SWITCHING

Link Street® 88E6065/88E6035 Datasheet

Low Power 6/3 Port 10/100 Switch with QoS, 802.1Q, Internal RAM and Transceivers

Doc. No. MV-S102922-00, Rev. --
March 21, 2006

MOVING FORWARD FASTER®

MARVELL®

Marvell Solutions

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. --	
Advance	Technical Publication: 1.0

OVERVIEW

The 88E6065/88E6035 devices are single-chip integrations of a complete 5+1 or 4+2 (88E6065)/2+1 or 1+2 (88E6035) port Fast Ethernet switch. These devices support 'Best-in-Class' Quality of Service (QoS) and the highest 'real-world' performance. They are uniquely suited for Routers and Gateways.

The 88E6065 device contains five 10BASE-T/100BASE-TX transceivers (PHYs), two that can be used to support 100BASE-FX, and six independent Fast Ethernet media access controllers (MACs).

The 88E6035 device contains two 10BASE-T/100BASE-TX transceivers (PHYs), one that can be used to support 100BASE-FX, and three independent Fast Ethernet media access controllers (MACs).

The 88E6065/88E6035 devices have a high-speed, non-blocking, four traffic class QoS switch fabric that uses the unique Marvell® Dynamic Queue Limit architecture. They contain a high-performance address lookup engine and a 512-Kbit frame buffer memory, and they are designed for cost-sensitive low port count, high performance, switch systems that require Quality of Service or other advanced features.

The 88E6065/88E6035 devices are designed to work in all environments. True Plug-n-Play is supported with Auto-Crossover, Auto-Polarity, and Auto-Negotiation in the PHYs, along with bridge loop prevention (using Port States).

The shared memory-based QoS switch fabric uses the latest Marvell QoS switch architecture that provides non-blocking switching performance in all traffic environments. Packets are directed into one of four traffic class queues based upon Port, IEEE 802.1p, IPv4's TOS or Diff-Serv, or IPv6's Traffic Class. Extensive priority override options are supported with separate frame and/or queue priority overrides.

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 1,024 active nodes to be connected.

The 88E6065/88E6035 devices support any 64 out of the 4,096 802.1Q VLANs, which can be enabled on a per-port basis. Three levels of 802.1Q security are supported with error frame trapping and logging.

The 88E6065/88E6035 devices' two RMII/MII/SNI interfaces support a direct connection to Management or Router CPUs with integrated MACs or to other devices. One interface is 'always on'. The second one is a multifunction port. They can be configured in either RMII mode, MII-PHY or MII-MAC mode, or SNI mode.

The 88E6065/88E6035 devices support multiple address databases, which allows packet routing without modification of the MAC address. This allows the same MAC address to exist multiple times in the MAC Address database with multiple port mappings, to completely isolate the WAN from the LAN database.

The 88E6065/88E6035 devices support a performance enhancing Marvell Header mode for wire speed steering of routed frames to either the LAN or the WAN ports. This Header places the IP portions of the frames on 32-bit aligned boundaries saving the CPU overhead of aligning the frame.

The PHY units in the 88E6065/88E6035 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 88E6065/88E6035 devices comply fully with the applicable sections of IEEE 802.3, IEEE 802.3u, and IEEE 802.3x standards.

The 88E6065/88E6035 devices' many operating modes can be configured using SMI (serial management interface - MDC/MDIO) and/or a low cost EEPROM. A stand-alone QoS mode is also supported.

HIGHLIGHTED FEATURES

- 'Best-in-Class' per port TCP/IP Ingress Rate Limiting along with independent Storm Prevention
- Non-Rate Limited frames based on VID, SA, or DA
- Wire Speed-Performance with Maximum Frame Size up to 2048 bytes
- Frame priority remapping selectable per port
- Frame priority overrides based on VID, SA or DA
- Queue priority overrides based on Port, VID, SA, DA, or ARP
- Strict or Weighted QoS selectable per port
- MAC SA based 802.1X Authentication
- Discard Frames from known bad SA
- Re-direct Frames to CPU from any SA
- Egress Provider Tag insertion (Q-in-Q)

- Recursive Ingress Provider Tag removal
- Provider Tag EtherType programmable per port
- IPv4 IGMP and IPv6 MLD Snooping
- ARP Monitoring to the CPU port
- 802.1s Per VLAN Spanning Tree
- 801.1w Rapid Spanning Tree
- Support for two-color LEDs (2-wire or 3-wire)

