are stored with the first write operation to SAM_TOGGLE2, SAM_SET2, or SAM_CLEAR2 respectively. The action for all SA entries is applied and the upper flags are cleared to zero with the second write operation to SAM_TOGGLE1, SAM_SET1, or SAM_CLEAR1 respectively.

Each SA flow can be enabled or disabled individually. If an SA flow is disabled, it will not match any packets.

When a previously enabled SA flow is disabled (by writing to the SAM_ENTRY_CLEAR1/2 or SAM_ENTRY_TOGGLE1/2 registers), the hardware loads the unsafe field in SAM_IN_FLIGHT register with the number of packets currently processed in the pipeline and the software must wait for the unsafe field to reach zero before it writes to the MACsec context or any of the registers belonging to that SA flow. This is necessary to make sure that all packets that might make use of the disabled flow or the associated MACsec context have left the engine.

3.6.5.2.3 Flow Actions

Each SA flow has a SAM_FLOW_CTRL_IGR/EGR register that specifies the action that must be taken when a frame is matched by that SA flow. The action is determined by one of the following four flow types.

- **Bypass** The frame is passed unchanged.
- **Drop.** The frame is dropped. The drop_action field specifies the action.
 - The packet can be forwarded with a corrupt CRC indication.
 - The packet can be forwarded with a bad packet (packet error) indication.

Note: In both cases, the frame abort signal is set towards the MAC and the drop behavior is the same.

- The packet can be dropped internally. The dropped packet does not appear on the output of the MACsec because the drop_internal decision is taken before the end of the packet is seen. This operation can drop packets received with CRC and/or packet errors.

3.6.5.2.4 MACsec Ingress and Egress Processing

MACsec ingress and egress processing includes performing the MACsec transform (adding/removing SecTag, encryption/decryption, and generating/verifying ICV), post-processing steps, and updating statistics counters. A properly configured MACsec block implements all per-packet steps of a compliant MACsec implementation.

The flow action also specifies the destination port of the packet (as defined in the 802.1AE standard) in a two-bit field that appears at the output of the data pipeline to PKT128to64 and will be used for statistics.

The following table shows the egress SA flow action related to a matching entry, as defined in the SAM_FLOW_CTRL_EGR register.

Table 8 • Egress SA Flow Actions

SA Flow Action	Description	Data Bits
Flow type	Bypass/Drop/Egress process	2
Dest_port	Destination port 00b: Common port 01b: Reserved port 10b: Controlled Port 11b: Uncontrolled port	2

Table 8 • Egress SA Flow Actions (continued)

SA Flow Action	Description	Data Bits
Drop_action	Defines the way drop operation is performed	2
protect_frame	1b: Enable frame protection 0b: Bypass frame through crypto-core	1
sa_in_use	MACsec SA is in use for the looked up SA	1
include_sci	Enables use of implicit/explicit SCI	1
use_es	Enable ES bit	1
use_scb	Enable SCB bit	1
Tag_bypass_size	The number of allowed tags to bypass MACsec (0/1/2)	2
Confidentiality offset	The number of bytes that must be authenticated but not encrypted after SecTAG	7
Confidentiality protect	Enables confidentiality protection	1

The following table shows the ingress SA flow action related to a matching entry, as defined in the SAM_FLOW_CTRL_IGR register.

Table 9 • Ingress SA Flow Actions

SA Match Action	Description	Data Bits
Flow Type	Bypass/Drop/Ingress process	2
Dest_port	Destination port 00b: Common port 01b: Reserved port 10b: Controlled Port 11b: Uncontrolled port	2
Drop_action	Defines the way drop operation is performed	2
Drop_non_reserved	Perform drop_action if packet is not from the reserved port	1
Replay_protect	Enable/Disable frame replay protection	1
sa_in_use	MACsec SA is in use for the looked up SA	1
validate_frames	Frame validation level for MACsec ingress processing (disable/check/strict)	2
Confidentiality offset	The number of bytes that must be authenticated but not decrypted after SecTAG	7

MACsec contexts, which store the sequence number, keys, SCI, and other information, are used for further transformation of frames for MACsec egress/ingress flow type processes.

3.6.5.2.5 Non-Matching Flows

The SAM_NM_FLOW_NCP/SAM_NM_FLOW_CP registers define how packets that did not match any of the SA match entries are handled. This is subdivided into eight packet type categories, split by whether or not the packet was classified as a control packet and the output of the MACsec tag classification logic (untagged/tagged/bad tag/KaY).

The actions specified for each flow are a subset of those specified for SA flows (only pass and drop are possible). Each of these flows can specify that a packet must be dropped or bypassed. It also specifies the destination port. The way a packet must be dropped can also be specified.

3.6.5.3 VLAN Tag and EoMPLS Header Bypass Modes

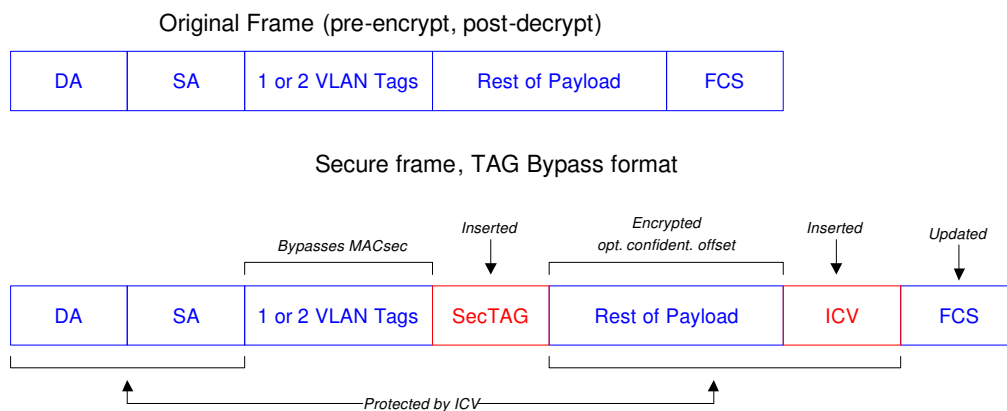
VLAN tag bypass and EoMPLS header bypass are advanced MACsec processing modes with the following classification extensions to the standard configuration.

- Handling of VLAN Tag bypass format (tag bypass).
- Handling of EoMPLS header bypass format (header bypass).
- Processing of packets with SecTAG appearing after one or two VLAN tags, where VLAN tags are not included in the cryptographic operations (Microsemi tag bypass format).
- Processing of packets with SecTAG (and C-SA & C-DA) appearing after an Ethernet Header (SA, DA, ET) with from 2 to 16 bytes of data, where the header and data is not included in the cryptographic operations (Microsemi header bypass format).
- Control packet detection for packets in these proprietary formats.
- Programmable match fields used in SA lookup for packets in these proprietary formats.

3.6.5.3.1 Tag Bypass Frame Format

Tag bypass is an extension to the standard MACsec frame that allows one or two VLAN tags in front of the SecTAG. These VLAN tags are fully excluded from MACsec protection and bypassed instead. The following illustration shows the format of the frame.

Figure 40 • VLAN Tag Bypass Format

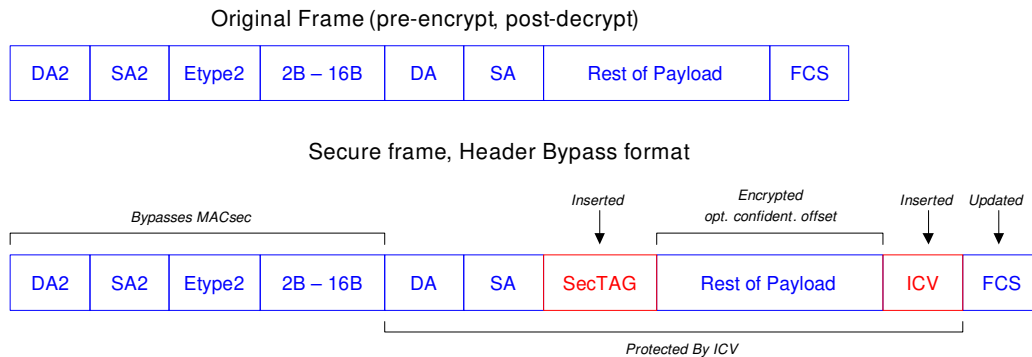


The following logic is used to process the tag bypass format.

- For egress processing, the number of bypassed VLAN tags for encryption is looked-up in the MACsec flow action (SAM_FLOW_CTRL_EGR::TAG_BYPASS_SIZE). If this value is zero, the standard MACsec protection is applied.
- For ingress processing, the number of bypassed VLAN tags is determined by the VLAN parser and position of the SecTAG. The VLAN parser does not look beyond the SecTAG.
- KaY packets (to be bypassed) are detected on both egress and ingress configurations (the VLAN parser defines SecTAG position).
- VLAN tags that bypass MACsec processing are fully excluded from the encryption and authentication, such that the receiver side must be able to remove the bypassed VLAN tags without breaking the MACsec packet.

3.6.5.3.2 EoMPLS Header Bypass Frame Format

EoMPLS Header bypass is an extension to the frame handling of the standard MACsec frame format that allows an additional proprietary header in front of the MAC frame. The following illustration shows the format of the frame.

Figure 41 • EoMPLS Header Bypass Format

The following restrictions are applied to the EoMPLS header bypass format.

- The mode is statically controlled by programming the size of the bypassed header. Size of zero indicates absence of the bypassed header.
- No other secure format is possible on the port when header bypass is enabled.
- 2B–16B field is always one size on a port, configurable to be 2, 4, 6...16 Bytes. A static configuration register specifies size of the field. For EoMPLS this is generally configured as multiple of 4B.

The following logic is used to process the EoMPLS header bypass format.

- **Control Packet Detection.** Based on Etype2 matching a configured (static) value.
 - If Etype2 matches, detect and process control packets using MAC_DA, MAC_SA, and parsed Etype (after MAC_SA) to detect EAPOL/MKA transported in MPLS tunnels.
 - Other Etype2 values, detect and process control packets using MAC_DA2, MAC_SA2, and Etype2 to detect any other MAC control frames.
- **SecTAG Position.** Determined by size of 2B–16B field and located right after it.
- **Egress SA Match.** Uses Etype2, up to first 64 bits of the 2B–16B field, MAC_DA, and MAC_SA. The 2B–16B match field in the SA is bit-maskable.
- **Ingress SA Match.** Uses Etype2, up to first 64 bits of the 2B–16B field, MAC_DA, MAC_SA, and SecTAG fields. The 2B–16B match field in the SA is bit-maskable.

3.6.5.4 MACsec Transform

The MACsec transform carries out the actual frame transformation. For egress MACsec operations it inserts the SecTAG, optionally encrypts the payload data, and appends the ICV. For ingress MACsec operations, it removes the SecTAG, optionally decrypts the payload data, and removes and validates the ICV. The MACsec transform stage can detect error conditions (such as sequence number and authentication errors) that cause the frame to be dropped by applying a flow define drop_action.

The MACsec transform does not detect errors in the SecTag that the MACsec classification logic can catch. Only packets that are classified as tagged (valid non-KaY tag) may be submitted to ingress MACsec processing.

The MACsec transform stage uses the MACsec crypto engine for the actual MACsec transform, which operates in the following two major modes.

- In MACsec mode, the crypto engine is active and MACsec transforms can be performed.
- In static bypass mode, the crypto engine is effectively bypassed, which leads to a lower system latency. In this mode, no MACsec transforms are possible. The classification, consistency checking, and MTU check logic are still functional and the MACsec block may still filter (pass or drop) frames.

The MISC_CONTROL register enables static bypass, controls the latency equalization function, allows MACsec-compliant handling of MACsec frames for which no MACsec SA is available, and controls the maximum size of transform record.

If a MACsec SecY receives a MACsec frame on the common port for which it has no SA and the frame payload is unchanged (authenticate-only operation, C = 0, E = 0), it can still forward the frame to the controlled port without checking the authentication simply by stripping the SecTAG and ICV. This will occur if all the following conditions are met.

- The frame is classified as tagged.
- The frame is not matched by any installed MACsec SAs. The frame may either match no SA flow at all or a non-MACsec SA flow.
- The flow type of SAM_NM_FLOW_CP/SAM_NM_FLOW_NCP (whichever is applicable to that packet) for tagged frames is set to bypass.
- The TCI field has C = 0, E = 0.
- The nm_macsec_en bit is set.
- The validate_frames setting is either disabled or check, but not strict.

3.6.5.4.1 MACsec Context and Transform Record

The MACsec block contains an array of MACsec transform records that correspond to the number of supported SAs. Each transform record is 20×32 -bit words (80 bytes) in size in the ingress direction and 24×32 -bit words (96 bytes) in size in the egress direction, and corresponds to the SA flow with the same index. The MACsec transform operation is fully specified by a combination of the contents of the SAM_FLOW_CTRL_IGR/EGR register and the contents of the transform record. It corresponds to the operation of a single MACsec SA.

Transform record refers to the data structure as stored in the array. MACsec context refers to the information contained in a transform record. Transform record data are stored in the XFORM_RECORD_REGS registers. The following tables show the format for each transform record.

Table 10 • Transform Record Format (Non-XPN)

128 Bit block	128 Bit AES Keys		256 Bit AES key	
	Egress	Ingress	Egress	Ingress
0	CTRL Word	CTRL Word	CTRL Word	CTRL Word
	Context ID	Context ID	Context ID	Context ID
	Key0	Key0	Key0	Key0
	Key1	Key1	Key1	Key1
1	Key2	Key2	Key2	Key2
	Key3	Key3	Key3	Key3
	HashKey0	HashKey0	Key4	Key4
	HashKey1	HashKey1	Key5	Key5
2	HashKey2	HashKey2	Key6	Key6
	HashKey3	HashKey3	Key7	Key7
	Seq	Seq	HashKey0	HashKey0
	IV0	Mask1 ¹	HashKey1	HashKey1
3	IV1	IV0	HashKey2	HashKey2
	(Zero)	IV1	HashKey3	HashKey3
	(Zero)	(Zero)	Seq	Seq
	(Zero)	(Zero)	IV0	Mask1
4			IV1	IV0
			(Zero)	IV1
			(Zero)	(Zero)
			(Zero)	(Zero)

1. For MACsec, MASK is an unsigned integer controlling a valid range of packet numbers.

Table 11 • Transform Record Format (XPN)

128 Bit block	128 Bit AES Keys		256 Bit AES key	
	Egress	Ingress	Egress	Ingress
0	CTRL Word	CTRL Word	CTRL Word	CTRL Word
	Context ID	Context ID	Context ID	Context ID
	Key0	Key0	Key0	Key0
	Key1	Key1	Key1	Key1
1	Key2	Key2	Key2	Key2
	Key3	Key3	Key3	Key3
	HashKey0	HashKey0	Key4	Key4
	HashKey1	HashKey1	Key5	Key5
2	HashKey2	HashKey2	Key6	Key6
	HashKey3	HashKey3	Key7	Key7
	Seq Low	Seq Low	HashKey0	HashKey0
	Sec High	Sec High	HashKey1	HashKey1
3	DUMMY	Sec Mask	HashKey2	HashKey2
	IS0 (Salt)	IV0 (Salt)	HashKey3	HashKey3
	IS1 (Salt)	IV1 (Salt)	Seq Low	Seq Low
	IS2 (Salt)	IV2 (Salt)	Sec High	Sec High
4	IV0 (SCI)		DUMMY	Sec Mask
	IV1 (SCI)		IS0 (Salt)	IV0 (Salt)
			IS1 (Salt)	IV1 (Salt)
			IS2 (Salt)	IV2 (Salt)
5			IV0 (SCI)	
			IV1 (SCI)	

All fields of the transform record must be populated by the host software before the corresponding SA flow can be enabled. The `ctx_size` bit in the `CONTEXT_CTRL` register controls the size of the context that must be fetched. For bypass and drop flows, the transform record is not used. The hardware only updates the sequence number field; it does not modify the other fields during MACsec egress and ingress processing.

The context control word is the first 32-bit word in each transform record. It specifies the type of operation. Only those settings that are relevant for MACsec operations need to be defined. The following table shows the fields in context control word.

Table 12 • Context Control Word Fields

Bits	Name	Description
3:0	ToP	Type of packet 0110b: Egress 1111b: Ingress All other values are invalid
4	Reserved	Write with zero and ignore on read
5	IV0	First word of IV present in context (SCI for MACsec) Must be set to 1b

Table 12 • Context Control Word Fields (continued)

Bits	Name	Description
6	IV1	Second word of IV present in context (SCI for MACsec) Must be set to 1b
7	IV2	Third word of IV present in context (use sequence number instead) Must be set to 0b
12:8	Reserved	Write with zero and ignore on read
13	Updated Seq	Update sequence number Must be set to 1b for MACsec
14	IV Format	If set, use sequence number as part of IV Must be set to 1b for MACsec
15	Encrypt Auth	If set, encrypt ICV Must be set to 1b for MACsec
16	Key	Load crypto key from context Must be set to 1b for MACsec
19:17	Crypto Algorithm	Algorithm for data encryption 101b: AES CTR 128 111b: AES CTR 256
20	Reserved	Write with zero and ignore on read
22:21	Digest Type	Type of digest key Only single digest key is supported, setting 10b
25:23	Auth Algorithm	Algorithm for authentication Only AES-GHASH is supported, setting 100b
27:26	AN	The two-bit Association Number inserted in the SecTag for egress operations Must be kept 00b for ingress
29:28	Seq type	Type of sequence number: only supported setting is 01b Use 32-bit sequence number, on ingress use the mask as a replay window size
30	Seq mask	Sequence mask is present in context 0b: Egress 1b: Ingress
31	Context ID	Context ID present: must be set to 1b

The following list shows the other fields of the transform record.

- **Context ID.** Unique identifier for each context. It is sufficient to give all transform records a different context ID, possibly by assigning them a number from 0 to maximum index.
- **Key 0 ... Key 7.** AES encryption key for the MACsec SA. Each word of the key is a 32-bit integer representing four bytes of the key in little-endian order. The number of words depends on AES key length.
- **H_Key 0, H_Key 1, H_Key 2, and H_Key 3.** 128-bit key for the authentication operation. It is represented in the same byte order as Key 0...Key 7. It is derived from Key 0...Key 7 as follows: $H_key = E(Key, 128'h0)$. This means performing a 128/256 bit AES-ECB block encryption operation with Key 0...Key 7 as the key and a block of 128 zero bits as the plaintext input. The cipher-text result of the AES block encryption is the 128-bit H_Key.
- **Sequence Number.** For egress MACsec, this is one less than the sequence number (PN) that is to be inserted into the MACsec frame. For a new SA, this must be initialized to 0. After each egress packet, this field is incremented by 1. If it rolls over from 0xFFFFFFFF to 0, a sequence number error occurs and the context is not updated, which means that the same error will occur again for any

subsequent egress packets with that context - the external system will forward these packets to the line with CRC/packet error. For ingress MACsec the sequence number must be initialized to 1.

- **Mask (replay window size).** Window size for ingress sequence number checking. By default, it is 0 (strict ordering enforced). It can be set to any integer value up to $2^{32}-1$, in which case any nonzero sequence number is accepted.
- **SCI 0 and SCI 1.** SCI that belongs to the specific MACsec SA. An SCI that depends on the source MAC address and the ES and SCB bits is defined, even in modes that do not explicitly transmit or receive the SCI with each packet. This is a 64-bit block, represented by two 32-bit integers in little-endian order. It is the same byte order in which SAM_SCI_MATCH_HI/SAM_SCI_MATCH_LO represent an SCI.

When the sequence number of an egress SA is about to roll over, it must be replaced by a new SA with different keys. It is not allowed to reset the sequence number of an egress SA to a lower value because doing so generally leads to sequence number checking failures at the receiving end of the connection.

For inbound frames, the PN is compared against the sequence number (PN) from the context, resulting in one of the following three cases:

- If the received number is above or equal to the number in the context:
{received_PN ≥ next_PN}

In this case the context sequence number (PN) is updated (if the update_seq bit is set to 1b). The updated value is the received number plus one.

- If the received number is below the number from the context, but within the replayWindow:
{received_PN < next_PN and received_PN ≥ (next_PN - replayWindow)}

In this case no context update is required.

- If the received number is below the number from the context, and outside the replayWindow:
{received_PN < (next_PN - replayWindow)}

In this case the sequence number check fails and error bit e10 is set in the result token. No context update is done.

3.6.5.4.2 MACsec Crypto Engine Interrupt Control/Status Register

The INTR_CTRL_STATUS register provides control and status for interrupts within the MACsec crypto engine only. The interrupt output pin controlled here is one of the inputs on the top-level Advanced Interrupt Controller (controlled using the AIC registers).

The following main interrupts are given by the Crypto engine.

- Bit 4 Outbound Sequence Number Threshold

This interrupt is triggered if a sequence number exceeds the programmed sequence number threshold (specified in SEQ_NUM_THRESH) due to an outbound sequence number increment.

- Bit 5 Outbound Sequence Number Roll-over

This interrupt is triggered if a sequence number rolls over (increment from maximum to zero) due to an outbound sequence number increment.

3.6.5.5 Ingress Consistency Checking

Consistency checking is used to verify that MACsec ingress packets satisfy certain properties after decryption. Packets are passed or dropped based on a set of rules. The number of rules is a fixed hardware parameter. As opposed to the static classification and flow lookup stages, consistency checking logic inspects the packet data after the MACsec transform.

Consistency checking logic contains a complete VLAN tag parser performing the same operations as the VLAN tag parser located in the input packet classification logic. The configuration of the parser is controlled by a separate set of registers (IG_CP_TAG, IG_PP_TAGS, IG_PP_TAGS2, and IG_CP_TAG2) similar to the input packet VLAN tag parser. It extracts the payload Ethertype from the second or third Ethertype location in the packet if that packet contains one and two VLAN tags respectively. The VLAN tag parser also extracts (and post-processes) the following fields:

- User Priority field from the first VLAN tag it encounters, to be used by the MTU checking logic and statistics counters update logic.

- VLAN ID and VLAN Up from the second VLAN tag in case of Q-in-Q.

Each consistency check rule can match on a set of mask-able match criteria. If the corresponding mask bit is cleared, the match criterion is not checked and packets can satisfy the rule regardless of the value in the packet. If the corresponding mask bit is set, a packet only satisfies the rule if the value in the packet matches the value in the rule. The following list shows the mask-able match criteria.

- **sai_hit (1b or 0b)**. The packet was matched by one of the SA flows during flow lookup.
- **sai_nr (range 0 to SAmx-1)**. The packet was matched by the specific SA flow (or is not matched by any SA flow and has a specific combination of control packet and MACsec tag classification). To match packets that were matched by a specific SA flow, also match on `sai_hit = 1`
- **vlan_valid (1b or 0b)**. The packet contains a valid VLAN tag.
- **vlan_id (12 bits value)**. The packet has the specified VLAN ID. A match on this criterion is only meaningful if also matched on `vlan_valid = 1`.
- **vlan_id_inner (12 bits value)**. The packet has the specified VLAN ID at second VLAN tag. A match on this criterion is only meaningful if also matched on `vlan_valid = 1` and Q-in-Q is detected.
- **vlan_up_inner (3 bits value)**. The packet has the specified VLAN Up at second VLAN tag. A match on this criterion is only meaningful if also matched on `vlan_valid = 1` and Q-in-Q is detected.
- **etype_valid (1b or 0b)**. The Ethertype is greater than `cp_etype_max_len`.
- **payload_e_type (16 bits value)**. The packet has a specific Ethertype if a VLAN packet is detected, this value is the Ethertype following the VLAN tag.
- **ctrl_packet (1b or 0b)**. The packet is a control packet.

If all mask bits are cleared, the rule will match every possible packet.

Each of the consistency check rules can be enabled or disabled individually. If a rule is disabled, it will not be selected for match checking. If more than one enabled rule is matched, the one with the highest priority (3 bit number from 0 to 7) is picked. The lowest numbered rule is picked from equal priority rules.

The rule that is eventually selected specifies either a pass or a drop action.

If no rules match, the default action is taken (pass or drop). It is possible to define different default actions for control and non-control packets. After reset, the default action for both of them is drop.

ICC rule configuration is controlled by the `IG_CC_PARAMS` and `IG_CC_PARAMS2` registers.

3.6.5.6 Output Post-Processor

The final stage of the pipeline is the output post-processor. It implements the post-processing decision tree that includes MACsec-compliant post-processing, as well as processing and MTU checking for non-MACsec frames. It can drop the frame due to error conditions detected by the MACsec transform stage (such as sequence number rollover and authentication failure), it checks for the correct combinations of port numbers, it checks the frame length against the MTU, and it updates all statistics counters. For ingress packets, the post-processor uses results of the consistency checking module's VLAN tag detection logic instead of the VLAN parser in front of the MACsec crypto-engine.

The post-process statistics updating is done in accordance with Fig. 10-4 and 10-5 for secure frame generation and secure frame verification management control and frame counters.

3.6.5.6.1 MTU Checking

Registers provide MTU limit values for VLAN tagged frames (per User Priority as provided by the consistency checking module) and one global MTU limit value for non-VLAN frames (detected by the consistency checking module). The limits programmed are also used for statistics counters that rely on an MTU value.

Ingress Frame MTU Checking

- The frame length is the size of the input frame (including header and excluding Ethernet preamble, start-of-frame byte, and CRC).
- The VLAN User Priority is extracted from the VLAN tag as parsed (and post-processed) by the VLAN tag parser implemented in the consistency checking logic.

Egress Frame MTU Checking

- The frame length is the size of the output frame (including header and excluding Ethernet preamble, start-of-frame byte, and CRC).

- The VLAN user priority is the one provided by the VLAN parsing logic in the static classification logic. MTU checking is configured using the VLAN_MTU_CHECK and NON_VLAN_MTU_CHECK registers.

3.6.5.6.2 Statistics

The following two types of statistics counters are used.

- Per-frame counters are 40 bits wide. They overflow after about 10^{12} frames. A MACsec block processing 10 Gbps traffic can process in the order of 10^7 frames per second so that the 40-bit counters only saturate after 10^5 seconds (one day).
- Per-octet counters are 80 bits wide. They overflow after about 10^{24} octets. Even for a system that processes in excess of 10^9 bytes per second, this means that they will never overflow during the expected lifetime of the system.

The statistics counters can be configured to be auto-cleared on read. Also they can be configured to saturate at maximum value instead of rolling over.

There are three classes of statistics counters, as follows:

- Global Statistics.** The MACsec block maintains global statistics counters to implement MACsec. Some global statistics are maintained per-SA, so they must be obtained by accumulating (summing) the per-SA statistics of the relevant SAs.
- Per-SA Statistics.** The MACsec block maintains all per-SA statistics for ingress and egress MACsec operations. Software maintains statistics for all four SAs that might belong to an SC. It keeps the per-SA statistics, even for SAs that it has deleted from the SA flow table. When an SA flow is deleted, its final SA statistics must be collected and added into the per-SA and per-SC statistics.
- Per-SC Statistics.** The MACsec block does not maintain any per-SC statistics. However, the per-SC statistics are the sum of per-SA statistics of the SAs belonging to that SC. Whenever the software reads per-SA statistics from the hardware, it must not only add them to the per-SA statistics administration, but to the per-SC statistics administration as well.

The following tables show the per SA (per SC), global (SecY), and per user priority egress statistics generated. Eight sets of user priority counters are implemented. If a frame is detected as VLAN it also increments user priority counters in addition to per-SA/global (SecY) counters.

Table 13 • Egress SA Counters

Egress SA STAT Counters	Size
sa.OutOctetsEncrypted/sa.OutOctetsProtected	80
sa.OutPktsEncrypted/sa.OutPktsProtected/sa.OutPktsHitDropReserved	40
sa.OutPktsTooLong (MTU check)	40
sa.ifOutBroadcast	40
sa.ifOutMulticast	40
sa.ifOutUnicast	40

Table 14 • Egress Global Counters

Egress Global Counters	Size
global.TransformErrorPkts	80
global.OutPktsCtrl	80
global.OutPktsUnknownSA	40
global.OutOverSizePkts (MTU check)	40
global.ifOutBroadcast	40
global.ifOutMulticast	40
global.ifOutUnicast	40

Table 14 • Egress Global Counters (continued)

Egress Global Counters	Size
global.ifOutOctets	40

Table 15 • Egress Per-User Global Counters

Egress Global Counters	Size
Vlan.OutOctetsUP	80
Vlan.OutPktsUP	40
Vlan.OutDroppedPktsUP	40
Vlan.OutOverSizePktsUP	40

The following tables show the per SA (per SC), global (SecY), and per user priority ingress statistics generated. Eight sets of user priority counters are implemented. If a frame is detected as VLAN, it also increments user priority counters in addition to per-SA/global (SecY) counters.

Table 16 • Ingress SA Counters

Ingress SA STAT Counters	Size
sa.InOctetsDecrypted/sa.InOctetsValidated	80
sa.InPktsUnchecked/sa.InPktsHitDropReserved	40
sa.InPktsDelayed	40
sa.InPktsLate	40
sa.InPktsOk	40
sa.InPktsInvalid	40
sa.InPktsNotValid	40
sa.InPktsAuthFail ¹	40
sa.InPktsNotUsingSA	40
sa.InPktsUnusedSA	40
sa.InPktsSAMiss ⁽¹⁾	40
sa.InPktsUntaggedHit	40
sa.ifInBroadcast	40
sa.ifInMulticast	40
sa.ifInUnicast	40

1. Implemented indirectly. sa.InPktsAuthFail is reported in software by adding sa.InPktsInvalid and sa.InPktsNotValid. sa.InPktsSAMiss is reported in software by adding sa.InPktsNotUsingSA and sa.InPktsUnusedSA.

Table 17 • Ingress Global Counters

Ingress Global Counters	Size
global.TransformErrorPkts	80
global.InPktsCtrl	80
global.InPktsNoTag	40
global.InPktsUntagged	40

Table 17 • Ingress Global Counters (continued)

Ingress Global Counters	Size
global.InPktsTagged	40
global.InPktsBadTag	40
global.InPktsUntaggedMiss	40
global.InPktsNoSCI	40
global.InPktsUnknownSCI	40
global.InPktsSCIMiss*	
global.InConsistCheckControlledNotPass	40
global.InConsistCheckUncontrolledNotPass	40
global.InConsistCheckControlledPass	40
global.InConsistCheckUncontrolledPass	40
global.InOverSizePkts	40
global.ifInBroadcast	40
global.ifInMulticast	40
global.ifInUnicast	40
global.ifInOctets	40

Table 18 • Ingress Per-User Global Counters

Egress Global Counters	Size
Vlan.OutOctetsUP	80
Vlan.OutPktsUP	40
Vlan.OutDroppedPktsUP	40
Vlan.OutOverSizePktsUP	40

3.6.5.7 Correlation with 802.1AE MACsec Statistics

The following table shows how the MACsec block statistics are derived from the MACsec standard.

Table 19 • 802.1AE Correlation

MACsec name (802.1AE)	Direction	Type	Microsemi MACsec register
Frame verification statistics (MACsec specification 10.7.9)			
InPktsUntagged	Ingress	Global	global.InPktsUntagged
InPktsNoTag	Ingress	Global	global.InPktsNoTag
InPktsBadTag	Ingress	Global	global.InPktsBadTag
InPktsUnknownSCI	Ingress	Global	global.InPktsUnknownSCI
InPktsNoSCI	Ingress	Global	global.InPktsNoSCI
InPktsOverrun	Ingress	Global	Not implemented, condition does not occur, report as zero.
InPktsUnchecked	Ingress	Per-SC	sa.InPktsUnchecked
InPktsDelayed	Ingress	Per-SC	sa.InPktsDelayed
InPktsLate	Ingress	Per-SC	sa.InPktsLate

Table 19 • 802.1AE Correlation (continued)

MACsec name (802.1AE)	Direction	Type	Microsemi MACsec register
InPktsOK	Ingress	Per-SC, per-SA	sa.InPktsOK
InPktsInvalid	Ingress	Per-SC, per-SA	sa.InPktsInvalid
InPktsNotValid	Ingress	Per-SC, per-SA	sa.InPktsNotValid
InPktsNotUsingSA	Ingress	Per-SC, per-SA	sa.InPktsNotUsingSA
InPktsUnusedSA	Ingress	Per-SC, per-SA	sa.InPktsUnusedSA
Frame validation statistics (MACsec specification 10.7.10)			
InOctetsValidated	Ingress	Global	Accumulate over each ingress SA with authentication only: sa.InOctetsDecrypted/Validated
InOctetsDecrypted	Ingress	Global	Accumulate over each ingress SA with encryption: sa.InOctetsDecrypted/Validated
Frame generation statistics (MACsec specification 10.7.18)			
OutPktsUntagged	Egress	Global	global.OutPktsUntagged
OutPktsTooLong	Egress	Global	Accumulate over each egress SA: sa.OutPktsTooLong
OutPktsProtected	Egress	Per-SC, per-SA	sa.OutPktsEncrypted/Protected if the SA is authenticate only.
OutPktsEncrypted	Egress	Per-SC, per-SA	sa.OutPktsEncrypted/Protected if the SA uses encryption
Frame protection statistics (MACsec spec 10.7.19)			
OutOctetsProtected	Egress	Global	Accumulate over each egress SA with authentication only: sa.OutOctetsEncrypted/Protected
OutOctetsEncrypted	Egress	Global	Accumulate over each egress SA with encryption: sa.OutOctetsEncrypted/Protected

3.6.5.8 Interrupts

The MACsec block can raise five interrupts from ingress and four from the egress block. The available interrupts are as follows:

3.6.5.8.1 MACsec Crypto-Core Interrupt

Indicates several errors detected by the MACsec crypto engine block. The software must read the INTR_CTRL_STATUS register of the MACsec crypto core to see which condition caused the interrupt. The software must then write the same bits to INT_CTRL_STATUS to clear the interrupt condition, as applicable.

- Input error (bit 0) may occur if the MACsec crypto core attempts to process certain malformed short MACsec packets where the packet is shorter than indicated by the SL field.
- Output error and fatal error (bits 1 and 14) indicate a hardware error.
- Processing error (bit 2) may indicate a hardware error, but more likely the flow type in SAM_FLOW_CTRL is inconsistent with the context control word in the transform record (MACsec ingress versus MACsec egress).
- Context error (bit 3) indicates an error in the transform record, probably the context control word, especially the settings for encryption and authentication algorithms.

- Sequence number threshold (bit 4) indicates that an egress flow has exceeded its sequence number threshold. The MACsec SA must be re-keyed to prevent a sequence number rollover. Exceeding the sequence number threshold will not affect packet processing; it is meant to be used as a warning for imminent sequence number rollover.
- Sequence number rollover (bit [5]) indicates that an egress flow has encountered a sequence number rollover. The software must look in the transform record table to see which active egress SA has a sequence number value of 0xFFFFFFFF in case of 32-bit Packet number or 0xFFFFFFFF_FFFFFFFF in case of 64-bit packet number. This egress SA flow must immediately be disabled and it must be re-keyed.

Use the following steps to make effective use of the sequence number threshold interrupt.

1. Set the SEQ_NUM_THRESHOLD register to an appropriate value. A suitable value might be 0xF0000000 for a 32-bit packet number.
2. Make sure the sequence number threshold interrupt is enabled.

Use the following steps if the sequence number threshold interrupt occurs.

1. Temporarily disable the sequence number threshold interrupt, then clear that interrupt bit.
2. Check all transform records of active egress SAs for a sequence number that is either over the threshold or close to it (any egress SA with a sequence number above 0xE0000000).
3. Start a re-keying procedure for all those SAs.
4. After re-keying has been completed (and new SAs are installed on both sides of the connection), re-enable the sequence number threshold interrupt.

3.6.5.8.2 Classification Drop Interrupt

Raised when a packet is dropped by the flow lookup logic where either the SA flow or the non-matching flow specifies a drop action.

3.6.5.8.3 Consistency Check Drop Interrupt (ingress only)

Raised when a packet is dropped by the ingress consistency checking logic.

3.6.5.8.4 Post-Processing Drop Interrupt

Raised when a packet is dropped by the post-processing stage for any other reason than MTU check failure. Ingress packets with an ICV check failure or sequence number check failure raise this interrupt.

3.6.5.8.5 MTU Check Drop Interrupt

Raised when a packet is dropped due to MTU check failure.

Note: Frequent packet dropping may indicate an attack attempt, a configuration error, or a software malfunction.

3.6.5.9 Updating the MACsec SA for Ingress

For synchronization purposes, the MACsec standard requires the lowestPN and the nextPN in an active SA to be updated to a greater value provided by the KaY (unless it is not already reached). This is achieved in the MACsec core by updating the sequence number in an active context to a greater value if the sequence number in the context did not reach this value. The lowest acceptable PN is implicitly updated assuming that the replay window size is not changed. The host must program next_pn_lower (and next_pn_upper for XPN flow) to the desired sequence number, must specify the flow for which the update should occur in next_pn_context_id register and should enable the update in enable_update register. MACsec core will clear this enable_update register once the transform record field is updated. If the sequence number is already equal or above the configured value, then no internal update is performed.

3.6.6 Debug Fault Code in FCS

Incrementing a counter for a packet may be a security failure in some cases. The SA_SECFAIL_MASK/GLOBAL_SECFAIL_MASK register can be used to configure which counter increments are regarded as security fail events. Debug functionality enables packets failing security check to be transmitted with corrupted FCS, which consists of a debug fault code to debug the security failing packet. The FCS of a frame failing security check is corrupted on the output. The corrupted FCS

field contains a fault code for debugging using a frame analyzer. The fault code uses 31 bits, with the last FCS bit reserved to make sure the FCS check fails.

The following table shows the FCS fault code for the 32 bits.

Table 20 • FCS Fault Codes

Bit	Description
31	Reserved to make sure that FCS check fails
30	SA hit
29:24	SA pointer If the SA-hit bit[30] is 0, then bits[29:27] are reserved, bit[26] indicates if the frame is classified as control frame, and bits[25:24] indicate the MACsec tag classification of the frame: 00b = untagged, 01b = tagged, 10b = bad tag, 11b = KaY tag
23:10	Global stat event vector
9:0	SA stat event vector

The following tables show the format of the ingress global and SA stat event vectors.

Table 21 • Ingress Global Stat Event Vector Format

Event Bit Position	Ingress Global Counter
0	global.TransformErrorPkts
1	global.InPktsCtrl
2	global.InPktsNoTag
3	global.InPktsUntagged
4	global.InPktsTagged
5	global.InPktsBadTag
6	global.InPktsUntaggedMiss
7	global.InPktsNoSCI
8	global.InPktsUnknownSCI
9	global.InConsistCheckControlledNotPass
10	global.InConsistCheckUncontrolledNotPass
11	global.InConsistCheckControlledPass
12	global.InConsistCheckUncontrolledPass
13	global.InOverSizePkts
14	global.ifInUcastPkts
15	global.ifInMulticastPkts
16	global.ifInBroadcastPkts
17	global.ifInOctets

Table 22 • Ingress SA Stat Event Vector Format

Event Bit Position	Ingress SA Stat Counter
0	sa.InOctetsDecrypted/InOctetsValidated
1	sa.InPktsUnchecked/InPktsHitDropReserved

Table 22 • Ingress SA Stat Event Vector Format (continued)

Event Bit Position	Ingress SA Stat Counter
2	sa.InPktsDelayed
3	sa.InPktsLate
4	sa.InPktsOk
5	sa.InPktsInvalid
6	sa.InPktsNotValid
7	sa.InPktsNotUsingSA
8	sa.InPktsUnusedSA
9	sa.InPktsUntaggedHit
10	sa.ifInUcastPkts
11	sa.ifInMulticastPkts
12	sa.ifInBroadcastPkts

The following tables show the format of the egress global and SA stat event vectors.

Table 23 • Egress Global Stat Event Vector Format

Event Bit Position	Egress Global Counter
0	global.TransformErrorPkts
1	global.OutPktsCtrl
2	global.OutPktsUnknownSA
3	global.OutPktsUntagged
4	global.OutOverSizePkts (MTU check)
5	global.ifOutUcastPkts
6	global.ifOutMulticastPkts
7	global.ifOutBroadcastPkts
8	global.ifOutOctets
13:9	Reserved: zeroes

Table 24 • Egress SA Stat Event Vector Format


Event Bit Position	Egress SA Stat Counter
0	sa.OutOctetsEncrypted/OutOctetsProtected
1	sa.OutPktsEncrypted/OutPktsProtected/OutPktsHitDropReserved
2	sa.OutPktsTooLong (MTU check)
3	sa.ifOutUcastPkts
4	sa.ifOutMulticastPkts
5	sa.ifOutBroadcastPkts
10:6	Reserved: zeroes

3.6.7 Capture FIFO

A 512-byte capture FIFO can be used to capture up to first 504 bytes for packets failing any security check. The security fail event can be used as a trigger. The FIFO can also be enabled to capture the first

packet of any given SA using the CAPT_DEBUG_TRIGGER_SA1/2 control. Multiple packets can also be captured and the maximum size of the packet to be captured is configured using CAPT_DEBUG_CTRL.MAX_PKT_SIZE. This FIFO can be programmed to capture frames from either egress or ingress direction (CAPT_DEBUG_CTRL.SIDE). Frames are captured after MACsec transformation. Software can view the FIFO as 32-bit wide and 128 deep. Each 32-bit location is accessible to CSR using CAPT_DEBUG_DATA (0 to 127). Each packet is captured in the FIFO with a 64-bit administration header. The following illustration shows the layout of multiple packets in the capture FIFO.

