SWITCHING

Marvell Solutions

Link Street® 88E6092/88E6095/ 88E6095F Datasheet

8 FE + 3 GE Stackable Ethernet Switch with QoS and 802.1Q

Doc. No. MV-S102629-00 Rev. C
June 9, 2006

MOVING FORWARD FASTER®

MARVELL®

Document Status	
Advance Information	This document contains design specifications for initial product development. Specifications may change without notice. Contact Marvell Field Application Engineers for more information.
Preliminary Information	This document contains preliminary data, and a revision of this document will be published at a later date. Specifications may change without notice. Contact Marvell Field Application Engineers for more information.
Final Information	This document contains specifications on a product that is in final release. Specifications may change without notice. Contact Marvell Field Application Engineers for more information.
Revision Code: Rev. C	
Advance	Technical Publication: 1.40

OVERVIEW

The Marvell® 88E6092, 88E6095, and 88E6095F devices are each single-chip 8-port 10/100 plus 3 port Gigabit Ethernet switches with support for Quality of Service (QoS), 802.1Q, and RMON counters. They contain eight 10BASE-T/100BASE-TX transceivers (PHYs), and three SERDES interfaces that can be used to connect to external Marvell 10/100/1000 triple speed Ethernet transceivers (PHYs).

The 88E6092 is designed for unmanaged or smart managed applications, while the 88E6095 is for fully managed applications with enhanced features for stacking. The 88E6095F device is also for fully managed applications and stacking, and includes fiber enable and controls on all ports.The devices integrate QoS with four levels of priority per port. The QoS architecture switches packets into one of four traffic class queues based upon Port, IEEE 802.1p, IPv4 Type of Service (TOS) or Differentiated Services (DiffServ), IPv6 Traffic Class, 802.1Q VLAN ID, DA MAC address or SA MAC address. Back-pressure and pause frame-based flow control schemes are included to support zero packet loss under temporary traffic congestion. The lookup engine allows for up to 8K active nodes to be connected with the switch.

The switches contain eleven independent 802.3 media access controllers (MACs), a high-speed, non-blocking four traffic class QoS switch fabric, a high-performance address lookup engine, and a 1 Mbit frame buffer memory. Up to 32 or more 88E6095 devices can be cascaded. One or more devices can be used to create larger switches, e.g 24 FE + 2 GE or 48 FE + 4 GE. The devices are designed for cost-sensitive low to high port count switch systems that require Quality of Service, Trunking, Stacking(88E6095 and 888E6095F only), and/or Spanning Tree.

The devices offer enhanced stacking capabilities that enable multiple devices and/or systems to act as a single unit. These features include the ability to Link Aggregate ports between boxes in a stack, monitor any port from any other port in the stack, and create port-based VLANs using any port(s) within the stack.

The ninth and tenth ports' optional (G)MII (or MII) interface supports a direct connection to Management or Router CPUs with integrated MACs. For the 88E6092/88E6095 devices, Port 9 is only MII-PHY or MII-MAC mode capable, while Port 10 can be configured in either GMII mode, MII-PHY or MII-MAC mode. Port 9 and Port 10 of the 88E6095F device are both capable of GMII/MII MAC or PHY mode. These interfaces, along with BPDU handling for IEEE 802.1D Spanning Tree Protocol, 802.1w Rapid Spanning Tree (88E6095 and 88E6095F only), 802.1s Multiple VLAN Spanning Tree (88E6095 and 88E6095F only), programmable per-port VLAN configurations, 802.1Q and Port States, support fully managed switches and truly isolated WAN vs. LAN firewall applications. The devices support 4,096 802.1Q VLAN IDs which can be enabled on a per port basis. Three levels of 802.1Q security is supported with error frame trapping and logging.

The devices provide an 8K MAC address lookup table which allows for up to 8K active nodes to be connected with the switch. Back-pressure and Pause frame-based flow control schemes are included to support zero packet loss under temporary traffic congestion. The MAC units in the devices comply fully with the applicable sections of IEEE 802.3.

The PHY and SERDES units in the devices are designed with the Marvell® cutting-edge mixed-signal processing technology for digital implementation of adaptive equalization and clock data recovery. Special power management techniques are used to facilitate low power dissipation and high port count integration. Both the PHY and MAC units in the devices comply fully with the applicable sections of IEEE 802.3, IEEE 802.3u, and IEEE 802.3x standards.

The devices' many operating modes can be configured using SMI (serial management interface - MDC/MDIO) and/or a low cost serial EEPROM (93C46, C56 or C66). A standalone QoS mode is also supported.

APPLICATIONS

- 24 FE + 2 GE switch using three 88E6095 devices (Stackable, Managed) or three 88E6092 devices (Smart, Unmanaged)
- 48 FE + 4 GE switch using six 88E6095 devices with a single 88E6185 (Non blocking, Stackable, Managed) or six 88E6092 devices and one 88E6182 device (Smart, Unmanaged)
- 24 FE + 4 GE switch using three 88E6095 devices with a single 88E6185 (10 port GE switch)—Non blocking, stackable, Managed or three 88E6092 and 88E6182 devices (Smart, Unmanaged)
- 8 FE + 2 1000BASE-X using a single 88E6095 device (Managed) or 88E6092 (Smart, Unmanaged)
- SOHO/SMB router with eight FE and two GE ports
- Fiber to the Curb cascade applications

FEATURES

- Single chip integration of an 11 port 8-FE + 3-GE QoS switch and memory in a 20 x 20 mm 176-pin TQFP package (88E6092/88E6095 device), or 24 x 24 mm 216 LQFP (88E6095F device)
- Enables stackable systems through standard Ethernet interfaces (88E6095 and 88E6095F devices only)
- Quality of Service support with four traffic classes
- QoS determined by Port, IEEE 802.1p tagged frames, IPv4's Type of Service (TOS) & Differentiated Services (DS), IPv6's Traffic Class 802.1Q VID, Destination MAC address, or Source MAC address
- 802.1X MAC based authentication
- Port Trunking across chips
- Port Monitoring/Mirroring across chips
- 802.1Q VLAN support for the full 4,096 VLAN IDs
- Egress tagging/untagging selectable per port or by 802.1Q VLAN ID
- Port based VLANs supported in any combination across multiple chips
- Port States & BPDU handling for Spanning Tree
- 28 32-bit and 2 64-bit RMON Counters per port
- Integrated with eight independent Auto-Crossover Fast Ethernet transceivers fully compliant with the applicable sections of IEEE802.3 and IEEE802.3u
- 100% pin compatible to the 88E6095 or 88E6095F device
- Integrated with three independent triple speed SERDES transceivers to interface with Marvell® Alaska® gigabit copper PHYs
- Ports 8, 9 & 10 can optionally be configured as fiber ports (1000BASE-X) with direct connection to lasers
- Ports 9 and 10 can support GMII Mode (full-duplex), MII-MAC Mode (Forward) or MII-PHY Mode (Reverse—full-duplex) interface options for management and firewall applications (Port 9 is MII only in the 88E6092/88E6095 devices)
- Each GE port works at 10 Mbps or 100 Mbps, full or half-duplex mode, or 1000 Mbps full-duplex mode
- Shared 1 Mbit on-chip memory-based switch fabric with true non-blocking switching performance
- High performance lookup engine with support for up to 8K MAC address entries with automatic learning and aging
- Flexible LED support for Link, Speed, Duplex Mode, Collision, and Tx/Rx Activities
- Supports a low-cost 25 MHz XTAL clock source
- Low power dissipation P_{AVE} = 1W
- Supports the Marvell® Distributed Switching Architecture (DSA)
- for STP, up to 32 cascaded devices, and CPU-directed packet processing
- 176-TQFP package (88E6092/88E6095 device) and 216-LQFP package (88E6095F device)
- 88E6095 and 88E6095F available in Commercial and Industrial grade

Table 1: 88E6092/88E6095/88E6095F Device Differences

	Description	88E6092 Smart/ Unmanaged	88E6095 Managed/ Stackable	88E6095F Managed/ Stackable
Features	FE PHY and GE SERDES	8 + 3	8 + 3	8+3
	FE Fiber Support	0	0	8
	# GMII/MII (Shared w/SERDES)	1/2	1/2	2/2
	Packet Buffer Memory	1 Mbit	1 Mbit	1 Mbit
	# MAC Addresses	8K	8K	8K
QoS	Queues per Port	4	4	4
	802.1p, Port, TOS/DS, IPv6, TC, MAC	Yes	Yes	Yes
	Priority Override (DA, SA, VID), Remap	No	Yes	Yes
VLAN	Port -based and 802.1Q VLANs	Yes	Yes	Yes
	Maximum number of Shared 802.1Q VLANs	4096	4096	4096
	Maximum number of Independent 802.1Q VLANs	256	256	256
	Double Tagging (Q in Q)	No	Yes	Yes
Cross-chip	Stacking Support	No	Yes	Yes
	Cross -chip Trunking (#Trunks/#Ports)	Yes (8/8)	Yes (16/8)	Yes (16/8)
	Cross-chip Port Based VLANs	Yes	Yes	Yes
	Cross-chip Flow Control	Yes	Yes	Yes
	Distributed Switching Architecture (DSA) Tag Support	Yes	Yes	Yes
Manage-ment	802.1D Spanning Tree	Yes	Yes	Yes
	802.1w/s (Rapid/Multiple) STP	No	Yes	Yes
	Port Mirroring	Yes	Yes	Yes
	IGMP/MLD Snooping	Yes	Yes	Yes
Other	Ingress Storm Limiting	Yes	Yes	Yes
	Egress Rate Shaping	No	Yes	Yes
	802.1X Port & MAC Authentication	Yes	Yes	Yes
	Industrial Grade	No	Yes	Yes
	Package	176-pin LQFP	176-pin LQFP	216-pin LQFP

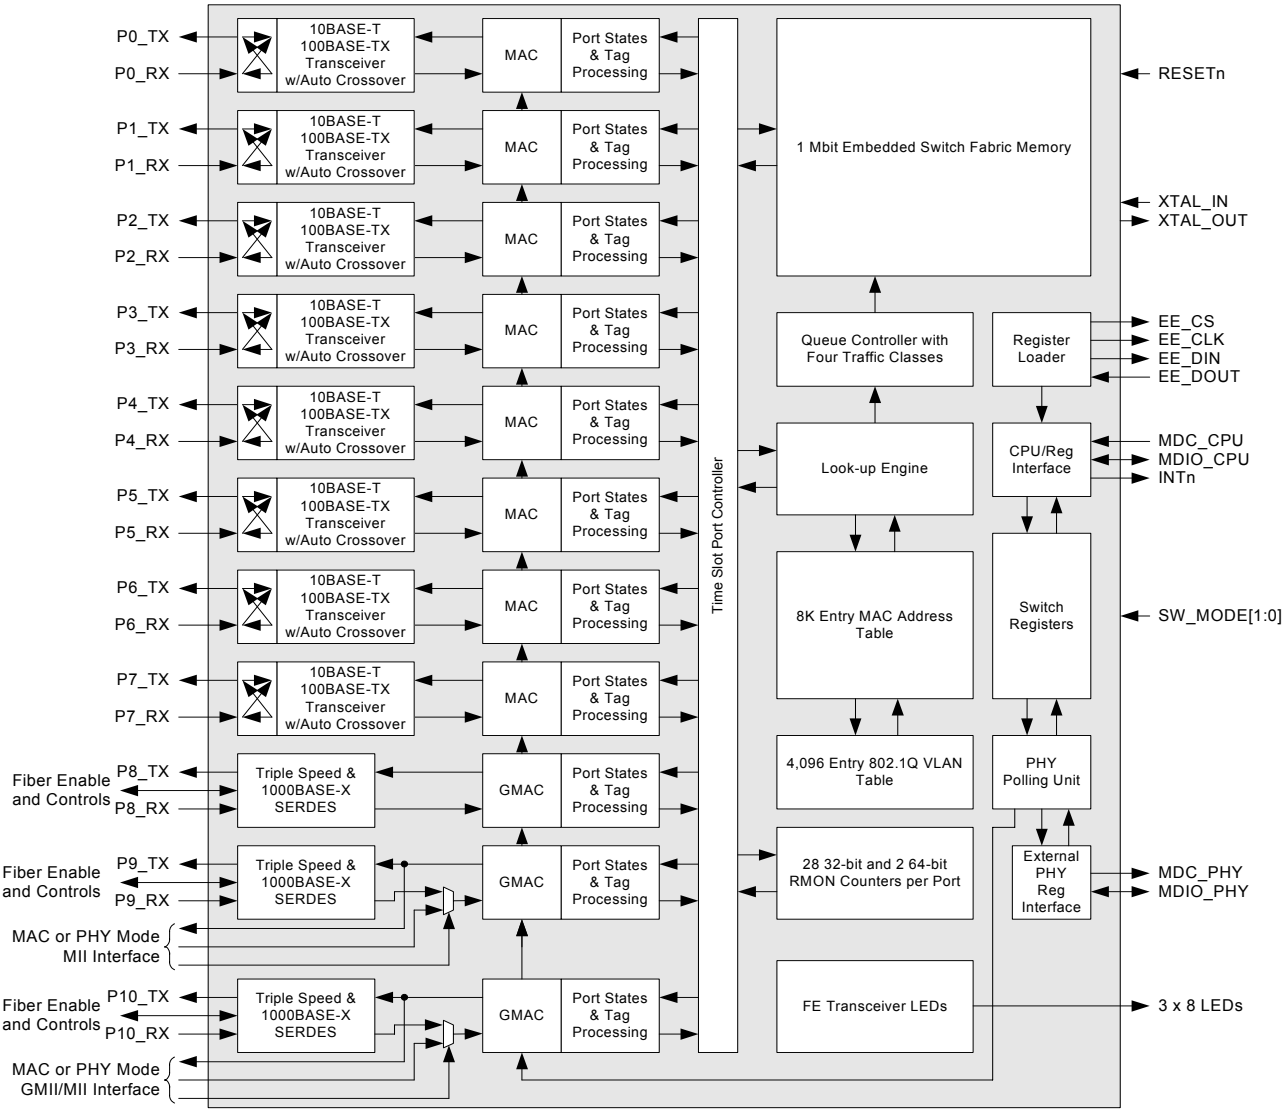

88E6092/88E6095 176 TQFP Top Level Block Diagram

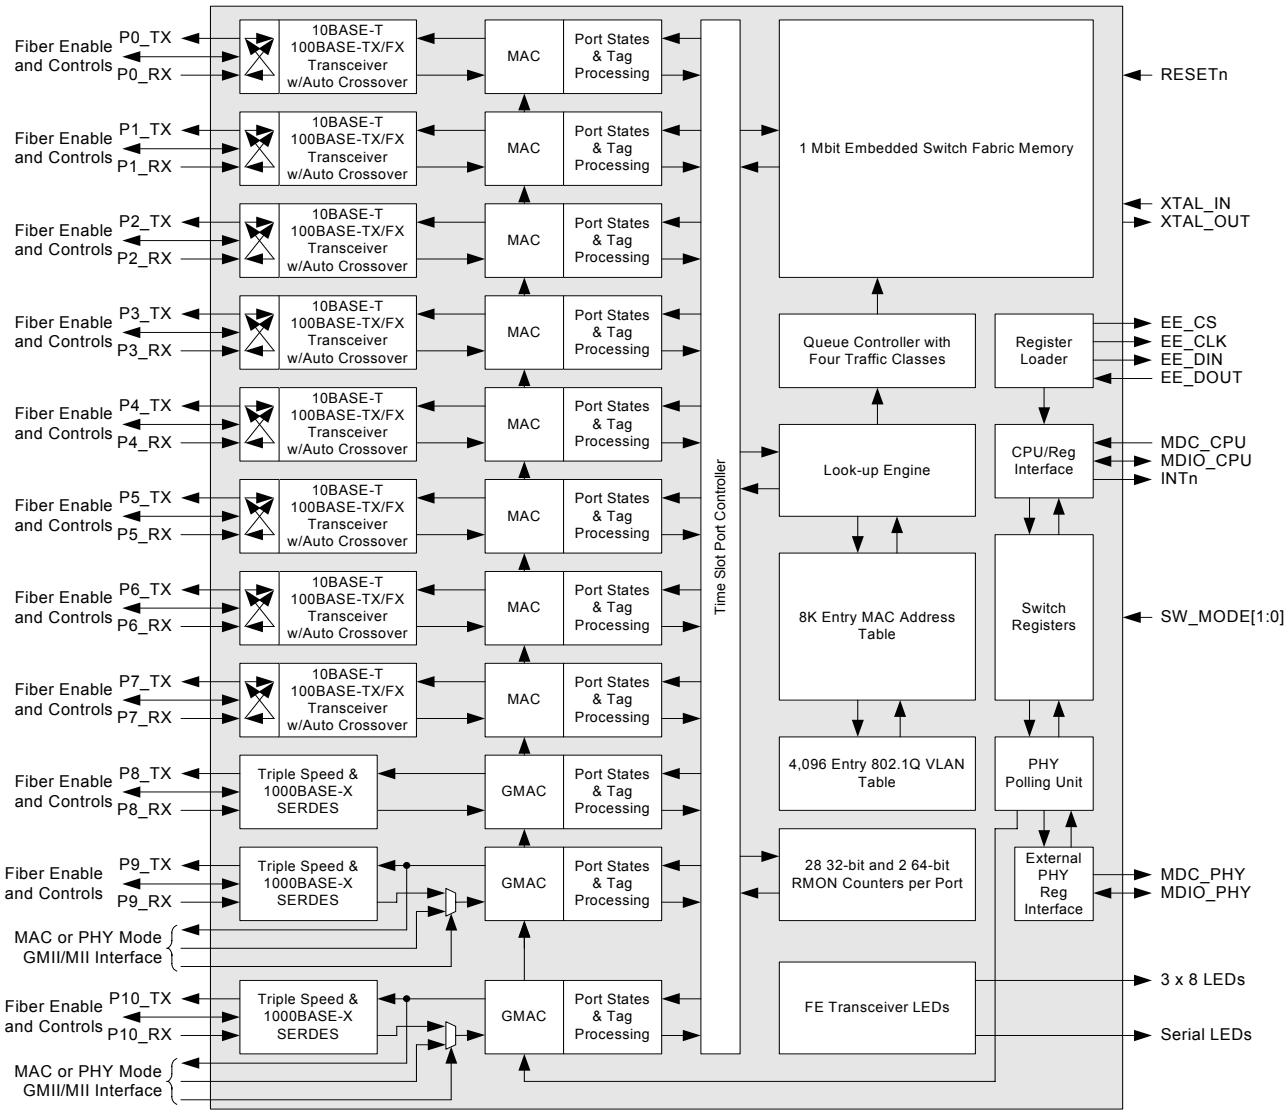

88E6095F 216 LQFP Top Level Block Diagram

Table of Contents

List of Tables

List of Figures

Section 1. Signal Description

Figure 1: 88E6092/88E6095 Device 176 TQFP EPAD Package Pinout

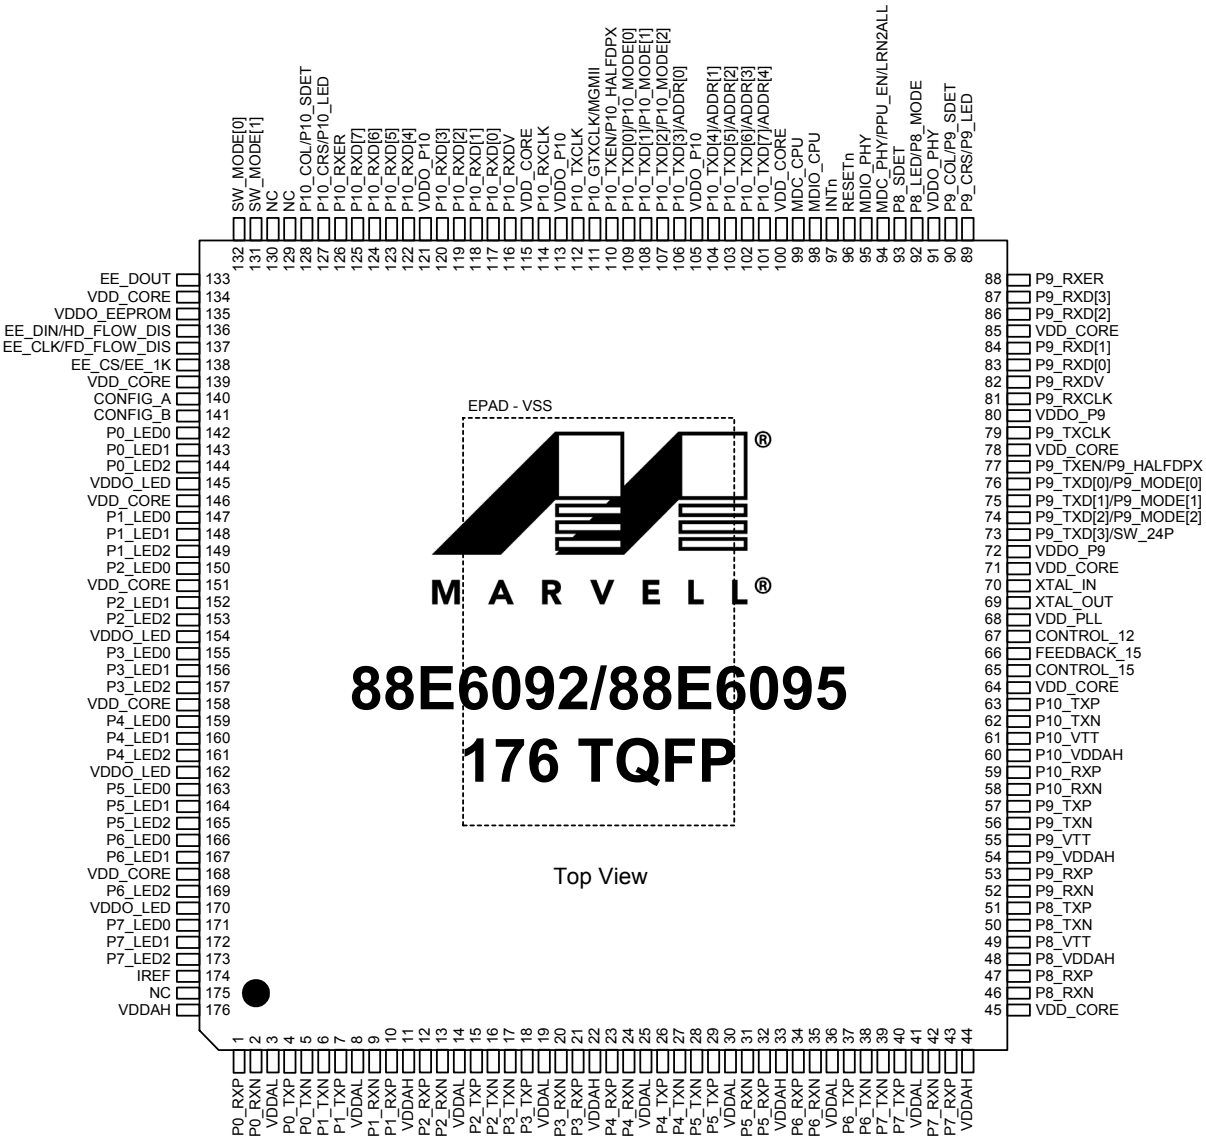

Figure 2: 88E6095F Device 216 LQFP EPAD Package Pinout

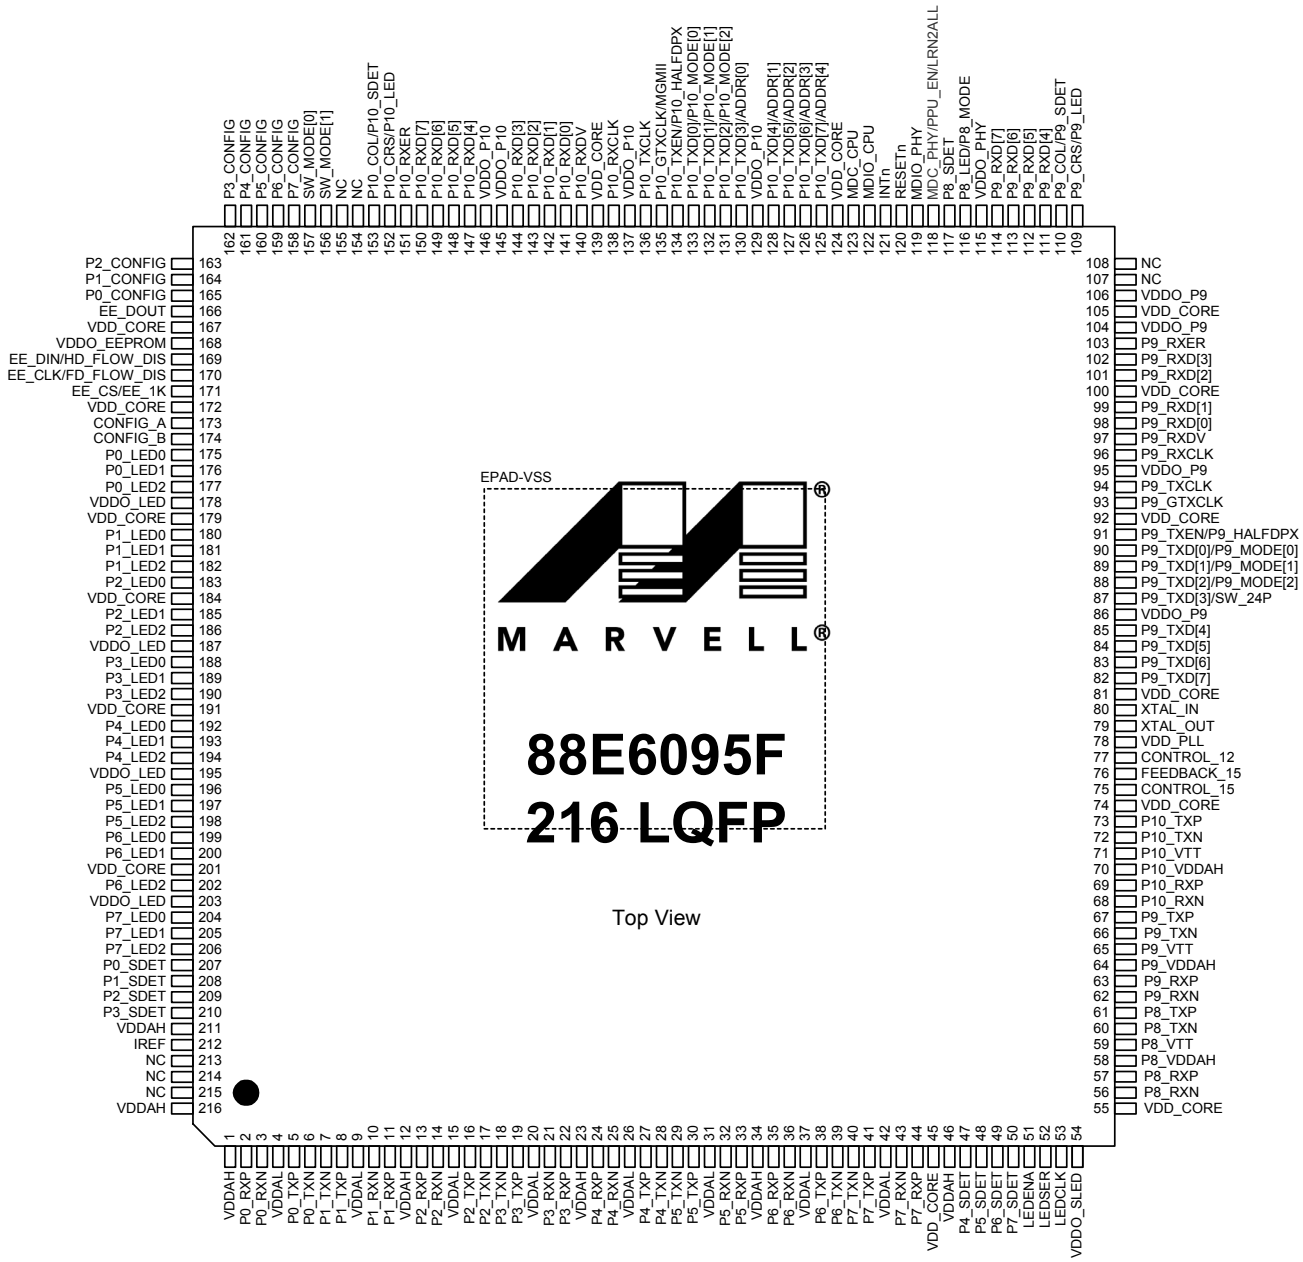

1.1 Pin Description

Table 2: Pin Type Definitions

Pin Type	Definition
H	Input with hysteresis
I/O	Input and output
I	Input only
O	Output only
PU	Internal pull-up
PD	Internal pull-down
D	Open drain output
Z	Tri-state output
mA	DC sink capability
Analog	Analog pin

Table 3: Network 10/100 PHY Interface (Ports 0 to 7)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
43 34 32 23 21 12 10 1	44 35 33 24 22 13 11 2	P7_RXP P6_RXP P5_RXP P4_RXP P3_RXP P2_RXP P1_RXP P0_RXP	Typically Input	Receiver input – Positive. P[7:0]_RXP connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode (available in the 88E6095F device only) RXP connects directly to the fiber-optic receiver's positive output. For lowest power, all unused port RXP pins should be tied to VSS. These pins can become outputs if Auto MDI/ MDIX Crossover is enabled.
42 35 31 24 20 13 9 2	43 36 32 25 21 14 10 3	P7_RXN P6_RXN P5_RXN P4_RXN P3_RXN P2_RXN P1_RXN P0_RXN	Typically Input	Receiver input – Negative. P[7:0]_RXN connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode (available in the 88E6095F device only) RXN connects directly to the fiber-optic receiver's negative output. For lowest power, all unused port RXN pins should be tied to VSS. These pins can become outputs if Auto MDI/ MDIX Crossover is enabled.
40 37 29 26 18 15 7 4	41 38 30 27 19 16 8 5	P7_TXP P6_TXP P5_TXP P4_TXP P3_TXP P2_TXP P1_TXP P0_TXP	Typically Output	Transmitter output – Positive. P[7:0]_TXP connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode (available in the 88E6095F device only) TXP connects directly to the fiber-optic transmitter's positive input. For lowest power, all unused port TXP pins should be tied to VSS. These pins can become inputs if Auto MDI/ MDIX Crossover is enabled.
39 38 28 27 17 16 6 5	40 39 29 28 18 17 7 6	P7_TXN P6_TXN P5_TXN P4_TXN P3_TXN P2_TXN P1_TXN P0_TXN	Typically Output	Transmitter output – Negative. P[7:0]_TXN connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode (available in the 88E6095F device only) TXN connects directly to the fiber-optic transmitter's negative input. For lowest power, all unused port TXN pins should be tied to VSS. These pins can become inputs if Auto MDI/ MDIX Crossover is enabled.
	50 49 48 47 210 209 208 207	P7_SDET P6_SDET P5_SDET P4_SDET P3_SDET P2_SDET P1_SDET P0_SDET	Input	Signal Detect input (88E6095F device only). If port 0 to 7 are configured for 100BASE-FX mode Px_SDET indicates whether a signal is detected by the fiber-optic transceiver. A positive level indicates that a signal is detected. If port 0 to 7 are configured for 10/100BASE-T mode Px_SDET is not used but it can not be left floating since these pins do not contain internal resistors. Px_SDET can be tied to VSS or VDDO either directly or through a 4.7kΩ resistor.

Table 4: PHY Configuration (Ports 0 to 7)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
	158 159 160 161 162 163 164 165	P7_CONFIG P6_CONFIG P5_CONFIG P4_CONFIG P3_CONFIG P2_CONFIG P1_CONFIG P0_CONFIG	Input	Port 7 to 0 Configuration (88E6095F device only). The Px_CONFIG pin is used to set the default configuration for Port 7 to 0 by connecting these pins to other device pins as follows: VSS = Auto-Negotiation enabled P0_LED1 = Forced 10BASE-T half-duplex P0_LED2 = Forced 10BASE-T full-duplex P1_LED0 = Forced 100BASE-TX half-duplex P1_LED1 = Forced 100BASE-TX full-duplex P1_LED2 = Forced 100BASE-FX half-duplex VDDO = Forced 100BASE-FX full-duplex Ports 7 to 0 default configuration is Auto-Negotiation enabled in the 88E6092/88E6095 device. Any port's default configuration can be modified by accessing the PHY registers by a CPU or a serial EEPROM. Fiber mode vs. copper mode cannot be configured in this way, however. Fiber vs. copper must be selected at Reset by using these pins. Px_CONFIG pins are configured after reset and contain internal pull-down resistors so they can be left floating.
140	173	CONFIG_A	Input	Global Configuration A. This global configuration pin is used to set the default LED mode and Far End Fault Indication (FEFI) mode by connecting these pins to other device pins as follows: VSS = LED Mode 0, FEFI disabled P0_LED0 = LED Mode 0, FEFI enabled P0_LED1 = LED Mode 1, FEFI disabled P0_LED2 = LED Mode 1, FEFI enabled P1_LED0 = LED Mode 2, FEFI disabled P1_LED1 = LED Mode 2, FEFI enabled P1_LED2 = LED Mode 3, FEFI disabled VDDO = LED Mode 3, FEFI enabled The CONFIG_A pin is configured after reset and does contain an internal pull-up resistor so it can be left floating.

Table 4: PHY Configuration (Ports 0 to 7) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	174	CONFIG_B	Input	Global Configuration B. This global configuration pin is used to set the default mode for Auto Crossover, the PHY driver type and Energy Detect by connecting these pins to other device pins as follows: VSS = No Crossover, Class A drivers, Energy Detect disabled P0_LED0 = No Crossover, Class A drivers, Energy Detect enabled P0_LED1 = No Crossover, Class B drivers, Energy Detect disabled P0_LED2 = No Crossover, Class B drivers, Energy Detect enabled P1_LED0 = Auto Crossover, Class A drivers, Energy Detect disabled P1_LED1 = Auto Crossover, Class A drivers, Energy Detect enabled P1_LED2 = Auto Crossover, Class B drivers, Energy Detect disabled VDDO = Auto Crossover, Class B drivers, Energy Detect enabled The CONFIG_B pin is configured after reset and does contain an internal pull-up resistor so it can be left floating.

Table 5: Port Status LEDs (Ports 0 to 7)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
173 169 165 161 157 153 149 144	206 202 198 194 190 186 182 177	P7_LED2 P6_LED2 P5_LED2 P4_LED2 P3_LED2 P2_LED2 P1_LED2 P0_LED2	Output	Parallel LED outputs—one for each port. This active low LED pin directly drives an led in parallel led mode. it can be configured to display many options. P[7:0]_LED2 are driven active low whenever RESETn is active low.
172 167 164 160 156 152 148 143	205 200 197 193 189 185 181 176	P7_LED1 P6_LED1 P5_LED1 P4_LED1 P3_LED1 P2_LED1 P1_LED1 P0_LED1	Output	Parallel LED outputs – one for each port. This active low LED pin directly drives an LED in Parallel LED mode. It can be configured to display many options. P[7:0]_LED1 are driven active low whenever RESETn is active low.
171 166 163 159 155 150 147 142	204 199 196 192 188 183 180 175	P7_LED0 P6_LED0 P5_LED0 P4_LED0 P3_LED0 P2_LED0 P1_LED0 P0_LED0	Output	Parallel LED outputs – one for each port. This active low LED pin directly drives an LED in Parallel LED mode. It can be configured to display many options. P[7:0]_LED0 are driven active low whenever RESETn is active low.
	52	LEDSER	Output	88E6095F device only; LEDSER outputs serial status bits that can be shifted into a shift register to be displayed via LEDs. LEDSER is output synchronously with LEDCLK.
	51	LEDENA	Output	88E6095F device only; LEDENA asserts High whenever LEDSER has valid status that is to be stored into the shift register. LEDENA is output synchronously with LEDCLK.
	53	LEDCLK	Output	88E6095F device only; LEDCLK is the reference clock for the serial LED signals.

Table 6: SERDES/1000BASE-X Interface (Ports 8 to 10)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
59 53 47	69 63 57	P10_RXP P9_RXP P8_RXP	Input	Receiver input – Positive. P[10:8]_RXP connects directly to the fiber-optic receiver's positive output or to another device's TXP (Transmitter output—positive) pins.
58 52 46	68 62 56	P10_RXN P9_RXN P8_RXN	Input	Receiver input – Negative. P[10:8]_RXN connects directly to the fiber-optic receiver's negative output or to another device's TXN (Transmitter output—negative) pins.
63 57 51	73 67 61	P10_TXP P9_TXP P8_TXP	Output	Transmitter output – Positive. P[10:8]_TXP connects directly to the fiber-optic transmitter's positive input or to another device's RXP (Receiver input—positive) pins.
62 56 50	72 66 60	P10_TXN P9_TXN P8_TXN	Output	Transmitter output – Negative. P[10:8]_TXN connects directly to the fiber-optic transmitter's negative input or to another device's RXN (Receiver input – Negative) pins.
93	117	P8_SDET	Input	[redacted]. A positive level indicates that a signal is detected. If Port 8 is configured for SERDES mode P8_SDET is not used but it can be left floating since this pin contains an internal pull-down resistor.

Table 6: SERDES/1000BASE-X Interface (Ports 8 to 10) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	116	P8_LED /P8_MODE	Typically Output	LED Output and Mode. If Port 8 is configured for 1000BASE-X mode Px_LED indicates whether a 'Link' signal is detected by the Px_SDET input. A low output level indicates that a positive input signal is detected by Px_SDET and the port's PCS has determines link. Px_LED indicates 'Activity' by blinking high when the port receives or transmits a frame. [redacted]. When reset is asserted, Px_LED become inputs and the Port's configuration is latched at the rising edge of RESETn as follows: Low = Port is configured for cross-chip SERDES mode, AutoNeg is off High = Port is configured for 1000BASE-X mode AutoNeg is on **NOTE:** MGMII mode overrides the port's configuration made by this pin. MGMII mode is selected when the PHY Polling Unit (PPU) detects an external PHY is connected to this port. See "P10_GTXCLK/MGMII" pin description in Table 10, "GMII/MII Transmit Interface (Port 10)," on page 34 for MGMII mode details. Px_LED is internally pulled low via a resistor so the pin can be left floating when unused or to select cross-chip SERDES mode. Connecting an external LED and its resistor to VDDO will configure the port for 1000BASE-X mode.

Table 7: GMII/MII Receive Interface (Port 9 is MII only - 88E6092/88E6095 devices only)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
81	96	P9_RXCLK	Input	Receive Clock. RXCLK is a reference for RXDV, RX_ER and RXD[7:0]. The speed of RXCLK is expected to be 125 MHz, 25 MHz or 2.5 MHz depending upon the speed of the port. RXCLK is internally pulled high so the pin can be left unconnected if not used.
    87 86 84 83	114 113 112 111 102 101 99 98	P9_RXD[7] P9_RXD[6] P9_RXD[5] P9_RXD[4] P9_RXD[3] P9_RXD[2] P9_RXD[1] P9_RXD[0]	Input	P9_RXD[7:0]—88E6095F device Only. P9_RXD[3:0]—88E6092/88E6095 devices Only. Receive Data RXD[7:0] receives the data octet or nibble to be sent into the switch. RXD[7:0] must be synchronous to RXCLK. In 1000BASE mode, RXD[7:0] is used. In 100BASE and 10BASE modes, RXD[3:0] is used and RXD[7:4] is ignored. RXD[7:0] are internally pulled low via resistor so the pins can be left unconnected when they are not used.
82	97	P9_RXDV	Input	Receive Data Valid. Receive Data Valid is used to indicate when RXD[7:0] (or RXD[3:0] where appropriate) contains frame information. RXDV must be synchronous to RXCLK. RXDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.
88	103	P9_RXER	Input	Receive Data Error. Receive Data Error is used to indicate when an external device (like a PHY) detects an error. When this signal is high, the receive MAC will discard the frame and it will be counted as a frame with a CRC error. RXER must be synchronous to RXCLK. RXER is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 7: GMII/MII Receive Interface (Port 9 is MII only - 88E6092/88E6095 devices only) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
89	109	P9_CRS /P9_LED	Input or Output	Carrier Sense. Carrier sense is used to indicate carrier has been detected on the line. CRS is not synchronous to RXCLK. CRS is used for half-duplex modes only and is ignored when the port is in full-duplex. When Port 9 is configured for 1000BASE-X operation (see P9_MODE[2:0]) Px_LED indicates whether a 'Link' signal is detected by the Px_SDET input. A low output level indicates that a positive input signal is detected by Px_SDET and the port's PCS has determines link. Px_LED indicates 'Activity' by blinking high when the port receives or transmits a frame. CRS is internally pulled low via resistor so the pin can be left unconnected when it is not used. The direction of this pin is determined by the value of the P9_MODE[2:0] pins at the rising edge of RESETn (during RESETn this pin is an input).
90	110	P9_COL /P9_SDET	Input	Collision. Collision is used to indicate both transmit and receive are occurring at the same time in half-duplex mode. The Collision signal is not synchronous to RXCLK. The Collision signal is used for half-duplex modes only and is ignored when the port is in full-duplex. When Port 9 is configured for 1000BASE-X operation (see P9_MODE[2:0]) Px_SDET is an input used to indicate whether a signal is detected by the fiber-optic transceiver. A positive level indicates that a signal is detected. P9_COL is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 8: GMII/MII Transmit Interface (Port 9)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
	93	P9_GTXCLK	Output	Transmit Clock; 88E6095F device only; GTXCLK is a reference for TXEN and TXD[7:0] when the port is in GMII mode. The speed of GTXCLK is 125 MHz and is normally only driven when the speed of the port is 1000BASE. GTXCLK can be configured to output a 25 MHz or 2.5 MHz clock so it can be used as a clock source for P9_TXCLK and P9_RXCLK when no other clock sources are available (see P9_MODE[2:0]). GTXCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.
79	94	P9_TXCLK	Input	Transmit Clock. TXCLK is a reference for TXEN and TXD[3:0] when the port is in MII mode. The speed of TXCLK is 25 MHz or 2.5 MHz depending the speed of the port. TXCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.

Table 8: GMII/MII Transmit Interface (Port 9) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
	82 83 84 85	P9_TXD[7] P9_TXD[6] P9_TXD[5] P9_TXD[4]	Typically Output	P9_TXD[7:0]/P9_MODE[2:0]—88E6095F device Only P9_TXD[3:0]/P9_MODE[2:0]—88E6092/88E6095 devices Only Transmit Data. TXD[7:0] outputs the data octet to be transmitted from the switch. TXD[7:0] is synchronous to GTXCLK in 1000BASE mode. In 100BASE and 10BASE modes TXD[3:0] is synchronous to TXCLK and TXD[7:4] is ignored. [redacted]. When reset is asserted, these pins become inputs and the configuration information below is latched at the rising edge of RESETn: [redacted] P9_MODE[2:0] sets Port 9's Mode of operation as follows: 000 = GMII with P9_GTXCLK = 125 MHz (1000BASE)[1] 001 = MII with P9_GTXCLK = 0 MHz (MII MAC Mode) 010 = MII with P9_GTXCLK = 25 MHz (100BASE PHY Mode) 011 = MII with P9_GTXCLK = 2.5 MHz (10BASE PHY Mode) 100 = Cross-chip SERDES port with all GMII/MII pins tri-stated - This mode configures the SERDES for 1000BASE-X operation, with Auto-Negotiation disabled. This mode is recommended for use when cascading multiple switch devices and is not recommended when connecting to PHY devices. 101 = 1000BASE-X with all GMII/MII pins tri-stated (except for P9_LED) 110 = SGMII or MGMII with all GMII/MII pins tri-stated. See "P10_GTXCLK/MGMII" pin description in Table 10, "GMII/MII Transmit Interface (Port 10)," on page 34 for MGMII mode details. 111 = Port 9 is disabled (with all GMII/MII pins tri-stated) SW_24P configures the device for cascading mode: • Cross-chip SERDES on all Ports 8, 9, and 10 interfaces, • Separate VLAN of Port 10 for CPU • Enabled configuration of Learn2All - see Learn2All section TXD[7:0] are internally pulled low via resistor so the pins can be left unconnected when they are not used. Use a 4.7 kohm resistor to VDDO for a configuration high.
73	87	[redacted]		
74	88	[redacted]		
75	89	[redacted]		
76	90	[redacted]		

Table 8: GMII/MII Transmit Interface (Port 9) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	91	P9_TXEN /P9_HALFDPX	Typically Output	Transmit Enable. Transmit enable is used to indicate when TXD[7:0] (or TDX[3:0] where appropriate) contains frame information. TXEN is synchronous to GTXCLK in 1000BASE mode and its synchronous to TXCLK in 100BASE and 10BASE modes. [redacted]. When reset is asserted, this pin become an input and the configuration information below is latched at the rising edge of RESETn: 0 = Sets Port 9 in full-duplex operation 1 = Sets Port 9 in half-duplex operation only if P9_MODE[2:0] = 01X TXEN is internally pulled low via resistor so the pin can be left unconnected to select full-duplex on Port 9 or when it is not used.

1. P9_MODE[2:0] = 000 is not supported in the 88E6092/88E6095 devices

Table 9: GMII/MII Receive Interface (Port 10)

88E6092/ 88e6095	88E6095F	Pin Name	Type	Description
114	138	P10_RXCLK	Input	Receive Clock. RXCLK is a reference for RXDV, RX_ER and RXD[7:0]. The speed of RXCLK is expected to be 125 MHz, 25 MHz or 2.5 MHz depending upon the speed of the port. RXCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.
125 124 123 122 120 119 118 117	150 149 148 147 144 143 142 141	P10_RXD[7] P10_RXD[6] P10_RXD[5] P10_RXD[4] P10_RXD[3] P10_RXD[2] P10_RXD[1] P10_RXD[0]	Input	Receive Data. RXD[7:0] receives the data octet or nibble to be sent into the switch. RXD[7:0] must be synchronous with RXCLK. In 1000BASE mode, RXD[7:0] is used. In 100BASE and 10BASE modes, RXD[3:0] is used and RXD[7:4] is ignored. RXD[7:0] are internally pulled low via resistor so the pins can be left unconnected when they are not used.
116	140	P10_RXDV	Input	Receive Data Valid. Receive Data Valid is used to indicate when RXD[7:0] (or RXD[3:0] where appropriate) contains frame information. RXDV must be synchronous to RXCLK. RXDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.
126	151	P10_RXER	Input	Receive Data Error. Receive Data Error is used to indicate when an external device (like a PHY) detects an error. When this signal is high the receive MAC will discard the frame and it will be counted as a frame with a CRC error. RXER must be synchronous to RXCLK. RXER is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 9: GMII/MII Receive Interface (Port 10) (Continued)

88E6092/ 88e6095	88E6095F	Pin Name	Type	Description
127	152	P10_CRS /P10_LED	Input or Output	Carrier Sense. Carrier sense is used to indicate that carrier has been detected on the line. CRS is not synchronous to RXCLK. CRS is used for half-duplex modes only and is ignored when the port is in full-duplex. When Port 10 is configured for 1000BASE-X operation (see P9/P10_MODE[2:0]) Px_LED indicates whether a 'Link' signal is detected by the Px_SDET input. A low output level indicates that a positive input signal is detected by Px_SDET and the port's PCS determines link. Px_LED indicates 'Activity' by blinking high when the port receives or transmits a frame. CRS is internally pulled low via resistor so the pin can be left unconnected when it is not used. The direction of this pin is determined by the value of the P10_MODE[2:0] pins at the rising edge of RESETn (during RESETn this pin is an input).
128	153	P10_COL /P10_SDET	Input	Collision. Collision is used to indicate both transmit and receive are occurring at the same time in half-duplex mode. The Collision signal is not synchronous to RXCLK. The Collision signal is used for half-duplex modes only and is ignored when the port is in full-duplex. When Port 10 is configured for 1000BASE-X operation (see P10_MODE[2:0]) Px_SDET is an input used to indicate whether a signal is detected by the fiber-optic transceiver. A positive level indicates that a signal is detected. P10_COL is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 10: GMII/MII Transmit Interface (Port 10)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	135	P10_GTXCLK /MGMII	Output	Transmit Clock. GTXCLK is a reference for TXEN and TXD[7:0] when the port is in GMII mode. The speed of GTXCLK is 125 MHz and is normally only driven when the speed of the port is 1000BASE. GTXCLK can be configured to output a 25 MHz or 2.5 MHz clock so it can be used as a clock source for P10_TXCLK and P10_RXCLK when no other clock sources are available (see P10_MODE[2:0]). [redacted] When reset is asserted, this pin becomes an input and the configuration information below is latched at the rising edge of RESETn: 0 = All SERDES ports run in SGMII mode 1 = All SERDES ports run in MGMII mode - MGMII mode is recommended when connecting to Marvell® PHY devices that support auto-media detect when connecting to SERDES ports. MGMII mode uses out of band communication for link, speed, and flow control status, and works specifically with Marvell PHYs that support MGMII. GTXCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.
112	136	P10_TXCLK	Input	Transmit Clock. TXCLK is a reference for TXEN and TXD[3:0] when the port is in MII mode. The speed of TXCLK is 25 MHz or 2.5 MHz depending upon the speed of the port. TXCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.

Table 10: GMII/MII Transmit Interface (Port 10) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
101	125	[redacted]	Typically Output	P10_TXD[7:0]/P10_MODE[2:0]/ADDR[4:0] Transmit Data. TXD[7:0] outputs the data octet to be transmitted from the switch. TXD[7:0] is synchronous to GTXCLK in 1000BASE mode. In 100BASE and 10BASE modes, TXD[3:0] is synchronous to TXCLK and TXD[7:4] is ignored. [redacted]. When reset is asserted, these pins become inputs and the configuration information is latched at the rising edge of RESETn: [redacted] P10_MODE[2:0] sets Port 10's mode of operation as follows: 000 = GMII with P10_GTXCLK = 125 MHz (1000BASE) 001 = MII with P10_GTXCLK = 0 MHz (MII MAC Mode) 010 = MII with P10_GTXCLK = 25 MHz (100BASE PHY Mode) 011 = MII with P10_GTXCLK = 2.5 MHz (10BASE PHY Mode) 100 = Cross-chip SERDES port with all GMII/MII pins tri-stated 101 = 1000BASE-X with all GMII/MII pins tri-stated (except for P10_LED) 110 = MGMII with all GMII/MII pins tri-stated - See "P10_GTXCLK/MGMII" pin description in Table 10, "GMII/MII Transmit Interface (Port 10)," on page 34 for MGMII mode details. 111 = Port 10 is disabled (with all GMII/MII pins tri-stated) ADDR[4:0] sets the device's SMI address. If ADDR[4:0] are all 0's the device is configured in single device addressing mode. TXD[7:0] are internally pulled low via resistor so the pins can be left unconnected when they are not used. Use a 4.7 kΩ resistor to VDDO for a configuration high.
102	126	[redacted]		
103	127	[redacted]		
104	128	[redacted]		
106	130	[redacted]		
107	131	[redacted]		
108	132	[redacted]		
109	133	[redacted]		

Table 10: GMII/MII Transmit Interface (Port 10) (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	134	P10_TXEN /P10_HALFDPX	Typically Output	Transmit Enable. Transmit enable is used to indicate when TXD[7:0] (or TXD[3:0] where appropriate) contains frame information. TXEN is synchronous to GTXCLK in 1000BASE mode and its synchronous to TXCLK in 100BASE and 10BASE modes. [redacted] When reset is asserted, this pin becomes an input and the configuration information is latched at the rising edge of RESETn: 0 = Sets Port 10 in full-duplex operation 1 = Sets Port 10 in half-duplex operation only if P10_MODE[2:0] = 01X TXEN is internally pulled low via resistor so the pin can be left unconnected to select full-duplex on Port 10 or when it is not used.

Table 11: Regulator & Reference

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
174	212	IREF	Analog	Current reference. A 2kΩ 1% resistor is placed between the IREF and VSS. This resistor is used to set an internal bias reference current.
67	77	CONTROL_12	Analog	Voltage control to external 1.2V regulator. This signal controls an external PNP transistor to generate the 1.2V power supply (from the 2.5V power rail) for the VDD_CORE pins. This pin is connected to the base of the PNP transistor, while the PNP transistor's collector will be connected to 1.2V and the emitter is connected to 2.5V
65	75	CONTROL_15	Analog	Voltage control to external 1.5V regulator. This signal controls an external PNP transistor to generate the 1.5V power supply (from the 2.5V power rail) for the VDDAH pins. This pin is connected to the base of the PNP transistor, while the PNP transistor's collector will be connected to 1.5V and the emitter is connected to 2.5V If this pin is not used, it must be left unconnected.
66	76	FEEDBACK_15	Analog	Voltage feedback for 1.5V regulator. This signal must be connected to the output of the external PNP transistor that is connected to the CONTROL_15 pin, if the CONTROL_15 pin is used. If the CONTROL_15 pin is not used, this pin must be left unconnected.

Table 12: System

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
70	80	XTAL_IN	Input	25 MHz system reference clock input provided from the board. The clock source can come from an external crystal or an external oscillator. This is the only clock required. Refer to Section 8.6.2 for XTAL_IN timing requirements.
69	79	XTAL_OUT	Output	System reference clock output provided to the board. This output can only be used to drive an external crystal. It cannot be used to drive external logic. If an external oscillator is used this pin should be left unconnected.
97	121	INTn	Open Drain Output	INTn is an active low, open drain pin that is asserted to indicate an unmasked interrupt event occurred. [redacted] The INTn pin will go active low if SW_MODE[1:0] (Table 15) is not 0b10 (stand alone mode) and if the EEPROM data has executed a Halt OpCode.
96	120	RESETn	Input	Hardware reset. Active low. The device is configured during reset. When RESETn is low, all configuration pins become inputs and the value seen on these pins is latched on the rising edge of RESETn or some time after. Refer to Section 8.6.1 for Reset and Configuration Timing details.

Table 13: Register Access Interface

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	118	MDC_PHY /PPU_EN /LRN2ALL	Typically Output	Management Data Clock, Master. MDC_PHY is the reference clock output for the serial management interface (SMI) that connects to an external SMI slave device, typically external PHYs. The Master SMI is used to access the registers in the external PHYs and it is available to the CPU to use as long as the PPU (PHY Polling Unit) is disabled. [redacted] When reset is asserted, MDC_PHY becomes an input and configuration information is latched into the device. If the SW_24P configuration pin (Table 8) is low at the rising edge of RESETn, the PHY Polling Unit's (PPU) mode is latched at the rising edge of RESETn as follows: Low = Disable the PHY Polling Unit High = Enable the PHY Polling Unit [redacted] Low = Do not send learning update frames to Distributed Switching Architecture (DSA) Tag ports High = Send learning update frames to DSA Tag ports MDC_PHY is internally pulled high via a resistor so it can be left floating when unused.
95	119	MDIO_PHY	I/O	Management Data I/O, Master. MDIO_PHY is used to transfer management data in and out of the device synchronously to MDC_PHY. [redacted] [redacted] Ω [redacted] Ω. The device uses Device Addresses 0x08, 0x09, and 0x0A to access the external PHYs for ports 8, 9, and 10 respectively. If the PPU (PHY Polling Unit) is disabled, all SMI activity from the CPU (see MDC_CPU and MDIO_CPU below) appears on the SMI Master pins. In this mode, all Device Addresses not used in the switch are available for use outside the switch. MDIO_PHY is internally pulled high via a resistor so it can be left floating when unused.

Table 13: Register Access Interface (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
99	123	MDC_CPU	Input	Management Data Clock, Slave. MDC_CPU is the reference clock input for the serial management interface (SMI) that connects to an external SMI master, typically a CPU. A continuous clock stream is not expected. The maximum frequency supported is 8.3 MHz. The CPU's SMI interface is used to access the registers in the switch and it is available in all combinations of SW_MODE[1:0] (Table 15). MDC_CPU is internally pulled high via a resistor so it can be left floating when unused.
98	122	MDIO_CPU	I/O	Management Data I/O, Slave. MDIO_CPU is used to transfer management data in and out of the device synchronously to MDC_CPU. [redacted] The device uses one or all of the 32 possible SMI port addresses (two modes are supported). The address(es) that are used are selectable using the P10_TXD[4:0]/ADDR[4:0] pins. MDIO_CPU is internally pulled high via a resistor so it can be left floating when unused.

Table 14: Serial EEPROM Interface

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
1[redacted]171		EE_CS /EE_1K	Typically Output	Serial EEPROM chip select. EE_CS is the serial EEPROM chip select referenced to EE_CLK. It is used to enable the external EEPROM (if present), and to delineate each data transfer. [redacted] When reset is asserted, EE_CS becomes an input and the desired EEPROM type configuration is latched at the rising edge of RESETn as follows: Low = Use 8-bit addresses (for 2K bit 93C56 & 4K bit 93C66) High = Use 6-bit addresses (for 1K bit 93C46) The external EEPROM must be configured in the x16 organization. EE_CS is internally pulled high via a resistor so the pin can be left floating when unused or to select support for 1K bit devices. Use a 4.7 kΩ resistor to VSS for a configuration low.

Table 14: Serial EEPROM Interface (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	170	EE_CLK /FD_FLOW_DIS	Typically Output	Serial EEPROM clock. EE_CLK is the serial EEPROM clock reference output by the device. It is used to shift the external serial EEPROM (if installed) to the next data bit so the default values of the internal registers can be overridden. [redacted] When reset is asserted, EE_CLK becomes an input and it becomes the full-duplex Flow Control disable as follows: Low = Enable advertisement of full-duplex flow control on all PHYs High = Disable full-duplex flow control on all full-duplex ports Full-duplex flow control requires support from the end station. It is supported on any full-duplex port that has Auto-Negotiation enabled, advertises that it supports Pause (i.e., FD_FLOW_DIS = Low at reset), and sees that the end station also supports Pause (from data returned during Auto-Negotiation). EE_CLK is internally pulled high via a resistor so the pin can be left unconnected for a configuration high. Use a 4.7 kΩ resistor to VSS for a configuration low.

Table 14: Serial EEPROM Interface (Continued)

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	169	EE_DIN /HD_FLOW_DIS	Typically Input Output During EEPROM Loading	Serial EEPROM data into the EEPROM device. EE_DIN is serial EEPROM data referenced to EE_CLK used to transmit the EEPROM command and address to the external serial EEPROM (if present). [redacted] Half-duplex flow control is active on all half-duplex ports whenever this pin is low. HD_FLOW_DIS is latched after reset. EE_DIN is internally pulled high via a resistor so it can be left floating to disable half-duplex flow control or when the pin is unused. Use a 4.7 kΩ resistor to VSS for a configuration low.
133	166	EE_DOUT	Input	Serial EEPROM data out from the EEPROM device. EE_DOUT is serial EEPROM data referenced to EE_CLK used to receive the EEPROM configuration data from the external serial EEPROM (if present). EE_DOUT is internally pulled high via a resistor so it can be left floating when the pin is unused.

Table 15: Switch Configuration Interface

88E6092/ 88E6095	88E6095F	Pin Name	Type	Description
[redacted]	156 157	SW_MODE[1] SW_MODE[0]	Input	[redacted]. Switch Mode pins work as follows: 1 0 Description 0 0 CPU attached mode – ports come up disabled[1] 0 1 Reserved 1 0 Stand alone mode – ports come up enabled – ignore EEPROM 1 1 EEPROM attached mode – EEPROM defined port states[2] The EEPROM attached mode (when both SW_MODE pins = high) can be used with a CPU and the CPU attached mode can be used with an EEPROM. In these modes, the INTn pin will go active low after the EEPROM is done initializing the internal registers. SW_MODE[1:0] are not latched on the rising edge of RESETn and they must remain static for proper device operation. They are internally pulled high via resistors so the pins can be left unconnected to enable the EEPROM attached mode.

1. The ports come up disabled in the CPU mode so the CPU can perform bridge loop detection on link up.
2. In EEPROM attached mode, the ports come up enabled unless the Port Control register is overwritten by the EEPROM data (see Section 6.7 for EEPROM programming format details).

Table 16: Power & Ground

88E6092/ 88E6095	88E6095F	PIN NAME	Type	Description
91	115	VDDO_PHY	Power	Power to the MDC_PHY and MDIO_PHY outputs (pin numbers: 91–95 in the 88E6092/88E6095 devices; 115–119 in the 88E6095F device). VDDO_PHY can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
72 80	86 95 104 106	VDDO_P9	Power	Power to Port 9's GMII/MII outputs (pin numbers: 71–90 in the 88E6092/88E6095 devices; 81–114 in the 88E6095F device). VDDO_P9 can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
105 113 121	129 137 145 146	VDDO_P10	Power	Power to Port 10's GMII/MII outputs (pin numbers: 96–128 in the 88E6092/88E6095 devices; 120–153 in the 88E6095F device). VDDO_P10 can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
135	168	VDDO_ EEPROM	Power	Power to the serial EEPROM outputs (pin numbers: 140–173 in the 88E6092/88E6095 devices; 173–206 in the 88E6095F device). VDDO_EEPROM can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
145 154 162 170	178 187 195 203	VDDO_LED	Power	Power to LED outputs (pin numbers: 129–139 in the 88E6092/88E6095 devices; 154–172 in the 88E6095F device) VDDO_LED can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
	54	VDDO_SLED	Power	Power to serial LED outputs (pin numbers: 51–54 in the 88E6095Fdevice). VDDO_SLED can be connected to 3.3V to support 3.3V I/O or 2.5V for 2.5V I/O.
68	78	VDD_PLL	Power	2.5V Power to the internal PLL blocks inside the device.
48 54 60	58 64 70	P8_VDDAH P9_VDDAH P10_VDDAH	Power	2.5V Power to analog core used to power the SERDES interface for Port 8, Port 9 and Port 10. If the port's SERDES is not used, Px_VDDAH can be tied to VSS to save power. If Port 9 or Port 10 is configured as a digital interface, P9_VDDAH and/or P10_VDDAH must be connected to 2.5V, respectively.

Table 16: Power & Ground (Continued)

88E6092/ 88E6095	88E6095F	PIN NAME	Type	Description
49 55 61	59 65 71	P8_VTT P9_VTT P10_VTT	Power	1.2V or 1.5V Power to analog core used to power each SERDES transmitter. Refer to Table 144 or Table 145 - “Transmitter DC Characteristics” for SERDES output drive levels and how they are affected by Px_VTT. Use 1.2V power for chip-to-chip connections. Use 1.5V power for 1000BASE-X or backplane applications. For lowest power, P9_VTT and/or P10_VTT may be left floating, respectively.
11 22 33 44 176	1 12 23 34 46 211 216	VDDAH	Power	2.5V Power to the analog core used to power each 10/100 PHY interface.
3 8 14 19 25 30 36 41	4 9 15 20 26 31 37 42	VDDAL	Power	1.2V Power to analog core used to power each 10/100 PHY interface.
45 64 71 78 85 100 115 134 139 146 151 158 168	45 55 74 81 92 100 105 124 139 167 172 179 184 191 201	VDD_CORE	Power	1.2V Power to digital core

Table 16: Power & Ground (Continued)

88E6092/ 88E6095	88E6095F	PIN NAME	Type	Description
EPAD	EPAD	VSS	Ground	Ground to device. The 88E6092/88E6095 devices are contained in the 176 TQFP package, while the 88E6095F device is contained in a 216 LQFP package. Both the 176 TQFP and 216 LQFP packages have an exposed die pad (E-PAD) at their base. The EPAD must be soldered to VSS as it is the only VSS connection on the device. The location and dimensions of the EPAD can be found in Figure 62 and Table 164 for the 176 TQFP, and Figure 63 and Table 165 for the 216 LQFP, respectively.
129 130 175	107 108 154 155 213 214 215	NC		No Connect. Do not connect these pins to anything.