FEATURES

- Marvell Header for increased Routing performance & Spanning Tree support
- 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), and IPv6's Traffic Class
- Full IEEE 802.1Q VLAN ID processing per port, dynamic VLAN membership, and VLAN tagging for any 64 of the 4,096 VLAN IDs; or port based VLANs supported in any combination
- Support for multiple address data bases on a per port or per VLAN ID basis, for transparent router applications
- 30 32-bit RMON counters per port
- 6 Policy counters per port
- Egress rate shaping on all ports
- Broadcast Storm Prevention
- DA based VLAN tunneling to off load IP Phone CPU's
- Port States & BPDU support for Spanning Tree
- High performance lookup engine with support for up to 1,024 MAC address entries with automatic learning and aging
- Back-pressure based flow control on half-duplex ports & Pause-frame based flow control on full-duplex ports
- Shared 512-Kbit on-chip memory-based switch fabric with true non-blocking switching performance
- Integrated with five/two independent Fast Ethernet transceivers fully compliant with the applicable sections of IEEE802.3 and IEEE802.3u
- Automatic MDI/MDIX crossover for 100BASE-TX and 10BASE-T ports
- Ports 0 and 1[1] can be configured as copper (100BASE-TX or 10BASE-T) or fiber (100BASE-FX)
- Virtual Cable Tester® (VCT™) on all PHY ports
- Port 5 has a dedicated, always on, MAC Mode (Forward) or PHY Mode (Reverse) RMII/MII/SNI interface for management and firewall applications
- Port 4 can be a 2nd RMII/MII/SNI interface or a 10/100 PHY interface
- Integrated with six/three independent media access controllers fully compliant with the applicable sections of IEEE802.3, IEEE802.3u, and IEEE802.3x
- Each port works at 10 Mbps or 100 Mbps, full-duplex or half-duplex mode (forced or auto-negotiated)
- Flexible LED support for Link, Speed, Duplex, Collision, and Activities
- Operates at 2.5V and 3.3V
- Supports 4-Wire 93C56 / 93C66 or 2-Wire 24C01/ 24C02 / 24C04 EEPROMs
- Supports a low cost 25 MHz XTAL clock source or a 25 MHz or 50 MHz OSC
- 14x20 mm 128-Pin LQFP package
- 88E6065 is pin compatible with the Marvell 88E6060, 88E6061 and 88E6063
- 88E6035 is pin compatible with the Marvell 88E6031.
- Low power dissipation P_{AVE} = Less than 700 mW
- I-temp available in 88E6065

APPLICATIONS

- Firewall Router QoS Switch with four 10/100BASE-T LAN ports & one 10/100BASE-T WAN port (88E6065)
- Firewall Router QoS Switch with four 10/100BASE-T LAN ports and an MII port for a WAN PHY (88E6065)
- Five port QoS Switch with Spanning Tree Support (88E6065)
- Firewall Router Switch supporting a Fiber WAN port
- Multi-Dwelling Unit interface Gateway
- Media Convertor for 100BASE-FX to 100BASE-T (88E6035)
- 2-port QoS Switch for Set Top Box applications (88E6035)
- 2-port QoS Switch for IP Phones (88E6035)