Figure 42 • Capture FIFO Layout



Each stored packet is preceded by a 64-bit administration header that contains the following information.

- **ADM_HDR0.** 22 bits reserved, 1 bit truncated, 9 bit pkt_size
- **Truncated (1 bit).** Indicates the packet is truncated and only a part of the packet is captured. The captured packet could be truncated because the packet could be bigger than the MAX_PKT_SIZE programmed by software to capture.
- **Pkt_size (9 bits).** Indicates the size of the captured packet in bytes.
- **ADM_HDR1.** 32-bit security fail debug code, see section 4.4.1.

The status of the capture FIFO can be accessed using the CAPT_DEBUG_STATUS register (PKT_COUNT, FULL, WR_PTR).

Use the following steps to capture frames.

1. Decide the SIDE and MAX_PKT_SIZE and program in CAPT_DEBUG_CTRL.
2. Enable the SA to capture the first packet. For enabling first packet capture on any SA, program CAPT_DEBUG_TRIGGER_SA1/SA2 = 0xFFFFFFFF. To enable first packet capture on SA index [0], program CAPT_DEBUG_TRIGGER_SA1 = 0x1
3. Enable the capture by programming CAPT_DEBUG_TRIGGER.ENABLE = 1.
4. Send frames.
5. Keep polling CAPT_DEBUG_STATUS to see if any frames have been captured (PKT_COUNT, FULL, WR_PTR).
6. If PKT_COUNT > 0, then frames have been captured, read CAPT_DEBUG_TRIGGER_SA1/SA2 to confirm if the packet for that SA has been captured. Bits will fall back to 0b automatically when a packet is captured for the SA.
7. Stop the capture by programming CAPT_DEBUG_TRIGGER.ENABLE = 0 to enable software to access the FIFO.
8. Read CAPT_DEBUG_DATA (0 to 127) to read the packet from the capture FIFO.


- The host MAC Tx can schedule a pause frame for transmission at the next frame boundary. The Tx FC buffer needs to be able to hold at least one jumbo frame until XOFF pause is scheduled so that it can continue to receive data downstream. The XOFF frame is then received by host/switch.
- The host device can only stop transmission at next frame boundary because it may have started transmitting a second jumbo frame.

The following configuration signals control the basic flow control mode.

- PAUSE_REACT_ENA.** Enables pause reaction and pause timer maintenance in egress flow control buffer. Set to 1.
- PAUSE_GEN_ENA.** Enables XON and XOFF pause frame signaling to host MAC based on XON and XOFF thresholds. Set to 1.
- INCLUDE_PAUSE_RCVD_IN_PAUSE_GEN.** Enables the optional OR and AND gate. Set to 1. If not enabled the pause gen signaling to host MAC is purely based on XOFF/XON thresholds.

The following illustration shows the sequence of events when a pause frame is received from host.

Figure 44 • Host Back-Pressure by Remote Link Partner




The following steps describe the sequence of events depicted in the illustration.

- Host experiences congestion in ingress and sends pause (XOFF) to line.
- Host MAC Rx receives pause frame. It is not enabled to react on received pause frames so it passes the pause frame to Tx FC buffer.
- Tx FC buffer maintains two logical queues, one for data and one for MAC control frames. If a data frame is already scheduled and in progress, it passes on MAC control frames at the next boundary to quickly relay MAC control frames to line, despite the presence of other data frames in the data queue.
- Tx FC buffer transmits any or all control frames in the control queue.
- Pause frame passes through the MACsec block. The MACsec egress block detects frame as a control frame and does not encrypt it. Frame eventually passes through the line MAC Tx block and the rest of the PHY blocks.

`TX_CTRL_QUEUE_ENA` determines if the control queue is enabled in the egress flow control buffer. This should be set to 1 in basic flow control mode. The physical memory of egress FC buffer can be partitioned between data and control queues using `TX_CTRL_QUEUE_START/END` and `TX_DATA_QUEUE_START/END` configuration fields.

3.6.8.2 Advanced Flow Control Handling

The following illustration shows the sequence of events when the PHY is configured to the advanced flow control mode of operation. `PAUSE_GEN_ENA` needs to be set to 1 and other configuration bits of FC buffer, such as `PAUSE_REACT_ENA`, `INCLUDE_PAUSE_RCVD_IN_PAUSE_GEN`, and `TX_CTRL_QUEUE_ENA`, need to be set to 0. All other configurations for this mode are part of line MAC and host MAC.

Figure 45 • Advanced Flow Control Handling

The following steps describe the sequence of events depicted in the illustration.

3.6.8.2.1 PHY Back-Pressured by Remote Link partner

1. Pause frame (XOFF) is received by PHY at line MAC Rx. This frame is internally consumed by MAC. Line MAC Rx signals line MAC Tx with pause received indication and pause quanta.
2. Line MAC Tx goes to pause state at the next frame boundary. Line MAC Tx stalls to pause the pipeline. Pause timer maintained by line MAC Tx is started only after it goes to pause state. The Tx FC buffer signals XOFF/XON to host MAC Tx based on XOFF/XON threshold.

3.6.8.2.2 Host Back-Pressuring Remote Link Partner

3. System pause is consumed by host MAC Rx. Pause timer maintained in host MAC Rx (instead of Tx) for egress direction to generate XOFF/XON pause gen signal for line MAC Tx.
4. Line MAC Tx stalls to send pause frame (either XOFF or XON). This path will work irrespective of whether line MAC is in pause state.

3.6.8.3 Frame Drop Statistics

The following 32-bit counters provide frame drop statistics. These counters roll over to 0 when the maximum value is reached.

3.6.8.3.1 TX_CTRL_QUEUE_OVERFLOW_DROP_CNT

Number of control frame drops due to overflow in the control queue of the egress flow control buffer.

3.6.8.3.2 TX_CTRL_QUEUE_UNDERFLOW_DROP_CNT

Number of control frame drops due to underflow in the control queue of the egress flow control buffer.

3.6.8.3.3 TX_CTRL_UNCORRECTED_FRM_DROP_CNT

Number of control frames aborted due to ECC check fail during reading from RAM in egress flow control buffer.

3.6.8.3.4 TX_DATA_QUEUE_OVERFLOW_DROP_CNT

Number of data frame drops due to overflow in the data queue of the egress flow control buffer.

3.6.8.3.5 TX_DATA_QUEUE_UNDERFLOW_DROP_CNT

Number of data frame drops due to underflow in the data queue of the egress flow control buffer.

3.6.8.3.6 TX_DATA_UNCORRECTED_FRM_DROP_CNT

Number of data frames aborted due to ECC check fail during reading from RAM in egress flow control buffer.

3.6.8.3.7 RX_OVERFLOW_DROP_CNT

Number of frame drops due to overflow in the ingress flow control buffer.

3.6.8.3.8 RX_UNDERFLOW_DROP_CNT.

Number of frame drops due to underflow in the ingress flow control buffer.

3.6.8.3.9 RX_UNCORRECTED_FRM_DROP_CNT.

Number of frames aborted due to ECC check fail during reading from RAM in ingress flow control buffer.

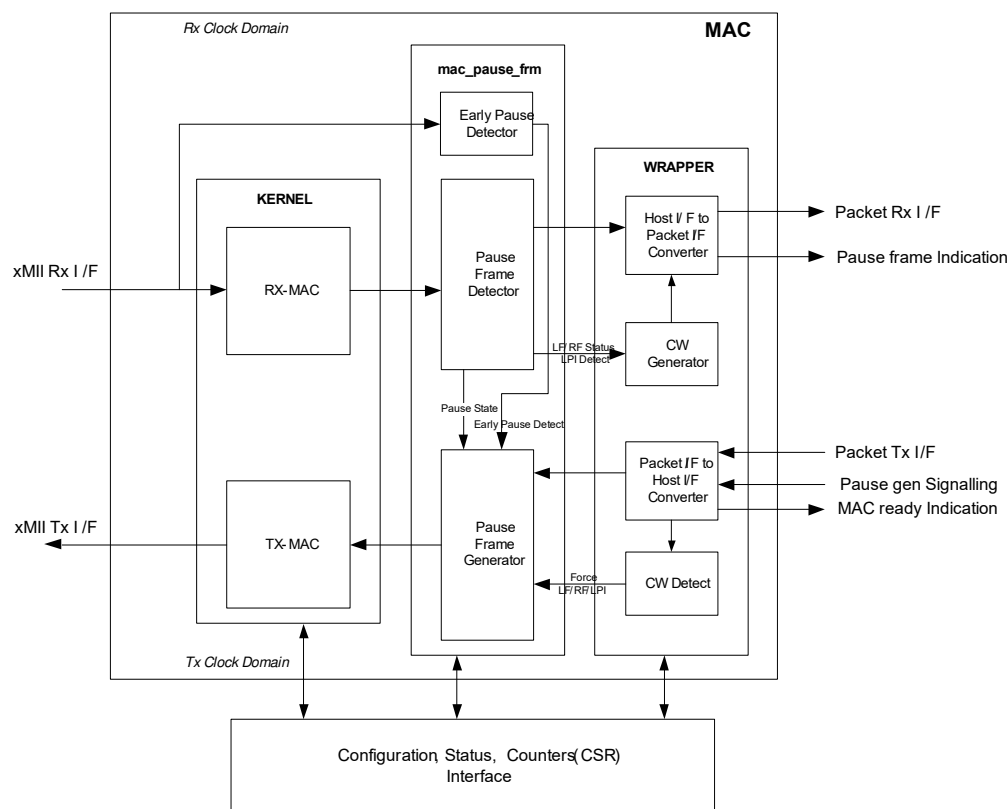
3.6.9 Media Access Control

This section describes the media access control sub layer (MAC) block. There are two instances of MAC block in each channel. One instance, which interfaces with MACsec and PCS/PMA, is called Line MAC and other instance which interfaces with FC Buffer and PHY XS is called Host MAC.

The MAC is defined in IEEE 802.3, clauses 3 and 4. The purpose of the MAC is to control the MACsec block access to the physical layer. In other words, it takes frames from the MACsec and converts those to a continuous byte stream on the xMII interface. In doing so, it is responsible for frame CRC generation and checking, preamble insertion and extraction, and pause frame generation and detection. The MAC block also contains the counters for an SNMP management information base (MIB) statistics module.

The MAC block supports frame sizes up to 10240 bytes in both receive and transmit directions. The maximum frame size is controlled by the host. The maximum frame size can also be set to the standard 1518 bytes or 1522 bytes, if desired. Maximum frame length restrictions are not enforced in the transmit direction. The following illustration shows the block diagram of MAC.

Figure 46 • MAC Block Diagram



3.6.9.1 MAC Transmit

The transmit section of the MAC contains three blocks, packet interface wrapper, pause frame generator, and MAC Tx kernel. All three blocks operate off the same clock, TX_MAC_CLK.

The MAC Tx kernel block handles the reconciliation sublayer functions as per IEEE 802.3.

- Calculates the CRC for pause frames generated by the pause frame generator.
- Converts MAC frames to the xMII format and adds control characters for framing as required by IEEE 802.3.
- Generates the interframe gap (IFG) on the xMII using the deficit idle count algorithm to achieve an average IPG of 12 bytes.
- Shapes all the traffic to go out with an average IPG of 12 bytes after MACsec frame expansion.
- Analyzes each packet and increments statistical counters used for RMON support.

The Pause Frame Generator (PFG) block performs the following two major functions.

- Requests packets from the upstream blocks, when packets are present and the Tx direction is not in the pause state (because a pause frame has been received in the Rx direction). They are forwarded to the MAC Tx kernel block for further processing.
- Generates flow control packets. Pause frames are generated based upon seeing the MAC_PAUSE_FRM_GEN signal. For the Host MAC this signal is generated by the FC buffer based upon programmable XOFF/XON threshold values in the FC buffer. In advance flow control mode of operation the line MAC can also generate pause frames based on MAC_PAUSE_FRM_GEN signal from Host MAC to relay pause frames that are deleted in Host MAC in this mode.

When the pause frame generator sees the MAC_PAUSE_FRM_GEN signal asserted, it generates pause frames using settings in configuration registers. Part of the pause frame is the pause value, which specifies how long the link partner (the network entity that the pause frame is destined for) stops sending traffic. The pause value specifies the requested delay in bit times and uses the equation $512 \times \text{PAUSE_VALUE}$.

After the PFG starts generating pause frames, it continues to generate pause frames at specified intervals until the de-assertion of the MAC_PAUSE_FRM_GEN signal. When this signal is deasserted, the PFG does one of two things, depending upon the configuration in MAC_TX_PAUSE_MODE. In normal mode, the PFG stops sending pause frames. This causes the link partner to start sending frames again after its pause frame timer has expired. In XON mode, the PFG generates a single pause frame with a pause value of 0 and sends it to the link partner. This causes the link partner to start sending frames again right away.

The PFG contains a configurable pause frame interval register, MAC_TX_PAUSE_INTERVAL. This register controls the time between generated pause frames when the FC buffer continues to request that pause frames be generated.

The packet interface wrapper handles the following functions:

- Provides the packet interfacing support to MACsec and FC buffer blocks. On this packet interface, frames are transported without preamble and FCS.
- Supports LF/RF/LPI generation on xMII interface through special control word received on packet interface. This special control word is received on packet interface if relaying of LF/RF/LPI is desired in MACsec subsystem.
- Padding of frames whose length is less than 64 bytes. This is required for padding of MACsec short length frames whose length is less than 64 bytes. This padding is enabled by configuring ENABLE_TX_PADDING in host MAC.
- Standard preamble insertion.
- FCS insertion.

3.6.9.2 MAC Receive

The receive section of the MAC contains three blocks, MAC Rx kernel, pause frame detector, and packet interface wrapper. All three blocks operate off the same clock, RX_MAC_CLK.

The MAC Rx kernel receives the byte stream from the xMII interface and handles the reconciliation sublayer processing to convert them to frames sent over the host interface. It checks the CRC of each frame for validity and aborts any frame with an invalid CRC. A variety of length checks are performed, including looking for short frames (less than 64 bytes), oversized, and jabber frames (longer than the configured maximum). VLAN tagging is supported up to three VLAN tags. Length checks are adjusted accordingly when VLAN tags are encountered. The Rx kernel supports counters in support of RMON statistics.

The pause frame detector (PFD) detects and reacts to valid pause frames received by the MAC from the xMII interface. The PFD reacts to PAUSE frames with a DMAC equal to either the multicast address (01-80-c2-00-00-01) or the address of the MAC (MAC_ADDRESS_LSB/MSB register value) in accordance with IEEE 802.3-2008, Annex 31B. Pause frames that are too short, or have invalid CRC, are abort marked and ignored by the PFD. Pause frames carry a pause value that indicates the desired pause time in units of pause quanta, where 1 pause-quantum equals 512 bit times. Because the data path in the MAC is 8 bytes (or 64 bits) wide, the extracted pause value is multiplied by 8 and stored in the pause counter. A signal from the PFG indicates if a packet is currently being transmitted.

After the current packet has completed or if there is no packet, the PFD tells the PFG to stop requesting packets (XOFF) and the pause counter is decremented by one for each MAC Rx clock cycle. When the counter reaches 0, the PFG is instructed that it may resume requesting packets from the upstream blocks. Pause frames must have a destination address equal to either the multicast address (01-80-c2-00-00-01) or the address of the MAC (MAC_ADDRESS_LSB/MSB register value). If there is no match, then the pause frame is ignored. If a pause frame is received while the Tx direction is already being paused (because a valid pause frame was already received and the pause counter had not yet counted down to 0), the pause counter is simply updated with the new value. If the received pause value is 0, then the state machine transitions immediately to END_PAUSE and frames are again requested from the upstream blocks.

The packet interface wrapper handles the following functions:

- Provides the packet interfacing support to MACsec and FC buffer blocks. On this packet interface, frames are transported without preamble and FCS.
- Supports LF/RF/LPI indication on packet interface through special control word. This special control word is relayed to other MAC if relaying of LF/RF/LPI is desired.
- Preamble strip on packet interface.
- FCS check and strip.

3.6.9.3 RMON Statistical Counters

The following counters count the number of bytes or frames received or transmitted. The counters count continuously and are only cleared if the device is reset or the counter is written with 0 through the CPU interface. These counters roll-over to 0 when the maximum value is reached. Unless specified otherwise, each counter is 32 bits.

- RX_IN_BYTES_CNT (40 bits) counts the total bytes received including preamble
- RX_OK_BYTES_CNT (40 bits) counts the number of bytes received in valid frames
- RX_BAD_BYTES_CNT counts the number of bytes received in invalid frames
- TX_OUT_BYTES_CNT (40 bits) counts the total number of bytes transmitted including preamble
- TX_OK_BYTES_CNT (40 bits) counts the number of bytes in successfully transmitted frames

The following counters are based on the type of frame received or transmitted.

- RX_PAUSE_CNT counts the number of pause frames received
- RX_UNSUP_OPCODE_CNT counts the number of control frames received with unsupported opcodes
- RX_UC_CNT counts the number of unicast frames received
- RX_MC_CNT counts the number of multicast frames received
- RX_BC_CNT counts the number of broadcast frames received
- TX_PAUSE_CNT counts the number of pause frames transmitted
- TX_UC_CNT counts the number of unicast frames transmitted
- TX_MC_CNT counts the number of multicast frames transmitted
- TX_BC_CNT counts the number of broadcast frames transmitted

The following error counters are provided.

- RX_SYMBOL_ERR_CNT counts the number of symbol errors received
- RX_CRC_ERR_CNT counts the number of frames received with CRC errors
- RX_UNDERSIZE_CNT counts the number of undersized frames received with valid CRC
- RX_FRAGMENTS_CNT counts the number of undersized frames received with invalid CRC
- RX_IN_RANGE_LENGTH_ERR_CNT counts the number of frames where the length field does not match the frame length

- RX_OUT_OF_RANGE_LENGTH_ERR_CNT counts the number of frames with an illegal length field
- RX_OVERSIZE_CNT counts the number of oversize frames with valid CRC
- RX_JABBERS_CNT counts the number of oversize frames with an invalid CRC
- RX_XGMII_PROT_ERR_CNT counts the number of XGMII protocol errors detected.

The following size histogram counters are provided for both transmit and receive directions.

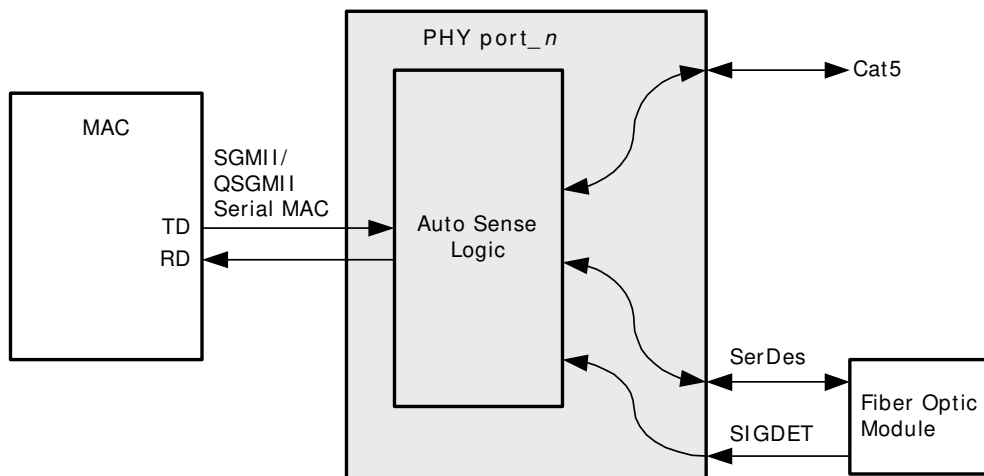
- Frames with 64-byte payloads
- Frames with 65-byte to 127-byte payloads
- Frames with 128-byte to 255-byte payloads
- Frames with 256-byte to 511-byte payloads
- Frames with 512-byte to 1023-byte payloads
- Frames with 1024-byte to 1518-byte payloads
- Frames with 1519-byte to maximum size payloads

Frame size counters also count invalid frames, as long as they are not short frames, fragments, long frames, or jabber frames. Long frames are defined as those greater than MAX_LEN bytes.

3.7 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 VSC8584-11 device.

Figure 47 • 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 non-preferred 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 25 • 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 14.

3.8 Reference Clock

The device reference clock supports both 25 MHz and 125 MHz clock signals. The IEEE 1588 differential input clock supports frequencies of 125 MHz to 250 MHz. Both reference clocks can be either differential or single-ended. If differential, they must be capacitively coupled and LVDS compatible.

3.8.1 Configuring the Reference Clock

The REFCLK_SEL2 pin configures the reference clock speed. The following table shows the functionality and associated reference clock frequency.

Table 26 • REFCLK Frequency Selection

REFCLK_SEL2	Frequency
0	25 MHz
1	125 MHz

3.8.2 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, with the clock centered at 1 V and a 500 mV peak-to-peak swing, are shown in the following illustrations.

Figure 48 • 2.5 V CMOS Single-Ended REFCLK Input Resistor Network




Figure 49 • 3.3 V CMOS Single-Ended REFCLK Input Resistor Network



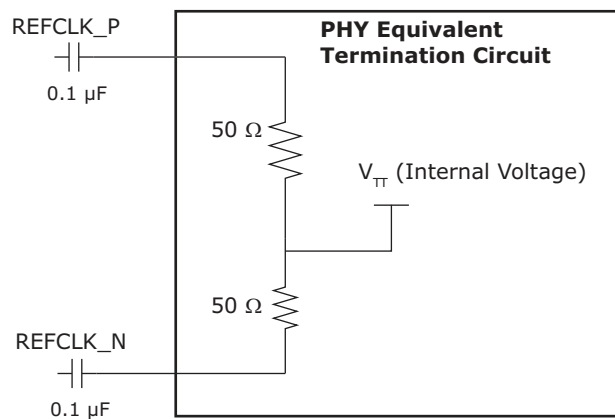


Figure 50 • 5 V CMOS Single-Ended REFCLK Input Resistor Network


Note: A single-ended 25 MHz reference clock is not guaranteed to meet requirements for QSGMII MAC operation.

3.8.3 Differential REFCLK Input

AC coupling is required when using a differential REFCLK. Differential clocks must be capacitively coupled and LVDS-compatible. The following illustration shows the configuration.

Figure 51 • AC Coupling for REFCLK Input


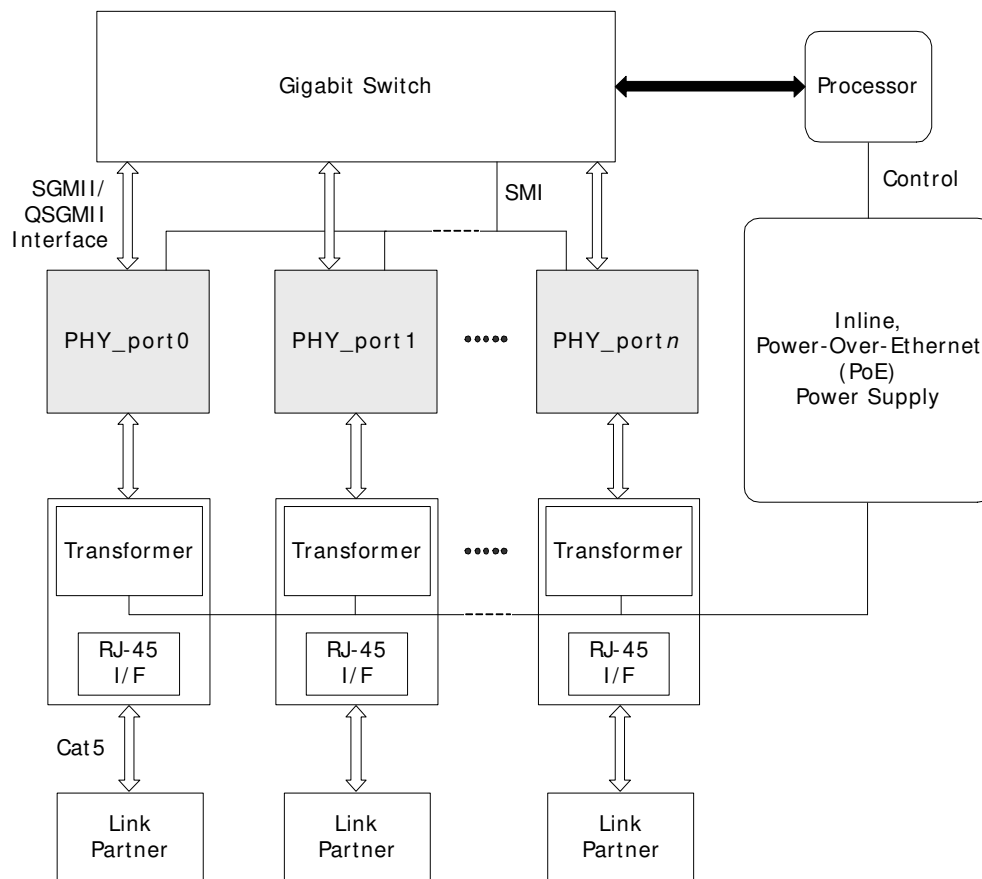
3.9 IEEE 1588 Reference Clock

The device IEEE 1588 reference clock input supports a continuum of frequencies between 125 MHz and 250 MHz. Both single-ended and differential clocks are supported, but differential clocks are preferred for better performance. If differential, they must be capacitively coupled and LVDS compatible. For more information about configuring the clock for single-ended operation, see [Reference Clock](#), page 57.

3.10 Ethernet Inline Powered Devices

The VSC8584-11 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 52 • 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 VSC8584-11 PHY using its serial management interface. Set register bit 23E1.10 to 1.
2. Ensure that the VSC8584-11 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 VSC8584-11 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 VSC8584-11 register bit 23E1.9:8 reads back 01. It can also be verified as an inline power detection interrupt by reading VSC8584-11 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, VSC8584-11 register bit 23E1.9:8 automatically resets to 10.
4. If the VSC8584-11 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 VSC8584-11 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 VSC8584-11 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 VSC8584-11 PHY disables its normal autonegotiation process and re-enables its inline powered device detection mode.

3.11 IEEE 802.3af PoE Support

The VSC8584-11 device is compatible with designs that are intended for use in systems that supply power to data terminal equipment (DTE) by means of the MDI or twisted pair cable, as described in IEEE 802.3af Clause 33.

3.12 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 VSC8584-11 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 VSC8584-11 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 53 • ActiPHY State Diagram



3.12.1 Low Power State

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

3.12.2 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 MDINT) functionality is provided.

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

3.12.3 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.

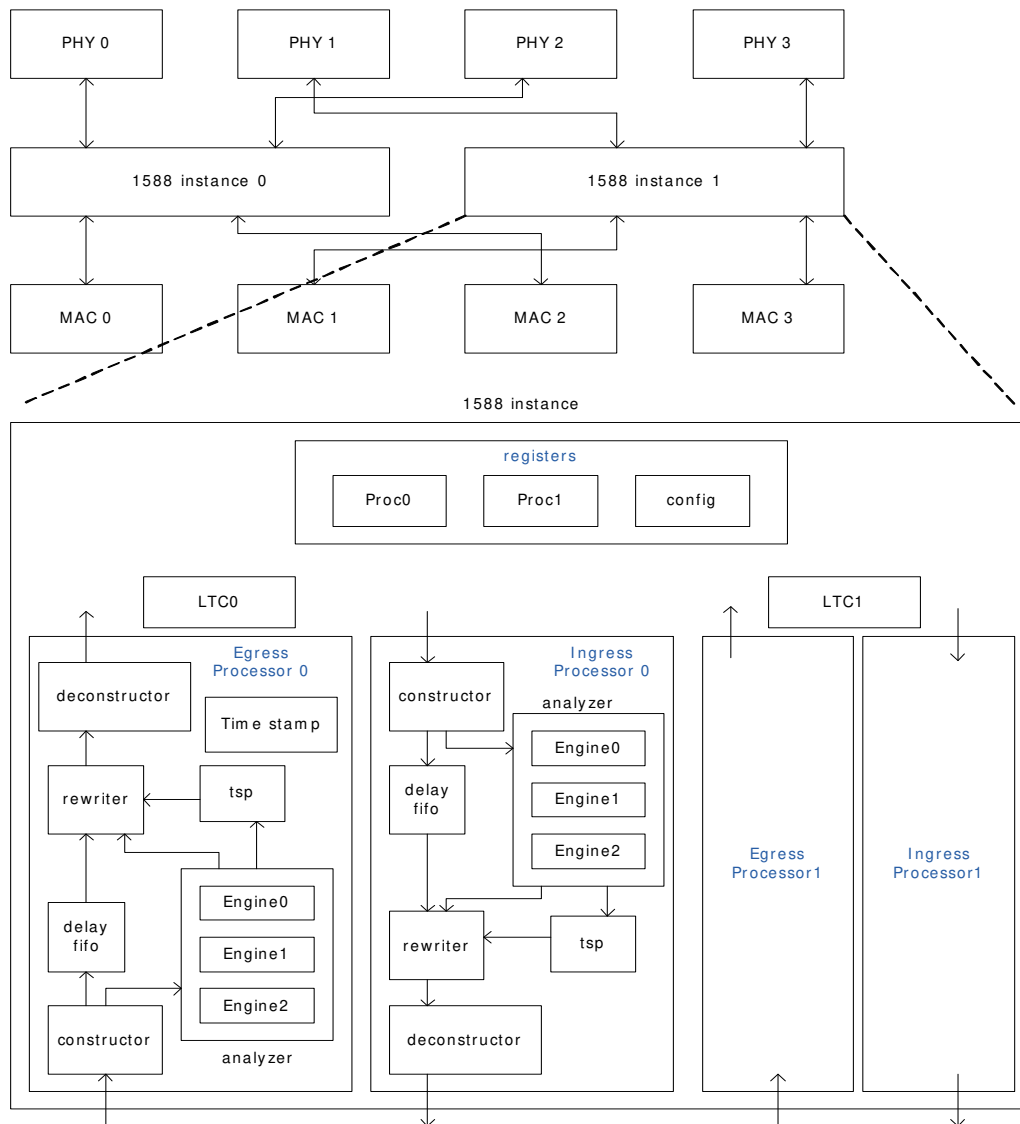
3.13 IEEE 1588 Block Operation

The VSC8584-11 device uses a second generation IEEE 1588 engine that is backward compatible with the earlier version of VeriTime™ (the Microsemi IEEE 1588 time stamping engine), both stand alone and in combination with MACsec. It is also compatible with the IEEE 1588 operations supported in Microsemi CE switches. The following list shows the new features of the Microsemi second generation IEEE 1588.

- MACsec support
- Increased time stamp accuracy and resolution
- Auto clear enables after system time is read or written
- Ability to load or extract the current system time in serial format
- Full 48-bit math support for incoming correction field
- Ability to add or subtract fixed offset from system time to synchronize between slaves
- Each direction of IEEE 1588 can be independently controlled and bypassed
- Support to extract frame signature in an IPv6 frame
- MPLS-TP OAM support in third analyzer engine
- Special mode where all frames traversing the system can be time stamped

The unique architecture of the MACsec and the second generation IEEE 1588 block combination provides for the lowest latency and maximum throughput on the channel. The following illustration shows a block diagram of the IEEE 1588 architecture in the VSC8584-11 device.

Figure 54 • IEEE 1588 Architecture



The following sections list some of the major IEEE 1588 applications.

3.13.1 IEEE 1588 Block

The IEEE 1588 engine may be configured to support one-step and two-step clocks as well as Ethernet and MPLS OAM delay measurement. It detects the IEEE 1588 frames in both the Rx and Tx paths, creates a time stamp, processes the frame, and updates them. It can add a 30/32-bit Rx time stamp to the 4-bytes reserved field of the PTP packet. It can also modify the IEEE 1588 correction field and update the CRC of changed frames. There are local time counters (reference for all time stamps) that can be preloaded and adjusted through the register interface.

A local time counter is used to hold the local time for Rx and Tx paths. A small FIFO delays frames to allow time for processing and modification. An analyzer detects the time stamp frames (PTP and OAM) and a time stamp block calculates the new correction field. The rewriter block replaces the correction field with an updated one and checks/calculates the CRC. For the Tx path, a time stamp FIFO saves Tx event time stamp plus frame identifier for use in some modes.

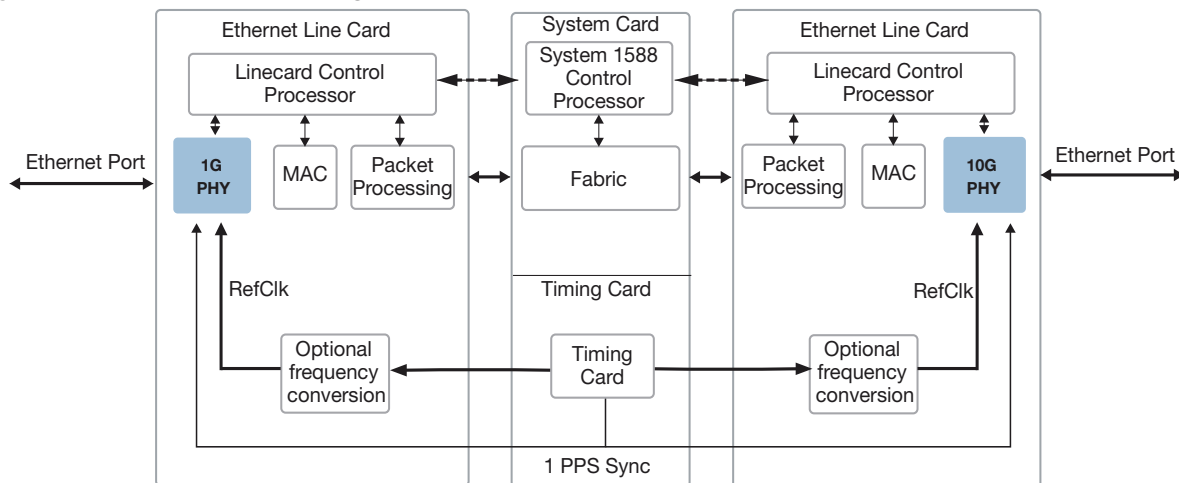
The IEEE 1588 engine's registers and time stamps are accessible through the MDIO or 4-pin SPI. To overcome the MDIO or 4-pin SPI speed limitations, the dedicated "push-out" style SPI output bus can be used for faster or large amounts of time stamp reads. This SPI output is used to push out time stamp

information to an external device only and does not provide read/write to the registers of the IEEE 1588 engine or registers of other blocks in the VSC8584-11 device. In addition, there is a LOAD/SAVE pin that is used to load the time in the PHYs and ensure that all the PHYs are in sync. The local time counter may come from any one of the following sources:

- Data path clock (varies according to mode)
- 250 MHz from host side PLL
- External clock from 1588_DIFF_INPUT_CLK_P/N pins

The local time counters contain two counters: nanosecond_counter and second_counter. The 1 PPS (pulse per second signal) output pin can be used for skew monitoring and adjustment. The following illustration shows an overview of a typical system using IEEE 1588 PHYs. The LOAD/SAVE and 1 PPS pins are signals routed to the GPIO pins. The following illustration shows how the PHY is embedded in a system.

Figure 55 • IEEE 1588 Block Diagram



The system card has to drive the REFCLK (125 MHz or 250 MHz timetick clock) to all the PHYs, including the VSC8584-11 device. The system clock may need local frequency conversion to match the required reference clock frequency. The system clock may be locked to a PRC by SyncE or by IEEE 1588. If locked by IEEE 1588, the central CPU recovers the PTP timing and adjusts the frequency of the system clock to match the PTP frequency. If the system clock is free running, the central CPU must calculate the frequency offset between the system clock and the synchronized IEEE 1588 clock and program the PHYs to make internal adjustments.

Other than the clock, the system card also provides a sync pulse to all PHYs, including the VSC8584-11 device, to the LOAD/SAVE pin. This signal is used to load the time to the PHYs and to ensure that all the PHYs are in sync. This may just be a centrally divided down system clock that gives a pulse at fixed time intervals. The delay from the source of the signal to each PHY must be known and taken into account when writing in the load time in the PHYs.

The VSC8584-11 device supports a vast variety of IEEE 1588 applications. In simple one-step end-to-end transparent clock applications, the VSC8584-11 device can be used without any central CPU involvement except for initial configuration. The IEEE 1588 block inside the VSC8584-11 device forwards Sync and Delay_req frames with automatic updates to the Correction field.

In other applications, the VSC8584-11 device enhances the performance by working with a central processor that runs the IEEE 1588 protocol. The VSC8584-11 device performs the accurate time stamp operations needed for all the different PTP operation modes. For example, at startup in a boundary clock application, the central CPU receives PTP sync frames that are time stamped by the ingress PHY and recovers the local time offset from the PTP master using the PTP protocol. It then sets the save bit in the VSC8584-11 device connected to the PTP master and later reads the saved time. The central CPU loads the expected time (time of the next LOAD/SAVE pulse, corrected by the offset to the recovered PTP time) into the PHY and sets the save bit. It checks that the time offset is 0. If not, it makes small adjustments to the time in the PHY by issuing add 1 ns or subtract 1 ns commands to the VSC8584-11

device through MDIO, until the time matches the PTP master. A save command is issued to the PHY connected to the PTP master and reads the saved time. The central CPU then writes the saved time plus the sync pulse interval plus any sync pulse latency variation (trace length difference compared to the PHY connected to the PTP master) to the other PHYs and sets the load bit in these VSC8584-11 devices.

The preceding sequence may be completed in several steps. Not all PHYs need to be loaded at once. The central CPU sets the save bit in all PHYs and reads back the values. They should all save the same value.

The central CPU continuously detects if the system time drifts off compared to the recovered PTP time. If needed, it can adjust each PHY for any known skew between PHYs without affecting the operation of the device. It can program the PHYs, including the VSC8584-11 device, to automatically add 1 ns or subtract 1 ns at specific time intervals.

3.13.2 IEEE 1588 One-Step E2E TC in Systems

Unique advantages for implementing IEEE 1588-2008 include:

- When several VSC8584-11 devices or Microsemi PHYs with integrated IEEE 1588 time stamping 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
- VSC8584-11 and other Microsemi PHYs with integrated IEEE 1588 time stamping blocks also work for pizza box solutions, where the switch/router can be upgraded to support IEEE 1588 E2E TC


Requirements for 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 pulse, that synchronizes the local time counters in each port.
- CPU access to each PHY to set up the required configuration. Can be MDC/MDIO or a dedicated CPU interface.

3.13.3 IEEE 1588 TC and BC in Systems

This is the same system as described previously, with the addition of a central IEEE 1588 engine (Boundary Clock). The IEEE 1588 engine is most likely a CPU system, possibly 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 IEEE 1588 engine for processing.

Figure 56 • TC and 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.

Requirements for 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 set up the required configuration. For one-step support this can be MDC/MDIO. For two-step support, a higher speed CPU interface (such as the SPI) might be required (depending on the number of time stamps 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 central IEEE 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 IEEE 1588 frames and redirect them to the central IEEE 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

3.13.4 Enhancing IEEE 1588 Accuracy for CE Switches and MACs

Connecting VSC8584-11 or other Microsemi PHYs that have integrated IEEE 1588 time stamping in front of the CE Switches and MACs improves the accuracy of the IEEE 1588 time stamp calculation. This is due to the clock boundary for the XAUI and SGMII/QSGMII interface. It will also add support for one-step TC and BC on the Jaguar-1 family of devices.

3.13.5 MACsec Support

MACsec is required when the physical link between two MACs must provide secure communication. MACsec PHYs such as the VSC8584-11 device are connected with CE switches to provide secure communication. PTP and OAM frames are recognizable only before or after encryption, meaning that the MACsec block must precede the IEEE 1588 block from the line inward.

Even though MACsec introduces large delay variation because of the insertion/removal of the MACsec header on all encrypted frames, the VSC8584-11 device provides the same accuracy with MACsec enabled as without. In all other aspects, the IEEE 1588 operation is as described in previous sections.

3.13.6 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 the other ports. On the master ports, Sync frames are transmitted from the IEEE 1588 engine that holds the Origin time stamp. These frames will have the correction field or the full Tx time stamp updated on the way out through 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 ingress through the PHY and the IEEE 1588 engine receives the Delay_req frame and generates a Delay_resp message. The Delay_resp messages are not event messages and are passed through the PHY as any other frame.