1.2 88E6092 and 88E6095 Device Pin Assignment List

Table 17: Package Pin List—Alphabetical by Signal Name

Pin Number	Pin Name
140	CONFIG_A
141	CONFIG_B
67	CONTROL_12
65	CONTROL_15
137	EE_CLK/FD_FLOW_DIS
138	EE_CS/EE_1K
136	EE_DIN/HD_FLOW_DIS
133	EE_DOUT
66	FEEDBACK_15
97	INTn
174	IREF
99	MDC_CPU
94	MDC_PHY/PPU_EN/LRN2ALL
98	MDIO_CPU
95	MDIO_PHY
129	NC
130	NC
175	NC
142	P0_LED0
143	P0_LED1
144	P0_LED2
2	P0_RXN
1	P0_RXP
5	P0_TXN
4	P0_TXP
147	P1_LED0
148	P1_LED1
149	P1_LED2
9	P1_RXN
10	P1_RXP
6	P1_TXN
7	P1_TXP
150	P2_LED0
152	P2_LED1
153	P2_LED2
13	P2_RXN
12	P2_RXP
16	P2_TXN
15	P2_TXP
155	P3_LED0
156	P3_LED1
157	P3_LED2
20	P3_RXN
21	P3_RXP
17	P3_TXN
18	P3_TXP
159	P4_LED0
160	P4_LED1
161	P4_LED2
24	P4_RXN
23	P4_RXP
27	P4_TXN

Pin Number	Pin Name
26	P4_TXP
163	P5_LED0
164	P5_LED1
165	P5_LED2
31	P5_RXN
32	P5_RXP
28	P5_TXN
29	P5_TXP
166	P6_LED0
167	P6_LED1
169	P6_LED2
35	P6_RXN
34	P6_RXP
38	P6_TXN
37	P6_TXP
171	P7_LED0
172	P7_LED1
173	P7_LED2
42	P7_RXN
43	P7_RXP
39	P7_TXN
40	P7_TXP
92	P8_LED/P8_MODE
46	P8_RXN
47	P8_RXP
93	P8_SDET
50	P8_TXN
51	P8_TXP

Pin Number	Pin Name
48	P8_VDDAH
49	P8_VTT
90	P9_COL/P9_SDET
89	P9_CRS/P9_LED
81	P9_RXCLK
83	P9_RXD[0]
84	P9_RXD[1]
86	P9_RXD[2]
87	P9_RXD[3]
82	P9_RXDV
88	P9_RXER
52	P9_RXN
53	P9_RXP
79	P9_TXCLK
76	P9_TXD[0]/P9_MODE[0]
75	P9_TXD[1]/P9_MODE[1]
74	P9_TXD[2]/P9_MODE[2]
73	P9_TXD[3]/SW_24P
77	P9_TXEN/P9_HALFDPX
56	P9_TXN
57	P9_TXP
54	P9_VDDAH
55	P9_VTT
128	P10_COL/P10_SDET
127	P10_CRS/P10_LED
111	P10_GTXCLK/MGMII
114	P10_RXCLK
117	P10_RXD[0]

Pin Number	Pin Name
118	P10_RXD[1]
119	P10_RXD[2]
120	P10_RXD[3]
122	P10_RXD[4]
123	P10_RXD[5]
124	P10_RXD[6]
125	P10_RXD[7]
116	P10_RXDV
126	P10_RXER
58	P10_RXN
59	P10_RXP
112	P10_TXCLK
109	P10_TXD[0]/P10_MODE[0]
108	P10_TXD[1]/P10_MODE[1]
107	P10_TXD[2]/P10_MODE[2]
106	P10_TXD[3]/ADDR[0]
104	P10_TXD[4]/ADDR[1]
103	P10_TXD[5]/ADDR[2]
102	P10_TXD[6]/ADDR[3]
101	P10_TXD[7]/ADDR[4]
110	P10_TXEN/P10_HALFDPX
62	P10_TXN
63	P10_TXP
60	P10_VDDAH
61	P10_VTT
96	RESETn
132	SW_MODE[0]
131	SW_MODE[1]

Pin Number	Pin Name
45	VDD_CORE
64	VDD_CORE
71	VDD_CORE
78	VDD_CORE
85	VDD_CORE
100	VDD_CORE
115	VDD_CORE
134	VDD_CORE
139	VDD_CORE
146	VDD_CORE
151	VDD_CORE
158	VDD_CORE
168	VDD_CORE
68	VDD_PLL
11	VDDAH
22	VDDAH
33	VDDAH
44	VDDAH
176	VDDAH
3	VDDAL
8	VDDAL
14	VDDAL
19	VDDAL
25	VDDAL
30	VDDAL
36	VDDAL
41	VDDAL
135	VDDO_EEPROM

Pin Number	Pin Name
145	VDDO_LED
154	VDDO_LED
162	VDDO_LED
170	VDDO_LED
72	VDDO_P9
80	VDDO_P9
105	VDDO_P10
113	VDDO_P10
121	VDDO_P10
91	VDDO_PHY
EPAD	VSS
70	XTAL_IN
69	XTAL_OUT

1.3 88E6095F Device Pin Assignment List

Table 18: Package Pin List—Alphabetical by Signal Name

Pin Number	Pin Name
173	CONFIG_A
174	CONFIG_B
77	CONTROL_12
75	CONTROL_15
170	EE_CLK/FD_FLOW_DIS
171	EE_CS/EE_1K
169	EE_DIN/HD_FLOW_DIS
166	EE_DOUT
76	FEEDBACK_15
121	INTn
212	IREF
53	LEDCLK
51	LEDENA
52	LEDSER
123	MDC_CPU
118	MDC_PHY/PPU_EN/LRN2ALL
122	MDIO_CPU
119	MDIO_PHY
107	NC
108	NC
154	NC
155	NC
213	NC
214	NC
215	NC
165	P0_CONFIG

Pin Number	Pin Name
175	P0_LED0
176	P0_LED1
177	P0_LED2
3	P0_RXN
2	P0_RXP
207	P0_SDET
6	P0_TXN
5	P0_TXP
164	P1_CONFIG
180	P1_LED0
181	P1_LED1
182	P1_LED2
10	P1_RXN
11	P1_RXP
208	P1_SDET
7	P1_TXN
8	P1_TXP
163	P2_CONFIG
183	P2_LED0
185	P2_LED1
186	P2_LED2
14	P2_RXN
13	P2_RXP
209	P2_SDET
17	P2_TXN
16	P2_TXP

Pin Number	Pin Name
162	P3_CONFIG
188	P3_LED0
189	P3_LED1
190	P3_LED2
21	P3_RXN
22	P3_RXP
210	P3_SDET
18	P3_TXN
19	P3_TXP
161	P4_CONFIG
192	P4_LED0
193	P4_LED1
194	P4_LED2
25	P4_RXN
24	P4_RXP
47	P4_SDET
28	P4_TXN
27	P4_TXP
160	P5_CONFIG
196	P5_LED0
197	P5_LED1
198	P5_LED2
32	P5_RXN
33	P5_RXP
48	P5_SDET
29	P5_TXN
30	P5_TXP
159	P6_CONFIG

Pin Number	Pin Name
199	P6_LED0
200	P6_LED1
202	P6_LED2
36	P6_RXN
35	P6_RXP
49	P6_SDET
39	P6_TXN
38	P6_TXP
158	P7_CONFIG
204	P7_LED0
205	P7_LED1
206	P7_LED2
43	P7_RXN
44	P7_RXP
50	P7_SDET
40	P7_TXN
41	P7_TXP
116	P8_LED/P8_MODE
56	P8_RXN
57	P8_RXP
117	P8_SDET
60	P8_TXN
61	P8_TXP
58	P8_VDDAH
59	P8_VTT
110	P9_COL/P9_SDET
109	P9_CRS/P9_LED
93	P9_GTXCLK

Pin Number	Pin Name
96	P9_RXCLK
98	P9_RXD[0]
99	P9_RXD[1]
101	P9_RXD[2]
102	P9_RXD[3]
111	P9_RXD[4]
112	P9_RXD[5]
113	P9_RXD[6]
114	P9_RXD[7]
97	P9_RXDV
103	P9_RXER
62	P9_RXN
63	P9_RXP
94	P9_TXCLK
90	P9_TXD[0]/P9_MODE[0]
89	P9_TXD[1]/P9_MODE[1]
88	P9_TXD[2]/P9_MODE[2]
87	P9_TXD[3]/SW_24P
85	P9_TXD[4]
84	P9_TXD[5]
83	P9_TXD[6]
82	P9_TXD[7]
91	P9_TXEN/P9_HALFDPX
66	P9_TXN
67	P9_TXP
64	P9_VDDAH
65	P9_VTT
153	P10_COL/P10_SDET

Pin Number	Pin Name
152	P10_CRS/P10_LED
135	P10_GTXCLK/MGMII
138	P10_RXCLK
141	P10_RXD[0]
142	P10_RXD[1]
143	P10_RXD[2]
144	P10_RXD[3]
147	P10_RXD[4]
148	P10_RXD[5]
149	P10_RXD[6]
150	P10_RXD[7]
140	P10_RXDV
151	P10_RXER
68	P10_RXN
69	P10_RXP
136	P10_TXCLK
133	P10_TXD[0]/P10_MODE[0]
132	P10_TXD[1]/P10_MODE[1]
131	P10_TXD[2]/P10_MODE[2]
130	P10_TXD[3]/ADDR[0]
128	P10_TXD[4]/ADDR[1]
127	P10_TXD[5]/ADDR[2]
126	P10_TXD[6]/ADDR[3]
125	P10_TXD[7]/ADDR[4]
134	P10_TXEN/P10_HALFDPX
72	P10_TXN
73	P10_TXP
70	P10_VDDAH

Pin Number	Pin Name
71	P10_VTT
120	RESETn
157	SW_MODE[0]
156	SW_MODE[1]
45	VDD_CORE
55	VDD_CORE
74	VDD_CORE
81	VDD_CORE
92	VDD_CORE
100	VDD_CORE
105	VDD_CORE
124	VDD_CORE
139	VDD_CORE
167	VDD_CORE
172	VDD_CORE
179	VDD_CORE
184	VDD_CORE
191	VDD_CORE
201	VDD_CORE
78	VDD_PLL
1	VDDAH
12	VDDAH
23	VDDAH
34	VDDAH
46	VDDAH
211	VDDAH
216	VDDAH
4	VDDAL

Pin Number	Pin Name
9	VDDAL
15	VDDAL
20	VDDAL
26	VDDAL
31	VDDAL
37	VDDAL
42	VDDAL
168	VDDO_EEPROM
178	VDDO_LED
187	VDDO_LED
195	VDDO_LED
203	VDDO_LED
86	VDDO_P9
95	VDDO_P9
104	VDDO_P9
106	VDDO_P9
129	VDDO_P10
137	VDDO_P10
145	VDDO_P10
146	VDDO_P10
115	VDDO_PHY
54	VDDO_SLED
EPAD	VSS
80	XTAL_IN
79	XTAL_OUT

Section 2. Application Examples

2.1 Examples using the 88E6095 Device

Figure 3: 16 FE + 2 GE Ports

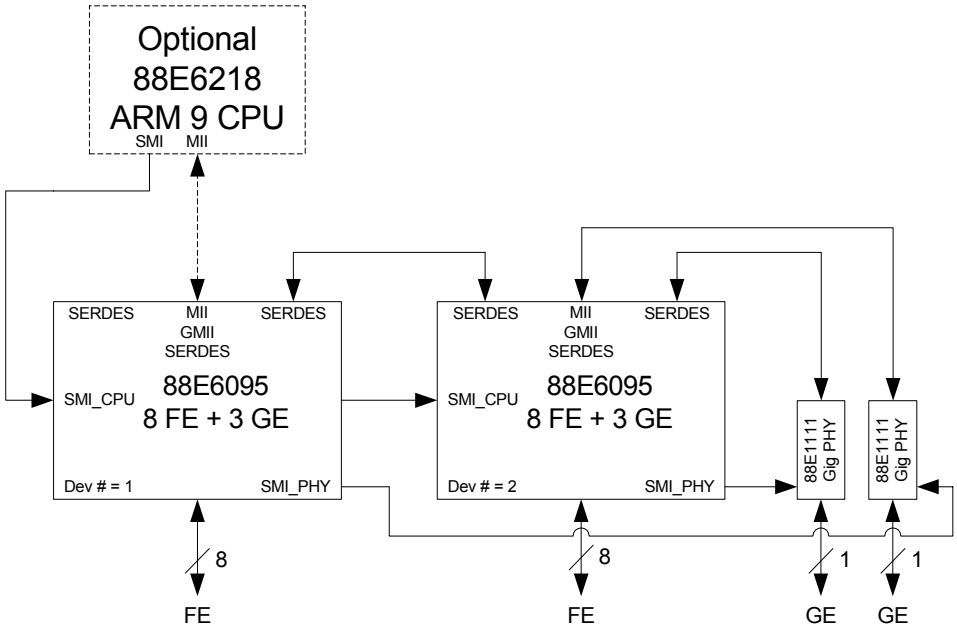

Figure 4: 24 FE + 2 GE with 88E6095 X 3

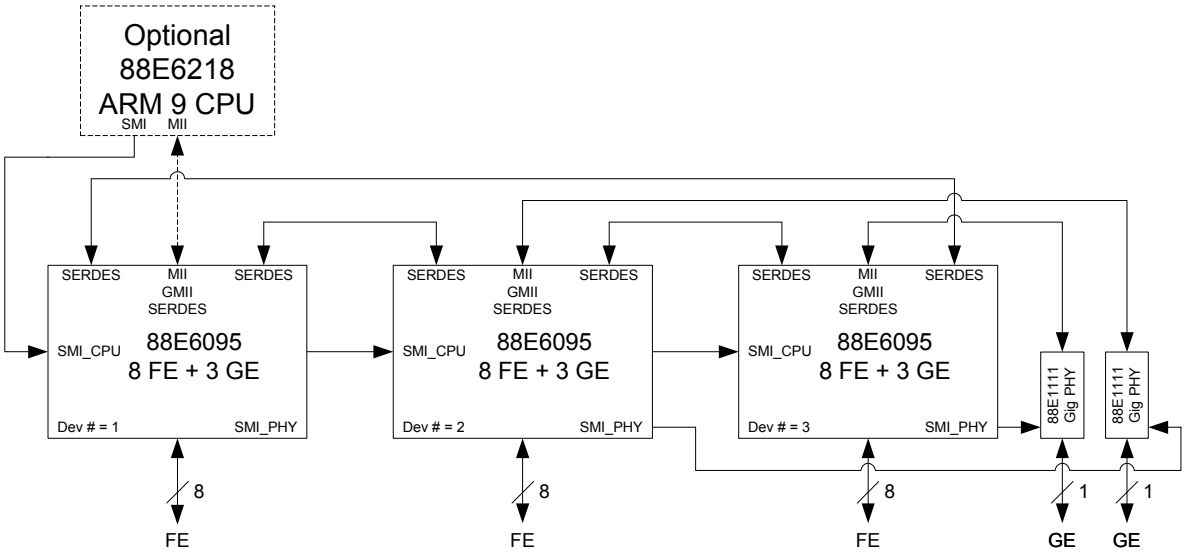

Figure 5: 48 FE + 4 GE wit 88E6095 X 6 + 88E6185

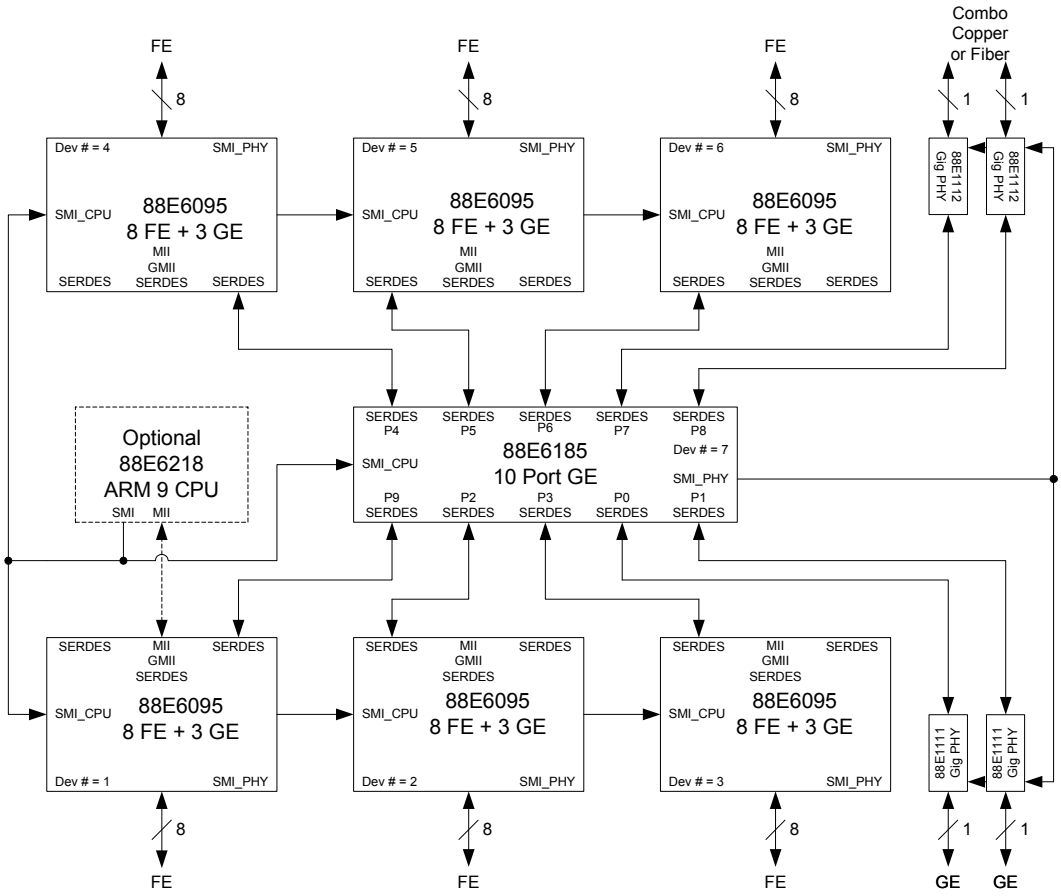

Figure 6: Fiber to the Curb

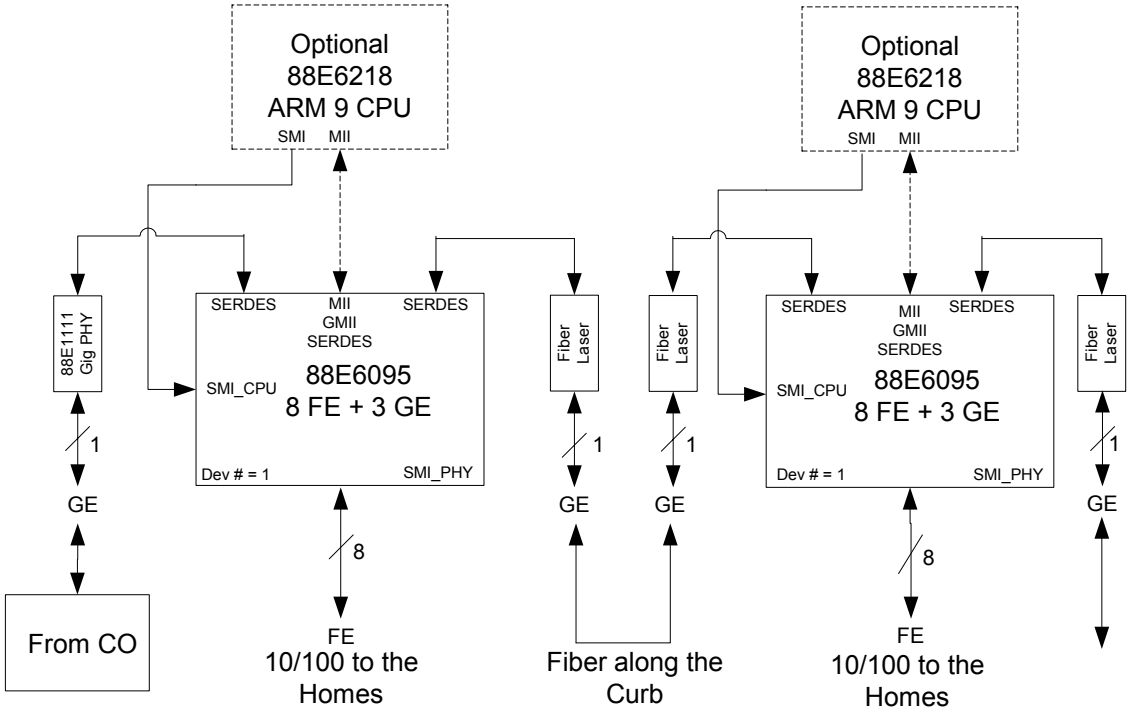

Figure 7: SOHO/SMB Router with FE and GE Ports

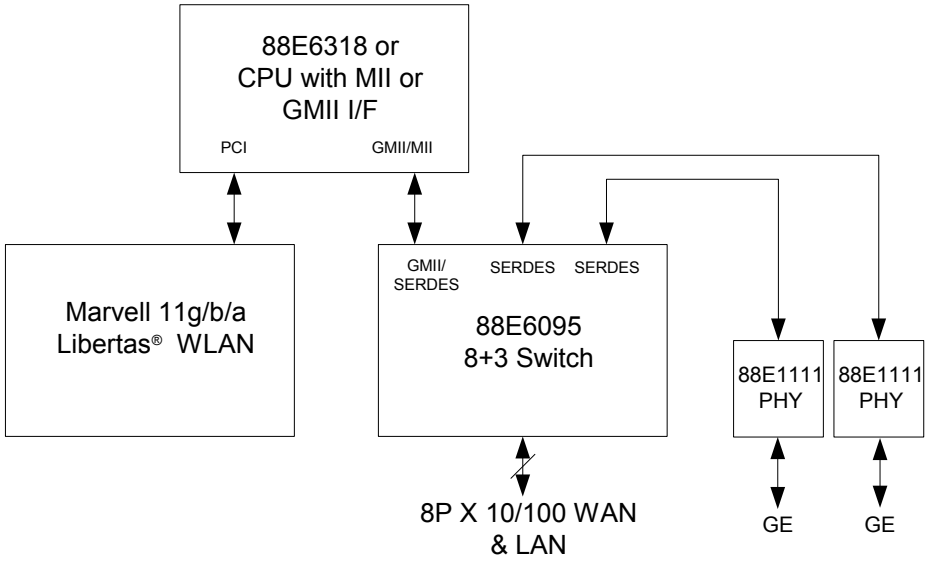

2.2 Device Physical Interfaces

The devices contain a number of interfaces that support both copper and fiber media. Table 19 lists the interfaces supported on each port of the 88E6092 and 88E6095 devices. Table 20 lists the interfaces supported on each port of the 88E6095F device. Refer to the diagrams further in this section for connection details.

Table 19: 88E6092/88E6095 Device Interfaces

Port	10BASE-T 100BASE-T[1]	1000BASE-T (w/ external PHY)	1000BASE-X SERDES	MAC or PHY Mode MII Only	MAC or PHY Mode GMII/MII
0-7	x				
8	x	x	x		
9	x	x	x	x	
10	x	x	x		x

1. Ports 8, 9, and 10 require an external PHY for these modes of operation.

Table 20: 88E6095F Device Interfaces

Port	10BASE-T 100BASE-T[1]	1000BASE-T (w/ external PHY)	1000BASE-X SERDES	100BASE-FX	MAC or PHY Mode GMII/MII
0-7	x			x	
8	x	x	x		
9	x	x	x		x
10	x	x	x		x

1. Ports 8, 9, and 10 require an external PHY for these modes of operation.

2.2.1 10/100 PHY Interface

Ports 0 to 7 on the devices support a 10/100 PHY interface. In the 88E6092/88E6095 devices this interface supports 10BASE-T and 100BASE-TX copper IEEE standards. Ports 0 to 7 on the 88E6095F device support a 100BASE-FX fiber option as well. The MAC inside the switch works the same way regardless of the external interface being used. Each PHY's Link, Speed, Duplex and Flow Control information is directly communicated to the MAC it is attached to so the MAC tracks, or follows, the mode the PHY links up in. A detailed description of the PHY is covered in Section 4, and the PHY registers are covered in Section 7.

2.2.2 SERDES or (G)MII Interface

Ports 8 to 10 in the devices are SERDES interfaces. The SERDES interfaces can be used for these options

- Connection to Marvell® triple speed 10/100/1000 copper PHYs
- Connection to 1000BASE-X fiber modules
- Cross-chip connection to other Marvell switch devices (another 88E6095 device) - i.e., cross-chip connection

The tenth and eleventh ports (Port 9 and Port 10), on the devices also support optional short-distance, industry-standard digital interfaces, referred to generically as the port's (G)MII or digital interface. Many interface modes and timings are supported so that a large number of external device types can be used. The 88E6092/ 88E6095 devices support an MII interface on Port 9 and an (G)MII interface on Port 10, while the 88E6095F device supports a (G)MII interface on Port 9 and Port 10. Refer to Section 2.2.3, "Digital Interface Options," on page 63 for details.

2.2.2.1 Triple Speed PHY SERDES Interface Option

Port 8, 9 and 10's SERDES can be configured to use a triple speed PHY interface to an external PHY. In this mode, the SERDES use the SGMII protocol. The in-band Link, Speed and Duplex signals in the SGMII protocol are ignored. The external PHY's Link, Speed, Duplex and Flow Control information must be transferred to the port's MAC so the MAC is in the correct mode. This can be done in software or automatically by enabling the PHY Polling Unit (PPU - Section 2.2.4)

The triple speed PHY interface can support Marvell PHYs with Auto-Media Detect™ for auto switching between copper and fiber. This can be supported in software or automatically in hardware by enabling the PHY Polling Unit (PPU) and by setting the port's MGMII bit to a one (in the port's Port Status Register - offset 0x00). If the port's MGMII bit is not set to a one, the PPU will support copper only and will not support Auto-Media Detect.

Port 8 is set to Triple Speed PHY mode if its Px_LED/Px_MODE pin is low at the rising edge of RESETn and if the port's PHYDetect bit is a 1 (in the port's Port Status Register - offset 0x00). PHYDetect will be set to a 1 on port 8 if the PPU is enabled (Section) and the PPU finds a PHY at SMI address 0x08.

Ports 9 and 10 are set to Triple Speed PHY mode if their Px_TXD[2:0]/Px_MODE[2:0] pins are set to 0x6 at the rising edge of RESETn. The automatic transfer of PHY status to its MAC requires that the port's PHYDetect bit be a 1 (in the port's Port Status Register - offset 0x00). PHYDetect will be set to a 1 on port 9 if the PPU is enabled (Section 2.2.4) and the PPU finds a PHY at SMI address 0x09. It will be set to a 1 on port 10 if the PPU finds a PHY at SMI address 0x0A.

2.2.2.2 IEEE 1000BASE-X SERDES Interface Option

Port 8, 9 and 10's SERDES can be configured in 1000BASE-X.

Port 8 is set to 1000BASE-X mode if its Px_LED/Px_MODE pin is high at the rising edge of RESETn. Connecting an LED to VDDO through a resistor to this pin will enable 1000BASE-X on the port. After RESETn the Px_LED/Px_MODE pin will become the Link/Activity LED for the port (off = no link, on = link, blink = activity).

Ports 9 and 10 are set to 1000BASE-X mode if their Px_TXD[2:0]/Px_MODE[2:0] pins are set to 0x5 at the rising edge of RESETn. In this mode, the Px_CRS/Px_LED pin become the Link/Activity LED for the port (off = no link, on = link, blink = activity) after RESETn.

The port enters 1000BASE-X mode, if configured, even if an external PHY is detected at the port's SMI address.

1000BASE-X mode uses a PCS to auto-negotiate with a link partner to determine if Flow Control should be supported or not (auto-negotiation can be disabled). Link will be automatically established if the port's Px_SDET is detected high and the port's PCS determines Sync is OK (sets the port's SyncOK to a 1). Link will automatically go down if either Px_SDET or SyncOK go to zero. Speed is always 1000 Mbps and Duplex is always full-duplex on 1000BASE-X ports.

2.2.2.3 Cross-chip Interface Option

Port 8, 9, and 10's SERDES can be configured in cross-chip mode. This mode is used to connect two or more Marvell® switch devices together to create a larger switch. 1000BASE-X protocol is used on the line and the port's Speed is locked to 1000 Mbps and its Duplex is fixed at full-duplex.

Port 8 is set to cross-chip mode if its Px_LED/Px_MODE pin is low at the rising edge of RESETn and if the port's PHYDetect bit is a 0 (in the port's Port Status Register - offset 0x00). PHYDetect will be set to a 0 on port 8 if the PPU is disabled (Section) or if the PPU cannot find a PHY at SMI address 0x08.

Note

Ports that are configured in cross-chip mode are initialized with their Link down (unless the SW_24P configuration option is being used - Section 2.2.5). This gives software time to initialize the switch's

configuration before software allows packets to flow by forcing the port's link up (using the port's ForcedLink bit in the port's PCS Control register - offset 0x01).

2.2.2.4 Port Status Registers

Each switch port of the devices has a status register that reports information about that port's MAC, SERDES or (G)MII interface. These registers can be used to check the current port configuration. See Table 61, Port Status Register, for more information.

2.2.3 Digital Interface Options

The (G)MII digital interface supports many different modes defined in the following sections. The mode to use is configured once at reset by external pull-up resistors connected to the P9_MODE[2:0] and P10_MODE[2:0] pins. If Port 9 or Port 10 is not connected to any device, the port should be disabled. If Port 9 or Port 10's(G)MII digital pins are unused but the port's SERDES interface is used, the P9_MODE[2:0] or P10_MODE[2:0] pins must be set accordingly. See Table 8 and Table 10 for more information.

For the 88E6092/88E6095 devices, Port 9 supports MII MAC and MII PHY Mode.

For the 88E6092 and 88E6095 devices Port 10 supports MII MAC mode, MII PHY mode and (G)MII mode options. For the 88E6095F device, Port 9 and Port 10 supports MII MAC mode, MII PHY mode and (G)MII mode options.

Note

(G)MII PHY mode and (G)MII MAC mode are discussed in the following sections. Electrically, there is no difference since the GMII Interface uses source synchronous clocks. Each concept is discussed separately since the port supports being connected to an external PHY (GMII MAC mode - where the port looks like a MAC supporting 10/100/1000 Mbps) or to an external MAC (GMII PHY mode where the port looks like a PHY supporting 1000 Mbps only).

Warning

A port's SERDES must be powered up to generate the GTX_CLK output for MII PHY modes of operation.

2.2.3.1 MII MAC Mode

The MII MAC Mode, sometimes called 'Forward MII', configures Port 9 or Port 10's MAC inside the devices to act as a MAC so it can be directly connected to an external MII-based PHY. In this mode, the devices receive the interface clocks (Px_TXCLK and Px_RXCLK) from the PHY and will work at any frequency from DC to 50 MHz (50 MHz supports 200 Mbps in each direction). The two clocks can be asynchronous with each other. Both full- and half-duplex modes are supported and need to be selected to match the mode of the link partner's MAC. The MII MAC mode is compliant with IEEE 802.3 clause 22. (**Note:** The MII requires only four data bits in each direction so only the lower four data bits are used). P9_MODE or P10_MODE should be set correctly at reset (see Table 8 and Table 10) to select this configuration and the PHY's SMI address must be set to 0x09 for Port 9 or 0x0A for Port 10 for auto-negotiation to operate correctly.

Figure 8: MII MAC Interface Pins

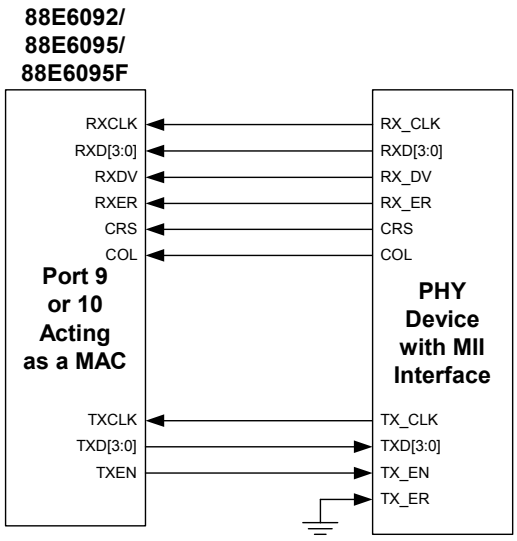

2.2.3.2 MII PHY Mode

The MII PHY Mode, sometimes called ‘Reverse MII’, configures Port 9 or Port 10’s MAC inside the devices to act as a PHY so that it can be directly connected to an external MAC. In this mode, the devices drive the interface clocks (RXCLK and TXCLK for both MACs) from its GTXCLK pin so the appropriate GTXCLK frequency must be selected. For the 88E6092/88E6095 devices, the GTX_CLK pin is only available on Port 10. The GTX_CLK pin is available on both Port 9 and Port 10 of the 88E6095F device. GTX_CLK is used as a generic asynchronous clock source, but it is recommended that there are not more than four loads on GTX_CLK. For connection to more than four loads (that is, connection for use with Port 9), buffer Port 10s GTX_CLK, or use a generic oscillator. Only full-duplex modes are supported (since CRS and COL are not driven by the devices outputs) and must match the mode of the link partner’s MAC. The MII PHY mode is compliant with IEEE 802.3 clause 22 in full-duplex mode (**Note**: The MII requires only four data bits in each direction so only the lower four data bits on the devices are used). At reset, P9_MODE and P10_MODE should be set for the appropriate speed —see Table 8 and Table 10. In this mode, there is no external PHY for Port 9 or Port 10, and so Port 9 or Port 10 is skipped by the PPU. In Reverse MII mode initially, the link status is down requiring the system software to force the port’s link up to enable the port.

Figure 9: MII PHY Interface Pins

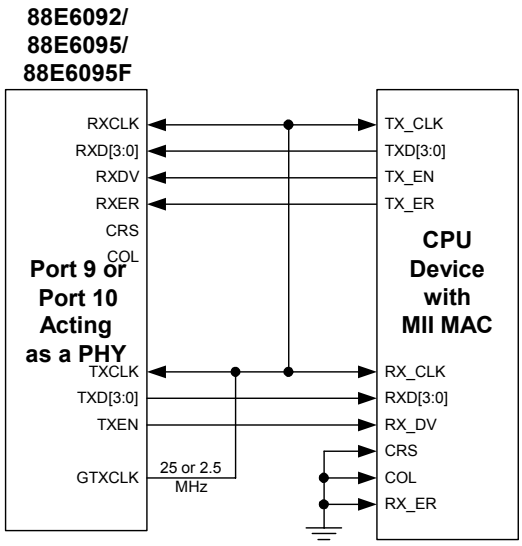

2.2.3.3 GMII MAC Mode

The GMII MAC Mode, sometimes called 'Forward GMII', configures Port 10's MAC (also Port 9's MAC for the 88E6095F device) inside the devices to act as a gigabit MAC (GMAC) so that it can be directly connected to an external GMII-based Gigabit PHY. In this mode, the devices receive the interface clocks (TXCLK and RXCLK) from the PHY but generates GTXCLK for the PHY. 10 Mbps, 100 Mbps or 1000 Mbps is supported in this configuration. Full- and half-duplex modes are supported at 10 Mbps or 100 Mbps. Full-duplex is supported at 1000 Mbps. The speed and mode in the external PHY's auto-negotiation must be restricted from advertising the 1000BASE, half-duplex case as the GMAC inside the devices do not support that mode. This is done automatically if the PHY Polling Unit (PPU) inside the devices is enabled. GMII MAC mode is compliant with IEEE 802.3 clause 28. P9_MODE and P10_MODE should be set to GMII mode at reset (see Table 8 and Table 10) for this configuration and the PHY's SMI address must be set to 0x09 for Port 9 or 0x0A for Port 10 for auto-negotiation to operate correctly.

A triple speed interface is supported in GMII MAC mode (i.e., 10, 100 and 1000). When the PHY completes auto-negotiation and brings the link up the auto-negotiated speed, duplex and flow control information must be moved from the PHY to the MAC so the MAC matches the PHY's settings. This is done automatically if the PPU is enabled. If the PPU is disabled software needs to move this information by reading the PHY's registers and then forcing the MAC into the appropriate mode (see PCS Control Register - offset 0x01). The interface pins will track the speed that the MAC is set to.

Figure 10: GMII MAC Interface Pins

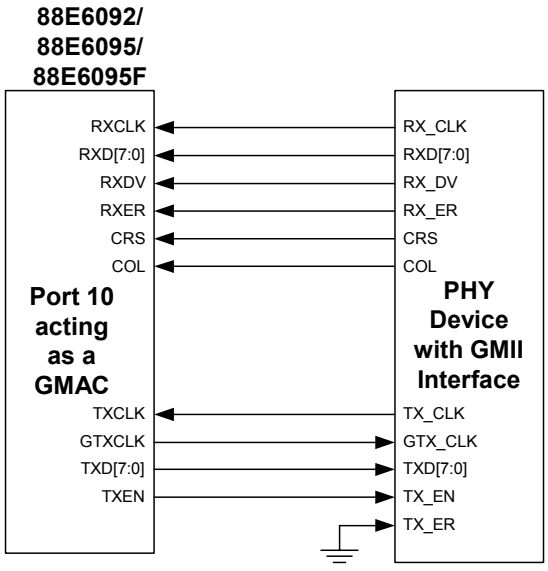

2.2.3.4 GMII PHY Mode

The GMII PHY Mode, sometimes called ‘Reverse GMII’, configures Port 10’s MAC (also Port 9’s MAC for the 88E6095F device) inside the devices to act as a gigabit PHY so that it can be directly connected to an external GMAC. In this mode, the devices drive the transmit interface clock (GTXCLK) and accept the receive interface clock (RXCLK). Only gigabit full-duplex mode is supported and must match the mode of the link partner’s GMAC. GMII PHY mode is compliant with IEEE 802.3 clause 28 in gigabit full-duplex. P9_MODE and P10_MODE should be set to GMII mode at reset (see Table 8 and Table 10). In this mode, there is no external PHY for Port 9 or Port 10, so Port 9 or Port 10 are skipped by the PHY Polling Unit (PPU). Initially, the link status is configured down requiring the system software to force the port’s link up to enable the port (in the PCS Control Register).

This configuration is identical to the GMII MAC Mode described above except that a CPU is connected instead of a PHY. The lack of an external PHY device restricts the interface to a gigabit speed only with the link initially being down. This allows the CPU time to initialize itself before it enables the switch port connected to it by forcing link up in the switch port’s MAC (in the port’s PCS Control Register - offset 0x01).

Figure 11: GMII PHY Interface Pins

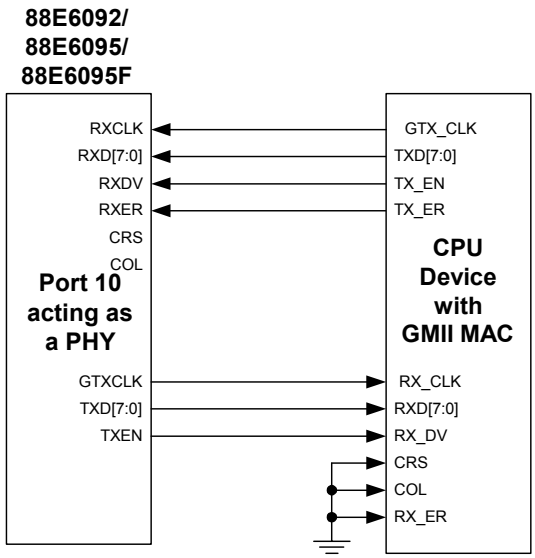

2.2.4 PHY Polling Unit (PPU)

The devices contain a PHY Polling Unit (PPU) to transfer Link, Speed, Duplex and Pause information from an external PHY to its associated MAC (the internal PHYs use a direct approach such that this information is transferred even if the PPU is disabled). The PPU can perform this job only if the SMI address of the external PHY matches the physical port number it is connected to in the switch (i.e., the PHY connected to Port 8 uses SMI address 0x08, the PHY connected to Port 9 uses SMI address 0x09, and the PHY connected to Port 10 uses SMI address 0x0A, etc.).

If the PPU is disabled, software must perform the job of setting the switch MAC's mode to the mode of the PHY (for the external PHYs) by forcing the MAC's link, speed, duplex and pause settings (in the port's PCS Control Register - offset 0x01) based upon what it sees in the PHY's registers. Link up must be the last mode register set and link down must be the first mode register cleared (i.e., the port's speed, duplex and pause modes must only be changed while the port's link is down).

If the PPU is enabled, it has full access to the external and internal PHY's registers so software cannot access any of the PHY registers unless it disables the PPU and insures the PPU is in the idle state. The PPU has four states that can be seen in the PPUState register (Switch Global Status Register - global offset 0x00) and software must use these states before changing the mode on the PPU (either enable or disable) or before it accesses any PHY register. The PPU can be re-enabled without it re-initializing the PHYs by setting the PPUEn bit to a 1 (in the Switch Global Control register - global offset 0x04). Software can temporarily disable the PPU, modify the PHY registers it needs to, and then re-enable the PPU to continue where it left off.

Note

Software will not be able to access any of the internal or external PHY registers as long as the PPU is enabled.

Warning

If the PPU is disabled at Reset it will not have a chance to initialize the external PHYs so software will need to perform this function. Two default PHY register settings are modified by the PPU. The 1st is the setting or the clearing of the PHY's full-duplex Flow Control advertisement bit. These bits are set or cleared based upon the setting of the EE_CLK/FD_FLOW_DIS pin at the rising edge of RESETn. The 2nd initialization is the prevention of any gigabit PHY from advertising gigabit half-duplex mode as this mode is not supported in the MACs.

2.2.5 24 Port Switch Configuration

The P9_TXD[3]/SW_24P configuration pin can be used to pre-configure Port 8 and Port 9 for a three-chip 24 port switch configuration without the need of a CPU or EEPROMs. This mode is intended for unmanaged designs only and should not be used if the design is different from Figure 6.

If SW_24P is pulled high at the rising edge of RESETn Ports 8 and 9 are effected in the following way:

- They are configured as Distributed Switching Architecture (DSA) Tag ports
- They are forced to an cross-chip interface (C_MODE = 0x4 - Table 62)
- Link is forced up, Duplex is forced to full-duplex and speed is forced to gigabit
- They are restricted from sending frames to each other (i.e., their VLANTable is set to 0x4FF on both ports)

The PPU is enabled and the MDC_PHY/PPU_EN/LRN2ALL configuration pin performs LRN2ALL configuration.

Section 3. Switch Core Functional Description

3.1 Switch Data Flow

The devices accept IEEE 802.3 frames and either discards them or transmits them out of one or more of the switch's ports. The decision on what to do with each frame is just one of the many tasks handled inside the switch. Figure 12 shows the data path inside the switch along with the major functional blocks that process the frame as it travels through the devices. Each of these blocks along with their register-controllable options and policy is described in the sections which follow.

This section focuses on the frame processing and policies that take place in the switch core (from MAC receive to MAC transmit) of a single device.

Figure 12: Switch Data Flow

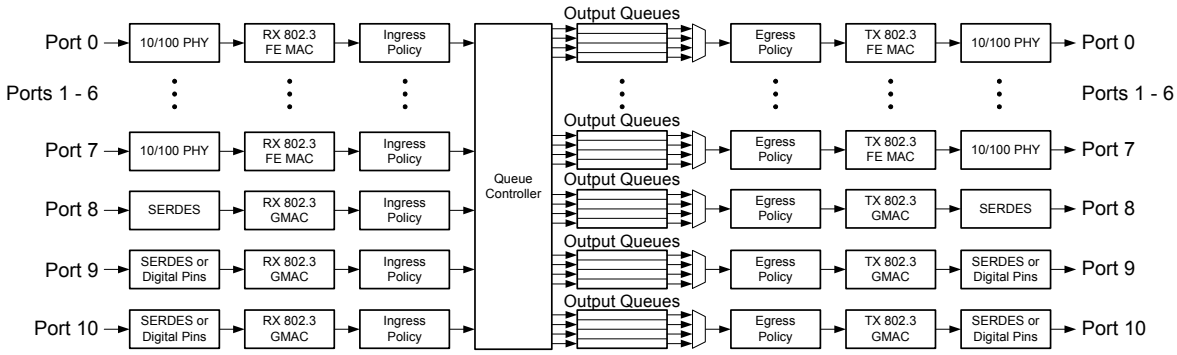

3.2 Physical Interface

Each port contains a physical interface of some sort to receive and transmit frames to and from the port's MAC. Some ports support many different physical interface options while others support only one. If a port supports many interface options only one option can be used at a time. The physical interface options that each port supports are covered in Section 2.2.

3.3 Media Access Controllers (MAC)

The devices contain eight independent 10/100 MACs and three gigabit MACs/SGMII. These MACs perform all of the functions in the 802.3 protocol including frame formatting, frame stripping, CRC checking, CSMA/CD enforcement, and collision handling. Each MAC supports 1 Gbps operation in full-duplex mode (Ports 8, 9 and 10 only), or 10/100 Mbps operation in either full-duplex or half-duplex mode.

The MAC receive block checks incoming packets and discards those with CRC errors, those with alignment errors, short packets (less than 64 bytes or less than 68 bytes for frames with 802.1Q tag)[1], or long packets (more than 1522 bytes)[2]. Each MAC constantly monitors its receive lines waiting for preamble bytes followed by the Start of Frame Delimiter (SFD). The first six bytes after the SFD are used as the packet's Destination Address (DA)[3] and the next six bytes after that are used as the packet's Source Address (SA). These two addresses are fundamental to the operation of the switch (see Section 3.7 for more information). The next two to sixty-six bytes are examined and may be used for QoS (Quality of Service) or snooping decisions made by the switch (see Section 3.5 for more information). The last four bytes of the packet contain the packet's Frame Check Sequence (FCS). The FCS must meet the IEEE 802.3 CRC-32 requirements for the packet or it will be discarded.

Before a packet can be sent out, the transmit block must check if the line is available for transmission. The transmit line is available all the time when the port is in full-duplex mode, but the line could be busy receiving a packet if the port is in half-duplex mode. If the line is busy, the transmitter waits by deferring its transmission. When the line is available the transmitter ensures that a minimum interpacket gap of at least 96 bits prior to transmitting a 56-bit preamble and an 8-bit Start of Frame Delimiter (SFD) ahead the frame. Actual transmission of the frame begins immediately after the SFD.

For half-duplex mode, the devices also monitor the collision signal while it is transmitting. If a collision is detected (i.e., both transmitter and receiver of a PHY are active at the same time, the MAC transmits a JAM pattern and then delays the re-transmission for a random time period determined by the IEEE 802.3 backoff algorithm. In full-duplex mode, the collision signal and backoff algorithm are ignored.

3.3.1 Backoff

In half-duplex mode, the devices MACs implement the truncated binary exponential backoff algorithm defined in the IEEE 802.3 standard. This algorithm starts with a randomly-selected small backoff time and follows by generating progressively longer random backoff times. The random times prevent two or more MACs from always attempting re-transmission at the same time. The progressively longer backoff times give a wider random range at the expense of a longer delay, giving congested links a better chance of finding a winning transmitter. Each MAC in the devices reset the progressively longer backoff time circuit after 16 consecutive retransmit trials. Each MAC then restarts the backoff algorithm with the shortest random backoff time and continues to retry and retransmit the frame. A packet that successively collides with a retransmit signal is re-transmitted until transmission is successful. This algorithm prevents packet loss in highly-congested environments.

3.3.2 Half-duplex Flow Control

Half-duplex flow control is used to throttle the throughput rate of an end station to avoid dropping packets during network congestion. It is enabled on all half-duplex ports via the EE_DIN/HD_FLOW_DIS pin (see Table 14). Flow control can be enabled or disabled on any particular port by forcing the port's Flow Control mode (FCValue and ForcedFC in the PCS Control Register, offset 0x01). The devices use a mixed carrier assertion and collision based scheme to perform half-duplex flow control. When the free buffer space is almost empty, the MAC issues

1. If Ingress Double Tagging (88E6095 and 88E6095F only) is enabled the minimum frame size for Tagged frames is 68 bytes.
2. A maximum frame size of 1632 bytes is supported by setting the MaxFrameSize bit in the Global Control register (Table 83).
3. The 1st six bytes of a frame are processed as the frame's DA unless the Marvell Header mode is enabled on the port (Table 65). If the Marvell Header mode is enable the 1st two bytes are processed as the Marvell Header and the next 6 bytes are processed as the frame's DA.

carrier and/or collision which prevents further incoming packets. Only the ports that are involved in the congestion are flow controlled. If the half-duplex flow control mode is not set and there is no packet buffer space available, the incoming packet is discarded.

Half-duplex flow control is not an IEEE defined feature. The IEEE defined full-duplex flow control is described in the next section.

3.3.3 Full-duplex Flow Control

IEEE 802.3 flow control mechanism requires two link partners to auto-negotiate and advertise their flow control capabilities. If both link partners are flow control capable, then flow control will be enabled in both link partners MACs. The PHYs are used to advertise the capability but the flow control itself is a function of the MAC. The IEEE flow control also requires full-duplex operation.

The purpose of full-duplex flow control is the same as in half-duplex – avoid dropping packets during congestion. Full-duplex flow control is enabled on all full-duplex ports via the EE_CLK/FD_FLOW_DIS pin (see Table 14), if Auto-Negotiation is enabled on the port, and if the link partner 'advertises' that it supports Pause during Auto-Negotiation. Basically, full-duplex flow control is automatically enabled on a port if:

- The port's PHY is advertising it supports flow control

and

- The port's PHY sees that its link partner is also advertising it supports flow control too (once link is established).

The EE_CLK/FD_FLOW_DIS pin (at the rising edge of RESETn) determines the initial flow control advertisement bit setting in the PHYs. The inverted value of this pin is moved to external PHYs by the PPU, if enabled (Section 2.2.4). Internal PHYs will have their flow control advertisement bit initialized even if the PPU is disabled.

When flow control is enabled using the EE_CLK/FD_FLOW_DIS or EE_DIN/HD_FLOW_DIS device pins, it is enabled for all ports of the same type (i.e., on all half-duplex ports or on all full-duplex ports that have a flow-controllable link partner). It may be required to have flow control enabled on only one or two ports and disable flow control on all other ports. In this case, flow control should be disabled via FD_FLOW_DIS and HD_FLOW_DIS device pins, which will disable all the ports. The ports chosen to have flow control enabled can then be configured to advertise flow control. This can be done by disabling the PPU and writing to the internal or external PHYs flow control advertisement bit. The PPU can be enabled to allow the MAC to determine the results of the auto-negotiation with the link partner.

If the full-duplex flow control mode is not set and if there is no packet buffer space available, the incoming packet is discarded.

In full-duplex mode, the devices MACs support flow control as defined in the IEEE 802.3 standard. This flow control mechanism enables stopping and restarting packet transmission at the remote node. The basic mechanism for performing full-duplex flow control is via a Pause frame. The format of the Pause frame is shown in Table 21

.

Table 21: Format of Pause

Destination Address (6 Bytes)	Source Address (6 Bytes)	Type (2 Bytes)	Op Code (2 Bytes)	Pause Time (2 Bytes)	Padding (42 Bytes)	FCS (4 Bytes)
01-80-C2-00-00-01	See text	88-08	00-01	See text	All zeros	Computed

Full-duplex flow control works as follows. When a port's free buffer space is almost empty the devices send out a Pause frame with the maximum pause time to stop the remote node from sending more frames into the switch. Only the ports that are involved in the congestion are Paused. When congestion on the port is relieved, the devices send out a Pause frame with pause time equal to zero, indicating that the remote node may resume transmission.

The devices also responds to the Pause command in the MAC receiving block. When the Pause frame is detected, the port responds within one slot time to stop transmission of new data for the amount of time defined in the pause time field of the received Pause frame.

The Source Address of a received Pause frame is not learned[1] since it may not represent the Source Address of the transmitting port. This is generally the case if the link partner is an unmanaged switch. The devices can be configured to transmit a unique Source Address on Pause frames (see switch MAC address register Table 80 through Table 82). A single fixed Source Address can be used for all ports, or a unique Source Address per port can be selected by the changing the value of the DiffAddr bit in the switch MAC Address register.

The MACs discard all IEEE 802.3 Pause frames received. This is always the case, even if full-duplex flow control is disabled or if the port is in half-duplex mode.

3.3.4 Forcing Flow Control in the MAC

Section 3.3.3 describes the IEEE defined flow control mechanism, which requires auto-negotiation with a link partner. Some ports may not have a PHY attached (e.g., Port 9 in GMII mode), or there may be a PHY attached without auto-negotiation. In this case, flow control can be enabled or disabled on any particular port by forcing the port's Flow Control mode (FCValue and ForcedFC in bit in the port's PCS Control Register, offset 0x01). Forcing flow control on will instruct the port's MAC to transmit Pause frames when needed and act on received Pause frames. It does not change the advertisement bits in the port's PHY[2].

If the port has a PHY connected it (either internal or external) with auto-negotiation enabled, it is best to not force flow control (by using FCValue and ForcedFC) if the port is in full duplex mode. Instead set the PHY's auto-negotiation flow control advertisement bit to allow flow control to occur automatically if the port's link partner agrees.

1. See Automatic Address Learning in Section 3.4.3.
2. In this case the port's link partner may not be supporting Pause frames because it does not see from the PHY that this port is advertising it wants to support Pause.

3.3.5 MAC Based RMON/Statistics Counters

The Statistics Counter logic maintains a set of 28, 32-bit counters and two 64-bit counters per port, that enable the user to monitor network performance remotely and to support RMON groups 1, 2, 3 and 9. These counters provide a set of Ethernet statistics for frames received on ingress and transmitted on egress. A register interface allows the CPU to capture, read, or clear the counter values (see the Stats Operations register, global offset 0x1D, for more information).

The counters are designed to support:

- RFC 2819 – RMON MIB (this RFC obsoletes 1757 which obsoletes 1271)
- RFC 2665 – Ethernet-like MIB (this RFC obsoletes 1643, 1623 and 1398)
- RFC 2233 – MIB II (this RFC obsoletes 1573 & 1213 with obsoletes 1229 & 1158)
- RFC 1493 – Bridge MIB (this RFC obsoletes 1286)

The complete description of each of the counters is contained in Table 22 and Table 23.

All CPU register interfaces are slow compared with the speed of Gigabit or even Fast Ethernet frames. For this reason the devices support a snapshot function to capture instantly and hold static any port's MAC Statistics counters. The capture function maintains a higher level of accuracy between the various counters in a port and also allows multiple counter values to be added together to get the required MIB. After capture, the CPU can take its time to read out the values of the counter or counters that it needs without concern for the values changing during the register read process.

The CPU interface supports the following operations on the Stat Counters:

- Clear all counters for all ports
- Clear all counters for a single port
- Capture all counters for a single port
- Read a captured counter (a Capture must be executed before a Read to the capture zone can be done).
- Read a counter directly (best used when reading only one counter on a port)

See the Stats Operation Register (global offset 0x1D) for more details.

Note

The Set 4 counters can be configured to be ingress only, egress only or both.

Table 22: Ingress Statistics Counters

Name	Offset Address	Description
Set 1		
InGoodOctetsLo	0x00	The lower 32-bits of the 64-bit InGoodOctets counter. The sum of lengths of all good Ethernet frames received, that is frames that are not bad frames.
InGoodOctetsHi	0x01	The upper 32-bits of the 64-bit InGoodOctets counter. See description above.
InBadOctets	0x02	The sum of lengths of all bad Ethernet frames received.
Set 2		
InUnicast	0x04	The number of good frames received that have a Unicast destination MAC address.
InBroadcasts	0x06	The number of good frames received that have a Broadcast destination MAC address.
InMulticasts	0x07	The number of good frames received that have a Multicast destination MAC address. **Note:** This does not include 802.3 Flow Control messages counted in InPause nor does it include Broadcast frames counted in InBroadcasts.
InPause	0x16	The number of good Flow Control frames received.
Set 3		
InUndersize	0x18	Total frames received with a length of less than 64 octets but with a valid FCS.
InFragments	0x19	Total frames received with a length of less than 64 octets and an invalid FCS.
InOversize	0x1A	Total frames received with a length of more than MaxSize octets but with a valid FCS.
InJabber	0x1B	Total frames received with a length of more than MaxSize octets but with an invalid FCS.
InRxErr	0x1C	Total frames received with an RxErr signal from the PHY.
InFCSErr	0x1D	Total frames received with a CRC error not counted in InFragments, InJabber or InRxErr.
Set 4		**These counters can be Ingress Only, Egress Only, or both**
64Octets	0x08	Total frames received (and/or transmitted) with a length of exactly 64 octets, including those with errors.
65to127Octets	0x09	Total frames received (and/or transmitted) with a length of between 65 and 127 octets inclusive, including those with errors.
128to255Octets	0x0A	Total frames received (and/or transmitted) with a length of between 128 and 255 octets inclusive, including those with errors.

Table 22: Ingress Statistics Counters (Continued)

Name	Offset Address	Description
256to511Octets	0x0B	Total frames received (and/or transmitted) with a length of between 256 and 511 octets inclusive, including those with errors.
512to1023Octets	0x0C	Total frames received (and/or transmitted) with a length of between 512 and 1023 octets inclusive, including those with errors.
1024toMaxOctets	0x0D	Total frames received (and/or transmitted) with a length of between 1024 and MaxSize[1] octets inclusive, including those with errors.

1. MaxSize is 1522 for non-tagged frames and for tagged frames if MaxFrameSize = 0 or MaxSize = 1632 if MaxFrameSize = 1 (Table 83).

Table 23: Egress Statistics Counters

Name	Offset Address	Description
Set 5		
OutOctetsLo	0x0E	The lower 32-bits of the 64-bit OutOctets counter. The sum of lengths of all Ethernet frames sent from this MAC.
OutOctetsHi	0x0F	The upper 32-bits of the 64-bit OutOctets counter. See description above.
Set 6		
OutUnicast	0x10	The number of frames sent that have a Unicast destination MAC address.
OutBroadcasts	0x13	The number of good frames sent that have a Broadcast destination MAC address.
OutMulticasts	0x12	The number of good frames sent that have a Multicast destination MAC address. **Note:** This does not include 802.3 Flow Control messages counted in OutPause nor does it include Broadcast frames counted in Out-Broadcasts.
OutPause	0x15	The number of Flow Control frames sent.
Set 7		
Deferred	0x05	The total number of successfully transmitted frames that experienced no collisions but are delayed because the medium was busy during the first attempt. This counter is applicable in half-duplex only.
Collisions	0x1E	The number of collision events seen by the MAC not including those counted in Single, Multiple, Excessive or Late. This counter is applicable in half-duplex only.
Single	0x14	The total number of successfully transmitted frames that experienced exactly one collision. This counter is applicable in half-duplex only.
Multiple	0x17	The total number of successfully transmitted frames that experienced more than one collision. This counter is applicable in half-duplex only.
Excessive	0x11	The number frames dropped in the transmit MAC because the frame experienced 16 consecutive collisions. This counter is applicable in half-duplex only and only of DiscardExcessive is a one (in Switch Global Control - global offset 0x04).
Late	0x1F	The number of times a collision is detected later than 512 bits-times into the transmission of a frame. This counter is applicable in half-duplex only.
OutFCSErr	0x03	The number of frames transmitted with an invalid FCS. Whenever a frame is modified during transmission (e.g., to add or remove a tag) the frame's original FCS is inspected before a new FCS is added to a modified frame. If the original FCS is invalid, the new FCS is made invalid too and this counter is incremented.

Table 23: Egress Statistics Counters (Continued)

Name	Offset Address	Description
Set 4		**These counters can be Ingress Only, Egress Only, or both**
64Octets	0x08	Total frames transmitted (and/or received) with a length of exactly 64 octets, including those with errors.
65to127Octets	0x09	Total frames transmitted (and/or received) with a length of between 65 and 127 octets inclusive, including those with errors.
128to255Octets	0x0A	Total frames transmitted (and/or received) with a length of between 128 and 255 octets inclusive, including those with errors.
256to511Octets	0x0B	Total frames transmitted (and/or received) with a length of between 256 and 511 octets inclusive, including those with errors.
512to1023Octets	0x0C	Total frames transmitted (and/or received) with a length of between 512 and 1023 octets inclusive, including those with errors.
1024toMaxOctets	0x0D	Total frames transmitted (and/or received) with a length of between 1024 and MaxSize[1] octets inclusive, including those with errors.

1. MaxSize is 1522 for non-tagged frames and for tagged frames if MaxFrameSize = 0 or MaxSize = 1632 if MaxFrameSize = 1 (Table 83).

3.3.6 Policy Based RMON/Statistics Counters

The Queue Controller logic maintains a set of policy counters (one 32-bit and two 16-bit) per port. that enable the user to monitor network performance by seeing where good frames have been dropped by the switch (bad frames that are dropped are counted in the MAC based counters). Some frames are dropped due to switch policy and others are due to excessive congestion in the switch.

The policy counters are:

- InDiscards - A 32-bit counter (16 bits in InDiscardsLo, offset 0x10, and 16 bits in InDiscardsHi, offset 0x11) that counts the number of good, non-filtered frames that normally would have been forwarded, but could not be due to a lack of buffer space.
- InFiltered - A 16-bit counter (offset 0x12) that counts the number of good frames that were filtered due to ingress switch policy rules. These rules include frames that are dropped due to 802.1X MAC authentication (DropOnLock - Port Control register, offset 0x04), 802.1Q Security checks (802.1QMode in Port Control 2 register, offset 0x08), DiscardTagged & DiscardUntagged (Port Control 2, offset 0x08), PortState other than Disabled (Port Control register, offset 0x04), and DA mappings back to the source port (normal switch filtering).
- OutFiltered - A 16-bit counter (offset 0x13) that counts the number of good frames that were filtered due to egress switch policy rules. These rules include frames that passed the ingress port's policy but are dropped due to the egress policy of this port including 802.1Q Security checks (802.1QMode in Port Control 2 register, offset 0x08) and PortState other than Disabled (Port Control register, offset 0x4).

The description of each of the counters is contained in Table 22 and Table 23.

These counters stop counting when the port's PortState is set to Disabled (Port Control register - offset 0x04) and they are all cleared when a Flush All Counters for this port or a Flush All Counter for All Ports command is issued to the MAC based counters (see the Stats Operation Register, global offset 0x1D, for more details).

3.4 Basic Switch Operation

The switch portion of the devices receive good packets from the MACs, processes them, and forwards them to the appropriate MACs for transmission. The primary task of the switch is to process frames and this activity involves the Ingress Policy, the Queue Controller, the Output Queues and the Egress Policy blocks shown in Figure 12. These blocks modify the normal or default packet flow through the switch and are discussed in Section 3.5 to Section 3.7.

The normal packet flow and processing through a switch involves learning how to switch packets to the correct MACs, and only to the correct ones. The switch learns what port an end station is connected to by remembering each packet's Source Address[1] along with the port number on which the packet arrived. Once a MAC address/port number mapping is learned, all future packets directed to that end station's MAC address (as defined in a frame's Destination Address field) are directed to the learned port number only. If a packet is directed to a new, currently unlearned, MAC address, the packet will be transmitted out of all the ports[2], except for the one on which it arrived[3]. This ensures that the packet is received by the correct end station (if it exists) and when the end station responds back its address is learned by the switch for the next series of packets.

All switches learn only a very small subset of the set of possible MAC addresses owing to the limits of physical memory. Switches learn only the currently 'active' MAC addresses. Sometimes end stations are moved from one port to another so that a new MAC address/port number association must be learned and the old one replaced. All of these issues are handled by what is called 'Aging' and 'Station Move Handling'. Basically, a MAC address/port number association is allowed to be 'active' for only a limited amount of time. This time limit is typically set to five minutes.

The following sections describe how the devices perform its basic switch functions.

3.4.1 Lookup Engine

The devices Lookup Engine or Address Translation Unit (ATU) uses the DA and SA fields from each frame received from each port. It performs all address searching, address learning, and address aging functions for all ports at 'wire speed' rates (i.e., a DA and an SA lookup/learn function can be performed for all ports in less time than it takes to receive a 64 byte frame on any port).

The address database uses a hashing technique for quick storage and retrieval. Hashing a 48-bit address into fewer bits results in some MAC addresses having the same hash address. This is called a hash collision and is solved in the devices by using four bins per hash location allowing for storage of up to four MAC addresses at each hash location. This allows the address database to be smaller while still holding the same number of active, random value MAC addresses.

The address database is stored in embedded SRAM and has a size of 8192 entries with a default aging time of about 300 seconds or 5 minutes. The age time can be modified in 15 second increments from 0 seconds (aging disabled) to 3825 seconds (almost 64 minutes). These options are set in the ATU Control register (global offset 0x0A).

3.4.2 Address Searching or Translation

The address search engine is used to search the address database to get the output port number(s), called the Destination Port Vector (DPV), for each frame's destination address so that it can switch the frame instead of flooding[4] it. It arbitrates destination address lookup requests from the ports and grants one lookup at a time. The address is hashed and then data is read from the SRAM table, looking for a MAC address match. Four addresses

1. The SA on switch management frames are not learned.
2. VLANs modify this operation – see Section 3.5.6.
3. The devices can be configured to transmit frames out the port they came in on – see Table 67.
4. Flooding refers to the action of sending frames out all the ports of the switch except for the port the frame came in on.

can be stored at each hash location. If a match is found, the Address Translation Unit (ATU) returns the Destination Port Vector (DPV) to the Ingress Policy block where it may be modified[1] before the packet is queued to the output ports. If no MAC address match is found the Ingress Policy block uses a unique default DPV for each ingress port[2], which typically floods the frame. If the destination address in the frame is a multicast address or broadcast address, the address is searched[3] in the same way as a unicast address and the frame is processed identically. This feature is used for multicast filtering. Multiple separate address databases are supported in the devices. The database that is searched is controlled by the port's default database number (DBNum in the Port Based VLAN Map register, offset 0x06, and the Port Control 1 register, offset 0x05) or the one assigned to the frame by the VTU (Section 3.5.11). MAC addresses that are not members of the port's or frame's DBNum cannot be found.

3.4.3 Automatic Address Learning

The address learning engine is used to learn the source address of ingressing frames. Up to 8192 MAC address/port number mappings can be stored in the address database (See Section 3.4.1 for more information). When the source address from an input frame can not be found in the address database, the ATU enters the self-learning mode and places the new MAC address/port number mapping into the database and refreshes its Age time[4]. If the MAC address is found to be already in the database, the port number and Age associated with the entry is updated and/or refreshed. The port number is updated in case the end station moved and the port number needs to be corrected. The entry's Age is refreshed since the MAC address is still 'active'. This prevents the MAC address/port number mapping from being removed as being 'inactive' prematurely.

When an address is added into the database it is hashed and stored in the first empty bin found at the hashed location. If all four address bins are full, a least recently used algorithm is used for looking at each entry's Age time (its EntryState field[5]). If all four address bins have the same Age time, then the first 'non-static' bin is used (see Section 3.4.10.1 for more information about locked or static addresses). If all four bins are 'static' the address is not learned and an ATUFull interrupt is generated (see the Switch Global Status register, global offset 0x00).

Multiple separate address databases are supported in the devices (see Section 3.4.7). The port's database number (DBNum in the Port Based VLAN Map register, offset 0x06, and the Port Control 1 register, offset 0x05) determines into which address database the MAC address is added if 802.1Q is disabled on the port. If 802.1Q is enabled on the port the DBNum associated with the frame's VID (VLAN ID field of Tagged frames, see Section 3.5.12.9) determines the address database into which the MAC address is stored. The same MAC address can be learned multiple times with different port mappings if different DBNum values are used.

Learning can be disabled on any individual port by clearing the port's PAV to all zeros (see Port Association Vector, offset 0x0B) or by setting the port's learn disable bit (in Port Based VLAN Map register, offset 0x06). Learning is also disabled on any port that has a PortState of Disabled or Blocking/Listening (See the Port Control register, offset 0x04).

Learning is never performed on switch management frames. This includes Pause frames (Section 3.3.3), BPDU, LAC and other management frames as long as they are determined to be MGMT frames (Section 3.4.8), and non-Forward Distributed Switching Architecture (DSA) Tag frames (Section 3.5.16).

1. The DPV returned from the ATU may be modified by the VLANTable data, VTU results, and/or the Trunk Mask Table.
2. The default DPV for each port is the list of ports that can egress multicast frames, if the frame is multicast (DefaultForward in Port Control 2 register, offset 0x08) or the list of ports that can egress unicast frames, if the frame is unicast (FordwardUnknown in Port Control register, offset 0x04).
3. Multicast addresses cannot be auto learned. Multicast addresses must be loaded manually with a CPU or EEPROM.
4. The Age time on a MAC Address entry is refreshed by setting its EntryState field to 0x7—see Table 25.
5. The EntryState field is described in Section 3.4.10.1.

3.4.4 Automatic Address Aging

The address aging process is used to ensure that if a node is disconnected from the network segment, or if it becomes inactive, its entry is removed in a timely manner from the address database. Aging makes room for new active addresses. An address is removed from the database after a programmable amount of time from the last time it appeared in an ingressing frame's Source Address. This programmable time is determined by the Age Time bits in the ATU Control register (global offset 0x0A).

The devices run the address aging process continuously (unless disabled by setting the AgeTime field to zeros). Aging is accomplished by a periodic sweeping of the address database. The speed of these sweeps determine the aging time. On each aging sweep of the database, the ATU reads each valid entry and updates its age time by decrementing its EntryState field[1] (as long as the entry is not locked – see Section 3.4.10.1). When the EntryState field reaches zero, the entry is considered invalid and purged from the database.

A new or just refreshed unicast MAC address has an EntryState value of 0x7 (see Section 3.4.3). A purged or invalid entry has an EntryState value of 0x0. The values from 0x6 to 0x1 indicate the Age time on the unicast MAC address with 0x1 being the oldest. This scheme results in seven age states on an entry allowing the Address Learning's least recently used replacement process (Section 3.4.3) to be more precise. An address is purged from the database within 1/7th of the programmed AgeTime value making the address's lifetime interval in the database more accurate as well.

3.4.5 CPU Directed Address Learning

Sometimes it is required to prevent automatic learning from occurring on a port and have a CPU direct the learning instead. The devices support CPU directed learning on a per port basis by setting the port's LockedPort bit to a one (in the Port Association Vector register - offset 0x0B). When a port is 'Locked' all frames received with an SA not found in the address database cause an SA Miss ATU Violation as long as learning is enabled on the port (i.e., the port's PAV, offset 0x0B, is non-zero). The SA Miss ATU Violation can be set to generate an interrupt—see the ATUProbIntEn bit of the Switch Global Control register (offset 0x04). One ATU Violation per port is held in the ATU. Once a violation is captured all subsequent violations are ignored until the first one is serviced by the CPU.

The CPU can retrieve the source MAC address and the source port information of the SA Miss Violation by issuing an ATU Get/Clear Violation Data ATUOp (Section 3.4.10.5). The CPU then decides if the address should be placed into the address database or not. If it should be, the CPU issues a Load ATUOp (Section 3.4.10.3). If the CPU loads the new address as 'non-static', the entry stays in the address database until it ages out. Its age time is determined by the EntryState's value, as addresses on Locked ports do not have their age time refreshed. The CPU receives no new interrupts from this address until it ages out or is purged out by the CPU.

If the CPU loads the new address as 'static', the entry stays in the address database until the CPU purges it. In this case, the CPU receives no new interrupts from this address as long as the address is never used as an SA on a port other than the original source port. If the MAC address is used on a different source port, the CPU can receive an ATU Member Violation interrupt if the IgnoreWrongData bit (see the Port Association Vector register, offset 0x0B) is cleared on the port where the frame just entered the switch. This interrupt may indicate that a station move just occurred or that an end station is masquerading by using another station's address.

3.4.6 802.1X Source MAC Address Checking

The devices support 802.1X source MAC address authentication using CPU Directed Address Learning (Section 3.4.5) along with the DropOnLock Ingress policy (Section 3.5.2). CPU Directed Address Learning is required for 802.1X[2] so that the requesting MAC address can be authorized by the authorization server. The DropOnLock policy causes all frames from unauthorized MAC addresses to be discarded.

1. The EntryState field is described in Section 3.4.10.1.
2. An 802.1X port needs to have its LockedPort and its DropOnLock bits each set to one (see the Port Association Vector, offset 0x0B, and the Port Control register, offset 0x04).

A side effect of authorization is that a CPU might become saturated from constant SA Miss Violations from a source that it has denied. This is prevented in the devices by masking denied MAC addresses. An address can be masked by loading it into the address database with a Destination Port Vector (DPV) of all zeros. The address appears in the database so it no longer causes a 'Miss' Violation. Any port trying to send a frame to this unauthorized address is discarded (since its DPV is all zeros) and any 802.1X port trying to use this unauthorized address has its frames discarded too (since the port's SA bit is not set in the ATU entry). But now the CPU will get SA Member Violation interrupts every time the unauthorized address is attempted to be used as an SA. These interrupts can be masked too by setting the IgnoreWrongData bit (in the Port Association Vector - offset 0x0B) on the 802.1X ports.

The CPU can mask unauthorized MAC addresses by loading them into the address database as static or non-static entries. If they are loaded non-static, the interrupts are masked until the entry ages out of the database or until the CPU purges the entry. This approach minimizes the number of interrupts the CPU needs to service while keeping the address database fluid. If the unauthorized address is loaded as static, the interrupts are masked until the CPU purges the entry. This requires the CPU to remember addresses it has loaded because at some time the CPU should purge these addresses to make room for new ones. The use of too many static addresses can also cause an ATU Full Violation, so it is best to mask addresses using the non-static approach. The DropOnLock feature prevents the reception of frames from all unauthorized MAC addresses regardless of whether the unauthorized address is currently being masked in the address database or not. The masking feature is intended to minimize the number of SA Miss Violation interrupts the CPU needs to service for addresses that it already has denied.

If a port is saturating a CPU by constantly using a new SA, the masking of unauthorized addresses is not applicable. Instead all SA Miss Violations can be masked on a port by disabling learning on the port (i.e., by setting the port's Port Association Vector bits to all zero, offset 0x0B or by setting the DropOnLock bit (Port Control Register Offset 0x04 or decimal 4 - to 1), which will drop the frame that caused the interrupt.). This configures the ingress port in a form of a Secure Port (Section 3.5.3) and will work as long as no new SA needs to be authorized on this port.

3.4.7 Multiple Address Database Support (DBNum)

The devices support up to 256 separate and independent address databases in the Address Translation Unit (ATU). Multiple address database are used to isolate MAC addresses by the VLAN so the same MAC address can appear multiple times in the address database with different port mappings. The devices use a database number (DBNum) mechanism to isolate the address databases from each other. Although the address database is isolated by DBNum value it is *not* divided up in equal segments by DBNum value. Each database number can hold none to all of the possible 8192 MAC addresses or any number in between. Any number of DBNum values can be used (from 1 to 256). Each database number (DBNum) uses only the MAC address entries it needs and leaves all the remaining ATU entries for all the other database numbers.

Each frame, as it ingresses a port, is assigned a DBNum. The frame's DBNum value, along with the frame's DA and SA, is sent to the Address Translation Unit (ATU) when the frame's MAC addresses are searched and/or learned. The frame's DBNum is determined in priority order by:

- The DBNum associated with the frame's VID[1] in the VLAN Translation Unit (VTU, Section 3.5.11). This requires the frame's VID to be valid in the VTU. All frames are assigned a VID, even untagged frames. This is true if 802.1Q is enabled on the port or not.
- The DBNum associated with the ingressing port (DBNum[7:4] in Port Control 1 register, offset 0x05 and DBNum[3:0] in Port Based VLAN Map register, offset 0x06).

If multiple address databases are not needed leave all the DBNum values at their reset value of 0x00 in all their occurrences in the registers (in the ports, ATU and VTU).

1. The frame's VID can be overridden. See Section 3.7.1.3.

A frame's DBNum is generally determined by the frames VID (via the VTU, Section 3.5.11). Multiple VID's can be mapped to the same DBNum allowing for shared VLAN address databases.

3.4.8 Management/802.1D BPDU Frame Detection

The devices support two methods of management/802.1D BPDU frame detection and mapping. These two mechanisms support IEEE industry standards like STP[1], LAC[2] and OAM[3], as well as any company proprietary protocol. Both of these mechanisms:

- Detect that a frame is special by examining the frame's DA.
- Assign these special frames to a category called MGMT (for management).
- Allow MGMT frames and only MGMT frames to ingress and egress Blocked ports (see Port States in Section 3.5.1).
- Set the priority on these MGMT frames overriding *all* other QoS decisions on the frame (Section 3.5.5).
- Map these MGMT frames to a port where a management CPU is directly or indirectly connected.

3.4.8.1 Reserved Multicast Address Support

The first mechanism for MGMT frame detection is optimized to support 802.1D's 16 reserved multicast addresses. Any or all of the 16 multicast addresses in the range of 01:80:C2:00:00:0x[4] can be treated as MGMT addresses in the devices. The Rsvd2Cpu register bits in the MGMT Enable register (global 2, offset 0x03) determines which of these 16 addresses are treated as MGMT and which are not as long as the Rsvd2Cpu bit in the Switch Management register (global 2 offset 0x05) is also be set to a one.

Any frame, regardless of its VLAN identifier (VID) or DBNum[5] assigned to it, whose DA matches an enabled reserved multicast address will be considered a MGMT frame. It will be given the priority defined in the MGMT_Pri bits (in the Switch Management register, global 2 offset 0x05) and mapped to the port defined by the source port's CPUPort register (in Port Control 2, offset 0x08).

3.4.8.2 New and Proprietary Protocol Support

The second mechanism for MGMT frame detection is optimized to support any new, or yet to be defined, standard and/or proprietary DA based protocol. It can also be used to map any of the 16 reserved multicast addresses where the VID of the frame must be considered in the MGMT determination[6]. Any multicast or unicast address can be treated as a MGMT address in the devices. The Address Translation Unit (ATU) is used in this case. The required MAC address must be loaded into the ATU with a MGMT EntryState value, the required priority for the frame and where the frame is to be mapped (see Section 3.4.10).

Any frame whose DA matches an ATU entry with a MGMT EntryState will be considered a MGMT frame. It will be given the priority defined in the ATU's entry and mapped to the port or ports defined by the entry's Destination Port Vector (DPV). The ATU supports multiple address databases (Section 3.4.7) so the DA must appear in the address database number (DBNum) assigned to the frame for the frame to be considered a MGMT frame. This feature allows a DA to be considered a MGMT address in some address databases and not in others. But it also requires that the DA be loaded multiple times, once for each address database that needs to use this address as a MGMT address.

1. STP - Spanning Tree Protocol
2. LAC - Link Aggregation Control
3. OAM - Operational, Administration and Maintenance
4. Frames with a DA of 01:80:C2:00:00:01 are always treated as Pause frames and are discarded and never mapped.
5. DBNum is an address database number assigned to each frame to support multiple address databases (see Section 3.4.7).
6. If the second mechanism is being used to map any of the 16 reserved multicast addresses the bit that corresponds to the required address in the Rsvd2CpuEnables must be cleared to a zero since the first mechanism takes priority over the second mechanism.