1. For 88E6065 only

Figure 1: 88E6065 Top Level Block Diagram

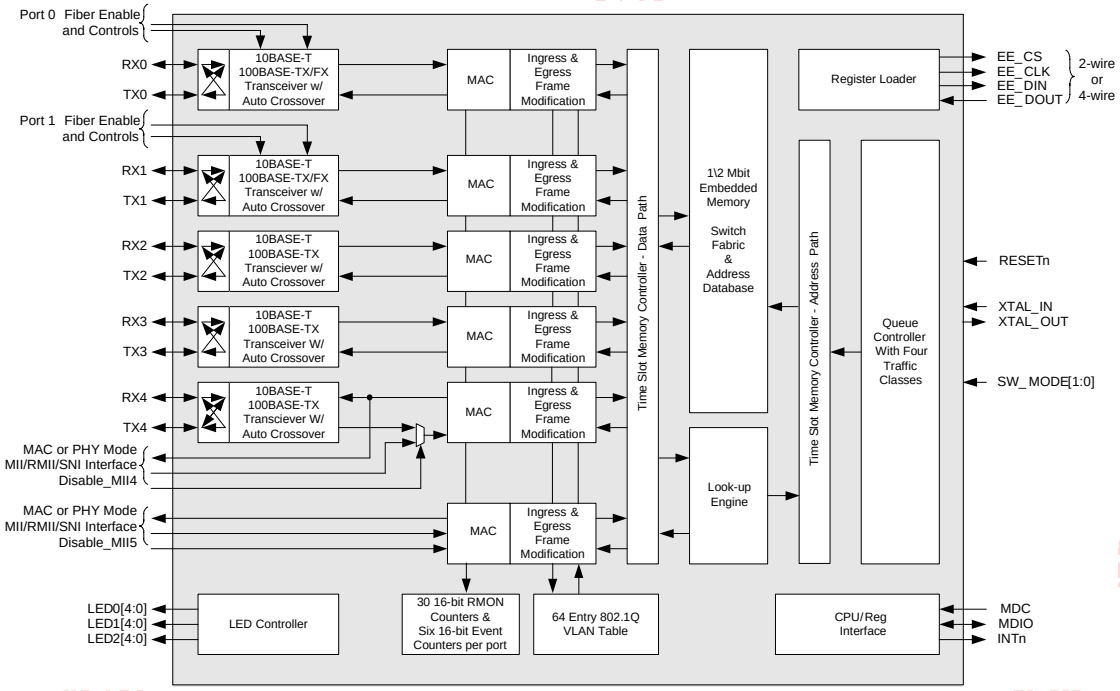

Figure 2: 88E6035 Top Level Block Diagram

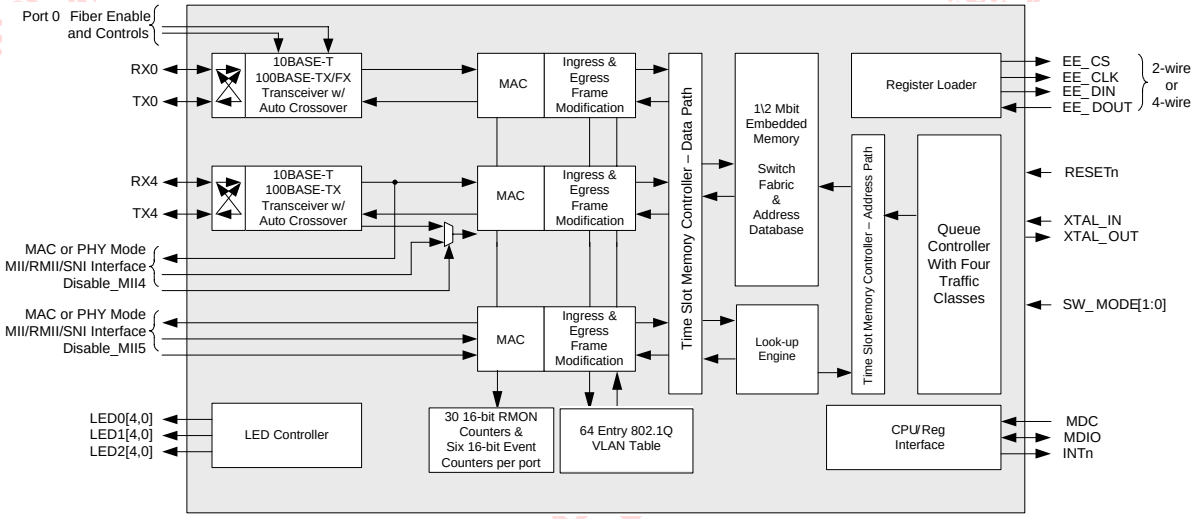

Table of Contents

List of Tables

List of Figures

Section 1. Signal Description

1.1 88E6065 Pin Diagram

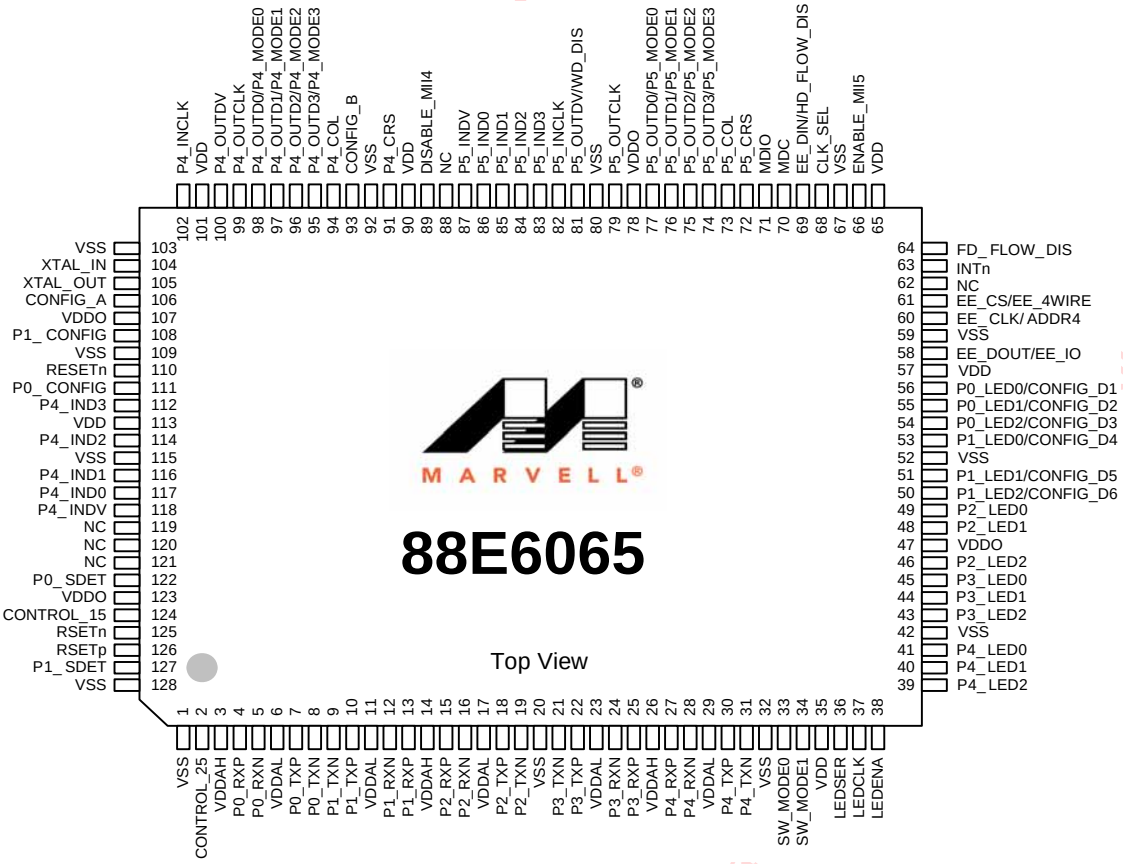

1.2 88E6035 Pin Diagram

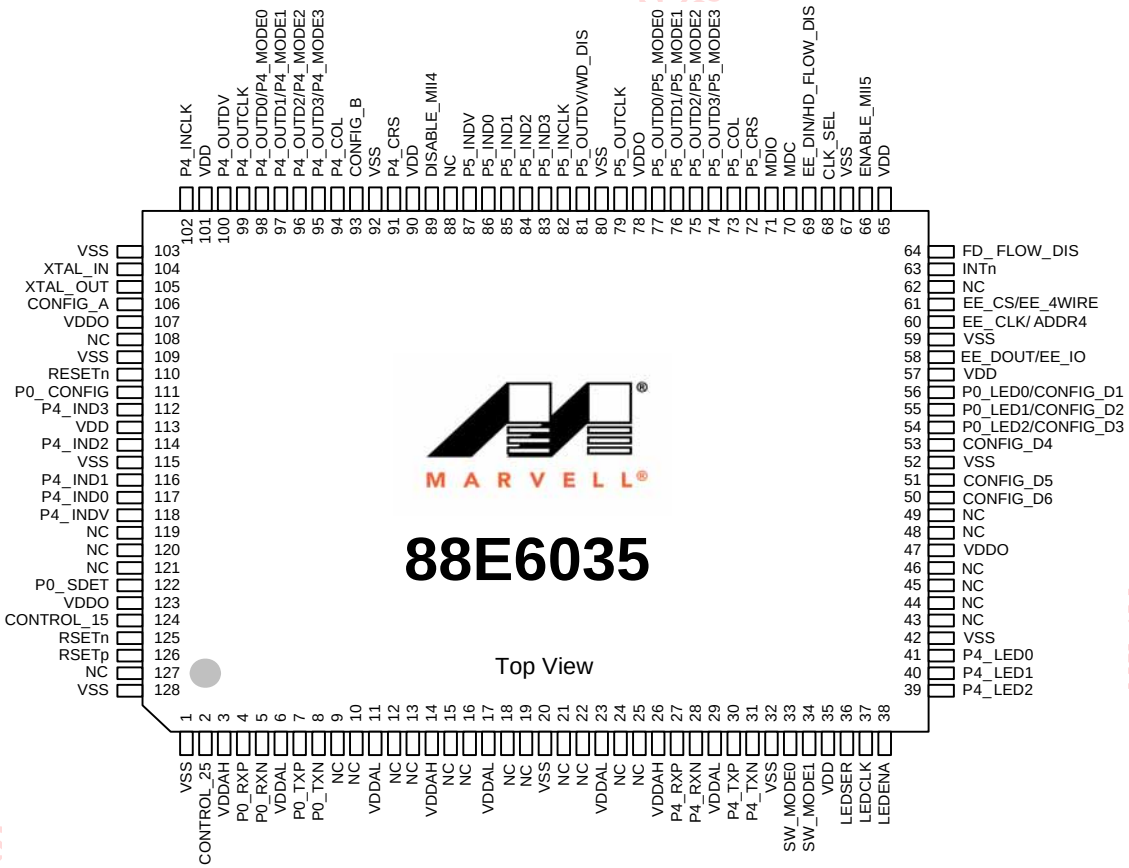

1.3 Pin Description

Table 1: Network Interface (Ports 0 to 4)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
27 25 15 13 4	27 - - - 4	P4_RXP P3_RXP P2_RXP P1_RXP P0_RXP	Typically Input	Receiver input – Positive. P[4:0]_RXP connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode (Ports 0 and 1 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-Crossover is enabled.
28 24 16 12 5	28 - - - 5	P4_RXN P3_RXN P2_RXN P1_RXN P0_RXN	Typically Input	Receiver input – Negative. P[4:0]_RXN connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode (Ports 0 and 1 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-Crossover is enabled.
30 22 18 10 7	30 - - - 7	P4_TXP P3_TXP P2_TXP P1_TXP P0_TXP	Typically Output	Transmitter output – Positive. P[4:0]_TXP connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode (Ports 0 and 1 only), P0_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-Crossover is enabled.