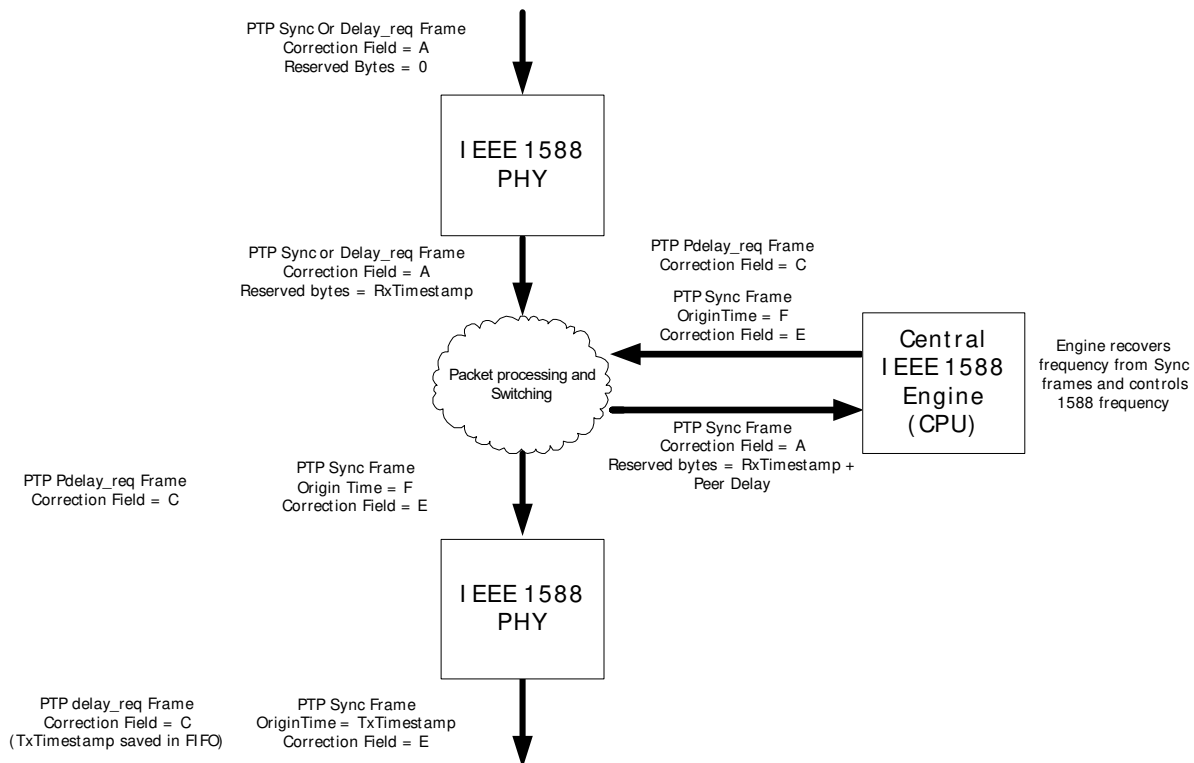
The port configured as slave receives Sync frames from its master. The Sync frames have an Rx time stamp added in the PHY and forwarded to the IEEE 1588 engine.

The IEEE 1588 engine also generates Delay_req frames that are sent on the port configured as slave port. Normally the transmit time for the Delay_req frames, t_3 , is saved in a time stamp FIFO in the PHYs, but when using Microsemi IEEE 1588 PHYs, a slight modification can be made to the algorithm to remove the CPU processing overhead of reading the t_3 time stamp.

To modify the algorithm, the IEEE 1588 engine should send the Delay_req message with a software generated t_3 value in the origin time stamp, the sub-second value of the t_3 time stamp in the reserved bytes of the PTP header and a correction field of 0. The software generated t_3 time stamp should be within a second before the actual t_3 time. The Egress PHY should then be configured to perform E2E TC egress operation, meaning calculate the "residence time" from the inserted t_3 time stamp to the actual t_3 time and insert this value in the correction field of the frame. When the local IEEE 1588 engine receives the corresponding Delay_resp frame back it can use the software generated t_3 value because the correction field of the Delay_resp frame contains a value that compensates for the actual t_3 transmission time.

Boundary clocks and ordinary clocks must also reply to Pdelay_req messages just as P2P TC using the same procedure for the P2P TC. For more information, see [Supporting One-Step Peer-to-Peer Transparent Clock](#), page 72.

Figure 57 • One-Step E2E BC



3.13.6.1 Ingress

Each time the PCS/PMA detects the start of a frame, it sends a pulse to the time stamp block, which saves the value of the Local_Time received from the Local Time counter. In the time stamp block, the programmed value in the local_correction register is subtracted from the saved time stamp. The local_correction register is programmed with the fixed latency from the measurement point to the place that the start of frame is detected in the PCS/PMA logic. The time stamp 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 small 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 an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain that the system is operating on. If so, it signals to the ingress time stamp block in the PHY 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 block in the PHY.

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

If the time stamp block gets the Write action, it delivers the value of the calculated time stamp for the frame to the rewriter block and the rewriter block adds this time stamp (ns part of it) to the four reserved bytes in the frame and recalculates FCS.

The rewriter block takes data out of the FIFO block continuously and feeds it to the system side PCS/PMA block using a counter to keep track of the byte positions of the frame. When the rewriter block receives a signal from the time stamp block to rewrite a specific position in the frame (that information comes from the analyzer block), it overwrites the position with the data from the time stamp block 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. If the frame is an IPv4 frame the rewriter ensures that the IP

checksum is 0. If the frame is IPv6 the rewriter keeps track of the modifications done to the frame and modifies a couple of bytes placed at the end of the PTP frame (for this specific purpose) so that the IP checksum stays correct.

The following full calculations are performed:

- Sync frames: $\text{Reserved_bytes} = (\text{Raw_Timestamp_ns} - \text{Local_correction})$ Correction field = Original Correction field + Asymmetry
- Delay_req frames: $\text{Reserved_bytes} = (\text{Raw_Timestamp_ns} - \text{Local_correction})$

3.13.6.2 Egress

When a frame is received from the system side 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 an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain that the system is operating on.

If the egress analyzer of the PHY detects that the frame is an IEEE 1588 Sync frame belonging to the PTP domain(s) of the system, it signals to the egress time stamp block which action to perform (Write). It also delivers the write offset and data size (location of the Tx time stamp inside the frame, 10 bytes wide) to the rewriter.

If the egress analyzer detects that the frame is an IEEE 1588 Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp block which action to perform (Write, Save). It also delivers the write offset and data size (location of the Tx time stamp inside the frame, 10 bytes wide) to the rewriter. It also outputs up to 16 bytes of frame identifier to the Tx time stamp FIFO, to be saved along with the Tx time stamp. The frame identifier bytes are selected information from the frame, configured in the analyzer.

If the time stamp block gets the (Write, Save) action, it delivers the calculated time stamp and signals to the time stamp FIFO block that it must save the time stamp along with the frame identifier data it received from the analyzer block.

The Tx time stamp FIFO block contains a buffer memory. It simply stores the Tx time stamp values that it receives from the time stamp block together with the frame identifier data it receives from the analyzer block and has a CPU interface that allows the IEEE 1588 engine to read out the time stamp sets (Frame identifier + New Tx time stamp).

The following full calculations are performed:

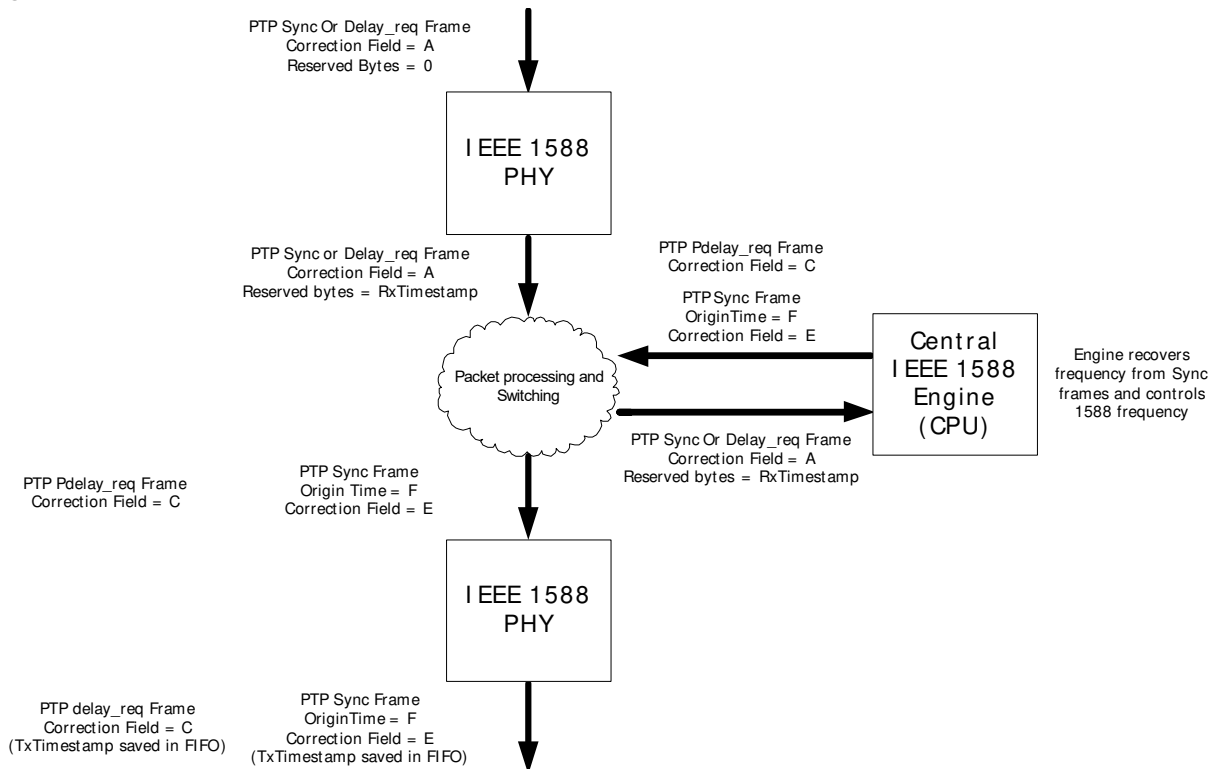
- Sync frames: $\text{OriginTimestamp} = (\text{Raw_Timestamp} + \text{Local_correction})$
- Delay_req frames: $\text{OriginTimestamp} = (\text{Raw_Timestamp} + \text{Local_correction})$ Correction field = Original Correction field + Asymmetry

3.13.7 Supporting Two- Step Boundary/Ordinary Clock

Two-step clocks are used in systems that cannot update the correction field on-the-fly and this requires more CPU processing than one-step.

Each time a Tx time stamp is sent in a frame, the IEEE 1588 engine reads the actual Tx transmission time from the time stamp FIFO and issues a follow-up message containing this time stamp. Even though the VSC8584-11 device supports one-step operation, thereby eliminating the need to run in two-step mode, support for this mode is provided for networks that include two-step-only implementations.

Figure 58 • Two-Step E2E BC



3.13.7.1 Ingress

If the ingress analyzer in the PHY detects that the frame is an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 time stamp block gets the Write action, it delivers the calculated time stamp to the rewriter block and the rewriter block adds this time stamp (ns part of it) to the four reserved bytes in the frame and recalculates FCS.

Note: When secure timing delivery is required, when using IPsec authentication for instance, the four reserved bytes must be reverted back to 0 before performing integrity check.

The following full calculations are performed:

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

3.13.7.2 Egress

If the egress analyzer detects that the frame is an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp block which action to perform (Write, Save). The analyzer outputs up to 15 bytes of frame identifier to the Tx time stamp FIFO to be saved along with the Tx time stamp. The frame identifier must include, at a minimum, the sequenceId field so the CPU can match the time stamp with the follow-up frame.

If the time stamp block gets the Write, Save action, it delivers the calculated time stamp to the time stamp FIFO and signals to the time stamp FIFO block that it must save the time stamp along with the frame identified 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)
Correction field = Original Correction field – Asymmetry

3.13.8 Supporting One-Step End-to-End Transparent Clock

End-to-end transparent clocks add the residence time (the time it takes to traverse the system from the input to the output port(s)) 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 IEEE 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 IEEE 1588 or other methods (SyncE), the error is only caused by sampling inaccuracies.


The VSC8584-11 device supports a number of different transparent clock modes that can be divided into two main modes, as follows:

- **Mode A** subtracts the ingress time stamp at ingress and adds the egress time stamp at egress. This mode can run in a number of sub-modes, depending on the format of the time stamp that is subtracted or added.
- **Mode B** saves the ingress time stamp in the reserved bytes of the PTP header (just as is done in BC and ordinary clock modes) and performs the residence time calculation at the egress PHY where the calculated residence time is added to the correction field of the PTP frame.

Mode B is recommended because it has a number of advantages, including the option to support TC and BC operation in the same system and on the same traffic and the ease of implementing syntonized TC operation.

When an E2E TC recovers frequency using IEEE 1588 and is using Mode A, it must either have a PHY with IEEE 1588 time stamping Mode A support or another way of adding the local time to the correction field placed in front of the IEEE 1588 engine. The IEEE 1588 engine is then able to receive Sync frames and adjust the local frequency to match the IEEE 1588 time.

If using Mode B, the IEEE 1588 engine can recover the frequency directly from the Sync frames because it can extract the ingress time stamp directly from the frames. The frequency adjustment can be done by adjusting the time counter in each PHY or by adjusting the global Timetick clock.

Figure 59 • One-Step E2E TC Mode A

When the system works in one-step E2E TC mode Sync and Delay_req frames must be forwarded through the system and the residence time = (Egress time stamp - Ingress time stamp) must be added to the correction field in the frame before it leaves the system.

The following sections describe the operation in Modes A and B.

3.13.8.1 Ingress, Mode A

If the analyzer detects that the frame is an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 time stamp block gets the Subtract action, it subtracts the time stamp converted to ns from the original correction field of the frame and outputs the value to the rewriter block.

As a result, the frame is sent towards the system with 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)

3.13.8.2 Egress, Mode A

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

If the analyzer detects that the frame is a IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 time stamp block and the rewriter block to ignore the frame (let it pass unmodified and flush the saved time stamp in the time stamp block).

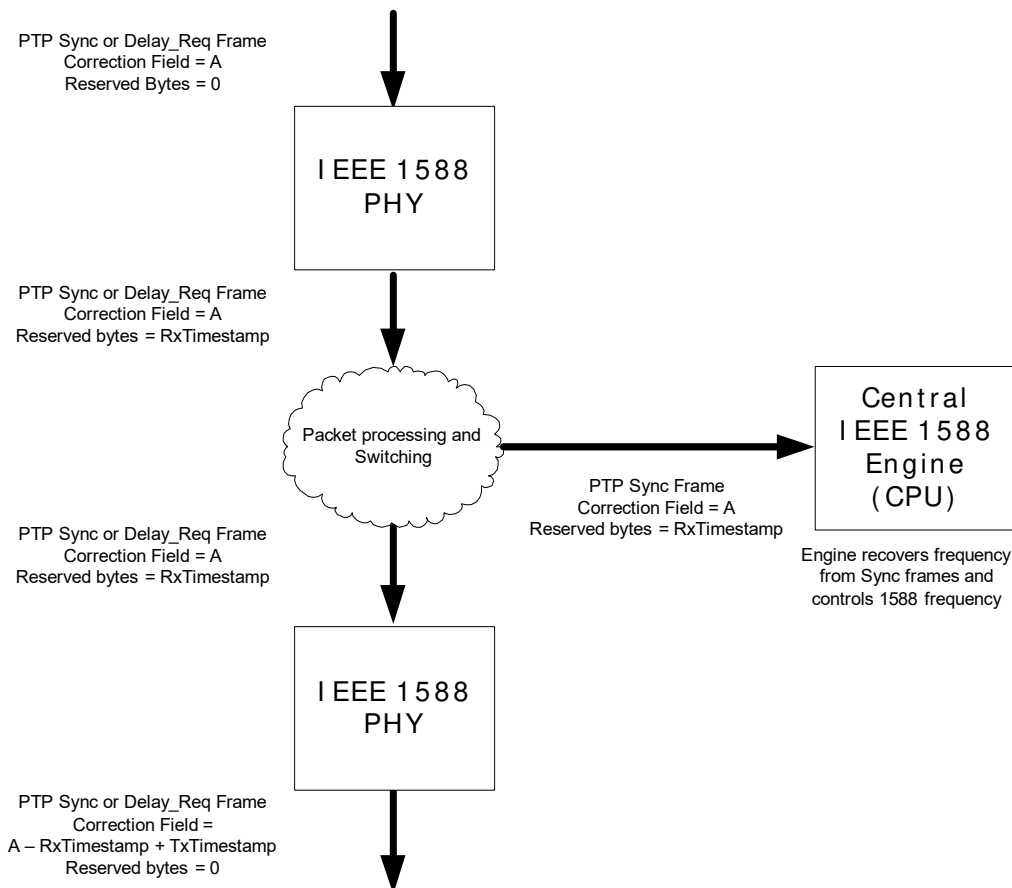
If the time stamp block gets the Add action, it adds the current value of the active_timestamp to the value of the correction field received from the analyzer and outputs the value to the rewriter block.

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 received from the time stamp block 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 60 • One-Step E2E TC Mode B



3.13.8.3 Ingress, Mode B

In ingress mode B, all calculations are performed at the egress port.

On the ingress side, when the analyzer detects Sync or Delay_req frames it adds the Rx time stamp to the four reserved bytes in the PTP frame.

The following full calculations are performed:

- Sync frames: $\text{Reserved_bytes} = \text{Raw_Timestamp_ns} - \text{Local_correction}$ Correction field = Original Correction field + Asymmetry
- Delay_req frames: $\text{Reserved_bytes} = \text{Raw_Timestamp_ns} - \text{Local_correction}$

3.13.8.4 Egress, Mode B

All 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 time stamp – Rx time stamp

The value of the Rx time stamp 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 time stamp – Rx time stamp

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

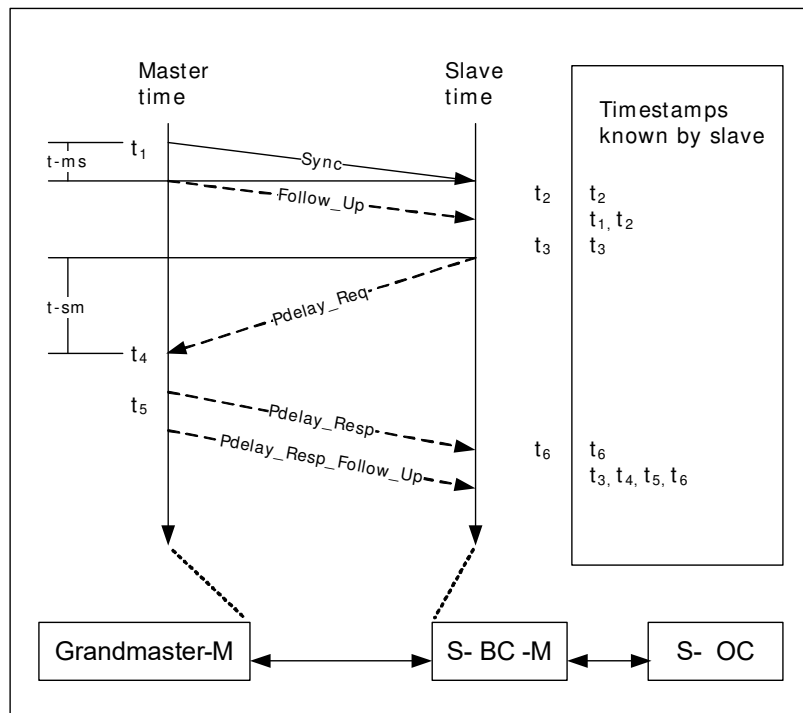
3.13.9 Supporting One-Step Peer-to-Peer Transparent Clock

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.

3.13.9.1 Peer Link Delay Measurement

In P2P TC, the P2P TC device actively sends and receives Pdelay_req and Pdelay_resp messages, and calculates the path delays to each neighbor node in the PTP network. The following illustration shows the delay measurements.

Figure 61 • Delay Measurements



To calculate the path delays on a link, the IEEE 1588 engine (located somewhere in the system) generates Pdelay_req messages on all ports. When transmitted, the actual Tx time stamp t_3 is saved for the CPU to read.

When a P2P TC, BC, or OC receives a Pdelay_req frame, it saves the Rx time stamp (t_4) and generates a Pdelay_resp frame, which adds $t_5 - t_4$ to the correction field copied from the received Pdelay_req frame, where t_5 is the time that the Pdelay_resp leaves the port (t_5).

When a P2P TC receives the Pdelay_resp frame, it saves the Rx time stamp (t_6) and then calculates the path delay as $(t_6 - t_3 - \text{the correction field of the frame})/2$. The time stamp corrections are combined into a single formula as follows:

$$\text{Path delay} = (t_6 - (t_3 + (t_5 - t_4)))/2 = (t_6 - t_3 - t_5 + t_4)/2 = ((t_4 - t_3) + (t_6 - t_5))/2$$

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 CPU processing overhead of reading the t_3 time stamp. To modify the algorithm, the IEEE 1588 engine should send the Pdelay_req message with a software generated t_3 value in the origin time stamp, the sub-second value of the t_3 time stamp in the reserved bytes of the PTP header, and a correction field of 0. The software generated t_3 time stamp should just be within a second before the actual t_3 time. The egress PHY should then be configured to perform E2E TC egress operation, meaning calculate the "residence time" from the inserted t_3 time stamp to the actual t_3 time and insert this value in the correction field of the frame. When the IEEE 1588 engine receives the corresponding Pdelay_resp frame back it can use the software generated t_3 value as the correction field of the Pdelay_resp frame will contain a value that compensates for the actual t_3 transmission time.

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

The following sections describe both the standard and modified methods for taking P2P measurements. As with E2E TC operations, the VSC8584-11 device also supports the different TC modes: mode A (with different time stamp formats) and mode B. Mode B is also the preferred method to implement P2P TC.

3.13.9.2 Ingress, Mode A

If the analyzer detects that the frame is an IEEE 1588 Sync frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 an IEEE 1588 Pdelay_req or Pdelay_resp frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 that is used to save the ns part of the Rx time stamp, 4 bytes wide) to the rewriter.

If the time stamp block gets the subtract_p2p action, it subtracts the value in the ingress time stamp from the correction_field data, adds the configured path delay value, and delivers the result to the rewriter block.

If the time stamp block gets the Write action, it outputs the value of the ingress time stamp register to the rewrite block and the rewriter block writes the sub-second value to the reserved bytes in the PTP header.

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
Correction Field = Original Correction field + Asymmetry

3.13.9.3 Egress, Mode A

If the analyzer detects that the frame is a IEEE 1588 Sync frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 an IEEE 1588 Pdelay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp block which action to perform (Sub_add), along with the original correction field of the frame (will have the value of 0) and the time stamp extracted from the reserved bytes. 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 user prefers to use the normal t3 handling where the t3 time stamp is saved in a time stamp FIFO, the following configuration should be used: If the analyzer detects that the frame is an IEEE 1588 Pdelay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 (sequenceld from the PTP header) to the time stamp FIFO, to save along with the Tx time stamp.

If the analyzer detects that the frame is an IEEE 1588 Pdelay_resp frame belonging to the PTP domain(s) of the system, it signals to the time stamp block which action to perform (Sub_add), along with the original correction field of the frame (will have the value of the CF received from the Pdelay_req frame) and the time stamp extracted from the reserved bytes. 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 time stamp block and the rewriter block to ignore the frame (let it pass unmodified and flush the saved time stamp in the time stamp 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) – Reserved_bytes – Asymmetry
- Pdelay_resp frames: Correction field = Original Correction field + (Raw_Timestamp_ns + Local_correction) – Reserved_bytes

3.13.9.4 Ingress, Mode B

If the analyzer detects that the frame is an IEEE 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the time stamp 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 an IEEE 1588 Pdelay_req frame belonging to the PTP domain(s) of system, it signals to the time stamp 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 time stamp, 4 bytes wide) to the rewriter.

If the analyzer detects that the frame is an IEEE 1588 Pdelay_resp frame belonging to the PTP domain(s) of system, it signals to the time stamp 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 time stamp, 4 bytes wide) to the rewriter.

If the time stamp 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 time stamp 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

3.13.9.5 Egress, Mode B

If the analyzer detects that the frame is an IEEE 1588 Sync frame belonging to the PTP domain(s) of system, it signals to the time stamp 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 an IEEE 1588 Pdelay_req frame belonging to the PTP domain(s) of system, it signals to the time stamp 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 time stamp FIFO, to save along with the Tx time stamp.

If the analyzer detects that the frame is an IEEE 1588 Pdelay_resp frame belonging to the PTP domain(s) of system, it signals to the time stamp block which action to perform (Add - this requires that the IEEE 1588 engine has subtracted the Rx time stamp 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 time stamp 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 time stamp 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 time stamp FIFO block contains an (implementation specific) amount of buffer memory. It simply stores the Tx time stamp values that it receives from the time stamp block together with the frame identifier data it receives from the Analyzer block and has a CPU interface that allows the IEEE 1588 engine to read out the time stamp sets (Frame identifier + New Tx time stamp).

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)


Figure 62 • One-Step P2P TC Mode B



3.13.10 Supporting Two-Step Transparent Clock

In two-step transparent clocks, the Rx and Tx time stamps are saved for the IEEE 1588 engine to read and the follow-up message is redirected to the IEEE 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 correction 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 63 • Two-Step E2E TC

3.13.10.1 Ingress

If the analyzer detects that the frame is an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 allows the frames to be copied to the IEEE 1588 engine, so that it can extract the Rx time stamp and that it knows that it needs to read the Tx time stamps to be ready for the follow up message. It is also possible to save the Rx time stamp value along with the Tx time stamp in the Tx time stamp FIFO.

If the time stamp block gets the Write action, it outputs the current time stamp to the rewriter and the rewriter writes the ns part of the time stamp into the reserved bytes and recalculates FCS.

The following full calculations are performed:

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

3.13.10.2 Egress

If the analyzer detects that the frame is an IEEE 1588 Sync or Delay_req frame belonging to the PTP domain(s) of the system, it signals to the time stamp 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 time stamp FIFO to be saved along with the Tx time stamp. The frame identifier must include, at minimum, the sequenceId field so the CPU can match the time stamp 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 time stamp value is cleared.

If the time stamp block gets the Write, Save action, it outputs the current time stamp value to the rewriter (and time stamp FIFO) and sets the save_timestamp bit. The time stamp 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 time stamp, so that the CPU now can read the set of Tx and Rx time stamp from the Tx time stamp FIFO.

The following full calculations are performed:

- Sync frames: FIFO = Raw_Timestamp_ns + Local_correction (reserved_bytes containing the Rx time stamp saved together with Tx time stamp)
- Delay_req frames: FIFO = Raw_Timestamp_ns + Local_correction – Asymmetry (reserved_bytes containing the Rx time stamp saved together with Tx time stamp)

3.13.11 Calculating OAM Delay Measurements

Frame delay measurements can be made as one-way and two-way delay measurements. Microsemi 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 time stamping 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.

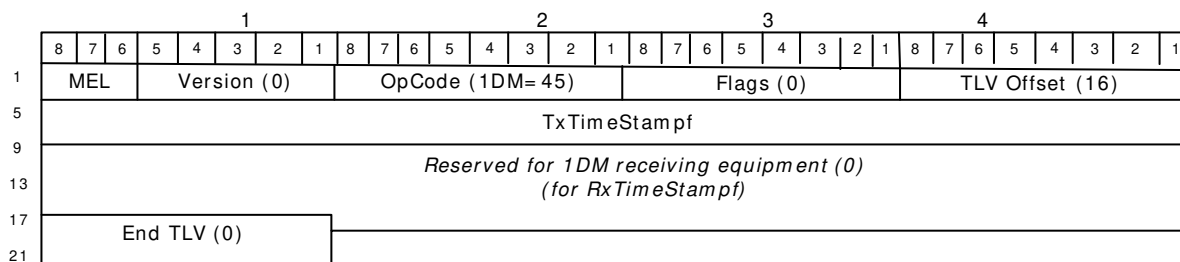
3.13.12 Supporting Y.1731 One-Way Delay Measurements

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.

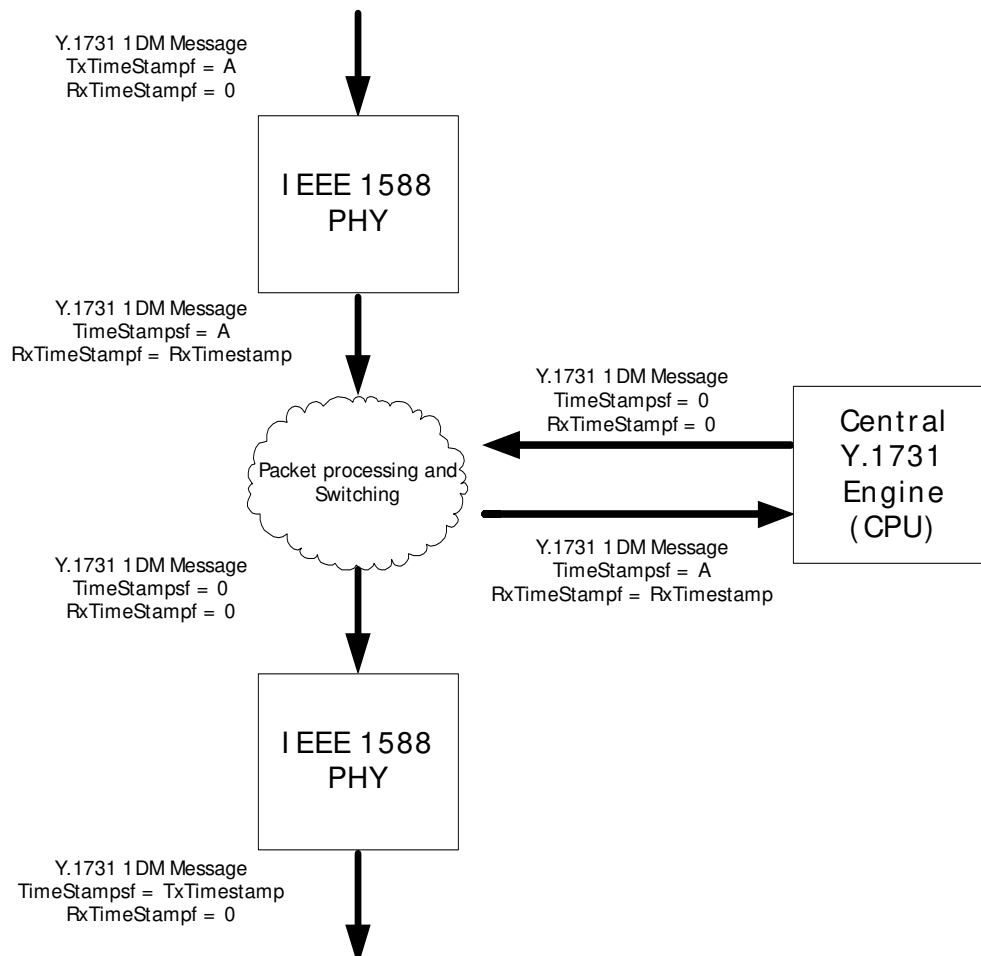
The receiver notes the time of reception of the 1DM frame and calculates the delay.

Figure 64 • 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 time stamp.
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 time stamp 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 65 • Y.1731 One-Way Delay



3.13.12.1 Ingress

If the analyzer detects that the frame is a Y.1731 1DM PDU frame belonging to the MEP, it signals to the time stamp 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 time stamp block gets the Write action, it delivers the time stamp to the rewriter block and the rewriter block adds this time stamp to the reserved bytes in the frame and recalculates FCS.

The following calculation is performed for 1DM frames:

$$RxTimeStampf = (Raw_Timestamp - Local_correction)$$

3.13.12.2 Egress

If the analyzer detects that the frame is a Y.1731 1DM PDU frame belonging to the MEP, it signals to the time stamp 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 time stamp block gets the Write action, it delivers the time stamp to the rewriter block and the rewriter block adds this time stamp to the reserved bytes in the frame and recalculates FCS.

The following calculation is performed for 1DM frames:

$$TxTimeStampf = (Raw_Timestamp + Local_correction)$$

3.13.13 Supporting Y.1731 Two-Way Delay Measurements

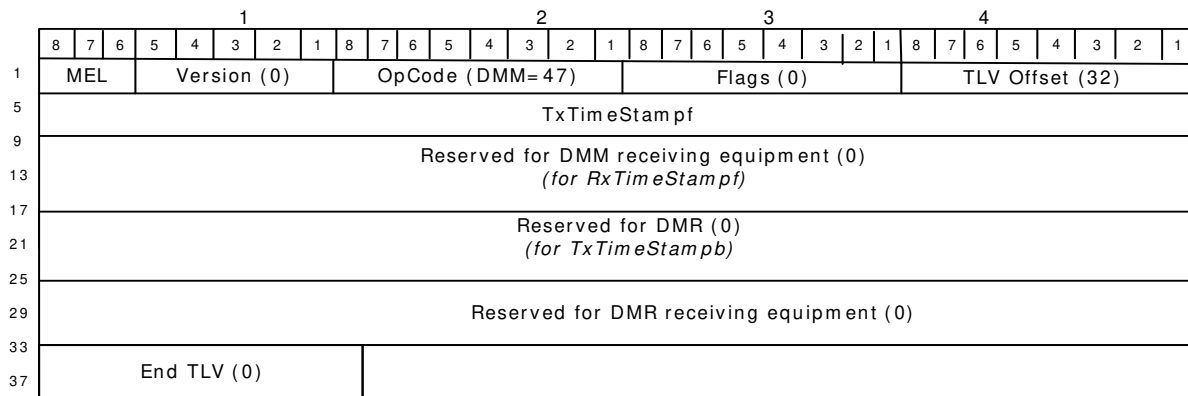
When performing two-way delay measurements, the initiating MEP transmits DMM frames containing a TxTimeStampf value. The receiving MEP replies with a DMR frame that is the same as the DMM frame, but with destination and source MAC address swapped and with a different OAMPDU opcode.

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 time stamps 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 time stamps allows for the exclusion of the processing time in the peer MEP, but it does not require that the two MEPs are synchronized.

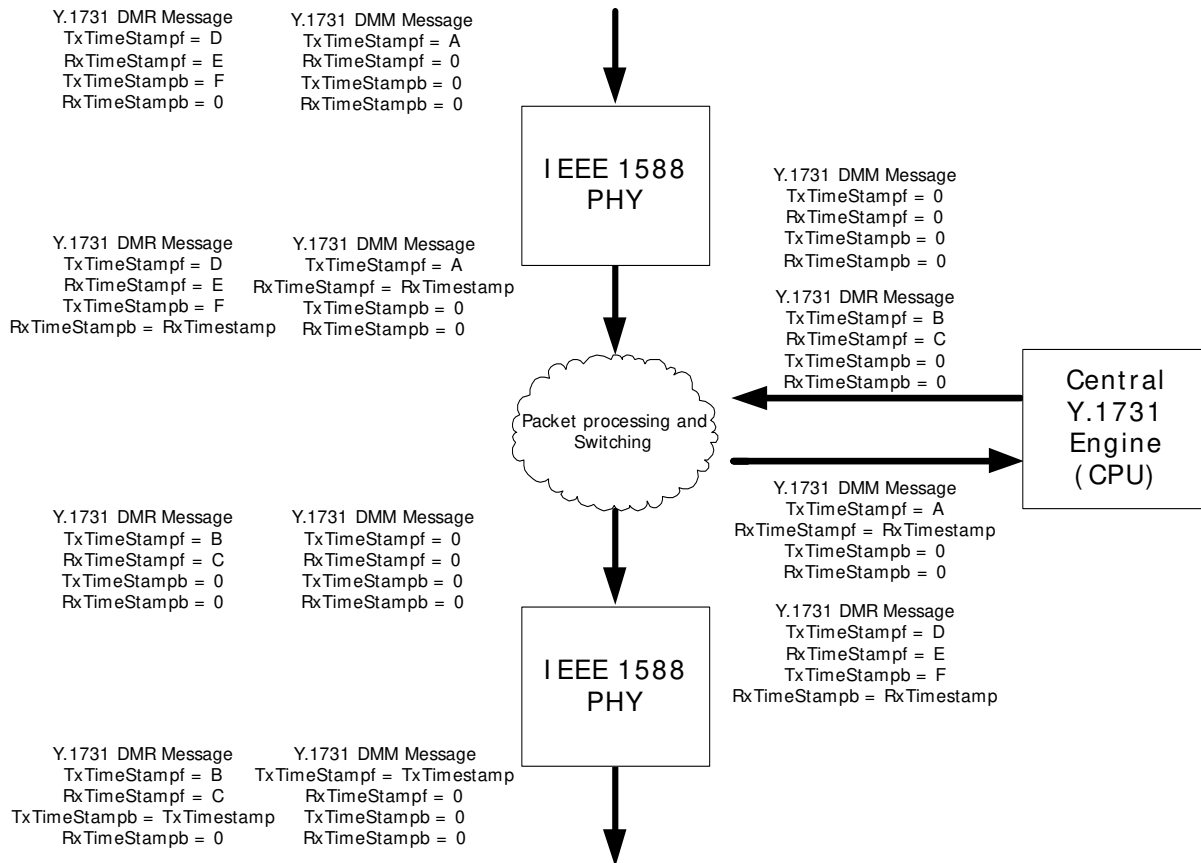
Figure 66 • Y.1731 DMM PDU Format



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 time stamp.
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 67 • Y.1731 Two-Way Delay



3.13.13.1 Ingress

If the analyzer detects that the frame is a Y.1731 DMM PDU frame belonging to the MEP, it signals to the time stamp 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 time stamp 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 time stamp block gets the Write action, it delivers the time stamp to the rewriter block and the rewriter block adds this time stamp to the reserved bytes in the frame and recalculates FCS.

The following calculations are performed:

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

3.13.13.2 Egress

If the analyzer detects that the frame is a Y.1731 DMM PDU frame belonging to the MEP, it signals to the time stamp 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 time stamp 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 time stamp block gets the Write action, it delivers the time stamp to the rewriter block and the rewriter block adds the time stamp to the reserved bytes in the frame and recalculates FCS as follows:

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

3.13.13.3 Supporting MPLS-TP One-Way and Two-Way Delay Measurements

MPLS-TP one- and two-way delay measurement are defined in RFC6374 (G.8113.2) and G.8113.1 (draft-bhh). These mechanisms are similar to the ones described for Y.1731 Ethernet OAM delay measurement except for the encapsulations. The following illustrations show the PDU formats.

Figure 68 • RFC6374 DMM/DMR OAM PDU Format

ETH (1)		14/18/22B
MPLS labels (2)		4/8/12/16B
ACH		4B
DMM/DMR OAM PDUs	OAM PDU Header	8B
	Time stamp 1	8B
	Time stamp 1	8B
	Time stamp 1	8B
	Time stamp 1	8B
	padding	(variable size)
FCS		4B

(1) 0, 1, or 2 VLAN tags
(2) Up to 4 MPLS labels

Figure 69 • Draft-bhh DMM/DMR/1DM OAM PDU Formats

DMM/DMR		1DM	
ETH (1)		14/18/22B	14/18/22B
MPLS labels (2)		4/8/12/16B	4/8/12/16B
ACH		4B	4B
DMM/DMR OAM PDUs	OAM PDU Header	8B	8B
	Time stamp 1	8B	8B
	Time stamp 1	8B	8B
	Time stamp 1	8B	8B
	Time stamp 1	8B	8B
	End TLV indicator	1B	1B
FCS		4B	4B

(1) 0, 1, or 2 VLAN tags
(2) Up to 4 MPLS labels

3.13.14 Device Synchronization for IEEE 1588 Support

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 IEEE 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 time stamp 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, Microsemi 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 IEEE 1588 time progress can be calculated. This PPM offset can be used to calculate how many local-time-clocks it 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 \pm 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.

3.13.15 Time Stamp Update Block

The IEEE 1588 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 time stamp updates/time stamps where necessary. For IEEE 1588 timing distribution the VSC8584-11 device supports ordinary clocks, boundary clocks, end-to-end transparent clocks, and peer-to-peer transparent clocks in a chassis based IEEE 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 and G.8113.1/G.8113.2. The TSU block implements part of the functionality required for full IEEE 1588 compliance.

The IEEE 588 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

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-TP (G.8113.1 (~draft-bhh-mpls-tp-oam-y1731) and G.8113.2 (RFC6374))


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 70 • PTP Packet Encapsulations



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

Figure 71 • OAM Packet Encapsulations



There is one TSU per channel in the VSC8584-11 device. The TSU detects and updates up to three 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 time stamp. A local time counter is maintained to provide the time stamps. Implementation of some of the IEEE 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 time stamp FIFO is only required in the egress direction because the CPU needs to know the actual time stamps of some of the transmitted PTP

frames. The CPU reads the time stamps and any associated frame information out of the time stamp FIFO. The FIFO saves the generated time stamps 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 72 • 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 time stamp 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 two channels share configuration registers and have identical setups.

The time stamp block waits for an SOF to be detected, captures a time stamp from the local time counter, and builds the new time stamp that is to be written into the PTP/OAM frame. Captured time stamps can be read by the CPU.

The rewriter block handles the actual writing of the new time stamp into the PTP/OAM frame. It is also able to clear parts of the frame such as the UDP checksum, if required, or it can update the frame to

ensure that the UDP checksum is correct (for IPv6 PTP frames). The block also calculates the new FCS to be written to the PTP frame after updating the fields with the new time stamp.

The VSC8584-11 device has variable latency in the PCS block. These variations are predicted and used to compensate/maximize the accuracy of the IEEE 1588 time stamp logic.

If the time stamp update function is not used, the block can be bypassed. 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.

Each direction of the IEEE 1588 can be bypassed individually by programming the INTERFACE_CTL.SPLIT_BYPASS bit. Bypass is then controlled by INTRERFACE_CTL.INGR_BYPASS and INTERFACE_CTL.EGR_BYPASS.