3.4.9 MUX'ing or Ignoring Address Translation

The devices support the ability to ignore the results of a frame's DA lookup in the Address Translation Unit. The use of the DA mapping results can be enabled or disabled on a port by port basis by changing the value of the port's MapDA bit (in Port Control 2, offset 0x08). DA lookups always take place, regardless of the setting of port's MapDA bit, however. This is done so that DA lookups that are found to be MGMT entries in the ATU[1] are always mapped (even if the MapDA bit is configured to ignore the results). The application of SA Learning (Section 3.4.3) DefaultForward (Section 3.5.4) and ForwardUnknown (Section 3.5.3) are not effected by the port's MapDA bit.

When DA mapping is disabled (the port's MapDA bit equals zero) all non-MGMT frames that enter the port will be mapped based on the VLAN rules that are applied to the frame (Section 3.5.6) along with the DefaultForward (Section 3.5.4) and ForwardUnknown (Section 3.5.3) rules. These bits can be configured in such a way that all non-MGMT frames that enter a port egress out a specific port (even the frame's source port, Section 3.5.12). This can be used for the following applications.

3.4.9.1 Passing Frames to a Router

The MapDA bit can be used to MUX all non-MGMT frames that enter a port to go out another port where a router can perform more processing on the frame. VLANs are used to get the frame to the router's port (Section 3.4.9). MGMT frames will be mapped as they normally would (Section 3.4.8).

The router may decide the frame can go where it normally would have gone and returns the frame back to the switch unmodified. Assuming the ingress port on which the router is attached to has its MapDA bit set (i.e., enabled), the frame will now be mapped using the address database information. If the frame's DA is unknown or is a multicast address, the frame will flood out all the ports of the frame's VLAN including the port the frame originally came in on if the port is also a member of the frame's VLAN. Use a different VID on the frame to prevent this.

3.4.9.2 Operational, Administration, and Maintenance (OAM) Loopback

The MapDA bit can be used to MUX all non-MGMT (i.e., non-OAM) frames that enter a port to go back out the port they came in on. VLANs are used to get the frame to go back out the original port and only the original port (Section 3.4.9). MGMT frames, including OAM frames, will be mapped as they normally would (Section 3.4.8).

For the VLANs to work for loopback, force all frames entering the loopback port to use the port's DefaultVID (Default Port VLAN ID & Priority register, offset 0x07) and set the port membership for this VID to the loopback port only. A currently unused VID must be used for this purpose. The devices support a VID of 0xFFF, which is an illegal VID to be used in a network, so this is a good VID to use in the switch for loopback. The VLANTable (in Port Based VLAN Map, offset 0x06) must be modified as well to allow frames to egress the port they came in on (Section 3.5.12).

1. See Section 3.4.8.2 and Section 3.4.10.

3.4.10 Address Translation Unit Operations

The Address Translation Unit (ATU) in the devices support user commands to access the contents of the MAC address database.

All ATU operations have the same user interface and protocol. Six global registers are used and are shown in Table 24. The protocol for an ATU operation is as follows:

- Ensure that the ATU is available by checking the ATUBusy bit in the ATU Operation register. The ATU can only perform one user command at a time.
- Load the ATU Data and ATU MAC registers if required by the selected operation.
- Start the ATU operation by defining the desired DBNum, ATUOp and setting the ATUBusy bit to a one in the ATU Operation register – this can all be done at the same time.
- Wait for the ATU operation to complete. This is done by polling the ATUBusy bit in the ATU Operation register or by receiving an ATUDone interrupt (see Switch Global Control, global offset 0x04, and Global Status, global offset 0x00).
- Read the results if appropriate.

Table 24: ATU Operation Registers

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define which address database to use (i.e., the high part of the DBNum field).	Used to indicate the returned MAC's DBNum[7:4].
ATU Operation	0x0B (decimal 11)	Used to define the MAC's Priority, the required operation (including which database to use, i.e., the low part of the DBNum field) and start the ATU operation.	Used to indicate the ATU's Busy status, the returned MAC Priority and its DBNum[3:0].
ATU Data	0x0C (decimal 12)	Used further to define the required operation and used as the required ATU Data that is to be associated with the MAC address below.	Returns the ATU Data that is associated with the resulting MAC address below.
ATU MAC (3 registers)	0x0D to 0x0F (decimal 13 to 15)	Used to define the required MAC address upon which to operate.	Returns the resulting MAC address from the desired operation.

3.4.10.1 Format of the ATU Database

Each MAC address entry in the ATU database is 67-bits in size. The lower 48 bits contains the 48-bit MAC address and the upper 19 bits contains information about the entry as shown in Figure 13. The database is accessed 16-bits at a time via the Switch Global registers shown in the figure.

Figure 13: Format of an ATU Entry

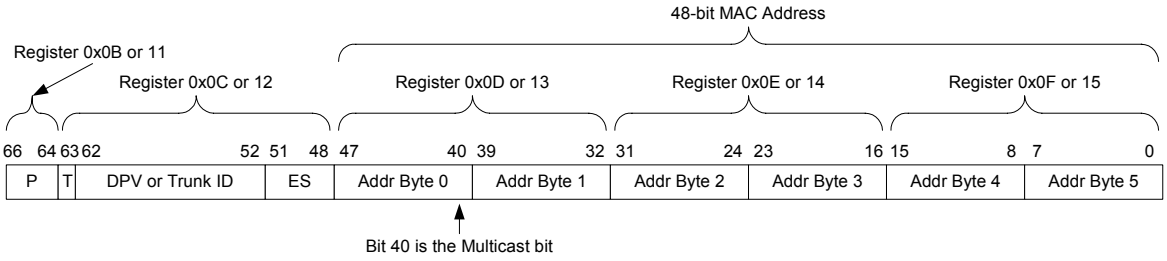

The right 48-bits in Figure 13 is the 48-bit MAC address associated with this ATU entry. The upper 19 bits is the data that is associated with the entry's MAC address. The upper 19 bits are defined as follows:

Table 25: ATU Data bits

Field	Bits	Description
P (88E6095 and 88E6095F only)	66:64	The MAC's Priority override value when enabled by the EntryState bits below. Used for priority override on ingressing frames (see Section 3.5.5). Enabling a priority on a MGMT MAC address (Multicast EntryState = 0xE or Unicast EntryState = 0xD) will override *all* priorities for these MGMT frames. Enabling a priority on a static, non-MGMT MAC address, will only override the frame's priority if the port's DAPriOverride and/or SAPriOverride bits are set to a one (in Port Control 2 register - offset 0x08).
T	63	The Trunk bit used to qualify the contents of the DPV or Trunk ID bits below. When this bit is zero bits 62:52 are the DPV (Destination Port Vector) associated with this MAC. When this bit is a one, bits 55:52 is the Trunk ID associated with this MAC (bits 62:56 must be zero in this case).
DPV or Trunk ID	62:52	The Destination Port Vector or Trunk ID. When the Trunk bit (bit 63 above) is a zero these bits indicate which port or ports are associated with this MAC address (i.e., where frames should be switched) when they are set to a one. A DPV of all zeros indicates frames with this DA should be discarded. Bit 52 is assigned to physical Port 0, 53 to Port 1, 54 to Port 2, etc. If more than one port's bit is set to a one frames mapped to this MAC address will attempt to egress out more than one port. This is used for multicast filtering. When the Trunk bit (bit 63 above) is a one bits 55:52 (the lower 4 bits) indicate the Trunk ID that is associated with this MAC address (in this case bits 62:56 must be zeros). The port or ports that frame with this DA MAC address are mapped to is determined by the contents of the Trunk Route Table (see Table 115 and global 2 offset 0x08).

Table 25: ATU Data bits (Continued)

Field	Bits	Description
EntryState	51:48	The EntryState field, together with the entry's Multicast bit (bit 40) is used to determine the entry's age or its type as follows: For unicast MAC addresses (bit 40 = 0): 0x0 = Invalid, empty or purged entry. 0x1 to 0x7 = Aging. Valid entry where the EntryState = the entry's age and the DPV indicates the port or ports or Trunk ID mapped to this MAC address. 0xD = MGMT Override. Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address. Frames with this MAC address are considered MGMT (management) frames and are handled specially inside the switch (Section 3.8) and the Priority bits (bits 66:64 above) indicate the priority to for this frame overriding *all* other priorities. Used to get unicast frames to the management CPU when the CPU is isolated with a VLAN barrier. 0xE = Static. Valid entry that is static and does not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address. 0xF = Priority Override (88E6095 and 88E6095F only). Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address and the Priority bits (bits 66:64 above) indicate the priority to use for this frame's DA or SA potentially overriding other priorities[1]. All other values are reserved for future use. For multicast MAC addresses (bit 40 = 1): 0x0 = Invalid, empty or purged entry. 0x5 = Non-Rate Limited. Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address. Frames with this MAC address are not limited by Ingress Rate Limiting (Section 3.5.15). 0x7 = Multicast Filtering. Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address. Used for multicast filtering. 0xE = MGMT Override. Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address. Frames with this MAC address are considered MGMT (management) frames and are handled specially inside the switch (Section 3.8) and the Priority bits (bits 66:64 above) indicate the priority to for this frame overriding *all* other priorities. Used to indicate a frame is a BPDU for Spanning Tree Support (Section 3.8). 0xF = Priority Override. (88E6095 and 88E6095F only) Valid entry that is static and will not age. The DPV indicates the port or ports or Trunk ID mapped to this MAC address and the Priority Override bits (bits 66:64 above) indicate the priority to use for this frame's DA or SA potentially overriding other priorities[1]. All other values are reserved for future use.

1. DA and SA priority overrides will only occur on frames that use this MAC address and only on ports that have their DAPriOverride and/or their SAPriOverride bits set (in Port Control 2 - offset 0x08).

3.4.10.2 Reading the Address Database

The contents of the address database can be dumped or searched. The dump operation is called Get Next since it returns the active contents of address database in ascending network byte order. A search operation can also be done using the Get Next operation. If multiple address databases are being used (see Section 3.4.7), set the DBNum field in the ATU Control register and the ATU Operation register to the database number to search when using the Get Next function.

The Get Next operation starts with the MAC address contained in the ATU MAC registers and returns the next higher active MAC address currently in the address database. Begin with an ATU MAC address of all ones to get the first or lowest active MAC address. The returned MAC address and its data is accessible in the ATU MAC and the ATU Data registers. To get the next higher active MAC address, the Get Next operation can be started again without setting the ATU MAC registers since they already contain the 'last' address. A returned ATU MAC address of all ones indicates that no higher active MAC addresses were found or that the Broadcast MAC address was found. In either case, the end of the database has been reached. If it were reached with a valid Broadcast address the entry's EntryState is returned with a non-zero value. A summary of how the Get Next operation uses the ATU's registers is shown in Table 26.

Table 26: ATU Get Next Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define which address database to search (i.e., the high part of the DBNum field).	No change.
ATU Operation	0x0B (decimal 11)	Used to define the required operation (including which database to search, i.e., the low part of the DBNum field) and start the ATU operation.	Used to indicate the ATU's Busy status and report the resulting MAC's priority.
ATU Data	0x0C (decimal 12)	Ignored.	Returns the ATU Data that is associated with the resulting MAC address below. If EntryState = 0x0 the returned data is not a valid entry.
ATU MAC (3 registers)	0x0D to 0x0F (decimal 13 to 15)	Used to define the starting MAC address to search. Use an address of all ones to find the 1st or lowest MAC address. Use the last address to find the next address (there is no need to write to this register in this case).	Returns the next higher active MAC address if found, or all ones are returned indicating the end of the table has been reached (all ones is a valid entry if EntryState ≠ 0x0).

To search for a particular MAC address, start the Get Next operation with a MAC address of one less than the target MAC address using the DBNum of the database to search. If the required MAC address is found it is returned in the ATU MAC registers along with its associated data in the ATU Data register. If the searched MAC address is not found active, the ATU MAC registers will not equal the target address.

3.4.10.3 Loading & Purging an Entry in the Address Database

Any MAC address (unicast or multicast) can be loaded into, or removed from, the address database by using the Load operation. An address is loaded into the database if the EntryState in the ATU Data register is non-zero. A value of zero indicates that the ATU operation is a purge.

The Load operation searches the address database indicated by the database number, DBNum (in the ATU Control register and the ATU Operation register), for the MAC address contained in the ATU MAC registers. If the address is found it is updated by the information found in the ATU Data register.

Note

A load operation becomes a purge operation if the ATU Data's EntryState equals zero. Also, static addresses can be modified without their needing to be purged first.

If the address is not found, and if the ATU Data register's EntryState does not equal zero, the address is loaded into the address database using the same protocol as automatic Address Learning (see Section 3.4.3). The 16-bits of the ATU Data register are written into bits 63:48 of the ATU entry (see Section 3.4.10.1). The 3 MACPri bits of the ATU Operation register are written into bits 66:64 of the ATU entry.

A summary of how the Load operation uses the ATU's registers is shown in Table 27.

Table 27: ATU Load/Purge Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define which address database to load or purge (i.e., the high part of the DBNum field).	No change.
ATU Operation	0x0B (decimal 11)	Used to define the operation (including the database to load or purge, i.e., the low part of the DBNum field and the entries priority bits) and start the ATU Operation.	Used to indicate the ATU's Busy status.
ATU Data	0x0C (decimal 12)	Used to define the associated data that is loaded with the MAC address below. When EntryState = 0, the load becomes a purge.	No change.
ATU MAC (3 registers)	0x0D to 0x0F (decimal 13 to 15)	Used to define the MAC address to load or purge.	No change.

3.4.10.4 Flushing Entries

All MAC addresses or just the unlocked MAC addresses can be purged from the entire set of address databases or from just a particular address database using single ATU operations. These ATU operations are:

- Flush All Entries
- Flush all Non-Static Entries
- Flush All Entries in a particular DBNum Database
- Flush all Non-Static Entries in a particular DBNum Database

The Flush requires that the EntryState bits in the ATU Data register be 0x0. The ATU MAC Address registers are not used for these operations and they are left unmodified. The DBNum field of the ATU Control and the ATU Operation register is used for the Flush operations that require a database number to be defined.

3.4.10.5 Servicing ATU Violations

The ATU captures ATU Full, SA Member Violation and SA Miss Violation data. An ATU Full violation occurs if an Automatic Address Learn (Section 3.4.3) or an ATU load operation (Section 3.4.10.3) could not enter the new MAC address into the address database owing to all four bins at the MAC address's hashed address being locked as static entries.

An SA Membership Violation occurs when a frame's SA is found in the address database and the entry is static or the port is locked (see the Port Association Vector register, offset 0x0B), and the port's bit is not set in the entry's Destination Port Vector (DPV). This interrupt can be masked on a per-port basis by setting the port's IgnoreWrongData bit (see the Port Association Vector register, offset 0x0B) to a one.

An SA Miss Violation occurs when a frame's SA is not found in the address database and the port is locked due to the port's LockedPort bit being set to a one (see the Port Association Vector register, offset 0x0B). This interrupt can be masked on a per-port basis by clearing the port's LockedPort bit to zero.

Captured ATU Violations and their associated interrupts are cleared by the Get/Clear Violation Data ATU Operation. This ATU Operation returns the source port that caused the violation in the EntryState/SPID field of the ATU Data register and returns the MAC address that caused the violation in the ATUByte[5:0] fields of the ATU MAC register including the DBNum that was associated with the frame (in the ATU Control and ATU Operations registers.

A summary of how the Get/Clear Violation Data operation uses the ATU's registers is shown in Table 28.

Table 28: ATU Get/Clear Violation Data Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (Decimal 10)	Ignored.	Used to indicate the violations' DBNum[7:4].
ATU Operation	0x0B (Decimal 11)	Used to define the desired operation and start it	Used to indicate the ATU's Busy status and the source of the violation and its DBNum[3:0].
ATU Data	0x0C (Decimal 12)	Ignored.	Used to indicate the SPID that was involved in the violation
ATU MAC (3 registers)	0x0D to 0x0F (Decimal 13 to 15)	Ignored.	Used to indicate the MAC that was involved in the violation

3.5 Ingress Policy

The Ingress Policy block is used to modify the normal packet flow through the switch. All ports have identical capabilities. The frame's source MAC address can be authenticated and the frame potentially discarded for 802.1X support. The content of each frame is examined in more detail for Quality of Service (QoS) priority information. Port based VLANs or 802.1Q VLANs are used to prevent a frame from going out of certain ports, and switch management Port States, 802.1s (per VLAN spanning tree) or 802.1Q are used to prevent the frame from entering the switch at all. An optional DSA Tag mode allows a management CPU to override the switch giving the CPU complete control over those frames it wants to control. The DSA Tag mode is used to support all switch management features across multiple devices under a central management CPU (cascading) and to support stacking across multiple boxes (88E6095 and 88E6095F only). An optional Marvell Header mode supports faster CPU routing. Ingress storm limiting is supported, along with a Double Tagging option (Q-in-Q) - see Section 3.5.18 "Ingress Double VLAN Tagging (88E6095 and 88E6095F Only)" - and IGMP/MLD snooping.

3.5.1 Port States for 802.1D Spanning Tree

The devices support four 802.1D Port States per port shown in Table 29 (802.1s per VLAN Port States are also supported - see Section 3.5.10). The 802.1D Port States are used by the Queue Controller (see Section 3.6) in the devices to adjust buffer allocation. They are used by the Ingress Policy blocks to control which frame types are allowed to enter and leave the switch so that Spanning Tree or bridge loop detection software can be supported. The PortState bits in the Port Control register (offset 0x04) determine each port's Port State and the bits can be modified at any time.

Table 29 lists the Port States and describes them. Two of the Port States require the detection of management (MGMT) frames. MGMT frames are defined in Section 3.4.8. Their primary purpose is to support the Spanning Tree Protocol (see Section 3.10) so these frames have the ability to tunnel through blocked ports. SA learning is not performed on MGMT frames. MGMT frames also ignore VLAN rules on ingress and egress (802.1Q and Port Based), IGMP snooping and Rate Limiting (unless MGMT frames are selected for limiting). This means they always go to the port indicated by the Destination Port Vector (DPV) assigned to the frame's DA in the address database or to the port's CPUPort (Port Control 2, offset 0x08) depending upon what MGMT detection mode was used (Section 3.4.8). MGMT frames are typically used for 802.1D Spanning Tree Bridge Protocol Data Units (BPDUs), but any multicast address can be used supporting new or proprietary protocols.

Table 29: Port State Options

Port State	Description
Disabled	Frames are not allowed to enter (ingress) or leave (egress) a Disabled port. Learning does not take place on Disabled ports.
Blocking/Listening	Only MGMT frames are allowed to enter (ingress) or leave (egress) a Blocked port. All other frame types are discarded. Learning is disabled on Blocked ports.
Learning	Only MGMT frames are allowed to enter (ingress) or leave (egress) a Learning port. All other frame types are discarded but learning takes place on all good non-MGMT frames that are not discarded owing to 802.1Q security violations (Section 3.5.7) or 802.1X authentication violations (Section 3.5.2).
Forwarding	Normal operation. All frames are allowed to enter (ingress) and leave (egress) a Forwarding port. Learning takes place on all good non-MGMT frames that are not discarded owing to 802.1Q security violations (Section 3.5.7) or 802.1X authentication violations (Section 3.5.2).

The default Port State for all the ports in the switch can be either Disabled or Forwarding depending upon the value of the SW_MODE pins (Table 15). The ports come up in the Forwarding Port State unless the SW_MODE is for CPU attached mode. This allows the CPU to bring up the ports slowly, running bridge loop detection software along the way.

Note

Managed switches must always use the CPU attached mode setting of the SW_MODE pins. But even these designs may want a jumper or test point on the PCB such that the switch can be made to reset to the Stand alone mode setting of the SW_MODE pins. This can help initial PCB debug and/or manufacturing test of the switch portion of the design as the switch will power up forwarding frames everywhere without any software.

3.5.2 802.1X MAC Address Authentication

All non-MGMT[1] (non-management) frames received on a port with an unauthorized Source MAC Address are discarded if the port's DropOnLock bit is set to a one (Port Control register, offset 0x04). In this mode, only frames with an authorized SA (or MGMT frames) are allowed into the switch to be processed further. An SA is considered authorized if it is present in the address database and the source port where the frame entered the switch has its bit set to a one in the SA's Destination Port Vector (the DPV in the SA's ATU entry – Section 3.4.10). The policy of how these addresses are entered into that address database is controlled by the ATU (Section 3.4.6). It is recommended that CPU Directed Address Learning (Section 3.4.5) be used on ports supporting MAC based 802.1X authentication (i.e., ports with their DropOnLock bit being set to a one).

3.5.3 Forward Unknown/Secure Port

The devices can be configured to prevent the forwarding of unicast frames with an unknown destination address (i.e., the address is not present in the address database – Section 3.4.1). The forwarding can be prevented on a per port basis (by clearing the port's ForwardUnknown bit in the Port Control register, offset 0x04) so that frames with unknown unicast addresses only go out from the port or ports where a server or router is connected.

This, together with the disabling of automatic address learning on a port (Section 3.4.3), allows a port to be configured as a Secure Port. A Secure Port allows communications to and from approved devices (by MAC address) only. In this mode all approved devices need to have their MAC address loaded as static into the address database (Section 3.4.10). When a new device tries to access the network through a Secure Port that new device's address is not automatically learned (learning must be disabled on Secure Ports) but its frame can progress to the server. When the server replies by sending a unicast frame back to the new device's address, that frame does not make it to the new device since the new device's address is not in the address database and frames with unknown unicast addresses are not allowed to egress the Secure Port. This effectively ends the communication between the un-approved new device and the rest of the network. Secure Port is similar to the 802.1X without the authentication server and the associated interrupts to the CPU.

3.5.4 Forward Unknown for Multicasts

The devices can be configured to prevent the forwarding of multicast frames with an unknown destination address (i.e., the address is not present in the address database – Section 3.4.1). The forwarding can be prevented on a per port basis (by clearing the port's DefaultForward bit in the Port Control 2 register, offset 0x08) so that frames with unknown multicast addresses only go out from the port or ports where a server or router is connected.

1. MGMT frames and their Ingress Rule changes are covered in Section 3.4.8.

3.5.5 Quality of Service (QoS) Classification

The Ingress Policy block has the task of determining the priority of each frame for the Queue Controller. The Ingress Policy block does not perform QoS switching policy, which is the task of the Queue Controller (Section 3.6). Instead, it has the job of determining the priority of each frame for the Queue Controller. The priority of a frame is determined by (in priority order):

1. The CPU's DSA Tag, if enabled on the port, and the frame is not a Forward type. In this case, the upper 2 bits of the DSA Tag's PRI field is used as the QoS priority for the frame. See Section 3.5.16 for more information.
2. The frame's DA is classified to make this frame a MGMT frame by either of two methods. The priority of the frame comes from the method that determined the frame was MGMT (Section 3.4.8).
3. The frame's DA is in the address database with a priority defined and DA priority override is enabled on the port. See Section 3.4.10 and DAPriOverride in Port Control 2, offset 0x08.
4. The frame's SA is in the address database, loaded as static, with a priority defined and with SA priority override enabled on the port. See Section 3.4.10 and SAPriOverride in Port Control 2, offset 0x08.
5. The frame's VID is in the VLAN database with a priority defined and VTU priority override is enabled on the port. See Section 3.5.11 and VTUPriOverride in Port Control 2, offset 0x08.
6. The IEEE 802.3ac Tag containing IEEE 802.1p priority information if enabled on the port. 802.3ac tagging is enabled by default on all of the ports. See Section 3.5.12 and UseTag in Port Control, offset 0x04.
7. The IPv4 Type of Service (TOS)/DiffServ field or IPv6 Traffic Class field if they are enabled on the port[1]. the default case is enabled on all of the ports. See Section 3.5.13 and UseIP in Port Control, offset 0x04.
8. The Port's default priority defined in DefPri (see the Default Port VLAN ID and Priority register, offset 0x07).

There are enables on each of the priority classification rules giving the designer any combination that is needed. For instance, if a, port based higher priority port is required for a switch design, priority classifications 1 to 7, above, can be disabled. This leaves priority classification 8 only (the Port's default) resulting in all Ingressed frames being assigned the same priority on that port.

The priority classification rules can be disabled separately on a per-port basis. This allows each port in the switch to be configured to use different rules for priory classification (see the UseTag and UseIP bits in the Port Control registers – offset 0x04, and the VTUPriOverride, SAPriOverride and DAPriOveride bit in the Port Control 2 registers – offset 0x08). Priority classification rules #6 (IEEE Tag) and #7 (IP) can be reversed on a per port basis as well (see the TagIfBoth bit in the Port Control register - offset 0x04). This allows the priority to be selected for frames that are both Tagged and IP.

The CPU's DSA Tag, if the Tag is a From_CPU, is used over any other priority that may be in the frame so the CPU can have ultimate control over every frame it transmits into the switch. The other non-Forward DSA Tag types maintain the QoS value passed in the frame. The DA address priority is used to set the highest priority on BPDU's or other management frames (called MGMT in this document) since these frames should never be dropped. A port's DAPriOverride bit does not need to be set for a DA priority override to occur on MGMT[2] frames determined with the ATU. SA priority and VID priority can be used to give frames a different priority depending upon where they came from or the VLAN they belong to (useful for IP phones).

IEEE Tagged frames that have their priority overridden and/or re-mapped[3] during Ingress, will egress with the overridden priority in the frame's IEEE Tag, if the frame egresses the switch tagged.

1. IPv4/IPv6's priority classification can be configured on a per port basis to have a higher priority than IEEE Tag by the port's TagIfBoth bit in Port Control, offset 0x04.
2. MGMT frames are determined by the frame's DA - Section 3.4.8.
3. Each port supports an 8x3 priority re-mapping table (offsets 0x18 and 0x19) that can be used to scale a port's priorities up or down separately from the other ports.

3.5.6 Port Based VLANs

The devices support a very flexible port based VLAN system that is used for all non-MGMT frames even if 802.1Q is enabled on the port.

Each Ingress port contains a register that restricts the output (or egress) ports to which it is allowed to send frames. This register is called the VLANTable register (offset 0x06). If bit 0 of a port's VLANTable register is set to a one, that port is allowed to send frames to Port 0. If bit 1 of a port's registers is set to a one, that port is allowed to send frames to Port 1. Bit 2 for Port 2, etc. At reset the VLANTable register for each port is set to a value of all one's, except for each port's own bit, which is cleared to a zero (this prevents frames from going back out of the port they came in on[1]). This default VLAN configuration allows all the ports to send frames to all the other ports as shown in Figure 14.

Figure 14: Switch Operation with VLANs Disabled

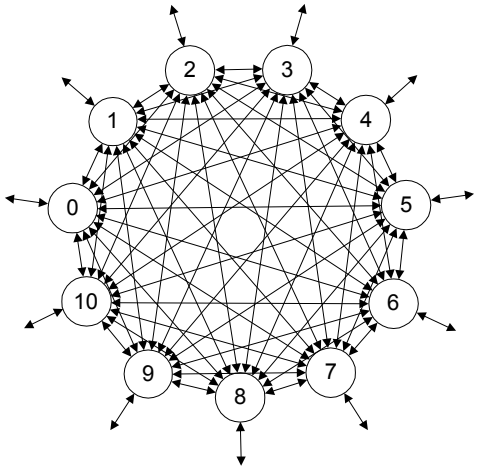

3.5.6.1 Port Based VLAN Router Examples

One of the main applications for port based VLAN support in the devices is to isolate a port or ports for firewall router applications. Figure 15 shows a typical VLAN configuration for a firewall router. Port 9 is used as the WAN port (since it can be either fiber, or copper). The data coming in from this WAN port must not go out to any of the LAN ports – but it must be able to go to the router CPU. All the LAN ports are able to send frames directly to each other without the need of CPU intervention – but they cannot send frames directly to the WAN port. The CPU is able to send frames to all of the ports so that routing can be accomplished. The use of the Marvell Header, (Section 3.5.17) enables a CPU to define dynamically which port or ports a particular frame is allowed to reach for purposes of WAN and LAN isolation on multicast traffic generated by the CPU.

1. The devices allow a port's own bit in its VLANTable to be set to a one – see Section 3.5.12

Figure 15: Switch Operation with a Typical Router VLAN Configuration

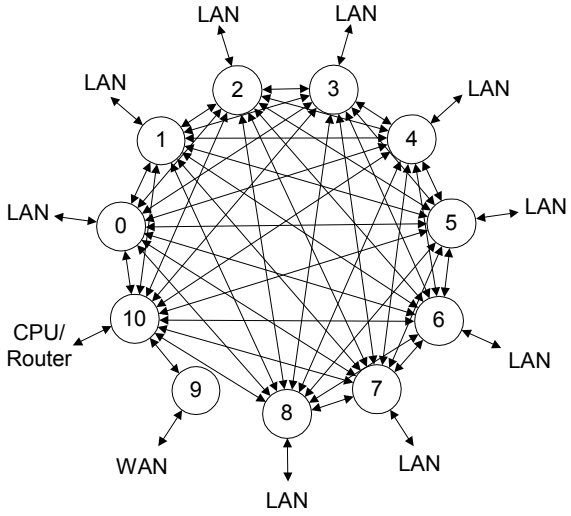

The example VLAN configuration shown in Figure 15 is achieved by setting the port's VLANTable registers as shown in Table 30:

Table 30: VLANTable Settings for Figure 15

Port #	Port Type	VLANTable Setting
0	LAN	0x5FE
1	LAN	0x5FD
2	LAN	0x5FB
3	LAN	0x5F7
4	LAN	0x5EF
5	LAN	0x5DF
6	LAN	0x5BF
7	LAN	0x57F
8	LAN	0x4FF
9	WAN	0x400
10	CPU	0x3FF

To show the flexibility of the devices VLAN configuration options, Figure 16 shows another example. In this case, the switch is divided into three independent VLANs connected to a common router.

Figure 16: Switch Operation with another Example VLAN Configuration

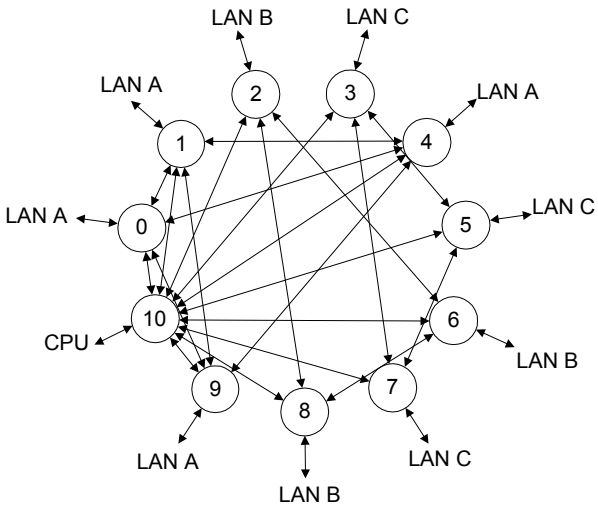

The example VLAN configuration shown in Figure 16 is accomplished by setting the port's VLANTable registers as shown in Table 31:

Table 31: VLANTable Settings for Figure 16

Port #	Port Type	VLANTable Setting
0	LAN A	0x612
1	LAN A	0x611
2	LAN B	0x540
3	LAN C	0x4A0
4	LAN A	0x603
5	LAN C	0x488
6	LAN B	0x504
7	LAN C	0x428
8	LAN B	0x444
9	LAN A	0x413
10	CPU	0x3FF

3.5.7 802.1Q VLANs

The devices support 802.1Q (and cross-chip port based VLANs, Section TBD) with the full set of 4,096 different VID (VLAN identifiers). Some or all of the VIDs can be used (i.e., software only needs to initialize the VIDs that are being used[1]). Since the device may be programmed with only a subset of the possible VIDs, and security requirements vary, the devices support 802.1Q in three different modes. The device's port-based VLAN feature (Section 3.5.6) is in effect for all 802.1Q modes described below.

3.5.7.1 Security & Port Mapping

The 802.1Q Security features of the devices support the discarding of ingressing frames that don't meet the security requirements and ensuring that those frames that do meet the requirements are sent to the allowed ports only. Three levels of security are supported and they can be set differently on each port. The security options are processed using the ingressing frame's VID or the ingress port's Default VID[2] as follows:

- Secure – The VID must be contained in the VTU and the Ingress port must be a member of the VLAN else the frame is discarded. The frame is allowed to exit only those ports that are both:
  - Members of the frame's VLAN *and*
  - Included in the source port's VLANTable (see Section 3.5.6)
- Check – The VID must be contained in the VTU or the frame is discarded (the frame will not be discarded if the Ingress port is not a member of the VLAN). The frame is allowed to exit only those ports that are both:
  - Members of the frame's VLAN *and*
  - Included in the source port's VLANTable (see Section 3.5.6)
- Fallback – Frames are not discarded if their VID is not contained in the VTU. If the frame's VID is contained in the VTU, the frame is allowed to exit only those ports that are both:
  - Members of the frame's VLAN *and*
  - Included in the source port's VLANTable (see Section 3.5.6)

  If the frame's VID is not contained in the VTU, the frame is allowed to exit only those ports that are:
  - Included in the source port's VLANTable (see Section 3.5.6)

Secure, Check, Fallback or 802.1Q Disabled modes for the port are controlled by the port's 802.1QMode bits (Port Control 2 register - offset 0x08).

3.5.7.2 Security Violations

If 802.1Q is enabled on a port, security violations are captured and an interrupt can be generated to the CPU (if unmasked by the VTUProbIntEn bit (Switch Global Control, global offset 0x04). This is true regardless of the 802.1Q mode (VTUProb interrupts will not occur from a port if the port's 802.1QMode is Disabled - Port Control 2 register, offset 0x08). The interrupts (up to one at a time) are captured by the VLAN Translation Unit (see VTU Operation register, global offset 0x05). Two kinds of security violations are captured. A MissViolation occurs if a frame's VID is not contained in the VTU. A MemberViolation occurs if a frame's VID is in the VTU but the source port of the frame is not a member of the frame's VLAN. The security violation captures the offending source port (SPID) and VID that caused the violation. This data is accessed by executing a Get/Clear Violation Data operation in the VTU.

3.5.7.3 Security Override of a Frame's VID

A Tagged frame's VID can be forced to the port's DefaultVID.

1. All ports are considered members for VIDs that are unused or invalid in the VTU.
2. The port's Default VID is used if the frame is not 802.3ac Tagged, or if the frame's VID is 0x000 or if the port's ForceDefaultVID bit is set (Default Port VLAN ID and Priority register - offset 0x07).

3.5.8 Switching Frames Back to their Source Port

The devices support the ability to send frames back out of the port on which they arrived. While this is not a standard way to handle Ethernet frames, some applications, like 802.3ah OAM loopback (Section 3.4.9.2), may require this ability on some ports. This feature can be enabled on a port by port basis by setting the port's own bit to a one in its VLANTable register (in the Port Based VLAN Map register, offset 0x06). This function is valid if 802.1Q is enabled on the port or not.

3.5.9 Tunneling Frames Through VLANs

Normally frames cannot pass between port-based VLANs nor 802.1Q VLANs. The devices can be configured to allow some frames to do so. Before a frame can tunnel through a VLAN barrier, its DA address must be loaded as static into the address database (see Section 3.4.10.1) and the VLANTunnel bit on the frame's Ingress port must be set to a one (see Port Control register, offset 0x04). When both of these conditions are true, the frame is sent out of the port or ports indicated in the static address's Destination Port Vector (the DPV field for the DA entry in the address database). The VLANTable and 802.1Q membership data is ignored in this case. This feature is enabled only on those ports that have their VLANTunnel bit set to a one.

3.5.10 802.1s Per VLAN Spanning Tree (88E6095 and 88E6095F Only)

The devices support per VLAN Port States for 802.1s (per VLAN Spanning Tree). Each VID entry in the VTU (Section 3.5.11) has two bits per port for 802.1s Port State information as shown in Table 32.

Table 32: 802.1s Port State Options

Port State	Description
802.1s Disabled	The port's PortState bits in the Port Control register (offset 0x04) are used for this port for frames with this VID. See Section 3.5.1.
Blocking/Listening	Only MGMT frames are allowed to enter (ingress) or leave (egress) this port for frames with this VID. All other frame types are discarded. Learning is disabled on Blocked ports.
Learning	Only MGMT frames are allowed to enter (ingress) or leave (egress) this port for frames with this VID. All other frame types are discarded but learning takes place on all good non-MGMT frames that are not discarded owing to 802.1Q security violations (Section 3.5.7) or 802.1X authentication violations (Section 3.5.2).
Forwarding	Normal operation. All frames are allowed to enter (ingress) and leave (egress) this port for frames with this VID. Learning takes place on all good non-MGMT frames that are not discarded owing to 802.1Q security violations (Section 3.5.7) or 802.1X authentication violations (Section 3.5.2).

If 802.1s is not being used, all VTU entries must use the 802.1s Disabled setting for all ports for all VIDs. If 802.1s is being used, all VTU entries must be configured with the required 802.1s Port State for each port for each VID, or, a mixture can be used. Some VID entries on some ports could be using the 802.1s Port States while the other ports in a VID entry, and/or the other VID entries, can use the 802.1s Disabled port state. Those ports/VIDs using the 802.1s Disabled port state will use the port's Port State setting instead (Section 3.5.1).

The 802.1s Port State options take precedence over the port's PortState bits settings (Section 3.5.1), with the exception of the port's Disabled Port State (set in the port's PortState bits, Port Control, offset 0x04). The port's Disabled Port State prevents all frames from entering and leaving the port so that has precedence over 802.1s Port States.

3.5.11 VLAN Translation Unit Operations

The VLAN Translation Unit (VTU) in the devices support user commands to access and modify the contents of the VLAN membership database.

All VTU operations have the same user interface and protocol. Global registers are used and are shown in Table 33. The protocol for an VTU operation is as follows:

- Ensure the VTU is available by checking the VTUBusy bit in the VTU Operation register. The VTU can only perform one user command at a time.
- Load the VTU Data and VTU VID registers if required by the desired operation.
- Start the VTU operation by defining the required DBNum, and VTUOp and setting the VTUBusy bit to a one in the VTU Operation register – this can be done with a single write operation.
- Wait for the VTU operation to complete. This can be done by polling the VTUBusy bit in the VTU Operation register or by receiving an VTUDone interrupt (see Switch Global Control, global offset 0x04, and Global Status, offset 0x00).
- Read the required results if appropriate.

Table 33: VTU Operation Registers

Register	Offset	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Used to define the required operation (including which database to associate with this VID) and start it.	Used to indicate the VTU's Busy status and violation status including source of the violation.
VTU VID	0x06	Used to further define the required operation and used as the VID that is to be operated on.	Returns the VID from the desired operation.
VTU Data (3 registers)	0x07 to 0x09	Used to define the required data that is to be associated with the required VID, including VTU Priority Override.	Returns the associated data from the desired operation.

3.5.11.1 Format of the VTU Database

Each VID entry in the VTU database contains:

- A 1-bit valid indicator (Valid),
- An 8-bit Database number (DBNum),
- 4-bits of VTU Priority Override data (VIDPRI[2:0] & VIDPRIOverride) and
- 4 bits of VTU Data per port

The format of a VTU entry is shown in Figure 17and Table 34. The database is accessed 16-bits at a time via the Switch Global registers shown in Figure 17 (not all the register bits are shown). For more information about these register see the VTU Operation register (global offset 0x05) and VTU Data registers for all ports (global offsets 0x06 to 0x09).

Figure 17: Format of a VTU Entry

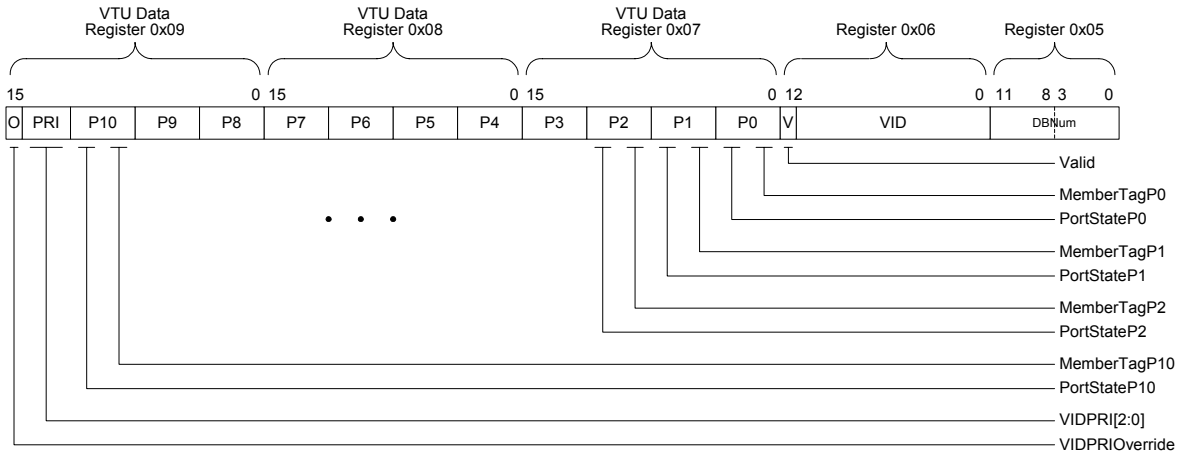

Table 34: VTU Entry Format

Field	Bits	Description
O (88E6095 and 88E6095F devices only)	15 in Reg 0x09	VID Priority Override. This bit is used to indicate that frames assigned with this VID can have their priority overridden with the PRI bits below if the port's VTUPriOverride bit is set (see Port Control 2 register, offset 0x08).
PRI (88E6095 and 88E6095F devices only)	14:12 in Reg 0x09	VIDPRI. These bits are used as a frame's priority for frames assigned with the entry's VID if the VIDPRIOverride bit above is set to a one (and if it is enabled on the port -see Port Control 2 register, offset 0x08).

Table 34: VTU Entry Format

Field	Bits	Description
MemberTag	**In Reg 0x07:** Port 0 1:0 Port 1 5:4 Port 2 9:8 Port 3 13:12 **In Reg 0x08:** Port 4 1:0 Port 5 5:4 Port 6 9:8 Port 7 13:12 **In Reg 0x09** Port 8 1:0 Port 9 5:4 Port 10 9:8	The lower two bits of each port's VTU data is called MemberTag. These bits are used to indicate which ports are members of the VLAN and if these VLANs frames should be tagged or untagged, or unmodified when exiting the port as follows: • 00 = Port is a member of this VLAN and frames are to egress unmodified • 01 = Port is a member of this VLAN and frames are to egress Untagged • 10 = Port is a member of this VLAN and frames are to egress Tagged • 11 = Port is not a member of this VLAN
PortState	**In Reg 0x07:** Port 0 3:2 Port 1 7:6 Port 2 11:10 Port 3 15:14 **In Reg 0x08:** Port 4 3:2 Port 5 7:6 Port 6 11:10 Port 7 15:14 **In Reg 0x09** Port 8 3:2 Port 9 7:6 Port 10 11:10	The upper two bits of each port's VTU data is called 802.1s PortState. These bits are used to support 802.1s per VLAN spanning tree as follows: • 00 = 802.1s Disabled. Use non-VLAN Port States (in Port Control, offset 0x04) for this port for frames with this VID • 01 = Blocking/Listening Port State for this port for frames with this VID • 10 = Learning Port State for this port for frames with this VID • 11 = Forwarding Port State for this port for frames with this VID **NOTE:** These 802.1s PortState bits take precedence over the port's Port State bits (in Port Control, offset 0x04) unless the port's Port State is Disabled (which prevents all frames from flowing).
Valid	12 in Reg 0x6	Valid bit. This bit is used to indicate that the below VID and its associated data is valid in the VTU's database and should be used. After a hardware reset, all 4096 entries in the table are considered invalid (the Valid bit on each entry is cleared).
VID	11:0 in Reg 0x6	VLAN ID. These bits indicate the VID number that is associated with the MemberTag data, 802.1s Port State data, VTU Priority and its override and the entry's address database number (DBNum).
DBNum	11:8 and 3:0 in Reg 0x5	DataBase Number. If separate address databases are used, these bits indicate the address database number to use for all frames assigned with the VID (see Section 3.4.7). All MAC DA look-ups and SA learning will refer to the address database number defined by the DBNum associated with the frame's VID. Multiple VID's can use the same DBNum. If separate address databases are not used DBNum must be written as zeros.

3.5.11.2 Reading the VLAN Database

The contents of the VLAN database can be dumped or searched. The dump operation is called Get Next since it returns the active contents of the VLAN database in ascending VID order. A search operation can also be done using the Get Next operation.

The Get Next operation starts with the VID contained in the VTU VID register and returns the next higher active VID in the VLAN database. Use a VID of all ones to get the first or lowest active VID. The returned VID and its data are accessible in the VTU Operation, VTU VID and the VTU Data registers. To get the next higher active VID, the Get Next operation can be started again without setting the VID registers since it already contains the last address. A returned VID of all ones indicates that no higher active VID was found or that the VID value of 0xFFF was found. In either case, it indicates that the end of the database has been reached. If it were reached with a valid VID of 0xFFF the entry's Valid bit is returned set to one. A summary of how the Get Next operation uses the VTU's registers is shown in Table 35.

Table 35: VTU Get Next Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Used to define the required operation and start it. Used to indicate the VTU's Busy status	Returns the DBNum associated with the VID
VTU VID	0x06	Used to define the starting VID to search. Use VID of all ones to find the first or lowest VID. Use the last address to find the next address (there is no need to write to this register in this case)	Returns the next higher active VID if found, or all ones are returned indicating the end of the table has been reached (all ones is a valid entry if the Valid bit = 1)
VTU Data (3 registers)	0x07 to 0x09	Ignored	Returns the VTU Data that is associated with the VID above. If the Valid bit = 0 the returned data is not a valid entry

To search for a particular VID, start the Get Next operation with a VID one less than the target VID. If the target VID is found, it is returned in the VTU VID register along with its associated data in the VTU Data register and VTU Operation register. If the target VID is not found active, the VID register contents do not equal the target VID.

3.5.11.3 Loading and Purging an Entry in the VLAN Database

Any VID can be loaded into or removed from the VLAN database by using the Load/Purge operation. A VID is loaded into the database if the Valid bit in the VTU VID register (global offset 0x06) is a one. A value of zero in the Valid bit indicates that the VTU operation is a purge and that the defined VID and its data are to be removed from the database (if they exist).

The Load operation accesses the VLAN database using the VID contained in the VTU VID register. If the VID in the database is found valid, it is updated by the information found in the VTU Data register and the VTU Operation register (the DBNum field only).

Note

A load operation becomes a purge operation if the VTU Valid bit equals zero causing the entry's Valid bit to be cleared.

If the VID in the database is not valid, and if the VTU Valid bit equals one, then the VID, along with its data, will be loaded into the VLAN database.

A summary of how the Load operation uses the VTU's registers is shown in Table 36.

Table 36: VTU Load/Purge Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Used to define the operation (including which database to associate with the VID on loads, i.e., the DBNum field) and start it.	Used to indicate the VTU's Busy status.
VTU VID	0x06	Used to define the VID to load or purge and to define if the operation is a load or a purge (Valid = 1 means load).	No change.
VTU Data (3 registers)	0x07 to 0x09	Used to define the associated data that will be loaded with the VID above.	No change.

3.5.11.4 Flushing Entries

All VID's in the VLAN database can be purged by a singe Flush All Entries VTU Operation. The VTU VID and VTU Data registers are not used by the Flush command.

3.5.11.5 Servicing VTU Violations

The VTU captures VID Member Violation and VID Miss Violation data. A VID Membership Violation occurs when an 802.1Q enabled port receives a frame whose VID is contained in the VLAN database (VTU) but where the source port is not a member of that VLAN. A VID Miss Violation occurs when an 802.1Q enabled port receives a frame whose VID is not contained in the VLAN database (VTU).

Captured VTU Violations and their associated interrupts are cleared by the Get/Clear Violation Data VTU Operation. This VTU Operation returns the source port number that caused the violation in the DBNum[3:0]/SPID field of the VTU Operation register (global offset 0x05) and returns the VID that caused the violation in the VID field of the VTU VID register (global offset 0x06).

A summary of how the Get/Clear Violation Data operation uses the VTU's registers is shown in Table 37.

Table 37: VTU Get/Clear Violation Data Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Used to define the desired operation and start it.	Used to indicate the VTU's Busy status, the type of violation and the source port of the violation.
VTU VID	0x06	Ignored	Used to indicate the VID that was involved in the violation
VTU Data (3 registers)	0x07 to 0x09	Ignored	No change

3.5.12 IEEE Tagged Frame Handling

The format of an IEEE tagged frame is shown in Figure 18.

Figure 18: IEEE Tag Frame Format

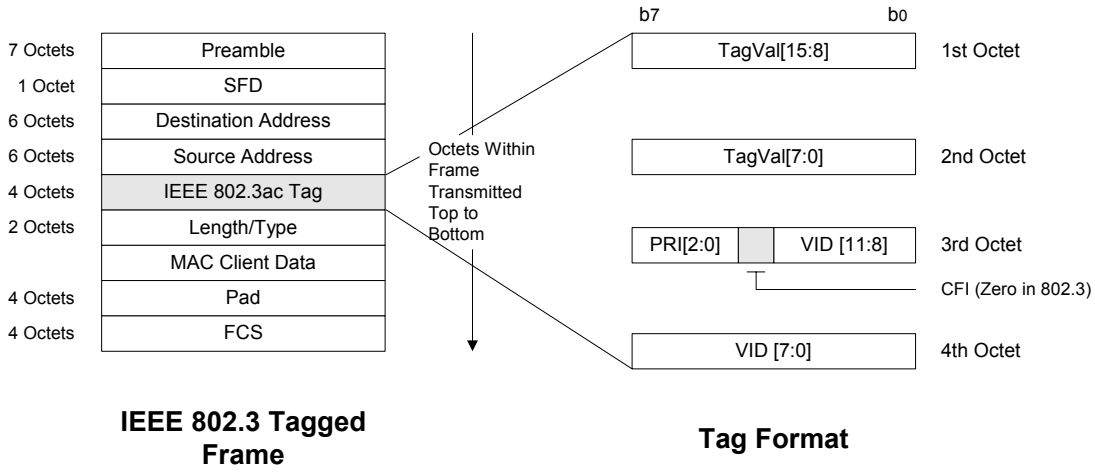

3.5.12.1 Determining if a Frame is Tagged

A frame is considered tagged in the devices if the two bytes after the frame's SA match TagVal. TagVal equals 0x8100 if the port's UseCoreTag bit is a zero (in Port Control, offset 0x04). If the port's UseCoreTag bit is a one then TagVal equals CoreTagType (in the Core Tag Type register, global offset 0x19).

Clearing UseCoreTag to zero is for normal IEEE 802.3ac tagged frame detection (frame's EtherType must be 0x8100 for it to be considered tagged). Setting UseCoreTag to one is for Provider Networks where the EtherType of Provider Tagged frame may be different from 0x8100. This feature supports Provider Network tags in the center of the network (where the Provider Tag is not removed - Section 3.5.12.1) or at the edge (where the Provider Tag is removed - Section 3.5.18.1).

3.5.12.2 Discarding UnTagged Frames

The devices support the discarding of frames that are *not* IEEE 802.3ac tagged on a port-by-port basis (DiscardUnTagged in Port Control 2, offset 0x08). A frame is considered tagged if the two bytes after the frame's SA equals TagVal which is defined in Section 3.5.12.1.

Priority only tagged frames (frames whose VID = 0x000) are considered tagged in this case and will *not* get discarded. But there is a way to discard these frames too if required. Priority only tagged frames are physically tagged (the data in the frame beyond the tag is offset by four bytes) but they are logically considered untagged in that the VID assigned to the frame is set to the port's DefaultVID (Default Port VLAN ID & Priority register, offset 0x07). If these frames egress the switch tagged their VID will be overwritten with the port's DefaultVID.

To discard priority only tagged frames along with all other untagged frames, set the ports DiscardUnTagged bit to a one, set the port's DefaultVID to 0x000, and add VID 0x000 to the VTU with all ports being non-members of the VLAN.

3.5.12.3 Discarding Tagged Frames

The devices support the discarding of IEEE 802.3ac tagged frames on a port by port basis (DiscardTagged in Port Control 2, offset 0x08). A frame is considered tagged if the two bytes after the frame's SA equals TagVal which is defined in Section 3.5.12.1.

Priority only tagged frames (frames whose VID = 0x000) are considered tagged in this case and *will* get discarded. But there is a way to keep these frames if required. Priority only tagged frames are physically tagged (the data in the frame beyond the tag is offset by four bytes) but they are logically considered untagged in that the VID assigned to the frame is set to the port's DefaultVID (Default Port VLAN ID & Priority register, offset 0x07). If these frames egress the switch tagged their VID will be overwritten with the port's DefaultVID.

To keep priority only tagged frames along with all untagged frames, clear the ports DiscardTagged bit to a zero, set the port's DefaultVID to 0x000, add VID 0x000 to the VTU with all ports being members of the VLAN with frames egressing UnTagged or Unmodified and add all other VID values (0x001 to 0xFFF) to the VTU with all ports being non-members.

3.5.12.4 Priority Extraction

The PRI[2:0] bits from tagged frames (Section 3.5.12.1) will be extracted if the port's UseTag bit is a one (in Port Control, offset 0x04). The devices re-map the frame's PRI bits and uses them as the frame's priority. The 8x3 re-mapping table is independent per port (see the Port IEEE Priority registers - offsets 0x18 and 0x19). IEEE tagged frames that get their priority re-mapped during ingress, will egress with the re-mapped priority in the frame's IEEE Tag, if the frame egresses the switch tagged, and if 802.1Q was enabled on the ingress port (Section 3.5.12.7).

The IEEE Tag supports eight priorities while the devices support four. So the Ingress Policy block takes the IEEE re-mapped PRI[2:0] bits and maps them into the two priority bits passed to the Queue Controller. This is done using the data found in the global IEEE-PRI register (global offset 0x18).

3.5.12.5 Priority Override (88E6095 and 88E6095F devices only)

If the port's UseTag bit is cleared to a zero (in Port Control, offset 0x04) the PRI bits from tagged frames are ignored. The priority of the frame is then determined by other means (Section 3.5.5). If all the other means to determine the priority of a frame are disabled, the port's DefPri (default priority in Default Port VLAN ID & Priority register, offset 0x07) is used instead. If these tagged frames egress tagged their PRI bits will be overwritten with the determined priority of the frame, if 802.1Q was enabled on the ingress port (Section 3.5.12.7).

3.5.12.6 CFI Bit

The CFI bit of the IEEE Tag is ignored and left unmodified by the devices.

3.5.12.7 VID Extraction

If any of the three supported 802.1Q modes in enabled on this port (Section 3.5.7) the VID read from tagged frames (as determined by Section 3.5.12.1) is assigned to the frame. If the frame's VID = 0x000 the port's Default-VID (from the Default Port VLAN ID & Priority register, offset 0x07) is assigned to the frame instead. If these tagged frames egress tagged their VID bits will be overwritten with the assigned value.

If 802.1Q is disabled on this port (see Section 3.5.7) the VID bits from tagged frames (as determined by Section 3.5.12.1) are ignored, and tagged frames are considered untagged for egress tag processing (i.e., the logic that determines if the frame will be transmitted Tagged, UnTagged, Unmodified or Double Tagged - Section 3.7.1). Tagged frames are still considered tagged for all other decisions, they just aren't considered tagged for the egress tagging rules. These frames are assigned the ingress port's DefaultVID (Default Port VLAN ID & Priority register, offset 0x07).

3.5.12.8 Security Override of a Frame's VID

The devices support a VID override function where a tagged frame's VID is ignored and the port's DefaultVID is assigned to the frame instead, even if 802.1Q is enabled on the port (see Section 3.5.7). This is a security feature that ensures that all frames that came from a specific ingress port (tagged or untagged) exit the switch with the ingress port's DefaultVID. This prevents an end user from masquerading by simply adding an improper tag to frames.

This feature is enabled on a per port basis by setting the port's ForceDefaultVID bit to a one (in Default Port VLAN ID and Priority register - offset 0x05).

3.5.12.9 VID Assigned to the Frame

Each frame entering the switch must have a VID (VLAN ID) assigned to it. This VID is used for 802.1Q, if enabled on the port, or it is used for cross-chip port based VLANs (Section TBD) if 802.1Q is disabled on the port. It is also used as the frame's VID if untagged frames are to egress the switch tagged.

If a frame entering the switch is untagged, it is assigned the port's DefaultVID during Ingress (see Default Port VLAN ID and Priority register, offset 0x05). If a frame is tagged its VID is generally used as the frame's VID unless the frame's VID is 0x000 or if the port's ForceDefaultVID is set.

A summary of how a VID is assigned to each frame is shown in Table 38:

Table 38: Example VID Assignment Summary

Frame's VID	802.1Q Mode	Force Default VID	Default VID	Assigned VID	Comments
Don't Care	Disabled	Don't Care	0x001	0x001	Use Default VID due to 802.1Q being disabled.
0x000	Enabled	Don't Care	0x001	0x001	Use Default VID due to frame VID = 0x000.
0x123	Enabled	Enabled	0x001	0x001	Use Default VID due to ForceDefaultVID = 1.
0x123	Enabled	Disabled	0x001	0x123	Use frame's VID.

3.5.13 Priority from IPv4 & IPv6 Frames

The format of an IPv4 TOS/DiffServ frame is shown in Figure 19 and an IPv6 Traffic Class frame is shown in Figure 20. The gray portions of the frame in the figures are the only portions of the frame examined for priority determination. If the IP portion of a frame is used as the frame's priority and the frame egresses tagged the upper 2 PRI bits in the tagged frame (PRI[2:1]) come from the egress queue the frame was mapped to and the lowest PRI bit (PRI[0]) comes from the source port's DefPri[0] (Default Port VLAN ID & Priority, offset 0x07).

Figure 19: IPv4 Priority Frame Format

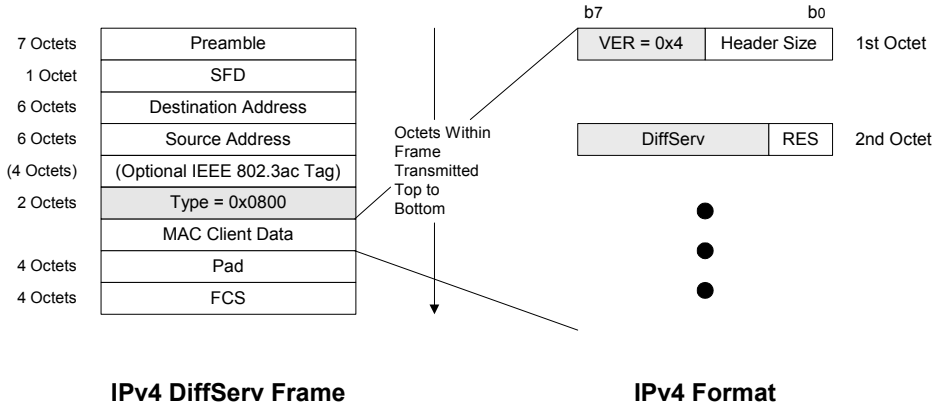

Figure 20: IPv6 Priority Frame Format

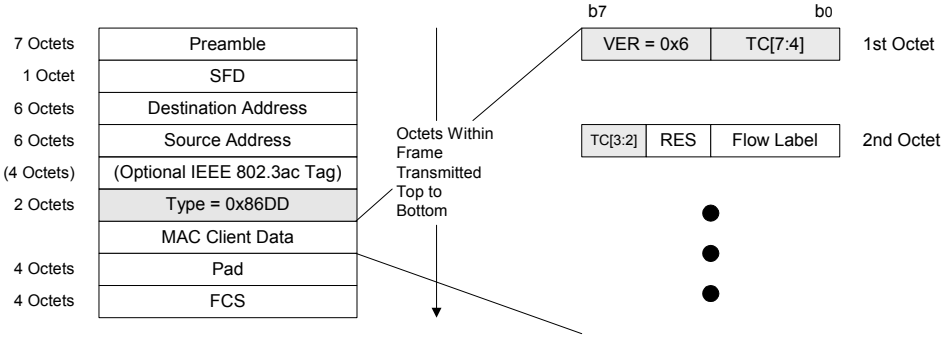

The devices capture the frame's DiffServ bits (if its an IPv4 frame) or the frame's TC[7:2] bits (if it is an IPv6 frame) and uses them as the frame's priority. The DiffServ/TC bits supports 64 priorities while the devices support four. So the Ingress Policy block takes DiffServ/TC bits and maps them into the two priority bits passed to the Queue Controller. This is done using the data found in the global IP-PRI registers (global offsets 0x10 to 0x17). The rest of the frame's IP bits are ignored by the devices unless IGMP/MLD snooping is enabled in the port (see Section 3.5.14). The device's default configuration is to capture and use IP frame priority data in the absence of IEEE Tag priority data on all ports.

3.5.14 IGMP/MLD Snooping

The devices support IPv4 IGMP snooping and IPv6 MLD snooping. It is used to direct certain frames to the CPU for processing. The required formats of these frames are shown in Figure 21 and Figure 22. Management (MGMT) frames bypass IGMP/MLD snooping (see Section 3.4.8). The gray portions of the frame in the figures are the only portions of the frame examined for this function.

Figure 21: IPv4 IGMP Snoop Format

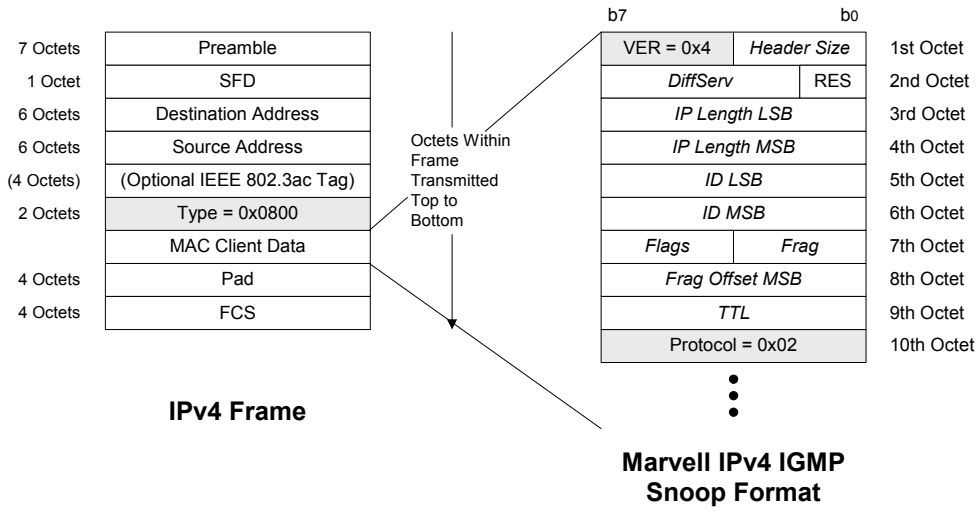

Figure 22: IPv6 MLD Snoop Format

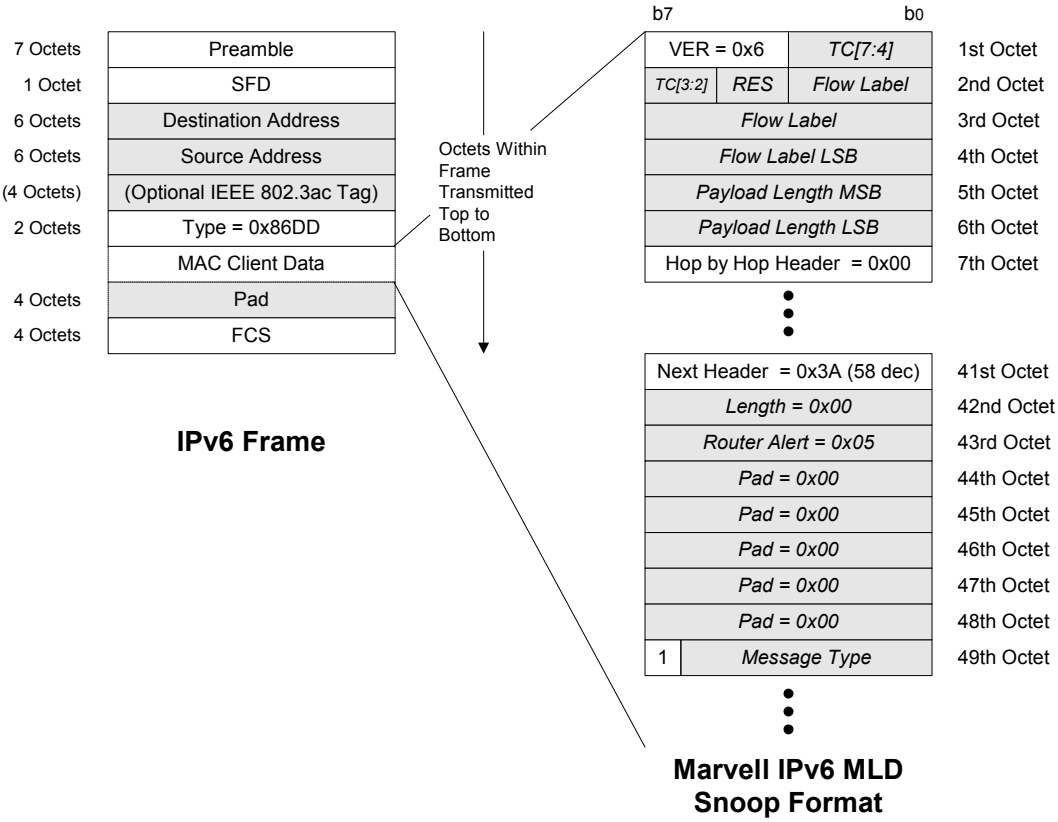

When one of these frames enters a port where IGMP/MLD snooping is enabled, the frame is sent to the CPU's port instead of where the Destination MAC Address directed it. 802.1Q VLAN membership and/or Port Based VLAN rules are still followed so the CPU port (in the local device, if cascaded devices are being used) must be a member of the frame's VLAN to receive the frame. IGMP/MLD snooping is enabled on a per port basis by setting the port's IGMP/MLD Snoop bit to a one (in the port's Port Control register, offset 0x04). The CPU's port is defined by the port's CPUPort field found in the port's Port Control 2 register (offset 0x08).

3.5.15 Ingress Rate Limiting

A switch design may need to limit the reception rate of multicast frames along with unknown or flooded unicast frames. It may need to limit the rate for all frames but still keep QoS. The devices support this facility on a per port basis by setting bits in the port's Rate Control and Rate Control 2 registers (offset 0x09 and 0x0A).

First, the kinds of frames to limit need to be determined and selected by setting the port's LimitMode bits (see Rate Control 2 register, offset 0x0A).

The devices can:

- Limit all frames, or
- Limit just multicast and flooded unicast frames (including broadcast frames), or
- Limit just multicast frames (including broadcast), or
- Limit just broadcast frames

Specific multicast addresses can override this limit setting if they are locked into the address database with an Entry_State value of 0x5 (see Section 3.4.10). Management [1] (MGMT) frames can also be specifically excluded by clearing the port's LimitMGMT bit to a zero in the Rate Control register. Any frame that is not limited by the above rules is ignored in the rate calculations (i.e., its size is not counted toward the limit total).

Second, the maximum rate needs to be selected and programmed. The devices support 4095 different rate limits from 62 Kbps to 256 Mbps (ingress rate limiting can be disabled by setting the Pri0Rate register bits to all zero). The programmed register value uses the equation listed in the Pri0Rate bits description (see Table 70, "Rate Control," on page 189 for details).

Third, the maximum rate for higher priority frames must be selected. devices support four different QoS priorities and each higher priority can be limited to the same rate as the next lower priority or its limit can be twice as much.

Fourth, the bytes to count for limiting needs to be determined. The default setting includes the frame's bytes from the beginning of the Preamble to the end of the FCS (frame check sequence) with an added minimum inter frame gap (IFG). The frame's preamble and SFD bytes (the eight bytes prior to the DA) can be excluded if the CountPre bit is cleared to a zero in the port's Rate Control 2 register. The frame's minimum IFG bytes (the 12 byte equivalent inter frame gap after the FCS) can be excluded if the CountIFG bit is cleared to a zero.

Note

Ingress Rate Limiting in the devices is not recommended for TCP/IP rate limiting. It is intended for storm prevention.

1. BPDU frames are MGMT frames (Section 3.4.8).

3.5.16 Ingress on Distributed Switching Architecture (DSA) Tag Ports

If a CPU needs to perform Spanning Tree or bridge loop detection, or if it must be able to send frames out of any specific port, the CPU's port needs to be configured for DSA Tag mode. Any port can be configured in this way by setting the DSA_Tag bit in the port's Port Control register (offset 0x04). All ports used for internal switch management must be set this way. This includes the port directly connected to a CPU and all the ports used to interconnect multiple switching devices together in a cascade.

When the DSA Tag mode is enabled on a port, the portion of the frame normally reserved for IEEE 802.3ac tags (the four bytes just after the SA) is used to control the switch. The Ingress Policy block examines the DSA Tag in every frame to determine what to do with the frame. This means the CPU must always send frames into the switch with a DSA Tag (if the DSA_Tag mode is enabled).

3.5.16.1 Forward DSA Tag

The Ingress DSA Tag gives the CPU the ability to override the action that the switch was to perform with the frame it received. It may not be convenient for the CPU to perform the switch processing all of the time. If the CPU wants the switch to process the frame based upon the switch's current Ingress policy then the CPU sets the DSA Tag data in the frame to the Forward Tag format as shown in Figure 23 and defined in Table 39.

Figure 23: Ingress Forward Tag Format

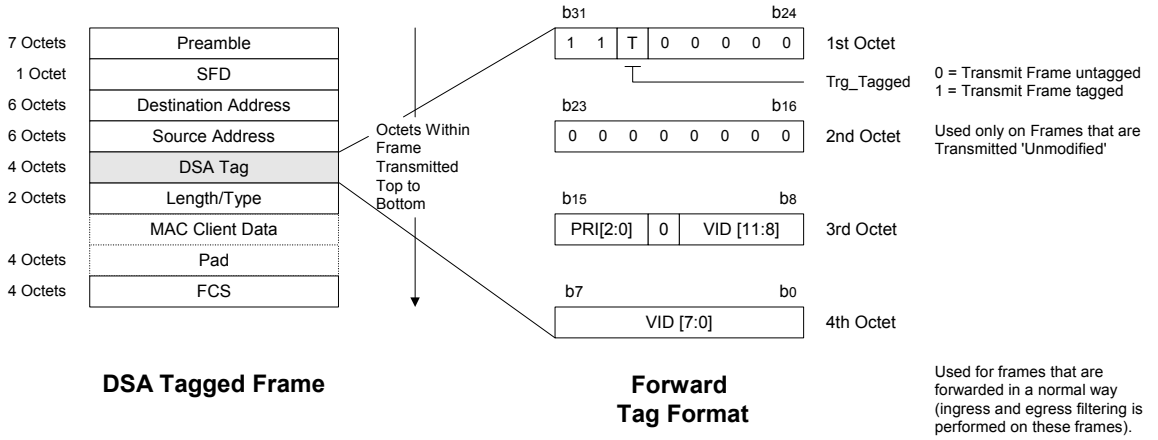

Table 39: Ingress Forward Tag Fields

Port State	Description
Trg_Tagged	Target Tag Mode. Egress ports can be configured to transmit frames 'Unmodified' (i.e., Tagged frames remain Tagged and Untagged frames remain Untagged). Since all DSA Tag frames are Tagged, this bit is used to inform the egress port to consider this frame as Tagged (or not) so it will be transmitted as Tagged (or not). This bit only applies to Egress Ports in the 'Unmodified' mode.
PRI[2:0]	The frame's priority. PRI is used as the frame's IEEE Tag priority (if UseTag is set to a one – see Section 3.5.5). Forward DSA Tag frames are processed as if they are IEEE Tag frames entering the port (except for the Trg_Tagged bit above).
VID[11:0]	The frame's VLAN identifier. VID is used as the frame's IEEE Tag VID. Forward DSA Tag frames are processed as if they are IEEE Tag frames entering the port. Normal VID extraction rules are followed on Forward DSA Tag frames (Section 3.5.12.7).