31 21 19 9 8	31 - - - 8	P4_TXN P3_TXN P2_TXN P1_TXN P0_TXN	Typically Output	Transmitter output – Negative. P[4:0]_TXN connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode (Ports 0 and 1 only), P0_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-Crossover is enabled.
127 122	- 122	P1_SDET P0_SDET	Input	Signal Detect input. If Port 0 and/or 1 is 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 and/or 1 is configured for 10/100BASE-T mode, Px_SDET is not used but 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 4.7 kΩ resistor.

Table 2: PHY Configuration

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
108 111	-- 111	P1_CONFIG P0_CONFIG	Input	Port 0 and 1 Configuration. The Px_CONFIG pin is used to set the default configuration for Port 0 and 1 by connecting these pins to other device pins as follows: VSS = Auto-Negotiation enabled CONFIG_D2 = Forced 10BASE-T half-duplex CONFIG_D3 = Forced 10BASE-T full-duplex CONFIG_D4 = Forced 100BASE-TX half-duplex CONFIG_D5 = Forced 100BASE-TX full-duplex CONFIG_D6 = Forced 100BASE-FX half-duplex VDDO = Forced 100BASE-FX full-duplex Port 2[1], 3, and 4's default configuration is Auto-Negotiation enabled. 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.
106	106	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 CONFIG_D1 = LED Mode 0, FEFI enabled CONFIG_D2 = LED Mode 1, FEFI disabled CONFIG_D3 = LED Mode 1, FEFI enabled CONFIG_D4 = LED Mode 2, FEFI disabled CONFIG_D5 = LED Mode 2, FEFI enabled CONFIG_D6 = LED Mode 3, FEFI disabled VDDO = LED Mode 3, FEFI enabled The CONFIG_A pin is configured after reset and contains an internal pull-up resistor.

Table 2: PHY Configuration (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
93	93	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, Back Plane, Energy Detect disabled CONFIG_D1 = No Crossover, Back Plane, Energy Detect enabled CONFIG_D2 = No Crossover, CAT 5, Energy Detect disabled CONFIG_D3 = No Crossover, CAT 5, Energy Detect enabled CONFIG_D4 = Auto-Crossover, Back Plane, Energy Detect disabled CONFIG_D5 = Auto-Crossover, Back Plane, Energy Detect enabled CONFIG_D6 = Auto-Crossover, CAT 5, Energy Detect disabled VDDO = Auto-Crossover, CAT 5, Energy Detect enabled The CONFIG_B pin is configured after reset and contains an internal pull-up resistor.
56 55 54 53 51 50	56 55 54 53 51 50	CONFIG_D1 CONFIG_D2 CONFIG_D3 CONFIG_D4 CONFIG_D5 CONFIG_D6	Output	These pins are used to configure the PHY portion of the device by connecting to CONFIG_A and/or CONFIG_B. These pins are configured after reset and contain an internal pull-up resistor. **NOTE:** On the 88E6035, these pins are used as LED pins. For further details, refer to Table 14.