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


3.13.16 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 73 • 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 time stamp block indicating that this frame does not need modification and that it should discard its time stamp.

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 88 and [MPLS Comparator](#), page 92.

Table 27 • Flows Per Engine Type

Comparator	Number of Flows	
	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

Encapsulation matches can be set independently in each direction by setting the ANALYZER_MODE.SPLIT_ENCAP_FLOW register. However strict and non-strict flow cannot be set independently for group A and group B of analyzer engine C.

Choice of strict flow or non-strict has to be made on each direction rather than on an engine by engine basis. Valid values for INGR_ENCAP_FLOW_ENA and EGR_ENCAP_FLOW_ENA are 3'b000 or 3'b111.

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 28 • 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 29 • 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/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).

3.13.16.1 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 ethertype 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.

3.13.16.1.1 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 74 • Type II Ethernet Basic Frame Format




3.13.16.1.2 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 75 • 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 76 • 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 77 • Ethernet Frame with VLAN Tags and SNAP



3.13.16.1.3 Provider Backbone Bridging (PBB) Support

The provider backbone bridging 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 78 • PBB Ethernet Frame Format (No B-Tag)





Figure 79 • PBB Ethernet Frame Format (1 B-Tag)



3.13.16.1.4 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.

3.13.16.1.5 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 30 • 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 Ethernertype 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 Ethernertype field in the packet is compared against this value.

3.13.16.1.6 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 31 • 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 31 • Ethernet Comparator (Flow) (continued)

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 IEEE 802.3 for more details. 0 = Match the full 48-bit address 1 = Match any unicast address 2 = Match any multicast address

Table 31 • Ethernet Comparator (Flow) (continued)

Parameter	Width	Description
Eth_VLAN_Tag1_Match	12 bit	Match field for the first VLAN tag (if configured to be present).
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.

3.13.16.2 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 80 • 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 81 • 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 82 • MPLS Labels and Control Word

CW=1	Destination Address (DA) 8 7 6 5 4 3 2 1 0	Source Address (SA) 8 7 6 5 4 3 2 1 0	Etype 3 2 1 0	Label (S=0) 3 2 1 0	Label (S=0) 3 2 1 0	Label (S=0) 3 2 1 0	Label (S=1) 3 2 1 0	Control 3 2 1 0	Payload
------	---	--	------------------	------------------------	------------------------	------------------------	------------------------	--------------------	---------

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 its own complete set of match registers.

Table 32 • MPLS Comparator: Next Word

Parameter	Width	Description
MPLS_Nxt_Comparator	3 bit	Pointer to the next comparator

Table 33 • 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 is performed from the end of the stack										
MPLS_STACK_DEPTH	4 bit	Each bit represents a possible stack depth, as shown in the following list. <table border="1" data-bbox="760 1465 1330 1612"> <thead> <tr> <th>MPLS_STACK_DEPTH Bit</th> <th>Allowed Stack Depth</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>1</td> </tr> <tr> <td>1</td> <td>2</td> </tr> <tr> <td>2</td> <td>3</td> </tr> <tr> <td>3</td> <td>4</td> </tr> </tbody> </table>	MPLS_STACK_DEPTH Bit	Allowed Stack Depth	0	1	1	2	2	3	3	4
MPLS_STACK_DEPTH Bit	Allowed Stack Depth											
0	1											
1	2											
2	3											
3	4											

Table 34 • 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

Table 34 • MPLS Range_Upper/Lower Label Map (continued)

Parameter	MPLS_REF_PNT = 0, top-of-stack referenced	MPLS_REF_PNT=1, end-of-stack referenced
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 35 • 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 36 • 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


3.13.16.3 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.

3.13.16.4 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 83 • 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.

Note: Checksum overwrite with 0 occurs on ingress only. PTP applications that generate 1588 frames with this format are responsible for creating IPv4/UDP frames with a zeroed checksum upon generation from the application.

3.13.16.5 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 84 • 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 at the end of the frame will be updated by the PHY so that the checksum field does not have to be modified.

Note: IPv6 extension headers are not supported.

3.13.16.6 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 85 • ACH Header Format





Figure 86 • ACH Header with Protocol ID Field




3.13.16.7 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 IEEE 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 time stamps in the TS FIFO are used to create follow-up messages. On ingress, the time stamp 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.

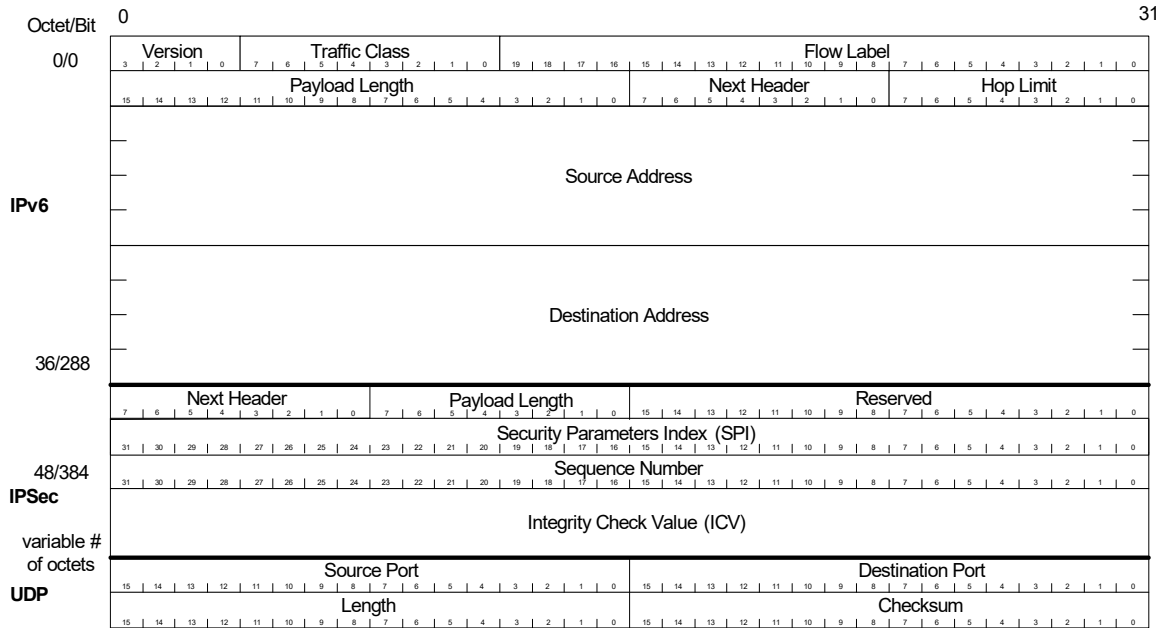
The following illustration shows the format of the IPSec frame. It normally appears between the IP header (IPv4 or IPv6) and the UDP header or at the start of the payload.

Figure 87 • 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 88 • IPv6 with UDP and IPSec



3.13.16.8 Comparator Field Summary

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

Table 37 • 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		

3.13.16.8.1 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 38 • 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 than IPv4 or 6, this register points to the beginning of the field that will be used for flow matching.

Table 38 • IP/ACH Next-Protocol Comparison (continued)

Parameter	Width	Description
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.
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 39 • 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

Table 39 • IP/ACH Comparator Flow Verification Registers (continued)

Parameter	Width	Description
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.


3.13.16.9 PTP/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 OAM frames.

3.13.16.10 PTP Frame Header

The following illustration shows the header of a PTP frame.

Figure 89 • 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.

3.13.16.11 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 90 • OAM 1DM Frame Header Format



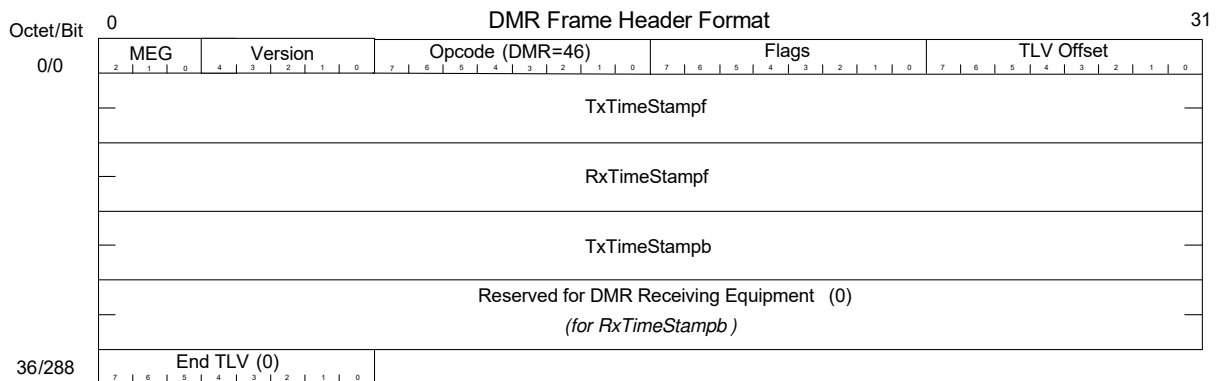
The following illustration shows a DMM frame header.

Figure 91 • OAM DMM Frame Header Format



The following illustration shows a DMR frame header.

Figure 92 • 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.

3.13.16.12Y.1731 OAM PDU

1DM, DMM, and DMR are the three supported G.8113.1 PDUs and DMM/DMR are the two supported RFC6374 PDUs. The following illustrations show the PDU formats.

Figure 93 • RFC6374 DMM/DMR OAM PDU Format

	ETH (1)	14/18/22B
	MPLS labels (2)	4/8/12/16B
	ACH	4B
DMM/DMR OAM PDUs	OAM PDU Header	8B
	Time stamp 1	8B
	Time stamp 1	8B
	Time stamp 1	8B
	Time stamp 1	8B
	padding	(variable size)
	FCS	4B

(1) 0, 1, or 2 VLAN tags
 (2) Up to 4 MPLS labels

Figure 94 • G8113.1/draft-bhh DMM/DMR/1DM OAM PDU Format

DMM/DMR		1DM	
	ETH (1)	14/18/22B	
	MPLS labels (2)	4/8/12/16B	
	ACH	4B	
DMM/DMR OAM PDUs	OAM PDU Header	8B	1DM OAM PDUs
	Time stamp 1	8B	
	Time stamp 1	8B	
	Time stamp 1	8B	
	Time stamp 1	8B	
	End TLV indicator	1B	
	FCS	4B	

(1) 0, 1, or 2 VLAN tags
(2) Up to 4 MPLS labels

As with PTP, there is no protocol validation for MPLS OAM; 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.

3.13.16.13 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.

Table 40 • 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
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

Table 40 • PTP Comparison (continued)

Parameter	Width	Description		
		Command Value	Mnemonic	Action
PTP_Action_Command	3 bit	0	NOP	Do nothing
		1	SUB	New correction field = Current correction field – Captured local time
		2	SUB_P2P	New correction field = Current correction field – Local latency + path_delay
		3	ADD	New correction field = Current correction field + Captured local time
		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 time stamp FIFO (only valid for egress ports).		
PTP_Correction_Field_Offset	5 bit	Points to the location of the correction field. 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 value in the delay asymmetry register is added to the correction field of the frame.		
PTP_Subtract_Delay_Asymmetry_Enable	1 bit	When enabled, the value in the delay asymmetry register is subtracted from the correction field of the frame.		
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.		

Table 40 • PTP Comparison (continued)

Parameter	Width	Description
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 time stamp
PTP_Rewrite_Offset	8 bits	Points to where in the frame relative to the SFD that the time stamp should be updated
PTP_New_CF_Loc	8 bits	Location where the updated correction field value is written 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 41 • 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 42 • 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.

3.13.16.14 Future Protocol Compatibility

Except for MPLS, 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 be 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.

3.13.16.15 Reconfiguration

There are three 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.
3. Disable a comparator.
Each comparator within the active encapsulation engine can be disabled. If an Ethernet header according to the configuration Type I or Type II with SNAP/LLC is not found then subsequent flows will not be matched. The ETH1 comparator can also be disabled so that all frames flowing through the IEEE 1588 block are time stamped.

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.

3.13.16.16 Frame Signature Builder

Along with time stamp and CRC updates, the analyzer outputs a frame signature that can be stored in the time stamp FIFO to help match frames with other info in the FIFO. This information is used by the CPU so that it can match time stamps in the time stamp 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 43 • 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
28-35	Selected address byte (select-28)

Table 44 • Frame Signature Address Source

Parameter	Width	Description
FSB_Map_Reg_0-15	6 bits	For each byte of the frame signature, use Table 43 , page 105 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.

Table 44 • Frame Signature Address Source (continued)

Parameter	Width	Description
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.

A frame signature can be extracted from frames matching in all the three engines. The frame signature address selection is limited to Ethernet Block1 because only a limited number of encapsulations are supported in the third engine, Engine C.

Engine C has two parts: part A and part B. Part A supports ETH1, ETH2, MPLS protocols while part B supports only ETH1 protocol. Selection of Ethernet block 1 or 2 is dependent on whether part A flow matches or part B flow matches.

If a frame matches part A flow configuration, then the frame signature as configured in ETH1_NXT_PROTOCOL_A and ETH2_NXT_PROTOCOL_A using FSB_ADR_SEL will be considered in computing the frame signature.

If a frame matches part B flow configuration, then the frame signature as configured in ETH1_NXT_PROTOCOL_A and FSB_ADR_SEL will be considered in computing the frame signature. In this configuration if FSB_ADR_SEL is set to 1, to select ETH2 then all zeros are padded as frame signature because ETH2 is not supported by part B.

3.13.16.17 Configuration Sharing

The analyzer configuration 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.

A 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.).

3.13.16.18 OAM-Optimized Engine

The OAM optimized engine, Engine C, supports a fewer set of encapsulations such as ETH1, ETH2, MPLS, and ACH. Engine C is enhanced with an ACH comparator to support the MPLS-TP OAM protocol. The MPLS-TP OAM protocol for Engine C is configured in the following registers.

- EGR2_ACH_PROT_MATCH_UPPER/LOWER_A
- EGR2_ACH_PROT_MASK_UPPER/LOWER_A
- EGR2_ACH_PROT_OFFSET_A

The ACH comparator will start the comparison operation right after the MPLS comparator.

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, 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.

3.13.17 Time Stamp Processor

The primary function of the time stamp processor block is to generate a new `Timestamp_field` or new `Correction_field` (Transparent clocks) for the rewriter block. The time stamp block generates an output that is either a snapshot of the corrected Local Time (struct time stamp) or a signed (two's complement) 64 bit `Correction_field`.

In the ingress direction, the time stamp block calculates a new time stamp 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 IEEE 1588 standard).

In the egress direction, the time stamp block calculates a new time stamp for the rewriter in time for the PCS block to transmit the new time stamp field in the frame. In this case the time stamp 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 time stamp block is software programmed based upon IEEE 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 (for two-step operation). However, the TSP only updates PTP event messages. Also, the IEEE 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 time stamp block generates an `Active_timestamp` from the captured/time stamped 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 time stamp 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:

$$\text{Active_timestamp} = \text{Raw_timestamp} - \text{Local_latency} - \text{Variable_latency}$$

In the egress direction:

$$\text{Active_timestamp} = \text{Raw_timestamp} + \text{Local_latency} + \text{Variable_latency}$$

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

$$\text{Active_timestamp_ns} = \text{Active_timestamp converted to nanoseconds}$$

$$\text{Active_timestamp_p2p_ns} = \text{active_timestamp_ns} + \text{path delay}$$

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

3.13.18 Time Stamp FIFO

The time stamp FIFO stores time stamps along with frame signature information. This information can be read out by a CPU or pushed out on a dedicated Serial Time Stamp Output Interface and used in 2-step processing mode to create follow-up messages. The time stamp FIFO is only present in the egress data path.

The time stamp FIFO takes a frame signature from the analyzer and the updated correction field, and the full data set for that time stamp 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 the EGR_TSFIFO_CSR.EGR_TS_SIGNAT_BYTES register. Only the indicated number of signature bytes is saved with each time stamp. The saved values are packed so that reducing the number of signature bytes allows more time stamps to be saved.

The packing of the time stamp data is done by logic before the write occurs to the FIFO. When no compression is used, each time stamp may contain 208 bits of information consisting of 128 bits of frame signature and 80 bits of time stamp data. Therefore a full sized time stamp is 26 bytes long. Compressing the frame signature can reduce this to as little as 10 bytes (or 4 bytes if EGR_TSFIFO_CSR.EGR_TS_4BYTES = 1) if no signature information is saved (EGR_TSFIFO_CSR.EGR_TS_SIGNAT_BYTES = 0). The value to store is built up in an internal register. When the register contains 26 valid bytes, that data is written to the time stamp FIFO. Data in the FIFO is packed end-to-end. It is up to the reader of the data to unpack the data.

The time stamps in the FIFO are visible and accessible for the CPU as a set of 32-bit registers. Multiple register reads are required to read a full time stamp if all bits are used. Bit 31 in register EGR_TSFIFO_0 contains the current FIFO empty flag, which can be used by the CPU to determine if the current time stamps are available for reading. If the bit is set, the FIFO is empty and no time stamps are available. The value that was read can be discarded because it does not contain any valid time stamp data. If the bit is 0 (deasserted), the value contains 16 valid data bits of a time stamp. The remaining bits should be read from the other registers in the other locations and properly unpacked to recreate the time stamp. Care should be taken to read the time stamps one at a time as each read of the last (7th) address will trigger a pop of the FIFO.

Time stamps are packed into seven registers named EGR_TSFIFO_0 to EGR_TSFIFO_6. If the time stamp FIFO registers are read to the point that the FIFO goes empty and there are remaining valid bytes in the internal packing register, then the packing register is written to the FIFO. In this case the registers may not be fully packed with time stamps. Flag bits are used to indicate where the valid data ends within the set of seven registers. The flag bits are in register EGR_TSFIFO_0.EGR_TS_FLAGS (together with the empty flag) and are encoded as follows:

- 000 = Only a partial time stamp is valid in the seven register set
- 001 = One time stamp begins in the current seven register set
- 010 = Two time stamps begin in the current seven register set.
- 011 = Three time stamps begin in the current seven register set (4-byte mode)
- 100 = Four time stamps begin in the current seven register set (4-byte mode)
- 101 = Five time stamps begin in the current seven register set (4-byte mode)
- 110 = Six time stamps begin in the current seven register set (4-byte mode)
- 111 = The current seven register set is fully packed with valid time stamp data

The FIFO empty bit is visible in the EGR_TSFIFO_0.EGR_TS_EMPTY register so the CPU can poll this bit to know when time stamps are available. There is also a maskable interrupt which will assert whenever the time stamp FIFO level reaches the threshold given in EGR_TSFIFO_CSR.EGR_TS_THRESH register. The FIFO level is also visible in the EGR_TSFIFO_CSR.EGR_TS_LEVEL register. If the time stamp FIFO overflows, writes to the FIFO are inhibited. The data in the FIFO is still available for reading but new time stamps are dropped.

Note: Time stamp FIFO exists only in the Egress direction. There is no time stamp FIFO in the Ingress direction

3.13.19 Serial Time Stamp Output Interface

For each IEEE 1588 Processor 0 and 1, time stamp information stored in the Egress direction can be read through either the register interface or through the Serial Time Stamp interface. These two ways to read registers are mutually exclusive. While enabling/disabling the serial interface is done on a Processor level, only one serial interface exists. This means the serial interface can be enabled for Processor 0, while the time stamp FIFO can be read through registers for Processor 1. If the serial interface is enabled for both Processor 0 and 1, then the serial interface will arbitrate between two Egress time stamp FIFOs in Processor 0 and 1 and push the data out.

The time stamp FIFO serial interface block writes, or pushes, time stamp/frame signature pairs that have been enqueued and packed into time stamp 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 serial interface (SPI) is enabled, the time stamp/frame signature pairs are dequeued from time stamp FIFO(s) and unpacked. Unpacked time stamp/frame signature pairs are then serialized and sent one at a time to the external interface. Unpacking shifts the time stamp/frame signature into alignment considering the configured size of the time stamps and frame signatures (a single SI write may require multiple reads from a time stamp FIFO). The time stamp FIFO serial interface is an alternative to the MDIO register interface described in the time stamp FIFO section. When the serial time stamp interface is enabled in register TS_FIFO_SI_CFG.TS_FIFO_SI_ENA, data read from the time stamp FIFO registers described in [Time Stamp FIFO](#), page 108 are invalid.

Time stamp/Frame signature pairs from two egress time stamp FIFOs are serialized one at a time and transmitted to the interface pins. The TS_FIFO_SI arbitrates in a round-robin fashion between the ports that have non-empty time stamp FIFOs. The port associated with each transmitted time stamp/frame signature pair is indicated in a serial address that precedes the data phase of the serial transmission. Because the time stamp FIFOs are instantiated in the per port clock domains, a small single entry asynchronous SI FIFO (per port) ensures that the time stamp/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 time stamp/frame signature performing any time stamp FIFO dequeues necessary. The SI FIFO goes empty following the completion of the last data bit of the serial transmission. Enabled ports (TS_FIFO_SI_CFG.TS_FIFO_SI_ENA) participate in the round-robin selection.

Register TS_FIFO_SI_TX_CNT accumulates the number of time stamp/frame signature pairs transmitted from the serial time stamp interface for each channel. Register EGR_TS_FIFO_DROP_CNT accumulates the number of time stamp/frame signature pairs that have been dropped per channel due to a time stamp FIFO overflow.


The SPI compatible interface asserts a chip select (SPI_CS) 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 SPI_CS is negated. Each write command corresponds to a single time stamp/frame signature pair. The length of the data phase depends upon the sum of the configured lengths of the time stamp and signature, respectively. The address phase is fixed at five bits. The SPI_CLK is toggled to transfer each SPI_DO bit (as well as the command and address bits). The Time Stamp 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 [Time Stamp FIFO](#), page 108 and [Figure 95](#), page 110.

The frequency of the generated output 1588_SPI_CLK can be flexibly programmed from 10 MHz up to 62.5 MHz using TS_FIFO_SI_CFG 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 clock, the CSR register would be set such that both SI_CLK_LO_CYCS and SI_CLK_HI_CYCS equal 3. Also, the number of CSR clocks after SPI_CS asserts before the first 1588_SPI_CLK is programmable (SI_EN_ON_CYCS), as is the number of clocks before SI_EN negates after the last 1588_SPI_CLK (SI_EN_OFF_CYCS). The number of clocks during which SI_EN is negated between writes is also programmable (SI_EN_DES_CYCS). The 1588_SPI_CLK may also be configured to be inverted (SI_CLK_POL).

Without considering de-selection between writes, if the PTP 16-byte SequenceID (frame signature) is used as frame identifier each 10 byte time stamp write take $2 + 55 + 10 \times 8 + 16 \times 8 = 265$ clocks (at

40 MHz) ~6625 ns. This corresponds to a time stamp bandwidth of > 0.15 M time stamp/second/port. The following illustration shows the serial time stamp/frame signature output.

Figure 95 • Serial Time Stamp/Frame Signature Output



3.13.20 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/EGR_RW_MODFRM_CNT, depending upon the direction. This counter accumulates the number of PTP frames to which a write was performed and includes errored frames.

3.13.20.1 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 time stamp 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 INGR_RW_FCS_ERR_CNT or EGR_RW_FCS_ERR_CNT register, depending upon the direction.

3.13.20.2 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 when enabled using the following register bits:

- INGR_IP1_UDP_CHKSUM_UPDATE_ENA
- INGR_IP2_UDP_CHKSUM_UPDATE_ENA
- EGR_IP1_UDP_CHKSUM_UPDATE_ENA
- EGR_IP2_UDP_CHKSUM_UPDATE_ENA


Based upon the analyzer command and the rewriter configuration, the rewriter writes the time stamp in one of the following ways:

- Using PTP_REWRITE_BYTES to choose four bytes write to PTP_REWRITE_OFFSET. This method is similar to other PTP frame modifications and the time stamp is typically written to the reserved field in the PTP header.
- Using PTP_REWRITE_BYTES and RW_REDUCE_PREAMBLE to select the mode of operation when writing Rx time stamps into the frame. In these modes, it cannot do both a time stamp write/append and a PTP operation in the same frame. If PTP_REWRITE_BYTES = 0xE and RW_REDUCE_PREAMBLE = 1, it does it by overwriting the existing FCS with the time stamp in the lowest four bytes of the calculated time stamp and generating a new FCS and appending it.

Because the rewriter cannot modify the IFG or change the size of the frame, if the original FCS is overwritten with time stamp data a new FCS needs to be appended and the frame shortened by reducing the preamble. The preamble length includes the /S/ character and all preamble characters up to but not including the SFD. In this mode, it is assumed that all incoming preambles are of sufficient (5 to 7-byte) length to delete four bytes and the preamble of every frame (not only PTP frames) will be reduced by four bytes by deleting four bytes of the preamble. Then, the new FCS is written at the end of the matched frame. For unmatched frames, or if the PTP_REWRITE_BYTES is anything but 0xE, the IFG is increased by adding four IDLE (/I/) characters after the /T/ which ends the packet.

To time stamp a frame in one of the modes, the actual length of the preamble is then checked and if the preamble is too short to allow a deletion of four bytes (if the preamble is not five bytes or more) then no operations are performed on the preamble, the FCS is not overwritten, and no time stamp is appended. For all such frames, a counter is maintained and every time an unsuccessful operation is encountered, the counter is incremented. This counter is read through register: INGR_RW_PREAMBLE_ERR_CNT/EGR_RW_PREAMBLE_ERR_CNT. The following illustration shows the deleted preamble bytes.

Figure 96 • Preamble Reduction in Rewriter



If PTP_REWRITE_BYTES = 0xF and RW_REDUCE_PREAMBLE = 0, the rewriter replaces the FCS of the frame with the four lowest bytes of the calculated time stamp and does not write the FCS to the frame. In this mode, all the frames have corrupted FCSs and the MAC needs to be configured to handle this case. In the case of a CRC error in the original frame, the rewriter writes all ones (0xFFFFFFFF) to the FCS instead of the time stamp. This indicates an invalid CRC to the MAC because this is reserved to indicate an invalid time stamp. In the rare case that the actual time stamp has the value 0xFFFFFFFF and the CRC is valid, the rewriter increments the time stamp to 0x0 and writes that value instead. This causes an error of 1 ns but is required to reserve the time stamp value of 0xFFFFFFFF for frames with an invalid CRC.

A flag bit may also be set in the PTP message header to indicate that the TSU has modified the frame (when set) or to clear the bit (on egress). The analyzer sends the byte offset of the flag byte to the rewriter in PTP_MOD_FRAME_BYTE_OFFSET and indicates whether the bit should be modified or not using PTP_MOD_FRAME_STATUS_UPDATE. The bit offset within the byte is programmed in the configuration register RW_FLAG_BIT. When the PTP frame is being modified, the selected bit is set to the value in the RW_FLAG_VAL. This only occurs when the frame is being modified by the rewriter; when the PTP frame matches and the command is not NOP.

3.13.21 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, 125 MHz, or some other frequency. The clock may be a line clock or the dedicated 1588_DIFF_INPUT_CLKP/N pins. The clock source is selected in register LTC_CTRL.LTC_CLK_SEL.

To support other frequencies, a flexible counter system is used that can convert almost any frequency in the 125–250 MHz range into a usable source clock. Supported frequencies of local time counter are 125 MHz, 156.25 MHz, 200 MHz, and 250 MHz. The frequency is programmed in terms of the clock period. Set the LTC_SEQUENCE.LTC_SEQUENCE_A register to the clock period to the nearest whole number of nanoseconds to be added to the local time counter on each clock cycle. Set LTC_SEQ.LTC_SEQ_E to the amount of error between the actual clock period and the