3.5.16.2 From_CPU DSA Tag

The CPU can override all of the switch's Ingress policy (including priorities, VLANs, PortStates, etc.) by setting the DSA Tag data in the frame to the From_CPU Tag format shown in Figure 24 and defined in Table 40. From_CPU frames are considered MGMT (or BPDU) frames that can tunnel through Blocked Ports (Section 3.5.1). MGMT frames are used primarily for the Spanning Tree Protocol (see Section 3.10). The devices support STP across multiple cascaded devices by giving the CPU direct control of defining the target port and the target device where this frame is to egress. The CPU also controls whether the frame is to Egress the target port Tagged or Untagged (this is necessary since all DSA Tagged frames are Tagged).

Figure 24: Ingress From_CPU Tag Format

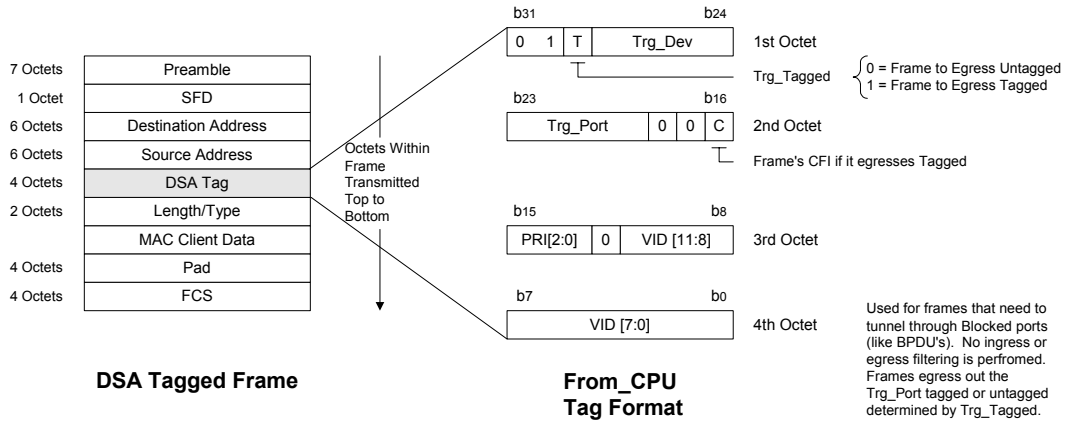

Note

The CFI bit in From_CPU DSA Tag frames is placed at b15. When this frame egresses out a normal network port the CFI bit will be moved to its correct position at b12.

Table 40: Ingress From_CPU Tag Fields

Port State	Description
Trg_Tagged	Target Tag Mode. This bit allows the CPU to define if this frame is to Egress the target port IEEE Tagged or not. If this bit is set to a one, the frame egresses the target port with a proper IEEE 802.3ac Tag. If this bit is cleared to a zero, the frame egresses the target port Untagged (i.e., the DSA Tag will be removed from the frame).
Trg_Dev	Target Device. These bits are used to define the target device's number. Use 0x00 for single chip switches (assuming the chip's DeviceNumber in Global Control 2, global offset 0x1C = 0x00). Use a non-zero value for multi chip cascaded switches. From_CPU frames pass from chip to chip (using the switch port defined by the CascadePort register or the Routing Table, global 2 offset 0x06) until the frame finds a chip where the frame's Trg_Dev field matches the chip's DeviceNumber register (CascadePort and DeviceNumber are in the Global Control 2 register, global offset 0x1C).
Trg_Port	Target Port. These bits are used to define the target port's number (on the target device – see Trg_Dev above). Use 0x00 for Port 0, 0x01 for Port 1, 0x02 for Port 2, etc. From_CPU frames will Egress the Target Port on the Target Device (see above).
C	CFI bit. This bit is use as the frame's IEEE tag CFI bit if the frame egresses the target port Tagged (as defined by the Trg_Tagged bit above).
PRI[2:0]	The frame's priority. PRI[2:1] are used to indicate which egress queue these frames are to be switched to. All three PRI bits are used as the frame's IEEE tag priority if the frame egresses the target port Tagged (as defined by the Trg_Tagged bit above).
VID[11:0]	The frame's VLAN identifier. These VID bits are ignored inside the switch on From_CPU frames. They are only used as the frame's IEEE VID if the frame egresses the target port Tagged (as defined by Trg_Tagged above).

3.5.17 Switch Ingress Header for Routers

The CPU in a router needs to perform many functions. One of those functions is to route IP frames from a WAN to or from LAN ports and another is to bridge frames between one VLAN and another VLAN. The devices Ingress Header mode increases the performance of both of these functions. Any port can be configured to support an Ingress Header by setting the Header bit[1] in the port's Port Control register (offset 0x04) but only the port directly connected to a CPU should be configured in this way.

The Ingress Header accelerates the CPU's performance when routing IP frames by aligning the IP portion of the frame to 32-bit boundaries. This is accomplished by prepending the frame with two extra bytes of data. Bridging between VLAN ports sometimes requires the switch to support multiple address databases (one for each VLAN) so that the same MAC address can be used on multiple VLANs. Since the CPU is generally a member of all VLANs, it must inform the switch which VLAN to use on a given frame (and thus which address database to use). This is accomplished by using an Ingress Header with a non-zero value as defined in Figure 25[2]. When the Ingress Header is seen with a non-zero value its contents are written to the port's Port Based VLAN Map register (offset 0x06) prior to the start of the rest of the frame. The frame is then processed by the switch using this new information. In this way, the CPU can direct which port based VLAN and address database to use on every frame at wire speed.

When the Ingress Header mode is enabled on a port, the first two bytes of the frame (just before the DA) are used to control the switch. The Ingress Policy block removes the Header from the frame, causing the frame to be two bytes smaller in size, and overwrites the frame's FCS with a new FCS. This adjustment makes the frame normal for the rest of the network since the Header's data is intended for the switch only. Frame size checking is performed on the adjusted framed size. This means the CPU must always add two bytes of data to the beginning of every frame it sends into the switch (if the Header mode is enabled).

The Ingress Header gives the CPU the ability to control which VLAN (port based), learning mode and address database to use on the frame that it just received. It may not always be convenient to use the CPU to manipulate the appropriate switch registers. If the CPU directs the switch to process the frame based upon the switch's current Ingress policy then the CPU sets the Header data in the frame to all zeros (i.e., it pre-pends the frame with two extra bytes of zeros). This zero padding indicates that the switch should ignore the Header's data and process the frame normally (after the Header's data is removed from the frame).

1. The Header bit enables the Marvell Header mode for both ingress and egress.
2. Reserved bits in the Marvell Header must always be zeros.

Figure 25: Ingress Header Format

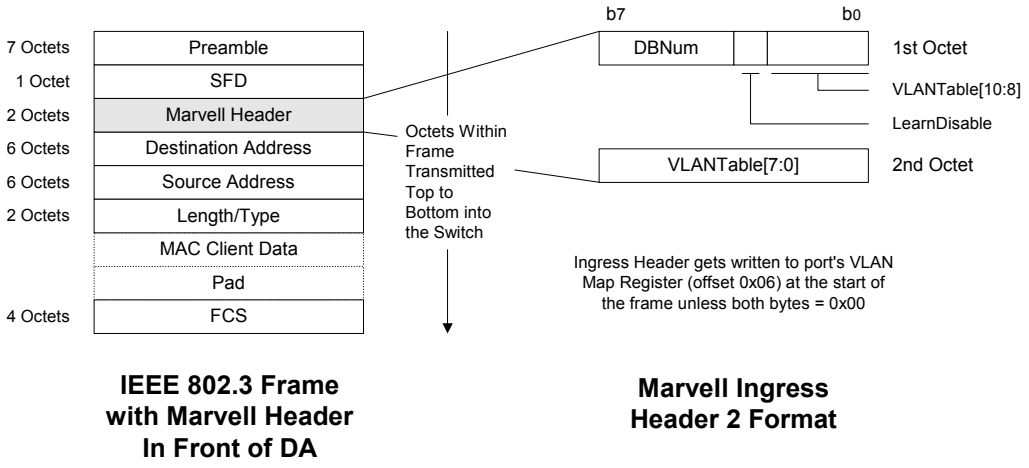

When an Ingress Header contains a non-zero value its contents are written directly to the port's Port Based VLAN Map Register. See Table 67 for a description of each of the Header's fields.

Note

The Marvell Header can only modify the lower 4-bits of the port's DBNum. The upper 4-bits come from the port's DBNum[7:4] register bits from Port Control 1 register, offset 0x05.

The Marvell Header can be used with or without the DSA Tag (see Section 3.5.16). In this mode, the Marvell Header's DBNum is used as the port's default DBNum which may be overwritten by the DSA Tag's VID lookup into the VTU (If the Tag's VID is contained in the VTU). The Marvell Header can be used in single chip or cascaded chip environments although some of its usefulness is reduced in a cascaded environment owing to the limited VLANTable size in the Header.

The DSA Tag can be used to force where a frame goes (by using the From_CPU format). The Marvell Header is used to limit where a frame goes (by using the VLANTable).

3.5.18 Ingress Double VLAN Tagging (88E6095 and 88E6095F Only)

Double Tagging is a way to isolate one IEEE 802.1Q VLAN from other IEEE 802.1Q VLANs in a hierarchical fashion that is compatible with IEEE 802.1Q aware switches as long as those switches support a maximum frame size of 1526 bytes or more. This method places an extra or Double Tag in front of a frame's normal tag (assuming the frame were already Tagged), increasing the frame size by 4 bytes. The Double Tag frame format is shown in Figure 26.

Ingress Double Tagging is selectable on a port by port basis by setting the port's DoubleTag bit in the appropriate Port Control register (offset 0x04). Typically, any port that has Ingress Double Tagging enabled also has Egress Double Tagging enabled.

An Ingress port that has Double Tagging enabled expects all Ingress frames to contain an extra tag that needs to be removed from the frame prior to performing the port's Ingress policy on the frame. In this mode, the Ingress port removes the first IEEE 802.3ac tag that appears after the Source Address in every frame. If the frame is untagged, it is not modified. If the frame has a single tag, it is removed. If the frame has two tags, the first tag is removed. The CRC on modified frames is updated.

The port's Ingress policies of frame size is performed after the first tag (if present) is removed. This requires that tagged frames must be at least 68 bytes in size so as not to be considered undersized frames after the tag is removed (frame padding is not performed on Ingress Double Tagging).

Figure 26: Double Tag Format

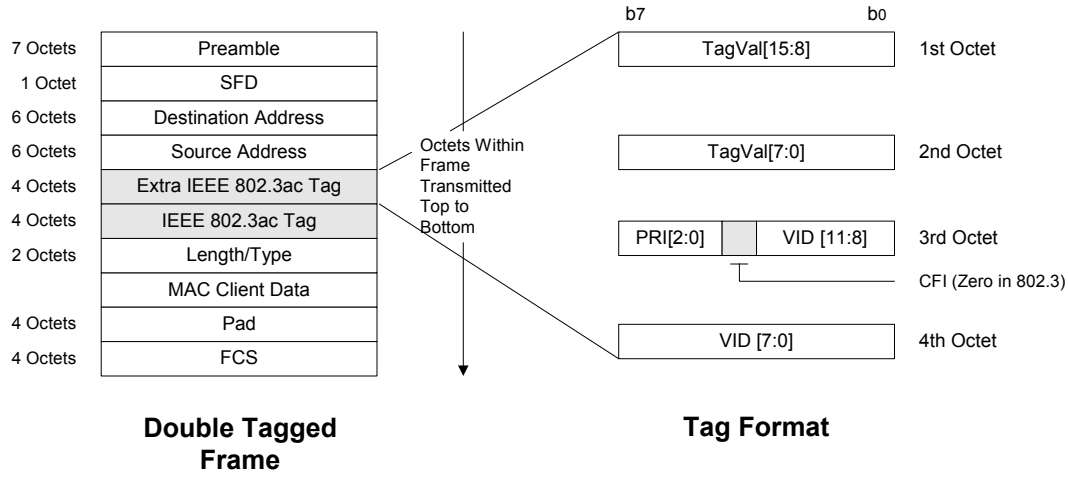

3.5.18.1 Determining if a Frame is Double Tagged

If the port's DoubleTag bit is a one (Port Control, offset 0x04) a frame is considered double tagged in the devices if the two bytes after the frame's SA match TagVal. TagVal equals 0x8100 if the port's UseCoreTag bit is a zero (in Port Control, offset 0x04). If the port's UseCoreTag bit is a one then TagVal equals CoreTagType (in the Core Tag Type register, global offset 0x19).

TagVal is becomes 0x8100 for the remainder of frame processing once the extra tag is removed from the frame regardless of the TagVal used to determine Double Tagged frames. This means the extra tag can use any EtherType value, but the inner tag must always be 0x8100 for the resulting frame to be considered tagged.

3.5.18.2 Using or Ignoring the Removed Double Tag's Data

Once the Double or extra tag is removed from the frame, the devices support ignoring the data that was removed (i.e., the extra tag's PRI and VID bits) or using it for ingress switching.

The data from the removed tag is 'ignored' if the UseDoubleTagData bit is zero (Switch Management register, global 2 offset 0x05). Use this mode if the switching of the frame is to be done using what is left in the frame after the first or extra tag is removed. This places the boundary between Customer ports and Provider ports at the ingress of the Double Tag port and requires that all the other ports of the switch belong to the same Customer.

The data from the removed tag is 'used' if the UseDoubleTagData bit is one (Switch Management register, global 2 offset 0x05). Use this mode if the switching of the frame is to be done using the contents of the first or extra tag before it is removed from the frame. This places the boundary between Customer ports and Provider ports at the egress of all the other ports in the switch and supports these other ports belonging to many different Customers. This mode supports more than one port per Customer as long as all the Customer ports belong to all the Customer's VLANs since switching is not done with the 2nd or standard IEEE tag in this mode.

3.6 Queue Controller

The devices queue controller uses an advanced non-blocking, four priority, output port queue architecture with Resource Reservation. As a result, the devices support definable frame latencies with guaranteed frame delivery (for high priority frames) without head-of-line blocking problems or non-blocked flow disturbances in any congested environment (for all frame priorities).

3.6.1 Frame Latencies

The devices can guarantee frame latencies owing to its unique, high performance, four priority queuing system. A higher priority frame is always the next frame to egress a port when a port is currently egressing frames of a lower priority. This is true regardless of the two priorities and the Scheduling[1] mode of the switch.

3.6.2 No Head-of-Line Blocking

An output port that is slow or congested never effects the transmission of frames to uncongested ports. The devices are designed to ensure that all uncongested flows traverse the switch without degradation regardless of the congestion elsewhere in the switch.

3.6.3 QoS with and without Flow Control

The Queue Controller is optimized for three modes of operation:

- Flow Control disabled on all ports
- Flow Control enabled on all ports
- Flow Control enabled on some ports and disabled on the rest

When flow control is enabled no frames are dropped and higher priority flows receive higher bandwidth through the switch (i.e., these flows are less constrained if there is congestion between a higher and a lower priority flow). The percentage of bandwidth that each flow receives is determined by the Scheduling mode see (Section 3.6.5). Flow control prevents frames from being dropped but it can greatly impact the available bandwidth on any network segment that is utilizing flow control. The latency of higher priority data on flowed-off network segments also increases since industry constrained standard flow control mechanisms stop all frames from being transmitted. Therefore, flow control may not be desirable in a QoS switch environment. For this reason, the devices are also optimized to work properly without flow control.

When flow control is disabled and congestion occurs for an extended period of time, frames will be dropped. This is true in all switches. The devices drop the correct frames, i.e., the lower priority frames. In this case, the higher priority flows get a higher percentage of the free buffers. The percentage of buffers they get is determined by the Scheduling mode (see Section 3.6.5).

For the mixed flow control case, frames are dropped if congestion occurs on the non-flow controlled ports while the flow controlled ports do not drop any frames. The percentage of bandwidth each port receives is priority based and fairness is maintained between all the ports in the switch (determined by the Scheduling mode).

1. The Scheduling mode selects either a Fixed priority or an 8-4-2-1 Weighted Fair Queuing – Section 3.6.3

3.6.4 Guaranteed Frame Delivery without Flow Control

The devices can guarantee high priority frame delivery[1], even with Flow Control disabled, due to its intelligent Resource Reservation system. Having an output queue with multiple priorities is not sufficient to support Quality of Service (QoS) if the higher priority frames cannot enter the switch due to a lack of buffers. The devices reserve buffers for higher priority frames so they can be received and then switched. These high priority buffers are the first to get replenished from the Free Queue which prepares the receiving port for the next high priority frame.

3.6.5 Fixed or Weighted Priority

The devices support either a fixed priority or weighted fair queuing schemes. The selection is made by the Scheduling bit in the Switch Global Control register (global offset 0x04). In the fixed priority scheme all top priority frames egress for a port until that priority's queue is empty, then the next lower priority queue's frames egress, etc. This approach can cause the lower priorities to be starved out preventing them from transmitting any frames but also ensures that all high priority frames Egress the switch as soon as possible. In the weighted fair scheme an 8, 4, 2, 1 weighting is applied to the four priorities. This approach prevents the lower priority frames from being starved out with only a slight delay to the higher priority frames.

3.6.6 The Queues

The queues in the devices are shown in Figure 27.

Figure 27: Switch Queues

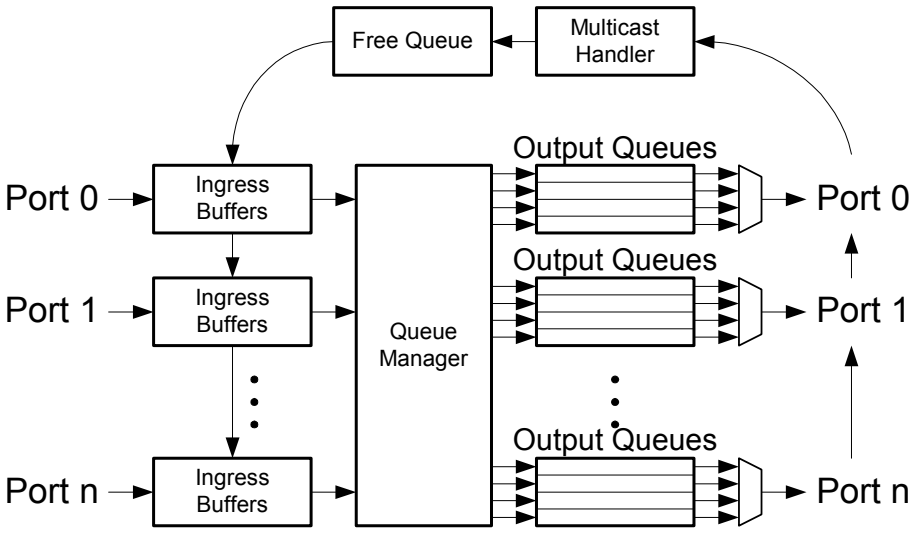

1. If all the frames entering a port are all high priority at wire speed their delivery cannot be guaranteed.

3.6.7 Queue Manager

At reset[1], the Queue Manager initializes the Free Queue by placing all the buffer pointers into it and ensures that all the other queues are empty. Then it takes the first available free buffer pointers from the Free Queue and assigns them to any Ingress port that is not Disabled[2] and whose link[3] is up. The switch is now ready to accept and switch packets. Whenever any port's Link goes down or if the port is set to the Disabled Port State the port's ingress Buffers and Output Queue buffers are immediately returned to the Free Queue. This prevents stale or lost buffers and allows the Free Queue to be large so larger bursts of momentary congestion can be handled. When a non-Disabled port's Link comes back up it gets its Ingress Buffers back so it can start receiving frames again.

When a MAC receives a packet it places it into the embedded memory at the address pointed to by the Input Pointers it received from the Queue Manager. When the packet is received, the MAC transfers the pointer(s) to the Queue Manager and requests new buffers from the Free Queue. If the Free Queue is empty the MAC does not receive any pointers until they become available. If the MAC starts to receive a packet when it has no pointers, the packet will be dropped. Flow control will be asserted before this occurs if it's enabled.

The Queue Manager uses the data returned from the Lookup Engine (see Section 3.4.1) and the Ingress Policy (see Section 3.5) to determine which Output Queue or Queues the packet's pointer should go to and at what priority. At this point, the Queue Manager modifies the desired mapping of the frame depending upon the mode of the switch and its level of congestion. Two modes are supported, with and without Flow Control. Both modes are handled at the same time and can be different per port (i.e., one port has Flow Control enabled and another has it disabled).

If Flow Control is enabled on an Ingress port the frame is switched to the desired Output Queues without modification. This is done so frames are not dropped. Instead, the Queue Manager carefully monitors which Output Queues are congested and enables or disabled Flow Control on the Ingress ports that are causing the congestion. This approach allows uncongested flows to progress through the switch without degradation.

If Flow Control is disabled on an Ingress port the frame may be discarded instead of being switched to the desired Output Queue(s). If a frame is destined to more than one Output Queue it may get switched to some and not others. The decisions are complex as the Queue Manager takes many pieces of information into account before the decision is made. The Queue Manager looks at the priority of the current frame, the current level of congestion in the Output Queue(s) the frame is being switched to and the current number of free buffers in the Free Queue. The result is uncongested flows transverse the switch unimpeded and higher priority frames get in and through the switch faster even if there is congestion elsewhere in the switch.

3.6.8 Output Queues

The Output Queues receive and transmit packets in the order received for any given priority. This is very important for some forms of Ethernet traffic. The Output Queues will be emptied as fast as they can – but they could empty at different rates. This could be due to a port being configured for a slower speed or it could be caused by network congestion (collisions or flow control).

Each port contains four independent Output Queues, one for each priority. The order the frames are transmitted out each port is controlled by the Scheduling bit in the Switch Global Control registers (global offset 0x04). A fixed or a weighted priority can be selected (see Section 3.6.5).

After a packet has been transmitted fully out to the MAC, the Output Queue passes the transmitted packet's pointer(s) to the Multicast Handler for processing. The MAC then begins transmitting the next packet.

1. The Queue Manager is reset by either the hardware RESETn pin or a software reset by the SWReset bit in the Switch Global Control register (global offset 0x04).
2. If a port is in the Disabled Port State (see Section 3.5.1) its Ingress buffers are left in the Free Queue for other ports to use.
3. SERDES-based ports need to get a Link up from the PHY, (G)MII based ports have Link up if they are enabled.

3.6.9 Multicast Handler

The Multicast Handler receives the pointers from all the packets that were transmitted. It looks up each pointer to see if this packet were directed to more than one Output Queue. If not, the pointers are returned to the Free Queue where they can be used again. If the frame were switched to multiple Output Queues the Multicast Handler ensures that the frame has egressed all of the ports to which it was switched before returning the pointer(s) to the Free Queue.

3.7 Egress Policy

If directed, the Egress Policy block modifies frames, as they exit the switch. IEEE or DSA Tags can be added or removed or converted from one form to the other. Specific switch information can be added to the frame for the switch's CPU as well.

3.7.1 Tagging and Untagging Frames

Egress tagging and untagging is supported dynamically using 802.1Q VLANs, or statically using Port Based VLANs. The mode that is used on a port is determined by the egress port's 802.1Q Mode bits (Port Control 2 register, offset 0x04) as follows:

- Secure or Check – The MemberTag bits (in VTU Data register, global offsets 0x07 to 0x09) associated with the VID assigned to the frame during ingress determines if the frame should egress unmodified, tagged or untagged.
- Fallback – The MemberTag bits associated with the VID assigned to the frame during ingress determine if the frame should egress unmodified, tagged or untagged if the VID is contained in the VLAN database (Section 3.5.11). If the VID is not found in the VLAN database, the frame egresses unmodified, tagged or untagged determined by the port's EgressMode bits (in Port Control register, offset 0x04).
- 802.1Q Disabled – The port's EgressMode bits determine if the frame will egresses unmodified, tagged or untagged.

Note

If the port's EgressMode bits are set to 'Always add a Tag' (or Egress Double Tag) then a tag is always added to the frame on egress regardless of the port's 802.1QMode.

The devices perform the following actions on the egressing frame depending upon the Egress tagging mode that was determined above:

- Transmit Unmodified[1] - UnTagged frames egress the port UnTagged. Tagged frames egress the port Tagged.
- Transmit UnTagged[2] - UnTagged frames egress the port unmodified. The IEEE Tag on Tagged frames is removed, the frame is zero padded if needed[3], and a new CRC is computed for the frame.
- Transmit Tagged[4] - Tagged frames egress the port unmodified. An IEEE Tag is added to UnTagged frames and a new CRC is computed. The contents of the added tag is covered in Section 3.7.1.1.
- Transmit all frames Double Tagged (see Section 3.8).

The format of an IEEE Tagged frame is shown in Figure 28.

1. This is the default setting so the switch acts as a transparent switch.
2. Needed when switching frames to end stations that don't understand Tags.
3. Tagged frames that are less than 68 bytes are padded with zero data to insure the UnTagged frame is at least 64 bytes in size.
4. Typically used when switching frames into the core or up to a server.

Figure 28: IEEE Tag Frame Format

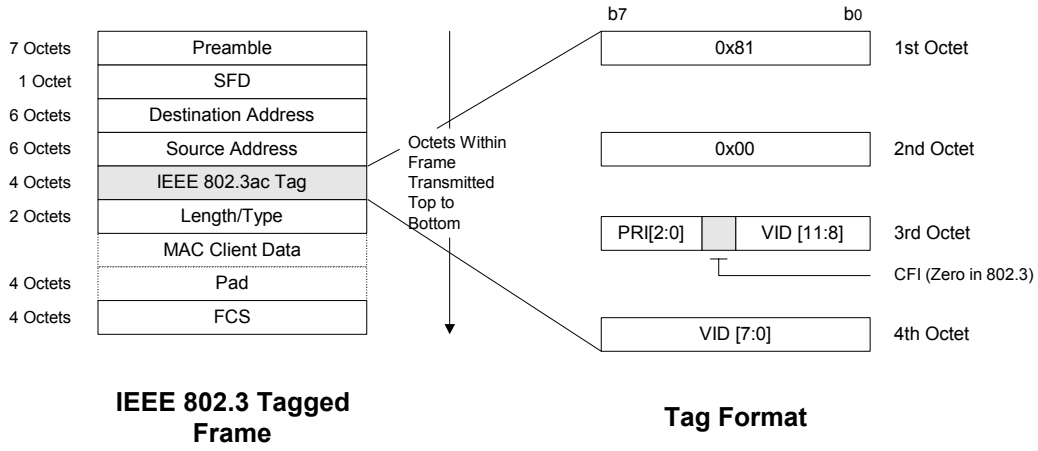

3.7.1.1 Adding a Tag to Untagged Frames

When a Tag is added to an UnTagged frame the Tag is inserted right after the frame's Source Address. The four bytes of added data are:

- The 1st Octet is always 0x81.
- The 2nd Octet is always 0x00.
- The PRI bits indicate the frame's priority determined during Ingress (see Section 3.5.5).
- The CFI bit is always set to a zero.
- The VID bits indicate the VID assigned to the frame during Ingress (see Section 3.5.12.9).

Note

A Tag that is added due to Egress Double Tagging is done differently. See Section 3.7.2.1

3.7.1.2 Priority Re-Map and Priority Override (88E6095 and 88E6095F devices Only)

When a Tagged frame egresses a port Tagged, the PRI bits in the tag are modified to reflect the frame's priority that was determined during Ingress (see Section 3.5.5). The PRI bits can be re-mapped by the ingress port's IEEE Priority Remapping registers (offsets 0x18 and 0x19) or the frame's PRI bits can be ignored and the ingress port's default priority used instead, or the frame's priority can be overridden.

3.7.1.3 VID 0x000 and VID Override

A Tagged frame egressing a port Tagged may have its VID bits modified. If the Ingressing frame's VID were 0x000, the Ingress port's DefaultVID (see Default Port VLAN ID and Priority register, offset 0x07) is assigned to the frame instead. The Ingress port's DefaultVID can be forced on all frames entering a port as the VID to use. See Section 3.5.12.8.

3.7.2 Egress Double VLAN Tagging (88E6095 and 88E6095F devices Only)

Double Tagging is a way to isolate one IEEE 802.1Q VLAN from other IEEE 802.1Q VLANs in a hierarchical fashion that is compatible with IEEE 802.1Q aware switches as long as those switches support a maximum frame size of 1526 bytes or more. This method places an extra or Double Tag in front of a frame's normal Tag (assuming the frame were already Tagged) increasing the frame size by four bytes. The Double Tag frame format is shown in Figure 29.

Egress Double Tagging is selectable on a port by port basis by setting the port's EgressMode bits in its Port Control register (offset 0x04). Typically any port that has Egress Double Tagging enabled also has Ingress Double Tagging enabled (see Section 3.5.18) but this is not always the case.

An Egress port that has Double Tagging enabled transmits all egress frames with an extra Tag. If a frame is UnTagged, it egresses Tagged. If a frame is Tagged, it egresses Double Tagged. The extra or Double Tag is inserted right after the frame's Source Address so this new Tag becomes the frame's first Tag.

Figure 29: Double Tag Format

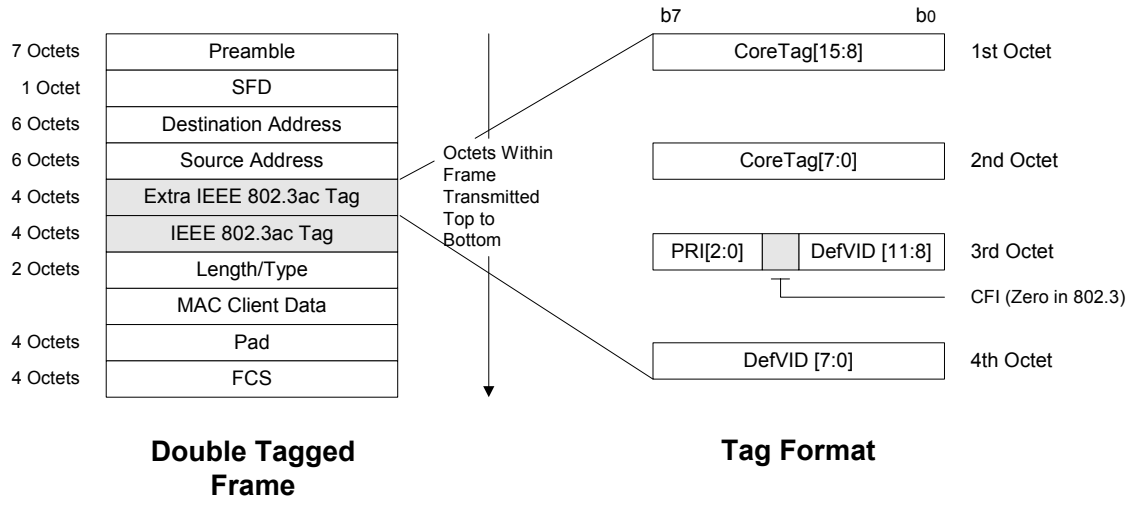

3.7.2.1 Always Adding a Tag or Double Tagging Frames

When a Tag is added because the port is in Egress Double Tagging mode, the Tag is inserted right after the frame's Source Address. The four bytes of added data are:

- The 1st Octet is always CoreTag[15:8] from the Core Tag Type register (global offset 0x19).
- The 2nd Octet is always CoreTag[7:0] from the Core Tag Type register.
- The PRI bits indicate the frame's priority determined during Ingress (see Section 3.5.5).
- The CFI bit is always set to a zero.
- The VID bits are always the source port's DefaultVID from the frame's source port's Default Port VLAN ID and Priority register (offset 0x07).

Note

A Tag that is added due to normal IEEE Tagging is done differently. See Section 3.7.1.1

3.7.3 Port States

The egress of frames from a port is controlled by the port's PortState bits (see Port Control register, offset 0x04) which supports 802.1D Port States defined in Section 3.5.1, and if 802.1Q is enabled on the port, the VTU (Section 3.5.11) which supports 802.1s Port States defined in Section 3.5.10.

3.7.4 Egress Rate Shaping (88E6095 and 88E6095F devices only)

A switch design may need to limit the transmission rate of selected egressing frames but still keep QoS. The devices support this on a per-port basis by setting bits in the port's Rate Control registers (offset 0x09 and 0x0A). Egress rate limiting is performed by shaping the output load.

First, the types of frames to limit, or shape, must be determined and then selected by setting the port's LimitMGMT bit (see the Rate Control register, offset 0x09).

The devices can:

- Limit all frames, or
- Limit all frames except for MGMT frames (see Section 3.4.8)

Management[1] (MGMT) frames are excluded by clearing the port's LimitMGMT bit to a zero in the Rate Control register. Any frame that is not limited is ignored in the rate calculations (i.e., their size is not counted toward the limit total).

Second, the required maximum rate must be selected and programmed. The devices support 4095 different rate speeds or shapes from 62 Kbps to 256 Mbps (Egress rate shaping can be disabled by setting the EgressRate bits to all zero). The programmed register value uses the equation listed in the Egress Rate bits description (see Table 71, "Rate Control 2," on page 190 for details).

Third, the bytes to count for shaping needs to be determined. The default setting includes the frame's bytes from the beginning of the Preamble to the end of the frame check sequence (FCS) with an added minimum inter frame gap (IFG). The frame's preamble bytes (the eight bytes prior to the DA) can be excluded if the CountPre bit is cleared to a zero in the port's Rate Control 2 register. The frame's minimum IFG bytes (the 12-byte inter frame gap after the FCS) can be excluded if the CountIFG bit is cleared to a zero.

1. BPDU frames are considered MGMT frames in this device (see Section 3.4.8).

3.7.5 Egress on DSA Tag Ports

If a CPU needs to perform Spanning Tree or bridge loop detection, it needs to be able to know the physical source port of the BPDU[1] frames it receives. The CPU receives this data whenever the CPU's port is configured for DSA Tag[2] mode. Any port can be configured in this way by setting the DSA_Tag bit in the port's Port Control register (offset 0x04). All ports used for switch management frames must be set in this way, including the port directly connected to a CPU and all of the ports used to interconnect multiple switching devices together in cascade.

When the DSA Tag mode is enabled on a port, the portion of the frame normally reserved for IEEE 802.3ac tags is used to give the CPU additional information about the frame (it's the four bytes just after the SA). The Egress Policy block will insure every egressing frame is a frame with a DSA Tag. This means the CPU will receive all of its frames from the switch with a DSA Tag (if the DSA_Tag mode is enabled).

Multiple DSA Tag types are supported and each is covered below.

3.7.5.1 Forward DSA Tag

The Forward DSA Tag is applied to a frame as it egresses a DSA Tag port if the frame is not a special frame (e.g., it is not a MGMT frame—see Section 3.4.8). Normal ingress and egress filtering rules apply to this frame[3]. Many of the frames the CPU receives will contain this Forward DSA Tag format, shown in Figure 30 and defined in Table 41.

1. BPDU = Bridge Protocol Data Unit
2. The DSA_Tag bit enables the DSA Tag mode for both Ingress and Egress.
3. Not only can Forward Tag frames be discarded but their PRI and VID bits can be modified/overridden (see Section 3.7.1.2 and Section 3.7.1.3) if they propagate through another devices before these frames reach a CPU.

Figure 30: Egress Forward Tag Format

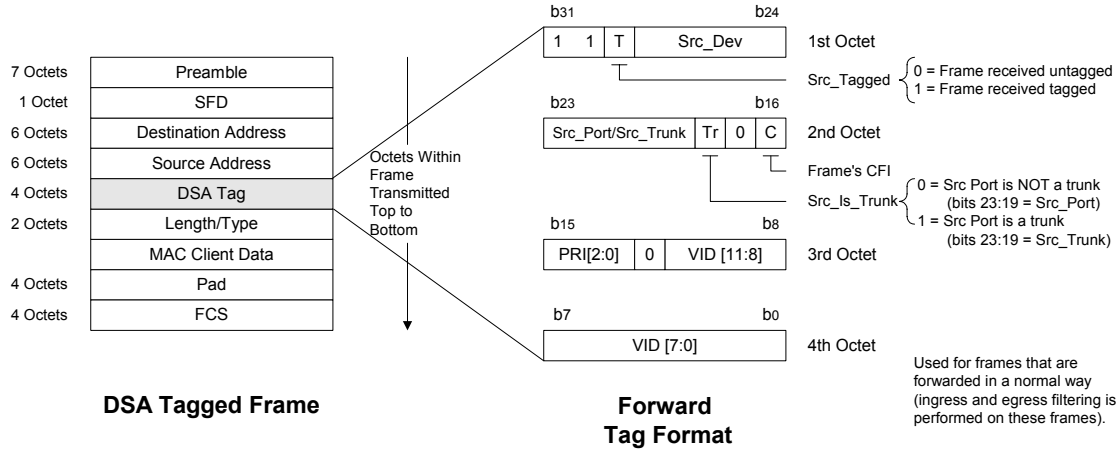

Table 41: Egress Forward Tag Fields

Port State	Description
Src_Tagged	Source Tag Mode. This bit informs the CPU whether the original frame ingressed the original source port with an IEEE Tag. If this bit is set to a one the frame contained a proper IEEE 802.3ac Tag and that Tag was converted into a DSA Tag. If this bit is cleared to a zero the frame was originally Untagged (i.e., the DSA Tag was added to the frame).
Src_Dev	Source Device. These bits used to define the original source device's number where the frame first ingressed (i.e., the first device where the frame Ingressed from a normal or Network port before being switched to an Internal port using the DSA Tag). These bits come from the device's DeviceNumber (Global Control 2 register, global offset 0x1C).
Src_Port/Src_Trunk	Source Port or Source Trunk. If the Tr bit, below, is zero, these bits are used to define the original source port's number (on the source device—see above). 0x00 indicates Port 0, 0x01 indicates Port 1, 0x02 for Port 2, 0x03 for Port 3 etc. If the Tr bit, below, is one, these bits are used to define the Trunk ID of the 1st trunk this frame entered or passed through.
Tr	Source is Trunk. When this bit is zero it indicates this frame originally entered a non-trunked port and this frame has never passed through a trunked port. In this case, the Src_Port/Src_Trunk bits are Src_Port. When this bit is one it indicates this frame originally entered a trunked port or this frame passed through a trunked port. In this case, the Src_Port/Src_Trunk bits are Src_Trunk.
C	The original frame's CFI bit if the frame was IEEE 802.3ac tagged when it originally entered the switch.
PRI[2:0]	The frame's priority as determined by the Ingress rules of the last devices that this frame entered (see Section 3.7.1.1 and Section 3.7.1.2).
VID[11:0]	The frame's VLAN identifier as determined by the Ingress rules of the last devices that this frame entered (see Section 3.7.1.1 and Section 3.7.1.3).

3.7.5.2 To_CPU DSA Tag

If the CPU needs to run the Spanning Tree Protocol (see Section 3.10) or to perform IGMP/MLD Snooping (see Section 3.5.14) or ARP Mirroring, the CPU must receive the special frames required to run these protocols. When devices are configured to detect these special frames and it is also configured to egress these frames from a DSA Tag port, the frame is modified with a To_CPU DSA Tag as it egresses the port. The format of the To_CPU Tag is shown in Figure 31 and defined in Table 42.

In an environment where more that one devices are connected together in a cascade, the devices in the middle will receive To_CPU DSA Tag frames. These frames are sent through the switch unmodified. They egress from the port defined by the ingressing port's CPU Port in the Port Control 2 register (offset 0x08).

Figure 31: Egress To_CPU Tag Format

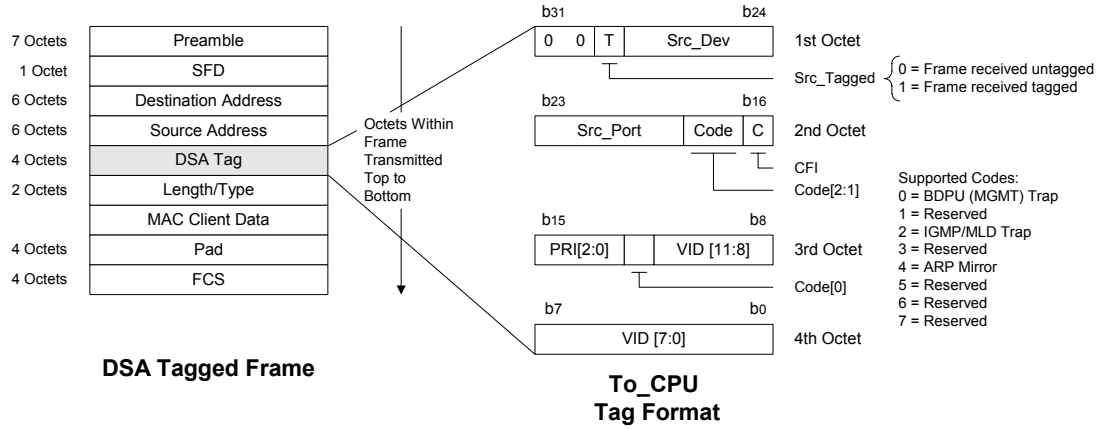

Table 42: Egress To_CPU Tag Fields

Port State	Description
Src_Tagged	Source Tag Mode. This bit informs the CPU whether the original frame entered the original source port with an IEEE Tag. If this bit is set to a one the frame contained a proper IEEE 802.3ac Tag and that Tag was converted into a DSA Tag. If this bit is cleared to a zero the frame was originally Untagged (i.e., the DSA Tag was added to the frame).
Src_Dev	Source Device. These bits are used to define the original source device's number where the frame first ingressed (i.e., the first device where the frame Ingressed from a normal or Network port before being switched to an internal port using the DSA Tag). These bits come from the device's DeviceNumber (Global Control 2 register, global offset 0x1C).
Src_Port	Source Port. These bits are used to define the original source port's number (on the source device – see Src_Dev). 0x00 indicates Port 0, 0x01 indicates Port 1, 0x02 for Port 2, 0x03 indicates Port 3 etc.
Code	To_CPU frame type code. These bits are set by the original Src_Dev (see above) to indicate the kind of To_CPU frame. The devices generate a frame type code of 0x0 for MGMT frames (see Section 3.4.8) and 0x2 for IGMP/MLD snooped frames (see Section 3.5.14) and 0x4 for ARP Mirrors.

Table 42: Egress To_CPU Tag Fields (Continued)

Port State	Description
C	The original frame's CFI bit if the frame was IEEE 802.3ac tagged when it originally entered the switch.
PRI[2:0]	The frame's priority as determined by the Ingress rules of the 1st or original device this frame entered (Section 3.7.1.1 and Section 3.7.1.2).
VID[11:0]	The frame's VLAN identifier as determined by the Ingress rules of the 1st or original device this frame entered (Section 3.7.1.1 and Section 3.7.1.2).

3.7.5.3 From_CPU DSA Tag

In an environment where devices are connected together in cascade, the devices in the middle will transmit From_CPU DSA Tag frames. These frames are sent through the switch unmodified if they are egressing out another DSA_Tag port.

From_CPU DSA Tag frames are defined in Section 3.5.16.2

3.7.6 Switch Egress Header for Routers

If a CPU wants to have the IP frame data of the Ethernet frame aligned to 32-bit boundaries for faster routing, the devices support a Marvell Header Mode that inserts two bytes into the frame just before the frame's Destination Address. Any port can be configured this way by setting the Header bit in the port's Port Control register (offset 0x04) but only the CPU's port should be configured this way[1].

When the Egress Header mode is enabled on a port, two extra bytes are added to the beginning of the frame just before the frame's DA and a new CRC is calculated for the frame. When the frame is received by the CPU the Header will be the first two bytes of the frame in memory. If the CPU's MAC needs to process the frame for filtering or for other reasons, the MAC must be aware that the frame data has been shifted down by two bytes.

The format of the Egress Header is shown in Figure 32 and its fields are defined in Table 43.

Figure 32: Egress Header Format

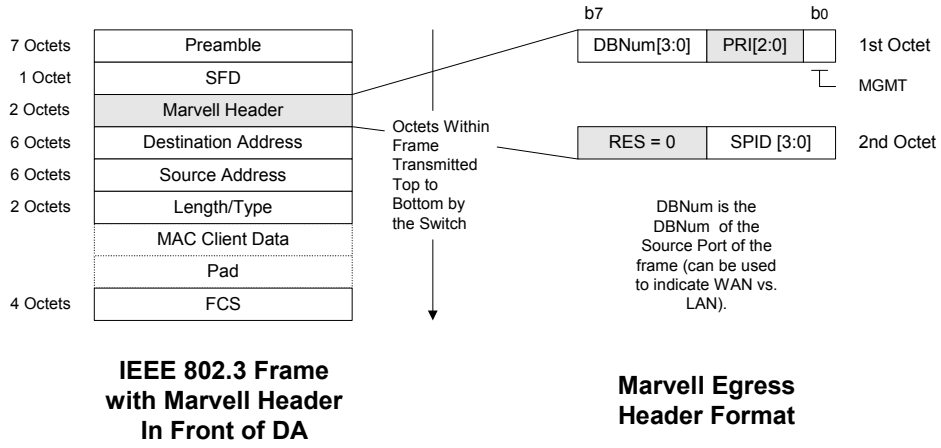

Table 43: Egress Header Fields

Port State	Description
DBNum[3:0]	Database Number. This field represents the lower 4 bits of the address database number assigned to this frame when it Ingressed into this device (i.e., it is assigned by the last device this frame entered). It can be used to indicate the logical port based VLAN number of the source port.
PRI[2:0]	The frame's priority as determined by the Ingress rules of the last device this frame entered (see Section 3.7.1.1 and Section 3.7.1.2). If the frame is a MGMT frame (see Section 3.4.8) the ingress rules effectively make this the frame's priority as determined by the ingress rules of the 1st or original device this frame entered.
RES	Reserved for future use. Currently set to 0x0.
SPID[3:0]	The Source Port ID. These bits indicate at which physical port the frame entered this device (i.e., it is assigned by the last physical device this frame entered). An SPID of all zeros indicates Port 0. An SPID of 0x1 indicates Port 1. 0x2 indicates Port 2, 0x03 indicates Port 3, etc.

1. The Header bit enables the Header mode for both Ingress and Egress

3.8 Port Monitoring Support

Port monitoring is supported by the devices with Egress only monitoring, Ingress only monitoring or Egress and Ingress monitoring. Egress monitoring sends any data that egresses out a particular port to a specific monitor port as well. Ingress monitoring sends any good data that ingresses in a particular port out to a specific monitor port as well as sending the frame where it normally would have gone.

This form of port monitoring is enabled by defining which port or ports are to be the Ingress Monitor Source (IMS) and/or the Egress Monitor Source (EMS). A port becomes an IMS and/or an EMS by setting the appropriate bit(s) to a one in the port's Port Control 2 register (offset 0x08). Both of these bits can be set at same time on the same port and multiple ports can have their bits set.

While many ports can be defined to be the monitor source only one destination port for the IMS frames and one destination port for the EMS frames per devices can be defined. Frames that are received on Ingress Monitor Source ports (IMS) are copied to the port defined as the Ingress Monitor Destination (IMD). Frames that are transmitted out Egress Monitor Source ports (EMS) are copied to the port defined as the Egress Monitor Destination (EMD). The IMD and EMD are defined in the Monitor Control register (global offset 0x1A). The IMD and EMD can point to the same physical port.

Each device that has a at least one monitor source port enabled must also have a monitor destination port defined (of the same type - i.e., an IMD for an IMS and an EMD for an EMS). If the destination port is a Network port connecting to the outside world, the frame will be copied there and that is the end of it. If the destination port is a DSA Tag port being used to connect to another device, the frame will be copied there, but the frame will egress with a To_Sniffer DSA Tag and the tag will indicate if the frame is from a IMS or an EMS.

The device that receives the To_Sniffer DSA Tag will map those frames to that device's IMD or EMD depending upon the indication in the DSA Tag frame. These frames are not learned from and they are not filtered in any way. They simply progress to the appropriate monitor destination port. If that destination port is another DSA Tag port, the frame egresses unmodified with its original To_Sniffer DSA Tag and the process continues. If the destination port is a Network port, the To_Sniffer DSA Tag is removed and the frame egresses looking as it originally looked at the monitor source port.

Cross-chip port monitoring requires the following:

- Any time a monitor source is enabled in a device, the associated monitor destination must also be defined.
- A monitor source cannot be a DSA Tag port unless it is the CPU's port. Basically it is best if the monitor source ports are Network ports and final monitor destination ports also Network ports.
- If the final monitor destination is cross-chip in another device, the monitor destinations in each device must form a complete path toward the final monitor destination port using DSA Tag enabled ports on both sides of the connections.
- All final Network monitor destination ports (both EMD and EMD) must be isolated from all the other local ports in the switch to prevent these ports from getting flooding frames from non-monitor source ports. The best way to isolate these ports is to use port based VLANs (Section 3.5.6).
- The final egress monitor destination port's 802.1Q mode and VID MemberTag information must be configured to match the egress monitor's source port's configuration. If this is not done, the frame will still egress the final EMD port with the correct data in the frame but its tag mode (i.e., egressing tagged or egressing untagged) may not match the way the frame egressed the original EMS port.

In the devices the following frames that ingress a port are not forwarded for Monitoring:

- Frames received by the MAC but were not stored by the switch due to a lack of memory.
- Frames received with a bad CRC (unless the CRC is fixed by the ForceFCS bit being set on the port).
- Pause frames received.
- Frames with a size less than 64 bytes or greater than the maximum size allowed.

- Frames discarded due to Ingress Rate Limiting

The following frames that ingress a port are forwarded only for Monitoring (i.e., they are not sent to any other port):

- Frames that are discarded for 802.1X Source MAC address security.
- Frames that are discarded for 802.1Q security.
- Tagged frames received when DiscardTagged is enabled.
- Untagged frames received when DiscardUntagged is enabled.
- Frames that are discarded due to the port's PortState setting.
- Frames that aren't mapped to any other port due to DA mapping and VLAN restriction.

3.9 Port Trunking Support

Port trunking is supported by the devices with any combinations of ports (in-chip and cross-chip). The ports that are to be associated with the trunk need to have all the port member's defined with the same TrunkID (in Port Control 1, offset 0x05) and have their Trunk Port bit set to a one (also in Port Control 1). Up to 16 trunk groups are supported with up to 8 ports per trunk group for the 88E6095/88E6095F device. Up to 8 trunks with 8 ports per trunk group are supported for the 88E6092 device. Load balancing is DA/SA based using the Trunk Mask Table (in global 2, offset 0x07).

3.10 Spanning Tree Support

802.1D Spanning Tree is inherently a cross-chip function as the CPU is always outside of the switch chip. It is supported in the devices with the help of an external CPU that runs the actual Spanning Tree algorithm. The devices support Spanning Tree by:

- Detection of BPDU[1] frames entering Network (external switch) ports. These frames are called MGMT (management) frames in the device (see Section 3.4.8). They are detected by loading the BPDU's multicast address (01:80:C2:00:00:00) into the address database with a MGMT Entry_State indicator (see Section 3.4.5) or by using the Rsvd2Cpu bits in MGMT Enable register (global 2 offset 0x03).
- Tunneling of BPDU frames through Blocked ports. Blocked ports are controlled by the Port's PortState bits (see Section 3.5.1 and Section 3.5.10). If a port is in the Blocked state, all frames are discarded except for frames with a DA address that is contained in the address database with a MGMT indicator.
- Redirection of BPDU frames. BPDU frames that enter a Network port need to go to the CPU only, even though they are multicast frames. This task is handled in the BPDU frame detection phase above by mapping the BPDU's multicast address to the CPU port directly or to the port that is to be used to cascade these frames to the CPU (the value of the DPV bits when the address is loaded) or by the source port's CPUPort register (in Port Control 2, offset 0x08). Cascaded BPDU frames egress the 1st device with a To_CPU DSA Tag (Section 3.7.5.2).
- Cascading of BPDU frames. BPDU frames that enter a DSA Tag port must enter it with a To_CPU DSA Tag. These frames are mapped directly, without modification to the source port's CPUPort (in Port Control 2 - offset 0x08). The CPUPort registers are used to form a path from all the Network ports in the switch to the CPU.
- Source Port information. The CPU needs to know the physical source port of the BPDU frame. This information is supplied in the frame's To_CPU DSA Tag (see Section 3.7.5.2) that is sent to the CPU.
- CPU transmission of BPDU frames. The CPU needs to be able to transmit BPDU frames from any physical port of the switch. This control is supported with the From_CPU DSA Tag data that is supplied by the CPU (Section 3.5.16.2). The CascadePort register (in Global Control 2 - global offset 0x1C) is used to route From_CPU frames that are not for this device to the next device.

The CPU and device hardware can support 802.1D Spanning Tree or it can be used to perform simpler bridge loop detection on new link. The only difference is the software that runs on the attached CPU.

Notes

- For Spanning Tree to work, all device-to-device and device-to-CPU interfaces must be configured in DSA Tag mode (Port Control - offset 0x04).
- Each DSA Tag port's CPUPort register must be configured pointing To_CPU DSA Tag frames toward the CPU.
- Likewise, the global CascadePort setting must be configured to route From_CPU DSA Tag frames out to each Network port. The CascadePort limits the configuration options of cascaded devices to run Spanning Tree. A single long chain is supported or a configuration with a single central star configured device is supported.

1. BPDU = Bridge Protocol Data Unit – the frame type used to run the Spanning Tree Protocol

3.11 Embedded Memory

The devices MACs interface directly to the embedded 1/2 Mb (2Kx256) Synchronous SRAM (SSRAM). The SSRAM is running at 125 MHz and the data bus is 256 bits wide. The memory interface provides up to 32 Gigabits per second bandwidth for packet reception/transmission. This memory bandwidth is enough for all of the ports running at full wire speed in full-duplex mode with minimum size frames.

3.12 Interrupt Controller

The devices contain an Interrupt Controller used to merge various interrupts into the device's INTn pin. The Switch Global Control register (Table 83) determines whether or not the INTn pin is asserted when an interrupt occurs. Each interrupt bit in the Switch Global Control register (StatsDone, VTUProb, VTUDone, ATUProb, ATUDone, PHYIntEn, or EEINTn) can be individually masked to enable a switch core interrupt. PHY interrupts are enabled through the PHYIntEn bit, which is in the Switch Global Control register as well.

The Switch Global Status Register (Table 79) reports the interrupt status. When an unmasked interrupt occurs and the INTn asserts, the CPU needs to read the Switch Global Status register to determine the source of the interrupt. If the interrupt came from the switch core (from StatsDone, VTUProb, VTUDone, ATUProb, ATUDone, PHYInt, or EEInt) the INTn pin deasserts after the Switch Global Status register is read (the interrupt status bits are clear on read). If the PHYIntEn bit in the Switch Global Control register is set to a one, the active interrupts enabled in PHY register 0x12 will drive the INTn pin low.

Section 4. Physical Interface (PHY) Functional Description (P0 to P7)

The devices contain IEEE 802.3 100BASE-TX and 10BASE-T compliant media-dependent interfaces on Port 0 to Port 7 for support of Ethernet over unshielded twisted pair (UTP) copper cable. DSP-based advanced mixed signal processing technology supports attachment of up 150 meters of CAT 5 cable to each of these interfaces. An optional, per port, automatic MDI/MDIX crossover detection function gives true “plug and play” capability without the need for confusing crossover cables or crossover ports.

Figure 33: Device Transmit Block Diagram

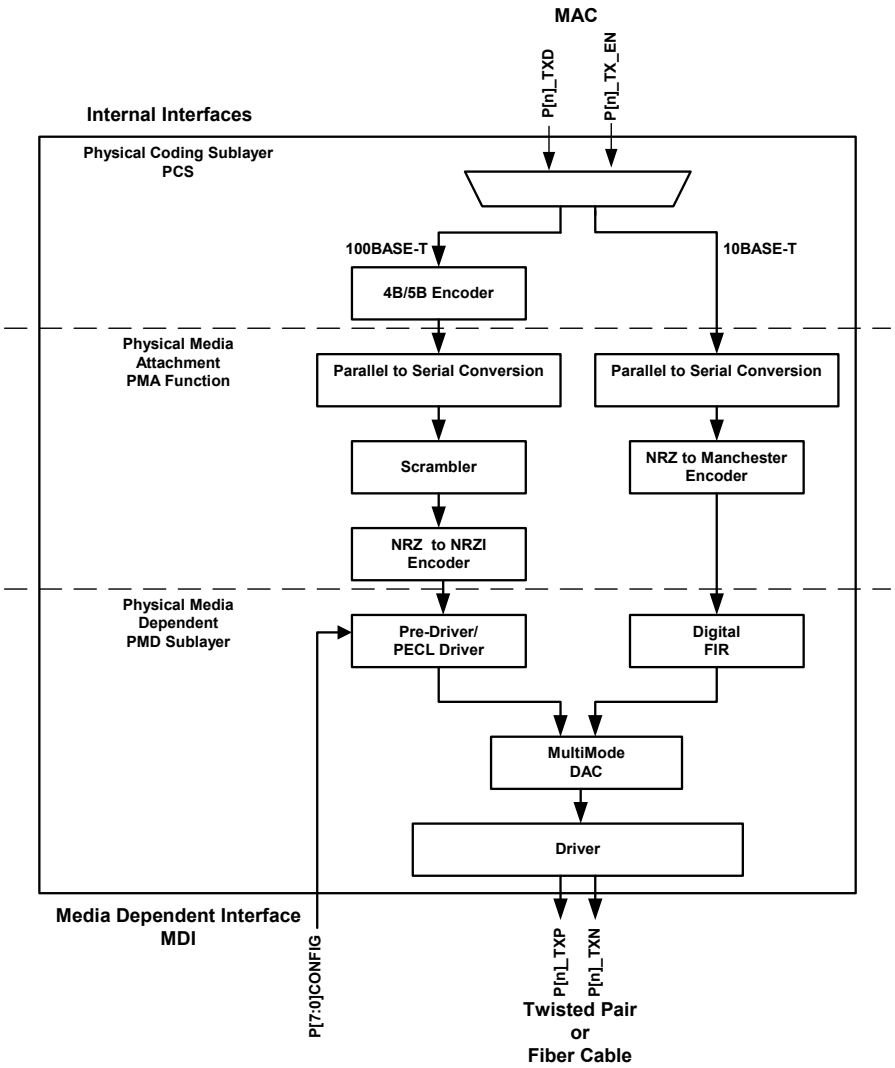

Figure 34: Device Receive Block Diagram

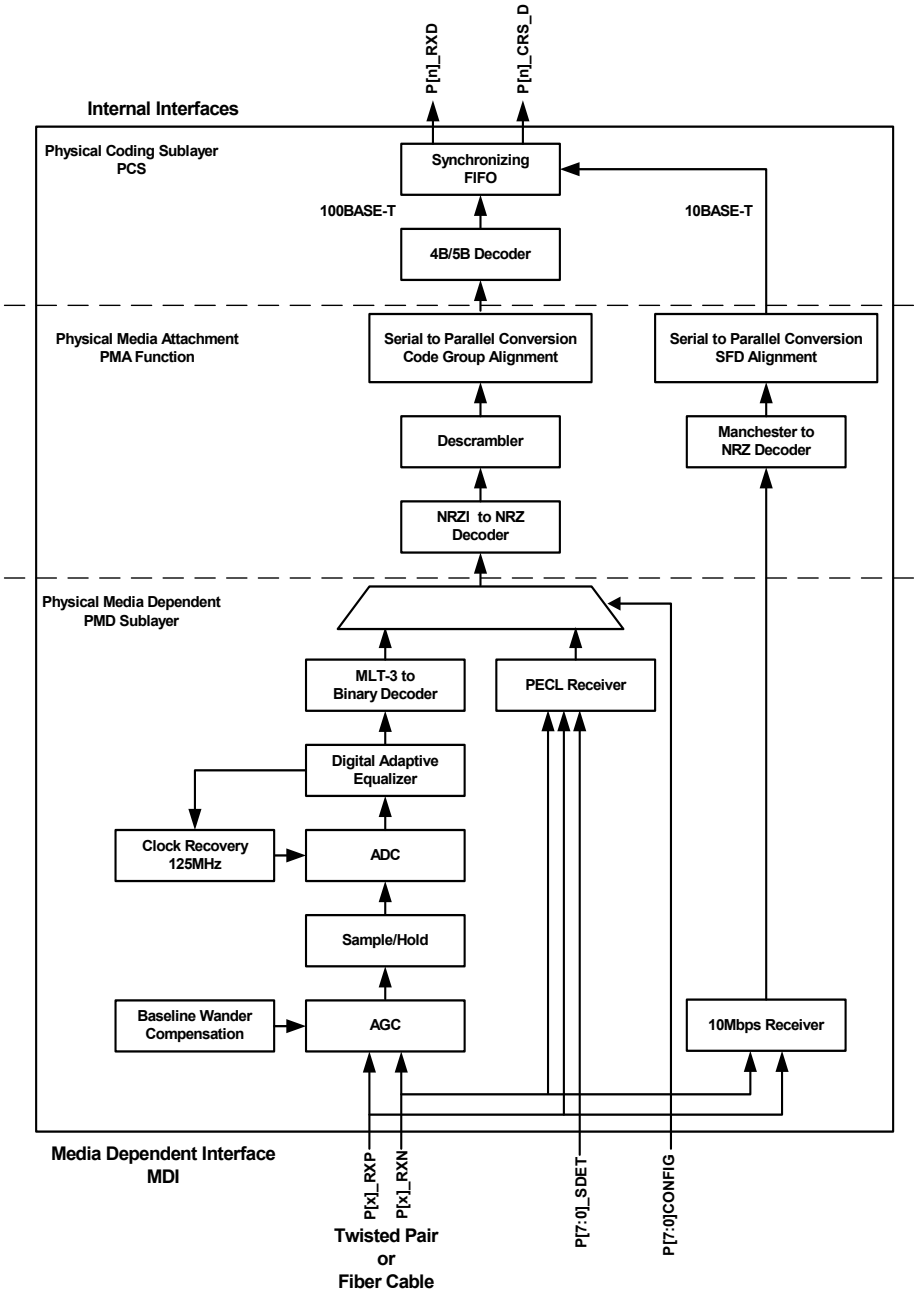

4.1 Transmit PCS and PMA

4.1.1 100BASE-TX Transmitter

The 100BASE-TX transmitter consists of several functional blocks that convert synchronous 4-bit nibble data to a scrambled MLT-3 125 Mbps serial data stream.

4.1.2 4B/5B Encoding

For 100BASE-TX mode, the 4-bit nibble is converted to a 5-bit symbol with /J/K/ start-of-stream delimiters and /T/R/ end-of-stream delimiters inserted as needed. The 5-bit symbol is then serialized and scrambled.

4.1.3 Scrambler

In 100BASE-TX mode, the transmit data stream is scrambled in order to reduce radiated emissions on the twisted pair cable. The data is scrambled by exclusive ORing the NRZ signal with the output of an 11-bit wide linear feedback shift register (LFSR), which produces a 2047-bit repeating pseudo-random sequence. The scrambler reduces peak emissions by randomly spreading the signal energy over the transmit frequency range and eliminating peaks at certain frequencies.

Note

The enabling and disabling of the scrambler and the far end fault generator are controlled in the same way as for the descrambler detection and far end fault detection on the receive side.

4.1.4 NRZ to NRZI Conversion

The data stream is converted from NRZ to NRZI.

4.1.5 Pre-Driver and Transmit Clock

The devices use an all-digital clock generator circuit to create the various receive and transmit clocks necessary for 100BASE-TX, 100BASE-FX, and 10BASE-T modes of operation.

For 100BASE-TX mode, the transmit data is converted to MLT-3-coded symbols. The digital time base generator (TBG) produces the locked 125 MHz transmit clock.

For 100BASE-FX mode, NRZI data is presented directly to the multimode DAC.

For 10BASE-T mode, the transmit data is converted to Manchester encoding. The digital time base generator (TBG) produces the 10 MHz transmit reference clock as well as the over-sampling clock for 10BASE-T waveshaping.

4.1.6 Multimode Transmit DAC

The multimode transmit digital to analog converter (DAC) transmits MLT-3 coded symbols in 100BASE-TX mode,NRZI symbols in 100BASE-FX mode, and Manchester-coded symbols in 10BASE-T mode. The transmit DAC utilizes a direct drive current driver which is well balanced to produce very low common mode transmit noise.

In 100BASE-TX mode, the multimode transmit DAC performs slew control to minimize high frequency EMI.

In 100BASE-FX mode, the pseudo ECL level is generated through external resistive terminations.

In 10BASE-T mode, the multimode transmit DAC generates the needed pre-equalization waveform. This pre-equalization is achieved by using a digital FIR filter.

4.2 Receive PCS and PMA

4.2.1 10-BASE-T/100BASE-TX Receiver

The differential RXP and RXN pins are shared by the 100BASE-TX, 100BASE-FX and Port 1 and 10BASE-T receivers.

The 100BASE-TX receiver consists of several functional blocks that convert the scrambled MLT-3 125 Mbps serial data stream to the synchronous 4-bit nibble data presented to the MAC interfaces.

4.2.2 AGC and Baseline Wander

In 100BASE-TX mode, after input to the AGC block the signal is compensated for baseline wander with a digitally controlled Digital to Analog converter (DAC). It automatically removes the DC offset from the input before it reaches the ADC input.

4.2.3 ADC and Digital Adaptive Equalizer

In 100BASE-T mode, an analog to digital converter (ADC) samples and quantizes the input analog signal and sends the result into the digital adaptive equalizer. The digital adaptive equalizer removes inter-symbol interference at the receiver. The digital adaptive equalizer takes unequalized signals from ADC output and uses a combination of feedforward equalizer (FFE) and decision feedback equalizer (DFE) for the best optimized signal-to-noise (SNR) ratio.

4.2.4 Digital Phased Locked Loop (DPLL)

In 100BASE-TX mode, the receive clock is locked to the incoming data stream and extracts a 125 MHz reference clock. The input data stream is quantized by the recovered clock and sent through to the digital adaptive equalizer from each port.

Digital interpolator clock recovery circuits are optimized for MLT-3, NRZI, and Manchester modes. A digital approach makes the device's receiver paths robust to the presence of variations in process, temperature, on-chip noise, and supply voltage.

4.2.5 NRZI to NRZ Conversion

In 100BASE-TX mode, the recovered 100BASE-TX NRZI signal from the receiver is converted to NRZ data, descrambled, aligned, parallelized, and 5B/4B decoded.

4.2.6 Descrambler

The descrambler is initially enabled upon hardware reset if 100BASE-TX is selected. The scrambler can be enabled or disabled via software by setting the DisScrambler bit (Table 127).

The descrambler "locks" to the descrambler state after detecting a sufficient number of consecutive idle code-groups. The receiver does not attempt to decode the data stream unless the descrambler is locked. Once locked, the descrambler continuously monitors the data stream to make sure that it has not lost synchronization.

The receiver descrambles the incoming data stream by exclusive ORing it with the output of an 11-bit wide linear feedback shift register (LFSR), which produces a 2047-bit non-repeating sequence.

The descrambler is always forced into the "unlocked" state when a link failure condition is detected or when insufficient idle symbols are detected.

4.2.7 Serial to Parallel Conversion and 5B/4B Code-Group Alignment

The Serial to Parallel /Symbol Alignment block performs serial to parallel conversion and aligns 5B code-groups to a nibble boundary.

4.2.8 5B/4B Decoder

The 5B/4B decoder translates 5B code-groups into 4B nibbles to be presented to the MAC interfaces. The 5B/4B code mapping is shown in Table 44.

4.2.8.1 FIFO

The 100BASE-X or 10BASE-T packet is placed into the FIFO in order to correct for any clock mismatch between the recovered clock and the reference clock REFCLK.

4.2.8.2 100BASE-FX Receiver

In 100BASE-FX mode, a pseudo-ECL (PECL) receiver is used to decode the incoming NRZI signal passed to the NRZI-NRZ decoder. The NRZI signal from the receiver is converted to NRZ data, aligned, parallelized, and 5B/4B decoded as in the 100BASE-TX mode.

4.2.8.3 Far End Fault Indication (FEFI)

When 100BASE-FX is selected and bit 0 of CONFIG_A is high at hardware reset, then the far end fault detect (FEFD) circuit is enabled. The FEFD enable state can be overridden by programming the DisFEFI bit (Table 127).

Note

The FEFI function is always disabled if 100BASE-TX is selected.

4.2.8.4 10BASE-T Receiver

In 10BASE-T mode, the recovered 10BASE-T signal is decoded from Manchester to NRZ and then aligned. The alignment is necessary to ensure that the start of frame delimiter (SFD) is aligned to the nibble boundary.

In 10BASE-T mode, a receiver is used to decode the differential voltage offset of the Manchester data. Carrier sense is decoded by measuring the magnitude of the voltage offset.

In this mode, the recovered 10BASE-T signal is decoded from Manchester to NRZ data. The data stream is converted from serial to parallel format and aligned. The alignment is necessary to insure that the start of frame delimiter (SFD) is aligned to a byte or nibble boundary.

For cable lengths greater than 100 meters, the incoming signal has more attenuation. Hence, the receive voltage threshold should be lowered via the Extended Distance bit in the PHY Specific Control Register (Table 127).

Table 44: 5B/4B Code Mapping

PCS Code-Group [4:0] 4 3 2 1 0	Name	TXD/RXD <3:0> 3 2 1 0	Interpretation
1 1 1 1 0	0	0 0 0 0	Data 0
0 1 0 0 1	1	0 0 0 1	Data 1
1 0 1 0 0	2	0 0 1 0	Data 2
1 0 1 0 1	3	0 0 1 1	Data 3
0 1 0 1 0	4	0 1 0 0	Data 4
0 1 0 1 1	5	0 1 0 1	Data 5
0 1 1 1 0	6	0 1 1 0	Data 6
0 1 1 1 1	7	0 1 1 1	Data 7
1 0 0 1 0	8	1 0 0 0	Data 8
1 0 0 1 1	9	1 0 0 1	Data 9
1 0 1 1 0	A	1 0 1 0	Data A
1 0 1 1 1	B	1 0 1 1	Data B
1 1 0 1 0	C	1 1 0 0	Data C
1 1 0 1 1	D	1 1 0 1	Data D
1 1 1 0 0	E	1 1 1 0	Data E
1 1 1 0 1	F	1 1 1 1	Data F
1 1 1 1 1	I	Undefined	IDLE; used as inter-stream fill code
1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
0 1 1 0 1	T	Undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
0 0 1 1 1	R	Undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
0 0 0 0 0	V	Undefined	Invalid code
0 0 0 0 1	V	Undefined	Invalid code
0 0 0 1 0	V	Undefined	Invalid code
0 0 0 1 1	V	Undefined	Invalid code
0 0 1 0 1	V	Undefined	Invalid code
0 0 1 1 0	V	Undefined	Invalid code
0 1 0 0 0	V	Undefined	Invalid code
0 1 1 0 0	V	Undefined	Invalid code
1 0 0 0 0	V	Undefined	Invalid code
1 1 0 0 1	V	Undefined	Invalid code

4.2.9 Setting Cable Characteristics

Since cable characteristics are different between unshielded twisted pair and shielded twisted pair cable, optimal receiver performance can be obtained in 100BASE-TX and 10BASE-T modes by setting the TPSelect bit in the PHY Specific Control Register (Table 127) for cable type.

4.2.10 Scrambler/Descrambler

The scrambler block is initially enabled upon hardware reset if 100BASE-TX is selected. If 100BASE-FX or 10BASE-T is selected, the scrambler is disabled by default. The scrambler is controlled by programming the DisScrambler bit in the PHY Specific Control Register (Table 127).