1. Port 1, 2, and 3 are valid for 88E6065 only.

Table 3: Regulator & Reference

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
126	126	RSETp	Analog	Resistor reference. A 2 kΩ 1% resistor is placed between RSETp and RSETn. This resistor is used to set an internal bias reference current.
125	125	RSETn	Analog	Resistor reference. A 2 kΩ 1% resistor is placed between RSETp and RSETn. This resistor is used to set an internal bias reference current.
124	124	CONTROL_15	Analog	Voltage control to external 1.5V regulator. This signal controls an external PNP transistor to generate the 1.5V power supply for the VDD and VDDAL 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 VDDO.
2	2	CONTROL_25	Analog	Voltage control to external 2.5V regulator. This signal controls an external PNP transistor to generate the 2.5V power supply 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 2.5V and the emitter is connected to VDDO.

Table 4: System

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
104	104	XTAL_IN	Input	25 MHz system reference clock input provided from the board. The frequency of this clock input is selected by the CLK_SEL pin. The clock source can come from an external crystal or an external oscillator. This is the only clock required as it is used for both the switch and the PHYs.
105	105	XTAL_OUT	Output	Output System reference clock output provided to the board. This output can only be used to drive an external crystal (25 MHz only). It cannot be used to drive external logic. If an external oscillator is used this pin should be left unconnected.

Table 4: System (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
68	68	CLK_SEL	Input	Input Clock frequency Select. Connect this pin to VSS if XTAL_IN is 25 MHz. Connect this pin to VDDO or leave it unconnected if XTAL_IN is 50 MHz. This pin must be stable before and after Reset. CLK_SEL is internally pulled high via a resistor so it can be left floating to select 50 MHz.
110	110	RESETn	Input	Hardware reset. Active low. The 88E6065/88E6035 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.

Table 5: Register Access Interface

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
70	70	MDC	Input	MDC is the management data clock reference for the serial management interface (SMI). A continuous clock stream is not expected. The maximum frequency supported is 8.3 MHz. The SMI is used to access the registers in the PHY and in the Switch and it is available in all combinations of SW_MODE[1:0] (Table 13). MDC is internally pulled high via a resistor so it can be left floating when unused.
71	71	MDIO	I/O	MDIO is the management data. MDIO is used to transfer management data in and out of the device synchronously to MDC. This pin requires an external pull-up resistor in the range of 1.5 kΩ to 10 kΩ. The 88E6065/88E6035 device uses 16 of the 32 possible SMI port addresses. The 16 that are used are selectable using the EE_CLK/ADDR4 pin. MDIO is internally pulled high via a resistor so it can be left floating when unused.
63	63	INTn	Open Drain Output	INTn is an active low, open drain pin that is asserted to indicate an unmasked interrupt event occurred. A single external pull-up resistor is required somewhere on this interrupt net for it to go high when it is inactive. The INTn pin will go active low if SW_MODE[1:0] (Table 13) are not 0x2 (i.e., the Stand Alone mode) and if the EEPROM data has been completely read into the device.

Table 6: Serial EEPROM Interface

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
61	61	EE_CS/ EE_4WIRE	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. EE_CS is a multi-function pin used to configure the 88E6065/88E6035 device during a hardware reset. 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 = It supports 2-WIRE (For 1K 24C01A, 2K bit 24C02 & 4K bit 24C04) • High = It supports 4-WIRE (for 2K bit 93C56 & 4K bit 93C66) The external EEPROM must be configured in the x16 organization using 8-bit addresses. EE_CS is internally pulled high via a resistor so the pin can be left floating when unused or to select support for 4-Wire EEPROMs. Use a 4.7 kΩ resistor to VSS for a configuration low.
60	60	EE_CLK/ ADDR4	Typically Output	Serial EEPROM clock. EE_CLK is the serial EEPROM clock reference output by the 88E6065/88E6035 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. EE_CLK is a multi-function pin used to configure the 88E6065/88E6035 device during a hardware reset. When reset is asserted, EE_CLK becomes an input and the desired SMI ADDR4 address space configuration is latched at the rising edge of RESETn as follows: • Low = Use SMI device addresses 0x00 to 0x0F • High = Use SMI device addresses 0x10 to 0x1F 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 6: Serial EEPROM Interface (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
69	69	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). After the serial EEPROM data is loaded EE_DIN becomes Half-duplex Flow Control disable as follows: • Low = Enable “forced collision” flow control on all half-duplex ports • High = Disable flow control on all half-duplex ports 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.
58	58	EE_DOUT/ EE_IO[1]	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.