LTC_SEQUENCE.LTC_SEQUENCE_A setting in femtoseconds. Register LTC_SEQ.LTC_SEQ_ADD_SUB indicates the direction of the error. An internal counter keeps track of the accumulated error. When the accumulated error exceeds 1 nanosecond, an extra nanosecond is either added or subtracted from the local time counter. Use the following as an example to program a 5.9 ns period:

```
LTC_SEQUENCE.LTC_SEQUENCE_A = 6 (6 ns)
LTC_SEQ.LTC_SEQ_E = 100000 (0.1 ns)
LTC_SEQ.LTC_SEQ_ADD_SUB = 0 (subtract an extra nanosecond, i.e 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 internal counter in nanoseconds is given in register

LTC_AUTO_ADJUST.LTC_AUTO_ADJUST_NS. At rollover, the next increment of the local time counter is increased by one additional or one less nanosecond as determined by the LTC_AUTO_ADJUST.LTC_AUTO_ADD_SUB_1NS register. When LTC_AUTO_ADJUST.LTC_AUTO_ADD_SUB_1NS is set to 0x1, an additional nanosecond is added to the local time counter. When it is set to 0x2, one less nanosecond is added to the local timer counter. No PPM adjustments are made when the register is set to 0x0 or 0x3.

PPM adjustments to the local time counter can be made on an as-needed basis by writing to the one-shot LTC_CTRL.LTC_ADD_SUB_1NS_REQ register. One nanosecond is added or subtracted from the local time counter each time LTC_CTRL.LTC_ADD_SUB_1NS_REQ is asserted. The LTC_CTRL.LTC_ADD_SUB_1NS register setting controls whether the local time counter adjustment is an addition or a subtraction.


The current time is loaded into the local time counter with the following procedure.

1. Configure the 1588_LOAD_SAVE pin.
2. Write the time to be loaded into the local time counter in registers LTC_LOAD_SEC_H, LTC_LOAD_SEC_L and LTC_LOAD_NS.
3. Program LTC_CTRL.LTC_LOAD_ENA to a 1.
4. Drive the 1588_LOAD_SAVE pin from low to high.

The time in registers LTC_LOAD_SEC_H, LTC_LOAD_SEC_L and LTC_LOAD_NS is loaded into the local time counter when the rising edge of the 1588_LOAD_SAVE strobe is detected. The LOAD_SAVE strobe is synchronized to the local time counter clock domain.

When the 1588_DIFF_INPUT_CLK_P/N pins are the clock source for the local time counter, and the LOAD_SAVE strobe is synchronous to 1588_DIFF_INPUT_CLK_P/N, the LTC_LOAD* registers are loaded into the local time counter, as shown in the following illustration.

Figure 97 • Local Time Counter Load/Save Timing



When the LOAD_SAVE strobe is not synchronous to the 1588_DIFF_INPUT_CLK_P/N pins or an internal clock drives the local time counter, there is some uncertainty as to when the local time counter is loaded, when higher accuracy circuit is turned off. This reduces the accuracy of the time stamping function by the period of the local time counter clock. When higher accuracy circuit is ON, any difference between the 1588_DIFF_INPUT_CLK_P and the rising edge of 1588_LOAD_SAVE is compensated within an error of 1 ns. This applies to both load and save operations.

There is a local time counter in each channel. The counter is initialized in both channels if the LTC_CTRL.LTC_LOAD_ENA register in each channel is asserted when the LOAD_SAVE strobe occurs.

When LTC_CTRL.LTC_SAVE_ENA register is asserted when the 1588_LOAD_SAVE input transitions from low to high, the state of the local time counter is stored in the LTC_SAVED_SEC_H, LTC_SAVED_SEC_L, and LTC_SAVED_NS registers.

The current local time can be stored in registers with the following procedure.


1. Configure the 1588_LOAD_SAVE pin.
2. Program LTC_CTRL.LTC_SAVE_ENA to a 1.
3. Set SER_TOD_INTF.LOAD_SAVE_AUTO_CLR to 1 if the operation is one-time save operation. This will clear LTC_CTRL.LTC_SAVE_ENA after the operation.
4. Drive the 1588_LOAD_SAVE pin from low to high.
5. Read the value from LTC_SAVED_SEC_H, LTC_SAVED_SEC_L, and LTC_SAVED_NS registers.

As with loading the local time counter, there is one clock cycle of uncertainty as to when the time is saved if the LOAD_SAVE strobe is not synchronous to the clock driving the counter.

3.13.22 Serial Time of Day

In addition to loading or saving as described in the preceding sections, it is possible to load or save LTC time in a serial fashion. For serial load, 1588_LOAD_SAVE has to send Time of Day (ToD) information in a specific format. For serial save, when the appropriate register bits is set, then PPS will drive out the ToD information. The following illustration shows the format for serial load and save.

Figure 98 • Standard PPS and 1PPS with TOD Timing Relationship



3.13.22.1 Pulse per Second Segment

In the preceding illustration, segment A is the pulse per second segment. The PPS signal is transmitted with high voltage. The rising edge of the PPS signal is aligned with the rising edge of the standard PPS signal. This segment lasts 1 μs. To obtain high accuracy, the response delay of the rising edge of the PPS signal should be considered.

3.13.22.2 Waiting Segment

In the preceding illustration, segment B is the Waiting segment. Due to the speed of operation, this segment is needed to make it easier for the receiver to obtain the following Time-of-Day information in current PPS cycle. The signal is in low voltage during this segment, which lasts 20 μs.

3.13.22.3 Time-of-Day Segment

In the preceding illustration, segment C is the Time-of-Day segment. The ToD information being carried in this segment indicates the time instant of the rising edge of the PPS signal transmitted in segment A of the current PPS cycle. The time instant is measured using the original network clock. In this segment, the ToD information is continuously transported and is represented in 16 octets. It consists of the following fields:


- Second field: 6 octets. It represents the time instant of the rising edge of the PPS signal in second. The value is defined as in IEEE 1588-2008.
- Date field: 6 octets. It represents the time instant of the rising edge of the PPS signal in year, month, day, hour, minute, and second. Each part is represented by one octet (the format of this field is 0xYYMMDDHHMMSS). In particular, only the lowest 2 decimal digits of the year are represented. The receiver can easily obtain the time instant of the rising edge of the PPS signal in this transparent format without additional circuitry to translate the value of the second field. It also has the significant

benefit of changing the value of this field when leap year or leap second occurs. (The Date field is ignored at the serial ToD input and is not generated at the serial ToD output.)

- Reserved field: 4 octets. Reserved for future use.

The ToD information is represented in units of octet, with each octet being transmitted with the low-order bit first. The following illustration shows an octet is transmitted between a start bit with high voltage and a stop bit with low voltage. The other octets are transmitted in the same manner. As a result, $(1+8+1) \times 1 \mu\text{s} = 10 \mu\text{s}$ are needed to transport one octet. This segment lasts $16 \times 10 \mu\text{s} = 160 \mu\text{s}$ to convey the ToD information.

Figure 99 • ToD Octet Waveform



The entire Time-of-Day segment should be detected. If the second 6 octets representing the Date field are not used by the upper layer, the Date field should still be detected and its value can be ignored.

3.13.22.4 Idle Segment

Segment D is the Idle segment in Figure 98, page 113. It follows segment C with high voltage until the end of the PPS cycle. The duration of the Idle segment is given by the following calculation.

$$1 \text{ s} - 0.5 \mu\text{s} - 20 \mu\text{s} - 160 \mu\text{s} = 999819.5 \mu\text{s}.$$

Use the following steps to enable Serial load.

1. Set SER_TOD_INTF.SER_TOD_INPUT_EN to 1
2. Set LTC_CTRL.LOAD_EN to 1.
3. Start the transmission of 1588_LOAD_SAVE conforming to the format.
4. To check the data transmission, enable serial save or save LTC time to check the registers.
5. To enable serial save, set SER_TOD_INTF.SER_TOD_OUTPUT_EN to 1.

The following table lists the different options to load or save LTC time.

Table 45 • LTC Time Load/Save Options

LTC_CTRL.LOAD_EN	SER_TOD_INTF.SER_TOD_INPUT_EN	LTC_CTR.SAVE_EN	Expected Operation
0	0	1	Parallel Save
0	1	1	Save
0	0	0	No operation
0	1	0	No operation
1	0	0	Parallel Load
1	1	0	Serial Load
1	0	1	Parallel Load and Save

Table 45 • LTC Time Load/Save Options (continued)

LTC_CTRL.LOAD_EN	SER_TOD_INTF.SER_TOD_INPUT_EN	LTC_CTR.SAVE_EN	Expected Operation
1	1	1	Serial Load and Save

When SER_TOD_INTF.SERIAL_ToD_OUTPUT_EN is set, the PPS output is driven with a serial ToD output based on the LTC timer value.

3.13.23 Programmable Offset for LTC Load Register

When a new LTC value is loaded into the system, a fixed offset may need to be added to the loaded value. Program SER_TOD_INTF.LOAD_PULSE_DLY and this value will be added to LTC counter whenever a new load occurs either through software, load_save pin or through serial ToD.

3.13.24 Adjustment of LTC counter

LTC counter value can be adjusted by about a second without reloading a new LTC value. LTC value can be programmed to tune the current value by adding or subtracting a specific value. The offset adjustment can be positive or negative, very similar to 1 ns adjustment being positive or negative. An adjustment every 232 ns can be performed using LTC_OFFSET_ADJ. Additionally, an adjustment every 220 ns can be performed using LTC_AUTO_M_x.

The purpose of this register is to add/subtract a programmable offset register of 30-bit width in ns, to the register block in order to cover the entire nanosecond portion of the 80-bit LTC. This offset control is independent of the LTC load control. The LTC timer is adjusted - added or subtracted as per the bit LTC_OFFSET_ADJ.LTC_ADD_SUB_OFFSET, by the value LTC_OFFSET_ADJ.LTC_OFFSET_VAL, when a load offset command is issued by the software (assertion of LTC_OFFSET_ADJ.LTC_OFFSET_ADJ). The hardware sets the status bit LTC_OFFSET_ADJ_STAT.LTC_OFFSET_DONE after completing the operation. However, in case the hardware cannot complete the operation because of the LTC value itself getting updated synchronously due to parallel or serial LTC load at the same time, it sets the bit LTC_OFFSET_ADJ_STAT.LTC_OFFSET_LOAD_ERR. The software on seeing either of these status bits set (LTC_OFFSET_ADJ_STAT.LTC_OFFSET_DONE or LTC_OFFSET_ADJ_STAT.LTC_OFFSET_LOAD_ERR), de-asserts the control bit LTC_OFFSET_ADJ.LTC_OFFSET_ADJ and might potentially retry the operation.

The maximum value in nanoseconds for the offset LTC_OFFSET_ADJ.LTC_OFFSET_VAL can be up to $10^9 - 1$. Thus, for addition operation, the maximum carry to the seconds counter is 2 because of the clock period addition to this maximum value present in the offset and LTC nanoseconds counter. For subtraction operation, if the resultant subtraction is negative or underflows the LTC timer would be set to wrong value. Therefore, such operations should never be allowed.

LTC_OFFSET_ADJ register (with LTC_OFFSET_VAL[29:0] and LTC_ADD_SUB_OFFSET) should be updated before asserting LTC_OFFSET_ADJ bit in LTC_OFFSET_ADJ register.

LTC_OFFSET_ADJ_STAT.LTC_OFFSET_DONE and LTC_OFFSET_ADJ_STAT.LTC_OFFSET_LOAD_ERR bits are set by the hardware and cleared by the software by writing a zero.

Should a conflict occur between LTC update due to parallel/serial load and LTC update due to offset adjustment, the load LTC takes precedence and the error condition is noted so that the polling software does not hang on the offset status bit assertion.

LTC counter could be adjusted for any known drift that occurs on every second. This feature will add or subtract one nanosecond every time LTC crosses over LTC_AUTO_ADJ_M_NS.

Example 1: If LTC_AUTO_ADJ_M_NS is 100 ns and LTC is started from reset with a value of 0 ns, then LTC counter will be added/subtracted 1 ns every time counter rolls over 100 ns.

Example 2: If LTC_AUTO_ADJ_M_NS is 100 ns and LTC is started from reset with a value of 0 ns, then LTC counter will be added/subtracted 1 ns every time counter rolls over. When counter is at 10 ns and LTC counter is loaded with 2 sec, 80 ns. Now 1 ns is adjusted when counter increments from 10 ns and rolls over 100 ns. It does not add/subtract when LTC timer rolls over 100 ns.

Example 3: LTC_AUTO_ADJ_M_NS value is loaded with 400 ns and after some time LTC_AUTO_ADJ_M_NS value is loaded with 500 ns. The AUTO_ADJ_M_COUNTER value when the new value is loaded is 333 ns. Then the next adjustment happens after 177 ns after load because the AUTO_ADJ_M_COUNTER continues to count until it reaches the newly loaded value 500 ns.

Example 4: LTC_AUTO_ADJ_M_NS value is loaded with 400 ns and after some time LTC_AUTO_ADJ_M_NS value is loaded with 100 ns. The AUTO_ADJ_M_COUNTER value when the new value is loaded is 333 ns. Then adjustment happens immediately because $333 > 100$ and the AUTO_ADJ_M_COUNTER is reset to zero after the adjustment

If LTC counter is loaded with a new value, set LTC_AUTO_ADJ_M_UPDATE bit to 1 and reload the LTC_AUTO_ADJ_M_NS value.

3.13.25 Pulse per Second Output

The local time counter generates a one pulse-per-second (1PPS) output signal with a programmable pulse width routed to GPIO pins. The pulse width of the 1PPS signal is determined by the LTC_1PPS_WIDTH_ADJ register.

When the LTC counter exceeds the value in PPS_GEN_CNT (both are in nanoseconds), the PPS signal is asserted. In default operation where PPS_GEN_CNT = 0 the LTC timer generates a PPS signal every time LTC crosses the 1 sec boundary. By writing a large value such as $10^9 - 60$ ns, the 1PPS pulse reaches its destination 60 ns away simultaneous with the LTC second wrap thus providing time-of-day synchronism between two systems.

The 1PPS output has an alternate mode of operation that increases the frequency of the pulses. This mode may be used for applications such as locking an external DPLL to the IEEE 1588 frequency. In the alternate mode the 1PPS signal is driven directly from a single bit of the nanosecond field counter of the local time counter. The pulse width can not be controlled in this alternate operation mode. The alternate mode is enabled with register LTC_CTRL.LTC_ALT_MODE_PPS_BIT.

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 1 MHz to 10 MHz) usable for external applications.

Table 46 • 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

In addition to the preceding frequencies, a specific frequency can be chosen by enabling the synthesizer on the PPS pin using the following steps.

1. Set LTC_FREQ_SYNTH.LTC_FREQ_SYNTH_EN to 1.
2. A toggle signal with the frequency specified will be pushed out onto PPS. The number of nanoseconds the signal stays high can be specified by LTC_FREQ_SYNTH.FREQ_HI_DUTY_CYC_NS. The number of nanoseconds the signal stays low can be specified by LTC_FREQ_SYNTH.FREQ_LO_DUTY_CYC_NS.
3. The above nanoseconds should be exactly divisible by clock frequency, otherwise the signal may have a jitter as high as the high duration/clock period or low duration/clock period.
4. To disable the this feature and revert back to PPS functionality, reset LTC_FREQ_SYNTH.LTC_FREQ_SYNTH_EN to 0

For example, to output a 10 MHz signal, set the `FREQ_HI_DUTY_CYC_NS` to 50 ns and `FREQ_LO_DUTY_CYC_NS` to 50 ns. On a 250 MHz LTC clock, this will make high time and low time of signal shift between 48 ns and 52 ns.

3.13.26 Accuracy and Resolution

The IEEE 1588 processor achieves time stamp resolution in any mode of operation of 1 ns utilizing special high-resolution circuitry. The accuracy of a device using high-resolution circuitry is improved more than 100% over the first generation IEEE 1588 engine. High accuracy for these devices will be supported regardless of the local time counter clock frequency supplied to the reference clock input. The timestamp accuracy is a system-level property and may depend upon oscillator selection, port type, and speed, system configuration, and calibration decisions. Supported frequencies of the local time counter are 125 MHz, 156.25 MHz, 200 MHz, and 250 MHz.

There are a total of five high resolution blocks per port to improve resolution for the following events:

- One pulse-per-second (1PPS) output signal
- 1588_PPS_RI input signal
- Start-of-frame in the egress direction
- Start-of-frame in the ingress direction
- 1588_LOAD_SAVE input (strobe) signal direction

Each of these blocks can individually be enabled using `ACC_CFG_STATUS`. Contact Microsemi with any questions regarding PTP accuracy calculations.

3.13.27 Loopbacks

Loopback options provide a means to measure the delay at different points to evaluate delays between on chip wire delays and external delays down to a nanosecond.

3.13.27.1 Loopback from PPS to PPS_RI pin

In this loopback, an external device will connect the PPS coming out of the IEEE 1588 to PPS_RI of the IEEE 1588 device. The external device could even process the PPS signal and then loopback at a far-end.

3.13.27.2 Loopback from LOAD_SAVE to PPS

When `LOAD_SAVE_PPS_LPBK_EN` is set, input `load_save` pin is connected to output PPS coming out of the IEEE 1588. In this mode, input `load_save` pin is taken as close to the pin as possible without going through any synchronization logic on the `load_save` pin.

3.13.27.3 Loopback of LOAD_SAVE pin

When `LOAD_SAVE_LPBK_EN` is set, one clock cycle before the PPS is asserted, an output enable for `load_save` pin is generated and PPS signal is pushed out on the `load_save` pin acting as an output pin. After two cycles, output enable is brought down and `load_save` will behave as an input pin.

3.13.27.4 Loopback from PPS to LOAD_SAVE pin

When `PPS_LOAD_SAVE_LPBK_EN` bit enabled, output pps signal is taken as close to the I/O as possible and looped back onto `load_save` input pin. This is to account for any delays from PPS generation block to the PPS output pin.

3.13.28 IEEE 1588 Register Access using SMI (MDC/MDIO)

The SMI mechanism is an IEEE defined register access mechanism (refer to Clause 22 of IEEE 802.3). The registers are arranged as 16 bits per register address with a 5 bit address field as defined by IEEE. However, Microsemi has extended this register address space by creating a register page key in register 31. When writing a particular key to register 31, a different set of 5 bit address space register bank can be accessed through the SMI mechanism. (extended page, GPIO page, etc).

The IEEE 1588 registers are organized on page 4. Setting register 31 to 4 provides a window to CSR registers through registers 16,17, and 18.

The IEEE 1588 IP registers are arranged as 32 bits of data. The access method through SMI is done by breaking up the 32 bits of each IEEE 1588 register into the high 16 bits into register 18 and lower 16 bits

into register 17. Then register 16 is used as a command register. Phy0 and Phy2 automatically read/write to engine A. Phy1 and Phy3 automatically read/write to engine B. For more information, see [Figure 54](#), page 62.

3.13.29 1588_DIFF_INPUT_CLK Configuration

The default configuration of the 1588_DIFF_INPUT_CLK_P/N pins sets the VSC8584-11 device to use an internal clock for the LTC. To configure these pins correctly to use an external clock for LTC, write 0xb71c to register 30E4 and 0x7ae0 to register 29E4. Set these two registers to 0x0 when an internal clock is used for LTC.

3.14 Daisy-Chained SPI Time Stamping

Registers 26E4–29E4 enable daisy-chaining multiple devices to reduce the number of pins required to transmit time stamping information to system ASICs gathering IEEE 1588 time stamps.


The VSC8584-11 device captures frame time stamps on the 1588_SPI_IN signals, arbitrates with the internal IEEE 1588 SPI time stamp outputs for the SPI output, and outputs the frame time stamps. Each device output acts as the SPI master while the input acts as the SPI slave. Up to eight devices can be daisy-chained to operate at up to 62.5 MHz. The following table shows the key throughput characteristics for daisy-chained SPI time stamping.

Table 47 • Daisy-Chain Parameters

SPI Bus Frequency	Maximum Time Stamps/Second per Port	Maximum SPI Bus Utilization
31.25 MHz	256	5.69%
31.25 MHz	16	0.71%

The following illustration shows the SPI time stamping format for clock polarity and phase of 0.

Figure 100 • SPI Time Stamping Format



3.15 SPI I/O Register Access

The VSC8584-11 device provides a bidirectional SPI I/O interface for register access to handle both IEEE 1588 and MACsec communication to the device. The device uses one slave select (SS) per slave for a simple slave design, and to share the SCLK, MOSI, and MISO signals. The SPI I/O port is fully independent of either the SPI time stamp input or output ports.

The following illustrations show the write and read cycle format supported by the VSC8584-11 device, with LSB_FIRST=0, BIG_ENDIAN=1, and PADDING_BYTES=3. No other formats are supported.

Figure 101 • SPI Write Cycles


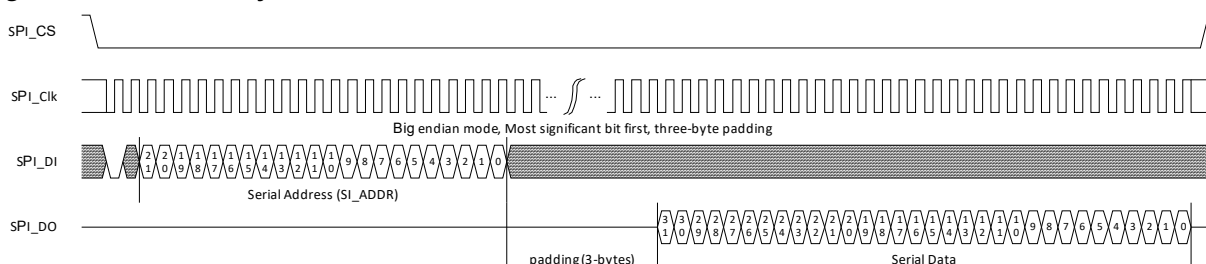


Figure 102 • SPI Read Cycle



A 25 MHz SPI operating rate is used to access the CSR address space.

The 22-bit address (indicated as SI_ADDR), is composed of a 2-bit ring select, an 8-bit Target ID, and a 12-bit register address. This register address represents a word address where a word is 32 bits. The SPI data is 32 bits and is consistent with this mapping.

Table 48 • SI_ADDR Mapping

Bit	Description
21:20	CSR ring select 00: Ring 0 01: Reserved 10: Reserved 11: Reserved
19:14	Target ID[7:2]
13:12	Target ID [1:0] for most targets CSR register address[13:12] for MACsec INGR/EGR Targets
11:0	CSR register address[11:0]

The 2-bit ring select (SI_ADDR[21:20]) selects the CSR ring. A coding of 00 selects ring 0. Coding of 01, 10, and 11 are reserved.

SI_ADDR[19:14] maps to Target ID[7:2] and SI_ADDR[11:0] maps to CSR register address bits 11:0. The Target ID[1:0] and CSR register address bits 13:12 depends on Target ID[5:3].

Target ID[5:3] bits set to 111 access the MACsec INGR or MACsec EGR registers. In this case, Target ID[1:0] is hard-coded to 00 and the CSR address bits 13:12 is supplied by SI_ADDR[13:12]. In all other cases, Target ID[1:0] is supplied by SI_ADDR[13:12] and the CSR address bits 13:12 is hard-coded to 00.

The Chip ID, Extended Chip ID, and Revision Code can be read at Target ID 0x40, address 0xFFFF.

3.16 Media Recovered Clock Outputs

For Synchronous Ethernet applications, the VSC8584-11 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 31.25 MHz or 125 MHz), and squelch conditions.

Note: When EEE is enabled on a link, the use of the recovered clock output is not recommended due to long holdovers occurring during EEE quiet/refresh cycles.

3.16.1 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.

Adjust the squelch level to enable 1000BASE-T master mode recovered clock for SyncE operation. This is accomplished by changing the 23G and 24G register bits 5:4 to 01. This setting also provides clock out for 10BASE-T operation. For 1000BASE-T master mode, the clock is based on the VSC8584-11 REFCLK input, which is a local clock.

3.16.2 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 VSC8584-11 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 controls the squelching of the clock. Both RCVRDCLK1 and RCVRDCLK2 are squelched when the CLK_SQUELCH_IN pin is high.

3.17 Serial Management Interface

The VSC8584-11 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 72](#), page 147 and [Table 148](#), page 189.

The SMI is a synchronous serial interface with input data to the VSC8584-11 on the MDIO pin that is clocked on the rising edge of the MDC pin. It is a multiple-target bus that incorporates open-collector drivers along with an external pull-up to share the MDIO data line between multiple PHY chips. 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-k Ω pull-up resistor is required on the MDIO pin.

3.17.1 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 103 • SMI Read Frame





Figure 104 • 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 VSC8584-11 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 VSC8584-11 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.

3.17.2 SMI Interrupt

The SMI includes an output interrupt signal, MDINT, for signaling the station manager when certain events occur in the VSC8584-11.

The MDINT pin is configured for open-drain (active-low). Tie the pin to a pull-up resistor to VDDIO. The following illustration shows the configuration.

Figure 105 • MDINT Configured as an Open-Drain (Active-Low) Pin



When a PHY generates an interrupt, the MDINT pin is asserted by driving low if the interrupt pin enable bit (MII register 25.15) is set.

3.18 LED Interface

The LED interface supports the following configurations: direct drive, basic serial LED mode, and enhanced serial LED mode. The polarity of the LED outputs is programmable and can be changed through register 17E2, bits 13:10. The default polarity is active low.

Direct drive mode provides four LED signals per port, LED0_[0:3] through LED3_[0:3]. The mode and function of each LED signal can be configured independently. When serial LED mode is enabled, the direct drive pins not used by the serial LED interface remain available.

In basic serial LED mode, all signals that can be displayed on LEDs are sent as LED_Data and LED_CLK for external processing. In enhanced serial LED mode, up to four LED signals per port can be sent as LED_Data, LED_CLK, LED_LD, and LED_Pulse. The following sections provide detailed information about the various LED modes.

Note: LED number is listed using the convention, LED<LED#>_<Port#>.

The following table shows the bit 9 settings for register 14G that are used to control the LED behavior for all the LEDs in VSC8584-11.

Table 49 • LED Drive State

Setting	Active	Not Active
14G.9 = 1 (default)	Ground	Tristate
14G.9 = 0 (alternate setting)	Ground	V _{DD}

3.18.1 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 50 • 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.
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 1000BASE-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).

Table 50 • LED Mode and Function Summary (continued)

Mode	Function Name	LED State and Description
12	Autonegotiation Fault	1: No autonegotiation fault present. 0: Autonegotiation fault occurred.
13	Serial Mode	Serial stream. See Basic Serial LED Mode , page 125. 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.

3.18.2 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 51 • 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 LED pin 0: Disable

3.18.3 LED Behavior

Several LED behaviors can be programmed into the VSC8584-11. Use the settings in register 30 and 19E1 to program LED behavior, which includes the following.

3.18.3.1 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.

3.18.3.2 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.

3.18.3.3 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.

3.18.3.4 LED Pulsing Enable

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

3.18.3.5 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.

3.18.3.6 Fiber LED Disable

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

3.18.3.7 Pulse Programmable Control

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

3.18.3.8 Fast Link Failure

For more information about this feature, see [Fast Link Failure Indication](#), page 126.

3.18.4 Basic Serial LED Mode

Optionally, the VSC8584-11 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 in register 29, bits 3:0 to 0xD. When serial LED mode is enabled, the LED0_0 pin becomes the serial data pin, and the LED1_0 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 following table shows the 48-bit serial output bitstream of each LED signal. The individual signals can be clocked in the following order.

Table 52 • LED Serial Bitstream Order

Output	PHY0	PHY1	PHY2	PHY3
Link/activity	1	13	25	37
Link1000/activity	2	14	26	38
Link100/activity	3	15	27	39
Link10/activity	4	16	28	40
Fiber link/activity	5	17	29	41
Duplex/collision	6	18	30	42
Collision	7	19	31	43
Activity	8	20	32	44

Table 52 • LED Serial Bitstream Order (continued)

Output	PHY0	PHY1	PHY2	PHY3
Fiber activity	9	21	33	45
Tx activity	10	22	34	46
Rx activity	11	23	35	47
Autonegotiation fault	12	24	36	48

3.18.5 Enhanced Serial LED Mode

VSC8584-11 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. In enhanced serial LED mode, the port 0 and port 1 LED output pins serve the following functions:

- LED0_0/LED0_1: LED_DATA
- LED1_0/LED1_1: LED_CLK
- LED2_0/LED2_1: LED_LD
- LED3_0/LED3_1: LED_PULSE

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. LED pin functionality is controlled by setting register 25G, bits 7:1.

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. The LED_PULSE duty cycle is controlled by setting register 25G, bits 15:8.

3.18.6 LED Port Swapping

For additional hardware configurations, the VSC8584-11 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. LED port swapping only applies to the direct-drive LEDs and not to any serial LED output modes.

3.19 Fast Link Failure Indication

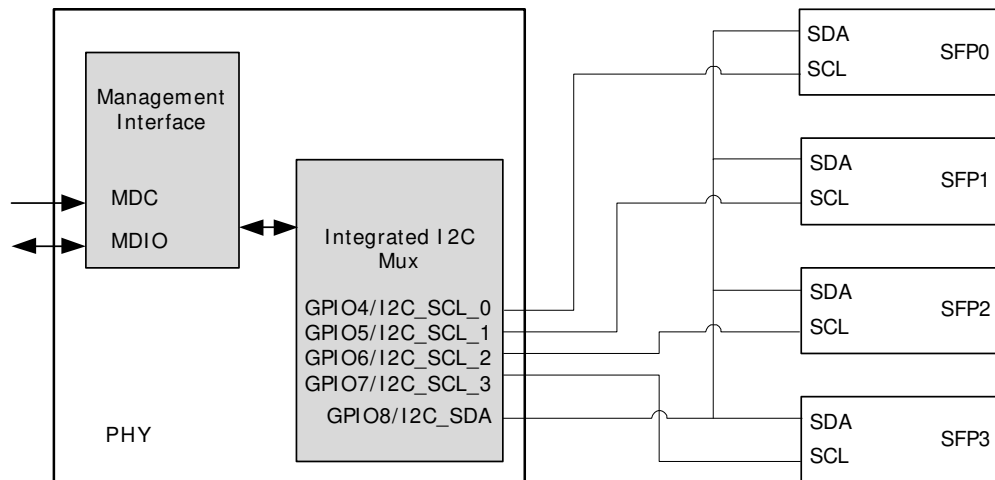
To aid Synchronous Ethernet applications, the VSC8584-11 can indicate the onset of a link failure in less than 1 ms (worst-case <3 ms). By comparison, the IEEE 802.3 standard establishes a delay of up to 750 ms before indicating that a 1000BASE-T link is no longer present. A fast link failure indication is critical to support ports used in a synchronization timing link application. The fast link failure indication works for all copper media speeds, but not for fiber media. 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.

Note: The Fast Link Failure Indication should not be used when EEE is enabled on a link.

3.20 Integrated Two-Wire Serial Multiplexer

The VSC8584-11 includes an integrated quad two-wire serial 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 two-wire serial 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 two-wire serial function. For more information about configuring the pins, see [Two-Wire Serial MUX Control 1](#), page 184. For SFP modules, VSC8584-11 can also provide control for the MODULE_DETECT and TX_DIS module pins using the multipurpose LED and GPIO pins.

Figure 106 • Two-Wire Serial MUX with SFP Control and Status


3.20.1 Read/Write Access Using the Two-Wire Serial MUX

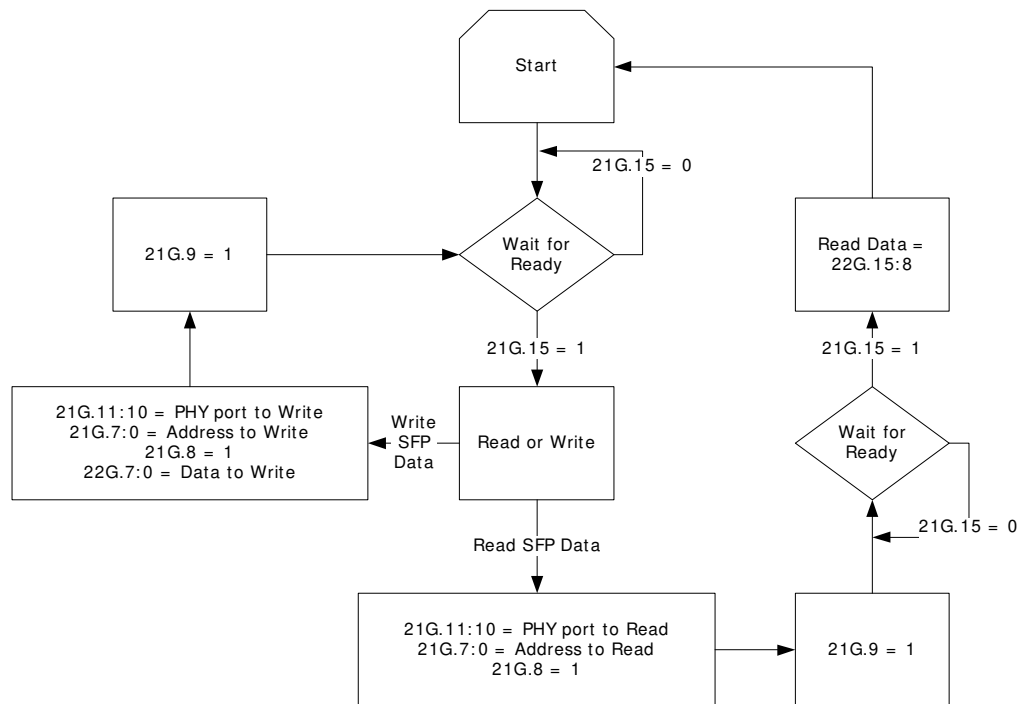
Using the integrated two-wire serial MUX, the VSC8584-11 device can read and write to an SFP or PoE module through the SCL and SDA pins. If the ability is required to write to the slave two-wire serial device, refer to the device's specific datasheet for more information.

Note: The VSC8584-11 device does not automatically increment the two-wire serial address. Each desired address must be intentionally set.

Main control of the integrated two-wire serial MUX is available through register 20G. The two-wire serial MUX pins are enabled or disabled using register 20G bits 3:0. Register 20G bits 15:9 set the two-wire serial 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.

Clock stretching is not supported so the connected devices must be able to operate at the selected serial frequency without wait states. The following illustration shows the read and write register flow.

Figure 107 • Two-Wire Serial MUX Read and Write Register Flow



To read a value from a specific address of the two-wire serial slave device:

1. Read the VSC8584-11 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 two-wire serial 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 two-wire serial slave device:

1. Read the VSC8584-11 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 two-wire serial bus, always wait until register 21G bit 15 changes to 1 before performing another two-wire serial read or write operation.

3.21 GPIO Pins

The VSC8584-11 provides 14 multiplexed general purpose input/output (GPIO) pins. All device GPIO pins and their behavior are controlled using registers. The following table shows an overview of the register controls for GPIO pins. For more information, see [General Purpose Registers](#), page 178.

Table 53 • Register Bits for GPIO Control and Status

GPIO Pin	GPIO_ctrl	GPIO Input	GPIO Output	GPIO Output Enable
GPIO0/SIGDET0	13G.1:0	15G.0	16G.0	17G.0
GPIO1/SIGDET1	13G.3:2	15G.1	16G.1	17G.1
GPIO2/SIGDET2	13G.5:4	15G.2	16G.2	17G.2
GPIO3/SIGDET3	13G.7:6	15G.3	16G.3	17G.3

Table 53 • Register Bits for GPIO Control and Status (continued)

GPIO Pin	GPIO_ctrl	GPIO Input	GPIO Output	GPIO Output Enable
GPIO4/I2C_SCL_0	13G.9:8.	15G.4	16G.4	17G.4
GPIO5/I2C_SCL_1	13G.11:10	15G.5	16G.5	17G.5
GPIO6/I2C_SCL_2	13G.13:12	15G.6	16G.6	17G.6
GPIO7/I2C_SCL_3	13G.15:14	15G.7	16G.7	17G.7
GPIO8/I2C_SDA	14G.1:0	15G.8	16G.8	17G.8
GPIO9/FASTLINK_FAIL	14G.3:2	15G.9	16G.9	17G.9
GPIO10/1588_LOAD_SAVE	14G.5:4	15G.10	16G.10	17G.10
GPIO11/1588_PPS_0	14G.7:6	15G.11	16G.11	17G.11
GPIO12/1588_SPI_CS	14G.15:14	15G.12	16G.12	17G.12
GPIO13/1588_SPI_DO	14G.15:14	15G.13	16G.13	17G.13

3.22 Testing Features

The VSC8584-11 device includes several testing features designed to facilitate performing system-level debugging and in-system production testing. This section describes the available features.

3.22.1 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 VSC8584-11, 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 VSC8584-11 is connected to a live network.

To enable the VSC8584-11 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.

3.22.2 CRC Counters

Cyclical redundancy check (CRC) counters are available in all PHYs in VSC8584-11. They monitor traffic on the copper and fiber media interfaces and on the MAC 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, register 21E3, 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.

3.22.2.1 Copper Interface CRC Counters

Two separate counters are available and reside at the output of the copper interface PCSs before any MACsec or IEEE 1588 packet processing. 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.

3.22.2.2 SerDes Fiber Media Receive CRC Counters

Two separate CRC counters are available and reside at the output of the SerDes media interface PCSs before any MACsec or IEEE 1588 packet processing. To select the SerDes Fiber media receive CRC counters, set register bits 29E3.15:14 to 00. 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.

3.22.2.3 SerDes Fiber Media Transmit Counters

Two fiber media transmit counters are available and reside at the input to the SerDes media interface PCSs after any MACsec or IEEE 1588 packet processing. To select the SerDes Fiber media transmit CRC counters, set register bits 22E3.15:14 to 00. Register bits 21E3.13:0 are the good CRC packet counters and register bits 22E3.7:0 are the CRC error counters.

3.22.2.4 MAC Interface Transmit CRC Counters

Two MAC interface transmit counters are available and reside at the output of the QSGMII/SGMII MAC-interface PCS before any MACsec or IEEE 1588 packet processing. To select these counters, set register bits 22E3.15:14 to 01. Register bits 21E3.13:0 are the good CRC packet counters and register bits 22E3.7:0 are the CRC error counters.

3.22.2.5 MAC Interface Receive CRC Counters

Two MAC interface receive counters are available and reside at the input of the QSGMII/SGMII MAC-interface PCS after any MACsec or IEEE 1588 packet processing. To select these counters, set register bits 29E3.15:14 to 01. Register bits 28E3.13:0 are the good CRC packet counters and register bits 29E3.7:0 are the CRC error counters.


3.22.3 Loopbacks

Loopbacks described herein are for test use only, and are not recommended for use on operational links. Furthermore, the 1588 TSU block should be bypassed whenever loopbacks are enabled or disabled. Changing the packet datapath with loopbacks while the 1588 engine is processing traffic is not recommended.

3.22.3.1 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 108 • Far-End Loopback Diagram



3.22.3.2 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 109 • Near-End Loopback Diagram



3.22.3.3 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 110 • 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.

3.22.3.4 SerDes Loops

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

Figure 111 • Data Loops of the SerDes Macro



3.22.3.4.1 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 0x3)

Bits 7:4: Loopback type

Bits 3:0: 0x2

where loopback type is:

0x0: No loopback

0x2: Input loopback
0x4: Facility loopback
0x8: Equipment loopback

3.22.3.4.2 QSGMII Mode

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

Bits 15:12: 0x9
Bits 11:8: Port address (0xC)
Bits 7:4: Loopback type
Bits 3:0: 0x2

where loopback type is:

0x0: No loopback
0x2: Input loopback
0x4: Facility loopback
0x8: Equipment loopback

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

3.22.3.4.3 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: Fiber0 port
0x2: Fiber1 port
0x4: Fiber2 port
0x8: Fiber3 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.

3.22.3.4.4 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.

3.22.3.4.5 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.

Note: After entering equipment loopback mode, the following workaround should be run with set= 1 option in case external signal is not present; when exiting equipment loopback mode, the set= 0 option should be run.

SGMII/QSGMII SerDes

```
PhyWrite(<phy>, 31, 0x52b5);
PhyWrite(<phy>, 16, 0xa68c);
tmp17 = PhyRead(<phy>,17);
if (set)
tmp17 |= 0x0010; //Set SigDet as desired, Set bit 4
else // clear SigDet
tmp17 &= 0xffef; //Clear SigDet, bit 4
PhyWrite(<phy>, 17, tmp17);
PhyWrite(<phy>, 16, 0x868c);
PhyWrite(<phy>, 31, 0x0);
```

Fiber Media SerDes