The scrambler setting is also controlled by hardware configuration at the end of hardware reset. Table 45 shows the effect of various configuration settings on the scrambler.

Table 45: Scrambler Settings

P[1:0]_CONFIG (If FX is selected)	DisScrambler Bit ()	Scrambler/ Descrambler
High	HW reset to 1	Disabled
Low	HW reset to 0	Enabled
X	User set to 1	Disabled
X	User set to 0	Enabled (in copper mode only)

4.2.11 Link Monitor

The link monitor is responsible for determining whether link is established with a link partner.

In 10BASE-T mode, link monitor function is performed by detecting the presence of the valid link pulses on the RXP/N pins.

In 100BASE-TX mode, the link is established by scrambled idles

In 100BASE-FX mode, the external fiber-optic receiver performs the signal detection function and communicates this information with the devices through SDET pins for all individual ports.

If Force Link Good is asserted (ForceLink bit is set high - PHY Specific Control Register, Table 127), the link is forced to be good, and the link monitor is bypassed. Pulse checking is disabled if Auto-Negotiation is disabled, and DisNLPCheck (PHY Specific Control Register, Table 127) is set high. If Auto-Negotiation is disabled and DisNLPGen (PHY Specific Control Register, Table 127) is set high, then the link pulse transmission is disabled.

4.2.12 Auto-Negotiation

Auto-Negotiation is initiated upon any of the following conditions:

- Power up reset
- Hardware reset
- Software reset (SWReset bit—PHY Control Register, Table 117)
- Restart Auto-Negotiation (RestartAneg bit—PHY Control Register, Table 117)
- Transition from power down to power up (PwrDwn bit—PHY Control Register, Table 117)
- Change from the linkfail state to the link-up state

If Auto-Negotiation is enabled, the devices negotiate with its link partner to determine the speed and duplex mode at which to operate. If the link partner is unable to Auto-Negotiate, the devices go into the parallel detect mode to determine the speed of the link partner. Under parallel detect mode, the duplex mode is fixed at half-duplex.

After hardware reset, Auto-Negotiation can be enabled and disabled via the AnegEn bit (PHY Control Register - Table 117). When Auto-Negotiation is disabled, the speed and duplex can be changed via the SpeedLSB and Duplex bits (PHY Control Register - Table 117), respectively. The abilities that are advertised can be changed via the Auto-Negotiation Advertisement Register (Table 121).

4.2.13 Register Update

Changes to the AnegEn, SpeedLSB, and Duplex bits (Table 117) do not take effect unless one of the following takes place:

- Software reset (SWReset bit - Table 117)
- Restart Auto-Negotiation (RestartAneg bit - Table 117)
- Transition from power down to power up (PwrDwn bit - Table 117)
- Loss of the link

The Auto-Negotiation Advertisement register (Table 121) is internally latched once every time Auto-Negotiation enters the ability detect state in the arbitration state machine. Hence, a write into the Auto-Negotiation Advertisement Register has no effect once the devices begin to transmit Fast Link Pulses (FLPs). This guarantees that a sequence of FLPs transmitted is consistent with one another.

The Next Page Transmit register (Table 125) is internally latched once every time Auto-Negotiation enters the next page exchange state in the arbitration state machine.

4.2.14 Next Page Support

The devices support the use of next page during Auto-Negotiation. By default, the received base page and next page are stored in the Link Partner Ability register - Base Page (Table 122). The devices have an option to write the received next page into the Link Partner Next Page register - Table 123 - (similar to the description provided in the IEEE 802.3ab standard) by programming the Reg8NxtPg bit (PHY Specific Control Register - Table 127).

4.2.15 Status Registers

Once the devices complete Auto-Negotiation, the various statuses in the PHY Status (Table 128), Link Partner Ability (Next Page)(Table 123), and Auto-Negotiation Expansion(Table 124) registers are updated. Speed, duplex, page received, and Auto-Negotiation completed statuses are also available in the PHY Specific Status (Table 128) and PHY Interrupt Status (Table 130) registers.

4.3 Power Management

The devices support advanced power management modes that conserve power. These features are only recommended for power-critical designs only.

4.3.1 Low Power Modes

Two low power modes are supported in the devices.

- IEEE 22.2.4.1.5 compliant power down
- Energy Detect+™

IEEE 22.2.4.1.5 power down compliance allows for the PHY to be placed in a low-power consumption state by register control.

Energy Detect+™ allows the devices to wake up when energy is detected on the wire with the additional capability to wake up a link partner. The 10BASE-T link pulses are sent once every second while listening for energy on the line.

4.3.2 MAC Interface and PHY Configuration for Low Power Modes

The devices have one CONFIG bit dedicated to support the low power modes. Energy Detect may be configured using the CONFIG_B pin, or through PHY Specific Control Register 0x10.

Low power modes are also register programmable. The EDet bit (Table 127) enables the user to turn on Energy Detect+™. When the low power mode is not selected, the PwrDwn bit (Table 117) can be used. If during the energy detect mode, the PHY wakes up and starts operating in normal mode, the EDet bit settings are retained. When the link is lost and energy is no longer detected, the devices return to the mode stored in the EDet bit.

Table 46 shows how these modes are entered

.

Table 46: Operating Mode Power Consumption

Power Mode	Est. Power	How to Activate Mode
IEEE Power down	See Section 8.	PwrDwn bit write (Table 117)
Energy Detect+™	Same as IEEE Power down mode—see Section 8.	Configuration option & register EDet bit write (Table 127)

4.3.3 IEEE Power Down Mode

The standard IEEE power down mode (22.2.4.1.5) is entered by setting the PwrDwn (Table 117) bit equal to one. In this mode, the PHY does not respond to any MAC interface signals except the MDC/MDIO. It also does not respond to any activity on the CAT 5 cable.

In this power down mode, the PHY cannot wake up on its own by detecting activity on the CAT 5 cable. It can only wake up by clearing the PwrDwn bit to 0.

4.3.3.1 Energy Detect +™

In this mode, the PHY sends out a single 10 Mbps NLP (Normal Link Pulse) every one second. If the devices are in Energy Detect+™ mode, it can wake a connected device. When ENA_EDET is 1, the mode of operation is Energy Detect+™. The devices also respond to MDC/MDIO.

4.4 Far End Fault Indication (FEFI)

Far-end fault indication provides a mechanism for transferring information from the local station to the link partner that a remote fault has occurred in 100BASE-FX mode.

A remote fault is an error in the link that one station can detect while the other one cannot. An example of this is a disconnected fiber at a station's transmitter. This station is receiving valid data and detects that the link is good via the link monitor, but is not able to detect that its transmission is not propagating to the other station.

A 100BASE-FX station that detects this remote fault modifies its transmitted idle stream pattern from all ones to a group of 84 ones followed by one zero. This is referred to as the FEFI idle pattern.

The FEFI function is controlled by CONFIG_A connection and the DisFEFI bit (Table 127).

Table 47 shows the various configuration settings affecting the FEFI function on hardware reset.

Table 47: FEFI Select

EN_FEFI Config_A) Conncection	FEFI	DisFEFI Bit (Table 127)
VSS	Disabled	HW reset to 1
P0_LED0	Enabled	HW reset to 0
P0_LED1	Disabled	HW reset to 1
P0_LED2	Enabled	HW reset to 0
P1_LED0	Disabled	HW reset to 1
P1_LED1	Enabled	HW reset to 1
P1_LED2	Disabled	HW reset to 0
VDDO	Enabled	HW reset to 1

4.5 Virtual Cable Tester® Feature

The device's Virtual Cable Tester (VCT™) feature uses Time Domain Reflectometry (TDR) to determine the quality of the cables, connectors, and terminations. Some of the possible problems that can be diagnosed include opens, shorts, cable impedance mismatch, bad connectors, termination mismatch, and bad magnetics.

The devices transmit a signal of known amplitude (+1V) down each of the two pairs of an attached cable. It will conduct the cable diagnostic test on each pair, testing the TX and RX pairs sequentially. The transmitted signal will continue down the cable until it reflects off of a cable imperfection. The magnitude of the reflection and the time it takes for the reflection to come back are shown in the VCT registers (Table 137 and Table 137) on the AmpRfln and DistRfln bits respectively.

Using the information from the VCT registers (Table 137 and Table 137), the distance to the problem location and the type of problem can be determined. For example, the time it takes for the reflection to come back, can be converted to distance using the cable fault distance trend line tables in Table 137 and Figure 47. The polarity and magnitude of the reflection together with the distance will indicate the type of discontinuity. For example, a +1V reflection will indicate an open close to the PHY and a -1V reflection will indicate a short close to the PHY.

When the cable diagnostic feature is activated by setting the ENVCT bit to one (Table 137), a pre-determined amount of time elapses before a test pulse is transmitted. This is to ensure that the link partner loses link, so that it stops sending 100BASE-TX idles or 10 Mbit data packets. This is necessary to be able to perform the TDR test. The TDR test can be performed either when there is no link partner or when the link partner is Auto-Negotiating or sending 10 Mbit idle link pulses. If the devices receive a continuous signal for 125 ms, it will declare test failure because it cannot start the TDR test. In the test fail case, the received data is not valid. The results of the test are also summarized in the VCTTst bits (Table 137 and Table 138).

- 11 = Test fail (The TDR test could not be run for reasons explained above)
- 00 = valid test, normal cable (no short or open in cable)
- 10 = valid test, open in cable (Impedance > 333 ohms)
- 01 = valid test, short in cable (Impedance < 33 ohms)

The definition for shorts and opens is arbitrary and the user can define it anyway they desire using the information in the VCT registers (Table 137 and Table 138). The impedance mismatch at the location of the discontinuity could also be calculated knowing the magnitude of the reflection. Refer to the App Note "Virtual Cable Tester -- How to use TDR results" for details.

4.6 Data Terminal Equipment (DTE) Detect

The devices support the Data Terminal Equipment (DTE) detect function. The IEEE 802.3af - 2003 DTE power scheme requires no role of the PHY to detect a DTE link partner that requires power; however, the devices offer the ability to detect a DTE link partner that requires power.

The DTE power function can be enabled after a hardware reset by writing to the Enable DTE Detect bit (Register 16.15), followed by a software reset. When DTE Detect is enabled, the devices will first monitor for any activity transmitted by the link partner. If the link partner is active, then the link partner has power and can link with the device. If there is no activity coming from the link partner, DTE Detect engages, and special pulses are sent to detect if the link partner requires DTE power. If the link partner has a low pass filter (or similar fixture) installed, the link partner will be detected as requiring DTE power.

The detection of the DTE power requirement can be reported to the devices in two ways.

- The DTE Detect Status bit (Register 17.15) immediately asserts as soon as the link partner is detected as a device requiring DTE power. The devices will continually send special link pulses to detect if the link partner requires DTE power.
- If the DTE Detect Status Change Interrupt bit (Register 18.15) is enabled, an interrupt can be generated if a DTE powered device is detected. The devices will then update the DTE Detect Status Drop bit (Register 22.15) and the DTE Detect Status bit (Register 17.15).

If a link partner that requires DTE power is unplugged, the DTE Detect Status bit (Register 17.15) will drop after a user controlled delay (default is 20 seconds - DTE Detect Status Drop bit (Register 22.15:14) = 0x04), since the lowpass filter (or similar fixture) is removed during power up.

A detailed description of the register bits used for DTE power detection for sthe device is shown in Table 48.

Note

Auto-Negotiation must be enabled to use DTE detect. The DTE detect must be turned off when performing the Virtual Cable Tester® (VCT™) test.

Table 48: Registers for DTE Detect

Register	Description
16.15 - **Enable DTE Detect**	1 = Enable DTE detect 0 = Disable DTE detect Default at HW reset: 0 At SW reset: Retain
17.15 DTE Detect status	1 = DTE detected 0 = DTE not detected Default at HW reset: 0 At SW reset: 0
18.15 - **DTE Detect state changed interrupt enable**	1 = Interrupt enable 0 = Interrupt disable Default at HW reset: 0 At SW reset: retain
19.15 DTE Detect state changed interrupt	1 = Changed 0 = No change Default at HW reset: 0 At SW reset: 0
22.15:12 DTE detect status drop	Once the devices no longer detect the link partner's DTE filter, the devices will wait a period of time before clearing the DTE detection status bit (17.15). The wait time is 5 seconds multiplied by the value of these bits. Default at HW reset: 0x4 At SW reset: retain

4.7 Auto MDI/MDIX Crossover

The devices automatically determine whether or not it needs to cross over between pairs so that an external crossover cable is not required. If the devices interoperate with a device that cannot automatically correct for crossover, the devices make the necessary adjustment prior to commencing Auto-Negotiation. If the devices interoperate with a device that implements MDI/MDIX crossover, a random algorithm as described in IEEE 802.3 Section 40.4.4 determines which device performs the crossover.

When the devices interoperate with legacy 10BASE-T devices that do not implement Auto-Negotiation, the devices follow the same algorithm as described above since link pulses are present. However, when interoperating with legacy 100BASE-TX devices that do not implement Auto-Negotiation (i.e. link pulses are not present), the devices use signal detect to determine whether or not to crossover.

The Auto MDI/MDIX crossover function can be disabled via the AutoMDI[X] bits (Table 127).

The devices are set to MDIX mode by default if auto MDI/MDIX crossover is disabled at hardware reset.

The pin mapping in MDI and MDIX modes is specified in Table 49.

Table 49: MDI/MDIX Pin Functions

Physical Pin	MDIX		MDI	
	100BASE-TX	10BASE-T	100BASE-TX	10BASE-T
TXP/TXN	Transmit	Transmit	Receive	Receive
RXP/RXN	Receive	Receive	Transmit	Transmit

4.8 Copper Line Loopback

Line loopback allows a link partner to send frames into the devices to test the transmit and receive data path. Frames from a link partner into the PHY, before reaching the MAC interface pins, are looped back and sent out on the copper line. They are also sent to the MAC. The packets received from the MAC are ignored during copper line loopback. Refer to Figure 35. This allows the link partner to receive its own frames.

Figure 35: Line Loopback Data Path

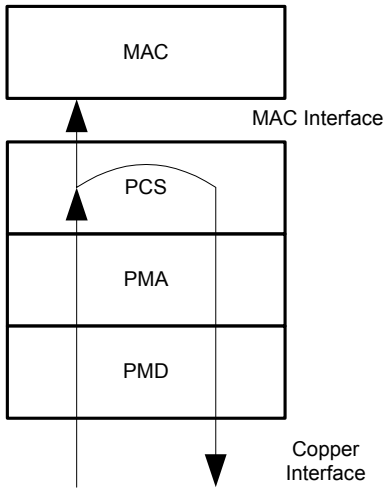

Before enabling the line loopback feature, the PHY must first establish link to another PHY link partner. If Auto-Negotiation is enabled, both link partners should advertise the same speed and full-duplex. If Auto-Negotiation is disabled, both link partners need to be forced to the same speed and full-duplex. Once link is established, enable the line loopback mode by writing to Register 28.4.

Line loopback mode works for 10BASE-T100BASE-TX, and 100BASE-FX modes.

- 28.4 = 1 (Enable line loopback)
- 28.4 = 0 (Disable line loopback)

4.9 LED Interface

The LEDs can either be controlled by the PHY or controlled externally, independent of the state of the PHY.

External control is achieved by writing to the PHY Manual LED Override register (Table 136). Any of the LEDs can be turned on, off, or made to blink[1] at variable rates independent of the state of the PHY. This independence eliminates the need for driving LEDs from the MAC or the CPU. If the LEDs are driven from the CPU located at the back of the board, the LED lines crossing the entire board will pick up noise. This noise will cause EMI issues. Also, PCB layout will be more difficult due to the additional lines routed across the board.

When the LEDs are controlled by the PHY, the activity of the LEDs is determined by the state of the PHY. Each LED can be programmed to indicate various PHY states, with variable blink rate.

Any one of the LEDs can be controlled independently of the other LEDs (i.e one LED can be externally controlled while another LED can be controlled by the state of the PHY).

The LED interface supports a 3-pin parallel interface for each port or a serial interface for all ports.

The devices can be programmed to display serial LED statuses in single or dual LED modes. Some of the statuses can be pulse stretched. Pulse stretching is necessary because the duration of these status events might be too short to be observable on the LEDs. The pulse stretch duration can be programmed via the PulseStretch bits (Table 135). The default pulse stretch duration is set to 170 to 340 ms. The pulse stretch duration applies to all applicable LEDs.

Some of the statuses indicate multiple events by blinking LEDs. The blink period can be programmed via the BlinkRate bits (Table 135). The default blink period is set to 84 ms. The blink rate applies to all applicable LEDs.

4.9.1 Parallel LED Interface

The parallel LED interface displays 3 different statuses for each port. LED2, LED1, and LED0 pins are used for each port. The LED Parallel Select Register specifies which single LED mode status to display on the LED pins. The defaults to display shown in Table 133 are based on the LED_DEF[1] and LED_DEF[0] values during hardware configuration through the CONFIG_A pin—see Table 4.

Table 50: Parallel LED Hardware Defaults

LED Mode —set by CONFIG_A at reset	P[7:0]_LED2	P[7:0]_LED1	P[7:0]_LED0
0	LINK	RX	TX
1	LINK	ACT	SPEED
2	LINK/RX	TX	SPEED
3	LINK/ACT	DUPLEX/COLX	SPEED

Table 133 shows additional display modes that can be set up by software after startup. Table 51 includes these extra modes that cannot be set up by hardware configuration pins at reset.

1. Energy Detect (Section 4.3) must be disabled on ports that are configured to blink an LED but don't have a link established.

Table 51: Parallel LED Display Interpretation

Status	Description
COLX	Low = Collision activity High = No collision activity This status is pulse stretched to 170 ms.
ERROR	Low = Jabber, received error, false carrier, or FIFO over/underflow occurred High = None of the above occurred This status is pulse stretched to 170 ms.
DUPLEX	Low = Full-duplex High = Half-duplex
LINK	Low = Link up High = Link down
RX	Low = Receive activity High = No receive activity This status is pulse stretched to 170 ms.
TX	Low = Transmit activity High = No transmit activity This status is pulse stretched to 170 ms.
ACT	Low = Transmit or received activity High = No transmit or receive activity This status is pulse stretched to 170 ms.
SPEED	Low = Speed is 100 Mbps High = Speed is 10 Mbps
LINK/RX	Low = Link up High = Link down Blink = Receive activity (blink rate is 84 ms active then 84ms inactive) The receive activity is pulse stretched to 84 ms.
LINK/ACT	Low = Link up High = Link down Blink = transmit or receive activity (blink rate is 84 ms active then 84 ms inactive) The transmit and receive activity is pulse stretched to 84 ms
DUPLEX/COLX	Low = Full-duplex High = Half-duplex Blink = collision activity (blink rate is 84 ms active then 84 ms inactive) The collision activity is pulse stretched to 84 ms.
ACT (blink mode)	Low = Full-duplex High = Half-duplex Blink = Collision activity (blink rate is 84 ms active then 84 ms inactive) The collision activity is pulse stretched to 84 ms.

4.9.2 Using Two Color LEDs

Two color LEDs are supported with the parallel LED pins on the devices. Each Px_LED[2:0] pin can be independently controlled to be active low (the default) or active high. They can also be independently selected to be active only if the port is in 10 Mbps mode or only 100 Mbps mode or either speed (the default). This can be configured by using the PHY Manual LED override register (Table 136).

The following lists some possible applications of two color LEDs:

Case One:

- Off - No Link
- Green - Link at 100 Mbps
- Yellow - Link at 10 Mbps
- Blink Green - Activity at 100 Mbps
- Blink Yellow - Activity at 10 Mbps

Case Two:

- Off - No Link
- Green - Link, or VCT™ Good
- Blink Green - Activity
- Red - VCT Short
- Blink Red - Testing VCT

Examples of each of these cases follow.

Figure 36: Possible Solutions for Case One

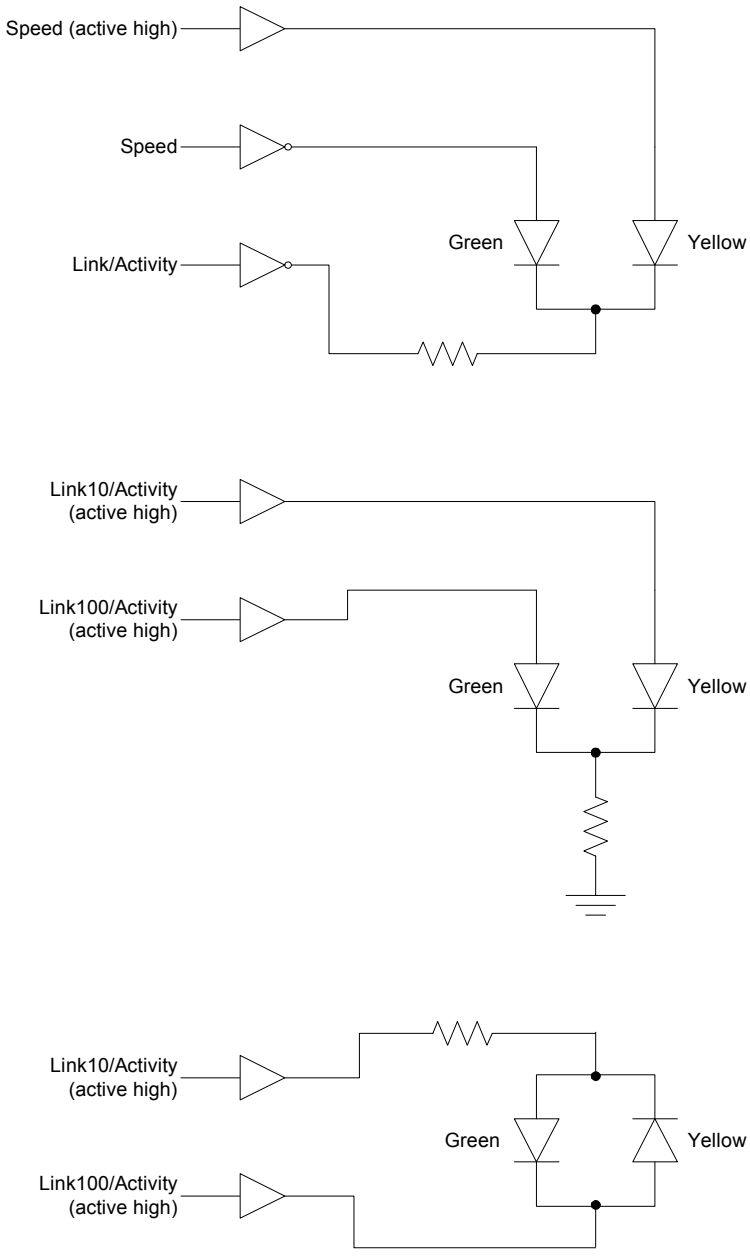

Figure 37: Possible Solutions for Case Two

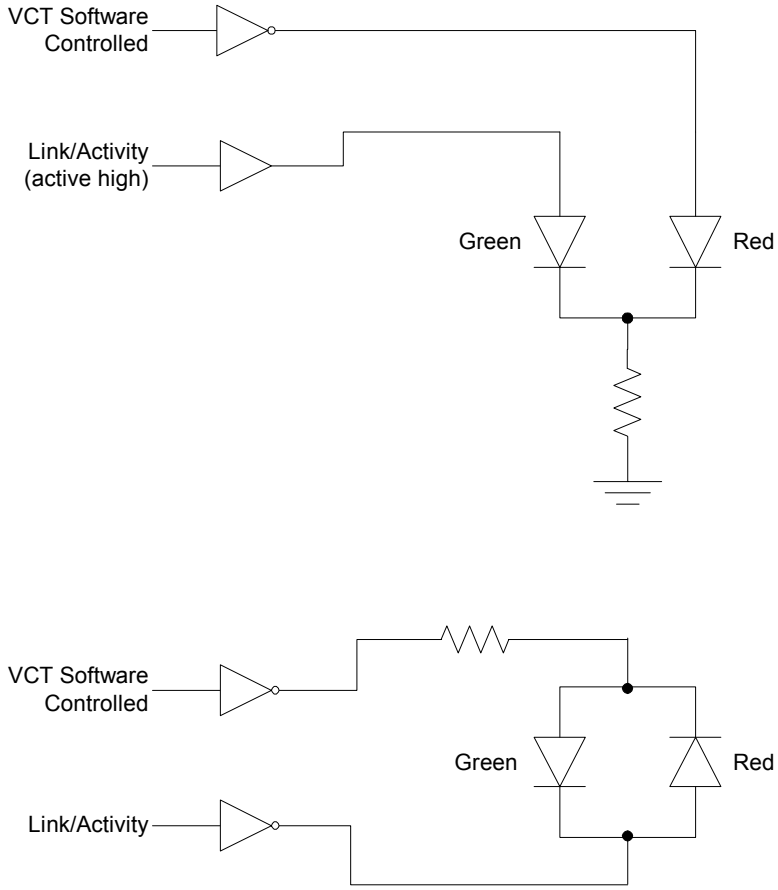

4.9.3 Serial LED Interface (88E6095F Device Only)

LEDSER, LEDENA, and LEDCLK pins of the 88E6095F device are used for the serial interface. The CONFIG_A pin is used to select 1 of 4 possible modes. The serial LED interface can display up to 11 different statuses in 100BASE-TX and 10BASE-T modes. Statuses to display, pulse stretching, and blink mode can be programmed via the LED Stream Select for Serial LEDs register and the PHY LED Control register bits 5:0.

4.9.4 Single and Dual LED Modes

4.9.4.1 Single LED Display Mode

In the single LED display mode, the same status is driven on both status 100 and status 10 positions in the bit stream. However, the LEDENA signal asserts only over the status that is set and de-asserts over the other position that is turned off in the bit stream. For example, DUPLEX shows the same status for DUPLEX100 and DUPLEX10. However, LEDENA signal is high over Duplex100 position only for one clock period. Refer to Figure 38 for more details.

4.9.4.2 Dual LED Display Mode

In the dual LED display mode, two LEDs are used: one for 10 Mbps, and one for 100 Mbps activity. A different status is driven on status 100 and status 10 positions in the bit stream. In this case, the LEDENA signal asserts over both 100 and 10 positions in the bit stream. For example, LEDENA signal asserts over COLX100 and COLX10 in Figure 38. LEDENA signal is high for two clock periods. If a particular status bit is turned off, then LEDENA is not asserted in both positions. Figure 38 illustrates single and dual LED modes.

Figure 38: Serial LEDENA High Clocking with COLX in Dual Mode, Error Off, and DUPLEX in Single Mode

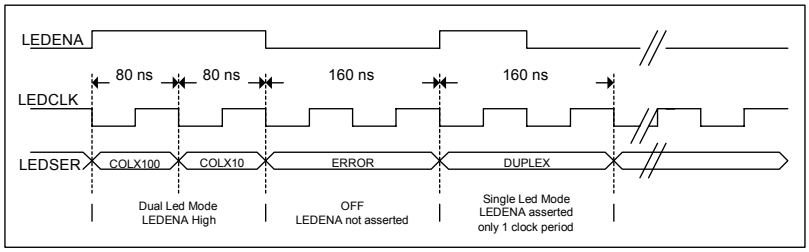

The bit stream on LEDSER can be clocked into a shift register with LEDENA as the shift enable signal as shown in Figure 39. The rate of update of the serial LED interface is controlled by programming the PHY LED Control Register bits 8:6. The default value is set to 42 ms.

Figure 39: Serial LED Conversion

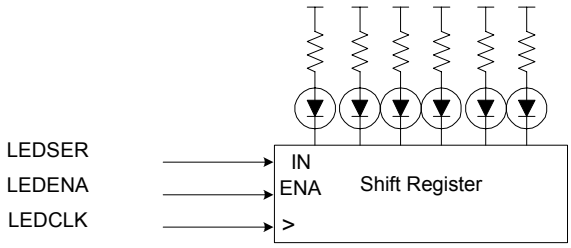

After the LED data is shifted into the correct position, the shift sequence is suspended to allow the appropriate LEDs to light or extinguish depending on status. The LED implementation used in the 88E6095F device is self-synchronizing. The default display options are given in Table 52.

Table 52: Serial LED Hardware Defaults (S = Single)

LED_DEF[1] (At Reset)	LED_DEF[0] (At Reset)	COLX	ERROR	DUPLEX	DUPLEX/ COLX	SPEED	LINK	TX	RX	ACT	LINK/RX	LINK/ACT
0	0	Off	S	Off	S	S	Off	Off	Off	Off	Off	S
0	1	S	S	S	Off	S	Off	S	Off	Off	S	Off
1	0	S	S	S	Off	S	S	Off	Off	S	Off	Off
1	1	S	S	S	Off	S	S	S	S	Off	Off	Off

The LED status bits are output in the order shown on the LEDSER pin synchronously to LEDCLK. All statuses for Port 0 are sent out first followed by those for Ports 1 through 7. Each bit in the stream occupies a period of 80 ns.

Table 53: Serial LED Display Order

<COLX100> → <COLX10> → <ERROR100> → <ERROR10> → <DUPLEX100> → <DUPLEX10>→ <DUPLEX/COLX100> → <DUPLEX/COLX10> → <SPEED100> → <SPEED10> → <LINK100>→ <LINK10> → <TX100> → <TX10> → <RX100> → <RX10> → <ACT100> → <ACT10> → <LINK/RX100>→ <LINK/RX10> → <LINK/ACT100> → <LINK/ACT10>

The following tables show the status events that can be displayed by programming the device in single and dual LED display modes

Table 54: Single LED Display Mode

Status	Description
COLX	Low = Collision activity High = No collision activity This status has a default pulse stretch duration of 170 ms.
ERROR	Low = Jabber, received error, false carrier, or FIFO over/underflow occurred High = None of the above occurred This status has a default pulse stretch duration of 170 ms.
DUPLEX	Low = Full-duplex High = Half-duplex
DUPLEX/COLX	Low = Full-duplex High = Half-duplex Blink = Collision activity (default blink rate is 84 ms active then 84 ms inactive) The collision activity has a default pulse stretch duration of 84 ms.
SPEED	Low = Speed is 100 Mbps High = Speed is 10 Mbps
LINK	Low = Link up High = Link down
TX	Low = Transmit activity High = No transmit activity This status has a default pulse stretch duration of 170 ms.
RX	Low = Receive activity High = No receive activity This status has a default pulse stretch duration of 170 ms.
ACT	Low = Transmit or received activity High = No transmit or receive activity This status has a default pulse stretch duration of 170 ms.
LINK/RX	Low = Link up High = Link down Blink = receive activity (blink rate is 84 ms active then 84 ms inactive) The receive activity has a pulse stretch duration of 84 ms.
LINK/ACT	Low = Link up High = Link down Blink = Transmit or receive activity (blink rate is 84 ms active then 84 ms inactive) The transmit and receive activity has a pulse stretch duration of 84 ms.

Table 55: Dual LED Display Mode

Event	Status
COLX100	1 = No 100 Mbps collision activity 0 = 100 Mbps collision activity This status has a default pulse stretch duration of 170 ms.
COLX10	1 = No 10 Mbps collision activity 0 = 10 Mbps collision activity This status has a default pulse stretch duration of 170 ms.
ERROR100	1 = None of the above occurred 0 = Received error, false carrier, or 100 Mbps FIFO over/underflow occurred This status has a default pulse stretch duration of 170 ms.
ERROR10	1 = None of the above occurred 0 = Jabber or 10 Mbps FIFO over/underflow occurred This status has a default pulse stretch duration of 170 ms.
DUPLEX100	1 = 100 Mbps half-duplex 0 = 100 Mbps full-duplex
DUPLEX10	1 = 10 Mbps half-duplex 0 = 10 Mbps full-duplex
DUPLEX/COLX100	1 = 100 Mbps half-duplex 0 = 100 Mbps full-duplex Blink = 100 Mbps collision activity The 100 Mbps collision activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.
DUPLEX/COLX10	1 = 10 Mbps half-duplex 0 = 10 Mbps full-duplex Blink = 10 Mbps collision activity The 10 Mbps collision activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.
SPEED100	1 = Speed is 10 Mbps 0 = Speed is 100 Mbps
SPEED10	1 = Speed is 100 Mbps 0 = Speed is 10 Mbps
LINK100	1 = 100 Mbps link down 0 = 100 Mbps link up
LINK10	1 = 10 Mbps link down 0 = 10 Mbps link up
TX100	1 = No 100 Mbps transmit activity 0 = 100 Mbps transmit activity This status has a default pulse stretch duration of 170 ms.

Table 55: Dual LED Display Mode

Event	Status
TX10	1 = No 10 Mbps transmit activity 0 = 10 Mbps transmit activity This status has a default pulse stretch duration of 170 ms.
RX100	1 = No 100 Mbps receive activity 0 = 100 Mbps receive activity This status has a default pulse stretch duration of 170 ms.
RX10	1 = No 10 Mbps receive activity 0 = 10 Mbps receive activity This status has a default pulse stretch duration of 170 ms.
ACT100	1 = No 100 Mbps transmit or 100 Mbps receive activity 0 = 100 Mbps transmit or 100 Mbps receive activity This status has a default pulse stretch duration of 170 ms.
ACT10	1 = No 10 Mbps transmit or 10 Mbps receive activity 0 = 10 Mbps transmit or 10 Mbps receive activity This status has a default pulse stretch duration of 170 ms.
LINK/RX100	1 = 100 Mbps link down 0 = 100 Mbps link up Blink = 100 Mbps receive activity The 100 Mbps receive activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.
LINK/RX10	1 = 10 Mbps link down 0 = 10 Mbps link up Blink = 10 Mbps receive activity The 10 Mbps receive activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.
LINK/ACT100	1 = 100 Mbps link down 0 = 100 Mbps link up Blink = 100 Mbps transmit or 100 Mbps receive activity The 100 Mbps receive or transmit activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.
LINK/ACT10	1 = 10 Mbps link down 0 = 10 Mbps link up Blink = 10 Mbps transmit or 10 Mbps receive activity This 10 Mbps receive or transmit activity has a default pulse stretch duration of 84 ms. The blink rate can be programmed.

Section 5. Serial Management Interface (SMI)

The device's serial management interface provides access to the internal registers via the MDC and MDIO signals and is compliant with IEEE 802.3u clause 22. MDC is the management data clock input whose frequency can run from DC to a maximum rate of 8.3 MHz. MDIO is the management data input/output which carries a bidirectional signal that runs synchronously with the MDC. The MDIO pin requires a pull-up resistor to pull the MDIO high during idle and turnaround times.

5.1 MDC/MDIO Read and Write Operations

All of the relevant serial management registers, as well as several optional registers, are implemented in the device's Switch Core. A description of these registers can be found in Section 6. "Switch Register Description" .

Figure 40: Typical MDC/MDIO Read Operation

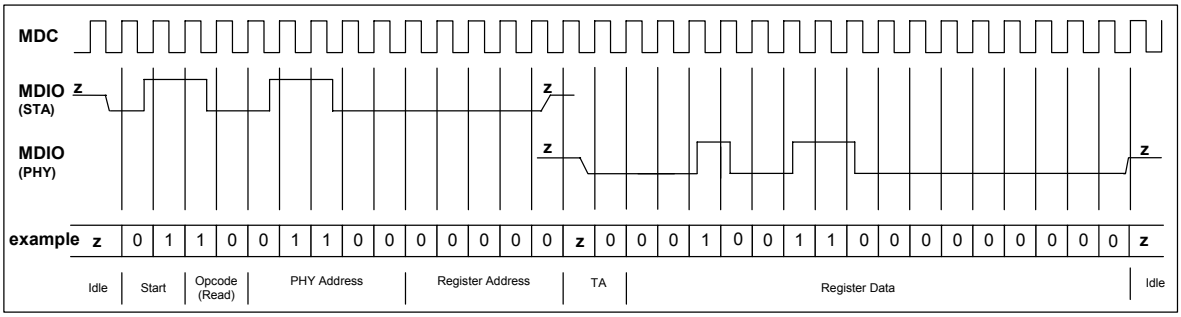

Figure 41: Typical MDC/MDIO Write Operation

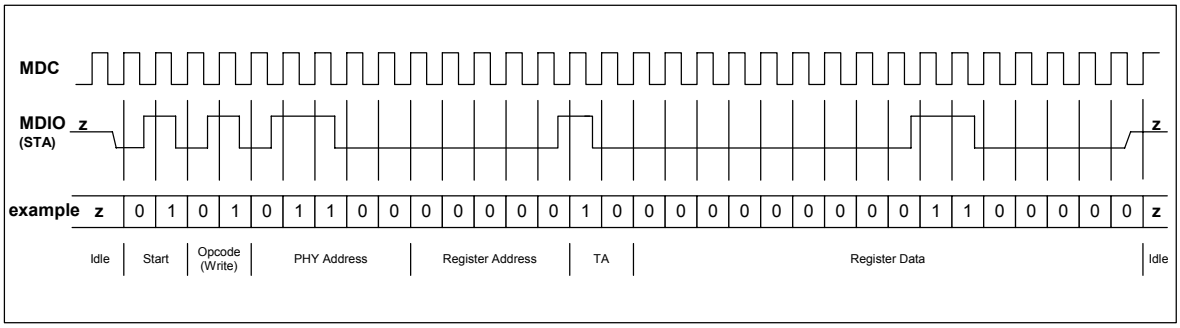

Table 56 shows an example of a read operation of PHY address 04, register 0, with data of 04C0.

Table 56: Serial Management Interface Protocol Example

32-Bit Preamble	Start of Frame	Opcode Read = 10 Write = 01	5-Bit Phy Device Address	5-Bit Phy Register Address	2-Bit Turnaround Read = z0 Write = 10	16-Bit Data Field	Idle
11111111	01	10	00100	00000	z0	0000010011000000	11111111

Section 6. Switch Register Description

The devices registers are accessible using the MDC_CPU and MDIO_CPU pins which support the IEEE Serial Management Interface (SMI – Clause 22) used for PHY devices (Refer to MDC/MDIO in Table 13). The devices support two kinds of SMI address usage models. One uses 1 of the 32 possible Device addresses (when in Multi-chip mode - 6.2 "Multi-chip Addressing Mode"). The other uses all of the 32 possible Device addresses (when in Single-chip mode - 6.3 "Single-chip Addressing Mode"). The device addresses and modes used are configurable at reset with the ADDR[4:0] configuration pins (Refer to "GMII/MII Transmit Interface," Table 10).

Table 57: Register Map—Multi-Chip Addressing Mode

Address	Description	Page #
00	“SMI Command Register”	page 169
01	“SMI Data Register”	page 170

Table 58: Register Map—Single-Chip Addressing Mode

Address	Description	Page #
Switch Per-port Registers		
00	“Port Status Register”	page 174
01	“PCS Control Register”	page 178
02	Reserved	
03	“Switch Identifier Register”	page 179
04	“Port Control Register”	page 180
05	Reserved	
06	“Port Based VLAN Map”	page 184
07	“Default Port VLAN ID & Priority,”	page 185
08	“Port Control 2 Register”	page 186
09	“Rate Control”	page 189
10	“Rate Control 2”	page 190
11	“Port Association Vector”	page 191
12-15	Reserved	
16	“InDiscards Low Counter”	page 192
17	“InDiscards High Counter”	page 192
18	“InFiltered Counter”	page 192
19	“OutFiltered Counter”	page 193

Table 58: Register Map—Single-Chip Addressing Mode (Continued)

Address	Description	Page #
20–23	Reserved	
24	"Port IEEE Priority Remapping Registers (88E6095 and 88E6095F devices only)"	page 194
25	"Port IEEE Priority Remapping Registers (88E6095 and 88E6095F devices only)"	page 195
26–31	Reserved	
Switch Global Registers		
00	"Switch Global Status Register"	page 197
01	"Switch MAC Address Register Bytes 0 & 1"	page 199
02	"Switch MAC Address Register Bytes 2 & 3,"	page 199
03	"Switch MAC Address Register Bytes 4 & 5"	page 199
04	"Switch Global Control Register"	page 200
05	"VTU Operation Register"	page 202
06	"VTU VID Register"	page 204
07	"VTU Data Register Ports 0 to 3"	page 205
08	"VTU Data Register Ports 4 to 7"	page 206
09	"VTU Data Register Port 8 to 10"	page 207
10	"ATU Control Register"	page 208
11	"ATU Operation Register"	page 209
12	"ATU Data Register"	page 211
13	"ATU MAC Address Register Bytes 0 & 1"	page 212
14	"ATU MAC Address Register Bytes 2 & 3"	page 212
15	"ATU MAC Address Register Bytes 4 & 5"	page 212
16	"IP-PRI Mapping Register 0"	page 213
17	"IP-PRI Mapping Register 1"	page 213
18	"IP-PRI Mapping Register 2"	page 214
19	"IP-PRI Mapping Register 3"	page 214
20	"IP-PRI Mapping Register 4"	page 215
21	"IP-PRI Mapping Register 5"	page 215
22	"IP-PRI Mapping Register 6"	page 216
23	"IP-PRI Mapping Register 7"	page 216
24	"IEEE-PRI Register"	page 217

Table 58: Register Map—Single-Chip Addressing Mode (Continued)

Address	Description	Page #
25	"Core Tag Type"	page 218
26	"Monitor Control"	page 219
27	Reserved	
28	"Global Control 2"	page 220
29	"Stats Operation Register"	page 221
30	"Stats Counter Register Bytes 3 & 2"	page 223
31	"Stats Counter Register Bytes 1 & 0"	page 223
Switch Global 2 Registers		
01 - 02	Reserved	
03	"MGMT Enable Register"	page 225
04	"Flow Control Delay Register"	page 225
05	"Switch Management Register"	page 226
06	"Device Mapping Table Register"	page 228
07	"Trunk Mask Table Register"	page 229
08	"Trunk Members Table Register"	page 230
09 - 31	Reserved	

6.1 Register Types

The registers in the devices are made up of one or more fields. The way in which each of these fields operate is defined by the field's Type. The function of each Type is described below.

Type	Description
LH	Register field with latching high function. If status is high, then the register bit is set to one and remains set until a read operation is performed through the management interface or a reset occurs.
LL	Register field with latching low function. If status is low, then the register bit is cleared to zero and remains zero until a read operation is performed through the management interface or a reset occurs.
RES	Reserved for future use. All reserved bits are read as zero unless otherwise noted.
RO	Read only.
R/W	Read and write with initial value indicated.
RWR	Read/Write reset. All field bits are readable and writable. After reset, register field is cleared to zero.
RWS	Read/Write set. All field bits are readable and writable. After reset, register field is set to a non-zero value specified in the text.
SC	Self-Clear. Writing a one to this register causes the desired function to be immediately executed, then the register field is automatically cleared to zero when the function is complete.
Update	Value written to the register field does not take effect until soft reset is executed; however the written value can be read even before the software reset.
Retain	Value written to the register does not take effect without a software reset and the computer maintains its value after a software reset.
WO	Write only. Reads to this type of register field return undefined data.

6.2 Multi-chip Addressing Mode

When Multi-chip addressing mode is used, the devices respond to only 1 of the 32 possible SMI device addresses so that it can share the SMI interface with multiple devices. In this mode, only two of the devices' registers are directly accessible, the SMI Command register (Table 59) and the SMI Data register (Table 60). These two registers are used to access all the other device registers indirectly (along with any PHY registers that may be attached to it).

Indirect accessing of the other devices registers is accomplished by setting the SMI Command register's DevAddr and RegAddr bits to point to the devices register to access. Use the DevAddr and RegAddr values defined for devices in the Single-chip Addressing mode (Section 6.3).

Multi-chip Addressing Mode is enabled when the ADDR[4:0] configuration pins (Table 10) carry a non-zero value at the rising edge of RESETn. The ADDR[4:0] configuration pins also define the single SMI address to which this device will respond to. To avoid conflicts, this requires that all devices on the same SMI interface use unique ADDR[4:0] values. An SMI address of 0x00 is not supported in this mode as this value on the ADDR[4:0] pins places the devices into Single-chip Addressing mode (Section 6.3).

Table 59: SMI Command Register[1]
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
15	SMIBusy	SC	Internal SMI Unit Busy. This bit must be set to a one to start an internal SMI operation (see SMIOp below). Only one SMI operation can be executing at one time so this bit must be zero before setting it to a one. When the requested SMI operation is completed, this bit is automatically be cleared to a zero.
14:13	Reserved	RES	Reserved for Future Use
12	SMIMode	RWR	Internal SMI Mode bit. This bit is used to define the SMI frame type to generate as follows: 0 = Generate IEEE 802.3 Clause 45 SMI frames[2] 1 = Generate IEEE 802.3 Clause 22 SMI frames[3]
11:10	SMIOp	RWR	Internal SMI Opcode. These bits are used to select the SMI opcode to operate on during SMI commands as follows: When the SMIMode bit = 1 then SMIOp = (IEEE 802.3 Clause 22): 11 = Reserved 10 = Read Data Register 01 = Write Data Register 00 = Reserved When the SMIMode bit = 0 then SMIOp = (IEEE 802.3 Clause 45): 11 = Read Data Register with post increment on the address register 10 = Read Data Register 01 = Write Data Register 00 = Write Address Register
9:5	DevAddr	RWR	Internal SMI Device Address bits. These bits are used to select the SMI device (Clause 22) or port (Clause 45) to operate on during SMI commands.
4:0	RegAddr	RWR	Internal SMI Register Address bits. These bits are used to select the SMI register (Clause 22) or device class (Clause 45) to operate on during SMI commands.

1. This register is accessible only when the device is in Multi-chip Addressing mode.
2. Clause 45 SMI frames can be used to access Clause 45 devices connected to the MDC_PHY and MDIO_PHY pins
3. Clause 22 SMI frames must be used to access the internal switch registers

Table 60: SMI Data Register[1]
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
15:0	SMIData	RWR	SMI Data register. During SMI writes these bits must be written with the SMI data to be written prior to starting the SMI operation (i.e., before setting SMIBusy to a one). During SMI reads these bits will contain the SMI data that was read after the SMI read operation is completed (i.e., SMIBusy returns to a zero). Writes to this register must not be done while SMIBusy is a one.

1. This register is accessible only when the device is in Multi-chip Addressing mode.

6.3 Single-chip Addressing Mode

Figure 42 shows the register map in Single-chip mode. In this mode, the devices respond to all 32 SMI device addresses so it must be the only device connected to a SMI Master (typically a CPU). The devices use 20 SMI device addresses internally (0x00 to 0x07 and 0x10 to 0x1B) and all unused SMI device addresses are ignored if the PHY Polling Unit (PPU) is enabled (Switch Global Control register, (Table 83). If the PPU is disabled, all unused SMI device addresses pass through the devices through the MDC_PHY and MDIO_PHY pins enabling the CPU to have direct access to any PHY (or other device) that may be attached to MDC_PHY and MDIO_PHY pins. If PHYs are attached to the devices Ports 8, 9, or 10, they must use SMI device addresses 0x08 (for Port 8) to 0x0A (for Port 10) so that the PPU can automatically poll the PHYs for link, speed, duplex and flow control status, and communicate the current PHY state to the correct MAC.

Single-chip Addressing Mode is enabled when the ADDR[4:0] configuration pins (Table 10) are all zeroes at the rising edge of RESETn.

Figure 42: Device Register Map

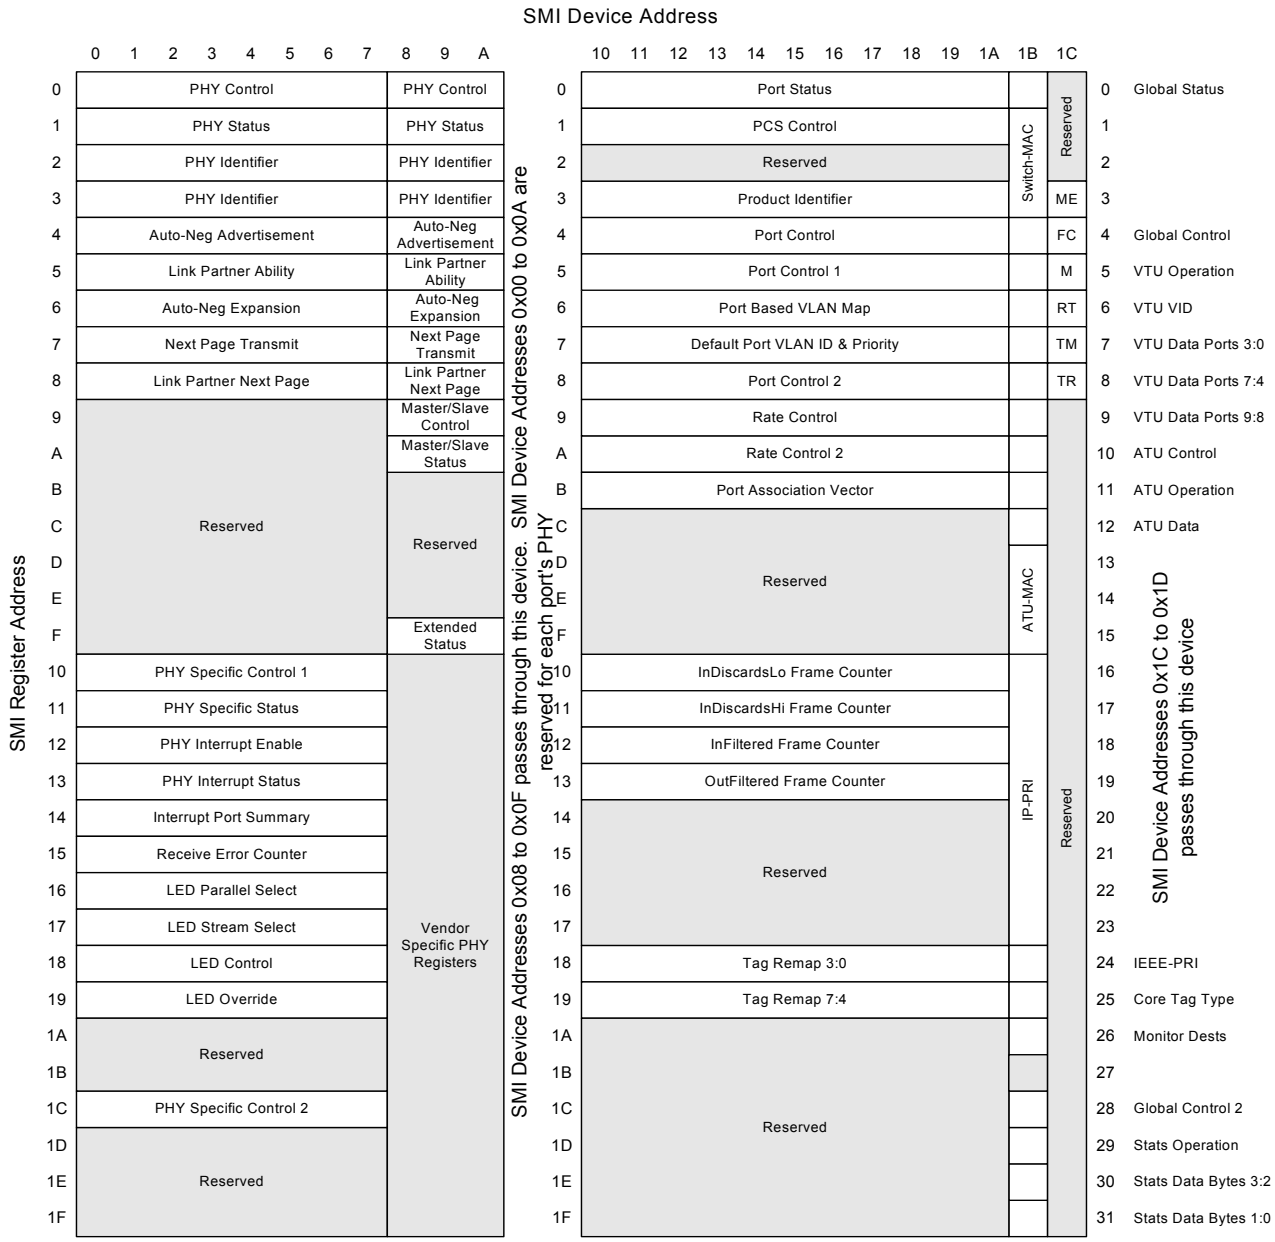

6.4 Switch Port Registers

Each Ethernet port in the devices contain their own per port registers. Each per port register is 16-bits wide and their bit assignments are shown in Figure 42.

Figure 43: Per Port Register Bit Map

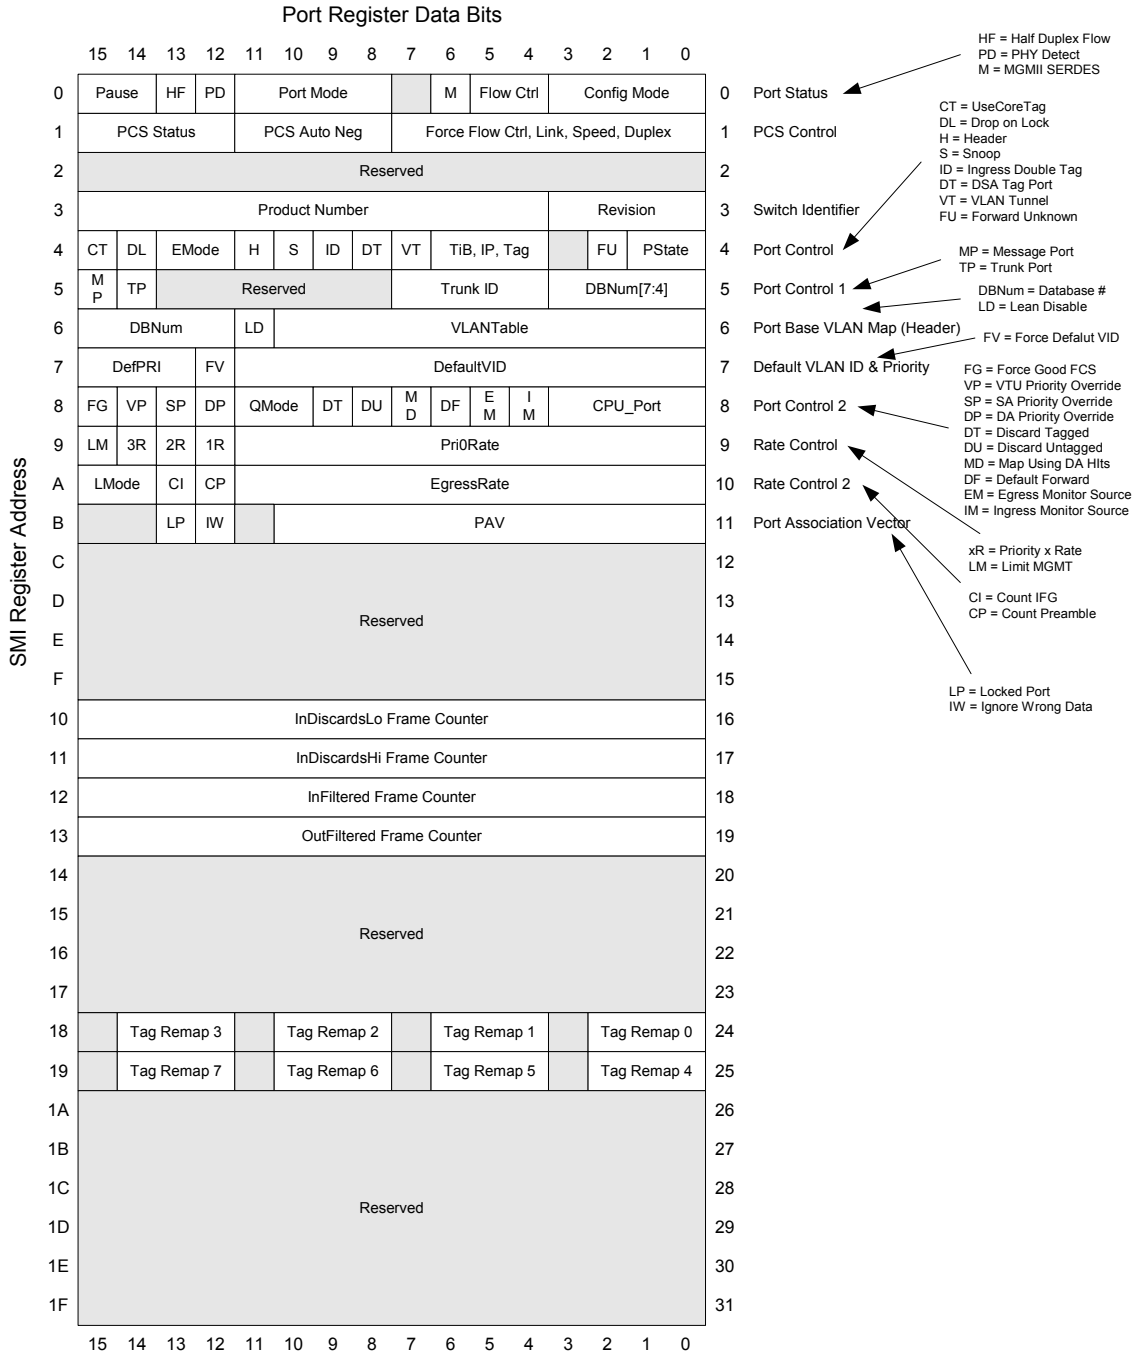

Table 61: Port Status Register[1]
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
15	PauseEn	RO	Pause Enabled bit. This indicates that full-duplex flow control will be used on this port if the port is in full-duplex mode. It is valid when the Link bit is a one. This bit reflects the Pause result from auto-negotiation only (i.e., the ForceFlowControl bit's value is not reflected in this bit). 0 = MAC Pause not implemented in the link partner or in MyPause 1 = MAC Pause is implemented in the link partner & in MyPause
14	MyPause	RWS or RWR	My Pause bit. This bit is sent to the PHY during PHY Polling Unit (PPU) initialization. This bit is not meaningful on ports that do not support Auto-Negotiation (i.e., internal ports). It is set high if FD_FLOW_DIS (Table 14) is low during RESETn. 0 = MAC Pause is not to be advertised as supported in this port 1 = MAC Pause is to be advertised as supported in this port
13	HdFlow	RO	Half-duplex Flow Control. This bit generally reflects the inverted value of the HD_FLOW_DIS bit during reset as follows: 0 = Half-duplex back pressure is not used on this port (unless the port's ForceFlowControl bit is set). 1 = Half-duplex back pressure is used on this port if this port is in a half-duplex mode.
12	PHYDetect	RWR	802.3 PHY Detected. This bits is set to a one if the PPU is enabled and it finds a non-all-one's value in either SMI registers 2 or 3 at the SMI device address 0x10 less than this address. 0 = An 802.3 PHY is not attached to this port 1 = An 802.3 PHY is attached to this port These bits are set at the end of the PPU's init routine (that is why this register must not be written to while the PPUState is Initializing – i.e., 01). The PPU poll routine uses these bits to determine which port's PHY to poll for Link, Duplex, Speed and Flow Control. If software changes these bits, the PPU changes its polling accordingly. A SWReset (Table 83) restarts the PPU's initialization and causes a re-write to this register if the PPU is enabled (also in Table 83).
11	Link	RO	Link Status. This bit indicates the link status of the port as follows: 0 = Link is down 1 = Link is up The port's Link bit mirrors the LinkValue bit if the link is being forced by ForcedLink being a one (Table 63). Otherwise the port's Link bit takes the value of the source defined in Table 62.

Table 61: Port Status Register[1] (Continued)
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
10	Duplex	RO	Duplex mode. This bit is valid when the Link bit, above, is set to a one. 0 = Half-duplex 1 = Full-duplex The port's Duplex bit takes the value of the DpxValue bit if the duplex is being forced by ForcedDpx being a one (Table 63). Otherwise the port's Duplex bit takes the value of the source defined in Table 62.
9:8	Speed	RO	Speed mode. These bits are valid when the Link bit, above, is set to a one. When the port is configured in (G)MII mode the Speed bits are determined by the Px_MODE[2:0] bits. 00 = 10 Mbps 01 = 100 Mbps 10 = 1000 Mbps 11 = Reserved for future use The port's Speed bits take the value of ForceSpeed bits if the speed is being forced by ForceSpeed being non-0x3 (Table 63). Otherwise the port's Speed bits take the value of source defined in Table 62.
7	Reserved	RES	Reserved for future use.
6	MGMII (valid on Ports 8 to 10 only)	RWR or RWS[2]	SERDES Interface mode. When this bit is cleared to a zero and a PHY is detected connected to this port, the SERDES interface between this port and the PHY is SGMII. When this bit is set to a one and a PHY is connected to this port, the SERDES interface between this port and the PHY is MGMII. When this bit is set to a one and no PHY is detected on this port and the SERDES interface is being used, it is configured in 1000BASE-X mode. MGMII mode is recommended when connecting to Marvell® PHY devices that support auto-media detect when connecting to SERDES ports. MGMII mode uses out of band communication for link, speed, and flow control status, and works specifically with Marvell PHYs that support MGMII.
5	TxPaused	RO	Transmitter Paused. This bit is set to a one whenever the Rx MAC receives a Pause frame with a non-zero Pause time that is used by the Tx MAC (i.e., the transmitter is paused off). If the port is in half-duplex mode, this bit is never a one since all Rx Pause frames are ignored. If the port is in full-duplex mode this bit is never a one if Rx Pause frames are ignored because flow control is disabled on this port.
4	FlowCtrl	RO	Flow Control. This bit is set to a one whenever the Rx MAC determines that no more data should be entering this port. If the port is in half-duplex mode, this bit's being a one indicates that the port is using back pressure. If the port is in full-duplex mode this bit's being a one indicates that the port is going to or has sent a Pause frame with a non-zero Pause time to its link partner.

Table 61: Port Status Register[1] (Continued)
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
3	C_Duplex (valid on Port 9, Port 10 only)	RO	Config Duplex. This bit returns the port's duplex configuration mode determined at reset as follows: 0 = Half-duplex operation 1 = Full-duplex operation This bit is valid on port 9 and port 10 only and it reflects the value of P9_HALFDPX or P10_HALFDPX after reset regardless of the value of the P9_MODE[2:0] or P10_MODE[2:0] bits respectively after reset (if this bit is a one the port is not in half-duplex mode unless C_MODE = 01X).
2:0	C_Mode	RO	Config Mode. These bits return the port's interface type configuration mode determined by the value of the Px_MODE[2:0] pins at reset as follows: 000 = FD GMII with Px_GTXCLK = 125 MHz (1000BASE – Port 9 or Port 10 only) 001 = FD MII with Px_GTXCLK = 0 MHz (Power Save – Port 9 or Port 10 only) 010 = FD or HD MII with Px_GTXCLK = 25 MHz (100BASE – Port 9 or Port 10 only) 011 = FD or HD MII with Px_GTXCLK = 2.5 MHz (10BASE – Port9 or Port 10 only) 100 = FD cross-chip SERDES port 101 = FD 1000BASE-X Port (Port 8, Port 9, or Port 10 only) 110 = PHY Port (duplex and speed determined by the PPU) 111 = Port is disabled (Port 9 or Port 10 only) **NOTE:** For the 88E6095 device GMII mode is not supported on Port 9.

1. This PortStatus register must not be written to if the PPUState in the Global Status register (Table 79) is 01.
2. The power on reset state of this bit on all ports is determined by the MGMII configuration pin.

The source of the port's Mode, Link, Speed and Duplex bits (assuming no forcing of the bits is occurring—Table 63) is defined in Table 62. Each of these bits, Link, Speed and Duplex, can be individually forced to any value by using the ForceXXX bits in Table 63.

Table 62: Port Configuration Matrix

Port #	Px_MODE[1]	PHY Detect[2]	C_Mode[3]	Port's Mode[4]	Link[5]	Speed[6]	Duplex[7]
0 to 7	No Pin	0	100	Disabled	0	100	FD
		1	110	AutoNeg	PPU	PPU	PPU
8	0	0	100	Cross-chip	0	1000	FD
		1	110	AutoNeg	PPU	PPU	PPU
	1	X	101	1000BASE	PCS	1000	FD
9 or 10	0x0[8]	0	000	GMII	0	1000	FD
		1	000	GMII	PPU	PPU	PPU
	0x1	0	001	MII	0	100	FD
		1	001	MII	PPU	PPU	PPU
	0x2	0	010	MII	0	100	P9_HALFDPX
		1	010	MII	PPU	PPU	PPU
	0x3	0	011	MII	0	10	P9_HALFDPX
		1	011	MII	PPU	PPU	PPU
	0x4	X	100	Cross-chip	0	1000	FD
	0x5	X	101	1000BASE	PCS	1000	FD
	0x6	X	110	PHY	PPU	PPU	PPU
	0x7	X	111	Disabled	0	1000	FD

1. Px_MODE is the single configuration pin for ports 8 or it is the three configuration pins for Port 9 and Port 10.
2. PHYDetect is the value of the PHYDetect bit for the port as seen in the port's Port Status register (Table 61).
3. C_Mode is the value that is seen in the C_Mode bits for this port (in the port's Port Status register - Table 61).
4. Port's Mode is the port's mode in name form as defined by the port's C_Mode bits.
5. Link is the source of the port's Link bit assuming Link is *not* being forced (Table 63). PPU means this bit comes from the PHY Polling Unit, PCS means this bit comes from the 1000BASE-X PCS block. Link bits that are '0' can be forced to a one by software (see the port's PCS Control register – Table 65). A port's Link bit as stored in the port's PPU Link register is cleared to zero whenever the port's PHYDetect transitions from a 1 to a 0. This insures that the PPU's link for this port does not start out as a one, but transitions to a one only if the port's PHY still sees a link.
6. Speed is the source of the port's Speed bits assuming Speed is *not* being forced (Table 63). PPU means these bits comes from the PHY Polling Unit. 10, 100, or 1000 means the port defaults to this speed.
7. Duplex is the source of the port's Duplex bit assuming Duplex is *not* being forced (Table 63). PPU means this bit comes from the PHY Polling Unit. P9_HALFDPX or P10_HALFDPX means the bit comes from inverting the value seen on the P9_HALFDPX or P10_HALFDPX configuration pin respectively. FD means this mode defaults to full-duplex.
8. Px_MODE of 0x0 is NOT supported on Port 9 in the 176-pin LQFP package and must not be selected.

Table 63: PCS Control Register
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
15[1]	PCSLink (Ports 8, 9 & 10 only)	RO	PCS Link up status. This bit b is a one whenever the PCS link is up.
14[1]	SyncOK (Ports 8, 9 & 10 only)	RO	This bit is set to a one when the PCS has detected a few comma patterns and is synchronized with its peer PCS layer.
13[1]	SyncFail (Ports 8, 9 & 10 only)	RO	This bit is set to a one if the PCS sees link. It will be cleared to a zero if the PCS has lost sync for 10 ms or more. Refer to IEEE 802.3 Clause 36.
12[1]	AnBypassed (Ports 8, 9 & 10 only)	RO	This bit indicates that Inband Auto-Negotiation was bypassed. If there is no reply during Auto-Negotiation, Bypass is activated and Link is set to 'up'. Speed is set to 1000 Mbps Speed and Duplex is set to Full-duplex. 1 = Auto-Negotiation is bypassed 0 = Auto-Negotiation is not bypassed
11[1]	AnBypass (Ports 8, 9 & 10 only)	RWS	In Band Auto-Negotiation Bypass Enable. This bit is relevant when Auto-Negotiation is done inband and not via the SMI interface (i.e., when InBandAn is set to a one). When this bit is set to a one and the link partner does not respond to the Auto-Negotiation process, the link is established by bypassing the Auto-Negotiation procedure (see AnBypassed bit). When this bit is cleared to a zero Auto-Negotiation cannot be bypassed.
10[1]	PCSAnEn (Ports 8, 9 & 10 only)	RWR or RWS	PCS Inband Auto-Negotiation Enable. PCS Inband Auto-Negotiation is done if the port is in 1000BASE-X for link and flow control support. This bit will be set to a one after reset mode if this port is in 1000BASE-X mode. When this port is in any other mode inband Auto-Negotiation is not done so this bit will be cleared to a zero after reset.
9[1]	RestartPCSAn (Ports 8, 9 & 10 only)	SC	Restart PCS Inband Auto-Negotiation. This bit is relevant when PCS Auto-Negotiation is done inband and not via the SMI interface (i.e., when InBandAn is set to a one). When this bit is set to a one the inband Auto-Negotiation restarts. This bit is reset to zero by the device immediately after restarting the Auto-Negotiation process.
8[1]	PCSAnDone (Ports 8, 9 & 10 only)	RO	This bit is cleared during the PSC Auto-Negotiation phase. It is set to one when PCS Auto-Negotiation is done (or if it were never done).
7	FCValue	RWR	Flow Control's Forced value. This bit is used to force flow control (if full-duplex) or backpressure (if half-duplex) to be enabled when the ForcedFC bit (below) is set to a one. Flow control/back pressure is forced enabled when this bit is set to a one. It is forced disabled when this bit is cleared to a zero. If the ForcedFC bit (below) is cleared to a zero, this bit has no effect.
6	ForcedFC		Force Flow Control. When this bit is set to a one flow control (if full-duplex) or backpressure (if half-duplex) for this port is forced to the value in the FCValue register (above) regardless of what the normal flow control value would be. In this way, flow control/backpressure can be forced to be enabled or disabled. When this bit is cleared to a zero, normal flow control detection occurs.

Table 63: PCS Control Register (Continued)
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
5	LinkValue	RWR	Link's Forced value. This bit is used to force the link up or down when the ForcedLink bit (below) is set to a one. The link will be forced up when this bit is set to a one. It will be forced down when this bit is cleared to a zero. If the ForcedLink bit (below) is cleared to a zero this bit has no effect.
4	ForcedLink	RWR	Force Link. When this bit is set to a one the link for this port is forced to the value in the LinkValue register (above) regardless of what the normal link's value would be. In this way, the link can be forced to be up or down. When this bit is cleared to a zero, normal link detection occurs.
3	DpxValue	RWR	Duplex's Forced value. This bit is used to force the link to full- or half-duplex mode, when the ForcedDpx bit (below) is set to a one. The link duplex is forced to full when this bit is set to a one. It will be forced to half when this bit is cleared to a zero (Do not try to force half-duplex mode in a 1000BASE link - it is not supported and results are unpredictable). If the ForcedDpx bit (below) is cleared to a zero this bit has no effect.
2	ForcedDpx	RWR	Force Duplex. When this bit is set to a one the duplex for this port will be forced to the value in the DpxValue register (above) regardless of what the normal duplex's value would be. In this way the duplex can be forced to be full or half. When this bit is cleared to a zero, normal duplex detection occurs.
1:0	ForceSpd	RWS to 0x3	Force Speed. These bits are used to force the speed on this port as follows: 00 = 10 Mbps 01 = 100 Mbps 10 = 1000 Mbps 11 = Speed is not forced. Normal speed detection occurs.

1. Bits 15:8 are irrelevant for the (G)MII mode.

Warning

The duplex and speed on a port must not be changed unless the link on the port is down.

Table 64: Switch Identifier Register
Offset: 0x03 or decimal 3

Bits	Field	Type	Description
15:4	Product Num	RO	Product Number or identifier: 88E6092 = 0x09788E6095 and 88E6095F = 0x095
3:0	Rev	RO	Revision Identifier. The initial version of the devices has a Rev of 0x0. This Rev field may change at any time. Contact Marvell FAEs for current information on the device revision identifier.