1. The address is sent out the EE_IO pin only for 2-Wire EEPROMs.

Table 7: Port 5's MII Interface Enable

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
66	66	ENABLE_MII5	Input	Enable MII Port 5. This pin is used to enable Port 5's MII Interface drivers (Table 8). A high enables Port 5's MII interface drivers and enables Link on Port 5. A low disables Port 5's MII interface drivers (i.e., they are tri-stated) and disables Link on Port 5. The function of this pin can be overwritten by the ForcedLink and LinkValue bits in the port's MAC Control Register (Offset 0x01). ENABLE_P5 is internally pulled high via a resistor so the pin can be left unconnected to enable Port 5's MII interface.

Table 8: Port 5's Input MII Interface - If ENABLE_P5 = High

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
82	82	P5_INCLK	I/O	Input Clock. P5_INCLK is a reference for P5_INDV, P5_INER, and P5_IND[3:0]. The direction and speed of P5_INCLK is determined by P5_MODE[3:0] (Table 9) at the end of RESETn. If the port is in PHY Mode, P5_INCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-TX mode, 2.5 MHz if the port is in 10BASE-T mode, and 50 MHz for RMII mode. If the port is in MAC Mode, P5_INCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25 MHz or 50 MHz although it should be 25 MHz for 100BASE-X mode and 2.5 MHz for 10BASE-T mode. 50 MHz is needed for RMII mode. P5_INCLK is tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down) and it is internally pulled high so the pin can be left unconnected if not used.
83 84 85 86	83 84 85 86	P5_IND3 P5_IND2 P5_IND1 P5_IND0	Input	Input Data. P5_IND[3:0] receives the data nibble to be transmitted into the switch in 100BASE-TX and 10BASE-T modes. P5_IND[3:0] is synchronous to P5_INCLK. These pins are inputs regardless of the port's mode (i.e., PHY mode or MAC mode). Only P5_IND0 is used when SNI mode is selected. P5_IND[1:0] are used when RMII mode is selected. P5_IND[3:0] are internally pulled high via resistor so the pins can be left unconnected when they are not used.
87	87	P5_INDV	Input	Input Data Valid. When P5_INDV is asserted high, data on P5_IND[3:0] along with P5_INER is encoded and transmitted into the switch. P5_INDV must be synchronous to P5_INCLK. P5_INDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 9: Port 5's Output MII Interface - If ENABLE_P5 = High

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
79	79	P5_OUTCLK	I/O	Output Clock. P5_OUTCLK is a reference for P5_OUTDV and P5_OUTD[3:0]. The direction and speed of P5_OUTCLK is determined by P5_MODE[3:0] at the end of RESETn. If the port is in PHY Mode, P5_OUTCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-TX mode, 2.5 MHz if the port is in 10BASE-T mode, and 50 MHz for RMII mode. If the port is in MAC Mode, P5_OUTCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25 MHz or 50 MHz although it should be 25 MHz for 100BASE-TX mode 2.5 MHz for 10BASE-T mode, and 50 MHz for RMII mode. P5_OUTCLK is tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down) and it is internally pulled high so the pin can be left unconnected if not used. P5_OUTCLK is tri-stated when ENABLE_P5 = Low.
74 75 76 77	74 75 76 77	P5_OUTD3/ P5_MODE3 P5_OUTD2/ P5_MODE2 P5_OUTD1/ P5_MODE1 P5_OUTD0/ P5_MODE0	Normally Output Input only when RESETn is low	Output Data. Data transmitted from the switch is decoded and presented on P5_OUTD[3:0] pins synchronous to P5_OUTCLK. These pins are outputs regardless of the port's mode (i.e., PHY or MAC mode). The Px_OUTD[3:0]/Px_MODE[3:0] pins are used to select MII/SNI/RMII modes. See Table 16 for RMII/MII/SNI Configuration options details. Only P5_OUTD0 contains meaningful data when SNI mode is selected. P5_OUTD[1:0] are used when RMII mode is selected. During reset, these pins are internally pulled high, and tri-stated, and are used to latch in the desired operating mode for the port. P5_OUTD is tri-stated when ENABLE_P5 = Low. These pins are tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down).