```
PhyWrite(<phy>, 31, 0x52b5);
PhyWrite(<phy>, 16, 0xa68a);
tmp17 = PhyRead(<phy>,17);
if (set)
tmp17 |= 0x0010; //Set SigDet as desired, Set bit 4
else // clear SigDet
tmp17 &= 0xffef; //Clear SigDet, bit 4
PhyWrite(<phy>, 17, tmp17);
PhyWrite(<phy>, 16, 0x868a);
PhyWrite(<phy>, 31, 0x0);
```

3.22.3.4.6 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 test only the analog parts of the SGMII interface because only the input and output buffer are part of this loop.

Note: When the enhanced SerDes macro is in input loopback, the output is inverted relative to the input.

The following table shows the SerDes macro address map.

Table 54 • SerDes Macro Address Map

SerDes Macro	Physical Address (s)	Interface Logical Type (p)	Address
SerDes0	0x0	Fiber0	0x1
SerDes1	0x1	SGMII1	0x1
SerDes2	0x2	Fiber1	0x2
SerDes3	0x3	SGMII2	0x2
SerDes4	0x4	Fiber2	0x4
SerDes5	0x5	SGMII3	0x3
SerDes6	0x6	Fiber3	0x8

3.22.4 VeriPHY Cable Diagnostics

VSC8584-11 includes a comprehensive suite of cable diagnostic functions that are available using SMI reads and writes. These functions enable 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 common faults that can occur on the Cat5 twisted pair cabling.

For the functional details of the VeriPHY suite and the operating instructions, see the ENT-AN0125, *PHY, Integrated PHY-Switch VeriPHY - Cable Diagnostics Feature Application Note*.

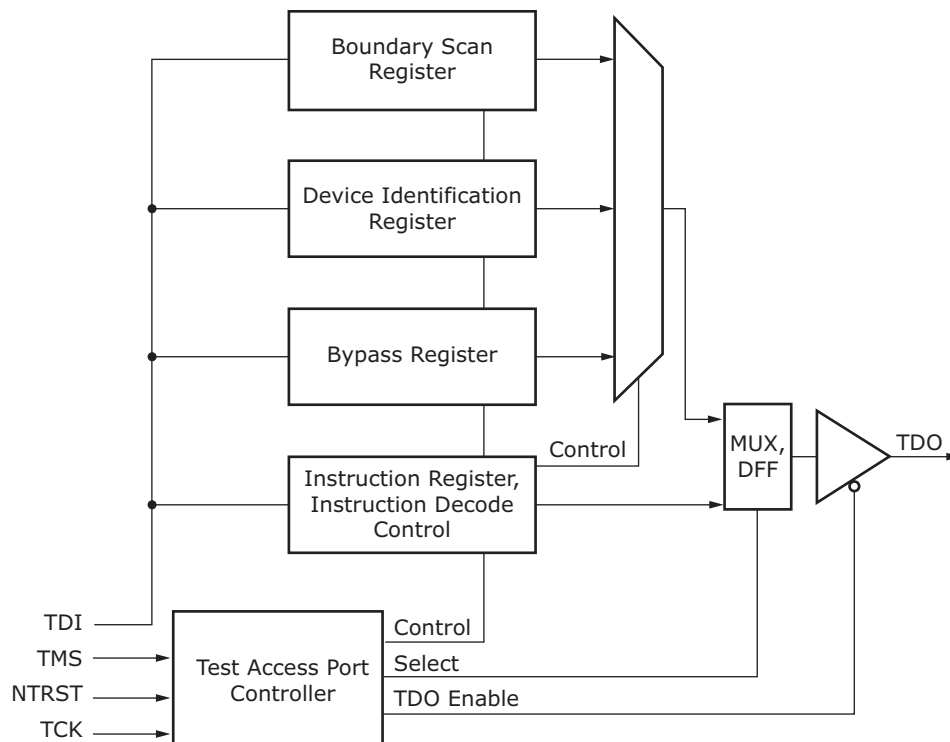
3.2.2.5 JTAG Boundary Scan

The VSC8584-11 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 VSC8584-11, 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 112 • 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'b100000 (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.

3.2.2.6 JTAG Instruction Codes

The VSC8584-11 supports the following instruction codes:

Table 55 • 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.

Table 55 • JTAG Instruction Codes (continued)

Instruction Code	Description
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.
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 56 • 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 1000 0100	000 0111 0100	1

Table 57 • 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	0001	1000 0101 1000 0100	000 0111 0100	1

The following table provides information about the location and IEEE compliance of the JTAG instruction codes used in the VSC8584-11. 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 58 • JTAG Instruction Code IEEE Compliance

Instruction	Code	Selected Register	Register Width	IEEE 1149.1
EXTEST	6'b000000	Boundary Scan	161	Mandatory
SAMPLE/PRELOAD	6'b000001	Boundary Scan	161	Mandatory
IDCODE	6'b100000	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

3.22.7 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 Microsemi Web site at www.microsemi.com.

3.23 100BASE-FX Far-End Fault Indication (FEFI)

The VSC8584-11 device implements Far-End Fault Indication (FEFI) generation and detection per IEEE 802.3-2005 clause 24.3.2.5 and 24.3.2.6 in 100BASE-FX.

FEFI enables stations on both ends of a pair of fibers to be informed when there is a problem with one of the fibers. Without this capability, it is impossible for a fiber interface to detect a problem that affects only its transmit fiber.

When FEFI is supported and enabled, a loss of receive signal (link) causes the transmitter to generate a Far End Fault pattern to inform the device at the far end of the fiber pair that a fault has occurred. When the local receiver again detects a signal, the local transmitter automatically returns to normal operation.

If a Far End Fault pattern is received by a fiber interface that supports Far End Fault and has the feature enabled, it causes link status “down.”

100BASE-FX far-end fault generation force/forceval pair forces the generation of a 100BASE-FX far-end fault or suppresses the automatic generation of the 100BASE-FX far-end fault. This is controlled by 23E3.1:0, which defaults to 00. Bit 1 forces 100BASE-FX far-end fault generation on when 23E3.0 is 1 or off when 23E3.0 is 0.

100BASE-FX far-end fault detection can be determined by reading status bit 27E3.13. This is a sticky bit that indicates the 100BASE-FX far-end fault has been detected since this register was last read. This register is cleared upon reading if 100BASE-FX far-end fault is no longer detected.

3.23.1 100BASE-FX Halt Code Transmission and Reception

The VSC8584-11 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.

3.24 Configuration

The VSC8584-11 can be configured by setting internal memory registers using the management interface. To configure the device, perform the following steps:

1. COMA_MODE active, drive high.
2. Apply power.
3. Apply RefClk and IEEE 1588 reference clock.
4. Release reset, drive high. Power and clock must be stable before releasing reset.
5. Wait 120 ms minimum.
6. Apply init scripts from PHY_API (required for production release, optional for board testing).
7. Configure register 19G for MAC mode (to access register 19G, register 31 must be 0x10). Read register 19G. Set bits 15:14, MAC configuration as follows:
 - 00: SGMII
 - 01: QSGMII
 - 10: Reserved
 - 11: Reserved
 Write new register 19G.
8. Configure register 18G for MAC on all PHY writes:
 - SGMII: 0x80F0
 - QSGMII: 0x80E0
9. Read register 18G until bit 15 equals 0.
10. If Fiber Media on all PHYs configure register 18G by writing:
 - Media 1000BASE-X, Protocol Transfer: 0x8FC1
 - Media 100BASE-FX: 0x8FD1
11. If Fiber Media read register 18G till bit 15 equals 0.
12. Configure register 23 for MAC and Media mode (to access register 23, register 31 must be 0). Read register 23. Set bits 10:8 as follows:
 - 000: Copper
 - 001: Protocol Transfer
 - 010: 1000BASE-X
 - 011: 100BASE-FX
 Write new register 23.
13. Software reset. Read register 0 (to access register 0, register 31 must be 0). Set bit 15 to 1. Write new register 0.
14. Read register 0 until bit 15 equals 0.
15. Apply Enhanced SerDes patch from PHY_API (required).
16. Release the COMA_MODE pin, drive low.

Note: All MAC interfaces must be the same — all QSGMII or SGMII.

3.24.1 Initialization

The COMA_MODE pin provides an optional feature that may be used to control when the PHYs become active. The typical usage is to keep the PHYs from becoming active before they have been fully initialized. For more information, see [Configuration](#), page 137. By not being active until after complete initialization keeps links from going up and down. Alternatively the COMA_MODE pin may be connected low (ground) and the PHYs will be fully active once out of reset.

4 Registers

This section provides information about how to configure the VSC8584-11 device using its internal memory registers and the management interface. The registers marked reserved and factory test should not be read or written to, because doing so may produce undesired effects.

The default value documented for registers is based on the value at reset; however, in some cases, that value may change immediately after reset.

The access type for each register is shown using the following abbreviations:

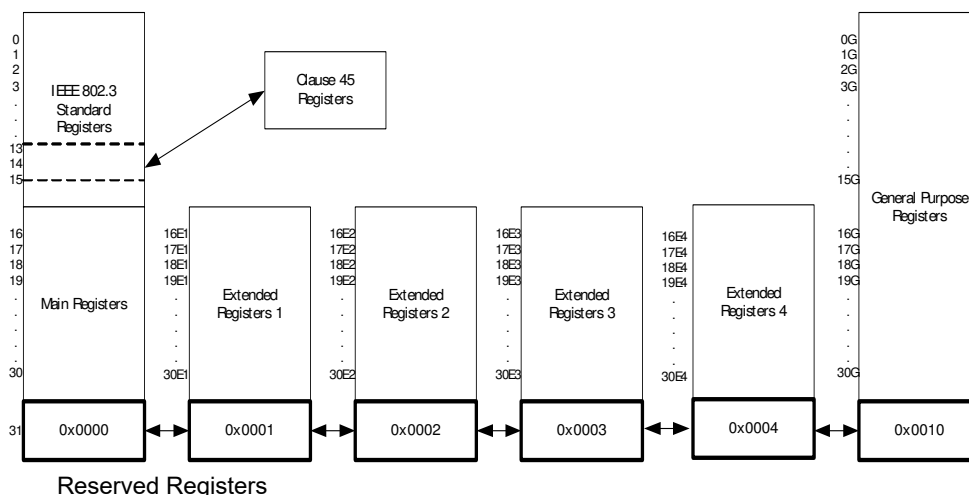
- RO: Read Only
- ROCR: Read Only, Clear on Read
- RO/LH: Read Only, Latch High
- RO/LL: Read Only, Latch Low
- R/W: Read and Write
- RWSC: Read Write Self Clearing

The VSC8584-11 uses several different types of registers:

- IEEE Clause 22 device registers with addresses from 0 to 31
- Four pages of extended registers with addresses from 16E1–30E1, 16E2–30E2, 16E3–30E3, and 16E4–30E4
- 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-2010 energy efficient Ethernet registers and IEEE 802.3bf-2011 registers

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

Figure 113 • Register Space Diagram



For main registers 16–31, extended registers 16E1–30E1, 16E2–30E2, 16E3–30E3, 16E4–30E4, 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 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.

4.2 IEEE 802.3 and Main Registers

In the VSC8584-11, the page space of the standard registers consists of the IEEE 802.3 standard registers and the Microsemi standard registers. The following table lists the names of the registers associated with the addresses as specified by IEEE 802.3.

Table 59 • 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 60 • 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

Table 60 • Main Registers (continued)

Address	Name
24	Extended PHY control 2
25	Interrupt mask
26	Interrupt status
27	Reserved
28	Auxiliary control and status
29	LED mode select
30	LED behavior
31	Extended register page access

4.2.1 Mode Control

The device register at memory address 0 controls several aspects of VSC8584-11 functionality. The following table shows the available bit settings in this register and what they control.

Table 61 • 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 1 μ s 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 ¹	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. ² 00: 10 Mbps. 01: 100 Mbps. 10: 1000 Mbps. 11: Reserved.	10

Table 61 • 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.	0
4:0	Reserved		Reserved.	00000

1. Half-duplex is not supported when the 1588 unit or MACsec unit is operating or if the legacy-MAC EEE feature is enabled.
2. Before selecting the 1000 Mbps forced speed mode, manually configure the PHY as master or slave by setting bit 11 in register 9 (1000BASE-T Control). Each time the link drops, the PHY needs to be powered down manually to enable it to link up again using the master/slave setting specified in register 9.11.

4.2.2 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 62 • 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

Table 62 • Mode Status, Address 1 (0x01) (continued)

Bit	Name	Access	Description	Default
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
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

4.2.3 Device Identification

All 16 bits in both register 2 and register 3 in the VSC8584-11 are used to provide information associated with aspects of the device identification. The following tables list the expected readouts.

Table 63 • Identifier 1, Address 2 (0x02)

Bit	Name	Access	Description	Default
15:0	Organizationally unique identifier (OUI)	RO	OUI most significant bits (3:18)	0x0007

Table 64 • Identifier 2, Address 3 (0x03)

Bit	Name	Access	Description	Default
15:10	OUI	RO	OUI least significant bits (19:24)	000001
9:4	Microsemi model number	RO	VSC8584-11 (0x3C)	111100
3:0	Device revision number	RO	Revision B	0001

4.2.4 Autonegotiation Advertisement

The bits in address 4 in the main registers space control the VSC8584-11 ability to notify other devices of the status of its autonegotiation feature. The following table shows the available settings and readouts.

Table 65 • 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

Table 65 • Device Autonegotiation Advertisement, Address 4 (0x04) (continued)

Bit	Name	Access	Description	Default
7	Advertise100BASE-TX HDX	R/W	1: Advertises 100BASE-TX HDX	1
6	Advertise10BASE-T FDX	R/W	1: Advertises 10BASE-T FDX	1
5	Advertise10BASE-T HDX	R/W	1: Advertises 10BASE-T HDX	1
4:0	Advertise selector	R/W		00001

4.2.5 Link Partner Autonegotiation Capability

The bits in main register 5 can be used to determine if the Cat5 link partner (LP) used with the VSC8584-11 is compatible with the autonegotiation functionality.

Table 66 • 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

4.2.6 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 67 • 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

4.2.7 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 68 • 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		0000000001

4.2.8 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 69 • 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

4.2.9 1000BASE-T/X Control

The VSC8584-11's 1000BASE-T/X functionality is controlled by the bits in register 9 of the main register space. The following table shows the settings and readouts available.

Table 70 • 1000BASE-T/X Control, Address 9 (0x09)

Bit	Name	Access	Description	Default
15:13	Transmitter test mode ¹	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 ¹	R/W	1: Master/slave manual configuration enabled	0

Table 70 • 1000BASE-T/X Control, Address 9 (0x09) (continued)

Bit	Name	Access	Description	Default
11	Master/slave value ¹	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/X FDX capability	R/W	1: PHY is 1000BASE-T/X FDX capable	1
8	1000BASE-T/X HDX capability	R/W	1: PHY is 1000BASE-T/X HDX capable	1
7:0	Reserved	R/W	Reserved	0x00

1. Applies to 1000BASE-T only.

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 1](#), page 151.

4.2.10 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 71 • 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 ¹	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
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

1. Indicates initial state and PCS scrambler in use. It does not change if Ring Resiliency is being used and the timing Master/Slave relationship is changed.

4.2.11 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-2010 Clause 45.

Table 72 • 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-2010 table 45–1

4.2.12 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-2010 Clause 45.

Table 73 • 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.

4.2.13 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 74 • 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
11:0	Reserved	RO	Reserved	0x000

4.2.14 100BASE-TX/FX Status Extension

Register 16 in the main registers page space of the VSC8584-11 provides additional information about the status of the device's 100BASE-TX/100BASE-FX operation.

Table 75 • 100BASE-TX/FX Status Extension, Address 16 (0x10)

Bit	Name	Access	Description	Default
15	100BASE-TX/FX Descrambler	RO	1: Descrambler locked	0
14	100BASE-TX/FX lock error	RO	Self-clearing bit. 1: Lock error detected	0

Table 75 • 100BASE-TX/FX Status Extension, Address 16 (0x10) (continued)

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

4.2.15 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 74](#), page 147.

Table 76 • 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
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

Table 76 • 1000BASE-T Status Extension 2, Address 17 (0x11) (continued)

Bit	Name	Access	Description	Default
4:0	Reserved	RO	Reserved	

4.2.16 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 77 • 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
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.

4.2.17 Error Counter 1

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

Table 78 • Extended Control and Status, Address 19 (0x13)

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

4.2.18 Error Counter 2

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

Table 79 • Extended Control and Status, Address 20 (0x14)

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

4.2.19 Error Counter 3

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

Table 80 • 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

4.2.20 Extended Control and Status

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

Table 81 • 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

Table 81 • Extended Control and Status, Address 22 (0x16) (continued)

Bit	Name	Access	Description	Default
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.

4.2.21 Extended PHY Control Set 1

The following table shows the settings available.

Table 82 • Extended PHY Control 1, Address 23 (0x17)

Bit	Name	Access	Description	Default
15:13	Reserved	R/W	Reserved	0
12	MAC interface mode	R/W	Super-sticky bit. 0: SGMII 1: 1000BASE-X. Note: Register 19G.15:14 must be = 00 for this selection to be valid.	0
11	AMS preference	R/W	Super-sticky bit. 1: Cat5 copper preferred. 0: SerDes fiber/SFP preferred.	0

Table 82 • Extended PHY Control 1, Address 23 (0x17) (continued)

Bit	Name	Access	Description	Default
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: Reserved.	000
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	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.

4.2.22 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 83 • 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)	000
12	PICMG 2.16 reduced power mode	R/W	Sticky bit. 1: Enabled	0
11:6	Reserved	RO	Reserved	

Table 83 • Extended PHY Control 2, Address 24 (0x18) (continued)

Bit	Name	Access	Description	Default
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.

4.2.23 Interrupt Mask

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

Table 84 • 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	Reserved	R/W	Reserved	0
5	Extended interrupt mask	R/W	Sticky bit. 1: Enabled.	0
4	AMS media changed mask ¹	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

1. If hardware interrupts are not used, the mask can still be set and the status polled for changes.

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 26.15 to be set.

4.2.24 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 85 • 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
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	Reserved	RO		0
5	Extended interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
4	AMS media changed status ¹	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

1. If hardware interrupts are not used, the mask can still be set and the status polled for changes.

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.

4.2.25 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 86 • Auxiliary Control and Status, Address 28 (0x1C)

Bit	Name	Access	Description	Default
15	Autonegotiation complete	RO	Duplicate of bit 1.5 when autonegotiation is enabled, otherwise this is the current link status	0

Table 86 • Auxiliary Control and Status, Address 28 (0x1C) (continued)

Bit	Name	Access	Description	Default
14	Autonegotiation disabled	RO	Inverted duplicate of bit 0.12 or AMS-enabled with 100BASE-FX operating mode selected	0
13 ¹	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
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: 2.3 seconds 01: 3.3 seconds 10: 4.3 seconds 11: 5.3 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: 2.3 seconds 01: 3.3 seconds 10: 4.3 seconds 11: 5.3 seconds	1
1:0	Media mode status	RO	00: No media selected 01: Copper media selected 10: SerDes (Fiber) media selected 11: Reserved	00

1. In 1000BT mode, if Force MDI crossover is performed while link is up, the 1000BT link must be re-negotiated in order for this bit to reflect the actual Auto-MDIX setting.

4.2.26 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 50](#), page 123. For information about enabling the extended LED mode bits in Register 19E1 bits 13 to 12, see [Table 51](#), page 124.

Table 87 • 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

4.2.27 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 88 • LED Behavior, Address 30 (0x1E)

Bit	Name	Access	Description	Default
15	Copper and fiber LED combine disable	R/W	Sticky bit 0: Combine enabled (Copper/Fiber on link/linkXXXX/activity LED) 1: Disable combination (link/linkXXXX/activity LED; indicates copper only)	0
14	Activity output select	R/W	Sticky bit 1: Activity LED becomes TX_Activity and fiber activity LED becomes RX_Activity 0: Tx and Rx activity both displayed on activity LEDs	0
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	Reserved	RO	Reserved	
8	LED3 pulse-stretch/blink select	R/W	Sticky bit 1: Pulse-stretch 0: Blink	0
7	LED2 pulse-stretch/blink select	R/W	Sticky bit 1: Pulse-stretch 0: Blink	0
6	LED1 pulse-stretch/blink select	R/W	Sticky bit 1: Pulse-stretch 0: Blink	0

Table 88 • LED Behavior, Address 30 (0x1E) (continued)

Bit	Name	Access	Description	Default
5	LED0 pulse-stretch/blink select	R/W	Sticky bit 1: Pulse-stretch 0: Blink	0
4	Reserved	RO	Reserved	
3	LED3 combine feature disable	R/W	Sticky bit 0: Combine enabled (link/activity, duplex/collision) 1: Disable combination (link only, duplex only)	0
2	LED2 combine feature disable	R/W	Sticky bit 0: Combine enabled (link/activity, duplex/collision) 1: Disable combination (link only, duplex only)	0
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.

4.2.28 Extended Page Access

To provide functionality beyond the IEEE 802.3-specified registers and main device registers, the VSC8584-11 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 VSC8584-11. Accessing the GPIO page register space is similar to accessing the extended page registers. The following table shows the settings available.

Table 89 • 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 0x0004: Registers 16-30 access extended register space 4 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 90 • Extended Registers Page 1 Space

Address	Name
16E1	SerDes Media Control
17E1	Reserved
18E1	Cu Media CRC good counter
19E1	Extended mode and SIGDET control
20E1	Extended PHY control 3 (ActiPHY)
21E1–22E1	Reserved
23E1	Extended PHY control 4 (PoE and CRC error counter)
24E1	Reserved
25E1	Reserved
26E1	Reserved
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 91 • 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 92 • 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 with valid CRCs modulo 10,000; this counter does not saturate and will roll over to zero on the next good packet received after 9,999.	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 93 • Extended Mode Control, Address 19E1 (0x13)

Bit	Name	Access	Description	Default
15	LED3 Extended Mode	R/W	1: See Extended LED Modes , page 124	0
14	LED2 Extended Mode	R/W	1: See Extended LED Modes , page 124	0
13	LED1 Extended Mode	R/W	1: See Extended LED Modes , page 124	0
12	LED0 Extended Mode	R/W	1: See Extended LED Modes , page 124	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
10:5	Reserved	RO	Reserved	0
4	Fast link failure	R/W	Enable fast link failure pin. This must be done from PHY0 only. 1: Enabled 0: Disabled (GPIO9 pin becomes general purpose I/O)	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	SIGDET pin polarity 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 94 • Extended PHY Control 3, Address 20E1 (0x14)

Bit	Name	Access	Description	Default
15	Disable carrier extension	R/W	1: Disable carrier extension in 1000BASE-T copper links	1
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	00
10	Slow MDC	R/W	1: Indicates that MDC runs at less than 10 MHz (use of this bit is optional and indicated when MDC runs at less than 1 MHz)	0
9	PHY address reversal	R/W	Reverse PHY address Enabling causes physical PHY 0 to have address of 3, PHY 1 address of 2, PHY 2 address of 1, and PHY 3 address of 0. Changing this bit to 1 should initially be done from PHY 0 and changing to 0 from PHY3 1: Enabled 0: Disabled Valid only on PHY0	0
8	Reserved	RO	Reserved	
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

Table 94 • Extended PHY Control 3, Address 20E1 (0x14) (continued)

Bit	Name	Access	Description	Default
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 VSC8584-11.

Table 95 • Extended PHY Control 4, Address 23E1 (0x17)

Bit	Name	Access	Description	Default
15:11	PHY address	RO	Internal PHY address. 00000: PHY 0 00001: PHY 1 00010: PHY 2 00011: PHY 3 others: Reserved	
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	

CRC 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 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 96 • 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

Table 96 • EPG Control Register 1, Address 29E1 (0x1D) (continued)

Bit	Name	Access	Description	Default
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	Bit times 1: 8,192 0: 96	0
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 source 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 VSC8584-11 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 F0 through 0xFF 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 F0 through 0xFF 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.

4.3.7 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 97 • 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 89](#), page 157.

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 98 • Extended Registers Page 2 Space

Address	Name
16E2	Cu PMD Transmit Control
17E2	EEE Control
18E2	Extended Chip ID
19E2	Entropy data
20E2-27E2	Reserved
28E2	Extended Interrupt Mask
29E2	Extended Interrupt Status
30E2	Ring Resiliency Control

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. For help with changing these values, contact your Microsemi representative.

Table 99 • Cu PMD Transmit Control, Address 16E2 (0x10)

Bit	Name	Access	Description	Default
15:12	1000BASE-T signal amplitude trim ¹	R/W	1000BASE-T signal amplitude 1111: -1.7% 1110: -2.6% 1101: -3.5% 1100: -4.4% 1011: -5.3% 1010: -7% 1001: -8.8% 1000: -10.6% 0111: 5.5% 0110: 4.6% 0101: 3.7% 0100: 2.8% 0011: 1.9% 0010: 1% 0001: 0.1% 0000: -0.8%	0000

Table 99 • Cu PMD Transmit Control, Address 16E2 (0x10) (continued)

Bit	Name	Access	Description	Default
11:8	100BASE-TX signal amplitude trim ²	R/W	100BASE-TX signal amplitude 1111: -1.7% 1110: -2.6% 1101: -3.5% 1100: -4.4% 1011: -5.3% 1010: -7% 1001: -8.8% 1000: -10.6% 0111 5.5% 0110: 4.6% 0101: 3.7% 0100: 2.8% 0011: 1.9% 0010: 1% 0001: 0.1% 0000: -0.8%	0010
7:4	10BASE-T signal amplitude trim ³	R/W	10BASE-T signal amplitude 1111: -7% 1110: -7.9% 1101: -8.8% 1100: -9.7% 1011: -10.6% 1010: -11.5% 1001: -12.4% 1000: -13.3% 0111: 0% 0110: -0.7% 0101: -1.6% 0100: -2.5% 0011: -3.4% 0010: -4.3% 0001: -5.2% 0000: -6.1%	1011
3:0	10BASE-Te signal amplitude trim	R/W	10BASE-Te signal amplitude 1111: -30.45% 1110: -31.1% 1101: -31.75% 1100: -32.4% 1011: -33.05% 1010: -33.7% 1001: -34.35% 1000: -35% 0111: -25.25% 0110: -25.9% 0101: -26.55% 0100: -27.2% 0011: -27.85% 0010: -28.5% 0001: -29.15% 0000: -29.8%	1110

- Changes to 1000BASE-T amplitude may result in unpredictable side effects.
- Adjust 100BASE-TX to specific magnetics.
- Amplitude is limited by V_{CC} (2.5 V).

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-2010) mode for debug.

Table 100 • EEE Control, Address 17E2 (0x11)

Bit	Name	Access	Description	Default
15	Enable 10BASE-Te	R/W	Sticky bit. Enable energy efficient (IEEE 802.3az-2010) 10BASE-Te operating mode.	0
14	Enable LED in fiber unidirectional mode	R/W	Sticky bit. 1: Enable LED functions in fiber unidirectional mode.	0
13:10	Invert LED polarity	R/W	Sticky bits. 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 126.	0000
9	Reserved	RO	Reserved.	
8	Link status	RO	1: Link is up.	0
7	1000BASE-T EEE enable	RO	1: EEE is enabled for 1000BASE-T.	0
6	100BASE-TX EEE enable	RO	1: EEE is enabled for 100BASE-TX.	0
5	Enable 1000BASE-T force mode	R/W	Sticky bit. 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 ¹	R/W	Sticky bit. 1: Enable the EPG to transmit LPI on the MDI, ignore data from the MAC interface. 0: Transmit idles being received from the MAC.	0
3	Inhibit 100BASE-TX transmit EEE LPI	R/W	Sticky bit. 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	Sticky bit. 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	Sticky bit. 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	Sticky bit. 1: Disable transmission of EEE LPI on receive path MAC interface in 1000BASE-T mode when receiving LPI from the MDI.	0

1. Register 17E2 bits 4:0 are for debugging purposes only, not for operational use.

4.4.3 Extended Chip ID, Address 18E2 (0x12)

The following table shows the register settings for the extended chip ID at address 18E2.

Table 101 • Extended Chip ID, Address 18E2 (0x12)

Bit	Name	Access	Description	Default
15	Industrial temperature capable	RO	VSC8584-11 1: Industrial temperature capable 0: Commercial temperature capable	0
15	Industrial temperature capable	RO	VSC8584-14 1: Industrial temperature capable 0: Commercial temperature capable	1
14	Quad/dual device	RO	1: Quad device 0: Dual device	1
13	1588 capable	RO	1: 1588 operation capable 0: Not 1588 operation capable	1
12	MACsec capable	RO	1: MACsec capable 0: Not MACsec capable	1
11	Reserved	RO	Reserved	0
10	Dual media device	RO	1: Dual media capable 0: Not dual media capable	1
9	1588 high-precision capable	RO	1: 1588 high-precision capable 0: 1588 low-precision capable only	0
8	MACsec 256-bit keys capable	RO	1: MACsec 256-bit key capable 0: MACsec 128-bit key capable	1
7:0	Extended chip ID	RO	Dash number of VSC8584-11 in BCD	0x11
7:0	Extended chip ID	RO	Dash number of VSC8584-14 in BCD	0x14

4.4.4 Entropy Data, Address 19E2 (0x13)

The following table shows the register settings for the entropy data at address 19E2.

Table 102 • Entropy Data, Address 19E2 (0x13)

Bit	Name	Access	Description	Default
15:0	Entropy data	RO	Random data that can be added to an entropy pool	

4.4.5 Extended Interrupt Mask, Address 28E2 (0x1C)

The following table shows the register settings for the extended interrupt mask at address 28E2.

Table 103 • Extended Interrupt Mask, Address 28E2 (0x1C)

Bit	Name	Access	Description	Default
15:11	Reserved	R/W	Reserved.	00000
10	Mem integrity ring control interrupt mask	R/W	Sticky bit. 1: Enabled.	0
9	MACsec egress interrupt mask	R/W	Sticky bit. 1: Enabled.	0
8	MACsec ingress interrupt mask	R/W	Sticky bit. 1: Enabled.	0
7	MACsec flow control buffer interrupt mask	R/W	Sticky bit. 1: Enabled.	0
6	MACsec line MAC interrupt mask	R/W	Sticky bit. 1: Enabled.	0

Table 103 • Extended Interrupt Mask, Address 28E2 (0x1C) (continued)

Bit	Name	Access	Description	Default
5	MACsec host MAC interrupt mask	R/W	Sticky bit. 1: Enabled.	0
4	RR switchover complete interrupt mask	R/W	Sticky bit. 1: Enabled.	0
3	EEE link fail interrupt mask	R/W	Sticky bit. 1: Enabled.	0
2	EEE Rx TQ timer interrupt mask	R/W	Sticky bit. 1: Enabled.	0
1	EEE wait quiet/Rx TS timer interrupt mask	R/W	Sticky bit. 1: Enabled.	0
0	EEE wake error interrupt mask	R/W	Sticky bit. 1: Enabled.	0

4.4.6 Extended Interrupt Status, Address 29E2 (0x1D)

The following table shows the register settings for the extended interrupt status at address 29E2.

Table 104 • Extended Interrupt Status, Address 29E2 (0x1D)

Bit	Name	Access	Description	Default
15:11	Reserved	RO	Reserved.	00000
10	Mem integrity ring control interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
9	MACsec egress interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
8	MACsec ingress interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
7	MACsec flow control buffer interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
6	MACsec line MAC interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
5	MACsec host MAC interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
4	RR switchover complete interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
3	EEE link fail interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
2	EEE Rx TQ timer interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
1	EEE wait quiet/Rx TS timer interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0
0	EEE wake error interrupt status	RO	Self-clearing bit. 1: Interrupt pending.	0

4.4.7 Ring Resiliency Control (0x1E)

The following table shows the register settings for the ring resiliency controls at address 30E2.

Table 105 • Ring Resiliency, Address 30E2 (0x1E)

Bit	Name	Access	Description	Default
15	Ring resiliency startup enable (master TR enable)	R/W	Sticky	0
14	Advertise ring resiliency	R/W	Sticky	0
13	LP ring resiliency advertisement	RO		0
12	Force ring resiliency enable (override autoneg)	R/W	Sticky	0
11:6	Reserved	RO	Reserved	000000
5:4	Ring resiliency status	RO	Ring resiliency status 00: Timing slave ¹ 10: Timing slave becoming master 11: Timing master ¹ 01: Timing master becoming slave	00
3:1	Reserved	RO	Reserved	000
0	Start switchover (only when not in progress)	RWSC		0

1. Reflects autoneg master/slave at initial link-up.

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 89](#), page 157.

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 106 • 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	Media/MAC SerDes Transmit Good Packet Counter
22E3	Media/MAC SerDes Transmit CRC Error Counter

Table 106 • Extended Registers Page 3 Space (continued)

Address	Name
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/MAC SerDes Receive SerDes status
28E3	Media/MAC SerDes Receive CRC Good Counter
29E3	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 107 • 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
10:9	SGMII input preamble for 100BASE-FX	R/W	This is a sticky bit. 00: No SGMII preamble required. 01: One-Byte SGMII preamble required. 10: Two-Byte SGMII preamble required. 11: Reserved.	00
8	SGMII output preamble	R/W	This is a sticky bit. 0: No SGMII preamble. 1: Two-Byte SGMII preamble.	1
7	MAC SerDes autonegotiation enable	R/W	This is a sticky bit. 1: MAC SerDes ANEG enable.	0
6	SerDes polarity at input of MAC	R/W	This is a sticky bit. 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.	0

Table 107 • MAC SerDes PCS Control, Address 16E3 (0x10) (continued)

Bit	Name	Access	Description	Default
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	Reserved	R/W	Reserved.	0
2	Inhibit MAC odd-start delay	R/W	This is a sticky bit. 1: Inhibits delay of 1 byte when receive packet begins on an odd-byte boundary (causes the first 0x55 byte of preamble to be removed on odd-byte alignment) 0: Allows delay on odd-byte aligned packets preserving all preamble bytes but introducing a delay variation between naturally even-byte aligned and odd-byte aligned packets.	1
1:0	Reserved	RO	Reserved.	0

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 108 • MAC SerDes PCS Status, Address 17E3 (0x11)

Bit	Name	Access	Description
15	MAC sync status failed	RO	1: Sync status on MAC SerDes has failed since last read
14	MAC cgbad received	RO	1: an invalid code-group was received on the MAC SerDes since last read
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
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 109 • MAC SerDes CI37 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) ¹	0x0000

1. The read value for this register is N/A for protocol transfer mode when 16E3.11 is not set.

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 110 • MAC SerDes CI37 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 111 • MAC SerDes Status, Address 20E3 (0x14)

Bit	Name	Access	Description
15	Comma realigned	RO	Self-clearing bit. Sticky bit. 1: MAC SerDes receiver comma was realigned.
14	SerDes signal detect	RO	Self-clearing bit. Sticky bit. 1: SerDes signal detection occurred.
13	QSGMII sync status	RO	Only applies on PHY0
12	MAC comma detect	RO	Self-clearing bit. Sticky bit. 1: Comma detected, cleared when comma not detected, reset to 1 upon read.
11:8	MAC comma position	RO	MAC comma-alignment position from 0 to 9. These bits are N/A for QSGMII.
7:0	Reserved	RO	Reserved.

4.5.6 Media/MAC SerDes Transmit Good Packet Counter

The register at address 21E3 consists of the bits that provide status of the media and MAC SerDes transmit good packet counter. The following table shows the settings available.

Table 112 • Media/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 Media/MAC SerDes Transmit CRC Error Counter

The register at address 22E3 consists of the bits that provide status of the media and MAC SerDes transmit packet count that had a CRC error. The following table shows the settings available.

Table 113 • Media/MAC SerDes Tx CRC Error Counter, Address 22E3 (0x16)

Bit	Name	Access	Description	Default
15:14	Tx counter select	R/W	Selects between fiber media and MAC SerDes transmit counters ¹ 00: Selects fiber media SerDes transmit counters 01: Selects MAC SerDes transmit counters others: Reserved	0
13	Tx preamble fix	R/W	Removes extraneous byte of preamble for egress frames. Only removes one byte if there are more than eight bytes present in the preamble.	0
12:8	Reserved	RO	Reserved	0
7:0	Tx CRC packet count	RO	Transmit CRC packet count (saturates at 255)	0

1. The counters are only operational when the media link state is up, irrespective of selecting media SerDes or MAC SerDes.

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 114 • Media SerDes PCS Control, Address 23E3 (0x17)

Bit	Name	Access	Description	Default
15:14	Remote fault to media	RO	Remote fault indication sent to media in most recent clause 37 auto-negotiation exchange.	
13	Media interface autonegotiation parallel-detection ¹	R/W	Sticky bit. 1: SerDes media autonegotiation parallel detect enabled.	0
12	Disable carrier extension	R/W	Disable carrier extension in 1000BASE-X fiber links.	1
11	Force advertised ability	R/W	1: Force 16-bit advertised ability from register 25E3.15:0.	0
10:7	Reserved	RO	Reserved.	
6	Polarity reversal input	R/W	This is a sticky bit. Media SerDes polarity reversal input. 0: No polarity reversal (default). 1: Polarity reversed.	0
5	Polarity reversal output	R/W	This is a sticky bit. Media SerDes polarity reversal output. 0: No polarity reversal (default). 1: Polarity reversed.	0

Table 114 • Media SerDes PCS Control, Address 23E3 (0x17) (continued)

Bit	Name	Access	Description	Default
4	Inhibit odd-start delay	R/W	This is a sticky bit. 1: Inhibits delay of one byte when transmit packet begins on an odd-byte boundary (causes the first 0x55 byte of preamble to be removed on odd-byte alignment). 0: Allows delay on odd-byte aligned packets preserving all preamble bytes but introducing a delay variation between naturally even-byte aligned and odd-byte aligned packets.	1
3	Reserved	RO	Reserved.	
2	100BASE-FX force HLS	R/W	1: Forces 100BASE-FX to transmit Halt Line State (HLS) continuously. 0: Normal 100BASE-FX transmit operation.	0
1	100BASE-FX force FEFI	R/W	1: Forces 100BASE-FX Far-End Fault Indication (FEFI) as specified by bit 0. 0: Normal automatic operation of FEFI in 100BASE-FX.	0
0	100BASE-FX FEFI force value	R/W	1: Forces FEFI on when bit 1 is asserted. 0: Suppresses FEFI when bit 1 is asserted.	0

1. Only applicable when clause 37 auto-negotiation is enabled. Enabling parallel detection along with clause 37 auto-negotiation functionality requires local PHY to advertise full-duplex operation.

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 115 • Media SerDes PCS Status, Address 24E3 (0x18)

Bit	Name	Access	Description
15	Sync status failed	RO	1: Sync status on fiber-media SerDes has failed since last read
14	cgbad received	RO	1: Invalid code-group was received on the fiber-media SerDes since last read
13	SerDes protocol transfer	RO	100 Mb or 100BASE-FX link status
12	SerDes protocol transfer	RO	10 Mb link status
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

Table 115 • Media SerDes PCS Status, Address 24E3 (0x18) (continued)

Bit	Name	Access	Description
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 116 • Media SerDes CI37 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. ¹	0x0000

1. The read value for this register is N/A for protocol transfer mode when 23E3.11 is not set.

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 117 • MAC SerDes CI37 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 118 • 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.

Table 118 • Media SerDes Status, Address 27E3 (0x1B) (continued)

Bit	Name	Access	Description
13	100BASE-FX FEFI detect	RO	1: 100BASE-FX far-end fault detected from link partner since last read.
12	Comma detect	RO	Self-clearing bit. Sticky bit. 1: Comma detected, cleared when comma not detected, reset to 1 upon read.
11:8	Comma position	RO	Fiber media SerDes comma-alignment position from 0 to 9.
7	100BASE-FX HLS detected	RO	1: 100BASE-FX Halt Line-State (HLS) detected since last read.
6:0	Reserved	RO	Reserved.

4.5.13 Media/MAC SerDes Receive 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 and MAC interfaces; the number of packets that have been received successfully. The following table shows the expected readouts.

Table 119 • Media/MAC SerDes Receive 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	Media/MAC SerDes Receive 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 to 0 when the count reaches 10,000 packets.	0x000

4.5.14 Media/MAC SerDes Receive 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 and MAC interfaces. The following table shows the expected readouts.

Table 120 • Media/MAC SerDes Receive CRC Error Counter, Address 29E3 (0x1D)

Bit	Name	Access	Description	Default
15:14	Rx counter select	RW	Selects between fiber media and MAC SerDes receive counters. ¹ 00: Selects fiber media SerDes receive counters. 01: Selects MAC SerDes receive counters. others: Reserved.	
13:8	Reserved	RO	Reserved.	
7:0	Media/MAC Receive CRC error counter	RO	Self-clearing bit. CRC error counter for packets received on the Fiber media or MAC interfaces. The value saturates at 0xFF and subsequently clears when read and restarts count.	0x00

1. The counters are only operational when the media link state is up, irrespective of selecting media SerDes or MAC SerDes.

4.6 Extended Page 4 Registers

To access the extended page 4 registers (16E4–30E4), enable extended register access by writing 0x0004 to register 31. For more information, see [Table 89](#), page 157.

When extended page 4 register access is enabled, reads and writes to registers 16–30 affect the extended registers 16E4–30E4 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 4 space. These registers are accessible only when the device register 31 is set to 0x0004.