Table 65: Port Control Register
Offset: 0x04 or decimal 4

Bits	Field	Type	Description
15	UseCoreTag	RWR	Use Core Tag. When this bit is cleared to a zero, ingressing frames are considered tagged if the 16-bits following the frame's Source Address are 0x8100. When this bit is set to a one, ingressing frames are considered tagged if the 16-bits following the frame's Source Address are equal to the CoreTag register value (Table 104). NOTE: If a frame is Double Tagged and the first tag is removed (Double Tag bit below) by matching the EtherType defined by this bit, the resulting frame will be considered tagged only if the next 16 bits in the frame equal 0x8100 (that is, the CoreTag register is NOT used after the first tag is removed).
14	DropOnLock	RWR	Drop On Lock. This bit is used to enable 802.1X MAC based authentication discards on this port. 1 = Ingressing frames are discarded if their SA field is not in the ATU's address database (i.e., its a new or unknown Source Address) or if this port's bit is not set in the PortVec bits for the frame's SA (i.e., this port is not the source port for that MAC address). 0 = Normal mode – ingressing frames are not discarded due to their SA field.
13:12	Egress Mode	RWR	Egress Mode. The first three Egress Modes are used as the default Egress mode for frames whose VID[1] is not valid in the VTU (Table 84) or if 802.1Q is disabled on the port. The fourth Egress Mode is applied to all frames all the time (if selected) even if 802.1Q is enabled on the port. 00 = Default to Normal mode – frames are transmitted unmodified[2] 01 = Default to Transmit all frames Untagged – remove the tag from any tagged frame 10 = Default to Transmit all frames Tagged – add a tag to any untagged frame 11 = Always add a Tag (or Egress Double Tag). This setting is *not* a default setting. It ignores the MemberTag data in the VTU for this port. The Ether Type used for these inserted tags come from the Global CoreTagType register (Table 104).
11	Header	RWR	Ingress & Egress Header Mode. When this bit is set to a one all frames egressing this port are pre-pended with the Marvell® 2-byte Egress Header just before the frame's DA field. Also, all frames ingressing this port are expected to be pre pended with the Mavell 2-byte Ingress Header just before the frame's DA field. On Ingress the 1st 2 bytes after the SFD are removed from the frame and the frame's CRC and size is recomputed. If the frame's Ingress Header is non-zero it is used to update the port's VLAN Map register value (at register offset 0x06). When this bit is cleared to a zero, normal Ethernet frames egress the switch and are expected to ingress the switch.

Table 65: Port Control Register (Continued)
Offset: 0x04 or decimal 4

Bits	Field	Type	Description
10	IGMP/MLD Snoop	RWR	IGMP and MLD Snooping. When this bit is set to a one and this port receives an IPv4 IGMP frame or an IPv6 MLD frame, the frame is switched to the CPU port[3] overriding the destination ports determined by the DA mapping[4]. When this bit is cleared to a zero IGMP/MLD frames are not treated specially.
9	DoubleTag (88E6095 and 88E6095F Only)	RWR	Ingress Double Tag Mode. When this bit is set to a one the ingressing frames are examined to see if they contain 802.3ac tag with the IEEE standard Ether Type of 0x8100) or if they contain a tag with an Ether_Type defined in the CoreTagType register (Table 104). The CoreTagType will be used if this port's UseCoreTag bit (bit 15 above) is a one, otherwise the standard Ether Type is used. If the frames do contain the expected tag, the tag is removed (and the CRC is recomputed). Essentially, untagged frames remain untagged, single tagged frames become untagged and double tagged frames become single tagged. These now single tagged frames will be considered tagged only if its Ether_Type value is 0x8100 no matter what the Ether Type value was found on the removed tag. No padding is performed so a tagged frame must be at least 68 bytes in size or it will be discarded. When this bit is set to a zero, tagged frames are not modified when received. The Tag data that is removed from the frame can be ignored (i.e., the frame is treated as if the removed tag never existed) or the removed Tag data can be used for switching and priority decisions (i.e., the removed Tag's VID and PRI bits are internally assigned to the frame). The action that is taken is controlled by the global UseDoubleTagData bit (Table 112).
8	DSA_Tag	RWR	DSA_Tag. When this bit is set to a one, it indicates this port is a interswitch port used to communicated with a CPU or to cascade with another switch device. Frames egressing this port are modified with the DSA Tag and frames ingressing this port are expected to always contain a DSA Tag. When this bit is cleared to a zero it indicates this port is a normal network port and DSA Tags are not used.
7	VLAN Tunnel	RWR	VLAN Tunnel. When this bit is cleared to a zero the port based VLANs defined in the VLANTable (Table 67), 802.1Q VLANs defined in the VTU (if 802.1Q is enabled - Table 65 and Table 84) and Trunk Masking (Table 114) are enforced for ALL frames. When this bit is set to a one the port based VLANTable masking, 802.1Q VLAN membership masking and the Trunk Masking is bypassed for any frame entering this port with a DA that is currently 'static' in the ATU. This applies to unicast as well as multicast frames.
6	TagIfBoth	RWS	Use IEEE 802.1p tag fields over IP fields. If the current frame is both IEEE 802.3ac tagged, and is an IPv4 or an IPv6 frame, and UseIP and UseTag bits are both enabled, then a priority selection between the two needs to be made. 0 = Use IP fields for priority data when both field types are present 1 = Use Tag fields for priority data when both field types are present

Table 65: Port Control Register (Continued)
Offset: 0x04 or decimal 4

Bits	Field	Type	Description
5	UseIP	RWS	Use IP for Priority. Use IPv4 TOS and/or DiffServ fields and/or IPv6 Traffic Class fields, if present, for priority data. 0 = Ignore IPv4 and IPv6 priority fields 1 = Use IPv4 fields if the frame is IPv4 and use IPv6 fields if the frame is IPv6 for priority data
4	UseTag	RWS	Use IEEE Tags. Use IEEE 802.1p Traffic Class field for priority data if the frame is an IEEE 802.3ac tagged frame. 0 = Ignore IEEE 802.1p tag fields even it the frame is tagged. In this mode, tagged frames that enter this port and egress the switch tagged are re-tagged with this port's default priority bits, if no other priorities are determined for the frame (i.e., the frame's IP priority and/or other priority overrides). This works as a Force Default Priority function on tagged frames. 1 = Use IEEE 802.1p tag Traffic Class for priority data if the frame is tagged.
3	Reserved	RES	Reserved for future use.
2	Forward Unknown	RWS	Forward Unknown. When this bit is set to a one normal switch operation occurs and unicast frames with unknown DA addresses are allowed to egress out this port (assuming the VLAN settings allow the frame to egress this port too). When this bit is cleared to a zero unicast frames with unknown DAs will not egress out of this port.
1:0	PortState	RWR Or RWS to 0b11[5]	Port State. These bits are used to manage a port to determine what kind of frames, if any, are allowed to enter or leave a port for simple bridge loop detection or 802.1D Spanning Tree. The state of these bits can be changed at any time without disrupting frames currently in transit. The Port States are: 00 = Disabled. The switch port is disabled and it will not receive or transmit any frames. The QC returns any pre-allocated ingress queue buffers when the port is in this mode. 01 = Blocking/Listening. The switch examines all frames without learning any SAs, and discards all frames. 10 = Learning. The switch will examine all frames, learning all SA address (except those from MGMT frames), and still discard all but MGMT frames. It will allow MGMT frames only to exit the port. 11 = Forwarding. The switch examines all frames, learning SAs from all good frames (except those from MGMT frames), and receives and transmits all frames as a normal switch.

1. The frame's VID is the VID that is contained in Tagged frames or the default VID assigned to an Untagged frame when it Ingresses into the switch.
2. If this port has 802.1Q disabled and Cross Chip Port Based VLANs are being used in the switch, this port's EgressMode must be Default to Normal mode (to ensure the frames egress the switch looking exactly how they entered the switch) or Always add a Tag (to insure the frames egress the switch with an extra tag compared to how the entered the switch).
3. The CPU port is determined by the CPUPort bits in Port Control 2 Register (Offset 0x08).
4. VLAN masking is performed on IGMP/MLD frames.
5. The PortState bits for all ports come up in the Forwarding state unless the SW_MODE[1:0] pins are set to 0b00, the CPU attached mode, in which case the ports come up Disabled.

Table 66: Port Control 1
Offset: Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
15	Message Port	RWR or RWS[1]	Message Port. When the Learn2All bit in the ATU is set to a one (Table 89) learning message frames will be generated. These frames will be sent out all ports whose Message Port is set to a one. If this feature is used it is recommended that all DSA Tag ports, except for the CPU's port, have their MessagePort bit set to a one. Ports that are not DSA Tag ports (i.e., normal Network ports) should not have their MessagePort bit set to a one.
14	Trunk Port	RWR	Trunk Port. When this bit is set to a one, this port is considered to be a member of a Trunk with the Trunk ID defined below. When this bit is set to a zero, the port is treated as an individual port.
13:8	Reserved	RES	Reserved for future use.
7:4(6:4 for 88E6092)	Trunk ID	RWR	Trunk ID. When the Trunk Port bit (above) is set to a one these bits define which trunk this port is to be associated with. All ports that are members of the same trunk must be assigned the same Trunk ID and each group of ports that form a trunk must be assigned unique Trunk IDs. Note: Bit 7 is used for the 88E6095 and 88E6095F devices, since those devices support up to 16 trunk groups per port. Bit 7 is ignored for the 88E6092 device since the 88E6092 device only supports up to 8 trunk groups per port.
3:0	DBNum [7:4]	RWR	Port's Default Database VLAN Number bits 7:4. This field can be used with non-overlapping VLANs to keep each VLAN's MAC address mapping database separate from the other VLANs. This allows the same MAC address to appear multiple times in the address database (at most one time per VLAN) with a different port mapping per entry. This field is overridden by the DBNum returned from a VTU hit and it should be zero if not used. It must be a unique number for each independent, non-overlapping, VLAN if used. The lower four bits of the port's default DBNum are contained in the Port Base VLAN Map register (Table 67).

1. The Message Port bit on Ports 8 and 9 will be set to a one if the P9_TXD[3]/SW_24P configuration pin is high at the rising edge of RESETn.

Table 67: Port Based VLAN Map
Offset: Offset: 0x06 or Decimal 6[1]

Bits	Field	Type	Description
15:12	DBNum[3:0]	RWR	Port's Default Database VLAN Number bits 3:0. This field can be used with non-overlapping VLANs to keep each VLAN's MAC address mapping database separate from the other VLANs. This allows the same MAC address to appear multiple times in the address database (at most one time per VLAN) with a different port mapping per entry. This field is overridden by the DBNum returned from a VTU hit and it should be zero if not being used. It needs to be a unique number for each independent, non-overlapping, VLAN, if used. The upper 4 bits of the port's default DBNum are contained in the Port Control 1 register (Table 66).
11	Learn Disable	RWR	Learn Disable. When this bit is cleared to a zero automatic learning on this port is controlled by the port's PAV bits (in the Port Association Vector Register at offset 0x0B—Table 72). When this bit is set to a one automatic learning does not occur for this port. This bit performs the same function as clearing the port's PAV bits but this bit is accessible by the CPU's Ingress Header so the CPU can enable and disable learning on a frame by frame basis.
10:0	VLANTable	RWS to all ones except for this port's bit	Port based VLAN Table. The bits in this table are use to restrict which output ports this input port can send frames to. The VLANTable bits are used for all frames even if 802.1Q is enabled on this port or if ProtectedPort is enabled on this port. These bits restrict where a port can send frames to (unless a VLANTunnel frame is being received – Table 65). To send frames to Port 0, bit 0 of this register must be a one. To send frames to Port 1, bit 1 of this register must be a one, etc. After reset, all ports are accessible since all the other port number bits are set to a one. This Port's bit is zero after reset. This prevents frames leaving the port on which they arrived. This Port's bit can to be set to a one in the devices, which allows frames to be switched back to the port on which they arrived. In view of this fact, care should be taken in writing code to manipulate these bits. This register is reset to 0x7FE for Port0 (SMI Device Address 0x10), and it resets to 0x7FD for Port1 (Addr 0x11),to 0x7FB for Port2 (Addr 0x12), etc. However, if the SW_24P configuration pin(on P9_TXD[3]) is set to one(24 Port Mode) Port 8 and 9 come up configured where they cannot communicate with each other (i.e., this register = 0x4FF for both Ports 8 & 9).

1. The contents of this register can be modified on a frame by frame basis if the port's Header Mode is enabled (Section 3.5.17). NOTE: Only the lower four bits of the DBNum field can be modified by the Header. Software that controls the DBNum field by using the Marvell® Header needs to take this into account. The DefaultVIDs used for Cross-Chip Port Based VLANs must be unique from the VIDs used for the 802.1Q VLANs currently active in the switch. Port Based VLAN ports need to have their frame's egress unmodified or the internal VID will be added to the frame if it is set to egress tagged.

Table 68: Default Port VLAN ID & Priority, Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15:13	DefPri	RWR	Default Priority. The bits of this register are used as the default ingress priority to use when no other priority information is available (neither is the frame IEEE Tagged, nor is it an IPv4 nor an IPv6 frame—or the frame is a priority type that is currently disabled (see UseIP and UseTag in Table 65) and no other priority overrides are active on this frame. The DefPri bits are re-mapped by the IEEE-PRI Register (offset 0x18—Table 103) prior to being used.
12	Force DefaultVID	RWR	Force to use Default VID. When 802.1Q is enabled on this port and this bit is set to a one, all ingress frames with IEEE 802.3ac Tags have their VID ignored and the port's DefaultVID below is used and replaced into the frame instead. When this bit is cleared to a zero all IEEE 802.3ac Tagged frames with a non-zero VID use the frame's VID unmodified. When 802.1Q is disabled on this port, this bit has no effect.
11:0	DefaultVID	RWS to 0x001	Default VLAN Identifier. When 802.1Q is enabled on this port the DefaultVID field is used as the IEEE Tagged VID added to untagged or priority tagged frames during egress that ingressed from this port. It is also used as a tagged frame's VID if the frame's VID was 0x000 (i.e., it is a priority tagged frame) or if the port's Force DefaultVID bit (above) is set to a one. When 802.1Q is disabled on this port, the DefaultVID field is assigned to all frames entering the port (if they are tagged or untagged). This assignment is used internal to the switch, so only that Cross Chip Port Based VLANs can be supported.

Table 69: Port Control 2 Register
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
15	IgnoreFCS	RWR	Force good FCS in the frame. When this bit is cleared to a zero frames entering this port must have a good CRC or else they are discarded. When this bit is set to a one the last four bytes of frames received on this port are overwritten with a good CRC and the frames are accepted by the switch (assuming that the frame's length is good and it has a destination).
14	VTUPri Override (88E6095 and 88E6095F only)	RWR	VTU Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one, then VTU priority overrides can occur on this port. A VTU priority override occurs when the determined VID of a frame[1] results in a VID whose VIDPRIOverride bit in the VLAN database is set to a one. When this occurs the VIDPRI value assigned to the frame's VID (in the VLAN database) is used to overwrite the frame's previously determined priority. If the frame egresses tagged the priority in the frame will be this new VIDPRI value. When used the upper two bits of the VIDPRI priority determine the Egress Queue the frame is switched into. The VTU Priority Override has higher priority than the port's Default Priority and the frame's IEEE and/or IP priorities. The priority determined by the frame's VID can be overridden, however, by the frame's SA and/or DA Priority Overrides (see bits 13:12 below).
13	SAPriOverride (88E6095 and 88E6095F only)	RWR	SA Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one, then SA ATU priority overrides can occur on this port. An SA ATU priority override occurs when the source address of a frame results in an ATU hit where the SA's MAC address returns an EntryState that indicates Priority Override. When this occurs the PRI value assigned to the frame's SA (in the ATU database) is used to overwrite the frame's previously determined priority. If the frame egresses tagged the priority in the frame is this new PRI value[2]. When used the two bits of the priority determine the Egress Queue the frame is switched into. The SA ATU Priority Override has higher priority than the port's Default Priority, the frame's IEEE and/or IP priorities, and the VTU Priority Override. The priority determined by the frame's SA can be overridden, however, by the frame's DA Priority Override (see bit 12 below).

Table 69: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
12	DAPriOverride (88E6095 and 88E6095F only)	RWR	DA Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one, then DA ATU priority overrides can occur on this port. A DA ATU priority override occurs when the destination address of a frame results in an ATU hit where the DA's MAC address returns an EntryState that indicates Priority Override. When this occurs, the PRI value assigned to the frame's DA (in the ATU database) is used to overwrite the frame's previously determined priority. If the frame egresses tagged, the priority in the frame will be this new PRI value[3]. When used, the two bits of the priority determine the Egress Queue the frame is switched into. The DA ATU Priority Override has highest priority over the port's Default Priority, the frame's IEEE and/or IP priorities, the VTU Priority Override and the SA Priority Override. **NOTE:** If a multicast's frame DA is contained in the ATU with a MGMT Entry State the frame's priority will be overriddwn regardless of the state of this bit.
11:10	802.1QMode	RWR	IEEE 802.1Q Mode for this port. These bits determine if802.1Q based VLANs are used along with port based VLANs for this Ingress port. It also determines the action to be taken if an 802.1Q VLAN Violation is detected. These bits work as follows: 00 = 802.1Q Disabled. Use Port Based VLANs only. The VLANTable bits and the DefaultVID assigned to the frame during ingress determine which Egress ports this Ingress port is allowed to switch frames to for all frames[4]. 01 = Fallback. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violations and use the VLANTable bits below if the frame's VID is not contained in the VTU (both errors are logged – Table 84). 10 = Check. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violation but discard the frame if its VID is not contained in the VTU (both errors are logged – Table 84). 11 = Secure. Enable 802.1Q for this Ingress port. Discard Ingress Membership violations and discard frames whose VID is not contained in the VTU (both errors are logged – Table 84).
9	Discard Tagged	RWR	Discard Tagged Frames. When this bit is set to a one all non-MGMT frames that are processed as tagged are discarded as they enter this switch port. This test is performed after Ingress Double Tag removal (if enabled). Priority only tagged frames (with a VID of 0x000) are considered tagged.
8	Discard Untagged	RWR	Discard Untagged Frames. When this bit is set to a one all non-MGMT frames that are processed as untagged are discarded as they enter this switch port. This test is performed after Ingress Double Tag removal (if enabled). Priority only tagged frames (with a VID of 0x000) are considered tagged.

Table 69: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
7	MapDA	RWS	Map using DA hits. When this bit is set to a one, normal switch operation occurs where a frame's DA is used to direct the frame out of the correct ports. When this bit is cleared to a zero the frame will be sent out of the ports defined by ForwardUnknown bits (Table 65) or the DefaultForward bits (below) even if the DA is found in the address database. **NOTE:** If a multicast or unicast frame's DA is contained in the ATU with a MGMT Entry State the frame will be mapped out the port(s) defined by the ATU entry (i.e., the setting of the MapDA bit is ignored for MGMT frames).
6	Default Forward	RWS	Default Forward. When this bit is set to a one normal switch operation occurs and multicast frames with unknown DA addresses are allowed to egress out this port (assuming the VLAN settings allow the frame to egress this port too). When this bit is cleared to a zero multicast frames with unknown DAs do not egress from this port. **NOTE:** The Forward Unknown bit (Table 65) performs this operation for unknown unicast frames and these two bits can be used together.
5	Egress Monitor Source	RWR	Egress Monitor Source Port. When this bit is cleared to a zero, normal network switching occurs. When this bit is set to a one any frame that egresses out this port will also sent to the EgressMonitorDest Port (Table 105). The 802.1Q mode and VTU entries on the Egress Monitor Destination Port must be set the same as they are on the Egress Monitor Source port so the frames egress with the same tagged or untagged information.
4	Ingress Monitor Source	RWR	Ingress Monitor Source Port. When this bit is cleared to a zero normal network switching occurs. When this bit is set to a one, any frame that ingresses this port is also sent to the IngressMonitorDest Port (Table 105). The frame is sent to the IngressMonitorDest Port even if it is discarded owing to switching policy (like VLAN membership, etc.) but the frame will not be forwarded if its contains an error (such as CRC, etc.).
3:0	CPUPort	RWR	CPU Port. When Snooping is enabled on this port or when this port is configured as an DSA Tag Port and it receives a To_Cpu frame, the switch needs to know what port on this device the frame should egress (i.e., where it should be sent to). The CPUPort bits indicate the port number on this device where the CPU is connected (either directly or indirectly through another Marvell switch device). **NOTE:** MGMT or BPDU frames entering a normal or Network port are directed to the correct port where the CPU is connected by ensuring the CPU port's bit is set in the frame's MGMT DA MAC address as stored in the ATU address database (Section 3.4.5).

1. The VID of a frame could be a tagged frame's VID or the port's DefaultVID.
2. The two bits in the address database are expanded to the three IEEE Tap priority bits as follows: 00 = 000 (PRI 0), 01 = 010 (PRI 2), 10 = 101 (PRI 5), 11 = 111 (PRI 7).
3. See footnote above.
4. The VLANTable is sufficient to define Port Based VLANs when only one devices are being used in a system (i.e., the VLANTable works for in-chip port based VLANs). When multiple devices are used in a system the DefaultVID along with the data assigned to it in the VTU supports cross-chip port based VLANs if 802.1Q is disabled on the port (even if 802.1Q is disabled on a port the VID that is assigned to the frame is still locked up into the VTU and the frame is mapped accordingly.)

Table 70: Rate Control
Offset: 0x09 or Decimal 9

Bits	Field	Type	Description
15	Limit MGMT (88E6095 and 88E6095F devices only)	RWR	Limit and count MGMT frame bytes. When this bit is set to a one, MGMT frames are included in Ingress and Egress rate limiting calculations and can be limited. When this bit is cleared to a zero MGMT bytes are not counted and are not limited.
14	Pri3Rate	RWR	Ingress data rate limit for priority 3 frames. Priority 3 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri2Rate 1 = Use twice the rate as Pri2Rate (but not faster than the port's speed)
13	Pri2Rate	RWR	Ingress data rate limit for priority 2 frames. Priority 2 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri1Rate 1 = Use twice the rate as Pri1Rate (but not faster than the port's speed)
12	Pri1Rate	RWR	Ingress data rate limit for priority 1 frames. Priority 1 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri0Rate 1 = Use twice the rate as Pri0Rate (but not faster than the port's speed)
11:0	Pri0Rate	RWR	Ingress data rate limit for priority 0 frames. Priority 0 frames are discarded after the ingress rate selected is reached or exceeded. The devices use the following formula to limit the ingress data rate: Pri0Rate = 8 bits/(32 ns * Ingress Rate bits/sec) Where Ingress Rate is in bits per second. For example for 128 Kbits/sec: Pri0Rate = 8 bits/(32 ns * 128000 bits/sec) = Pri0Rate = 8 bits/(0.004096 bits) ¨(0.000032 * 128) = 1953 or 0x7A1 The value programmed into this register This register supports a range of 256 Mbps to 62 kbps. Ingress rate limiting is not recommended for TCP/IP rate limiting. It is designed for Storm Prevention.

Table 71: Rate Control 2
Offset: 0x0A or Decimal 10

Bits	Field	Type	Description
15:14	Limit Mode	RWR	Ingress Limit Mode. These bits determine what kinds of frames are limited and counted against Ingress limiting as follows: 00 = Limit and count all frames 01 = Limit and count Broadcast, Multicast and flooded unicast frames 10 = Limit and count Broadcast and Multicast frames only 11 = Limit and count Broadcast frames only If a frame is not limited by the above setting, its size is not counted against the limit for the other frames.
13	Count IFG	RWS	Count IFG bytes. When this bit is set to a one, each frame's minimum inter frame gap (IFG) bytes (12 per frame) are included in Ingress and Egress rate limiting calculations. When this bit is cleared to a zero IFG bytes are not counted.
12	Count Pre	RWS	Count Preamble bytes. When this bit is set to a one, each frame's preamble bytes (8 per frame) are included in Ingress and Egress rate limiting calculations. When this bit is cleared to a zero preamble bytes are not counted.
11:0	Egress Rate (88E6095 and 88E6095F devices only)	RWR	Egress data rate limit. The EgressRate bits modify this port's effective transmission rate. The 88E6095 and 88E6095F devices use the following formula to limit the Egress data rate: EgressRate = 8 bits/(32 ns * Egress Rate bits/sec) Where Egress Rate is in bits per second. For example for 256 Kbits/sec: Pri0Rate = 8 bits/(32 ns * 256000 bits/sec) = Pri0Rate = 8 bits/(0.008192 bits) ¨ (0.000032 * 256) = 976 or 0x3D0 - The value programmed into this register This register supports a range of 256 Mbps to 62 kbps.

Table 72: Port Association Vector
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
15:14	Reserved	RES	Reserved for future use.
13	LockedPort	RWR	Locked Port. When this bit is cleared to a zero, normal address learning will occur. When this bit is set to a one CPU directed learning (needed for 802.1X MAC authentication) is enabled on this port. In this mode, an ATU Miss Violation interrupt occurs when a new SA address is received in a frame on this port. Automatically SA learning and refreshing is disabled in this mode.
12	Ignore WrongData	RWR	Ignore Wrong Data. All frame's SA addresses are searched for in the ATU's address database. If the frame's SA address is found in the database and if the entry is 'static' (see Section 3.4.10.1) or if the port is 'locked' (see bit 13 above), the source port's bit is checked to ensure the SA has been assigned to this port. If the SA is NOT assigned to this port it is considered an ATU Member Violation. If the IgnoreWrongData bit is cleared to a zero, an ATU Member Violation interrupt will be generated. If the IgnoreWrongData bit is set to a one the ATU Member Violation error is masked and ignored.
11	Reserved	RES	Reserved for future use
10:0	PAV	RWS to all zeros except for this port's bit	Port Association Vector for ATU learning. The value in these bits is used as the port's DPV on automatic ATU Learning or Entry_State refresh whenever these bits contain a non-zero value. When these bits are all zero, automatic Learning and Entry_State refresh is disabled on this port. For normal switch operation, this port's bit should be the only bit set in the vector. These bits must only be changed when frames are not entering the port (see PortState bits in Port Control – Table 65). The PAV bits can be used to set up port trunking (along with the VLANTable bits (Table 67). For the two ports that form a trunk, set both of their port's bits in both port's PAV registers, then use the VLANTable to isolate the two ports from each other and to steer the traffic from the other ports down the desired trunk line of the pair using DA/SA Load Balancing.

Table 73: InDiscards Low Counter
Offset: 0x10 or Decimal 16

Bits	Field	Type	Description
15:0	InDiscardsLo	RO	InDiscards Low Frame Counter. This counter contains the lower 16-bits of the 32-bit InDiscards counter. This 32-bit counter increments each time a good, non-filtered, frame is received but it cannot be forwarded owing to a lack of buffer memory. The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (see PortState Table 65). This register will be cleared by a Flush All Counters for this port or all ports StatsOp command (Table 107).

Table 74: InDiscards High Counter
Offset: 0x11 or Decimal 17

Bits	Field	Type	Description
15:0	InDiscardsHi	RO	InDiscards High Frame Counter. This counter contains the upper 16-bits of the 32-bit InDiscards counter. This 32-bit counter increments each time a good, non-filtered, frame is received but it cannot be forwarded due to a lack of buffer memory. The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (see PortState Table 65). This register is cleared by a Flush All Counters for this port or all ports StatsOp command (Table 107).

Table 75: InFiltered Counter
Offset: 0x12 or Decimal 18

Bits	Field	Type	Description
15:0	InFiltered	RO	InFiltered Frame Counter. This 16-bit counter gets incremented each time a good frame enters this port that was not forwarded due to filtering rules. The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (see PortState Table 65). This register will be cleared by a Flush All Counters for this port or all ports StatsOp command (Table 107).

Table 76: OutFiltered Counter
Offset: 0x13 or Decimal 19

Bits	Field	Type	Description
15:0	OutFiltered	RO	OutFiltered Frame Counter. This 16-bit counter gets incremented each time a good frame enters this port that was not forwarded due to filtering rules. The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (see PortState Table 65). This register will be cleared by a Flush All Counters for this port or all ports StatsOp command (Table 107).

Table 77: Port IEEE Priority Remapping Registers (88E6095 and 88E6095F devices only)
Offset: 0x18 or Decimal 24

Bits	Field	Type	Description
15	Reserved	RES	Reserved for future use.
14:12	TagRemap3	RWS to 0x3	Tag Remap 3. All IEEE tagged frames with a priority of 3 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
11	Reserved	RES	Reserved for future use.
10:8	TagRemap2	RWS to 0x2	Tag Remap 2. All IEEE tagged frames with a priority of 2 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
7	Reserved	RES	Reserved for future use.
6:4	TagRemap1	RWS to 0x1	Tag Remap 1. All IEEE tagged frames with a priority of 1 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
3	Reserved	RES	Reserved for future use.
2:0	TagRemap0	RWR	Tag Remap 0. All IEEE tagged frames with a priority of 0 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.

Table 78: Port IEEE Priority Remapping Registers (88E6095 and 88E6095F devices only)
Offset: 0x19 or Decimal 25

Bits	Field	Type	Description
15	Reserved	RES	Reserved for future use.
14:12	TagRemap7	RWS to 0x7	Tag Remap 7. All IEEE tagged frames with a priority of 7 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
11	Reserved	RES	Reserved for future use.
10:8	TagRemap6	RWS to 0x6	Tag Remap 6. All IEEE tagged frames with a priority of 6 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
7	Reserved	RES	Reserved for future use.
6:4	TagRemap5	RWS to 0x5	Tag Remap 5. All IEEE tagged frames with a priority of 5 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.
3	Reserved	RES	Reserved for future use.
2:0	TagRemap4	RWS to 0x4	Tag Remap 4. All IEEE tagged frames with a priority of 4 get this register's value as the frame's new priority inside the switch. If a tagged frame egresses the switch tagged, this new priority is written to the frame's tag.

6.5 Switch Global Registers

The devices contain global registers that affect all Ethernet ports in the device. Each global register is 16-bits wide. Global registers' bit assignments are shown in Figure 44.

Figure 44: Global Register Bit Map

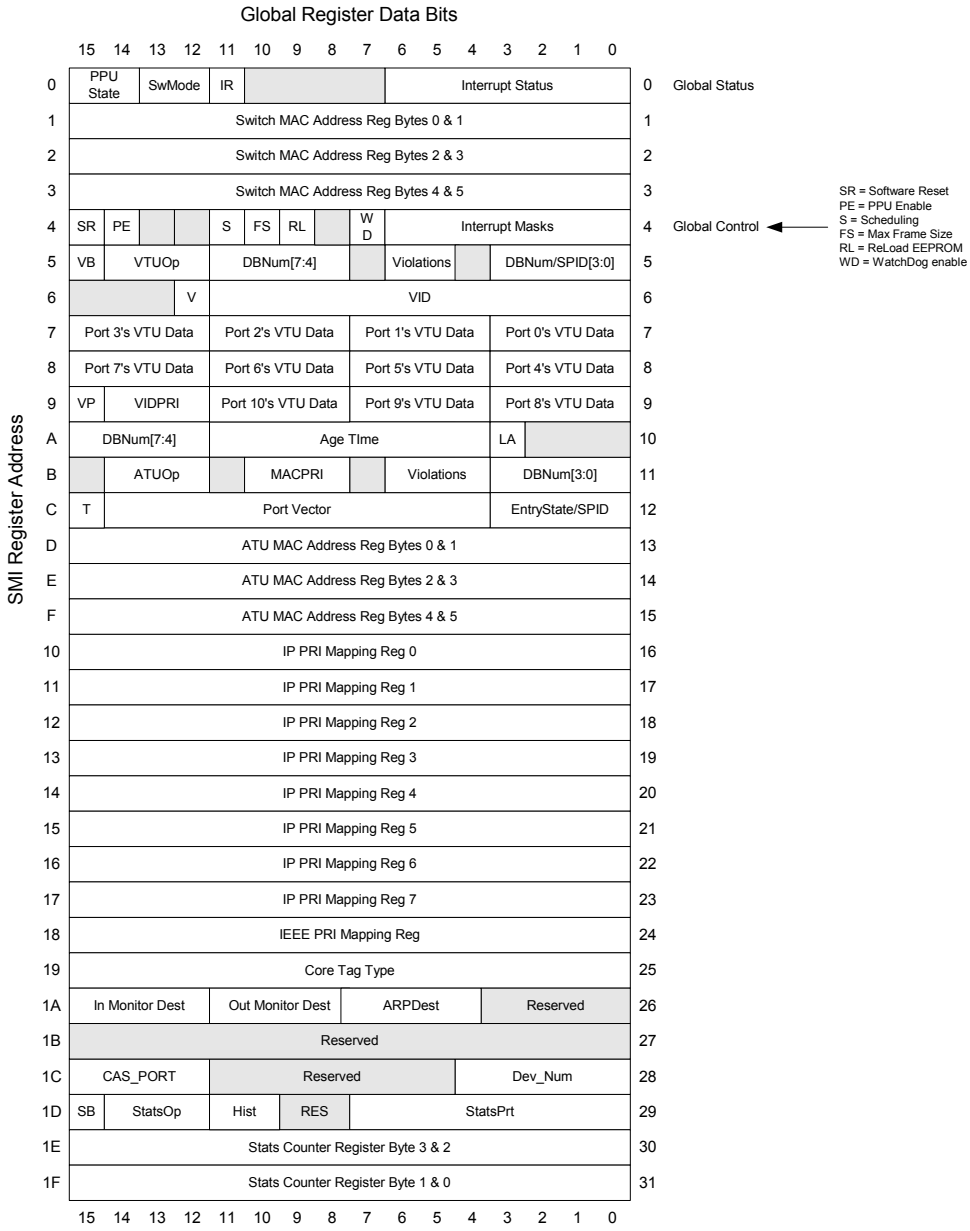

Table 79: Switch Global Status Register
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
15:14	PPUState	RO	PHY Polling Unit State. These bits indicate the state of the PPU as follows: 00 = PPU Disabled at Reset. This state indicates the PPU was disabled at Reset and it has not initialized or detected any external PHYs. Software can write to the PortStatus registers, i.e., the port's PHYDetect bits can be set by software (Table 61) and software can access the external PHYs registers. 01 = PPU is Active detecting and initializing external PHYs. The PortStatus registers (Table 61) must not be written by software and software must not try to access the external (PHY) registers. 10 = PPU Disabled after Initialization. This state indicates the PPU was enabled at some time and completed the PHY detection and initialization. Software can write to the PortStatus registers (Table 61) and it can access the external PHY's registers. 11 = PPU Polling. This indicates the PPU is Active polling the external PHYs. Software can write to the PortStatus registers (Table 61) but it must not try to access the external (PHY) registers.
13:12	SW_Mode	RO	Switch Mode. These bits return the value of the SW_MODE[1:0] pins.
11	InitReady	RO	SwitchReady. This bit is set to a one when the Address Translation Unit, the VLAN Translation Unit, the Queue Controller and the Statistics Controller complete their initialization and are ready to accept frames.
10:7	Reserved	RES	Reserved for future use.
6	StatsDone	LH	Statistics Done Interrupt. This bit is set to a one whenever the STATBusy bit (Table 107) transitions from a one to a zero. It is automatically cleared when read. This bit being high causes the device's INTn pin to go low if the STATDoneIntEn bit in the Global Control register (Table 83) is set to a one.
5	VTUProb	RO	VLAN Table Problem/Violation Interrupt. This bit is set to a one if a VLAN Violation is detected. It is automatically cleared when all of the pending VTU Violations have been serviced by the VTU Get/Clear Violation Data operation (Table 84). This bit being high causes the device's INTn pin to go low if the VTUProbIntEn bit in Global Control (Table 83) is set to a one.
4	VTUDone	LH	VTU Done Interrupt. This bit is set to a one whenever the VTUBusy bit (Table 84) transitions from a one to a zero. It is automatically cleared when read. This bit being high causes the device's INTn pin to go low if the VTUDoneIntEn bit in Global Control (Table 83) were set to a one.
3	ATUProb	RO	ATU Problem/Violation Interrupt. This bit is set to a one if the ATU cannot load or learn a new mapping due to all the available locations for an address being static or if an ATU Violation is detected. It is automatically cleared when all the pending ATU Violations have been serviced by the ATU Get/Clear Violation Data operation (Table 90). This bit being high causes the device's INTn pin to go low if the ATUProbIntEn bit in Global Control (Table 84) is set to a one.

Table 79: Switch Global Status Register (Continued)
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
2	ATUDone	LH	ATU Done Interrupt. This bit is set to a one whenever the ATUBusy bit (Table 90) transitions from a one to a zero. It is automatically cleared when read. This bit being high causes the device's INTn pin to go low if the ATUDoneIntEn bit in Global Control (Table 83) is set to a one.
1	PHYInt	RO	PHY Interrupt. This bit is set to a one when the internal PHYs interrupt logic has at least one active interrupt (from ports 0 to 7). This bit being high causes the device's INTn pin to go low if the PHYIntEn bit in Global Control (Table 83) is set to a one. This bit does not clear until all the PHY interrupts are serviced.
0	EEInt	LH	EEPROM Done Interrupt. This bit is set to a one after the EEPROM is done loading registers and it is automatically cleared when read. This bit being high causes the device's INTn pin to go low if the EEIntEn bit in Global Control (Table 83) is set to a one.

Table 80: Switch MAC Address Register Bytes 0 & 1
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
15:9	MACByte0	RWR	MAC Address Byte 0 (bits 47:41) used as the switch's source address (SA) in transmitted full-duplex Pause frames. Since bit 0 of byte 0 (bit 40) is the multicast bit (it is the first bit down the wire), it is always transmitted as a zero, and its value cannot be changed.
8	DiffAddr	RWR	Different MAC Addresses per Port. This bit is used to have all ports transmit the same or different source addresses in full-duplex Pause frames. 0 = All ports transmit the same SA 1 = Each port uses a different SA where bits 47:4 of the MAC address are the same, but bits 3:0 are the port number (Port 0 = 0, Port 1 = 1, Port 2 = 2, Port 3 = 3, etc.)
7:0	MACByte1	RWR	MAC Address Byte 1 (bits 39:32) used as the switch's source address (SA) in transmitted full-duplex Pause frames.

Table 81: Switch MAC Address Register Bytes 2 & 3,
Offset: 0x02 or Decimal 2

Bits	Field	Type	Description
15:8	MACByte2	RWR	MAC Address Byte 2 (bits 31:24) used as the switch's source address (SA) in transmitted full-duplex Pause frames.
7:0	MACByte3	RWR	MAC Address Byte 3 (bits 23:16) used as the switch's source address (SA) in transmitted full-duplex Pause frames.

Table 82: Switch MAC Address Register Bytes 4 & 5
Offset: 0x03 or Decimal 3

Bits	Field	Type	Description
15:8	MACByte4	RWR	MAC Address Byte 4 (bits 15:8) used as the switch's source address (SA) in transmitted full-duplex Pause frames.
7:0	MACByte5	RWR	MAC Address Byte 5 (bits 7:0) used as the switch's source address (SA) in transmitted full-duplex Pause frames. Note: Bits 3:0 of this register are ignored if DiffAddr, above, is set to a one.

Table 83: Switch Global Control Register
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
15	SWReset	SC	Switch Software Reset. Writing a one to this bit causes the QC, the MAC state machines and the PPU state machines in the switch to be reset. Register values are not modified, except for the PortState bits (Table 65) and the PHYDetect bit (Table 61). The EEPROM is not re-read. The ATU and VTU are not affected by this bit, but the PHYs are (when the PPU performs its PHY detect and initialization). When the reset operation is complete, this bit is cleared to a zero automatically. The reset occurs immediately. To prevent transmission of CRC frames, set all of the ports to the Disabled state (Table 65), and wait for 2 ms. (i.e., the time for a maximum frame to be transmitted at 10 Mbps) before setting the SWReset bit to a one.
14	PPUEn	RWR or RWS[1]	PHY Polling Unit Enabled. When this bit is set to a one the PPU is Active. At the appropriate time after reset the PPU detects and initializes any external PHYs that it finds and then it polls those PHYs before polling for Link, Speed, Duplex and Pause data that is reflected in the PortStatus register (Table 61). When this bit is cleared to a zero the PPU finishes what is was doing and then frees the MDC_PHY and MDIO_PHY pins so software can directly access the external PHY registers. Refer to the PPUState bits in the Global Status register (Table 79) to determine the state of the PPU and to see if it is safe to access the external (PHY) registers. Setting this bit back to a one re-connects the PPU to the MDC_PHY and MDIO_PHY pins with the PHY re-entering its polling function without performing PHY detection and initialization. To re-start the PPU including re-doing PHY detection and initialization, set this bit to a one along with setting the SWReset bit above to a one.
13	Discard Excessive	RWR	Discard frames with Excessive Collisions. When this bit is set to a one frames that encounter 16 consecutive collisions are discarded. When this bit is cleared to a zero Egress frames are never discarded and the backoff range is reset after 16 consecutive collisions on a singe frame.
12	Reserved	RES	Reserved for future use.
11	Scheduling	RWR	Scheduling mode. This bit is used to select the Queue Controller's scheduling mode as follows: 0 = Use an 8, 4, 2, 1 weighted fair queuing scheme 1 = Use a strict priority scheme

Table 83: Switch Global Control Register (Continued)
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
10	MaxFrame Size	RWR	Maximum Frame Size allowed. The Ingress block discards all frames less than 64 bytes in size. It also discards all frames greater than a certain size (regardless of whether or not the frame is IEEE 802.3ac tagged) as follows: 0 = Max size is 1522 for IEEE tagged frames, 1518 for untagged frames 1 = Max size is 1632
9	ReLoad	SC	Reload the registers using the EEPROM. When this bit is set to a one, the contents of the external EEPROM are used to load the registers just as if a reset had occurred. When the reload operation finishes, this bit is cleared to a zero automatically and the EEInt interrupt bit is set.
8:7	Reserved	RES	Reserved for future use.
6	StatsDone IntEn	RWR	Statistics Operation Done Interrupt Enable. This bit must be set to a one to allow the Stat Done interrupt to drive the device's INTn pin low.
5	VTUProb IntEn	RWR	VLAN Problem/Violation Interrupt Enable. This bit must be set to a one to allow the VTUProblem interrupt to drive the device's INTn pin low.
4	VTUDone IntEn	RWR	VLAN Table Operation Done Interrupt Enable. This bit must be set to a one to allow the VTUDone interrupt to drive the device's INTn pin low.
3	ATUProb IntEn	RWR	ATU Problem/Violation Interrupt Enable. This bit must be set to a one to allow the ATU Problem interrupt to drive the device's INTn pin low.
2	ATUDone IntEn	RWR	ATU Operation Done Interrupt Enable. This bit must be set to a one to allow the ATU Done interrupt to drive the device's INTn pin low.
1	PHYIntEn	RWR	PHY Interrupt Enable. This bit must be set to a one to allow active interrupts enabled in PHY register 0x12 to drive the device's INTn pin low.
0	EEIntEn	RWS	EEPROM Done Interrupt Enable. This bit must be set to a one to allow the EEPROM Done interrupt to drive the device's INTn pin low.

1. PPUEn will be initialized set if the P9_TXD[3]/SW_24P configuration pin is low and the MDC_PHY/PPU_EN/LRN2ALL configuration pin is high at the rising edge of RESETn - or - if the P9_TXD[3]/SW_24P configuration pin is high at the rising edge of RESETn.

Table 84: VTU Operation Register
Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
15	VTUBusy	SC	VLAN Table Unit Busy. This bit must be set to a one to start a VTU operation (see VTUOp below). Only one VTU operation can be executing at one time so this bit must be zero before setting it to a one. When the requested VTU operation completes, this bit will automatically be cleared to a zero. The transition of this bit from a one to a zero can be used to generate an interrupt (Table 83).
14:12	VTUOp	RWR	VLAN Table Unit Table Opcode. The devices support the following VTU operations (all of these operations can be executed while frames are transiting through the switch): 000 = No Operation 001 = Flush All Entries 010 = No Operation 011 = Load[1] or Purge[2] an Entry 100 = Get Next[3] 101 = Reserved 110 = Reserved 111 = Get/Clear Violation Data[4]
11:8	DBNum[7:4]	RWR	VTU MAC Address Database Number bits 7:4. On all VTUOps except for Get Violation Data, this field is DBNum[7:4] and it is used to separate MAC address databases by a frame's VID. If multiple address databases are not being used, these bits must remain zero. If multiple address databases are being used, these bits are used to set the desired database number that is associated with VID value on Load operations (or used to read the currently assigned DBNum on Get Next operations). The lower four bits of the VTU's DBNum are in bits 3:0 of this register.
7	Reserved	RES	Reserved for future use
6	Member Violation	RO	Source Member Violation. On Get/Clear Violation Data VTUOps, this bit is returned set to a one if the Violation being serviced is due to an 802.1Q Member Violation. A Member Violation occurs when an 802.1Q enabled Ingress port accesses the VTU with a VID that is contained in the VTU but whose Membership list does not include this Ingress port. Only the first Member Violation or Miss Violation (below) will be saved until cleared.
5	Miss Violation	RO	VTU Miss Violation. On Get/Clear Violation Data VTUOps this bit is returned set to a one if the Violation being serviced was due to an 802.1Q Miss Violation. A Miss Violation occurs when an 802.1Q enabled Ingress port accesses the VTU with a VID that is not contained in the VTU. Only the first Miss Violation or Member Violation (above) is saved until cleared.

Table 84: VTU Operation Register (Continued)
Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
4	Reserved	RES	Reserved for future use
3:0	DBNum[3:0]/ SPID	RWR	VTU MAC Address Database Number bits 3:0 or Source Port ID. On Load and GetNext VTUOps, this field is DBNum[3:0] and it is used to separate MAC address databases by a frame's VID. If multiple address databases are not being used, these bits must remain zero. If multiple address databases are being used, these bits are used to set the required address database number that is associated with a VID value on Load operations (or used to read the currently assigned DBNum on Get Next operations). The upper four bits of the VTU's DBNum are in bits 11:7 of this register. On the Get Violation Data VTUOp, this field returns the Source Port ID of the port that caused the violation. If SPID 0xF the source of the violations was the CPU register interface (i.e., the VTU was full during a CPU Load operation).

1. An Entry is Loaded when the Valid bit (Table 85) is set to a one. Load is the only VTUOp that uses the DBNum field and it uses it as data to be loaded along with the desired VID.
2. An Entry is Purged when the Valid bit (Table 85) is cleared to a zero.
3. A Get Next operation finds the next higher VID currently in the VTU's database. The VID value (Table 85) is used as the VID to start from. To find the lowest VID set the VID field to ones. When the operation is done the VID field contains the next higher VID currently active in the VTU. To find the next VID simply issue the Get Next opcode again. If the VID field is returned set to all one's with the Valid bit cleared to zero, no higher VID's were found. To Search for a particular VID, perform a Get Next operation using a VID field with a value one less than the one being searched for.
4. When the VTUProb bits is set to a one (Global Status—Table 79) the Get/Clear Violation VTUOp can be used to retrieve the data associated with the Violation. It will return the source port of the violation in the SPID field of this registers (bits 3:0) and it will return the VID of the violation in the VID field of the VTU VID register (Table 85). When all Violations currently pending in the VTU have been serviced the VTUProb bit in Global Status will be cleared to a zero.

Table 85: VTU VID Register
Offset: 0x06 or Decimal 6

Bits	Field	Type	Description
15:13	Reserved	RES	Reserved for future use.
12	Valid	RWR	Entry's Valid bit. At the end of Get Next operations, if this bit is set to a one it indicates the VID value below is valid. If this bit is cleared to a zero and the VID is all ones, it indicates the end of the VID list was reached with no new valid entries found. On Load or Purge operations, this bit indicates the desired operation of a Load (when set to a one) or a Purge (when cleared to a zero).
11:0	VID	RWR	VLAN Identifier. This VID is used in all the VTUOp commands (except Get/Clear Violation Data) and it is the VID that is associated with the VTU data below (Table 86) or the VID that caused the VTU Violation.

Table 86: VTU Data Register Ports 0 to 3
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15:14	PortStateP3	RWR	Per VLAN Port States for Port 3. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP0 below.
13:12	Member TagP3	RWR	Membership and Egress Tagging for Port 3. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP0 below.
11:10	PortStateP2	RWR	Per VLAN Port States for Port 2. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP0 below.
9:8	Member TagP2	RWR	Membership and Egress Tagging for Port 2. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP0 below.
7:6	PortStateP1	RWR	Per VLAN Port States for Port 1. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero of 802.1s is not used. See PortStateP0 below.
5:4	Member TagP1	RWR	Membership and Egress Tagging for Port 1. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP0 below.
3:2	PortStateP0	RWR	Per VLAN Port States for Port 0. These bits are used to support 802.1s (per VLAN spanning tree) and should be cleared to zero if 802.1s is not used. The Per VLAN Port States are: 00 = 802.1s Disabled. Use non-VLAN Port States for this port from frames with this VID. 01 = Blocking/Listening Port State for this port for frames with this VID. 10 = Learning Port State for this port for frames with this VID. 11 = Forwarding Port State for this port for frames with this VID.
1:0	Member TagP0	RWR	Membership and Egress Tagging for Port 0. These bits are used to support 802.1Q membership and Egress Tagging as follows: 00 = Port is a member of this VLAN and frames are to egress unmodified. 01 = Port is a member of this VLAN and frames are to egress Untagged. 10 = Port is a member of this VLAN and frames are to egress Tagged. 11 = Port is not a member of this VLAN. Any frames with this VID[1] are discarded at ingress and are not allowed to egress this port.

1. The VID used is the VID assigned to the frame during ingress (Section 3.5.7.3.)

Table 87: VTU Data Register Ports 4 to 7
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
15:14	PortStateP7	RWR	Per VLAN Port States for Port 7. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP4 below.
13:12	Member TagP7	RWR	Membership and Egress Tagging for Port 7. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP4 below.
11:10	PortStateP6	RWR	Per VLAN Port States for Port 6. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP4 below.
9:8	Member TagP6	RWR	Membership and Egress Tagging for Port 6. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP4 below.
7:6	PortStateP5	RWR	Per VLAN Port States for Port 5. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP4 below.
5:4	Member TagP5	RWR	Membership and Egress Tagging for Port 5. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP4 below.
3:2	PortStateP4	RWR	Per VLAN Port States for Port 4. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. The Per VLAN Port States are: 00 = 802.1s Disabled. Use non-VLAN Port States for this port for frames with this VID. 01 = Blocking/Listening Port State for this port for frames with this VID. 10 = Learning Port State for this port for frames with this VID. 11 = Forwarding Port State for this port for frames with this VID.
1:0	Member TagP4	RWR	Membership and Egress Tagging for Port 4. These bits are used to support 802.1Q membership and Egress Tagging as follows: 00 = Port is a member of this VLAN and frames are to egress unmodified. 01 = Port is a member of this VLAN and frames are to egress Untagged. 10 = Port is a member of this VLAN and frames are to egress Tagged. 11 = Port is not a member of this VLAN. Any frames with this VID[1] are discarded at ingress and are not allowed to egress this port.

1. The VID used is the VID assigned to the frame during ingress (Section 3.5.7.3.)

Table 88: VTU Data Register Port 8 to 10
Offset: 0x09 or Decimal 9

Bits	Field	Type	Description
15	VIDPRI Override(88E6095 and 88E6095F devices only)	RWR	VID Priority Override. When this bit is set to a one the VIDPRI bits (below) are used to override the priority on any frame associated with this VID as long as the port's VTUPriOverride bit is set (Table 69).
14:12	VIDPRI	RWR	VID Priority bits. These bits are used to override the priority on any frames associated with this VID value, if the VIDPRIOverride bit (above) is set to a one.
11:10	PortStateP10	RWR	Per VLAN Port States for Port 10. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP8 below.
9:8	MemberTagP10	RWR	Membership and Egress Tagging for Port 10. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP8 below.
7:6	PortStateP9	RWR	Per VLAN Port States for Port 9. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. See PortStateP8 below.
5:4	MemberTagP9	RWR	Membership and Egress Tagging for Port 9. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP8 below.
3:2	PortStateP8	RWR	Per VLAN Port States for Port 8. These bits are used to support 802.1s (per VLAN Spanning Tree) and should be cleared to zero if 802.1s is not used. The Per VLAN Port States are: 00 = 802.1s Disabled. Use non-VLAN Port States for this port for frames with this VID. 01 = Blockout/Listening Port State for this port for frames with this VID. 10 = Learning Port State for this port for frames with this VID. 11 = Forwarding Port State for this port for frames with this VID.
1:0	MemberTagP8	RWR	Membership and Egress Tagging for Port 8. These bits are used to support 802.1Q membership and Egress Tagging as follows: 00 = Port is a member of this VLAN and frames are to egress unmodified. 01 = Port is a member of this VLAN and frames are to egress Untagged. 10 = Port is a member of this VLAN and frames are to egress Tagged. 11 = Port is not a member of this VLAN. Any frames with this VID[1] are discarded at ingress and are not allowed to egress this port.

1. The VID used is the VID assigned to the frame during ingress (Section 3.5.7.3.)

Table 89: ATU Control Register
Offset: 0x0A or Decimal 10

Bits	Field	Type	Description
15:12	DBNum[7:4]	RWR	ATU MAC Address Database Number bits 7:4. If multiple address databases are not being used these bits must remain zero. If multiple address databases are being used these bits are used to set the desired address database number that is to be used to set the desired address database number that is to be used on the Database supported commands (ATUOps 0x3, 0x4, 0x5, and 0x6 above). On Get/Clear Violation Data ATUOps these bits return the DBNum[7:4] value associated with the ATU violation that was just serviced. The lower four bits of the ATU's DBNum are in bits 3:0 of the ATU Operation register at offset 0x0B.
11:4	AgeTime	RWS to 0x14	ATU Age Time. These bits determine the time that each ATU Entry remains valid in the database, since its last access as a source address, before being purged. The value in this register times 15 is the age time in seconds. For example: The default value of 0x14 is 20 decimal. 20 x 15 = 300 seconds or 5 minutes. The minimum age time is 0x1 or 15 seconds. The maximum age time is 0xFF or 3825 seconds or almost 64 minutes. If the AgeTime is set to 0x0 the Aging function is disabled and all learned addresses will remain in the database forever.
3	Learn2All	RWR or RWS[1]	Learn to All devices in a Switch. When more than one Marvell device is used to for a single 'switch' it may be desirable for all devices in the 'switch' to learn any address this device learns[2]. When this bit is set to a one all other devices in the 'switch' learn the same addresses this device learns. When this bit is cleared to a zero only the devices that actually receive frames will learn from those frames. This mode typically supports more active MAC addresses at one time as each device in the switch does not need to learn addresses it may never use.
2:0	Reserved	RES	Reserved for future use.

1. Learn2All is initialized set if both the P9_TXD3/SW_24P configuration pin and the MDC_PHY/PPU_EN/LRN2ALL configuration pin are high at the rising edge of RESETn.
2. Learn2All message learning frames will be sent out a port if that port's MessagePort bit is set to a one (Table 66). If this frame is used it is recommended that all DSA Tag ports, except for the CPU's port, have their MessagePort but set to a one. Ports that are not DSA Tag ports (i.e., normal Network ports) should not have their MessagePort bit set to a one.

Table 90: ATU Operation Register
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
15	ATUBusy	SC	Address Translation Unit Busy. This bit must be set to a one to start an ATU operation (see ATUOp below). Only one ATU operation can be executing at one time so this bit must be zero before setting it to a one. When the requested ATU operation completes, this bit is automatically be cleared to a zero. The transition of this bit from a one to a zero can be used to generate an interrupt (Table 83).
14:12	ATUOp	RWR	Address Translation Unit Opcode. The devices support the following ATU operations. (All of these operations can be executed while frames are passing through the switch): 000 = No Operation 001 = Flush[1] All Entries 010 = Flush all Non-Static[2] Entries 011 = Load[3] or Purge[4] an Entry in a particular DBNum Database 100 = Get Next[5] from a particular DBNum Database 101 = Flush All Entries in a particular DBNum Database 110 = Flush all Non-Static Entries in a particular DBNum Database 111 = Get/Clear Violation Data[6]
11	Reserved	RES	Reserved for future use.
10:8	MACPri (88E6095 and 88E6095F devices only)	RWR	MAC Priority bits. These bits are used to override the priority on any frames associated with this MAC value, if the EntryState bits indicate MAC Priority can be used – see Section 3.4.10.1) and the port's SA and/or DA priority overrides are enabled (in Port Control 2 – Table 69).
7	Reserved	RES	Reserved for future use.
6	Member Violation	RO	Source Port Violation. On Get/Clear Violation Data ATUOps this bit is returned set to a one if the Violation being serviced is due to a Source Address look-up that resulted in a Hit but where the ATUData[8:0] bits does not contain the frame's Ingress port bit set to a one (i.e., a station move occurred). This violation can be masked on a per port basis by setting the port's IngoreWrongData bit. Only the first Member Violation, Miss Violation (below) or Full Violation (below) is be saved per port until cleared.
5	Miss Violation	RO	ATU Miss Violation. On Get/Clear Violation Data ATUOps this bit is returned set to a one if the Violation being serviced is due to a Source Address look-up that resulted in a Miss on ports that are Locked (i.e., CPU directed learning is enabled on the port). Only the first Miss Violation, Member Violation (above) or Full Violation (below) is saved per port until cleared.

Table 90: ATU Operation Register (Continued)
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
4	ATUFull Violation	RO	ATU Full Violation. On Get/Clear Violation Data ATUOps this bit is set to a one if the Violation being serviced is due to a Load ATUOp or automatic learn that could not store the desired entry. This only occurs if all available locations for the desired address contain other MAC addresses that are loaded Static. Only the first Full Violation, Member Violation (above) or Miss Violation (above) is saved per port until cleared.
3:0	DBNum[3:0]	RWR	ATU MAC Address Database Number bits 3:0. If multiple address databases are not being used these bits must remain zero. If multiple address databases are being used these bits are used to set the desired address database number that is to be used on the Database supported commands (ATUOps 0x3, 0x4, 0x5 and 0x6 above). On Get/Clear Violation Data ATUOps these bits return the DBNum[3:0] value associated with the ATU violation that was just serviced. The upper four bits of the ATU's DBNum are in bits 15:12 of the ATU Control register at Offset 0x0A.

1. A Flush occurs when the EntryState (Table 91) is zero.
2. A Non-Static entry is any unicast address with an EntryState less than 0x8. All unicast frames flood until new addresses are learned.
3. An Entry is Loaded when the EntryState (Table 91) is non-zero.
4. An Entry is Purged when the EntryState (Table 91) is zero.
5. A Get Next operation finds the next higher MAC address currently in a particular ATU database (defined by the DBNum field). The ATUByte[5:0] values (Table 92) are used as the starting address. To find the lowest MAC address set ATU[5:0] to ones. When the operation is done, ATUByte[5:0] contains the next higher MAC address. To find the next address, simply issue the Get Next opcode again. If ATUByte[5:0] is returned set to all one's with an EntryState of 0x0, no higher MAC address was found. If ATUByte[5:0] is returned set to all one's with a non-zero EntryState, the highest MAC address was found (i.e., the Broadcast address) and the end of the table was reached. To search for a particular address, perform a Get Next operation using a MAC address with a value one less than the one being searched for.
6. When the ATUProb bit is set to a one (Global Status - Table 79), the Get/Clear Violation ATUOp can be used to retrieved the data associated with the violation. When all violations currently pending in the ATU have been serviced the ATUProb bit in the Global Status is cleared to a zero.

Table 91: ATU Data Register
Offset: 0x0C or Decimal 12

Bits	Field	Type	Description
15	Trunk	RWR	Trunk Mapped Address. When this bit is set to a one the data bits 7:4 below (PortVec bits [3:0]) is the Trunk ID assigned to this address. PortVec bits [10:4] must be written as zero when this bit is set to a one. When this bit is cleared to a zero the data in bits 14:4 below (PortVec bit[10:0]) is the port vector assigned to this address.
14:4	PortVec/ TrunkID	RWR	Port Vector. If the Trunk bit, above, is zero, these bits are used as the input Port Vector for ATU Load operations and it's the resulting Port Vector from ATU Get Next operations. If the Trunk bit, above, is one, these bits are used as the input TrunkID for ATU load operations and it is the resulting TrunkID from ATU GetNext operations.
3:0	EntryState/ SPID	RWR	ATU Entry State. These bits are used as the Entry State for ATU Load/ Purge or Flush/Move operations and it is the resulting Entry State from ATU Get Next operations (GetNext is the only ATU operation supported in the devices). If these bits equal 0x0 then the ATUOp is a Purge or a Flush. If these bits are not 0x0 then the ATUOp is a Load or a Move (a Move ATUOp requires these bits to be 0xF). On Get/Clear Violation Data ATUOps, these bits return the Source Port ID (SPID) associated with the ATU violation that was just serviced. If SPID = 0xF the source of the violation was the CPU's register interface (i.e., the ATU was full during a CPU Load operation).

Table 92: ATU MAC Address Register Bytes 0 & 1
Offset: 0x0D or Decimal 13

Bits	Field	Type	Description
15:8	ATUByte0	RWR	ATU MAC Address Byte 0 (bits 47:40) used as the MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. Bit 0 of byte 0 (bit 40) is the multicast bit (it is the first bit down the wire). Any MAC address with the multicast bit set to a one is considered Static by the ATU. On Get/Clear Violation Data ATUOps these bits return ATUByte0 associated with the ATU violation that was just serviced.
7:0	ATUByte1	RWR	ATU MAC Address Byte 1 (bits 39:32) used as the input MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. On Get/Clear Violation Data ATUOps, these bits return ATUByte1 associated with the ATU violation that was just serviced.

Table 93: ATU MAC Address Register Bytes 2 & 3
Offset: 0x0E or Decimal 14

Bits	Field	Type	Description
15:8	ATUByte2	RWR	ATU MAC Address Byte 2 (bits 31:24) used as the input MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. On Get/Clear Violation Data ATUOps, these bits return ATUByte2 associated with the ATU violation that was just serviced.
7:0	ATUByte3	RWR	ATU MAC Address Byte 3 (bits 23:16) used as the input MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. On Get/Clear Violation Data ATUOps, these bits return ATUByte3 associated with the ATU violation that was just serviced.

Table 94: ATU MAC Address Register Bytes 4 & 5
Offset: 0x0F or decimal 15

Bits	Field	Type	Description
15:8	ATUByte4	RWR	ATU MAC Address Byte 4 (bits 15:8) used as the input MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. On Get/Clear Violation Data ATUOps, these bits return ATUByte4 associated with the ATU violation that was just serviced.
7:0	ATUByte5	RWR	ATU MAC Address Byte 5 (bits 7:0) used as the input MAC address for ATU Load, Purge or Get Next operations and it is the resulting MAC address from ATU Get Next operations. On Get/Clear Violation Data ATUOps, these bits return ATUByte5 associated with the ATU violation that was just serviced.

Table 95: IP-PRI Mapping Register 0
Offset: 0x10 or Decimal 16

Bits	Field	Type	Description
15:14	IP_0x1C	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x1C.
13:12	IP_0x18	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x18.
11:10	IP_0x14	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x14.
9:8	IP_0x10	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x10.
7:6	IP_0x0C	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x0C.
5:4	IP_0x08	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x08.
3:2	IP_0x04	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x04.
1:0	IP_0x00	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x00.

Table 96: IP-PRI Mapping Register 1
Offset: 0x11 or Decimal 17

Bits	Field	Type	Description
15:14	IP_0x3C	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x3C.
13:12	IP_0x38	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x38.
11:10	IP_0x34	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x34.
9:8	IP_0x30	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x30.
7:6	IP_0x2C	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x2C.
5:4	IP_0x28	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x28.
3:2	IP_0x24	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x24.
1:0	IP_0x20	RWR	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x20.

Table 97: IP-PRI Mapping Register 2
Offset: 0x12 or Decimal 18

Bits	Field	Type	Description
15:14	IP_0x5C	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x5C.
13:12	IP_0x58	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x58.
11:10	IP_0x54	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x54.
9:8	IP_0x50	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x50.
7:6	IP_0x4C	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x4C.
5:4	IP_0x48	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x48.
3:2	IP_0x44	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x44.
1:0	IP_0x40	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x40.

Table 98: IP-PRI Mapping Register 3
Offset: 0x13 or Decimal 19

Bits	Field	Type	Description
15:14	IP_0x7C	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x7C.
13:12	IP_0x78	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x78.
11:10	IP_0x74	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x74.
9:8	IP_0x70	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x70.
7:6	IP_0x6C	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x6C.
5:4	IP_0x68	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x68.
3:2	IP_0x64	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x64.
1:0	IP_0x60	RWS to 0x1	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x60.

Table 99: IP-PRI Mapping Register 4
Offset: 0x14 or Decimal 20

Bits	Field	Type	Description
15:14	IP_0x9C	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x9C.
13:12	IP_0x98	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x98.
11:10	IP_0x94	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x94.
9:8	IP_0x90	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x90.
7:6	IP_0x8C	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x8C.
5:4	IP_0x88	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x88.
3:2	IP_0x84	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x84.
1:0	IP_0x80	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0x80.

Table 100: IP-PRI Mapping Register 5
Offset: 0x15 or Decimal 21

Bits	Field	Type	Description
15:14	IP_0xBC	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xBC.
13:12	IP_0xB8	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xB8.
11:10	IP_0xB4	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xB4.
9:8	IP_0xB0	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xB0.
7:6	IP_0xAC	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xAC.
5:4	IP_0xA8	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xA8.
3:2	IP_0xA4	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xA4.
1:0	IP_0xA0	RWS to 0x2	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xA0.

Table 101: IP-PRI Mapping Register 6
Offset: 0x16 or Decimal 22

Bits	Field	Type	Description
15:14	IP_0xDC	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xDC.
13:12	IP_0xD8	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xD8.
11:10	IP_0xD4	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xD4.
9:8	IP_0xD0	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xD0.
7:6	IP_0xCC	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xCC.
5:4	IP_0xC8	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xC8.
3:2	IP_0xC4	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xC4.
1:0	IP_0xC0	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xC0.

Table 102: IP-PRI Mapping Register 7
Offset: 0x17 or Decimal 23

Bits	Field	Type	Description
15:14	IP_0xFC	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xFC.
13:12	IP_0xF8	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xF8.
11:10	IP_0xF4	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xF4.
9:8	IP_0xF0	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xF0.
7:6	IP_0xEC	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xEC.
5:4	IP_0xE8	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xE8.
3:2	IP_0xE4	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xE4.
1:0	IP_0xE0	RWS to 0x3	IPv4 and IPv6 mapping. The value in this field is used as the frame's priority if bits (7:2) of its IP TOS/DiffServ/Traffic Class value is 0xE0.

Table 103: IEEE-PRI Register
Offset: 0x18 or Decimal 24

Bits	Field	Type	Description
15:14	Tag_0x7	RWS to 0x3	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 7.
13:12	Tag_0x6	RWS to 0x3	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 6.
11:10	Tag_0x5	RWS to 0x2	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 5.
9:8	Tag_0x4	RWS to 0x2	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 4.
7:6	Tag_0x3	RWS to 0x1	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 3.
5:4	Tag_0x2	RWR	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 2.
3:2	Tag_0x1	RWR	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 1.
1:0	Tag_0x0	RWS to 0x1	IEEE 802.1p mapping. The value in this field is used as the frame's priority if its IEEE Tag has a value of 0.

Table 104: Core Tag Type
Offset: 0x19 or decimal 25

Bits	Field	Type	Description
15:0	CoreTag Type (88E6095 and 88E6095F devices only)	RWS to 0x8100	Core Tag Ether Type. This Ether Type is added to frames that egress the switch as Double Tagged frames. It is also the Ether Type expected during Ingress to determine if a frame is Tagged or not on ports configured in UseCoreTag mode (Table 65) as follows: For Ingress: UseCoreTag=0, DoubleTag=0 Tag Type = 0x8100 UseCoreTag=0, DoubleTag=1 Tag Type = 0x8100 UseCoreTag=1, DoubleTag=0 Tag Type = CoreTagType UseCoreTag=1, DoubleTag=1 1st Tag Type = CoreTagType then 2nd Tag Type in the frame = 0x8100[1] For Egress: Normal Add a Tag (standard IEEE rule) Added Tag Type = 0x8100 Always Add a Tag (Egress DoubleTag) Added Tag Type = CoreTagType

1. In this mode, a frame is considered Tagged only if the first Ethertype just after the frame's SA equals the CoreTagType and then this Tag is removed from the frame. The resulting frame is considered Tagged only if the remaining Ethertype equals 0x8100 (placed after the frame's SA subsequent to the previous Tag's being removed).

Table 105: Monitor Control
Offset: 0x1A or decimal 26

Bits	Field	Type	Description
15:12	Ingress Monitor Dest	RWS	Ingress Monitor Destination Port. Frames that are targeted toward an Ingress Monitor Destination go out the port number indicated in these bits. This includes frames received on a DSA Tag port with the Ingress Monitor type, and frames received on a Network port that is enabled to be the Ingress Monitor Source Port (Table 65). If the Ingress Monitor Destination Port resides in this device these bits should point to the Network port where these frames are to egress. If the Ingress Monitor Destination Port resides in another device these bits should point to the DSA Tag port in this device that is used to get to the device that contains the Ingress Monitor Destination Port.
11:8	Egress Monitor Dest	RWS	Egress Monitor Destination Port. Frames that are targeted toward an Egress Monitor Destination go out of the port number indicated in these bits. This includes frames received on a DSAI Tag port with the Egress Monitor type, and frames transmitted on a Network port that is enabled to be the Egress Monitor Source Port (Table 69). If the Egress Monitor Destination port resides in this device, these bits should point to the Network port where these frames are to egress. If the Egress Monitor Destination Port resides in another device, these bits should point to the DSA Tag port in this device that is used to reach the device that contains the Egress Monitor Destination Port.
7:4	ARPDest	RWS	ARP Monitor Destination Port. Tagged or Untagged Frames that ingress Network ports[1] that have the Broadcast Destination Address with an Ethertype of 0x0806 are mirrored to this port. The ARPDest should point to the port that directs these frames to the switch's CPU that will process ARPs. This target port should be a DSA Tag port so the frames will egress with a To_CPU DSA Tag with a CPU Code of ARP. To CPU DSA Tag frames with a CPU Code off ARP that ingress a DSA Tag port will be sent to the port number defined in ARPDest. If ARPDest = 0xF ARP Monitoring is disabled and ingressing To_CPU ARP frames will be discarded.
3:0	Reserved	RES	Reserved for future use.

1. Network ports are ports whose DSA_Tag bit (Table 66) is cleared to a zero.

Table 106: Global Control 2
Offset: 0x1C or Decimal 28

Bits	Field	Type	Description
15:12	Cascade Port	RWR	Cascade Port number. In multi-chip (Cascade) systems, frames coming from a CPU (From_CPU frames) need to know when they have reached their destination chip (i.e. when the frame's Trg_Dev = this device's DeviceNumber, below). Frames that have not reached their destination chip are sent out from this chip using the these bits as follows: A value of 0x0 to 0xA will transmit these frames out Port 0 to Port 10 respectively. Use a value of 0xE if the devices have no Cascade (i.e., it is the last device in the Cascade). The From_CPU frames will be discarded if they were not intended for this device. Use a value of 0xF when multiple devices are cascaded off a single port of this device. In this mode the frame's Trg_Dev value is used as an index into this device's Routing Table (Global 2 offset 0x06) and the From_CPU frames are sent to the port defined there.
11:5	Reserved	RES	Reserved for future use.
4:0	DeviceNumber	RWS to 0xXX[1]	Device Number. In multi-chip systems, frames coming from a CPU (From_CPU frames) need to know when they have reached their destination chip. From_CPU frames whose Dev_Num field matches these bits have reached their destination chip and are sent out from this chip using the port number indicated in the frame's Trg_Port field. The DeviceNumber value must be unique for each chip in a Multi-chip system. These bits are set at reset by the ADDR[4:0] configuration pins.