Table 121 • Extended Registers Page 4 Space

Address	Name
16E4–20E4	CSR Access Controls and Status
21E4	1588_PPS_0/1 Mux Control
22E4–25E4	Reserved
26E4–28E4	SPI Daisy-Chain Controls and Status
29E4–30E4	1588 RefClk Input Buffer Control

4.6.1 CSR Access Controls and Status

The following tables show the CSR ring access controls and status registers.

Table 122 • CSR Access Control, Address 16E4

Bit	Access	Description
15	RWSC	Command bit. 1: Must be set to execute the command. It is set back to 1 when done. 0: Command busy, do not do any write to register 16. Register 17 and 18 maintain previous write values.
14	RW	1: Execute a read on the CSR registers. 0: Execute a write on the CSR registers.
13:11	RW	Target block code. 000: Analyzer 0 Ingress 001: Analyzer 0 Egress 010: Analyzer 1 Ingress 011: Analyzer 1 Egress 100: Analyzer 2 Ingress 101: Analyzer 2 Egress 110: Processor 0 111: Processor 1
10:0	R/W	CSR register address[10:0]

Table 123 • CSR Buffer, Address 17E4

Bit	Access	Description
15:0	RWSC	CSR Data_LSB[15:0]

Table 124 • CSR Buffer, Address 18E4

Bit	Access	Description
15:0	RWSC	CSR Data_MSB[31:16]

Table 125 • CSR Access Control, Address 19E4

Bit	Access	Description
15	RWSC	Command bit. 1: Must be set to execute the command. It is set back to 1 when done. 0: Command busy, do not do any write to register 19. Register 17 and 18 maintain previous write values.
14	RW	1: Execute a read on the CSR registers. 0: Execute a write on the CSR registers.
13:12	RW	Target ID [1:0] for most targets CSR register address[13:12] for MACsec INGR/EGR Targets
11:0	R/W	CSR register address[11:0]

Table 126 • CSR Status, Address 20E4

Bit	Access	Description
15:13	RO	CSR status 000: REQUEST_OK 001: TGT_BUSY 010: UTM 011: NO_ACTION 100: WD_DROP 101: WD_DROP_ORG
12:4	RO	Reserved
3:0	RW	Target ID[5:2]

4.6.2 1588_PPS_0/1 Mux Control

The 1588_PPS_0 mux control register controls the PHY used to access 1588_PPS_0. The following table shows the settings available. For more information, see [Figure 54](#), page 62.

Table 127 • 1588_PPS_0 Mux Control, Address 21E4

Bit	Name	Access	Description	Default
15:2	Reserved	RO	Reserved	
1:0	1588_PPS_0 control	R/W	00: PPS_0 from Phy0 01: PPS_0 from Phy1 10: PPS_0 from Phy2 11: PPS_0 from Phy3	00

4.6.3 SPI Daisy-Chain Controls and Status

The following tables show the SPI daisy-chain controls and status registers.

Register 26

Table 128 • SPI Daisy-Chain Control, Address 26E4

Bit	Access	Description
15	RW	Enable SPI daisy-chain input port
14	RW	Enable SPI daisy-chain output port
13	RW	Output SI clock phase control
12	RW	Output SI clock polarity control
11:8	RW	Number of CSR clock periods SI_CS negates between writes
7:4	RW	Threshold (units 1/16 of FIFO size) which enables SPI daisy-chain as highest priority
3:0	RW	Threshold (units 1/16 of FIFO size) which enables SPI daisy-chain as equal priority

Table 129 • SPI Daisy-Chain Status, Address 27E4

Bit	Access	Description
15:12	RW	Number of CSR clock periods after last clock edge before SI_CS goes high at end-of-frame
11:8	RW	Number of CSR clock periods before first clock edge after SI_CS goes low at start-of-frame
7:4	RW	Number of CSR clock periods that the SI_CLK is high
3:0	RW	Number of CSR clock periods that the SI_CLK is low

Table 130 • SPI Daisy-Chain Counter, Address 28E4

Bit	Access	Description
15:14	RW	Selects SPI daisy-chain counter
13:10	RO	Reserved
9:0	RO/SC	Reads out the selected counter (clear-on-read, if appropriate)

Table 131 • 1588 RefClk Input Buffer Control (LSW), Address 29E4

Bit	Access	Description
15:0	RW	REFCLK_1588_IB_CTRL[15:0]

Table 132 • 1588 RefClk Input Buffer Control (MSW), Address 30E4

Bit	Access	Description
15:0	RW	REFCLK_1588_IB_CTRL[31:16]

4.7 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.

This register space is global in nature to all four PHYs in the VSC8584-11 device.

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. All general purpose register bits are super-sticky.

Table 133 • 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	Micro Command
19G	MAC Mode and Fast Link Configuration
20G	Two-Wire Serial MUX Control 1
21G	Two-Wire Serial MUX Control 2
22G	Two-Wire Serial MUX Data Read/Write
23G	Recovered Clock 1 Control
24G	Recovered Clock 2 Control
25G	Enhanced LED Control
26G	Reserved
27G	Reserved
28G	Reserved
29G	Global Interrupt Status
30G	Reserved

4.7.1 Reserved General Purpose Address Space

The bits in registers 0G to 12G and 30G of the general purpose register space are reserved.

4.7.2 LED/SIGDET/GPIO Control

The LED control bits configure the LED[3:0]_[31:0] 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 134 • LED/SIGDET/GPIO Control, Address 13G (0x0D)

Bit	Name	Access	Description	Default
15:14	GPIO7/I2C_SCL_3	R/W	00: SCL for PHY3 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
13:12	GPIO6/I2C_SCL_2	R/W	00: SCL for PHY2 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00

Table 134 • LED/SIGDET/GPIO Control, Address 13G (0x0D) (continued)

Bit	Name	Access	Description	Default
11:10	GPIO5/I2C_SCL_1	R/W	00: SCL for PHY1 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
9:8	GPIO4/I2C_SCL_0	R/W	00: SCL for PHY0 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
7:6	GPIO3/SIGDET3 control	R/W	00: SIGDET operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
5:4	GPIO2/SIGDET2 control	R/W	00: SIGDET operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
3:2	GPIO1/SIGDET1 control	R/W	00: SIGDET operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00
1:0	GPIO0/SIGDET0 control	R/W	00: SIGDET operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	00

4.7.3 GPIO Control 2

The GPIO control 2 register configures the functionality of the COMA_MODE and 1588 control input pins, and provides control for possible GPIO pin options.

Table 135 • GPIO Control 2, Address 14G (0x0E)

Bit	Name	Access	Description	Default
15:14	GPIO12/1588_SPI_CS and GPIO13/1588_SPI_DO	R/W	GPIO12/1588_SPI_CS and GPIO13/1588_SPI_DO control. 00: 1588_SPI_CS/1588_SPI_DO operation. 01: Reserved. 10: Reserved. 11: GPIO12/GPIO13 operation. Controlled by MII registers 15G to 17G.	
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.	

Table 135 • GPIO Control 2, Address 14G (0x0E) (continued)

Bit	Name	Access	Description	Default
10	Tri-state enable for two-wire serial bus	R/W	1: Tri-states two-wire serial 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 two-wire serial bus output signals to high and low values as appropriate.	1
9	Tri-state enable for LEDs	R/W	1: Tri-state LED output signals instead of driving them high. This allows the signals to be pulled above VDDIO using an external pull-up resistor. 0: Drive LED bus output signals to high and low values.	1
8	PPS 1-3 output enable	R/W	PPS 1-3 output enable (must be 0 to allow SPI daisy-chain input).	0
7:6	GPIO11/1588_PPS_0	R/W	GPIO11/1588_PPS_0 control. 00: 1588_PPS_0 operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	
5:4	GPIO10/1588_LOAD_SAVE	R/W	GPIO10/1588_LOAD_SAVE control. 00: 1588_LOAD_SAVE operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	
3:2	GPIO9/FASTLINK_FAIL	R/W	GPIO9/FASTLINK_FAIL control. 00: FASTLINK_FAIL operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	
1:0	GPIO8/I2C_SDA	R/W	GPIO8/I2C_SDA control. 00: I2C_SDA operation 01: Reserved 10: Reserved 11: Controlled by MII registers 15G to 17G	

4.7.4 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 136 • GPIO Input, Address 15G (0x0F)

Bit	Name	Access	Description	Default
15:14	Reserved	RO	Reserved	
13	GPIO13/1588_SPI_DO	R/W	GPIO13/1588_SPI_DO input	0
12	GPIO12/1588_SPI_CS	R/W	GPIO12/1588_SPI_CS input	0
11	GPIO11/1588_PPS_0	R/W	GPIO11/1588_PPS_0 input	0
10	GPIO10/1588_LOAD_SAVE	R/W	GPIO10/1588_LOAD_SAVE input	0
9	GPIO9/FASTLINK_FAIL	R/W	GPIO9/FASTLINK_FAIL input	0

Table 136 • GPIO Input, Address 15G (0x0F) (continued)

Bit	Name	Access	Description	Default
8	GPIO8/I2C_SDA	R/W	GPIO8/I2C_SDA input	0
7	GPIO7/I2C_SCL_3	R/W	GPIO7/I2C_SCL_3 input	0
6	GPIO6/I2C_SCL_2	R/W	GPIO6/I2C_SCL_2 input	0
5	GPIO5/I2C_SCL_1	R/W	GPIO5/I2C_SCL_1 input	0
4	GPIO4/I2C_SCL_0	R/W	GPIO4/I2C_SCL_0 input	0
3	GPIO3/SIGDET3	R/W	GPIO3/SIGDET3 input	0
2	GPIO2/SIGDET2	R/W	GPIO2/SIGDET2 input	0
1	GPIO1/SIGDET1	R/W	GPIO1/SIGDET1 input	0
0	GPIO0/SIGDET0	R/W	GPIO0/SIGDET0 input	0

4.7.5 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 137 • GPIO Output, Address 16G (0x10)

Bit	Name	Access	Description	Default
15:14	Reserved	RO	Reserved	
13	GPIO13/1588_SPI_DO	R/W	GPIO13/1588_SPI_DO output	0
12	GPIO12/1588_SPI_CS	R/W	GPIO12/1588_SPI_CS output	0
11	GPIO11/1588_PPS_0	R/W	GPIO11/1588_PPS_0 output	0
10	GPIO10/1588_LOAD_SAVE	R/W	GPIO10/1588_LOAD_SAVE output	0
9	GPIO9/FASTLINK_FAIL	R/W	GPIO9/FASTLINK_FAIL output	0
8	GPIO8/I2C_SDA	R/W	GPIO8/I2C_SDA output	0
7	GPIO7/I2C_SCL_3	R/W	GPIO7/I2C_SCL_3 output	0
6	GPIO6/I2C_SCL_2	R/W	GPIO6/I2C_SCL_2 output	0
5	GPIO5/I2C_SCL_1	R/W	GPIO5/I2C_SCL_1 output	0
4	GPIO4/I2C_SCL_0	R/W	GPIO4/I2C_SCL_0 output	0
3	GPIO3/SIGDET3	R/W	GPIO3/SIGDET3 output	0
2	GPIO2/SIGDET2	R/W	GPIO2/SIGDET2 output	0
1	GPIO1/SIGDET1	R/W	GPIO1/SIGDET1 output	0
0	GPIO0/SIGDET0	R/W	GPIO0/SIGDET0 output	0

4.7.6 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 138 • GPIO Input/Output Configuration, Address 17G (0x11)

Bit	Name	Access	Description	Default
15:14	Reserved	RO	Reserved	
13	GPIO13/1588_SPI_DO	R/W	GPIO13/1588_SPI_DO output enable	0

Table 138 • GPIO Input/Output Configuration, Address 17G (0x11) (continued)

Bit	Name	Access	Description	Default
12	GPIO12/1588_SPI_CS	R/W	GPIO12/1588_SPI_CS output enable	0
11	GPIO11/1588_PPS_0	R/W	GPIO11/1588_PPS_0 output enable	0
10	GPIO10/1588_LOAD_SAVE	R/W	GPIO10/1588_LOAD_SAVE output enable	0
9	GPIO9/FASTLINK_FAIL	R/W	GPIO9/FASTLINK_FAIL output enable	0
8	GPIO8/I2C_SDA	R/W	GPIO8/I2C_SDA output enable	0
7	GPIO7/I2C_SCL_3	R/W	GPIO7/I2C_SCL_3 output enable	0
6	GPIO6/I2C_SCL_2	R/W	GPIO6/I2C_SCL_2 output enable	0
5	GPIO5/I2C_SCL_1	R/W	GPIO5/I2C_SCL_1 output enable	0
4	GPIO4/I2C_SCL_0	R/W	GPIO4/I2C_SCL_0 output enable	0
3	GPIO3/SIGDET3	R/W	GPIO3/SIGDET3 output enable	0
2	GPIO2/SIGDET2	R/W	GPIO2/SIGDET2 output enable	0
1	GPIO1/SIGDET1	R/W	GPIO1/SIGDET1 output enable	0
0	GPIO0/SIGDET0	R/W	GPIO0/SIGDET0 output	0

4.7.7 Microprocessor Command

Register 18G is a command register. Bit 15 tells the internal processor to execute the command. When bit 15 is cleared the command has completed. Software needs to wait until bit 15 = 0 before proceeding with the next PHY register access. Bit 14 = 1 typically indicates an error condition where the squelch patch was not loaded. Use the following steps to execute the command:

1. Write desired command
2. Check bit 15 (move existing text)
3. Check bit 14 (if set, then error)

Commands may take up to 25 ms to complete before bit 15 changes to 0.

Note: All MAC interfaces must be the same — all QSGMII or SGMII.

Table 139 • Microprocessor Command Register, Address 18G

Command	Setting
Enable four MAC SGMII ports	0x80F0
Enable four MAC QSGMII ports	0x80E0
Enable four Media 1000BASE-X ports	0x8FC1 ¹
Enable four Media 100BASE-FX ports	0x8FD1 ¹

1. The “F” in the command has a bit representing each of the four PHYs. To exclude a PHY from the configuration, set its bit to 0. For example, the configuration of PHY 3 and PHY 2 to 1000BASE-X would be 1100 or a “C” and the command would be 0x8CC1.

4.7.8 MAC Configuration and Fast Link

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 140 • MAC Configuration and Fast Link Register, Address 19G (0x13)

Bit	Name	Access	Description	Default
15:14	MAC configuration	R/W	Select MAC interface mode 00: SGMII 01: QSGMII 10: Reserved 11: Reserved	00
13:4	Reserved	RO	Reserved	
3:0	Fast link failure port setting	R/W	Select fast link failure PHY source 0000: Port0 0001: Port1 0010: Port2 0011: Port3 1100–1111: Output disabled	0xF

4.7.9 Two-Wire Serial MUX Control 1

The following table shows the settings available to control the integrated two-wire serial MUX.

Table 141 • Two-Wire Serial MUX Control 1, Address 20G (0x14)

Bit	Name	Access	Description	Default
15:9	Two-wire serial 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	Two-wire serial SCL clock frequency	R/W	00: 50 kHz 01: 100 kHz 10: 400 kHz 11: 2 MHz	01
3	Two-wire serial MUX port 3 enable	R/W	1: Enabled. 0: Two-wire serial disabled. Becomes GPIO pin.	0
2	Two-wire serial MUX port 2 enable	R/W	1: Enabled. 0: Two-wire serial disabled. Becomes GPIO pin.	0
1	Two-wire serial MUX port 1 enable	R/W	1: Enabled. 0: Two-wire serial disabled. Becomes GPIO pin.	0
0	Two-wire serial MUX port 0 enable	R/W	1: Enabled. 0: Two-wire serial disabled. Two-wire serial MUX port 0 becomes GPIO pin if serial LED function is enabled, regardless of the settings of this bit.	0

4.7.10 Two-Wire Serial MUX Control 2

Register 21G is used to control the two-wire serial MUX for status and control of two-wire serial slave devices.

Table 142 • Two-Wire Serial MUX Interface Status and Control, Address 21G (0x15)

Bit	Name	Access	Description	Default
15	Two-wire serial MUX ready	RO	1: Two-wire serial MUX is ready for read or write	
14:12	Reserved	RO	Reserved	
11:10	PHY port Address	R/W	Specific PHY port being addressed.	00
9	Enable two-wire serial MUX access	R/W	Self-clearing bit. 1: Execute read or write through the two-wire serial MUX based on the settings of register bit 21G.8	0
8	Two-wire serial MUX read or write	R/W	1: Read from two-wire serial MUX 0: Write to two-wire serial MUX	1
7:0	Two-wire serial MUX address	R/W	Sets the address of the two-wire serial MUX used to direct read or write operations.	0x00

4.7.11 Two-Wire Serial MUX Data Read/Write

Register 22G in the extended register space enables access to the two-wire serial MUX.

Table 143 • Two-Wire Serial MUX Data Read/Write, Address 22G (0x16)

Bit	Name	Access	Description	Default
15:8	Two-wire serial MUX read data	RO	Eight-bit data read from two-wire serial MUX; requires setting both register 21G.9 and 21G.8 to 1.	
7:0	Two-wire serial MUX write data	R/W	Eight-bit data to be written to two-wire serial MUX.	0x00

4.7.12 Recovered Clock 1 Control

Register 23G in the extended register space controls the functionality of the recovered clock 1 output signal.

Table 144 • Recovered Clock 1 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

Table 144 • Recovered Clock 1 Control, Address 23G (0x17) (continued)

Bit	Name	Access	Description	Default
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. When the CLK_SQUELCH_IN pin is set high, it squelches the recovered clocks regardless of bit settings.	
3	Reserved	RO	Reserved.	
2:0	Clock selection for specified PHY	R/W	000: Serial media recovered clock 001: Copper PHY recovered clock 010: Copper PHY transmitter TCLK 011–111: Reserved	000

4.7.13 Recovered Clock 2 Control

Register 24G in the extended register space controls the functionality of the recovered clock 2 output signal.

Table 145 • Recovered Clock 2 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

Table 145 • Recovered Clock 2 Control, Address 24G (0x18) (continued)

Bit	Name	Access	Description	Default
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. When the CLK_SQUELCH_IN pin is set high, it squelches the recovered clocks regardless of bit settings.	
3	Reserved	RO	Reserved	
2:0	Clock selection for specified PHY	R/W	000: Serial media recovered clock 001: Copper PHY recovered clock 010–111: Reserved	000

4.7.14 Enhanced LED Control

The following table contains the bits to control advanced functionality of the parallel and serial LED signals.

Table 146 • 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	Port 1 enhanced serial LED output enable	R/W	Enable the enhanced serial LED output functionality for port 1 LED pins. 1: Enhanced serial LED outputs 0: Normal function	0

Table 146 • Enhanced LED Control, Address 25G (0x19) (continued)

Bit	Name	Access	Description	Default
6	Port 0 enhanced serial LED output enable	R/W	Enable the enhanced serial LED output functionality for port 0 LED pins. 1: Enhanced serial LED outputs 0: Normal function	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 52 , page 125.	
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 126.	

4.7.15 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 147 • Global Interrupt Status, Address 29G (0x1D)

Bit	Name	Access	Description
15:12	Reserved	RO	Reserved
11	PHY3 1588 ¹	RO	PHY3 1588 interrupt source indication 0: PHY3 1588 caused the interrupt 1: PHY3 1588 did not cause the interrupt
10	PHY2 1588 ¹	RO	PHY 2 1588 interrupt source indication 0: PHY2 1588 caused the interrupt 1: PHY2 1588 did not cause the interrupt
9	PHY1 1588 ¹	RO	PHY 1 1588 interrupt source indication 0: PHY1 1588 caused the interrupt 1: PHY1 1588 did not cause the interrupt
8	PHY0 1588 ¹	RO	PHY 0 1588 interrupt source indication 0: PHY0 1588 caused the interrupt 1: PHY0 1588 did not cause the interrupt
7:4	Reserved	R	Reserved
3	PHY3 interrupt source ²	RO	PHY3 interrupt source indication 0: PHY3 caused the interrupt 1: PHY3 did not cause the interrupt
2	PHY2 interrupt source ²	RO	PHY2 interrupt source indication 0: PHY2 caused the interrupt 1: PHY2 did not cause the interrupt

Table 147 • Global Interrupt Status, Address 29G (0x1D) (continued)

Bit	Name	Access	Description
1	PHY1 interrupt source ²	RO	PHY1 interrupt source indication 0: PHY1 caused the interrupt 1: PHY1 did not cause the interrupt
0	PHY0 interrupt source ²	RO	PHY0 interrupt source indication 0: PHY0 caused the interrupt 1: PHY0 did not cause the interrupt

1. This bit is set to 0 when the corresponding PHY's 1588 interrupt is asserted and is set to 1 when the corresponding PHY's 1588 interrupt is cleared.
2. This bit is set to 1 when the corresponding PHY's Interrupt Status register 26 (0x1A) is read.

4.8 Clause 45 Registers to Support Energy Efficient Ethernet and 802.3bf

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. When the link is down, 0 is the value returned for the x.180x addresses.

Table 148 • Clause 45 Registers Page Space

Address	Name
1.1	PMA/PMD status 1
1.1800	TimeSync PMA/PMD capability
1.1801	Tx maximum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1803	Tx minimum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1805	Rx maximum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1807	Rx minimum delay through PHY (PMA/PMD/PCS, until MACsec block)
3.1	PCS status 1
3.1800	TimeSync PCS capability
3.1801	Tx maximum delay through 1588 and MACsec
3.1803	Tx minimum delay through 1588 and MACsec
3.1805	Rx maximum delay through 1588 and MACsec
3.1807	Rx minimum delay through 1588 and MACsec
3.20	EEE capability
3.22	EEE wake error counter
4.1800	TimeSync PHY XS Capability
4.1801	Tx maximum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1803	Tx minimum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1805	Rx maximum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1807	Rx minimum delay through xMII (SGMII, QSGMII, including FIFO variations)
7.60	EEE advertisement
7.61	EEE link partner advertisement

4.8.1 PMA/PMD Status 1

The following table shows the bit descriptions for the PMA/PMD Status 1 register.

Table 149 • PMA/PMD Status 1

Bit	Name	Access	Description
15:3	Reserved	RO	Reserved
2	PMD/PMA receive link status	RO/LL	1: PMA/PMD receive link up 0: PMA/PMD receive link down
1:0	Reserved	RO	Reserved

4.8.2 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 150 • 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/LL	1: PCS receive link up 0: PCS receive link down
1:0	Reserved	RO	Reserved

4.8.3 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 151 • 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

4.8.4 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 152 • 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

4.8.5 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 153 • 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	

4.8.6 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 154 • 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

The following table shows the bit assignments for the 802.3bf registers. When the link is down, 0 is the value returned. Register 1.1801 would be device address of 1 and register address of 1801.

Table 155 • 802.3bf Registers

Register	Name	Function
1.1800	PMA/PMD Time Sync capable	Bit 1: PMA/PMD Time Sync Tx capable Bit 0: PMA/PMD Time Sync Rx capable
1.1801	PMA/PMD delay Tx max	Tx maximum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1803	PMA/PMD delay Tx min	Tx minimum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1805	PMA/PMD delay Rx max	Rx maximum delay through PHY (PMA/PMD/PCS, until MACsec block)
1.1807	PMA/PMD delay Rx min	Rx minimum delay through PHY (PMA/PMD/PCS, until MACsec block)
3.1800	PCS Time Sync capable	Bit 1: PCS Time Sync Tx capable bit 0: PCS Time Sync Rx capable
3.1801	PCS delay Tx max low	Tx maximum delay through 1588 and MACsec lower bits
3.1802	PCS delay Tx max high	Tx maximum delay through 1588 and MACsec upper bits
3.1803	PCS delay Tx min low	Tx minimum delay through 1588 and MACsec lower bits
3.1804	PCS delay Tx min high	Tx minimum delay through 1588 and MACsec upper bits
3.1805	PCS delay Rx max low	Rx maximum delay through 1588 and MACsec lower bits
3.1806	PCS delay Rx max high	Rx maximum delay through 1588 and MACsec upper bits
3.1807	PCS delay Rx min low	Rx minimum delay through 1588 and MACsec lower bits
3.1808	PCS delay Rx min high	Rx minimum delay through 1588 and MACsec upper bits
4.1800	PHY XS Time Sync capable	Bit 1: PHY XS Time Sync Tx capable Bit 0: PHY XS Time Sync Rx capable
4.1801	PHY XS delay Tx max	Tx maximum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1803	PHY XS delay Tx min	Tx minimum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1805	PHY XS delay Rx max	Rx maximum delay through xMII (SGMII, QSGMII, including FIFO variations)
4.1807	PHY XS delay Rx min	Rx minimum delay through xMII (SGMII, QSGMII, including FIFO variations)

5 Electrical Specifications

This section provides the DC characteristics, AC characteristics, recommended operating conditions, and stress ratings for the VSC8584-11 device.

5.1 DC Characteristics

This section contains the DC specifications for the VSC8584-11 device.

5.1.1 VDD25 and VDDMDIO (2.5 V)

The following table shows the DC specifications for the pins referenced to V_{VDD25} and $V_{VDDMDIO}$ when it is set to 2.5 V. The specifications listed in the following table are valid only when $V_{VDD1} = 1.0$ V, $V_{VDD1A} = 1.0$ V, and $V_{VDD25A} = 2.5$ V.

Table 156 • VDD25 and VDDMDIO

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, LVTTTL	V_{OH_TTL}	2.0	2.8	V	$I_{OH} = -1$ mA
Output high voltage, open drain	V_{OH_OD}	2.0	2.8	V	$I_{OH} = -100$ μ A
Output low voltage	V_{OL}	-0.3	0.4	V	$I_{OL} = 4$ mA
Input high voltage	V_{IH}	1.85	3.3	V	Except SMI pins
Input high voltage	V_{IH}	1.88	3.3	V	SMI pins
Input low voltage	V_{IL}	-0.3	0.7	V	
Input leakage current	I_{ILEAK}	-32	32	μ A	Internal resistor included (except GPIO, LED, and COMA_MODE)
Input leakage current	I_{ILEAK}	-76	32	μ A	Internal resistor included (GPIO, LED, and COMA_MODE)
Output leakage current	I_{OLEAK}	-32	32	μ A	Internal resistor included (except GPIO, LED, and COMA_MODE)
Output leakage current	I_{OLEAK}	-76	32	μ A	Internal resistor included (GPIO, LED, and COMA_MODE)

5.1.2 VDDMDIO (1.2 V)

The following table shows the DC specifications for the pins referenced to $V_{VDDMDIO}$ when it is set to 1.2 V. The specifications listed in the following table are valid only when $V_{VDD1} = 1.0$ V, $V_{VDD1A} = 1.0$ V, $V_{VDD25} = 2.5$ V, $V_{VDD25A} = 2.5$ V, and $V_{VDDMDIO} = 1.2$ V.

Table 157 • VDDMDIO

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, open drain	V_{OH}	1.0	1.5	V	$I_{OH} = -100$ μ A
Output low voltage, open drain	V_{OL}	-0.3	0.25	V	$I_{OL} = 4$ mA
Input high voltage	V_{IH}	0.9	1.5	V	
Input low voltage	V_{IL}	-0.3	0.36	V	

Table 157 • VDDMDIO (continued)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Input leakage current	I_{LEAK}	-32	32	μA	Internal resistor included
Output leakage current	I_{OLEAK}	-32	32	μA	Internal resistor included

5.1.3 Supply Voltage

The following table shows the supply voltage specifications.

Table 158 • Supply Voltage Specifications

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Power supply voltage for digital logic	V_{DD1}	0.95	1	1.05	V
Power supply voltage for analog logic	V_{DD1A}	0.95	1	1.05	V
Power supply voltage for digital supply	V_{DD25}	2.375	2.5	2.625	V
Power supply voltage for analog supply	V_{DD25A}	2.375	2.5	2.625	V
1.2 V MDIO internal supply	V_{DDMDIO}	1.19	1.2	2.625	V
2.5 V MDIO internal supply	V_{DDMDIO}	1.19	2.5	2.625	V

5.1.4 LED and GPIO

The following table shows the DC specifications for the LED and GPIO pins.

Table 159 • LED and GPIO

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage for LED pins, LVTTTL	V_{OH}	1.7	2.8	V	$V_{VDD25} = 2.5\text{ V}$ $I_{OH} = -24\text{ mA}$
Output low voltage for LED pins, LVTTTL	V_{OL}	-0.3	0.6	V	$V_{VDD25} = 2.5\text{ V}$ $I_{OL} = 24\text{ mA}$
Output high voltage for GPIO pins, LVTTTL	V_{OH}	1.7	2.8	V	$V_{VDD25} = 2.5\text{ V}$ $I_{OH} = -12\text{ mA}$
Output low voltage for GPIO pins, LVTTTL	V_{OL}	-0.3	0.6	V	$V_{VDD25} = 2.5\text{ V}$ $I_{OL} = 12\text{ mA}$

5.1.5 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 218.

All internal pull-up resistors are connected to their respective I/O supply.

Table 160 • Internal Pull-Up or Pull-Down Resistors

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Internal pull-up resistor (GPIO, LED, and COMA_MODE)	R_{PU1}	33	53	90	$\text{k}\Omega$
Internal pull-up resistor, all others	R_{PU2}	96	120	144	$\text{k}\Omega$
Internal pull-down resistor	R_{PD}	96	120	144	$\text{k}\Omega$

5.1.6 Reference Clock

The following table shows the DC specifications for a differential reference clock input signal.

Table 161 • Reference Clock DC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Input voltage range	V_{IP}, V_{IN}	-25		1260	mV
Input differential peak-to-peak voltage	$ V_{ID} $	150 ¹		1200	mV
Input common-mode voltage	V_{ICM}	0		1200 ²	mV
Differential input impedance	R_I		100		Ω

1. To meet jitter specifications, the minimum $|V_{ID}|$ must be 400 mV. When using a single-ended clock input, the REFCLK_P low voltage must be less than $V_{DDA} - 200$ mV, and the high voltage level must be greater than $V_{DDA} + 200$ mV.
2. The maximum common-mode voltage is provided without a differential signal. The common-mode voltage is only limited by the maximum and minimum input voltage range and by the differential amplitude of the input signal.

5.1.7 1588 Reference Clock

The following table shows the DC specifications for a differential 1588 reference clock input signal.

Table 162 • 1588 Reference Clock DC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Input voltage range	V_{IP}, V_{IN}	-25		1260	mV
Input differential peak-to-peak voltage	$ V_{ID} $	150		1200	mV
Input common-mode voltage	V_{ICM}	0		1200 ¹	mV
Differential input impedance	R_I		100		Ω

1. The maximum common-mode voltage is provided without a differential signal. The common-mode voltage is only limited by the maximum and minimum input voltage range and by the differential amplitude of the input signal.

5.1.8 SerDes Interface (SGMII)

The SerDes output drivers are designed to operate in SGMII/LVDS mode. 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 163 • SerDes Driver DC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Output high voltage, V_{OA} or V_{OB}	V_{OH}		1050	mV	$R_L = 100 \Omega \pm 1\%$
Output low voltage, V_{OA} or V_{OB}	V_{OL}	0		mV	$R_L = 100 \Omega \pm 1\%$
Output differential peak voltage	$ V_{OD} $	350	450	mV	$V_{DD_VS} = 1.0$ V $R_L = 100 \Omega \pm 1\%$
Output differential peak voltage, fiber media 1000BASE-X	$ V_{OD} $	350	450	mV	$V_{DD_VS} = 1.0$ V $R_L = 100 \Omega \pm 1\%$
Output offset voltage ¹	V_{OS}	420	580	mV	$V_{DD_VS} = 1.0$ V $R_L = 100 \Omega \pm 1\%$

Table 163 • SerDes Driver DC Specifications (continued)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
DC output impedance, differential	R_O	80	140	Ω	$V_C = 1.0\text{ V}$ See Figure 116, page 196
R_O mismatch between A and B, SGMII mode ²	ΔR_O		10	%	$V_C = 1.0\text{ V}$ See Figure 116, page 196
Change in $ V_{OD} $ between 0 and 1, SGMII mode	$\Delta V_{OD} $		25	mV	$R_L = 100\ \Omega \pm 1\%$
Change in V_{OS} between 0 and 1, SGMII mode	ΔV_{OS}		25	mV	$R_L = 100\ \Omega \pm 1\%$
Output current, driver shorted to GND, SGMII mode	$ I_{OSA} $, $ I_{OSB} $		40	mA	
Output current, drivers shorted together, SGMII mode	$ I_{OSAB} $		12	mA	

1. Requires AC-coupling for SGMII compliance.
2. Matching of reflection coefficients. For more information about test methods, see IEEE Std 1596.3-1996.

Figure 114 • SGMII DC Transmit Test Circuit






Figure 115 • SGMII DC Definitions



$$\Delta|V_{OD}| = | |V_{OA_H} - V_{OB_L}| - |V_{OB_H} - V_{OA_L}| |$$

$$\Delta V_{OS} = | \frac{1}{2}(V_{OA_H} + V_{OB_L}) - \frac{1}{2}(V_{OA_L} + V_{OB_H}) |$$

Figure 116 • SGMII DC Driver Output Impedance Test Circuit




The following table lists the DC specifications for the SGMII receivers.

Table 164 • 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 peak voltage	$ V_{ID} $	50	1000	mV	
Input common-mode voltage ¹	V_{ICM}	0	V_{DD_A} ²	mV	Without any differential signal
Receiver differential input impedance	R_I	80	120	Ω	
Input differential hysteresis, SGMII mode	V_{HYST}	25		mV	

1. SGMII compliancy requires external AC-coupling. When interfacing with specific Microsemi devices, DC-coupling is possible. For more information, contact your local Microsemi sales representative.
2. The common-mode voltage is only limited by the maximum and minimum input voltage range and the input signal's differential amplitude.

Figure 117 • SGMII DC Input Definitions



5.1.9 Enhanced SerDes Interface (QSGMII)

All DC specifications for the enhanced SerDes interface operating in QSGMII mode meet or exceed the requirements specified for CEI-6G-SR according to OIF-CEI-02.0 requirements where applicable.

The enhanced SerDes interface supports the following operating modes: SGMII, QSGMII, and SFP. The following table shows the DC specifications for the enhanced SerDes driver.

Table 165 • Enhanced SerDes Driver DC Specifications

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Signaling speed	T_{BAUD}	5.0 -100 ppm		5.0 + 100 ppm	Gbps	Signaling speed
Differential peak-to-peak output voltage	V_{OD}			30	mV	Tx disabled
Differential peak output voltage, SFP mode	$ V_{ODp} $	250		400	mV	$V_{DD_VS} = 1.0\text{ V}$ $R_L = 100\ \Omega \pm 1\%$ maximum drive
Differential peak output voltage, QSGMII mode	$ V_{ODp} $	400		900	mV	
Differential peak output voltage, SGMII mode ¹	$ V_{ODp} $	150		400	mV	$V_{DD_VS} = 1.0\text{ V}$ $R_L = 100\ \Omega \pm 1\%$
DC output impedance, differential	R_O	80	100	140	Ω	$V_C = 1.0\text{ V}$ See Figure 116, page 196

Table 165 • Enhanced SerDes Driver DC Specifications (continued)

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
R_O mismatch between A and B, SGMII mode ²	ΔR_O			10	%	$V_C = 1.0\text{ V}$ See Figure 116, page 196
Change in $ V_{OD} $ between 0 and 1, SGMII mode	$\Delta V_{OD} $			25	mV	$R_L = 100\ \Omega \pm 1\%$
Change in V_{OS} between 0 and 1, SGMII mode	$\Delta V_{OS} $			25	mV	$R_L = 100\ \Omega \pm 1\%$
Output current, drivers shorted to ground, SGMII and QSGMII modes	$ I_{OSA} $, $ I_{OSB} $			40	mA	
Output current, drivers shorted together, SGMII and QSGMII modes	$ I_{OSAB} $			12	mA	

1. Voltage is adjustable in 64 steps.
2. Matching of reflection coefficients. For more information about test methods, see IEEE Std 1596.3-1996.

The following table lists the DC specifications for the enhanced SerDes receiver.

Table 166 • Enhanced SerDes Receiver DC Specifications

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Input voltage range, V_{IA} or V_{IB}	V_I	-0.25		1.20	V	
Input differential peak-to-peak voltage	$ V_{ID} $	100		1600	mV	
Input common-mode voltage	V_{ICM}	$V_{DDA} - 100$	V_{DDA}	$V_{DDA} + 100$	mV	Load-type 2 (DC-coupled)
Receiver differential input impedance	R_I	80	100	120	Ω	

5.1.10 Current Consumption

The following table shows the estimated current consumption values for each mode, assuming the 1588 and MACsec functions are disabled. Add significant margin above the values for sizing power supplies. Add values from tables for the 1588 and MACsec blocks to calculate total typical and maximum current for each power supply with those functions enabled.

Table 167 • Current Consumption (1588 and MACsec Disabled)

Mode	Typical				Maximum				Unit	Condition
	1 V Digital	1 V Analog	2.5 V Digital	2.5 V Analog	1 V Digital	1 V Analog	2.5 V Digital	2.5 V Analog		
Reset	105	130	9	2	910	190	11	4	mA	
Power down	175	200	9	21	1015	260	11	23	mA	
1000BASE-T	460	230	10	445	1890	270	15	500	mA	4-port SGMII
100BASE-TX	255	215	10	290	1525	245	15	310	mA	4-port SGMII
10BASE-T	200	210	10	230	1420	240	15	245	mA	4-port SGMII
10BASE-Te	200	210	10	215	1420	240	15	210	mA	4-port SGMII

Table 167 • Current Consumption (1588 and MACsec Disabled) (continued)

Mode	Typical				Maximum				Unit	Condition
	1 V Digital	1 V Analog	2.5 V Digital	2.5 V Analog	1 V Digital	1 V Analog	2.5 V Digital	2.5 V Analog		
1000BASE-X	225	275	10	21	1430	300	15	25	mA	4-port SGMII
100BASE-FX	200	270	10	21	1400	290	15	25	mA	4-port SGMII
1000BASE-T	450	185	10	445	1885	225	15	515	mA	4-port QSGMII
100BASE-TX	250	165	9	290	1520	200	15	325	mA	4-port QSGMII
10BASE-T	200	165	10	230	1415	195	15	260	mA	4-port QSGMII
10BASE-Te	195	165	10	215	1415	195	15	225	mA	4-port QSGMII
1000BASE-X	225	230	11	21	1425	255	15	40	mA	4-port QSGMII
100BASE-FX	200	225	10	22	1395	245	15	40	mA	4-port QSGMII

The following tables show the 1588 and MACsec current consumption values for each mode.

Table 168 • 1588 Current Consumption

Mode	1 V Digital	Unit	Condition
1000BASE-T	80	mA	4-port SGMII/QSGMII
100BASE-TX	50	mA	4-port SGMII/QSGMII
10BASE-T	50	mA	4-port SGMII/QSGMII
10BASE-Te	50	mA	4-port SGMII/QSGMII
1000BASE-X	80	mA	4-port SGMII/QSGMII
100BASE-FX	35	mA	4-port SGMII/QSGMII

Table 169 • MACsec Current Consumption

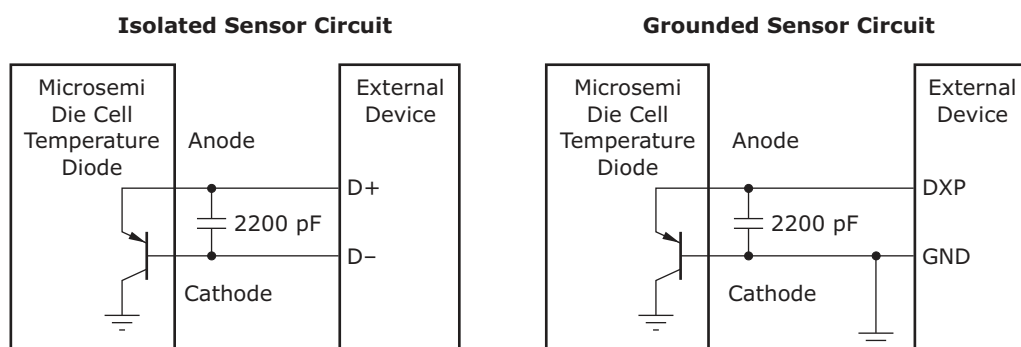
Mode	1 V Digital	Unit	Condition
1000BASE-T	160	mA	4-port SGMII/QSGMII
100BASE-TX	120	mA	4-port SGMII/QSGMII
10BASE-T	120	mA	4-port SGMII/QSGMII
10BASE-Te	120	mA	4-port SGMII/QSGMII
1000BASE-X	160	mA	4-port SGMII/QSGMII
100BASE-FX	120	mA	4-port SGMII/QSGMII

5.1.11 Thermal Diode

The VSC8584-11 device includes an on-die diode and internal circuitry for monitoring die temperature (junction temperature). The operation and accuracy of the diode is not guaranteed and should only be used as a reference.

The on-die thermal diode requires an external thermal sensor, located on the board or in a stand-alone measurement kit. Temperature measurement using a thermal diode is very sensitive to noise. The following illustration shows a generic application design.

Figure 118 • Thermal Diode



Note: Microsemi does not support or recommend operation of the thermal diode under reverse bias.

The following table provides the diode parameter and interface specifications with the pins connected internally to VSS in the device.

Table 170 • Thermal Diode Parameters

Parameter	Symbol	Typical	Maximum	Unit
Forward bias current	I_{FW}	See note ¹	1	mA
Diode ideality factor	n	1.008		

1. Typical value is device dependent.

The ideality factor, n , represents the deviation from ideal diode behavior as exemplified by the diode equation.

$$I_{FW} = I_S \times \left(e^{V_d \times \frac{q}{nkT}} - 1 \right)$$

where, I_S = saturation current, q = electron charge, V_d = voltage across the diode, k = Boltzmann Constant, and T = absolute temperature (Kelvin).

5.2 AC Characteristics

This section provides the AC specifications for the VSC8584-11 device.

5.2.1 Reference Clock

The following table shows the AC specifications for a 125 MHz differential reference clock source. Performance is guaranteed for 125 MHz differential clocks only; however, 125 MHz single-ended clocks are also supported for QSGMII interfaces.

25 MHz clock implementations are available but are limited to SGMII interfaces. For more information, contact your Microsemi representative.

Table 171 • Reference Clock AC Characteristics for QSGMII 125 MHz Differential Clock

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Reference clock frequency, REFCLK_SEL2 = 1	f		125.00		MHz	± 100 ppm Jitter < 1 ps RMS
Duty cycle	DC	40	50	60	%	
Rise time and fall time	t_R, t_F			1.5	ns	20% to 80% threshold


Table 171 • Reference Clock AC Characteristics for QSGMII 125 MHz Differential Clock (continued)

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
RefClk input RMS jitter requirement, bandwidth between 12 kHz and 500 kHz ¹				20	ps	Meets jitter generation of 1G output data per IEEE 802.3z
RefClk input RMS jitter requirement, bandwidth between 500 kHz and 15 MHz ¹				4	ps	Meets jitter generation of 1G output data per IEEE 802.3z
RefClk input RMS jitter requirement, bandwidth between 15 MHz and 40 MHz ¹				20	ps	Meets jitter generation of 1G output data per IEEE 802.3z
RefClk input RMS jitter requirement, bandwidth between 40 MHz and 80 MHz ¹				100	ps	Meets jitter generation of 1G output data per IEEE 802.3z
Jitter gain from RefClk to SerDes output, bandwidth between 0.1 MHz and 0.1 MHz				0.3	dB	
Jitter gain from RefClk to SerDes output, bandwidth between 0.1 MHz and 7 MHz			1	3	dB	
Jitter gain from RefClk to SerDes output, bandwidth above 7 MHz		1–20 × log (f/7 MHz)		3–20 × log (f/7 MHz)	dB	

1. Maximum RMS sinusoidal jitter allowed at the RefClk input when swept through the given bandwidth.

5.2.2 Recovered Clock

This section provides the AC characteristics for the recovered clock output signals. The following illustration shows the test circuit for the recovered clock output signals.

Figure 119 • Test Circuit for Recovered Clock Output Signals

The following table shows the AC specifications for the RCVRDCLK1 and RCVRDCLK2 outputs.

Table 172 • 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	

Table 172 • Recovered Clock AC Characteristics (continued)

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
Recovered clock frequency	f		25.00		MHz	
Recovered clock cycle time	t_{RCYC}		8.0		ns	
Recovered clock cycle time	t_{RCYC}		32.0		ns	
Recovered clock cycle time	t_{RCYC}		40.0		ns	
Duty cycle	DC	45	50	55	%	
Clock rise time and fall time	t_R, t_F			1.0	ns	20% to 80%
Peak-to-peak jitter, copper media interface, 1000BASE-T slave mode	JPP_{CLK_Cu}			400	ps	10k samples
Peak-to-peak jitter, fiber media interface, 100BASE-FX	JPP_{CLK_FIFX}			1.2	ns	10k samples
Peak-to-peak jitter, fiber media interface, 1000BASE-X	JPP_{CLK_FIX}			200	ps	10k samples

5.2.3 SerDes Outputs

The values listed in the following table are valid for all configurations, unless otherwise noted.

Table 173 • SerDes Outputs AC Specifications

Parameter	Symbol	Minimum	Maximum	Unit	Condition
V_{OD} ringing compared to V_S , SGMII mode	V_{RING}		± 10	%	$RL = 100 \Omega \pm 1\%$
V_{OD} rise time and fall time, SGMII mode	t_R, t_F	100	200	ps	20% to 80% of V_S $RL = 100 \Omega \pm 1\%$
Differential peak-to-peak output voltage	V_{OD}		30	mV	Tx disabled
Differential output return loss, 50 MHz to 625 MHz	R_{LO_DIFF}	≥ 10		dB	$RL = 100 \Omega \pm 1\%$
Differential output return loss, 625 MHz to 1250 MHz	R_{LO_DIFF}	$10 - 10 \times \log(f/625 \text{ MHz})$		dB	$RL = 100 \Omega \pm 1\%$
Common-mode return loss, 50 MHz to 625 MHz	RL_{OCM}	6		dB	
Interpair skew, SGMII mode	t_{SKEW}		20	ps	

5.2.4 SerDes Driver Jitter

The following table lists the jitter characteristics for the SerDes output driver.

Table 174 • SerDes Driver Jitter Characteristics

Parameter	Symbol	Maximum	Unit	Condition
Total jitter	TJ _O	192	ps	Measured according to IEEE 802.3.38.5
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 175 • SerDes Input AC Specifications

Parameter	Maximum	Unit	Condition
Differential input return loss, 50 MHz to 625 MHz	≥10	dB	RL = 100 Ω ±1%
Differential input return loss, 625 MHz to 1250 MHz	10–10 × log (f/625 MHz)	dB	RL = 100 Ω ±1%

5.2.6 SerDes Receiver Jitter Tolerance

The following table lists jitter tolerances for the SerDes receiver.

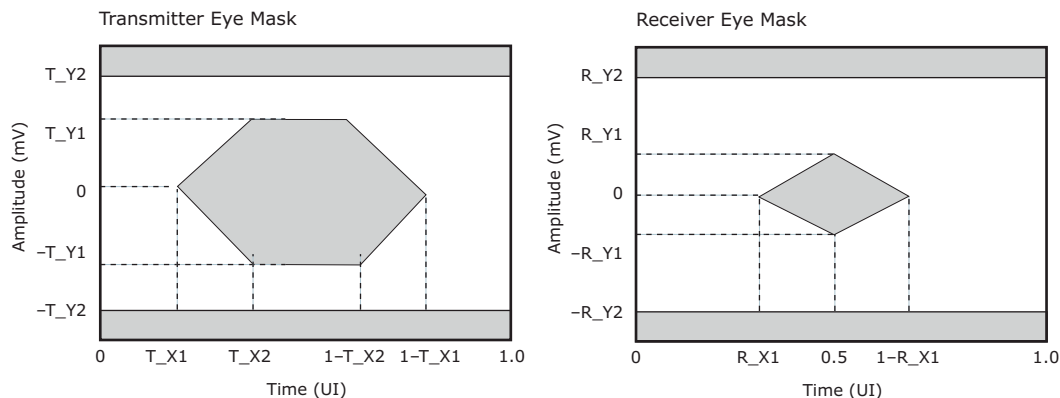
Table 176 • SerDes Receiver Jitter Tolerance

Parameter	Symbol	Minimum	Unit	Condition
Total jitter tolerance, greater than 637 kHz, SFP mode	TJT _I	600	ps	Measured according to IEEE 802.3 38.6.8
Deterministic jitter tolerance, greater than 637 kHz, SFP mode	DJT _I	370	ps	Measured according to IEEE 802.3 38.6.8
Cycle distortion jitter tolerance, 100BASE-FX mode	JT _{CD}	1.4	ns	Measured according to ISO/IEC 9314-3:1990
Data-dependent jitter tolerance, 100BASE-FX mode	DDJ	2.2	ns	Measured according to ISO/IEC 9314-3:1990
Random peak-to-peak jitter tolerance, 100BASE-FX mode	RJT	2.27	ns	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 the following modes of operation: SGMII, QSGMII, and SFP. The values in the tables in the following sections apply to the QSGMII modes listed in the condition column and are based on the test circuit shown in [Figure 114](#), page 196. 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 120 • QSGMII Transient Parameters

5.2.7.1 Enhanced SerDes Outputs

The following table provides the AC specifications for the enhanced SerDes outputs in SGMII mode.

Table 177 • Enhanced SerDes Outputs AC Specifications, SGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Unit interval, 1.25G mode	UI				800 ps