1. The ADDR[4:0] configuration pins are used to set the initial value of this register. The ADDR[4:0] pins are also used to select between Multi-chip addressing mode or Single-chip addressing mode. Changing the value in this register after reset does *not* change the device's addressing mode nor its SMI address.

Table 107: Stats Operation Register
Offset: 0x1D or Decimal 29

Bits	Field	Type	Description
15	StatsBusy	SC	Statistics Unit Busy. This bit must be set to a one to start a Stats operation (See StatsOp below). Only one Stats operation can be executing at one time so this bit must be zero before setting it to a one. When the requested Stats operation completes, this bit automatically is cleared to a zero. The transition of this bit from a one to a zero can be used to generate an interrupt (Table 83).
14:12	StatsOp	RWR	Statistics Unit Opcode. The devices support the following Stats operations (all of these operations can be executed while frames are transiting through the switch): 000 = No Operation 001 = Flush (clear) All Counters for all Ports 010 = Flush (clear) All Counters for a Port 011 = Reserved 100 = Read a Captured Counter 101 = Capture All Counters for a Port 11x = Reserved
11:10	Histogram Mode	RES to 0x3	Histogram Counters Mode. The Histogram mode bits control how the Histogram counters work as follows: 00 = Reserved 01 = Count received frames only 10 = Count transmitted frames only 11 = Count receive and transmitted frames

Table 107: Stats Operation Register (Continued)
Offset: 0x1D or Decimal 29

Bits	Field	Type	Description
9:6	Reserved	RES	Reserved for future use
5:0	StatsPtr	RWR	Statistics Pointer. This field is used as a parameter for the above StatsOp commands. It must be set to the desired Port number for the Capture All Counters for a Port (0x5) and Flush All Counters for a Port (0x2) StatsOps. Use 0x00 for Port 0, 0x01 for Port 1, etc. StatsPtr must be set to the desired counter to read for the Read a Captured Counter (0x4) StatsOp (valid range is 0x00 to 0x1F). A Capture All Counters for a Port StatsOp must be done prior to using the Read A Captured Counter StatsOp. The counter that is read is defined as follows: **Ingress Counters[1]** / **Egress Counters** 0x00 – InGoodOctetsLo / 0x0E – OutOctetsLo[2] 0x01 – InGoodOctetsHi / 0x0F – OutOctetsHi 0x02 – InBadOctets 0x04 – InUnicast / 0x10 – OutUnicast 0x06 – InBroadcasts / 0x13 – OutBroadcasts 0x07 – InMulticasts / 0x12 – OutMulticasts 0x16 – InPause / 0x15 – OutPause 0x18 – InUndersize 0x19 – InFragments 0x1A – InOversize 0x1B – InJabber 0x1C – In RxErr / 0x11 – Excessive 0x1D – InFCSErr / 0x1E – Collisions 0x05 – Deferred 0x14 – Single 0x17 – Multiple 0x03 – OutFCSErr 0x1F – Late **Histogram Counters[3]** 0x08 – 64Octets 0x09 – 65 to 127Octets 0x0A – 128 to 255Octets 0x0B – 256 to 511Octets 0x0C – 512 to 1023Octets 0x0D – 1024 to MaxOctets

1. If Marvell Header mode is used on Ports 0 to 7 the extra two bytes in the frame are not included in the InGoodOctet nor the InBadOctet counts.
2. OutOctets may not accurately count the bytes transmitted on frames that encounter a collision.
3. If Marvell Header mode is used on Ports 0 to 7 the extra two bytes in the frame are not included in the count before determining which Histogram Counter to increment.

Table 108: Stats Counter Register Bytes 3 & 2
Offset: 0x1E or Decimal 30

Bits	Field	Type	Description
15:8	StatsByte3	RO	Statistics Counter Byte 3. These bits contain bits 31:24 of the last stat counter requested to be read by the CPU (by using the Read a Counter StatsOp— Table 107).
7:0	StatsByte2	RO	Statistics Counter Byte 2. These bits contain bits 23:16 of the last stat counter requested to be read by the CPU (by using the Read a Counter StatsOp — Table 107).

Table 109: Stats Counter Register Bytes 1 & 0
Offset: 0x1F or Decimal 31

Bits	Field	Type	Description
15:8	StatsByte1	RO	Statistics Counter Byte 1. These bits contain bits 15:8 of the last stat counter requested to be read by the CPU (by using the Read a Counter StatsOp — Table 107).
7:0	StatsByte0	RO	Statistics Counter Byte 0. These bits contain bits 7:0 of the last stat counter requested to be read by the CPU (by using the Read a Counter StatsOp — Table 107).

6.6 Switch Global 2 Registers

The devices contain a second set of global registers that effect all the Ethernet ports in the device. Each Global 2 register is 16-bits wide and their bit assignment are shown in Figure 45 below.

Figure 45: Global 2 Register bit Map (Device Addr 0x1C)

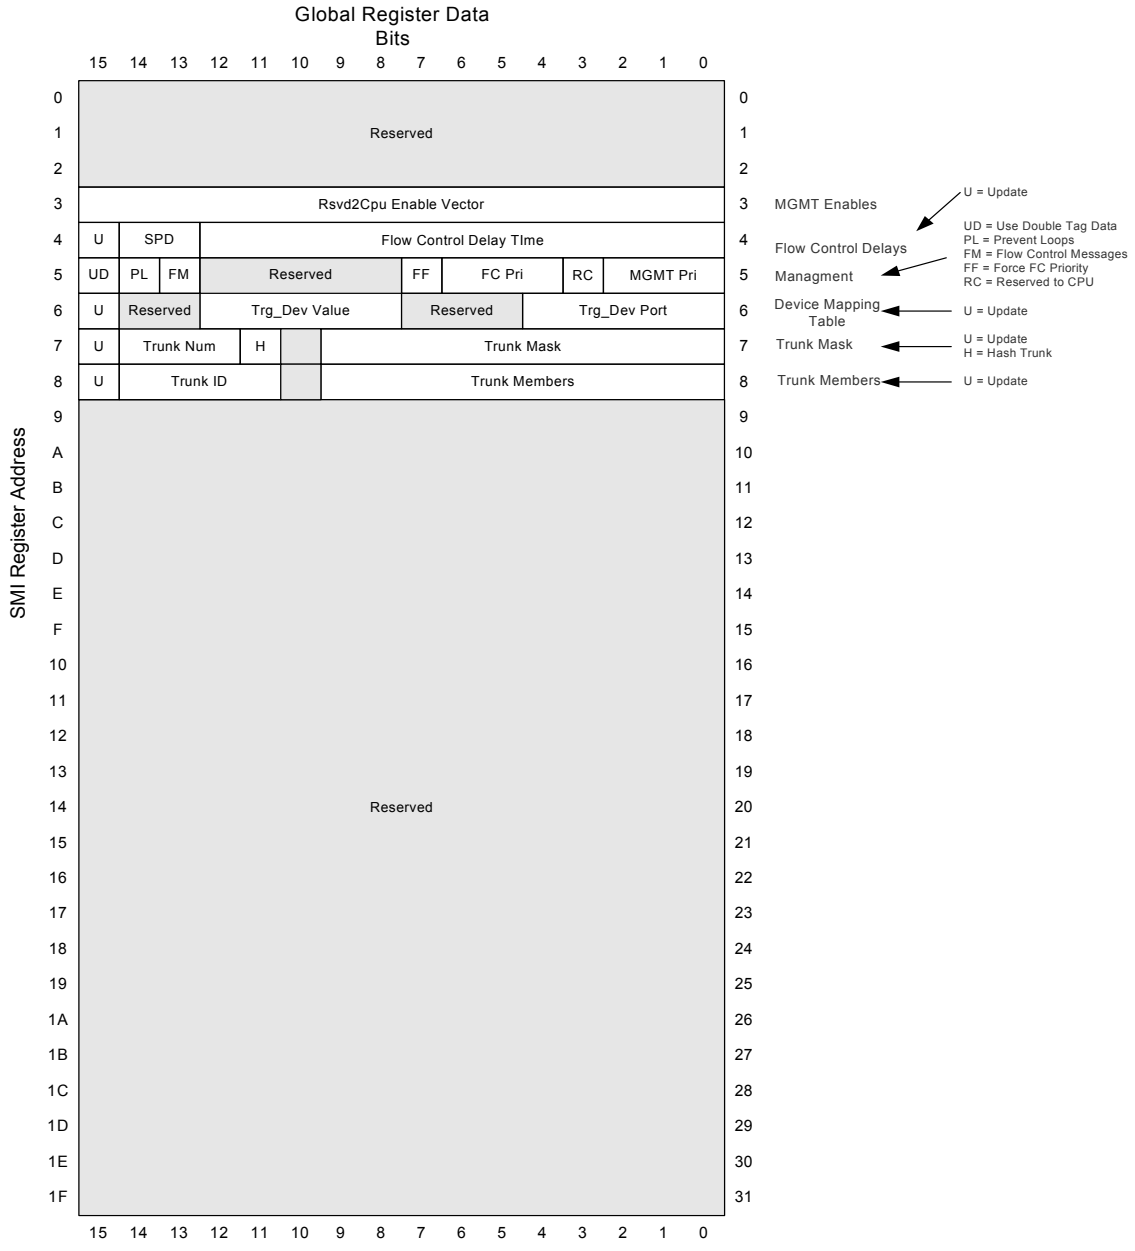

Table 110: MGMT Enable Register
Offset: 0x03 or Decimal 3

Bits	Field	Type	Description
15:0	Rsvd2CPU Enables	RWS	Reserved DA Enables. When the Rsvd2Cpu bit (Table 112) is set to a one the 16 reserved multicast DA addresses whose bit in this register are also set to a one, are treated as MGMT[1] frames. All the reserved DA's take the form 01:80:C2:00:00:0x. When x = 0x0, bit 0 of this register is tested. When x = 0x2 bit 2 of this field is tested and so on with x = 0xF bit 15 of this register is tested. If the tested bit in this register is cleared to a zero, the frame will be treated as a normal (non-MGMT) frame. This register allows some or all of the 16 reserved multicast addresses to be treated as MGMT frames[2]. If the Rsvd2Cpu bit (Table 112) is cleared to a zero these bits will have no effect.

1. MGMT, or management, frames are used for management switch protocols like Spanning Tree (STP) and Link Aggregation (LAC). The switch processes MGMT frames differently (Section 3.4.8).
2. Frames with a DA of 01:80:C2:00:00:01 (the Pause frame DA) are always treated as MAC control frames and cannot be treated as MGMT frames.

Table 111: Flow Control Delay Register
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
15	Update	SC	Update FC Delay TIme data. When this bit is set to a one the data written to bits 12:0 will be loaded into the FC Delay Time register selected by the SPD bits below. After the write has taken place this bit self clears to zero.
14:13	SPEED	RWR	Speed Number. These bits select one of three possible FC Delay Time register for both read and write operations. A write operation occurs if the Update bit is a one. Otherwise a read operation occurs.
12:0	FC Delay Time	RWS	Flow Control Delay Time. These bits are used to cause a MAC to assert Flow Control for the delay amount times 2.048 uSecs. The register used is determined by the Flow Control DSA Tag frame's SPD bits that was directed at this MAC. Three FC Delay Time registers are accessed by using the SPD bits above. SPD 0b00 is assigned for as the Flow Control delay to use when talking to 10 Mbit ports. SPD 0b01 is assigned for 100 Mbit ports and SPD 0b10 is assigned for 1000 Mbit ports (SPD of 0b11 is reserved for future use and should not be accessed). The default value for each of these registers is shown below. SPD 0b00 resets to 0x0191 (401 decimal) SPD 0b01 resets to 0x0029 (41 decimal) SPD 0b10 resets to 0x0005 (5 decimal)

Table 112: Switch Management Register
Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
15	Use DoubleTag Data (88E6095 and 88E6095F only)	RWR	Use Double Tag Data. This bit is used to determine if Double Tag data that is removed from a Double Tag frame is used or ignored when making switching decisions on the frame. 0 = Ignore removed Tag data 1 = Use removed Tag data This bit has no effect if the port's DoubleTag bit is zero (Table 65).
14	Prevent Loops	RWR	Prevent Loops. When a DSA Tag port receives a Forward DSA Tag whose Src_Dev field equals this device's DeviceNumber (Table 106) the following action will be taken depending upon the value of this bit: 0 = The frame will be prevented from going out its original source port as defined by the frame's Src_Port field. 1 = The frame will be discarded. The 1st mode supports external layer 3 and above routers, the 2nd mode prevent loops if one is made and not broken by Port States (Table 65).
13	Flow Control Message	RWR or RWS[1]	Enable Flow Control Messages. When this bit is set to a one DSA Tag Flow Control messaged will be generated when an output queue becomes congested and received DSA Tag Flow Control messages will pause MACs inside this device. When this bit is cleared to a zero DSA Tag Flow Control messages will not be generated and any received will be ignored at the target MAC.
12:8	Reserved	RES	Reserved for future use.
7	ForceFlow ControlPri	RWS	Force Flow Control Priority. When this bit is set to a one the PRI[2:0] bits of generated DSAl Tag Flow Control frames will be set to the value of the FC Pri bits below. When this bit is cleared to a zero generated DSA Tag Flow Control frames will retain the PRI[2:0] bits from the frame that caused the congestion. This bit will have no effect if the FlowControlMessage bit (above) is cleared to a zero.
6:4	FC Pri	RWS to 0x7	Flow Control Priority. These bits are used as the PRI[2:0] bits on generated DSA Tag Flow Control frames if the ForceFlowControlPri bit above is set to a one.

Table 112: Switch Management Register (Continued)
Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
3	Rsvd2CPU		Reserved multicast frames to CPU. This device supports two ways to support protocols that use multicast addresses. The first way is to enter the multicast address into the address database with a MGMT Entry_State, mapping it toward the CPU's port (Table 90). This allows proprietary protocols to be supported while also supporting standard protocols. If multiple address databases are used each multicast address will need to be added to the database for each database. The second way is to set this bit to a one. When this bit is a one frames with a Destination Address in the range 01:80:C2:00:00:0x, regardless of their VLAN membership, will be considered MGMT frames and sent to the port's CPUPort (Table 69) as long as the associated Rsvd2Cpu Enable bit for the frame's DA is also set to a one (Table 110). The MGMT Pri field (below) is used as the priority on these frames.
2:0	MGMT Pri	RWS to 0x7	MGMT Priority. These bits are used as the priority to use on Rsvd2CPU frames (above).

1. The FlowControlMessage bit will set to a one (enabled) if the SW_24P configuration pin is high, the HD_FLOW_DIS configuration pin in high and the FD_FLOW_DIS configuration pin is low at the rising edge of RESETn. This combination of configuration pins enables Cross Chip Flow Control on all Network ports when these ports are in either full or half duplex mode of operation.

Table 113: Device Mapping Table Register
Offset: 0x06 or Decimal 6

Bits	Field	Type	Description
15	Update	SC	Update Target Device Routing data. When this bit is set to a one the data written to bits 3:0 will be loaded into the Target Device entry selected by the Trg_DevValue bits below. After the write has taken place this bit self clears to zero.
14:13	Reserved	RES	Reserved for future use.
12:8	Trg_Dev Value	RWR	Target Device Value. These bits select one of 32 (only 8 for the 88E6092 device) possible Target Device Port register for both read and write operations to the Mapping Table. A write operation occurs if the Update bit is a one. Otherwise a read operation occurs.
7:4	Reserved	RES	Reserved for future use.
3:0	Trg_Dev Port	RWS	Target Device Port number. These bits point to the physical port on this device where From_CPU frames will be routed by using the frame's Trg_Dev as an index into this table (when the Cascade Port, Global Control 2, Offset 0x1C, is set to a value of 0xF). In this way a physical mapping, or Routing Table, of the interconnection of the devices that make up the switch box or boxes in a stack is defined. When a write occurs to this register with the Update bit being a one these bits are written to the Trg_Dev Port selected by the Trg_Dev Value bits. When a write occurs to this register with the Update bit being a zero these bits are not written anywhere (this allow the Trg_Dev Value bits to be written to for read operations). When a read occurs to this register these bits reflect the Target Device Port data found for the entry selected by the Trg_Dev Value bits. The Routing Table is reset to the following values: Trg_Dev Value / Trg_Dev Port / Trg_Dev Value / Trg_Dev Port 0x00 / 0x0 / 0x10 / 0x0 0x01 / 0x1 / 0x11 / 0x1 0x02 / 0x2 / 0x12 / 0x2 0x03 / 0x3 / 0x13 / 0x3 0x04 / 0x4 / 0x14 / 0x4 0x05 / 0x5 / 0x15 / 0x5 0x06 / 0x6 / 0x16 / 0x6 0x07 / 0x7 / 0x17 / 0x7 0x08 / 0x8 / 0x18 / 0x8 0x09 / 0x9 / 0x19 / 0x9 0x0A / 0xA / 0x1A / 0xA 0x0B / 0xF / 0x1B / 0xF 0x0C / 0xF / 0x1C / 0xF 0x0D / 0xF / 0x1D / 0xF 0x0E / 0xF / 0x1E / 0xF 0x0F / 0xF / 0x1F / 0xF

Table 114: Trunk Mask Table Register
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15	Update	SC	Update Trunk Mask data. When this bit is set to a one the data written to bits 10:0 will be loaded into the Trunk Mask selected by the MaskNum bits below. After the write has taken place this bit self clears to zero.
14:12	MaskNum	RWR	Mask Number. These bits select one of eight possible Trunk Mask vectors for both read and write operations. A write operation occurs if the Update bit is a one. Otherwise a read operation occurs.
11	HashTrunk	RWR	Hash DA & SA for TrunkMask selection. Trunk load balancing is accomplished by using the frame's DA and SA fields to access one of eight Trunk Masks. When this bit is set to a one the hashed computed for address table lookups is used for the TrunkMask selection. When this bit is cleared to a zero the lower 3 bits of the frame's DA and SA are XOR'ed together to select the TrunkMask to use.
10:0	TrunkMask	RWS	Trunk Mask bits. Bit 0 controls trunk masking for port 0, bit 1 for port 1, etc. When a write occurs to this register with the Update bit being a one these bits are written to the Trunk Mask selected by the MaskNum bits. When a write occurs to this register with the Update bit being a zero these bits are not written anywhere (this allow the MaskNum bits to be written to for read operations). When a read occurs to this register these bits reflect the Trunk Mask data found for the entry selected by the MaskNum bits. The TrunkMask is reset to all ones for all MaskNum entries.

Table 115: Trunk Members Table Register
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
15	Update	SC	Update Trunk Routing data. When this bit is set to a one the data written to bits 10:0 will be loaded into the Trunk Route selected by the Trunk ID bits below. After the write has taken place this bit self clears to zero.
14:12	Trunk ID	RWR	Trunk Identifier. These bits select one of sixteen possible Trunk ID routing vectors for both read and write operations. A write operation occurs if the Update bit is a one. Otherwise a read operation occurs.
11	Reserved	RES	Reserved for future use.
10:0	Trunk Members	RWR	Trunk Member bits. Bit 0 controls trunk routing for port 0, bit 1 for port 1, etc. When a write occurs to this register with the Update bit being a one these bits are written to the Trunk Member selected by the Trunk ID bits. When a write occurs to this register with the Update bit being a zero these bits are not written anywhere (this allow the Trunk ID bits to be written to for read operations). When a read occurs to this register these bits reflect the Trunk Member data found for the entry selected by the Trunk ID bits.

6.7 EEPROM Programming Format

The devices support an optional external serial EEPROM device for programming its internal registers. The EEPROM data is read in once after RESETn is deasserted unless the Stand Alone Switch Mode is selected (SW_MODE[1:0] = 0b10—Table 15), in which the EEPROM is not read in at all.

The devices support 1K bit (93C46), 2K bit (93C56) or 4K bit (93C66) EEPROM devices (the size of the device is selected by the EE_CS/EE_1K pin at reset – Table 15). The external EEPROM device must be configured in x16 data organization mode.

No matter what device is attached, the EEPROM device is read and processed in the same way:

1. Start at EEPROM address 0x00.
2. Read in the 16-bits of data from the current address, this is called the Command.
3. If the just read in Command is all one's, terminate the serial EEPROM reading process, go to step 8
4. Increment the address by 1 (to the next address). If the Command does not need any data from the EEPROM, process the Command and go to step 2.
5. Read in the 16-bits of data from the next address, this is called RegData, and increment the address by 1.
6. Write RegData to the location or locations defined by the previous Command.
7. Go to step 2.
8. Set the EEInt bit in Global Status to a one (Table 79) generating an Interrupt (if enabled).
9. Done.

The 16-bit Command determines which register or registers inside the devices are updated as follows (refer to Figure 46):

1. A Switch Global register is written to if the upper 8 bits of the Command are 0x7F. The Global space written to is determined by bit 6, the G bit. If G=0, Global 1 at Device Address 0x1B is written to. If G=1 Global 2 at Device Address 0x1C is written to. Bits 4 to 0 determine the register to load.
2. If Bit 15=1 the lower ten Switch MAC port registers are written to (SMI Device Addresses 0x10 to 0x19). A Device Vector (Command bits 14 to 5) is used to determine which port or ports MAC is to be written to. One Command can be used to write the same data to all 10 MACs or to just some of them. Bits 4 to 0 determine the register to load.
3. The devices contain 11 ports so the previously described command cannot be used to reach MAC 10. Instead, a separate command is available for this. If the upper 10 bits of the Command is 0x7CC, then Port 10's MAC will be written to if a Primer command has been executed first. Bits 4 to 0 determine the register to load. The Primer for MAC 10 Command (0x731A) only needs to be executed once, but it must be executed before any Commands to write to Port 10's MAC.
4. When the upper 3 bits of the Command is 0x0, the lower eight PHY registers are written to (SMI Device Addresses 0x00 to 0x07). A Device Vector (Command bits 12 to 5) is used to determine which port or ports PHY is to be written to. One Command can be used to write the same data to all 8 PHYs or to just one of them. Bits 4 to 0 determine the register to load. The PPU must be disabled or the writes to the PHYs will not occur.
5. The devices contain 11 ports so the previously described command cannot be used to reach PHYs 8 to 10. Instead, separate commands are available for this. If the upper 10 bits of the Command is 0x7C0, then Port 8's PHY will be written to if a Primer command has been executed first. Bits 4 to 0 determine the register to load. The Primer for PHY 8 Command (0x7008) only needs to be executed once, but it must be executed before any Commands to write to Port 8's PHY. Port 9's PHY and Port 10's PHY have separate Commands to write to the PHY (0x7C4x and 0x7C8x, respectively) and separate Primer commands (0x7109 and 0x720A respectively). The PPU must be disable or the writes to the PHYs will not occur.
6. The EEPROM updates the internal register settings and then releases an internal Reset when it reads a Command of all 1's (a Halt Command). This allows register setting changes to be made prior to letting packets flow through the switch (required to prevent VLAN leaks); however, the ATU and VTU cannot be loaded with the data from the EEPROM unless the internal Reset is released. A Command of 0x7D01 will release the internal Reset and continue EEPROM processing so the ATU or VTU can be loaded from the EEPROM. This will also allow packets to start flowing through the switch so placement of this Command is important.

Figure 46: EEPROM Data Format

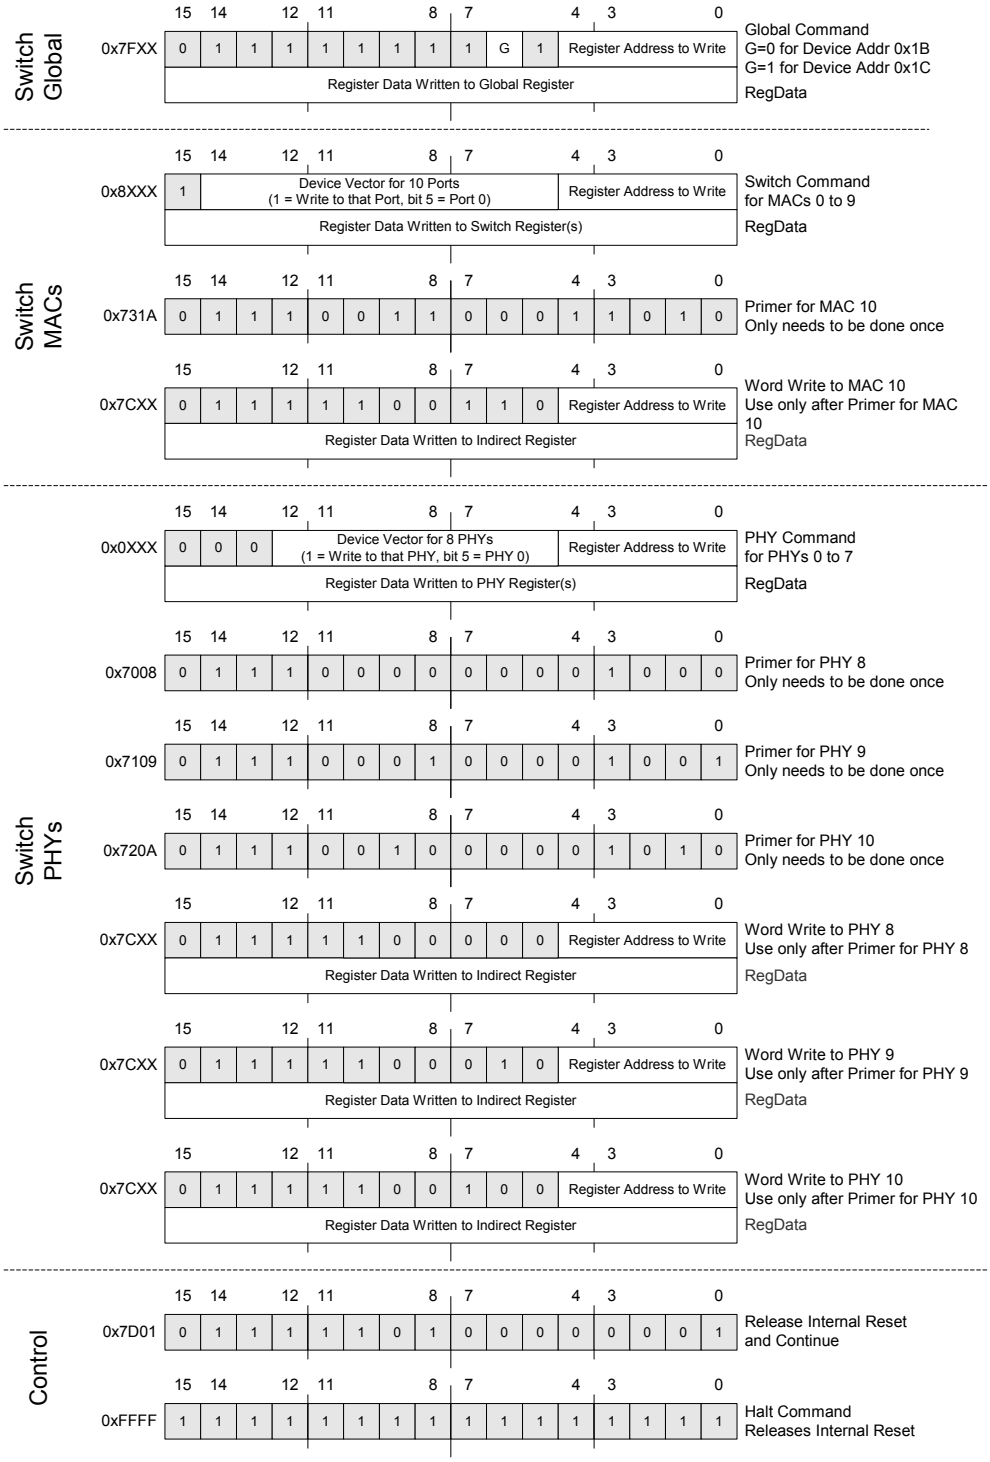

Section 7. PHY Register Description

The device contains eight physical layer devices (PHYs). These devices are accessible using SMI device addresses 0x00 to 0x07 (or 0x10 to 0x14) depending upon the value of the P8_OUTD[3:0]/P8_MODE[3:0] pins at reset. The PHYs are fully IEEE 802.3 compliant including their register interface.

The PHYs in the devices are identical to those in the Marvell® 88E3082 Octal Transceiver.

Table 116: PHY Register Map

Description	Offset Hex	Offset Decimal	Page Number
PHY Control Register	0x00	0	page 234
PHY Status Register	0x01	1	page 237
PHY Identifier	0x02	2	page 239
PHY Identifier	0x03	3	page 239
Auto-Negotiation Advertisement Register	0x04	4	page 240
Link Partner Ability Register (Base Page)	0x05	5	page 242
Link Partner Ability Register (Next Page)	0x05	5	page 243
Auto-Negotiation Expansion Register	0x06	6	page 244
Next Page Transmit Register	0x07	7	page 245
Link Partner Next Page Register	0x08	8	page 246
Reserved Registers	0x09-0x0F	9 - 15	page 246
PHY Specific Control Register I	0x10	16	page 247
PHY Specific Status Register	0x11	17	page 250
PHY Interrupt Enable	0x12	18	page 252
PHY Interrupt Status	0x13	19	page 254
PHY Interrupt Port Summary	0x14	20	page 256
PHY Receive Error Counter	0x15	21	page 257
LED Parallel Select Register	0x16	22	page 258
LED Stream Select for Serial LEDs (Global Register) - 88E6095F Only	0x17	23	page 260
PHY LED Control Register	0x18	24	page 262
PHY Manual LED Override Register	0x19	25	page 263
VCT™ Control Register	0x1A	26	page 265
VCT™ Status Register	0x1B	27	page 266
PHY Specific Control Register II	0x1C	28	page 267
Reserved Registers	0x1D to 0x1F	29 - 31	page 267

Table 117: PHY Control Register
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	SWReset	R/W, SC	0x0	0	PHY Software Reset Writing a 1 to this bit causes the PHY state machines to be reset. When the reset operation is done, this bit is cleared to 0 automatically. The reset occurs immediately. 0 = Normal operation 1 = PHY reset
14	Loopback	R/W	0x0	Retain	Enable Loopback Mode When loopback mode is activated, the transmitter data presented on TXD is looped back to RXD internally. The PHY has to be in forced 10 or 100 Mbps mode. Auto-Negotiation must be disabled. 0 = Disable loopback 1 = Enable loopback
13	SpeedLSB	R/W	See Descr	Update	Speed Selection (LSB) The speed selection (MSB or LSB) determines the forced speed if Auto-Negotiation is disabled. If Auto-Negotiation is enabled, the speed ability advertisement is located in Register 4, and 0.13 and 0.6 have a "don't care" definition. A write to this register bit has no effect unless any one of the following also occurs: Software reset is asserted (bit 15) or Power down (bit 11) transitions from power down to normal operation. 0 = 10 Mbps 1 = 100 Mbps After hardware reset, the default value is as follows: 88E6092/88E6095 device = 1 88E6095F device = Determined by P[7:0]_CONFIG

Table 117: PHY Control Register (Continued)
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
12	AnegEn	R/W	See Descr	Update	Auto-Negotiation Enable If Auto-Negotiation is enabled (0.12 = 1), the speed and duplex ability advertisement are located in Register 4, and register bits 0.13, 0.6, and 0.8 have a "don't care" definition. A write to this register bit has no effect unless any one of the following also occurs: Software reset is asserted (bit 15, above), Power down (bit 11, below), or the PHY transitions from power down to normal operation. If Auto-Negotiation is already enabled, any link drop event will cause the Auto-Negotiation process to start again. 0 = Disable Auto-Negotiation Process 1 = Enable Auto-Negotiation Process After hardware reset the default value is as follows:88E6092/88E6095 device 0.12 = 1 88E6095F device = Determined by P[7:0]_CONFIG
11	PwrDwn	R/W	0x0	Retain	Power Down Mode When the port is switched from power down to normal operation, software reset and restart Auto-Negotiation are performed even when bits Reset (bit 15, above) and Restart Auto-Negotiation (bit 9, below) are not set by the user. 0 = Normal operation 1 = Power down
10	Isolate	RO	Always 0	Always 0	Isolate Mode Will always be 0. The Isolate function is not available, since full MII is not implemented. 0 = Normal operation
9	RestartAneg	R/W, SC	0x0	Self Clear	Restart Auto-Negotiation Auto-Negotiation automatically restarts after hardware or software reset regardless of whether or not the restart bit is set. 0 = Normal operation 1 = Restart Auto-Negotiation Process

Table 117: PHY Control Register (Continued)
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
8	Duplex	R/W	See Descr	Update	Duplex Mode Selection This bit determines full-duplex or half-duplex mode if Auto-Negotiation is disabled. If Auto-Negotiation is enabled, the speed and duplex ability advertisement are located in Register 4, and Register bit 0.8 have a "don't care" definition. A write to these registers has no effect unless any one of the following also occurs: Software reset is asserted (bit 15), Power down (bit 11), or transitions from power down to normal operation. 0 = Half-duplex 1 = Full-duplex After hardware reset this bit is set as follows: If copper mode, this bit is 0.If 100BASE-FX mode, this bit is 1.
7	ColTest	RO	Always 0	Always 0	Collision Test Mode Will always be 0. The Collision test is not available, since full MII is not implemented. 0 = Disable COL signal test
6	SpeedMSB	RO	Always 0	Always 0	Speed Selection Mode (MSB) Will always be 0. 0 = 100 Mbps or 10 Mbps
5:0	Reserved	RO	Always 0	Always 0	Will always be 0.

Table 118: PHY Status Register
Offset: 0x01 (Hex), or 1 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	100T4	RO	Always 0	Always 0	100BASE-T4 This protocol is not available. 0 = PHY not able to perform 100BASE-T4
14	100FDX	RO	Always 1	Always 1	100BASE-T and 100BASE-X full-duplex 1 = PHY able to perform full-duplex
13	100HDX	RO	Always 1	Always 1	100BASE-T and 100BASE-X half-duplex 1 = PHY able to perform half-duplex
12	10FDX	RO	Always 1	Always 1	10BASE-T full-duplex 1 = PHY able to perform full-duplex
11	10HPX	RO	Always 1	Always 1	10BASE-T half-duplex 1 = PHY able to perform half-duplex
10	100T2FDX	RO	Always 0	Always 0	100BASE-T2 full-duplex. This protocol is not available. 0 = PHY not able to perform full-duplex
9	100T2HDX	RO	Always 0	Always 0	100BASE-T2 half-duplex This protocol is not available. 0 = PHY not able to perform half-duplex
8	ExtdStatus	RO	Always 0	Always 0	Extended Status 0 = No extended status information in Register 15
7	Reserved	RO	Always 0	Always 0	Must always be 0.
6	MFPreSup	RO	Always 1	Always 1	MF Preamble Suppression Mode Must be always 1. 1 = PHY accepts management frames with preamble suppressed
5	AnegDone	RO	0x0	0	Auto-Negotiation Complete 0 = Auto-Negotiation process not completed 1 = Auto-Negotiation process completed
4	RemoteFault	RO, LH	0x0	0	Remote Fault Mode 0 = Remote fault condition not detected 1 = Remote fault condition detected
3	AnegAble	RO	Always 1	Always 1	Auto-Negotiation Ability Mode 1 = PHY able to perform Auto-Negotiation

Table 118: PHY Status Register (Continued)
Offset: 0x01 (Hex), or 1 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
2	Link	RO, LL	0x0	0	Link Status Mode This register indicates when link was lost since the last read. For the current link status, either read this register back-to-back or read RTLink in Table 128 on page 250. 0 = Link is down 1 = Link is up
1	JabberDet	RO, LH	0x0	0	Jabber Detect 0 = Jabber condition not detected 1 = Jabber condition detected
0	ExtdReg	RO	Always 1	Always 1	Extended capability mode. 1 = Extended register capabilities

Table 119: PHY Identifier
Offset: 0x02 (Hex), or 2 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	Organization-ally Unique Identifier Bits 3:18	RO	0x0141	0x0141	Marvell® OUI is 0x005043 0000 0000 0101 0000 0100 0011 ^ ^ bit 1.....bit 24 Register 2.[15:0] show bits 3 to 18 of the OUI. 0000000101000001 ^ ^ bit 3.....bit 18

Table 120: PHY Identifier
Offset: 0x03 (Hex), or 3 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:10	OUI LSb	RO	Always 000011	Always 000011	Organizationally Unique Identifier bits 19:24 00 0011 ^.....^ bit 19...bit 24
9:4	ModelNum	RO	Always 001000	Always 001000	Model Number = 001000
3:0	RevNum	RO	Varies	Varies	Revision Number Contact Marvell® FAEs for information on the device revision number.

Table 121: Auto-Negotiation Advertisement Register
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	AnegAd NxtPage	R/W	0x0	Retain	Next Page 0 = Not advertised 1 = Advertise Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
14	Ack	RO	Always 0	Always 0	Must be 0.
13	AnegAd ReFault	R/W	0x0	Retain	Remote Fault Mode 0 = Do not set Remote Fault bit 1 = Set Remote Fault bit Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
12:11	Reserved	R/W	0x0	Retain	Must be 00. Reserved bits are R/W to allow for forward compatibility with future IEEE standards. Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
10	AnegAd Pause	R/W	0x0	Retain	Pause Mode 0 = MAC PAUSE not implemented 1 = MAC PAUSE implemented Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
9	AnegAd 100T4	R/W	0x0	Retain	100BASE-T4 mode 0 = Not capable of 100BASE-T4 Must be 0.
8	AnegAd 100FDX	R/W	0x1	Retain	100BASE-TX full-duplex Mode 0 = Not advertised 1 = Advertise Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.

Table 121: Auto-Negotiation Advertisement Register (Continued)
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
7	AnegAd 100HDX	R/W	0X1	Retain	100BASE-TX half-duplex Mode 0 = Not advertised 1 = Advertise Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
6	AnegAd 10FDX	R/W	0X1	Retain	10BASE-TX full-duplex Mode 0 = Not advertised 1 = Advertise Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
5	AnegAd 10HDX	R/W	0X1	Retain	10BASE-TX half-duplex Mode 0 = Not advertised 1 = Advertise Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 117 on page 234) or link goes down.
4:0	AnegAd Selector	R/W	Always 0x01	Always 0x01	Selector Field Mode 00001 = 802.3

Table 122: Link Partner Ability Register (Base Page)
Offset: 0x05 (Hex), or 5 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	LPNxt Page	RO	0x0	0	Next Page Mode Base page will be overwritten if next page is received and if Reg8NxtPg (Table 127 on page 247) is disabled. When Reg8NxtPg (Table 127 on page 247) is enabled, then next page is stored in the Link Partner Next Page register (Table 126 on page 246), and the Link Partner Ability Register (Table 122 on page 242) holds the base page. Received Code Word Bit 15 0 = Link partner not capable of next page 1 = Link partner capable of next page
14	LPAck	RO	0x0	0	Acknowledge Received Code Word Bit 14 0 = Link partner did not receive code word 1 = Link partner received link code word
13	LPRemote Fault	RO	0x0	0	Remote Fault Received Code Word Bit 13 0 = Link partner has not detected remote fault 1 = Link partner detected remote fault
12:5	LPTechAble	RO	0x00	0x00	Technology Ability Field Received Code Word Bit 12:5
4:0	LPSelector	RO	00000	00000	Selector Field Received Code Word Bit 4:0

Table 123: Link Partner Ability Register (Next Page)
Offset: 0x05 (Hex), or 5 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	LPNxtPage	RO	--	--	Next Page Mode Base page will be overwritten if next page is received and if Reg8NxtPg (Table 127 on page 247) is disabled. When Reg8NxtPg (Table 127 on page 247) is enabled, then next page is stored in the Link Partner Next Page register (Table 126 on page 246), and Link Partner Ability Register (Table 122 on page 242) holds the base page. Received Code Word Bit 15
14	LPAck	RO	--	--	Acknowledge Received Code Word Bit 14
13	LPMessage	RO	--	--	Message Page Received Code Word Bit 13
12	LPack2	RO	--	--	Acknowledge 2 Received Code Word Bit 12
11	LPToggle	RO	--	--	Toggle Received Code Word Bit 11
10:0	LPData	RO	--	--	Message/Unformatted Field Received Code Word Bit 10:0

Table 124: Auto-Negotiation Expansion Register
Offset: 0x06 (Hex), or 6 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:5	Reserved	RO	Always 0x000	Always 0x000	Reserved. Must be 00000000000. The Auto-Negotiation Expansion Register is not valid until the AnegDone (Table 118 on page 237) indicates completed.
4	ParFaultDet	ROC/ LH	0x0	0x0	Parallel Detection Level 0 = A fault has not been detected via the Parallel Detection function 1 = A fault has been detected via the Parallel Detection function
3	LPNxtPg Able	RO	0x0	0x0	Link Partner Next Page Able 0 = Link Partner is not Next Page able 1 = Link Partner is Next Page able
2	LocalNxtPg Able	RO	Always 0x1	Always 0x1	Local Next Page Able This bit is equivalent to AnegAble (Table 118 on page 237). 1 = Local Device is Next Page able
1	RxNewPage	RO/LH	0x0	0x0	Page Received 0 = A New Page has not been received 1 = A New Page has been received
0	LPAnegAble	RO	0x0	0x0	Link Partner Auto-Negotiation Able 0 = Link Partner is not Auto-Negotiation able 1 = Link Partner is Auto-Negotiation able

Table 125: Next Page Transmit Register
Offset: 0x07 (Hex), or 7 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	TxNxtPage	R/W	0x0	0x0	A write to the Next Page Transmit Register implicitly sets a variable in the Auto-Negotiation state machine indicating that the next page has been loaded. Transmit Code Word Bit 15
14	Reserved	RO	0x0	0x0	Reserved Transmit Code Word Bit 14
13	TxMessage	R/W	0x1	0x1	Message Page Mode Transmit Code Word Bit 13
12	TxAck2	R/W	0x0	0x0	Acknowledge2 Transmit Code Word Bit 12
11	TxToggle	RO	0x0	0x0	Toggle Transmit Code Word Bit 11
10:0	TxData	R/W	0x001	0x001	Message/Unformatted Field Transmit Code Word Bit 10:0

Table 126: Link Partner Next Page Register
Offset: 0x08 (Hex), or 8 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	RxNxtPage	RO	0x0	0x0	If Reg8NxtPg (Table 127 on page 247) is enabled, then next page is stored in the Link Partner Next Page register (Table 126 on page 246); otherwise, theLink Partner Next Page register (Table 126 on page 246) is cleared to all 0's. Received Code Word Bit 15
14	RxAck	RO	0x0	0x0	Acknowledge Received Code Word Bit 14 0 = Link partner not capable of next page 1 = Link partner capable of next page
13	RxMessage	RO	0x0	0x0	Message Page Received Code Word Bit 13
12	RxAck2	RO	0x0	0x0	Acknowledge 2 Received Code Word Bit 12
11	RxToggle	RO	0x0	0x0	Toggle Received Code Word Bit 11
10:0	RxData	RO	0x000	0x000	Message/Unformatted Field Received Code Word Bit 10:0

Note

Registers 0x09 through 0x0F (hexadecimal (9 through 15 decimal) are reserved. Do not read or write to these registers.

Table 127: PHY Specific Control Register
Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Enable DTE Detect	R/W	0x0	Retain	Enable DTE Detect 0 = Disable DTE Detect 1 = Enable DTE Detect
14	EDet	R/W	See Descr.	Retain	Energy Detect 0 = Disable 1 = Enable with sense and pulse Enable with sense only is not supported Enable Energy Detect takes on the appropriate value defined by the CONFIG_B pin at hardware reset. See Table 4 for details.
13	DisNLP Check	R/W	0x0	0x0	Disable Normal Linkpulse Check Linkpulse check and generation disable have no effect, if Auto-Negotiation is enabled locally. 0 = Enable linkpulse check 1 = Disable linkpulse check
12	Reg8NxtPg	R/W	0x0	0x0	Enable the Link Partner Next Page register (Table 123 on page 243) to store Next Page. If set to store next page in the Link Partner Next Page register (Table 123 on page 243), then 802.3u is violated to emulate 802.3ab. 0 = Store next page in the Link Partner Ability Register (Base Page) register (Table 122 on page 242). 1 = Store next page in the Link Partner Next Page register (Table 123 on page 243)
11	DisNLPGen	R/W	0x0	0x0	Disable Linkpulse Generation. Linkpulse check and generation disable have no effect, when Auto-Negotiation is enabled locally. 0 = Enable linkpulse generation 1 = Disable linkpulse generation
10	ForceLink	R/W	0x0	0x0	Force Link Good When link is forced to be good, the link state machine is bypassed and the link is always up. 0 = Normal operation 1 = Force link good
9	DisScrambler	R/W	See Descr[1]	Retain	Disable Scrambler For 100BASE-TX, the scrambler is enabled by default and can be disabled by writing to this bit. For all other modes the scrambler is disabled regardless of the state of this bit. 0 = Enable scrambler 1 = Disable scrambler

Table 127: PHY Specific Control Register (Continued)
Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
8	DisFEFI	R/W	CONFIG_A	Retain	Disable FEFI In 100BASE-FX mode, Disable FEFI takes on the appropriate value defined by the CONFIG_A pin at hardware reset. FEFI is automatically disabled regardless of the state of this bit if copper mode is selected. 0 = Enable FEFI 1 = Disable FEFI
7	ExtdDistance	R/W	0x0	0x0	Enable Extended Distance When using cable exceeding 100 meters, the 10BASE-T receive threshold must be lowered in order to detect incoming signals. 0 = Normal 10BASE-T receive threshold 1 = Lower 10BASE-T receive threshold
6	TPSelect	R/W	See Descr	Update	(Un)Shielded Twisted Pair This setting can be changed by writing to this bit followed by software reset. 0 = Unshielded Twisted Pair - default 1 = Shielded Twisted Pair (Un)Shielded Twisted Pair selection applies to the copper modes of operation only.
5:4	AutoMDI[X]	R/W	See Descr	Update	MDI/MDIX Crossover This setting can be changed by writing to these bits followed by software reset. 00 = Transmit on pins RXP/RXN, Receive on pins TXP/TXN 01 = Transmit on pins TXP/TXN, Receive on pins RXP/RXN 1x = Enable Automatic Crossover The default MDI/MDIX crossover setting is determined by CONFIG_B. If Auto-Crossover is enabled by CONFIG_B, then by default bits 5:4 = 11. MDI/MDIX crossover applies to the copper modes of operation only.

Table 127: PHY Specific Control Register (Continued)
Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
3:2	RxFIFO Depth	R/W	0x0	0x0	Receive FIFO Depth 00 = 4 Bytes 01 = 6 Bytes 10 = 8 Bytes 11 = Reserved
1	AutoPol	R/W	0x0	00	Polarity Reversal If Automatic polarity is disabled, then the polarity is forced to be normal in 10BASE-T mode. Polarity reversal has no effect in 100BASE-TX mode. 0 = Enable automatic polarity reversal 1 = Disable automatic polarity reversal
0	DisJabber	R/W	0x0	00	Disable Jabber Jabber has no effect in full-duplex or in 100BASE-X mode. 0 = Enable jabber function 1 = Disable jabber function

Table 128: PHY Specific Status Register
Offset: 0x11 (Hex), or 17 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	DTE Detect Status	RO	0x0	0x0	DTE Status 0 = DTE not detected 1 = DTE detected
14	ResSpeed	RO	See Descr	Retain	Resolved Speed The values are updated after the completion of Auto-Negotiation and should be read if link is established. This bit is valid only if the resolved bit 11 is set to 1. 0 = 10 Mbps 1 = 100 Mbps If Auto-Negotiation is disabled the forced speed is already determined, so this bit is a "don't care" definition.
13	ResDuplex	RO	See Descr	Retain	Resolved Duplex Mode This value is updated after the completion of Auto-Negotiation and should be read if link is established. This bit is valid only after the resolved bit 11 is set to 1. 0 = Half-duplex 1 = Full-duplex If Auto-Negotiation is disabled the duplex is already determined, so this bit is a "don't care" definition.
12	RcvPage	RO, LH	0x0	0x0	Page Receive Mode 0 = Page not received 1 = Page received
11	Resolved	RO	0x0	0x0	Speed and Duplex Resolved. Speed and duplex bits (14 and 13) are valid only after the Resolved bit is set. The Resolved bit is set after the completion of Auto-Negotiation and should be read if link is established. 0 = Not resolved 1 = Resolved If Auto-Negotiation is disabled the speed and duplex is already determined, so this bit is a "don't care" definition.
10	RTLink	RO	0x0	0x0	Link (real time) 0 = Link down 1 = Link up
9:7	Reserved	RES	Always 000	Always 000	Always 000.
6	MDI/MDIX	RO	0x0	0x0	MDI/MDIX Crossover Status 0 = Transmit on pins TXP/TXN, Receive on pins RXP/RXN 1 = Transmit on pins RXP/RXN, Receive on pins TXP/TXN

Table 128: PHY Specific Status Register (Continued)
Offset: 0x11 (Hex), or 17 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
5	Reserved	RES	Always 0	Always 0	Always 0.
4	Sleep	RO	0x0	0x0	Energy Detect Status 0 = Chip is not in sleep mode (Active) 1 = Chip is in sleep mode (No wire activity)
3:2	Reserved	RES	Always 00	Always 00	Always 00.
1	RTPolarity	RO	0x0	0x0	Polarity (real time) 0 = Normal 1 = Reversed
0	RTJabber	RO	0x0	Retain	Jabber (real time) 0 = No Jabber 1 = Jabber

Table 129: PHY Interrupt Enable
Offset: 0x12 (Hex), or 18 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	DTE Detect State Changed Interrupt Enable	R/W	0x0	Retain	DTE Detect State Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
14	SpeedIntEn	R/W	0x0	Retain	Speed Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
13	DuplexIntEn	R/W	0x0	Retain	Duplex Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
12	RxPageIntEn	R/W	0x0	Retain	Page Received Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
11	AnegDone IntEn	R/W	0x0	Retain	Auto-Negotiation Completed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
10	LinkIntEn	R/W	0x0	Retain	Link Status Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
9	SymErrIntEn	R/W	0x0	Retain	Symbol Error Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
8	FlsCrsIntEn	R/W	0x0	Retain	False Carrier Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
7	FIFOErrInt	R/W	0x0	Retain	FIFO Over/Underflow Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
6	MDI[x]IntEn	R/W	0x0	0x0	MDI/MDIX Crossover Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
5	Reserved	RES	0x0	Retain	Must be 0.
4	EDetIntEn	R/W	0x0	Retain	Energy Detect Interrupt Enable 0 = Disable 1 = Enable
3:2	Reserved	RES	0x0	Retain	Must be 00.

Table 129: PHY Interrupt Enable (Continued)
Offset: 0x12 (Hex), or 18 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
1	PolarityIntEn	R/W	0x0	Retain	Polarity Changed Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable
0	JabberIntEn	R/W	0x0	Retain	Jabber Interrupt Enable 0 = Interrupt disable 1 = Interrupt enable

Table 130: PHY Interrupt Status
Offset: 0x13 (Hex), or 19 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	DTE Detect State Changed Interrupt	RO, LH	0x0	0x0	DTE detect state changed interrupt. 0 = DTE detect state not changed 1 = DTE detect state changed
14	SpeedInt	RO, LH	0x0	0x0	Speed Changed 0 = Speed not changed 1 = Speed changed
13	DuplexInt	RO, LH	0x0	0x0	Duplex Changed 0 = Duplex not changed 1 = Duplex changed
12	RxPageInt	RO, LH	0x0	0x0	0 = Page not received 1 = Page received
11	AnegDoneInt	RO, LH	0x0	0x0	Auto-Negotiation Completed 0 = Auto-Negotiation not completed 1 = Auto-Negotiation completed
10	LinkInt	RO, LH	0x0	0x0	Link Status Changed 0 = Link status not changed 1 = Link status changed
9	SymErrInt	RO, LH	0x0	0x0	Symbol Error 0 = No symbol error 1 = Symbol error
8	FlsCrsInt	RO, LH	0x0	0x0	False Carrier 0 = No false carrier 1 = False carrier
7	FIFOErrInt	RO, LH	0x0	0x0	FIFO Over /Underflow Error 0 = No over/underflow error 1 = Over/underflow error
6	MDIMDIXInt	RO, LH	0x0	0x0	MDI/MDIX Crossover Changed 0 = MDI/MDIX crossover not changed 1 = MDI/MDIX crossover changed
5	Reserved	RO	Always 0	Always 0	Always 0
4	EDetChg	RO, LH	0x0	0x0	Energy Detect Changed 0 = No Change 1 = Changed
3:2	Reserved	RO	Always 00	Always 00	Always 00

Table 130: PHY Interrupt Status (Continued)
Offset: 0x13 (Hex), or 19 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
1	PolarityInt	RO	0x0	0x0	Polarity Changed 0 = Polarity not changed 1 = Polarity changed
0	JabberInt	RO, LH	0x0	0x0	Jabber Mode 0 = No Jabber 1 = Jabber

Table 131: PHY Interrupt Port Summary (Global[1])
Offset: 0x14 (Hex), or 20 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:8	Reserved	RO	0x0	0x0	Must be 00000000.
7	Port7Int Active	RO	0x0	0x0	Port 7 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
6	Port6Int Active	RO	0x0	0x0	Port 6 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
5	Port5Int Active	RO	0x0	0x0	Port 5 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
4	Port4Int Active	RO	0x0	0x0	Port 4 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
3	Port3Int Active	RO	0x0	0x0	Port 3 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
2	Port2Int Active	RO	0x0	0x0	Port 2 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt

Table 131: PHY Interrupt Port Summary (Global[1]) (Continued)
Offset: 0x14 (Hex), or 20 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
1	Port1Int Active	RO	0x0	0x0	Port 1 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt
0	Port0Int Active	RO	0x0	0x0	Port 0 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 0 = Port does not have active interrupt 1 = Port has active interrupt

1. A Global register is accessible for writing from any port.

Table 132: Receive Error Counter
Offset: 0x15 (Hex), or 21 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	RxErrCnt	RO	0x0000	0x0000	Receive Error Count This register counts receive errors on the media interface. When the maximum receive error count reaches 0xFFFF, the counter will roll over.

Table 133: LED Parallel Select Register (bits 11:0 are Global[1] bits)
Offset: 0x16 (Hex), or 22 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:12	DTE Detect Status Drop	R/W	0x4	Retain	DTE detect status drop. Once the devices no longer detect the link partner's DTE filter, the devices will wait a period of time before clearing the DTE detection status bit (17.15). The wait time is 5 seconds multiplied by the value of these bits. Example: 5*0x4 = 20 seconds
11:8	LED2	R/W	LED[1:0] 00 = LINK 01 = LINK 10 = LINK/RX 11= LINK/ ACT	Retain	LED2 Control. This is a global[1] setting. The parallel LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 0000 = COLX 0001 = ERROR 0010 = DUPLEX 0011 = DUPLEX/COLX 0100 = SPEED 0101 = LINK 0110 = TX 0111 = RX 1000 = ACT 1001 = LINK/RX 1010 = LINK/ACT 1011 = ACT (BLINK mode) 1100 = TX (Blink Mode) 1101 = RX (Blink Mode) 1110 = COLX (Blink Mode) 1111 = Force to 1 (inactive)
7:4	LED1	R/W	LED[1:0] 00 = RX 01 = ACT 10 = TX 11= DUPLEX /COLX	Retain	LED1 Control. This is a global[1] setting. The parallel LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 0000 = COLX 0001 = ERROR 0010 = DUPLEX 0011 = DUPLEX/COLX 0100 = SPEED 0101 = LINK 0110 = TX 0111 = RX 1000 = ACT 1001 = LINK/RX 1010 = LINK/ACT 1011 = ACT (BLINK mode) 1100 = TX (Blink Mode) 1101 = RX (Blink Mode) 1110 = COLX (Blink Mode) 1111 = Force to 1 (inactive)

Table 133: LED Parallel Select Register (Continued) (bits 11:0 are Global[1] bits)
Offset: 0x16 (Hex), or 22 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
3:0	LED0	R/W	LED[1:0] 00 = TX 01 = SPEED 10 = SPEED 11= SPEED	Retain	LED0 Control. This is a global[1] setting. The parallel LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 0000 = COLX 0001 = ERROR 0010 = DUPLEX 0011 = DUPLEX/COLX 0100 = SPEED 0101 = LINK 0110 = TX 0111 = RX 1010 = LINK/ACT 1011 = ACT (BLINK mode) 1100 = TX (Blink Mode) 1101 = RX (Blink Mode) 1110 = COLX (Blink Mode) 1111 = Force to 1 (inactive)

1. Global register bits are used to control features and functions that are common to all ports in the device.

Table 134: LED Stream Select for Serial LEDs (Global Register) - 88E6095F Only
Offset: 0x17 (Hex), or 23 (Decimal)

Bits	Function	Mode	HW Rst	SW Rst	Description
15:14	LEDLnkActy	R/W	LED[1:0] 00 = Single 01 = Off 10 = Off 11 = Off	Retain	LED Link Activity The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
13:12	LEDRcvLnk	R/W	LED[1:0] 00 = Off 01 = Single 10 = Off 11 = Off	Retain	LED Receive Link The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
11:10	LEDActy	R/W	LED[1:0 00 = Off 01 = Off 10 = Single 11 = Off	Retain	LED Activity The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
9:8	LEDRcv	R/W	LED[1:0] 00 = Off 01 = Off 10 = Off 11 = Single	Retain	LED Receive The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
7:6	LEDTx	R/W	LED[1:0] 00 = Off 01 = Single 10 = Off 11 = Single	Retain	Transmit The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
5:4	LEDLnk	R/W	LED[1:0] 00 = Off 01 = Off 10 = Single 11 = Single	Retain	Link The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single

Table 134: LED Stream Select for Serial LEDs (Global Register) - 88E6095F Only (Continued)
Offset: 0x17 (Hex), or 23 (Decimal)

Bits	Function	Mode	HW Rst	SW Rst	Description
3:2	LEDSpd	R/W	11	Retain	Speed The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single
1:0	LEDDx/ COLX	R/W	LED[1:0] 00 = Single 01 = Off 10 = Off 11 = Off	Retain	LED Duplex/ COLX The serial LED settings take on the appropriate value defined by the CONFIG_A pin at hardware reset. 00 = Off 01 = Reserved 10 = Dual 11 = Single

Table 135: PHY LED Control Register (bits 14:0 are Global[1] bits)
Offset: 0X18 (Hex), 0r 24 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RO	Always 0	Always 0	Must be 0.
14:12	PulseStretch	R/W	0x4	Retain	Pulse stretch duration. This is a global setting. Default Value = 100. 000 = No pulse stretching 001 = 21 ms to 42 ms 010 = 42 ms to 84 ms 011 = 84 ms to 170 ms 100 = 170 ms to 340 ms 101 = 340 ms to 670 ms 110 = 670 ms to 1.3s 111 = 1.3s to 2.7s
11:9	BlinkRate	R/W	0x1	Retain	Blink Rate. This is a global setting. Default Value = 001 000 = 42 ms 001 = 84 ms 010 = 170 ms 011 = 340 ms 100 = 670 ms 101 to 111 = Reserved
8:6	SrStrUpdate	R/W	0x2	Retain	Serial Stream Update. This is a global setting. 000 = 10 ms 001 = 21 ms 010 = 42 ms 011 = 84 ms 100 = 170 ms 101 = 340 ms 110 to 111 = Reserved
5:4	Duplex	R/W	LED[1:0] 00 = Off 01 = Single 10 = Single 11 = Single	Retain	00 = Off 01 = Reserved 10 = Dual 11 = Single
3:2	Error	R/W	11	Retain	00 = Off 01 = Reserved 10 = Dual 11 = Single
1:0	COLX	R/W	LED[1:0] 00 = Off 01 = Single 10 = Single 11 = Single	Retain	00 = Off 01 = Reserved 10 = Dual 11 = Single

1. Global register bits are used to control features and functions that are common to all ports in the device.

Table 136: PHY Manual LED Override
Offset: 0x19 (Hex), or 25 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	R/W	0x0	Retain	0
14	InvLED2	R/W	0x0	Retain	Invert LED2. This bit controls the active level of the LED2 pin for all LED options (those controlled below and those controlled by the LED Parallel Select Register at offset 0x16, decimal 22). 0 = Active Low LED2 1 = Active High LED2
13	InvLED1	R/W	0x0	Retain	Invert LED1. This bit controls the active level of the LED1 pin for all LED options (those controlled below and those controlled by the LED Parallel Select Register at offset 0x16, decimal 22). 0 = Active Low LED1 1 = Active High LED1
12	InvLED0	R/W	0x0	Retain	Invert LED0. This bit controls the active level of the LED0 pin for all LED options (those controlled below and those controlled by the LED Parallel Select Register at offset 0x16, decimal 22). 0 = Active Low LED1 1 = Active High LED1
11:10	SpLED2	R/W	0x0	Retain	LED2 Speed Select Register for LED Parallel Select register at offset 0x16, decimal 22. 00 = Normal 01 = Active for 10 Mbps Speed only 10 = Active for 100 Mbps Speed only 11 = Reserved
9:8	SpLED1	R/W	0x0	Retain	LED1 Speed Select Register for LED Parallel Select register at offset 0x16, decimal 22. 00 = Normal 01 = Active for 10 Mbps Speed only 10 = Active for 100 Mbps Speed only 11 = Reserved
7:6	SpLED0	R/W	0x0	Retain	LED0 Speed Select Register for LED Parallel Select register at offset 0x16, decimal 22. 00 = Normal 01 = Active for 10 Mbps Speed only 10 = Active for 100 Mbps Speed only 11 = Reserved
5:4	ForceLED2	R/W	0x0	Retain	00 = Normal 01 = Blink[1] 10 = LED Off 11 = LED On

Table 136: PHY Manual LED Override
Offset: 0x19 (Hex), or 25 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
3:2	ForceLED1	R/W	0x0	Retain	00 = Normal 01 = Blink[1] 10 = LED Off 11 = LED On
1:0	ForceLED0	R/W	0x0	Retain	00 = Normal 01 = Blink[1] 10 = LED Off 11 = LED On

1. Energy Detect (Section 4.3) must be disabled on ports that are configured to blink an LED but don't have a link established.

Table 137: VCT™ Register for TXP/N Pins
Offset: 0x1A[1] (Hex), or 26 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	EnVCT	R/W, SC	0x0	0x0	Enable VCT The devices must be in forced 100 Mbps mode before enabling this bit. 0 = VCT completed 1 = Run VCT After running VCT once, bit 15 = 0 indicates VCT completed. The cable status is reported in the VCTTst bits in registers 26 and 27. Refer to the "Virtual Cable Tester®" feature on page 147.
14:13	VCTTst	RO	0x0	Retain	VCT Test Status These VCT test status bits are valid after completion of VCT. 00 = valid test, normal cable (no short or open in cable) 01 = valid test, short in cable (Impedance < 33 ohm) 10 = valid test, open in cable (Impedance > 333 ohm) 11 = Test fail
12:8	AmpRfln	RO	0x0	Retain	Amplitude of Reflection The amplitude of reflection is stored in these register bits. These amplitude bits range from 0x07 to 0x1F. 0x1F = Maximum positive amplitude 0x13 = Zero amplitude 0x07 = Maximum negative amplitude These bits are valid after completion of VCT (bit 15) and if the VCT test status bits (bits 14:13) have not indicated test failure.
7:0	DistRfln	RO	0x0	Retain	Distance of Reflection These bits refer to the approximate distance (± 1m) to the open/short location, measured at nominal conditions (room temperature and typical VDDs) These bits are valid after completion of VCT (bit 15) and if the VCT test status bits (bit 14:13) have not indicated test failure.

1. The results stored in this register apply to the Tx pin pair.

Table 138: VCT™ Register for RXP/N pins
Offset: 0x1B[1] (Hex), or 27 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RO	Always 0	Always 0	Reserved
14:13	VCTTst	RO	0	Retain	VCT Test Status The VCT test status bits are valid after completion of VCT. 00 = valid test, normal cable (no short or open in cable) 01 = valid test, short in cable (Impedance < 33 ohm) 10 = valid test, open in cable (Impedance > 333 ohm) 11 = Test fail
12:8	AmpRfln	RO	0	Retain	Amplitude of Reflection The amplitude of reflection is stored in these register bits. These amplitude bits range from 0x07 to 0x1F. 0x1F = Maximum positive amplitude 0x13 = Zero amplitude 0x07 = Maximum negative amplitude These bits are valid after completion of VCT (bit 15) and if VCT test status bits (bit 14:13) have not indicated test failure.
7:0	DistRfln	RO	0	Retain	Distance of Reflection These bits refer to the approximate distance (+/- 1m) to the open/short location, measured at nominal conditions (room temperature and typical VDDs) These bits are valid after completion of VCT (bit 15) and if VCT test status bits (bits 14:13) have not indicated test failure.

1. The results stored in this register apply to the Rx pin pair.

Figure 47: Cable Fault Distance Trend Line

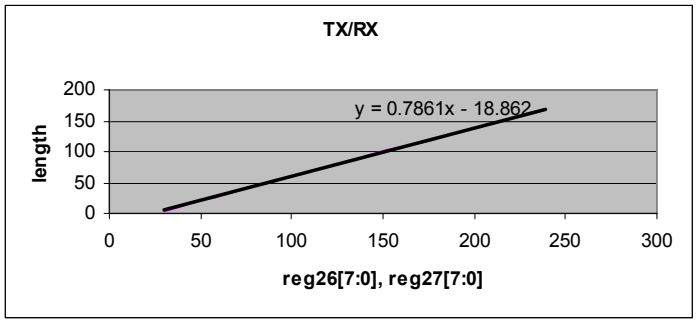

Table 139: PHY Specific Control Register II
Offset: 0x1C (Hex), or 28 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:5	Reserved	R/W	0x0	0x0	Must be 00000000000
4	EnLineLpbk	R/W	0	Retain	0 = Disable Line Loopback 1 = Enable Line Loopback
3	SoftwareMedia-Select	R/W	0	Update	0 = Select Copper Media 1 = Select Fiber Media **NOTE:** AutoMDI and Autoneg Enable take on the values set by the hardware configuration default.
2	TDRWaitTime	R/W	0x0	Retain	0 = Wait time is 1.5s before TDR test is started 1 = Wait time is 25 ms before TDR test is started
1	EnRXCLK	R/W	0x1	Update	0 = Disable MAC interface clock (RXCLK) in sleep mode 1 = Enable MAC interface clock (RXCLK) in sleep mode
0	SelClsA	R/W	SEL_ CLASS A/B[1]	Update	0 = Select Class B driver (typically used in CAT 5 applications) 1 = Select Class A driver - available for 100BASE-TX mode only (typically used in Backplane or direct connect applications, but may be used with CAT 5 applications)

1. This is the PHY SelClsA /B bit Hardware Reset name. The CONFIG_B pin of the 88E6092/88E6095/88E6095F devices contains an internal pull-up resistor, setting the default SelClsA PHY bit Hardware Reset to Class B drivers.

Note

Registers 0x1D through 0x1F (hexadecimal 29 through 31 decimal) are reserved. Do not read or write to these registers.

Section 8. Electrical Specifications

8.1 Absolute Maximum Ratings

Stresses above those listed in Absolute Maximum Ratings may cause permanent device failure. Functionality at or above these limits is not implied. Exposure to absolute maximum ratings for extended periods may affect device reliability.

Table 140: Absolute Maximum Ratings

Symbol	Parameter	Min	Typ	Max	Units
$V_{DD(3.3)}$	Power Supply Voltage on any 3.3V supply with respect to VSS	-0.5	3.3	+3.6	V
$V_{DD(2.5)}$	Power Supply Voltage on any 2.5V signal with respect to VSS	-0.5	2.5	+3.6 or $V_{DD(3.3)}$ +0.5[1] whichever is less	V
$V_{DD(1.5)}$	Power Supply Voltage on any 1.5V supply with respect to VSS	-0.5	1.5	+3.6 or $V_{DD(2.5)}$ +0.5[2] whichever is less	V
$V_{DD(1.2)}$	Power Supply Voltage on any 1.2V supply with respect to VSS	-0.5	1.2	+3.6 or $V_{DD(1.5)}$ +0.5[3] whichever is less	V
V_{PIN}	Voltage applied to any input pin with respect to VSS	-0.5		+3.6 or V_{DDO_PIN}[4] +0.5[5] whichever is less	V
$T_{STORAGE}$	Storage temperature	-55		+125[6]	°C

1. VDD(2.5) must never be more than 0.5V greater than VDD(3.3) or damage will result. Power must be applied to VDD(3.3) before or at the same time as VDD(2.5).
2. VDD(1.5) must never be more than 0.5V greater than VDD(2.5) or damage will result. Power must be applied to VDD(2.5) before or at the same time as VDD(1.5).
3. VDD(1.2) must never be more than 0.5V greater than VDD(1.5) or damage will result. Power must be applied to VDD(1.5) before or at the same time as VDD(1.2).
4. The VDDO pad ring has separate I/O power supply options. Therefore, the voltage applied to a group of I/O pins must follow what is defined in Section 1.
5. VPIN must never be more than 0.5V greater than VDDO or damage will result.
6. 125°C is the re-bake temperature. For extended storage time greater than 24 hours, +85°C should be the maximum.

8.2 Recommended Operating Conditions

Table 141: Recommended Operating Conditions

Symbol	Parameter	Condition	Min	Typ	Max	Units
$V_{DD(3.3)}$	3.3V power supply	For any 3.3V supply pin	3.135	3.3	3.465	V
$V_{DD(2.5)}$	2.5V power supply	For any 2.5V supply pin[1]	2.375	2.5	2.625	V
$V_{DD(1.5)}$	1.5V power supply	For any 1.5V supply pin	1.425	1.5	1.590	V
$V_{DD(1.2)}$	1.2V power supply	For any 1.2V supply pin	1.140	1.2	1.272	V
T_A	Ambient operating temperature[2]	Commercial parts	0		70	°C
		Industrial parts[3]	-40		85	°C
T_J	Maximum junction temperature				125[2]	°C
IREF	Internal bias reference	External resistor value required to be placed between IREF and VSS pins	1980	2000	2020	Ω

1. Some VDDO pins can be set to either 2.5V or 3.3V. To guarantee proper operation they must be set within the appropriate ranges in this table. VDDO voltages between 2.625V and 3.135V are not supported.
2. The important parameter is maximum junction temperature. As long as the maximum junction temperature is not exceeded, the device can be operated at any ambient temperature. Refer to White Paper on "TJ Thermal Calculations" for more information.
3. Industrial part numbers have an "I" following the commercial part numbers. See “Ordering Part Numbers and Package Markings” on page 295.