V_{OD} ringing compared to V_S	V_{RING}		± 10	%	$R_L = 100 \Omega \pm 1\%$
V_{OD} rise time and fall time	t_R, t_F	100	200	ps	20% to 80% of V_S $R_L = 100 \Omega \pm 1\%$
Differential output return loss, 50 MHz to 625 MHz	RL_{O_DIFF}	≥ 10		dB	$R_L = 100 \Omega \pm 1\%$
Differential output return loss, 625 MHz to 1250 MHz	RL_{O_DIFF}	$10 - 10 \times \log(f/625 \text{ MHz})$		dB	$R_L = 100 \Omega \pm 1\%$
Common-mode return loss, 50 MHz to 625 MHz	RL_{OCM}	6		dB	
Intrapair skew	t_{SKEW}		20	ps	

The enhanced SerDes transmitter operating in QSGMII mode complies with the AC characteristics specified for CEI-6G-SR interfaces according to OIF-CEI-02.0 with some modifications as specified by Cisco's QSGMII specification.

Table 178 • Enhanced SerDes Outputs AC Specifications, QSGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Signaling speed	T_{BAUD}	5.0 – 100 ppm	5.0 + 100 ppm	Gbps	
Differential output return loss	RLO_{SDD22}	8		dB	100 MHz to 2.5 GHz $R_L = 100 \Omega \pm 1\%$
Differential output return loss	RLO_{SDD22}	$8 - 16.6 \times \log(f/2.5)$		dB	2.5 GHz to 5 GHz $R_L = 100 \Omega \pm 1\%$
Common-mode output return loss	RLO_{CM}	6		dB	100 MHz to 2.5 GHz $R_L = 100 \Omega \pm 1\%$
Transition time	t_{TR}, t_{TF}	30		ps	20% to 80%
Random jitter	RJ		0.15	$UI_{P,P}$	
Deterministic jitter	DJ		0.15	$UI_{P,P}$	Measured according to OIF-CEI-02.0/CEI-6G-SR.

Table 178 • Enhanced SerDes Outputs AC Specifications, QSGMII Mode (continued)

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Duty cycle of distortion (part of DJ)	DCD		0.05	UI _{P,P}	
Total jitter	TJ		0.30	UI _{P,P}	Measured according to OIF-CEI-02.0/CEI-6G-SR.
Eye mask X1	X1		0.15	UI _{P,P}	Near-end
Eye mask X2	X2		0.40	UI _{P,P}	Near-end
Eye mask Y1	Y1	200		mV	Near-end
Eye mask Y2	Y2		450	mV	Near-end

5.2.7.2 Enhanced SerDes Inputs

The enhanced SerDes operating in QSGMII mode complies to the AC characteristics as specified for CEI-6G-SR interfaces according to OIF-CEI-02.0 with some modifications as specified by Cisco's QSGMII specification. The following table lists the AC specifications for the enhanced SerDes inputs in SGMII mode.

Table 179 • Enhanced SerDes Input AC Specifications, SGMII Mode

Parameter	Symbol	Minimum	Unit	Condition
Unit interval, 1.25G	UI		ps	800 ps
Differential input return loss, 50 MHz to 625 MHz	RL _{I,DIFF}	10	dB	R _L = 100 Ω ±1%
Common-mode input return loss, 50 MHz to 625 MHz	RL _{ICM}	6	dB	

The following table lists the AC specifications for the enhanced SerDes inputs in QSGMII mode.

Table 180 • Enhanced SerDes Inputs AC Specifications, QSGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Signaling speed	T _{BAUD}	5.0 – 100 ppm	5.0 + 100 ppm	Gbps	
Differential input return loss	RL _{I,SDD11}	8		dB	100 MHz to 2.5 GHz
Differential input return loss	RL _{I,SDD11}	8-16.6 x log(f/2.5)		dB	2.5 GHz to 5 GHz
Common-mode input return loss	RL _{SCC11}	6		dB	100 MHz to 2.5 GHz
Bounded high-probability jitter	RBHPJ		0.45	UI _{P,P}	Uncorrelated bounded high-probability jitter (0.15 UI) + correlated bounded high-probability jitter (0.3 UI)
Sinusoidal jitter max	SJ _{MAX}		5	UI _{P,P}	
Sinusoidal jitter, HF	SJ _{HF}		0.05	UI _{P,P}	
Total jitter	TJ		0.60	UI _{P,P}	Does not include sinusoidal jitter, link operates at BER of 10 ⁻¹⁵
Eye mask X1	R_X1		0.30	UI _{P,P}	
Eye mask Y1	R_Y1		50	mV	

Table 180 • Enhanced SerDes Inputs AC Specifications, QSGMII Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Eye mask Y2	R_Y2		450	mV	

5.2.7.3 Enhanced SerDes Receiver Jitter Tolerance

The following table lists the jitter tolerance for the enhanced SerDes receiver in QSGMII mode.

Table 181 • Enhanced SerDes Receiver Jitter Tolerance

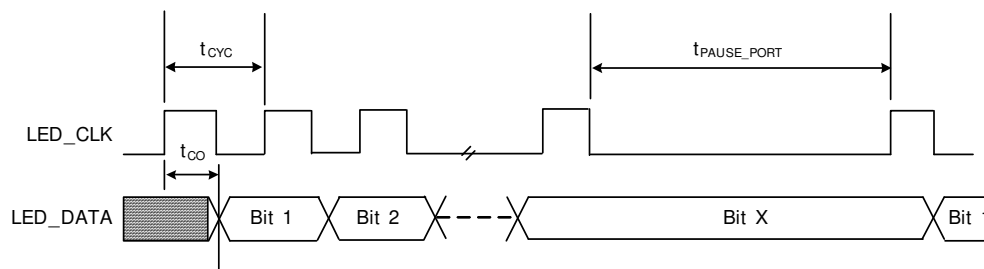
Parameter	Symbol	Minimum	Unit	Condition
Total jitter tolerance, greater than 637 kHz, SFP mode	TJT _I	600	ps	Measured according to IEEE 802.3 38.6.8
Deterministic jitter tolerance, greater than 637 kHz, SFP mode	DJT _I	370	ps	Measured according to IEEE 802.3 38.6.8
Cycle distortion jitter tolerance, 100BASE-FX mode	JT _{CD}	1.4	ns	Measured according to ISO/IEC 9314-3:1990
Data-dependent jitter tolerance, 100BASE-FX mode	DDJ	2.2	ns	Measured according to ISO/IEC 9314-3:1990
Random peak-to-peak jitter tolerance, 100BASE-FX mode	RJT	2.27	ns	Measured according to ISO/IEC 9314-3:1990

5.2.8 Basic Serial LEDs

This section contains the AC specifications for the basic serial LEDs.

Table 182 • Basic Serial LEDs AC Characteristics

Parameter	Symbol	Typical	Unit
LED_CLK cycle time	t _{CYC}	1024	ns
Pause between LED port sequences	t _{PAUSE_PORT}	3072	ns
Pause between LED bit sequences	t _{PAUSE_BIT}	25.541632	ms
LED_CLK to LED_DATA	t _{CO}	1	ns

Figure 121 • Basic Serial LED Timing


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 183 • Enhanced Serial LEDs AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit
LED_CLK cycle time	t_{CYC}		256		ns
Pause between LED_DATA bit sequences	t_{PAUSE}	0.396		24.996	ms
LED_CLK to LED_DATA	t_{CO}		127		ns
LED_CLK to LED_LD	t_{CL}		256		ns
LED_LD pulse width	t_{LW}		128		ns
LED_PULSE cycle time	t_{PULSE}	199		201	μ s

Figure 122 • Enhanced Serial LED Timing



5.2.10 Serial CPU Interface (SI) for Slave Mode

All serial CPU interface (SI) slave mode timing requirements are specified relative to the input low and input high threshold levels. The following illustrations show the timing parameters and measurement points for SI input and output data.

Figure 123 • SI Input Data Timing Diagram for Slave Mode





Figure 124 • SI Output Data Timing Diagram for Slave Mode



All SI signals comply with the specifications shown in the following table. The SI input timing requirements are requested at the pins of the device.


Table 184 • SI Timing Specifications for Slave Mode

Parameter	Symbol	Minimum	Maximum	Unit	Condition
Clock frequency	f		25	MHz	
Clock cycle time	t_C	40		ns	
Clock time high	$t_{W(CH)}$	16		ns	
Clock time low	$t_{W(CL)}$	16		ns	
Clock rise time and fall time	t_R, t_F		10	ns	Between $V_{IL(MAX)}$ and $V_{IH(MIN)}$
DI setup time to clock	$t_{SU(DI)}$	4		ns	
DI hold time from clock	$t_{H(DI)}$	4		ns	
Enable active before first clock	t_{LEAD}	10		ns	
Enable inactive after clock (input cycle) ¹	t_{LAG1}	25		ns	
Enable inactive after clock (output cycle)	t_{LAG2}	See note ⁽²⁾		ns	
Enable inactive width	$t_{W(EH)}$	20		ns	
DO valid after clock	$t_{V(C)}$		20	ns	$C_L = 30 \text{ pF}$
DO hold time from clock	$t_{H(DO)}$	0		ns	$C_L = 0 \text{ pF}$
DO disable time ³	t_{DIS}		15	ns	See the following illustration

- t_{LAG1} is defined only for write operations to the device, not for read operations.
- The last rising edge on the clock is necessary for the external master to read in the data. The lag time depends on the necessary hold time on the external master data input.
- Pin begins to float when a 300 mV change from the loaded V_{OH} or V_{OL} level occurs.

The following illustration shows the test circuit for the SI_DO disable time.

Figure 125 • Test Circuit for SI_DO Disable



5.2.11 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 185 • 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	$C_L = 10$ pF
TDO hold time from TCK falling	$t_{H(TDO)}$	0		ns	$C_L = 0$ pF
TDO disable time ¹	t_{DIS}		30	ns	
TRST time low	$t_{W(TL)}$	30		ns	

1. The pin begins to float when a 300 mV change from the actual V_{OH}/V_{OL} level occurs.

Figure 126 • JTAG Interface Timing Diagram





Figure 127 • Test Circuit for TDO Disable Time



5.2.12 Serial Management Interface

This section contains the AC specifications for the serial management interface (SMI).

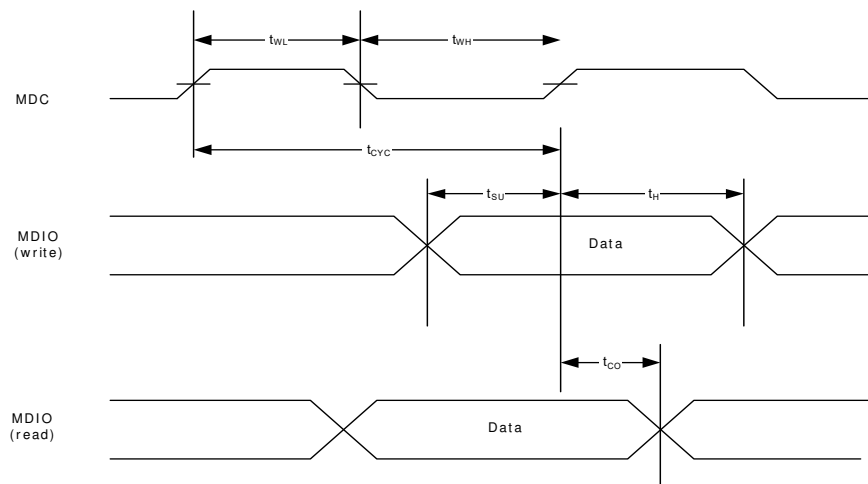
Table 186 • Serial Management Interface AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
MDC frequency ¹	f_{CLK}		2.5	12.5	MHz	
MDC cycle time	t_{CYC}	80	400		ns	
MDC time high	t_{WH}	20	50		ns	
MDC time low	t_{WL}	20	50		ns	
Setup to MDC rising	t_{SU}	10			ns	
Hold from MDC rising	t_H	10			ns	
MDC rise time	t_R			100 $t_{CYC} \times 10\%^1$	ns	MDC = 0: 1 MHz MDC = 1: MHz – f_{CLK} maximum
MDC fall time	t_F			100 $t_{CYC} \times 10\%^1$	ns	

Table 186 • Serial Management Interface AC Characteristics (continued)

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
MDC to MDIO valid	t_{CO}		10	300	ns	Time-dependent on the value of the external pull-up resistor on the MDIO pin

- For f_{CLK} above 1 MHz, the minimum rise time and fall time is in relation to the frequency of the MDC clock period. For example, if f_{CLK} is 2 MHz, the minimum clock rise time and fall time is 50 ns.

Figure 128 • Serial Management Interface Timing

5.2.13 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 187 • Reset 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}		105	ms
NRESET pulse width	$t_{W(RL)}$	100		ns
Wait time between NRESET de-assert and access of the SMI interface	t_{WAIT}	105		ms

5.2.14 IEEE 1588 Timing Specifications

This section contains the AC specifications for the IEEE 1588 clock pins.

Table 188 • IEEE 1588 Timing Specifications AC Characteristics

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
1588 reference clock frequency ¹	f		125 156.25 200 250		MHz	± 100 ppm Jitter < 10 ps RMS
Duty cycle	DC	40	50	60	%	
Rise time and fall time	t_R, t_F			1.5	ns	20% to 80% threshold

1. Only the listed nominal frequency values are supported.

5.2.15 Serial Timestamp Interface

This section contains information about the AC specifications for the SPI interface.


Table 189 • SPI Timing

Parameter	Symbol	Minimum	Typical	Maximum	Unit	Condition
SPI_CLK frequency	f			62.5	MHz	
SPI_CLK duty cycle	t_C	40		60	%	
SPI_DO clock-to-Q timing	$t_{CLK-to-Q}$	-5		3	ns	
SPI_CS clock-to-Q timing	$t_{CLK-to-Q}$	-5		3	ns	

The following illustration shows the SPI interface timing.


Note: Data changes state on a falling SPI_CLK edge in the default configuration. SPI_CLK can be inverted by setting the 1588 register bit TS_FIFO_SI_CFG:SI_CLK_PHA.

Figure 129 • SPI Interface Timing



5.2.16 Local Time Counter Load/Save Timing


This section contains information about the AC specifications for the local time counter load/save signal.

Figure 130 • Local Time Counter Load/Save Timing Diagram**Table 190 • Local Time Counter Load/Save Timing Specifications**

Parameter	Symbol	Minimum	Maximum	Unit
Clock frequency	f		250	MHz
Clock cycle time	t_C	4		ns
DI setup time to clock	$t_{SU(DI)}$	0.6		ns
DI hold time from clock	$t_{H(DI)}$	3.3		ns

5.2.17 Daisy-Chained SPI Timestamping Inputs

This section contains information about the AC specifications for the daisy-chained SPI timestamping interface.

Figure 131 • Daisy-Chained SPI Timestamping Input Timing Diagram**Table 191 • Daisy-Chained SPI Timestamping Input Timing Specifications**

Parameter	Symbol	Minimum	Maximum	Unit
Clock frequency	f		62.5	MHz
Clock cycle time	t_C	16		ns
DI setup time to clock	$t_{SU(DI)}$	3		ns
DI hold time from clock	$t_{H(DI)}$	3		ns

The following table shows the PHY latency in IEEE 1588 bypass mode, measured between the media interface and SGMII MAC interface pins.

Table 192 • PHY Latency in IEEE 1588 Timing Bypass Mode

Mode	Transmit (egress)			Receive (ingress)			Unit
	Minimum	Typical	Maximum	Minimum	Typical	Maximum	

Table 192 • PHY Latency in IEEE 1588 Timing Bypass Mode (continued)

Mode	Transmit (egress)			Receive (ingress)			
	Min	Typ	Max	Min	Typ	Max	Unit
1000BASE-T	212 – 16	212	212 + 16	321 – 16	321	321 + 16	ns
100BASE-TX	581 – 80	581	581 + 80	485 – 40	485	485 + 40	ns
10BASE-T	4075 – 400	4075	4075 + 400	3375 – 200	3375	3375 + 200	ns
1000BASE-X	210 – 16	210	210 + 16	192 – 16	192	192 + 16	ns
100BASE-FX	531 – 40	531	531 + 40	430 – 16	430	430 + 16	ns

5.3 Operating Conditions

The following table shows the recommended operating conditions for the device.

Table 193 • Recommended Operating Conditions

Parameter	Symbol	Minimum	Typical	Maximum	Unit
Power supply voltage for core supply	V_{VDD1}	0.95	1.00	1.05	V
Power supply voltage for analog circuits	V_{VDD1A}	0.95	1.00	1.05	V
Power supply voltage for digital I/O	V_{VDD25}	2.38	2.50	2.62	V
Power supply voltage for analog circuits	V_{VDD25A}	2.38	2.50	2.62	V
2.5 V Power supply voltage for SMI	V_{VDD_MDIO}	2.38	2.50	2.62	V
1.2 V Power supply voltage for SMI	V_{VDD_MDIO}	1.14	1.2	1.26	V
VSC8584-11 operating temperature ¹	T	0		125	°C
VSC8584-14 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 VSC8584-11 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 194 • Stress Ratings

Parameter	Symbol	Minimum	Maximum	Unit
Power supply voltage for core supply	V_{VDD1}	–0.3	1.10	V
Power supply voltage for analog circuits	V_{VDD1A}	–0.3	1.10	V
Power supply voltage for analog circuits	V_{VDD25A}	–0.3	2.75	V
Power supply voltage for digital I/O	V_{VDD25}	–0.3	2.75	V
Power supply voltage for SMI	V_{VDD_MDIO}	–0.3	2.75	V
Input voltage for GPIO and logic input pins			3.3	V
Storage temperature	T_S	–55	125	°C
Electrostatic discharge voltage, charged device model, 1588_DIFF_INPUT_CLK_N and 1588_DIFF_INPUT_CLK_P pins	V_{ESD_CDM}	–200	200	V

Table 194 • Stress Ratings (continued)

Parameter	Symbol	Minimum	Maximum	Unit
Electrostatic discharge voltage, charged device model, all pins except the 1588_DIFF_INPUT_CLK_N pin and 1588_DIFF_INPUT_CLK_P pin	V_{ESD_CDM}	-250	250	V
Electrostatic discharge voltage, human body model, VDD_MDIO pin	V_{ESD_HBM}	-1000	1000	V
Electrostatic discharge voltage, human body model, all pins except the VDD_MDIO pin	V_{ESD_HBM}	See note ¹		V

1. This device has completed all required testing as specified in the JEDEC standard JESD22-A114, *Electrostatic Discharge (ESD) Sensitivity Testing Human Body Model (HBM)*, and complies with a Class 2 rating. The definition of Class 2 is any part that passes an ESD pulse of 2000 V, but fails an ESD pulse of 4000 V.

Warning This device can be damaged by electrostatic discharge (ESD) voltage. Microsemi 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 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 device. The following table provides notations for definitions of the various pin types.

Table 195 • Pin Type Symbol Definitions

Symbol	Pin Type	Description
3V	Power	3.3 V-tolerant pin.
ABIAS	Analog bias	Analog bias pin.
ADIFF	Analog differential	Analog differential signal pair.
I	Input	Input without on-chip pull-up or pull-down resistor.
I/O	Bidirectional	Bidirectional input or output signal.
NC	No connect	No connect pins must be left floating.
O	Output	Output signal.
OD	Open drain	Open drain output.
OS	Open source	Open source output.
PD	Pull-down	On-chip pull-down resistor to VSS.
PU	Pull-up	On-chip pull-up resistor.
ST	Schmitt-trigger	Input has Schmitt-trigger circuitry.

6.2 Pin Diagram

The following illustration shows the pin diagram for the device. For clarity, the device is shown in two halves, the top left and top right.

Figure 132 • Top-Left Pin Diagram

	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_PHY0	LED1_PHY0	LED2_PHY0	LED3_PHY0	VDD1_5	VSS_27	VSS_28	VSS_29
H	LED0_PHY1	LED1_PHY1	LED2_PHY1	LED3_PHY1	VDD1_7	VSS_33	VSS_34	VSS_35
J	LED0_PHY2	LED1_PHY2	LED2_PHY2	LED3_PHY2	VDD1_9	VSS_39	VSS_40	VSS_41
K	LED0_PHY3	LED1_PHY3	LED2_PHY3	LED3_PHY3	VDD1_11	VSS_45	VSS_46	VSS_47
L	RESERVED_5	<small>1588_PPS_31588_SPL_IN_D</small>	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	<small>1588_PPS_111588_SPL_IN_CLK</small>	<small>1588_PPS_211588_SPL_IN_CS</small>	VDD1_17	VSS_63	VSS_64	VSS_65
P	RESERVED_8	MDC	VDD_MDIO	RESERVED_4	VDD25A_8	VDD1A_5	VDD1A_6	VDD1A_7
R	VSS	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 133 • Top-Right Pin Diagram

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	TXVND_0	TXVPD_0	D
VSS_17	VSS_18	VSS_19	VDD1_2	VDD25A_7	1588_PPS_RI	CLK_SQUELCH_IN	1588_SPI_CLK	E
VSS_24	VSS_25	VSS_26	VDD1_4	PHYADD1	PHYADD4	SPI_IO_DO	RCVRDCLK1	F
VSS_30	VSS_31	VSS_32	VDD1_6	PHYADD2	PHYADD3	SPI_IO_DI	RCVRDCLK2	G
VSS_36	VSS_37	VSS_38	VDD1_8	VDD25_1	1588_SPI_DO/GPIO13	SPI_IO_CLK	VSS	H
VSS_42	VSS_43	VSS_44	VDD1_10	SPI_IO_CS	1588_SPI_CS/GPIO12	1588_DIFF_INPUT_CLK_P	1588_DIFF_INPUT_CLK_N	J
VSS_48	VSS_49	VSS_50	VDD1_12	GPI08/I2C_SDA	GPI09/FASTLINK-FAL	1588_LOAD_SAVE/GPI0	1588_PPS_0/GPI01	K
VSS_54	VSS_55	VSS_56	VDD1_14	GPI04/I2C_SCL_0	GPI05/I2C_SCL_1	GPI06/I2C_SCL_2	GPI07/I2C_SCL_3	L
VSS_60	VSS_61	VSS_62	VDD1_16	VDD25_3	SIGDET1/GPI01	SIGDET2/GPI02	SIGDET3/GPI03	M
VSS_66	VSS_67	VSS_68	VDD1_18	SerDes_Rext_1	SIGDET0/GPI00	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 device.

6.3.1 1588 Support

The following table lists the 1588 support pins.

Table 196 • 1588 Support Pins

Name	Pin	Type	Description
1588_DIFF_INPUT_CLK_N	J16	ADIFF	Differential reference clock input pair.
1588_DIFF_INPUT_CLK_P	J15		
1588_PPS_1/1588_SPI_IN_CLK	N3	I/O, PU, 3 V	1588 local timer 1 PPS fixed to local timestamp counter PHY1. 1588 serial peripheral interface input clock.
1588_PPS_2/1588_SPI_IN_CS	N4	I/O, PU, 3 V	1588 local timer 2 PPS fixed to local timestamp counter PHY2. 1588 serial peripheral interface input chip select.
1588_SPI_CLK	E16	I	1588 SPI clock.
1588_PPS_3/1588_SPI_IN_DI	L2	I/O, PU, 3 V	1588 local timer 3 PPS fixed to local timestamp counter PHY3. 1588 serial peripheral interface input data input.

Table 196 • 1588 Support Pins (continued)

Name	Pin	Type	Description
1588_PPS_RI	E14	I/O, PU, 3 V	PPS return input signal.

6.3.2 1588 Support and GPIO

The following table lists the 1588 support and GPIO pins.

Table 197 • 1588 Support and GPIO Pins

Name	Pin	Type	Description
1588_LOAD_SAVE/GPIO10	K15	I/O, PU, 3 V	Sync signal to load the time to the 1588 engine. Rising edge triggered. Can be configured to serve as General Purpose Input/Output (GPIO).
1588_PPS_0/GPIO11	K16	I/O, PU, 3 V	1588 local timer 0 PPS configurable to local timestamp counter PHY0 through PHY3. Can be configured to serve as General Purpose Input/Output (GPIO).
1588_SPI_CS/GPIO12	J14	I/O, PU, 3 V	1588 SPI chip select. Can be configured to serve as General Purpose Input/Output (GPIO).
1588_SPI_DO/GPIO13	H14	I/O, PU, 3 V	1588 SPI data output. Can be configured to serve as General Purpose Input/Output (GPIO).

6.3.3 GPIO and Two-Wire Serial

The following table lists the GPIO and two-wire serial pins.

Table 198 • GPIO and Two-Wire Serial Pins

Name	Pin	Type	Description
GPIO4/I2C_SCL_0	L13	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller pins can be configured to serve as GPIOs.
GPIO5/I2C_SCL_1	L14	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller pins can be configured to serve as GPIOs.
GPIO6/I2C_SCL_2	L15	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller pins can be configured to serve as GPIOs.
GPIO7/I2C_SCL_3	L16	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller pins can be configured to serve as GPIOs.
GPIO8/I2C_SDA	K13	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller pins can be configured to serve as GPIOs.
GPIO9/FASTLINK-FAIL	K14	I/O, PU, 3 V	General purpose input/output (GPIO). The two-wire serial controller and fast link fail pins can be configured to serve as GPIOs.

6.3.4 JTAG

The following table lists the JTAG test pins.

Table 199 • JTAG Pins

Name	Pin	Type	Description
TCK	F3	I, PU, ST, 3 V	JTAG test clock input.
TDI	F2	I, PU, ST, 3 V	JTAG test serial data input.
TDO	F1	O	JTAG test serial data output.
TMS	E2	I, PU, ST, 3 V	JTAG test mode select.
TRST	E3	I, PU, ST, 3 V	JTAG reset.

Note: When JTAG is not in use, this pin must be tied to ground with a pull-down resistor for normal operation.

6.3.5 Miscellaneous

The following table lists the miscellaneous pins.

Table 200 • Miscellaneous Pins

Name	Pin	Type	Description
CLK_SQUELCH_IN	E15	I, PU, 3 V	Input control to squelch recovered clock.
COMA_MODE	L3	I, PU, 3 V	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. For more information, see Initialization , page 138. For more information about a typical bring-up example, see Configuration , page 137.
LED0_PHY[0:3]	G1, H1, J1, K1	O	LED direct-drive outputs. All LEDs pins are active-low. A serial LED stream can also be implemented. For more information, see LED Mode Select , page 155.
LED1_PHY[0:3]	G2, H2, J2, K2		
LED2_PHY[0:3]	G3, H3, J3, K3		
LED3_PHY[0:3]	G4, H4, J4, K4		
NC_[1:4]	A1, A16, T1, T16	NC	No connect.
NRESET	M3	I, PD, ST, 3 V	Device reset. Active low input that powers down the device and sets all register bits to their default state.
PHYADD1	F13	I, PU, 3V	Device SMI address bit 1. Normally tied to VSS unless an address offset of 0x2 is used by a specific system's station manager. For more information, see PHY Addressing , page 15.
PHYADD[2:4]	G13, G14, F14	I, PD, 3 V	Device SMI address bits 4:2. For more information, see PHY Addressing , page 15.

Table 200 • Miscellaneous Pins (continued)

Name	Pin	Type	Description
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	ABIAS	Reference filter connects to an external 1 μ F capacitor to analog ground.
REF_REXT_A	D4	ABIAS	Reference external connects to an external 2 k Ω (1%) resistor to analog ground.
REFCLK_N REFCLK_P	C1 D1	I, ADIFF	125 MHz or 25 MHz reference clock input pair. Must be capacitively coupled and LVDS compatible.
REFCLK_SEL2	E1	I, PU, 3 V	Selects the reference clock speed. 0: 25 MHz (VSS) 1: 125 MHz (2.5 V) Use 125 MHz for typical applications.
RESERVED_[1:8]	C10, D13, L4, P4, L1, M1, N1, P1		Reserved. Leave unconnected.
THERMDA	C3	A	Thermal diode anode.
THERMDC_VSS	D2	A	Thermal diode cathode connected to device ground. Temperature sensor must be chosen accordingly.

6.3.6 Power Supply and Ground

The following table lists the power supply and ground 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. No power supply sequencing is required. However, clock and power must be stable before releasing the NRESET pin.

Table 201 • Power Supply and Ground Pins

Name	Pin	Description
VDD_MDIO	P3	1.2 V or 2.5 V power for SMI pins.
VDD1_[1:18]	E5, E12, F5, F12, G5, G12, H5, H12, J5, J12, K5, K12, L5, L12, M5, M12, N5, N12	1.0 V internal digital logic.
VDD1A_[1:10]	C7, C8, C9, C13, P6, P7, P8, P9, P10, P11	1.0 V analog power requiring additional PCB power supply filtering. Associated with the QSGMII/SGMII MAC receiver output pins.
VDD25_[1:3]	H13, M4, M13	2.5 V general digital power supply. Associated with the LED, GPIO, JTAG, twisted pair interface, reference filter, reference external supply connect, and recovered clock pins.
VDD25A_[1:10]	C2, C4, C6, C11, C14, E4, E13, P5, P12, P13	2.5 V general analog power supply.

Table 201 • Power Supply and Ground Pins (continued)

Name	Pin	Description
VSS	D14, H16, R1	Ground.
VSS_[1:4]	B1, B16, C5, C12	Ground.
VSS_[6:68]	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	Ground.
VSS_70	R16	Ground.

6.3.7 SerDes MAC Interface

The following table lists the SerDes MAC interface pins. Leave unused SerDes transceiver pairs unconnected.

Table 202 • SerDes MAC Interface Pins

Name	Pin	Type	Description
RDN_0	P16	O, ADIFF	PHY0 QSGMII/SGMII/SerDes MAC receiver output pair.
RDP_0	P15		
RDN_[1:3]	T12, T8, T4	O, ADIFF	SGMII/SerDes MAC receiver output pair.
RDP_[1:3]	R12, R8, R4		
SerDes_Rext_0	P14	ABIAS	SerDes bias pins. Connect a 620 Ω 1% resistor across these pins.
SerDes_Rext_1	N13		
TDN_0	N16	I, ADIFF	PHY0 QSGMII/SGMII/SerDes MAC transmitter input pair.
TDP_0	N15		
TDN_[1:3]	T13, T9, T5	I, ADIFF	SGMII/SerDes MAC transmitter input pair.
TDP_[1:3]	R13, R9, R5		

6.3.8 SerDes Media Interface

The following table lists the SerDes media interface pins. Leave unused SerDes transceiver pairs unconnected.

Table 203 • SerDes Media Interface Pins

Name	Pin	Type	Description
FIBRIN_[0:3]	T15, T11, T7, T3	I, ADIFF	SerDes media receiver input pair.
FIBRIP_[0:3]	R15, R11, R7, R3	I, ADIFF	SerDes media receiver input pair.
FIBRON_[0:3]	T14, T10, T6, T2	O, ADIFF	SerDes media transmitter output pair.
FIBROP_[0:3]	R14, R10, R6, R2	O, ADIFF	SerDes media transmitter output pair.

6.3.9 Serial Management Interface

The following table lists the serial management interface (SMI) pins. The SMI pins are referenced to VDD_MDIO and can be set to either 1.2 V or 2.5 V. This interface must be set to the appropriate voltage that VDD_MDIO is set to.

Table 204 • SMI Pins

Name	Pin	Type	Description
MDC ¹	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	I/O, OD	Management interrupt signal. These pins can be tied together in a wired-OR configuration with a single pull-up resistor.
MDIO ^{1, 2}	N2	I/O, 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.

1. 3.3 V input tolerant when supply VDD_MDIO is at 2.5 V, and 2.5 V input tolerant when VDD_MDIO is at 1.2 V.
2. When the PHY drives read data on MDIO, it can only source 2.5 V to the MDIO pin.

6.3.10 SIGDET/GPIO

The following table lists the GPIO and signal detect pins.

Table 205 • SIGDET/GPIO Pins

Name	Pin	Type	Description
SIGDET0/GPIO0	N14	I/O, PU, 3 V	General purpose input/output (GPIO). The multipurpose SIGDET and fast link fail pins can be configured to serve as GPIOs.
SIGDET1/GPIO1	M14	I/O, PU, 3 V	General purpose input/output (GPIO). The multipurpose SIGDET and fast link fail pins can be configured to serve as GPIOs.
SIGDET2/GPIO2	M15	I/O, PU, 3 V	General purpose input/output (GPIO). The multipurpose SIGDET and fast link fail pins can be configured to serve as GPIOs.
SIGDET3/GPIO3	M16	I/O, PU, 3 V	General purpose input/output (GPIO). The multipurpose SIGDET and fast link fail pins can be configured to serve as GPIOs.

6.3.11 SPI Interface

The following table lists the SPI interface pins.

Table 206 • SPI Interface Pins

Name	Pin	Type	Description
SPI_IO_CLK	H15	I/O, PU, 3 V	Serial peripheral interface clock input from external device
SPI_IO_CS	J13	I/O, PU, 3 V	Serial peripheral interface chip select
SPI_IO_DI	G15	I/O, PU, 3 V	Serial peripheral interface data input
SPI_IO_DO	F15	I/O, PU, 3 V	Serial peripheral interface data output

6.3.12 Twisted Pair Interface

The following table lists the twisted pair interface pins.

Table 207 • Twisted Pair Interface Pins

Name	Pin	Type	Description
TXVNA_[0:3]	B14, B10, B6, B2	ADIFF	TX/RX channel A negative signal
TXVNB_[0:3]	B15, B11, B7, B3	ADIFF	TX/RX channel B negative signal
TXVNC_[0:3]	C15, B12, B8, B4	ADIFF	TX/RX channel C negative signal
TXVND_[0:3]	D15, B13, B9, B5	ADIFF	TX/RX channel D negative signal
TXVPA_[0:3]	A14, A10, A6, A2	ADIFF	TX/RX channel A positive signal
TXVPB_[0:3]	A15, A11, A7, A3	ADIFF	TX/RX channel B positive signal
TXVPC_[0:3]	C16, A12, A8, A4	ADIFF	TX/RX channel C positive signal
TXVPD_[0:3]	D16, A13, A9, A5	ADIFF	TX/RX channel D positive signal

7 Package Information

VSC8584XKS-11 and VSC8584XKS-14 are packaged in a lead-free (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-free products from Microsemi 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 device.

7.1 Package Drawing

The following illustration shows the package drawing for the device. The drawing contains the top view, bottom view, side view, dimensions, tolerances, and notes.

Figure 134 • 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 208 • Thermal Resistances

Symbol	°C/W	Parameter
θ_{JCTop}	5.1	Die junction to package case top
θ_{JB}	10.5	Die junction to printed circuit board
θ_{JA}	19.6	Die junction to ambient
θ_{JMA} at 1 m/s	16.3	Die junction to moving air measured at an air speed of 1 m/s
θ_{JMA} at 2 m/s	14.2	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

This device is rated moisture sensitivity level 4 as specified in the joint IPC and JEDEC standard IPC/JEDEC J-STD-020. For more information, see the IPC and JEDEC standard.

8 Design Considerations

This section provides information about design considerations for the VSC8584-11 device.

8.1 Clause 45 register address post-increment

Clause 45 register address post-increment only works when reading registers and only when extended page access register 31 is set to 0.

The workaround is to access registers individually.

8.2 Clause 45, register 7.60

Clause 45, register 7.60 bit 10 always reads back value 1. However per IEEE 802.3z, reserved bits should read back value 0.

This has a minor implication on software that needs to ignore bit 10 in the read-back value of clause 45 register 7.60.

8.3 10BASE-T signal amplitude

10BASE-T signal amplitude can be lower than the minimum specified in IEEE 802.3 paragraph 14.3.1.2.1 (2.2 V) at low supply voltages. Additionally, associated templates may be marginal or have failures.

This issue is not estimated to present any system level impact. Performance is not impaired with cables up to 130 m with various link partners.

8.4 10BASE-T Half-Duplex linkup after initial reset from power up

After initial power-on, register 0 Mode Control of port 0 of the device may contain all zeros and thus operate in forced 10BASE-T half-duplex mode.

The workaround is to perform a software reset of port 0 using register 0 bit 15 to restore the power-on default values. For more information, see [Mode Control](#), page 141.

8.5 Link performance in 100BASE-TX and 1000BASE-T modes

PHY ports may exhibit sub-optimal performance under certain environmental and cabling conditions without proper initialization.

Furthermore, under worst-case operating conditions while in 100BASE-TX mode, the PHY cannot compensate for the limits of attenuation and phase distortion introduced at maximum cable lengths, or compensate for worst-case baseline wander introduced by 100BASE-TX “killer” packets.

Contact Microsemi for a script that needs to be applied during system initialization.

8.6 Clause 36 PCS incompatibilities in 1000BASE-X media mode

While operating in 1000BASE-X media mode, certain Media Access Controllers may reject Carrier Extension patterns in ingress frame traffic, in violation of the Clause 36 PCS specification.

The workaround to ensure no frame errors with such link-partner MACs is to disable Carrier Extensions in 1000BASE-X mode by setting register 23E3 bit 12.

8.7 1000BASE-X parallel detect mode with Clause 37 auto-negotiation enabled

When connected to a forced-mode link partner and attempting auto-negotiation, the PHY in 1000BASE-X parallel detect mode requires a minimum 250 millisecond IDLE stream in order to establish a link. If the PHY port is programmed with 1000BASE-X parallel detect-enabled (MAC-side register 16E3 bit 13, or media-side register 23E3 bit 13), then a forced-mode link partner sending traffic with an inter-packet gap less than 250 milliseconds will not allow the local device's PCS to transition from a link-down to link-up state.

8.8 Near-end loopback non-functional in protocol transfer mode

Near-end loopback does not work correctly when the device is configured in protocol transfer mode.

This is a debug feature and does not have any effect on the normal operation of the device.

8.9 LED Duplex/Collision function not working when in protocol transfer mode 10/100 Mbps

When a PHY port is configured for protocol transfer mode of operation, the Duplex/Collision function of an LED does not properly indicate if the link is full-duplex.

Because protocol transfer mode is used with copper SFP modules, the PHY-copper SFP link will always operate in full-duplex mode, and thus the PHY LED duplex indicator has little real-world significance.

8.10 Fast Link Failover indication delay when using interrupts

Whenever the Fast Link Failure interrupt mask is enabled (register 25, bit 7), MDINT will assert at the onset of a link failure in less than 6 ms typical (8 ms, worst-case).

8.11 Anomalous Fast Link Failure indication in 1000BT Energy Efficient Ethernet mode

When a port is linked in 1000BT in EEE mode, the Fast Link Failure indication may falsely assert while processing bursting traffic. EEE should be disabled when FLF indication is in use.

8.12 Anomalous PCS error indications in Energy Efficient Ethernet mode

When a port is processing traffic with EEE enabled on the link, certain PCS errors (such as false carriers, spurious start-of-stream detection, and idle errors) and EEE wake errors may occur. There is no effect on traffic bit error rate, but some error indications should not be used while EEE is enabled. These error indications include false carrier interrupts (interrupt status, register 26 bit 3), receive error interrupts (interrupt status, register 26 bit 0), and EEE wake error interrupts.

Contact Microsemi for a script that needs to be applied during system initialization if EEE will be enabled.

8.13 EEE allowed only for MACs supporting IEEE 802.3az-2010 in 1588 applications

When the 1588 processor is enabled and the PHY is advertising EEE mode of operation, the PHY requires a 802.3az-2010 MAC to control EEE features.

EEE operation with legacy MACs that do not support 802.3az-2010 is not allowed whenever the 1588 time stamping engine is used.

8.14 Auto-Negotiation Management Register Test failures

MII register 10 bit 15 (Master/Slave Configuration Fault) does not self-clear before a subsequent 1000BT link up attempt. This is a minor bug, and can be worked around by reading the register twice to ensure bit 15 status shows the correct 1000BT auto-negotiation result.

MII register 4 bit field 4:0 (Reserved Selector Fields) will accept values other than IEEE 802.3 Ethernet code values to be transmitted by the PHY during auto-negotiation base page exchange. This is a minor bug because it requires deliberate misconfiguration of the selector field in register 4, which the PHY API does not perform.

8.15 1588 bypass switch may drop packets

Up to two 1588 packets in the 1588 engine pipeline could be dropped if the bypass switch is activated on the fly during live 1588 traffic. Disabling the bypass switch at any time is acceptable.

Whenever the MACsec datapath is enabled, the 1588 bypass feature should never be used.

8.16 1588 bypass shall be enabled during engine reconfiguration

When the 1588 datapath is enabled, the 1588 bypass feature shall be enabled before reprogramming 1588 configuration registers. It is recommended to disable 1588 bypass before live traffic begins flowing through the re-provisioned port. Whenever the MACsec datapath is enabled, no traffic should be flowing through the port before enabling or disabling the 1588 bypass feature.

8.17 Time stamp errors due to IEEE 1588 reference clock interruption

Interruption of the IEEE 1588 reference clock after releasing device hardware reset will corrupt the local time counter value and may cause loss of 1588 processor time stamp coherency. After clock interruption, a local time counter reload and execution of the 1588 out-of-sync (OOS) recovery routine within the Unified API is required.

Contact Microsemi for the Unified API v4.67.03 or later with these capabilities, and for additional usage information on the OOS recovery routine.

8.18 Soft power-down procedure

The Unified API is required to use the power-down control bit (device register 0, bit 11).

Contact Microsemi to obtain Unified API-4.67.03 or later, which performs the correct procedures involving the power-down control bit.

8.19 Egress time stamp variation in PHY Packet Mode of operation

When the MACsec datapath is enabled with static bypass set, and the 1588 processor is also enabled, the PHY will time stamp outgoing PTP frames with up to 32 bit-times of error.

A software workaround is to disable the MACsec processor rather than set the static bypass feature whenever it is desired to skip MACsec processing.

8.20 Changing speed from 1 Gbps to 100 Mbps hangs MACsec engine

When MACsec is enabled and traffic is flowing, the flow control buffer inside the MACsec engine may hang when the PHY media speed is changed. The workaround is to poll the flow control buffer overflow status bit from the MACsec engine, and if the buffer overflow bit is set, command a soft reset of the MACsec engine and then reconfigure it.

8.21 VTSS_MACSEC_uncontrolled_counters_get show incorrect counter values

The counter value displayed by `vtss_macsec_uncontrolled_counters_get` is incorrect. When `protectframes == false`, the controlled port counter will not include untagged frames. This is considered a minor bug because the `protectframes == false` control is provided to facilitate MACsec deployment and is not the usual MACsec mode of operation.

The Microsemi API will return NULL values for counter return when `protectframes == false` to indicate that there is a bug in this function and this function is not available.

8.22 Controlled port counter `if_in_octets` does not get set correctly

When `validateFrames == Disabled`, `InOctetsValidated/Decrypted` is not incremented (as per standard). As a result, the `if_in_octets` calculation of a controlled port is incorrect. This is considered a minor bug because the `validateFrames == Disabled` control is provided to facilitate MACsec deployment and is not the usual MACsec mode of operation.

The Microsemi API will return NULL values for `if_in_octets` calculations to indicate there is a bug in this function.

8.23 Far-end loopback requires disabling Tx preamble fix for 1000M traffic

Any traffic flowing at 1000 Mbps on the port when far-end loopback is enabled may be dropped if Tx preamble fixing (22E3 bit 13) is also enabled.

The workaround for this issue is to clear bit 22E3.13 before using far-end loopback. Once loopback operation is disabled, then 22E3 bit 13 should be set again before flowing traffic.

8.24 Non-standard preamble frames discarded by MACsec engine

When MACsec is enabled, the PHY only processes incoming frames with preamble lengths less than or equal to 8 bytes. Frames with longer preambles are discarded by the MACsec IP block and do not appear at the external MAC interface of the device.

8.25 Pause control frames `pause_timer` may be exceeded by MACsec engine

When MACsec is enabled, a pause control frame's `pause_timer` value may be marginally exceeded before the PHY resumes transmission of MAC client frames.

This issue has no real-world impact to observable delays in traffic from the PHY device because the threshold of failure is on the order of microseconds.

8.26 MACsec engine may shrink the minimum 100M IPG during wire-speed transmission

When MACsec is enabled and transmitting in 100 Mbps mode, the PHY may transmit frames with `InterPacketGap` as low as 72 bit times during frame transmissions at line speed. However, the PHY maintains an average `InterPacketGap` of 96 bit times as per the IEEE 802.3 standard. Thus, there is marginal impact on frequency offset tolerance between link partner PHYs when operating at 100 Mbps.

8.27 Operate MACsec with flow control to handle bandwidth expansion

Because MACsec processing results in frame expansion, flow control must be enabled to manage bandwidth at the host interface. When MACsec is enabled, the Unified API enables advanced flow control handling (see [Advanced Flow Control Handling](#), page 51).

The host MAC to which the PHY is connected must also have flow control enabled in order to react to PAUSE commands from the PHY. If flow control is disabled, egress traffic originating from the PHY client at wire speed may result in truncated (invalid) frames.

This is only a practical issue during testing scenarios because this invalid frame error only occurs when flow control is disabled, which would not be useful in a real-world network.

The workaround and recommended method to operate MACsec is with pause control enabled such that the PHY pauses host egress traffic when utilization exceeds the available bandwidth of the line.

8.28 Frames with nibbles after the FCS may be discarded by MACsec engine

When MACsec is enabled, ingress traffic in 10 Mbps or 100 Mbps mode received with an extra nibble following a valid FCS field may be dropped. This scenario only happens during MACsec operation in a half-duplex collision-domain. In the rare event that collision occurs at the byte after FCS, the IP packet is dropped by the PHY and retransmission requested.

8.29 10BASE-T ingress time stamp errors when MACsec enabled

When MACsec is enabled and operating in 10BASE-T mode, the PHY may time stamp PTP frames with up to $\pm 6.4 \mu\text{s}$ of error.

8.30 MACsec datapath switch may cause large time stamp errors in the 1588 engine

Static bypass during MACsec initialization must be toggled prior to 1588 engine initialization, or there may be a loss of timing coherency between the MACsec datapath and the 1588 processor unit.

Contact Microsemi to obtain Unified API-4.67.03 or later for this implementation.

8.31 Packet mode datapath while processing 1588 traffic

When the MACsec datapath is enabled and the 1588 processor is also enabled, the PHY requires the host and line MACs in the device to prevent undersize frames from being processed by the 1588 engine, using the `UNDERSIZED_FRAME_DROP` feature of the host and line MACs.

Contact Microsemi for the Unified API v4.67.03 or later with these capabilities.

9 Ordering Information

The device is offered with two operating temperature ranges. The range for VSC8584-11 is 0 °C ambient to 125 °C junction. The range for VSC8584-14 is –40 °C ambient to 125 °C junction.

VSC8584XKS-11 and VSC8584XKS-14 are packaged in a lead-free (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-free products from Microsemi 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.

Table 209 • Ordering Information

Part Order Number	Description
VSC8584XKS-11	Lead-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 ¹ .
VSC8584XKS-14	Lead-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 ¹ .

1. For carrier class applications, the maximum operating temperature is 110 °C junction.