8.3 Thermal Conditions for 88E6092/88E6095 devices 176-pin TQFP Package

Symbol	Parameter	Condition	Min	Typ	Max	Units
θ_{JA}	Thermal resistance[1] - junction to ambient of the 88E6092/88E6095 device 176-Pin TQFP package $\theta_{JA} = (T_J - T_A)/ P$ P = Total Power Dissipation	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		20.30		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/ sec air flow		16.80		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/ sec air flow		15.80		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/ sec air flow		15.20		°C/W
ψ_{JT}	Thermal characteristic parameter[1] - junction to top center of the 88E6092/ 88E6095 device 176-Pin TQFP package $\psi_{JT} = (T_J - T_{TOP})/P$. T_{TOP} = Temperature on the top center of the package	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		0.13		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/ sec air flow		0.21		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/ sec air flow		0.27		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/ sec air flow		0.31		°C/W
θ_{JC}	Thermal resistance[1] - junction to case of the 88E6092/ 88E6095 device 176-Pin TQFP package $\theta_{JC} = (T_J - T_C)/ P_{Top}$ P_{Top} = Power Dissipation from the top of the package	JEDEC with no air flow		5.20		°C/W
θ_{JB}	Thermal resistance[1] - junction to board of the 88E6092/ 88E6095 device 176-Pin TQFP package $\theta_{JB} = (T_J - T_B)/ P_{bottom}$ P_{bottom} = power dissipation from the bottom of the package to the PCB surface.	JEDEC with no air flow		10.60		°C/W

1. Refer to white paper on TJ Thermal Calculations for more information.

8.4 Thermal Conditions for 88E6095FF device 216-pin LQFP Package

Symbol	Parameter	Condition	Min	Typ	Max	Units
θ_{JA}	Thermal resistance[1] - junction to ambient of the 88E6095F device 216-Pin LQFP package $\theta_{JA} = (T_J - T_A)/ P$ P = Total Power Dissipation	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		19.60		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/ sec air flow		16.00		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/ sec air flow		15.10		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/ sec air flow		14.50		°C/W
ψ_{JT}	Thermal characteristic parameter[1] - junction to top center of the 88E6095F device 216-Pin LQFP package $\psi_{JT} = (T_J - T_{TOP})/P$. T_{TOP} = Temperature on the top center of the package	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		0.33		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/ sec air flow		0.51		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/ sec air flow		0.62		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/ sec air flow		0.70		°C/W
θ_{JC}	Thermal resistance[1] - junction to case of the 88E6095F device 216-Pin LQFP package $\theta_{JC} = (T_J - T_C)/ P_{Top}$ P_{Top} = Power Dissipation from the top of the package	JEDEC with no air flow		7.60		°C/W
θ_{JB}	Thermal resistance[1] - junction to board of the 88E6095F device 216-Pin LQFP package $\theta_{JB} = (T_J - T_B)/ P_{bottom}$ P_{bottom} = power dissipation from the bottom of the package to the PCB surface.	JEDEC with no air flow		10.00		°C/W

1. Refer to white paper on TJ Thermal Calculations for more information.

8.5 DC Electrical Characteristics

8.5.1 Current Consumption

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 142: Current Consumption

Pins	Parameter	Condition	Min	Typ	Max	Units
V_{DDO}[1]	Outputs with 3.3V power applied	All ports (FE and GE) idle		96		mA
		All ports (FE and GE) active		96		mA
	Outputs with 2.5V power applied	All ports (FE and GE) idle		41		mA
		All ports (FE and GE) active		41		mA
V_{DDAH}[2]	2.5V analog power to PHY (Energy Detect Disabled)	No link on any FE PHY		86		mA
		All FE PHYs 10 Mbps linked and idle		212		mA
		All FE PHYs 10 Mbps and active		215		mA
		All FE PHYs 100 Mbps		180		mA
V_{CT}	2.5V power to FE Magnetics Center Tap	No link on any FE PHY		2		mA
		All FE PHYs 10 Mbps linked and idle		2		mA
		All FE PHYs 10 Mbps and active		475		mA
		All FE PHYs 100 Mbps		159		mA
Px_V_{DDAH}	2.5V analog power to SERDES	All GE SERDES active		80		mA
		All GE SERDES power down		0		mA
Px_V_{TT}	1.5V analog power to SERDES	All GE SERDES active		64		mA
		All GE SERDES power down		0		mA
	1.2V analog power to SERDES	All GE SERDES active		39		mA
		All GE SERDES power down		0		mA
V_{DDAL}	1.2V analog power to PHY	No link on any FE PHY		4		mA
		All FE PHYs 10 Mbps linked and idle		4		mA
		All FE PHYs 10 Mbps and active		6		mA
		All FE PHYs 100 Mbps		50		mA
V_{DD_CORE}	1.2V analog power to core	All ports (FE and GE) idle		252		mA
		All ports (FE and GE) active		290		mA

1. VDDO includes VDDO_P9, VDDO_P10, VDDO_LED, VDDO_PHY, VDDO_EEPROM
2. VDDAH includes VDDAH and VDD_PLL

8.5.2 Digital Operating Conditions

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 143: Digital Operating Conditions

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
V_{IH}	High level input voltage	All pins	VDDO = 3.135V	2.0			V
			VDDO = 2.375V	1.7			V
V_{IL}	Low level input voltage	All pins	VDDO = 3.135V	-0.3		0.8	V
			VDDO = 2.375V	-0.3		0.7	V
V_{OH}	High level output voltage	LED pins	I_{OH} = -8 mA	VDDO - 0.4			V
		All others (except INTn [1])	I_{OH} = -4 mA	VDDO - 0.4			V
V_{OL}	Low level output voltage	INTn and LED pins	I_{OL}= 8 mA			0.4	V
		All others	I_{OL}= 4 mA			0.4	V
I_{ILK}	Input leakage current	With pull-up resistor	$0<V_{IN}<V_{DD}$			+ 10 - 50	µA
		With pull-down resistor	$0<V_{IN}<V_{DD}$			+ 50 - 10	µA
		All others	$0<V_{IN}<V_{DD}$			±10	µA
C_{IN}	Input capacitance	All pins				5	pF

1. The INTn is an active low, open drain pin. See INTn description in the Signal Description.

8.5.3 SERDES Electrical Specifications

Table 144: Transmitter DC Characteristics with Px_VTT = 1.5V Typical

Symbol	Parameter[1]	Min	Typ	Max	Units
V_{OH}	Output Voltage High			1500	mV
V_{OL}	Output Voltage Low	800			mV
V_{RING}	Output Ringing			10	mV
$\|V_{OD}\|$	Output Voltage Swing (single - ended, peak-to-peak)	100		700	mV peak
V_{OS}	Output Offset Voltage	1175		1425	mV
R_O	Output Impedance (single-ended) (50 ohm termination)	40		60	Ωs
Delta R_O	Mismatch in a pair			10	%
Delta V_{OD}	Change in VOD between 0 and 1			25	mV
Delta V_{OS}	Change in VOS between 0 and 1			25	mV
I_{SA}, I_{SB}	Output current on short to VSS			30	mA
I_{SAB}	Output current when a, b are shorted			12	mA
I_{XA}, I_{XB}	Power off leakage current			5	mA

1. All parameters measured at RLOAD = 100 Ω ± 1% load (P[9:0]_TXP/N).

Table 145: Transmitter DC Characteristics with Px_VTT = 1.2V Typical

Symbol	Parameter[1]	Min	Typ	Max	Units
V_{OH}	Output Voltage High			1200	mV
V_{OL}	Output Voltage Low	800			mV
V_{RING}	Output Ringing			10	mV
$\|V_{OD}\|$	Output Voltage Swing (single - ended, peak-to-peak)	100		400	mV peak
V_{OS}	Output Offset Voltage	TBD		TBD	mV
R_O	Output Impedance (single-ended) (50 ohm termination)	40		60	Ωs
Delta R_O	Mismatch in a pair			10	%
Delta V_{OD}	Change in VOD between 0 and 1			25	mV
Delta V_{OS}	Change in VOS between 0 and 1			25	mV
I_{SA}, I_{SB}	Output current on short to VSS			24	mA
I_{SAB}	Output current when a, b are shorted			12	mA
I_{XA}, I_{XB}	Power off leakage current			5	mA

1. All parameters measured at RLOAD = 100 Ω ± 1% load (P[9:0]_TXP/N).

Table 146: Receiver DC Characteristics

Symbol	Parameter	Min	Typ	Max	Units
V_I	Input Voltage range a or b	675		1725	mV
V_{IDTH}[1]	Input Differential Threshold |RXP - RXN|	50			mV
V_{HYST}[1]	Input Differential Hysteresis	25			mV
R_{IN}	Receiver 100Ω Differential Input Impedance	80		120	Ω

1. Receiver is at high level when VRXP - VRXN is greater than VIDTH (min) + VOFFSET and is at low level when VRXP - VRXN is less than -VIDTH (min) + VOFFSET. A minimum hysteresis of VHYST is present between -VIDTH and +VIDTH as shown in the figure. When the fiber link is down, an offset of VOFFSET is applied to prevent false signal detect due to noise. When the fiber link is up, the offset circuit is disabled.

Figure 48: Input Differential Hysteresis

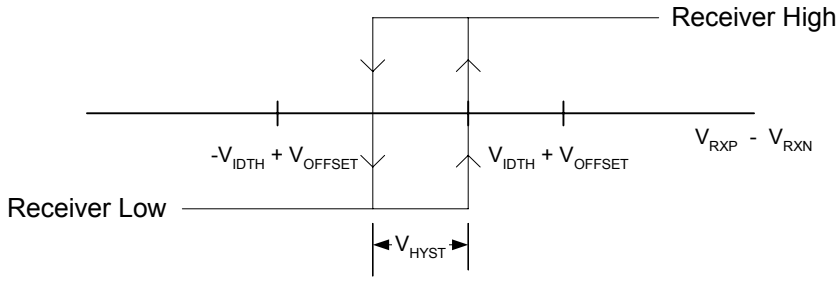

8.5.4 IEEE DC Transceiver Parameters

IEEE tests are typically based on templates and cannot simply be specified by a number. For an exact description of the template and the test conditions, refer to the IEEE specifications:
-10BASE-T IEEE 802.3 Clause 14
-100BASE-TX ANSI X3.263-1995
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 147: IEEE DC Transceiver Parameters

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
V_{ODIFF}	Absolute peak differential output voltage	TXP/N[7:0]	10BASE-T no cable	2.2	2.5	2.8	V
		TXP/N[7:0]	10BASE-T cable model	585[1]			mV
		TXP/N[1:0]	100BASE-FX mode	0.4	0.8	1.2	V
		TXP/N[7:0]	100BASE-TX mode	0.950	1.0	1.05	V
	Overshoot[2]	TXP/N[7:0]	100BASE-TX mode	0		5%	V
	Amplitude Symmetry (positive/ negative)	TXP/N[7:0]	100BASE-TX mode	0.98x		1.02x	V+/V-
V_{IDIFF}	Peak Differential Input Voltage accept level	RXP/N[7:0]	10BASE-T mode	585[3]			mV
		RXP/N[1:0] P[1:0]_SDET	100BASE-FX mode	200			mV
	Signal Detect Assertion	RXP/N[1:0]	100BASE-TX mode	1000	460[4]		mV peak-peak
	Signal Detect De-assertion	RXP/N[1:0]	100BASE-TX mode	200	360[5]		mV peak-peak

1. IEEE 802.3 Clause 14, Figure 14.9 shows the template for the “far end” wave form. This template allows as little as 495 mV peak differential voltage at the far end receiver.
2. ANSI X3.263-1995 Figure 9-1.
3. The input test is actually a template test. IEEE 802.3 Clause 14, Figure 14.17 shows the template for the receive wave form.
4. The ANSI TP-PMD specification requires that any received signal with peak-to-peak differential amplitude greater than 1000 mV should turn on signal detect (internal signal in 100BASE-TX mode). The devices will accept signals typically with 460 mV peak-to-peak differential amplitude.
5. The ANSI TP-PMD specification requires that any received signal with peak-to-peak differential amplitude less than 200 mV should be de-assert signal detect (internal signal in 100BASE-TX mode). The devices will reject signals typically with peak-to-peak differential amplitude less than 360 mV.

8.6 AC Electrical Specifications

8.6.1 Receiver AC Characteristics

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified.)

Table 148: Reset and Configuration Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{PU_RESET}	Valid power to RESETn de-asserted or RESETn assertion time	At power up or subsequent resets after power up	10			ms
T_{SU_CLK}	Number of valid REFCLK cycles prior to RESETn de-asserted		10			Clks
T_{SU}	Configuration data valid prior to RESETn de-asserted[1]		200			ns
T_{HD}	Config data valid after RESETn de-asserted		0			ns

1. When RESETn is low all configuration pins become inputs, and the value seen on these pins is latched on the rising edge of RESETn. All configuration pins that become outputs during normal operation will remain tri-stated for 40 ns after the rising edge of RESETn.

Figure 49: Reset and Configuration Timing

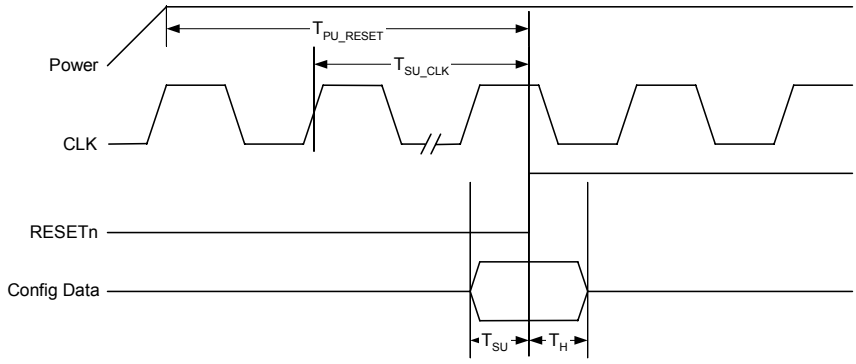

8.6.2 Clock Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified.)

Table 149: Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P[1]	XTAL_IN period		40 -50 ppm	40	40 +50 ppm	ns
T_H	XTAL_IN high time		16			ns
T_L	XTAL_IN low time		16			ns
T_R	XTAL_IN rise				3	ns
T_F	XTAL_IN fall				3	ns

1. 25.000 MHz

Figure 50: Oscillator Clock Timing

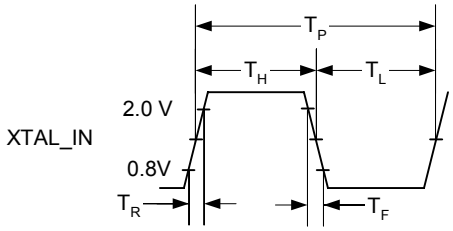

8.7 GMII Interface Timing

8.7.1 GMII Transmit Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 150: GMII Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{SU\_GMII\_GTX\_CLK}$	GMII output to clock		2.0			ns
$T_{HD\_GMII\_GTX\_CLK}$	GMII clock to output		0			ns
$T_{H\_GMII\_GTX\_CLK}$	GTX_CLK High		2.5[1]		5.5	ns
$T_{L\_GMII\_GTX\_CLK}$	GTX_CLK Low		2.5[1]		5.5	ns
$T_{P\_GMII\_GTX\_CLK}$	GTX_CLK Period		7.5[1]	8.0		ns
$T_{R\_GMII\_GTX\_CLK}$	GTX_CLK Rise Time				1.0	ns
$T_{F\_GMII\_GTX\_CLK}$	GTX_CLK Fall Time				1.0	ns
$T_{RSLEW\_GMII\_GTX\_CLK}$	GTX_CLK Rising Slew Rate		0.6[2]			V/ns
$T_{FSLEW\_GMII\_GTX\_CLK}$	GTX_CLK Falling Slew Rate		0.6[2]			V/ns

1. GTX_CLK numbers not guaranteed during transition between 10/100/1000BASE-T operation.
2. Instantaneous change during internal VIH_GMII (Min.) and VIL_GMII (Max.).

Figure 51: GMII Transmit Timing

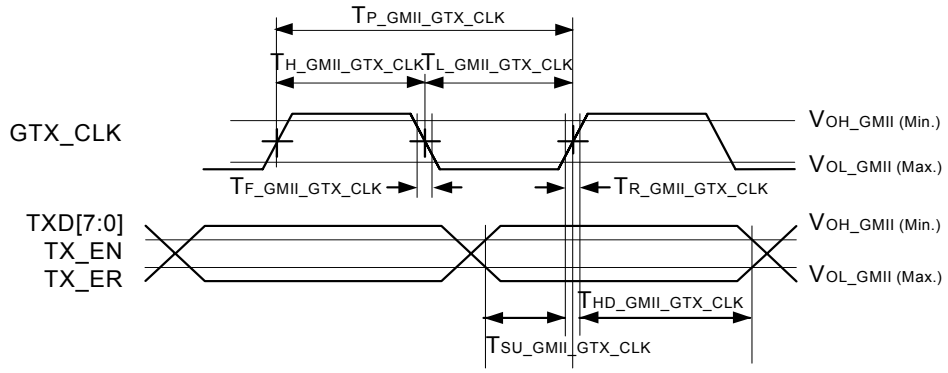

8.7.2 GMII Receive Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 151: GMII Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{SU\_GMII\_RX\_CLK}$	GMII Setup Time		2.0			ns
$T_{HD\_GMII\_RX\_CLK}$	GMII Hold Time		0			ns
$T_{H\_GMII\_RX\_CLK}$	RX_CLK High		2.5[1]			ns
$T_{L\_GMII\_RX\_CLK}$	RX_CLK Low		2.5[1]			ns
$T_{P\_GMII\_RX\_CLK}$	RX_CLK Period		7.5[1]	8.0	8.5	ns
$F_{GMII\_RX\_CLK}$	RX_CLK Frequency		125[1] -100 ppm		125 +100 ppm	MHz
$T_{R\_GMII\_RX\_CLK}$	RX_CLK Rise Time				1.0	ns
$T_{F\_GMII\_RX\_CLK}$	RX_CLK Fall Time				1.0	ns

1. RX_CLK toggle rate is “don’t care” if link is down, or if not in 1000BASE-T mode.

Figure 52: GMII Receive Timing

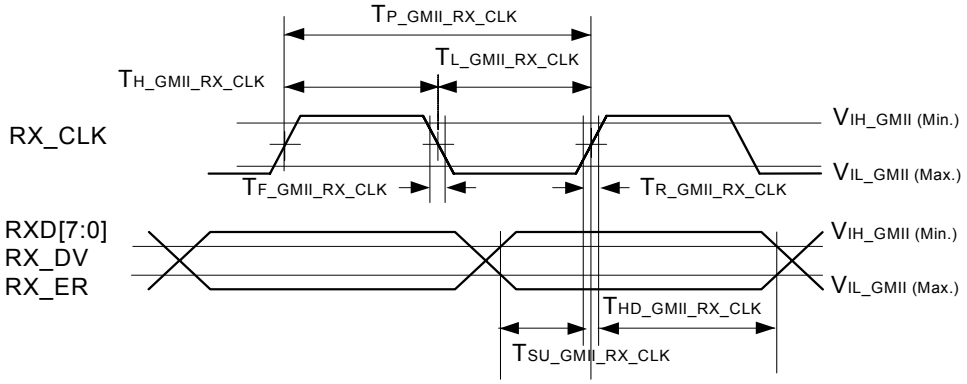

8.8 MII MAC Mode Clock Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

In MII MAC mode, the P9_RXCLK, P10_RXCLK and P9_TXCLK, P10_TXCLK pins are inputs.

Table 152: MII MAC Mode Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P	MACCLK_IN period		20[1]	40 or 400		ns
T_H	MACCLK_IN high time		8			ns
T_L	MACCLK_IN low time		8			ns
T_R	MACCLK_IN rise				3	ns
T_F	MACCLK_IN fall				3	ns

1. This value applies for 200 Mbps mode

Figure 53: MAC Clock Timing

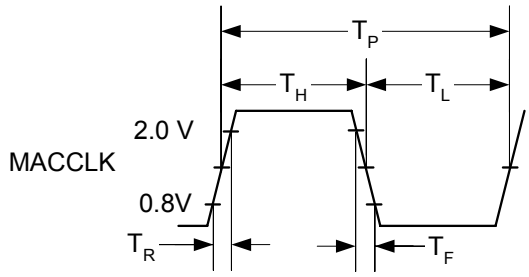

8.8.1 MII Receive Timing - MAC Mode

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 153: MII Receive Timing—MAC Mode 100 Mbps Operation

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{SU_RX}	MII inputs (P9_RXD[3:0], P9_RXDV, and P9_RXER) valid prior to P9_RXCLK going high, or MII inputs (P10_RXD[3:0], P10_RXDV, and P10_RXER) valid prior to P10_RXCLK going high	With 10 pF load	10			ns
T_{HD_RX}	MII inputs (P9_RXD[3:0], P9_RXDV, and P9_RXER) valid after P9_RXCLK going high, or MII inputs (P10_RXD[3:0], P10_RXDV, and P10_RXER) valid after P10_RXCLK going high	With 10pF load	10			ns

Table 154: MII Receive Timing—MAC Mode 200 Mbps Operation

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{SU_RX}	MII inputs (P9_RXD[3:0], P9_RXDV, and P9_RXER) valid prior to P9_RXCLK going high, or MII inputs (P10_RXD[3:0], P10_RXDV, and P10_RXER) valid prior to P10_RXCLK going high	With 10 pF load	5			ns
T_{HD_RX}	MII inputs (P9_RXD[3:0], P9_RXDV, and P9_RXER) valid after P9_RXCLK going high, or MII inputs (P10_RXD[3:0], P10_RXDV, and P10_RXER) valid after P10_RXCLK going high	With 10pF load	2			ns

Figure 54: MAC Mode MII Receive Timing

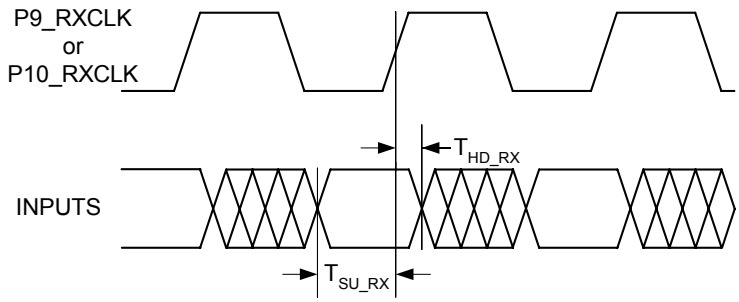

8.8.2 MII Transmit Timing - MAC Mode

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 155: MII Transmit Timing—MAC Mode 100 Mbps Operation

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{TXCLK}	P9_TXCLK period P10_TXCLK period	10BASE mode[1]		400		ns
		100BASE mode		40		ns
T_{H_TXCLK}	P9_TXCLK high P10_TXCLK high	10BASE mode		200		ns
		100BASE mode		20		ns
T_{L_TXCLK}	P9_TXCLK low P10_TXCLK low	10BASE mode		200		ns
		100BASE mode		20		ns
T_{CQ_MAX}	P9_TXCLK to outputs (P9_TXD[3:0], P9_TXEN) valid P10_TXCLK to outputs (P10_TXD[3:0], P10_TXEN) valid	With 10 pF load			25	ns
T_{CQ_MIN}	P9_TXCLK to outputs P9_TXD[3:0], P9_TXEN) invalid P10_TXCLK to outputs P10_TXD[3:0], P10_TXEN) invalid	With 10 pF load	0			ns

1. 2.5 MHz for 10 Mbps or 25 MHz for 100 Mbps

Table 156: MII Transmit Timing—MAC Mode 200 Mbps Operation

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{TXCLK}	P9_TXCLK period P10_TXCLK period	200 Mbps Operation mode[1]		20		ns
T_{H_TXCLK}	P9_TXCLK high P10_TXCLK high	200 Mbps Operation		10		ns
T_{L_TXCLK}	P9_TXCLK low P10_TXCLK low	200 Mbps Operation		10		ns
T_{CQ_MAX}	P9_TXCLK to outputs (P9_TXD[3:0], P9_TXEN) valid P10_TXCLK to outputs (P10_TXD[3:0], P10_TXEN) valid	With 10 pF load			15	ns
T_{CQ_MIN}	P9_TXCLK to outputs P9_TXD[3:0], P9_TXEN) invalid P10_TXCLK to outputs P10_TXD[3:0], P10_TXEN) invalid	With 10 pF load	3			ns

1. 2.5 MHz for 10 Mbps or 25 MHz for 100 Mbps

Figure 55: MAC Mode MII Transmit Timing

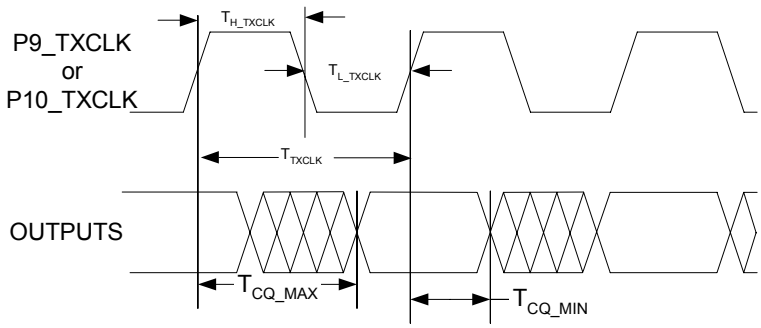

8.8.3 Serial Management Interface Clock Timing (CPU Set)

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 157: Serial Management Interface Clock Timing (CPU Set)

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	MDC period		120			ns	8.33 MHz
T_H	MDC high time		48			ns	
T_L	MDC low time		48			ns	
T_R	MDC rise				6	ns	
T_F	MDC fall				6	ns	

Figure 56: Serial Management Interface Clock Timing (CPU Set)

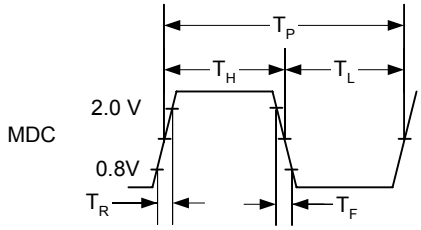

8.8.4 Serial Management Interface Data Timing (CPU Set)

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 158: Serial Management Interface Data Timing (CPU Set)

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{DLY_MDIO}	MDC to MDIO (Output) delay time		5		12.5	ns	
T_{SU}	MDIO (Input) to MDC setup time		10			ns	
T_{HD}	MDIO (Input) to MDC hold time		10			ns	

Figure 57: Serial Management Interface Data Timing

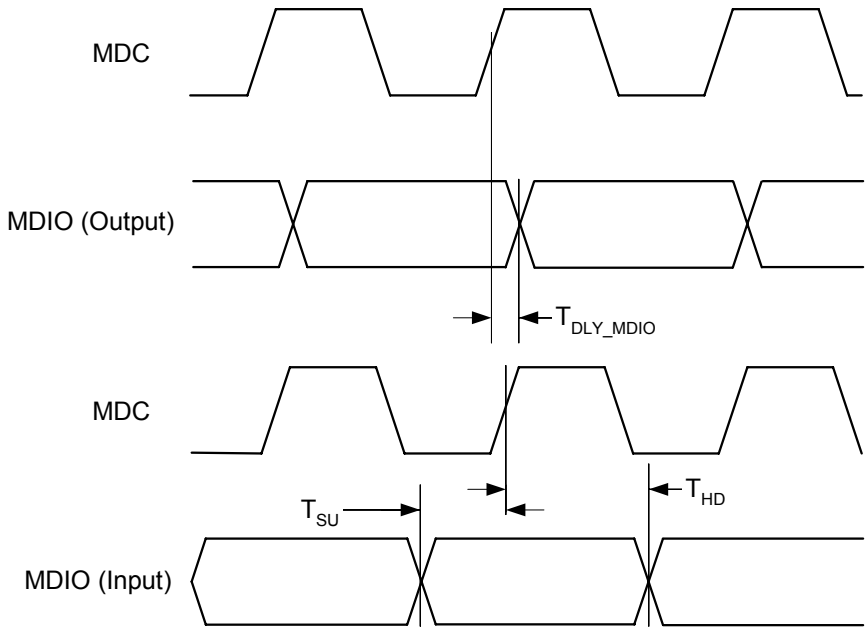

8.8.5 Serial Management Interface Timing (PHY Set)

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 159: Serial Management Interface Timing (PHY Set)

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	MDC period		120[1]			ns	8.33 MHz
T_H	MDC high time		48			ns	
T_L	MDC low time		48			ns	
T_R	MDC rise				6	ns	
T_F	MDC fall				6	ns	
T_{SU}	MDIO setup time		15			ns	2
T_{HD}	MDIO hold time		5			ns	2
T_{DLY_MDIO}	MDC to MDIO (Output) delay time		0		5	ns	3

1. MDC_PHY will track MDC_CPU when the PPU is disabled. When the PPU is enabled the MDC_PHY period will be 240 ns.
2. MDIO input setup and hold time is intentionally sampled with respect to the MDC falling edge.
3. MDIO data is intentionally clocked out on the falling edge of MDC.

Figure 58: Serial Management Interface Timing Output (PHY Mode)

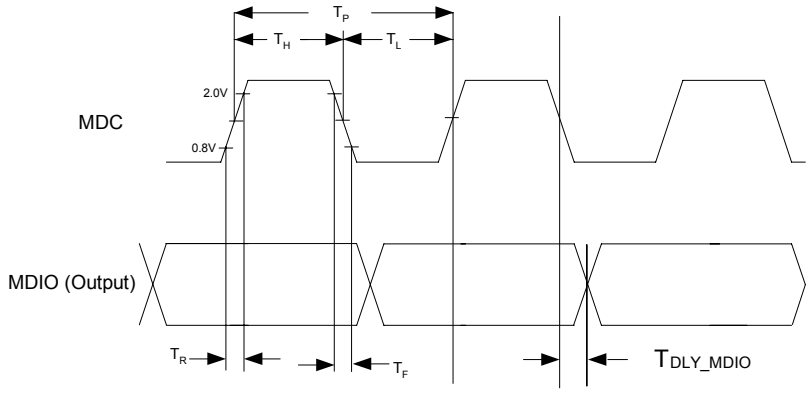

Figure 59: Serial Management Interface Timing Input (PHY Mode)

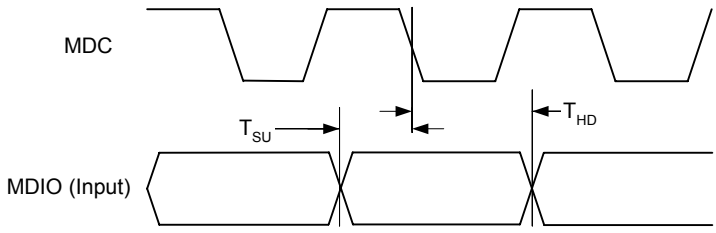

8.8.6 EEPROM Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 160: EEPROM Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	EE_CLK period			5120		ns	
T_H	EE_CLK high time			2560		ns	
T_L	EE_CLK low time			2560		ns	
T_{CQCSMX}	Serial EEPROM chip select valid	Referenced to EE_CLK			5	ns	
T_{CQCSMN}	Serial EEPROM chip select invalid				5	ns	
T_{CQDMX}	Serial EEPROM data transmitted to EEPROM valid				10	ns	
T_{CQDMN}	Serial EEPROM data transmitted to EEPROM invalid		3			ns	
T_S	Setup time for data received from EEPROM		10			ns	
T_H	Hold time for data received from EEPROM		10			ns	

Figure 60: EEPROM Timing

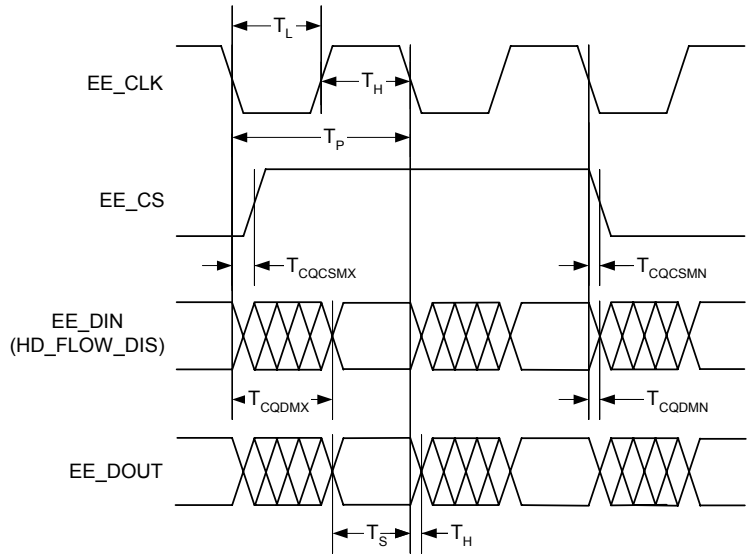

8.9 SERDES (Serial Interface) Timing

Table 161: SERDES (Serial Interface) Transmitter AC Characteristics

Symbol	Parameter	Min	Typ	Max	Units
CLOCK	Clock signal duty cycle @ 125 MHz	45		55	%
T_{FALL}	V_{OD} Fall time (20% - 80%)	100		1000	ps
T_{RISE}	V_{OD} Rise time (20% - 80%)	100		1000	ps

Table 162: SERDES (Serial Interface) Receiver AC Characteristics

Symbol	Parameter	Min	Typ	Max	Units
$T_{SOUTPUT}$[1]	SERDES output	360	400	440	ps

1. Measured at 50% of the transition.

Figure 61: SERDES Rise and Fall Times

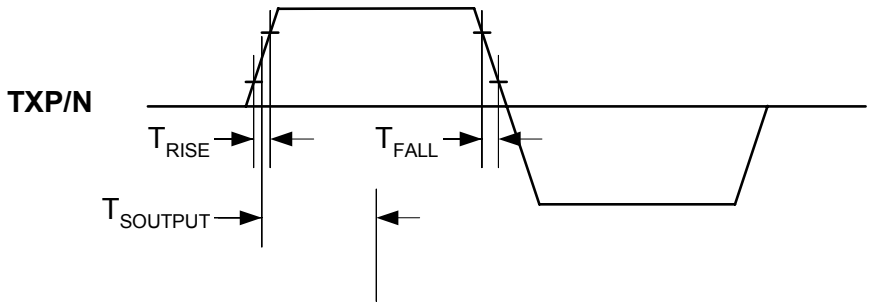

8.9.1 IEEE AC Parameters (Ports 0–7)

IEEE tests are typically based on templates and cannot simply be specified by number. For an exact description of the templates and the test conditions, refer to the IEEE specifications:

-10BASE-T IEEE 802.3 Clause 14-2000

-100BASE-TX ANSI X3.263-1995

-1000BASE-T IEEE 802.3ab Clause 40 Section 40.6.1.2 Figure 40-26 shows the template waveforms for transmitter electrical specifications.

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 163: IEEE AC Parameters

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
T_{RISE}	Rise time	TXP/N[7:0]	100BASE-TX	3.0	4.0	5.0	ns
T_{FALL}	Fall time	TXP/N[7:0]	100BASE-TX	3.0	4.0	5.0	ns
T_{RISE}/ T_{FALL} Symmetry		TXP/N[7:0]	100BASE-TX	0		0.5	ns
DCD	Duty cycle distortion	TXP/N7:0]	100BASE-TX	0		0.5[1]	ns, peak-peak
Transmit Jitter		TXP/N[7:0]	100BASE-TX	0		1.4	ns, peak-peak

1. ANSI X3.263-1995 Table 3.

Section 9. Package Mechanical Dimensions

Figure 62: 88E6092/88E6095 176-pin TQFP EPAD Package Mechanical Drawings

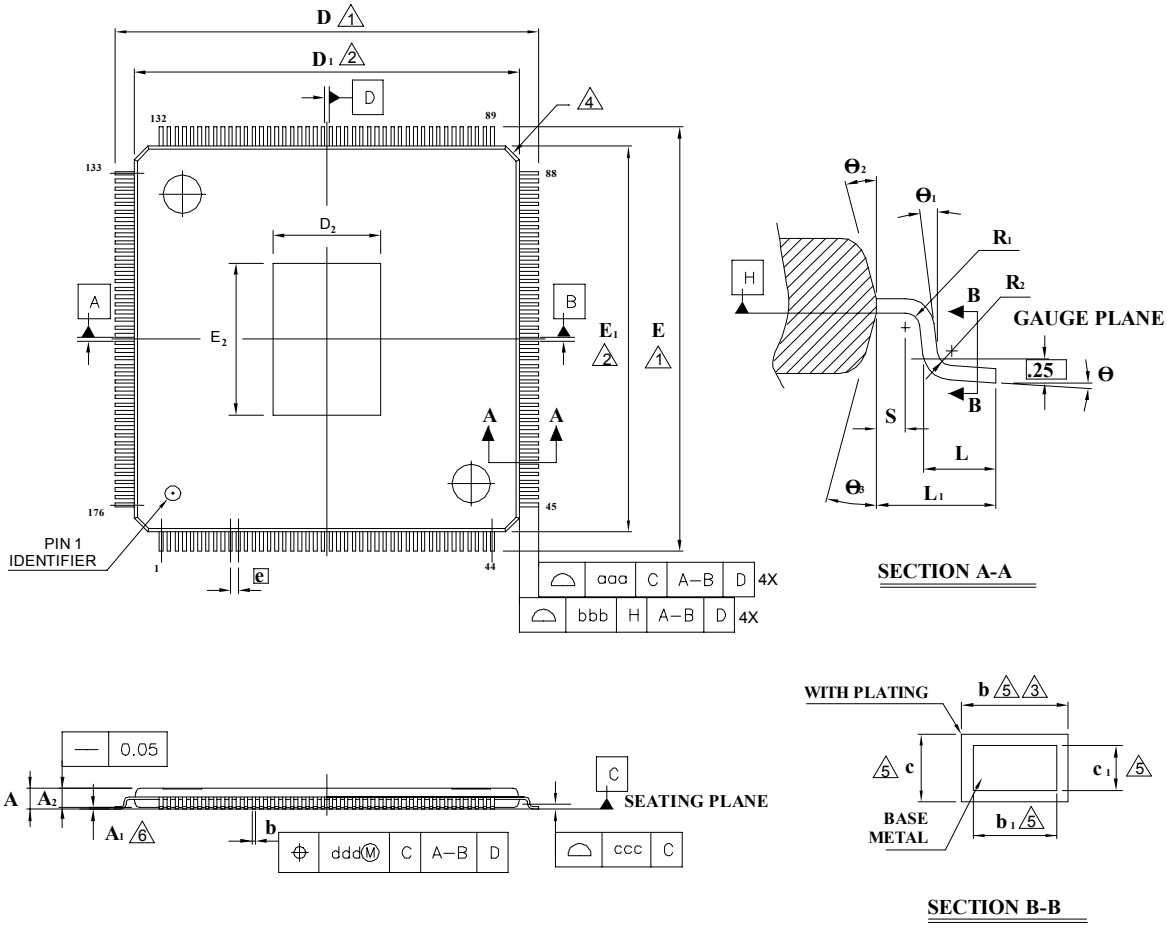

NOTE :

1. TO BE DETERMINED AT SEATING PLANE C .
2. DIMENSIONS D1 AND E1 DO NOT INCLUDE MOLD PROTRUSION D1 AND E1 ARE MAXIMUM PLASTIC BODY SIZE DIMENSIONS INCLUDING MOLD MISMATCH.
3. DIMENSION b DOES NOT INCLUDE DAMBAR PROTRUSION. DAMBAR CAN NOT BE LOCATED ON THE LOWER RADIUS OR THE FOOT.
4. EXACT SHAPE OF EACH CORNER IS OPTIONAL.
5. THESE DIMENSIONS APPLY TO THE FLAT SECTION OF THE LEAD BETWEEN 0.10 mm AND 0.25 mm FROM THE LEAD TIP.
6. A1 IS DEFINED AS THE DISTANCE FROM THE SEATING PLANE TO THE LOWEST POINT OF THE PACKAGE BODY.
7. CONTROLLING DIMENSION : MILLIMETER.

Table 164: 88E6092/88E6095 176-Pin TQFP EPAD Package Dimensions in mm

Symbol	Dimension in mm		
	Min	Nom	Max
A	1.00	1.10	1.20
A_1	0.05	0.10	0.15
A_2	0.95	1.00	1.05
b	0.13	0.18	0.23
b_1	0.13	0.16	0.19
c	0.09	--	0.20
c_1	0.09	--	0.16
D	22.00 BSC		
D_1	20.00 BSC		
E	22.00 BSC		
E_1	20.00 BSC		
e	0.40 BSC		
L	0.45	0.60	0.75
L_1	1.00 REF		
R_1	0.08	--	--
R_2	0.08	--	0.20
S	0.20	--	--
D_2	7.11		
E_2	7.87		
q	0°	3.5°	7°
θ_1	0°	--	--
θ_2	11°	12°	13°
θ_3	11°	12°	13°
aaa	0.20		
bbb	0.20		
ccc	0.08		
ddd	0.07		

Figure 63: 88E6095F 216-pin LQFP EPAD Package Mechanical Drawings

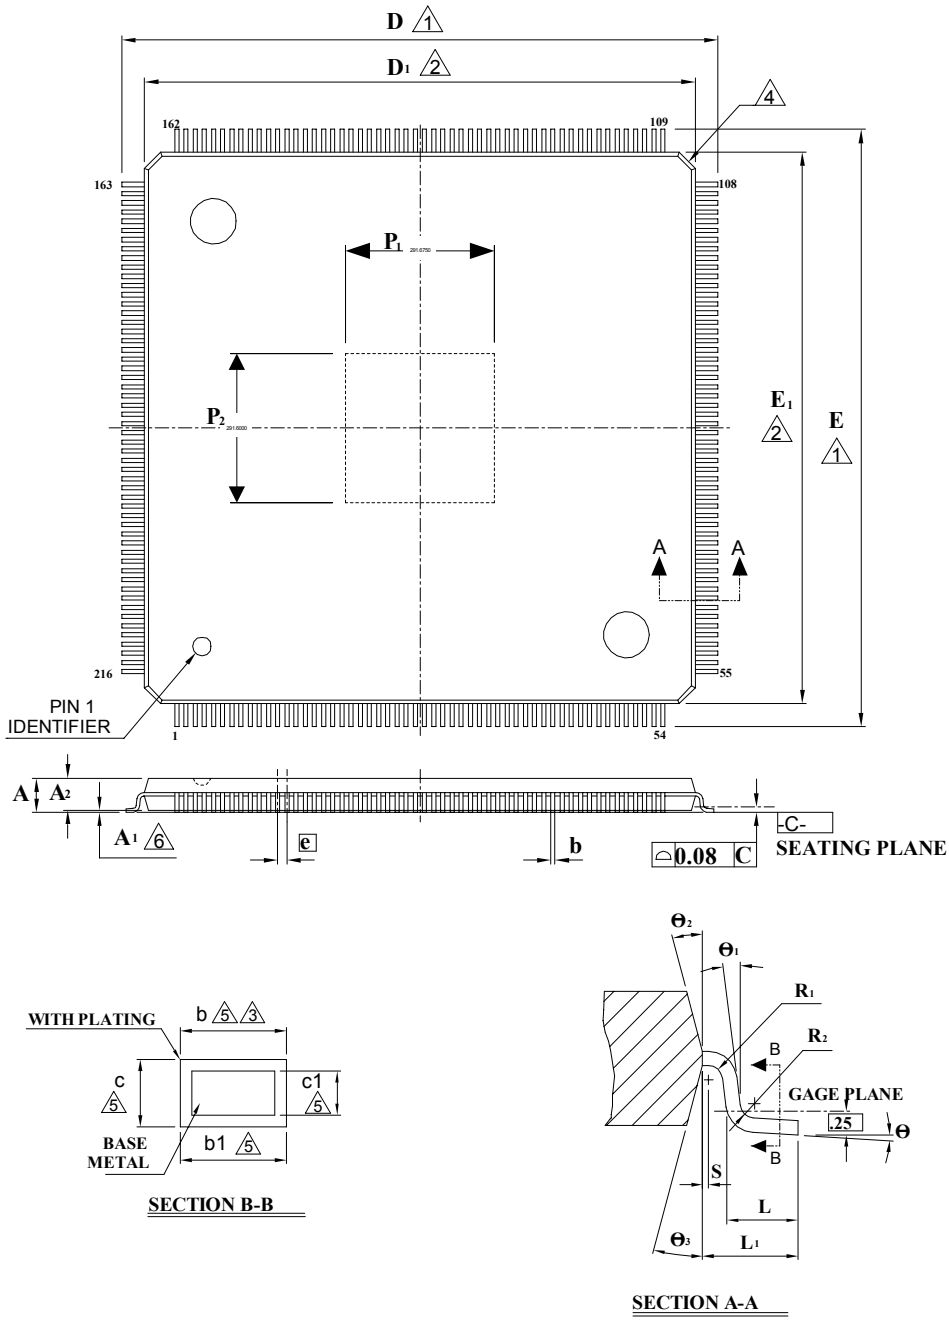

Table 165: 88E6095F 216-pin LQFP EPAD Package Dimensions

Symbol	Dimension in mm		
	Min	Nom	Max
A	--	--	1.60
A_1	0.05	--	0.15
A_2	1.35	1.40	1.45
b	0.13	0.18	0.23
b_1	0.13	0.16	0.19
c	0.09	0.14	0.20
c_1	0.09	0.12	0.16
D	25.60	26.00	26.40
D_1	--	24.00	--
E	25.60	26.00	26.40
E_1	--	24.00	--
e	0.40 BSC		
L	0.45	0.60	0.75
L_1	1.00 REF		
R_1	0.08	--	--
R_2	0.08	--	--
S	0.20	--	--
P_1	7.87		
P_2	7.87		
θ	0°	3.5°	7°
θ_1	0°	--	--
θ_2	11°	12°	13°
θ_3	11°	12°	13°

Notes:

1. TO BE DETERMINED AT SEATING PLANE -C-.
2. DIMENSIONS D1 AND E1 DO NOT INCLUDE MOLD PROTRUSION. D1 AND E1 ARE MAXIMUM PLASTIC BODY SIZE DIMENSIONS INCLUDING MOLD MISMATCH.
3. DIMENSION b DOES NOT INCLUDE DAMBAR PROTRUSION. DAMBAR CAN NOT BE LOCATED ON THE LOWER RADIUS OR THE FOOT.
4. EXACT SHAPE OF EACH CORNER IS OPTIONAL.
5. THESE DIMENSIONS APPLY TO THE FLAT SECTION OF THE LEAD BETWEEN 0.10 mm AND 0.25 mm FROM THE LEAD TIP.
6. A1 IS DEFINED AS THE DISTANCE FROM THE SEATING PLANE TO THE LOWEST POINT OF THE PACKAGE BODY.
7. CONTROLLING DIMENSION: MILLIMETER.

Section 10. Ordering Information

10.1 Ordering Part Numbers and Package Markings

Figure 64 shows the ordering part numbering scheme for the devices. Contact Marvell® FAEs or sales representatives for complete ordering information.

Figure 64: Sample Part Number

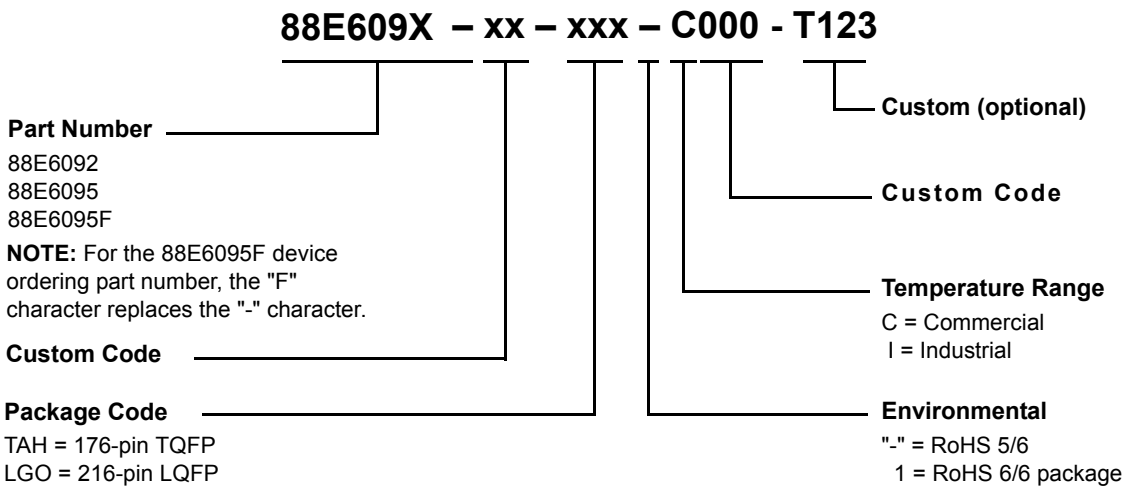

Table 166: 88E6092/88E6095/88E6095F Part Order Options - RoHS 5/6 Compliant Package

Package Type	Part Order Number
88E6092 176-pin TQFP - Commercial	88E6092-XX-TAH-C000
88E6095 176-pin TQFP - Commercial	88E6095-XX-TAH-C000
88E6095F 216-pin LQFP - Commercial	88E6095FXX-LGO-C000
88E6095 176-pin TQFP - Industrial	88E6095-XX-TAH-I000
88E6095F 216-pin LQFP - Industrial	88E6095FXX-LGO-I000

Table 167: 88E6092/88E6095/88E6095F Part Order Options - RoHS 6/6 Compliant Package

Package Type	Part Order Number
88E6092 176-pin TQFP - Commercial	88E6092-XX-TAH1C000
88E6095 176-pin TQFP - Commercial	88E6095-XX-TAH1C000
88E6095F 216-pin LQFP - Commercial	88E6095FXX-LGO1C000
88E6095 176-pin TQFP - Industrial	88E6095-XX-TAH1I000
88E6095F 216-pin LQFP - Industrial	88E6095FXX-LGO1I000

10.1.1 RoHS 5/6 Marking Examples

Figure 65 is an example of the package marking and pin 1 location for the 88E6092 176-pin TQFP Commercial RoHS 5/6 compliant package.

Figure 65: 88E6092 176-pin TQFP Commercial RoHS 5/6 Compliant Package Marking and Pin 1 Location

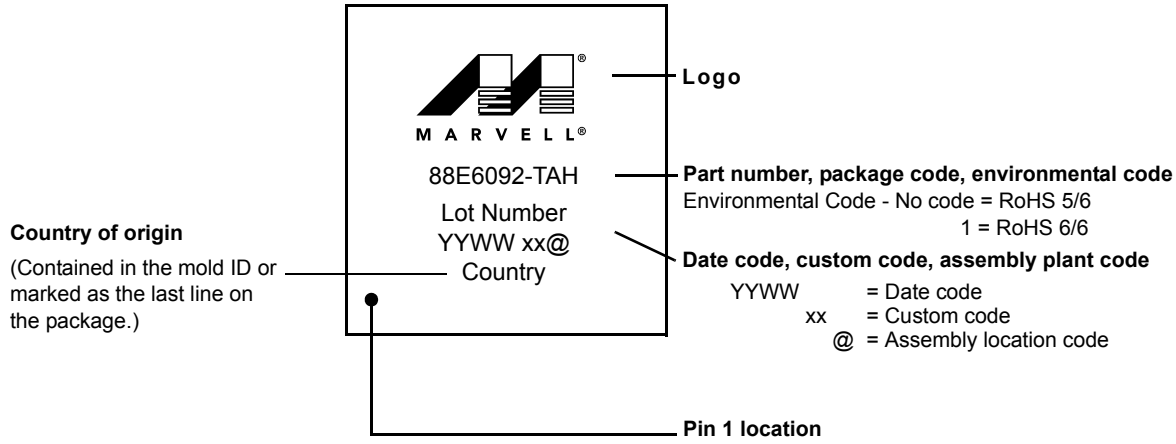

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 66 is an example of the package marking and pin 1 location for the 88E6095 176-pin TQFP Commercial RoHS 5/6 compliant package.

Figure 66: 88E6095 176-pin TQFP Commercial RoHS 5/6 Compliant Package Marking and Pin 1 Location

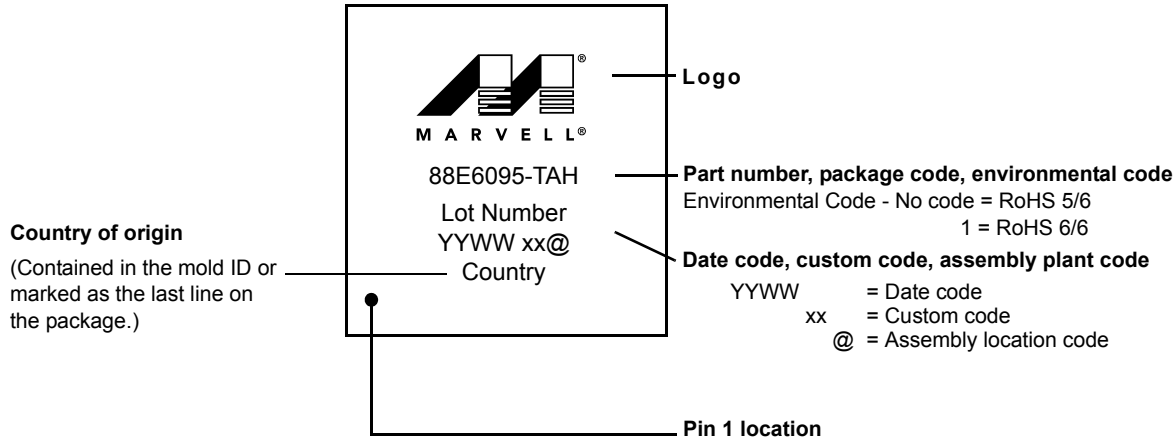

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 67 is an example of the package marking and pin 1 location for the 88E6095F 216-pin LQFP Commercial RoHS 5/6 compliant package.

Figure 67: 88E6095F 216-pin LQFP Commercial RoHS 5/6 Compliant Package Marking and Pin 1 Location

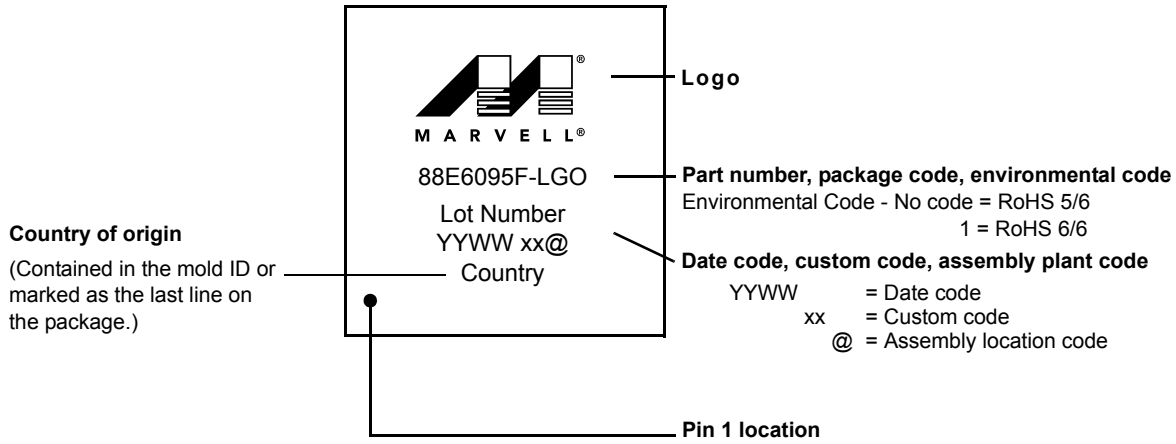

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 68 is an example of the package marking and pin 1 location for the 88E6095 176-pin TQFP Industrial RoHS 5/6 compliant package.

Figure 68: 88E6095 176-pin TQFP Industrial RoHS 5/6 Compliant Package Marking and Pin 1 Location

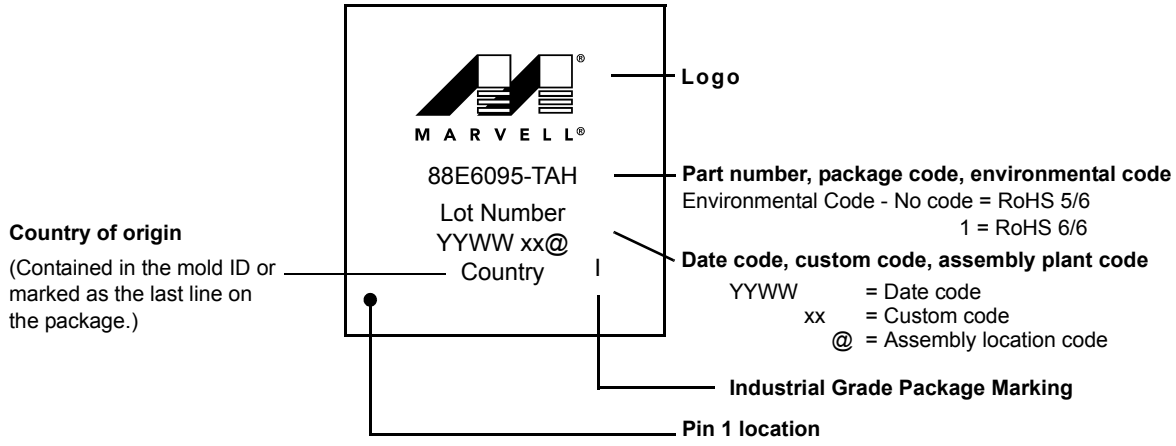

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 69 is an example of the package marking and pin 1 location for the 88E6095F 216-pin LQFP Industrial RoHS 5/6 compliant package.

Figure 69: 88E6095F 216-pin LQFP Industrial RoHS 5/6 Compliant Package Marking and Pin 1 Location

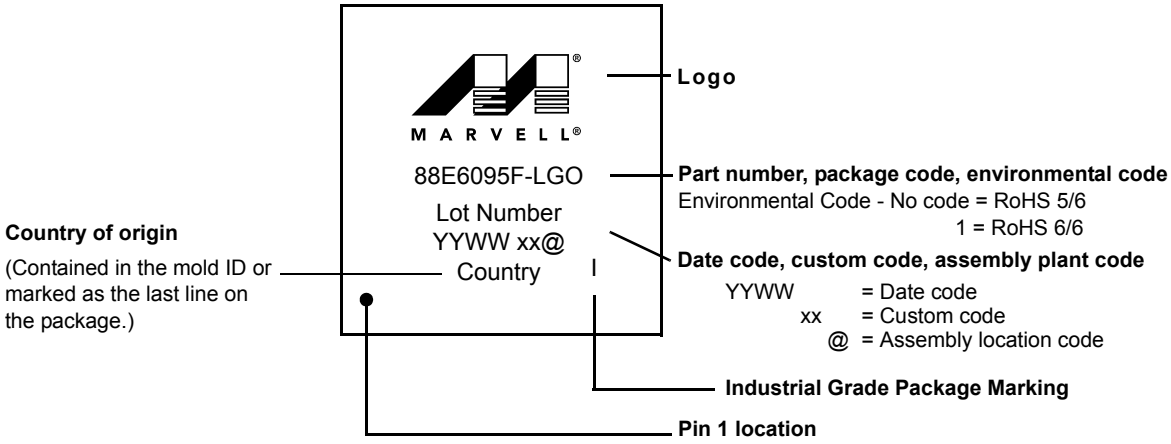

Note: The above example is not drawn to scale. Location of markings is approximate.

10.1.2 RoHS 6/6 Compliant Marking Examples

Figure 70 is an example of the package marking and pin 1 location for the 88E6092 176-pin TQFP Commercial RoHS 6/6 compliant package.

Figure 70: 88E6092 176-pin TQFP Commercial RoHS 6/6 Compliant Package Marking and Pin 1 Location

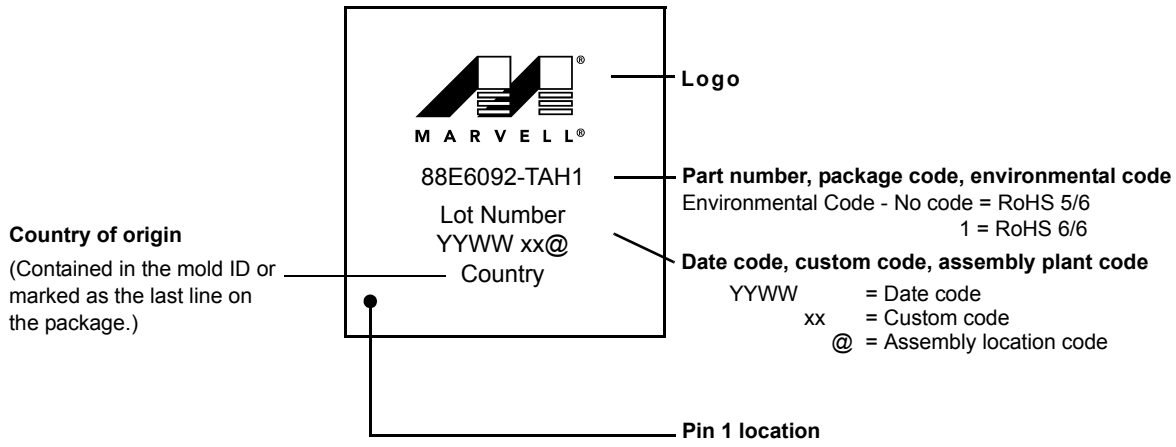

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 71 is an example of the package marking and pin 1 location for the 88E6095 176-pin TQFP Commercial RoHS 6/6 compliant package.

Figure 71: 88E6095 176-pin TQFP Commercial RoHS 6/6 Compliant Marking and Pin 1 Location

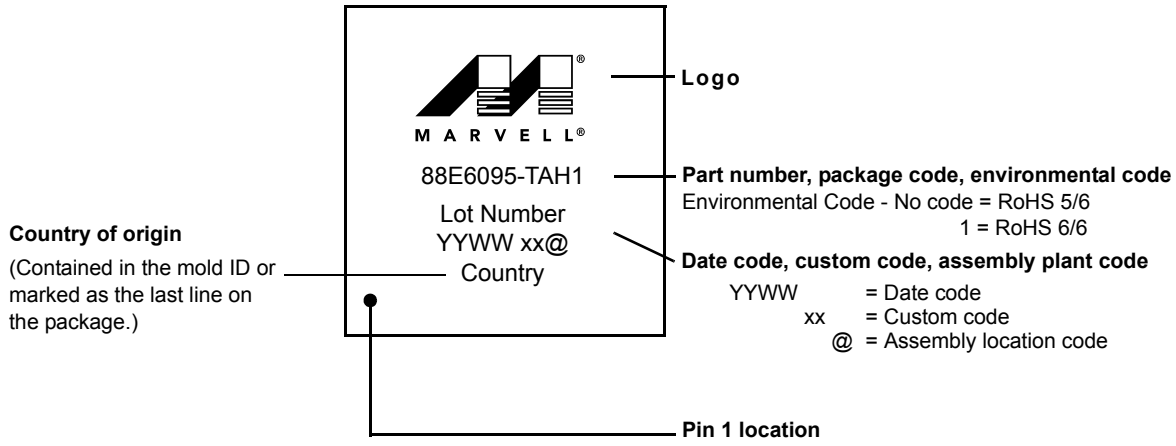

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 72 is an example of the package marking and pin 1 location for the 88E6095F 216-pin LQFP Commercial RoHS 6/6 compliant package.

Figure 72: 88E6095F 216-pin LQFP Commercial RoHS 6/6 Compliant Package Marking and Pin 1 Location

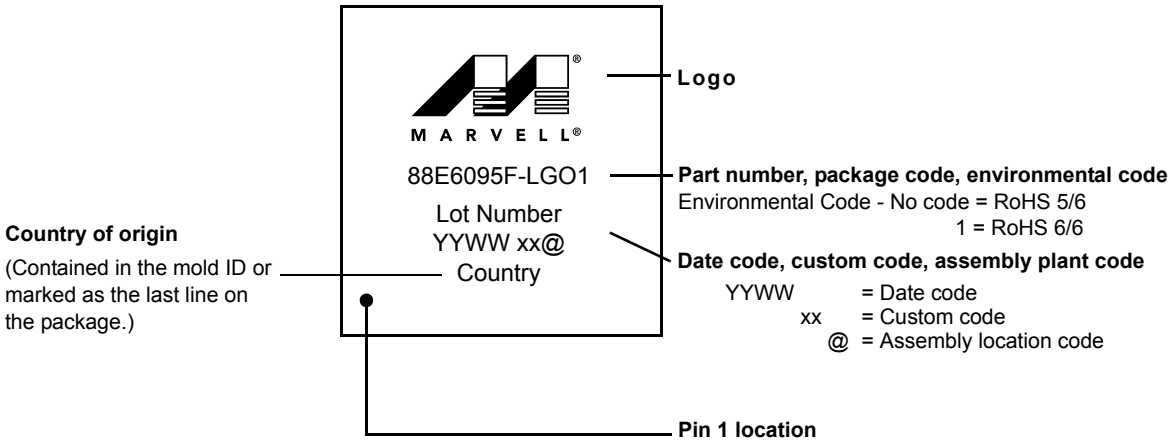

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 73 is an example of the package marking and pin 1 location for the 88E6095 176-pin TQFP Industrial RoHS 6/6 compliant package.

Figure 73: 88E6095 176-pin TQFP Industrial RoHS 6/6 Compliant Package Marking and Pin 1 Location

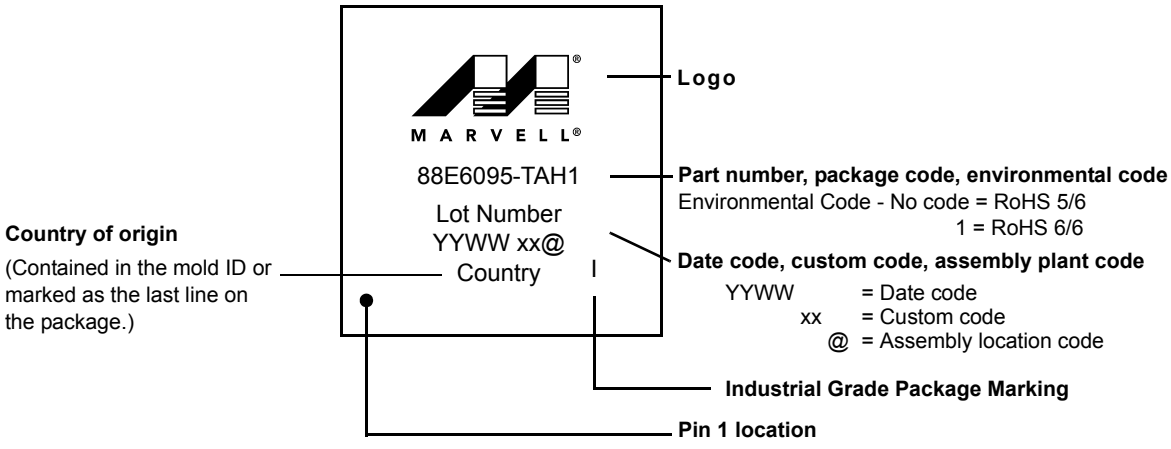

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 74 is an example of the package marking and pin 1 location for the 88E6095F 216-pin LQFP Industrial RoHS 6/6 compliant package.

Figure 74: 88E6095F 216-pin LQFP Industrial RoHS 6/6 Compliant Package Marking and Pin 1 Location

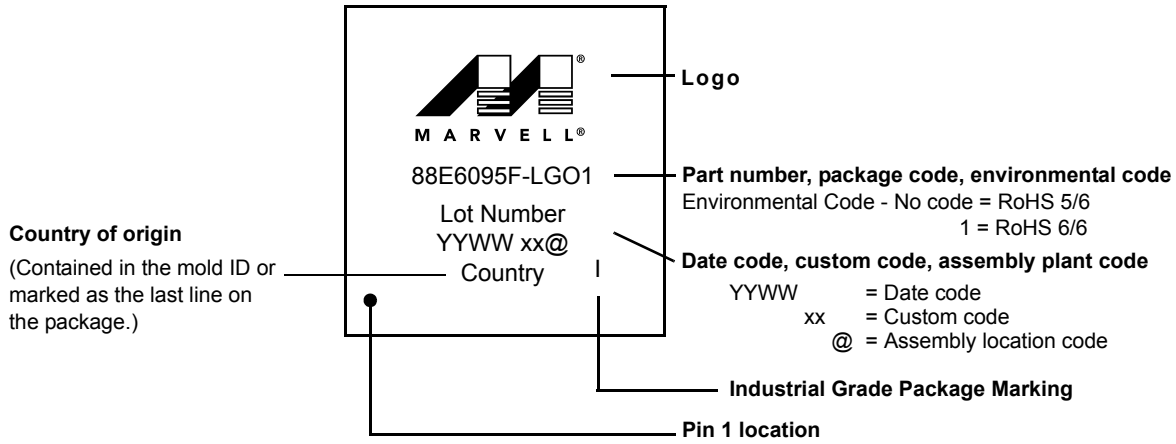

Note: The above example is not drawn to scale. Location of markings is approximate.

MARVELL®

MOVING FORWARD FASTER®

Marvell Semiconductor, Inc.

5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
Fax: 1.408.752.9028

www.marvell.com

Worldwide Corporate Offices

Marvell Semiconductor, Inc.
5488 Marvell Lane
Santa Clara, CA 95054, USA
Tel: 1.408.222.2500

Marvell Asia Pte, Ltd.
151 Lorong Chuan, #02-05
New Tech Park, Singapore 556741
Tel: 65.6756.1600
Fax: 65.6756.7600

Marvell Japan K.K.
Shinjuku Center Bldg. 44F
1-25-1, Nishi-Shinjuku, Shinjuku-ku
Tokyo 163-0644, Japan
Tel: 81.(0).3.5324.0355
Fax: 81.(0).3.5324.0354

Marvell Semiconductor Israel, Ltd.
6 Hamada Street
Mordot HaCarmel Industrial Park
Yokneam 20692, Israel
Tel: 972.(0).4.909.1500
Fax: 972.(0).4.909.1501

Marvell Semiconductor Korea, Ltd.
Rm. 603, Trade Center
159-2 Samsung-Dong, Kangnam-Ku
Seoul 135-731, Korea
Tel: 82.(0).2.551-6070/6079
Fax: 82.(0).2.551.6080

Radlan Computer Communications, Ltd.
Atidim Technological Park, Bldg. #4
Tel Aviv 61131, Israel
Tel: 972.(0).3.645.8555
Fax: 972.(0).3.645.8544

Worldwide Sales Offices

Western US
Marvell
5488 Marvell Lane
Santa Clara, CA 95054, USA
Tel: 1.408.222.2500
Fax: 1.408.752.9028
Sales Fax: 1.408.752.9029

Central US
Marvell
9600 North MoPac Drive, Suite #215
Austin, TX 78759, USA
Tel: 1.512.343.0593
Fax: 1.512.340.9970

Eastern US/Canada
Marvell
Parlee Office Park
1 Meeting House Road, Suite 1
Chelmsford, MA 01824 , USA
Tel: 1.978.250.0588
Fax: 1.978.250.0589

Europe
Marvell
5 Marchmont Gate
Boundary Way
Hemel Hempstead
Hertfordshire, HP2 7BF
United Kingdom
Tel: 44.(0).1442.211668
Fax: 44.(0).1442.211543

Israel
Marvell
6 Hamada Street
Mordot HaCarmel Industrial Park
Yokneam 20692, Israel
Tel: 972.(0).4.909.1500
Fax: 972.(0).4.909.1501

China
Marvell
5J1, 1800 Zhongshan West Road
Shanghai, PRC 200233
Tel: 86.21.6440.1350
Fax: 86.21.6440.0799

Marvell
Rm. 1102/1103, Jintian Fudi Mansion
#9 An Ning Zhuang West Rd.
Qing He, Haidian District
Beijing, PRC 100085
Tel: 86.10.8274.3831
Fax: 86.10.8274.3830

Japan
Marvell
Shinjuku Center Bldg. 44F
1-25-1, Nishi-Shinjuku, Shinjuku-ku
Tokyo 163-0644, Japan
Tel: 81.(0).3.5324.0355
Fax: 81.(0).3.5324.0354

Taiwan
Marvell
2Fl., No.1, Alley 20, Lane 407, Sec. 2
Ti-Ding Blvd., Nei Hu District
Taipei, Taiwan, 114, R.O.C
Tel: 886.(0).2.8177.7071
Fax: 886.(0).2.8752.5707

Korea
Marvell
Rm. 603, Trade Center
159-2 Samsung-Dong, Kangnam-Ku
Seoul 135-731, Korea
Tel: 82.(0).2.551-6070/6079
Fax: 82.(0).2.551.6080

For more information, visit our website at: www.marvell.com