Table 9: Port 5's Output MII Interface - If ENABLE_P5 = High (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
81	81	P5_OUTDV/ WD_DIS	Normally Output Input only when RESETn is low	Output Data Valid. When P5_OUTDV is asserted high, data transmitted from the switch is decoded and presented on P5_OUTD[3:0]. P5_OUTDV is synchronous to P5_OUTCLK. During reset this internally pulled high pin is used to latch the desired WatchDog mode. If this pin is pulled low during RESETn, the automatic WatchDog reset monitor is enabled. This pin is tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down).
72	72	P5_CRS	I/O	Carrier Sense. After reset, P5_CRS becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P5_CRS asserts (or is expected to be asserted) when the receive data path is non-idle. In half-duplex mode P5_CRS is also asserted (or is expected to be asserted) during transmission. P5_CRS is asynchronous to P5_OUTCLK and P5_INCLK. P5_CRS is tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down) and it is internally pulled low so the pin can be left unconnected if not used. P5_CRS is tri-stated when ENABLE_P5 = Low.
73	73	P5_COL	I/O	Collision. After reset, P5_COL becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P5_COL asserts (or is expected to be asserted) only in half-duplex mode when both the transmit and receive paths are non-idle. In full-duplex mode, P5_COL is always low (or it is ignored). P5_COL is asynchronous to P5_OUTCLK and P5_INCLK. P5_COL is tri-stated during RESETn (by ENABLE_MII5 or by forcing the port's link down) and it is internally pulled low so the pin can be left unconnected if not used. P5_COL is tri-stated when ENABLE_P5 = Low.

Table 10: Port 4's MII Interface Enable

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
89	89	DISABLE_MII4	Input	Disable MII Port 4. This pin is used to disable Port 4's MII Interface drivers (Table 11). A high disables Port 4's drivers (i.e., they are tri-stated) and enables Port 4's PHY interface (its MDI pins). A low enables Link up on Port 4 and enables Port 4's MII Interface and its drivers and disables Port 4's PHY interface. The low function of this pin (described above) can be overwritten by the ForcedLink and LinkValue bits in the port's MAC Control Register (Offset 0x01). DISABLE_P4 is internally pulled high via a resistor so the pin can be left unconnected to disable Port 4's MII Interface.

Table 11: Port 4's Input MII Interface - If DISABLE_P4 = Low

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
102	102	P4_INCLK	I/O	Input Clock. P4_INCLK is a reference for P4_INDV, P4_INER, and P4_IND[3:0]. The direction and speed of P4_INCLK is determined by P4_MODE[3:0] (Table 12) at the end of RESETn. If the port is in PHY Mode, P4_INCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-TX mode, 2.5 MHz if the port is in 10BASE-T mode, and 50 MHz for RMII mode. If the port is in MAC Mode, P4_INCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25 MHz or 50 MHz although it should be 25 MHz for 100BASE-TX mode and 2.5 MHz for 10BASE-T mode. 50 MHz is needed for RMII mode. P4_INCLK is tri-stated during RESETn (by DISABLE_MII4 or by forcing the port's link down) and it is internally pulled high so the pin can be left unconnected if not used.
112 114 116 117	112 114 116 117	P4_IND3 P4_IND2 P4_IND1 P4_IND0	Input	Input Data. P4_IND[3:0] receives the data nibble to be transmitted into the switch in 100BASE-X and 10BASE-T modes. P4_IND[3:0] is synchronous to P4_INCLK. These pins are inputs regardless of the port's mode (i.e., PHY mode or MAC mode). Only P4_IND0 is used when SNI mode is selected. P4_IND[1:0] are used when RMII mode is selected. P4_IND[3:0] are internally pulled high via resistor so the pins can be left unconnected when they are not used.
118	118	P4_INDV	Input	Input Data Valid. When P4_INDV is asserted high, data on P4_IND[3:0] along with P4_INER is encoded and transmitted into the switch. P4_INDV must be synchronous to P4_INCLK. P4_INDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.

Table 12: Port 4's Output MII Interface - If DISABLE_P4 = Low

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
99	99	P4_OUTCLK	I/O	Output Clock. P4_OUTCLK is a reference for P4_OUTDV and P4_OUTD[3:0]. The direction and speed of P4_OUTCLK is determined by P4_MODE[3:0] at the end of RESETn. If the port is in PHY Mode, P4_OUTCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-X mode, 2.5 MHz if the port is in 10BASE-T mode, and 50 MHz for RMII mode. If the port is in MAC Mode, P4_OUTCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25 MHz or 50 MHz although it should be 25 MHz for 100BASE-X mode, 2.5 MHz for 10BASE-T mode, and 50 MHz for RMII mode. P4_OUTCLK is tri-stated during RESETn (by DISABLE_MII4 or by forcing the port's link down) and it is internally pulled high so the pin can be left unconnected if not used. P4_OUTCLK is tri-stated when DISABLE_P4 = High.
95 96 97 98	95 96 97 98	P4_OUTD3/ P4_MODE3 P4_OUTD2/ P4_MODE2 P4_OUTD1/ P4_MODE1 P4_OUTD0/ P4_MODE0	Normally Output Input only when RESETn is low	Output Data. Data transmitted from the switch is decoded and presented on the P4_OUTD[3:0] pins synchronous to P4_OUT_CLK. These pins are outputs regardless of the port's mode (i.e., PHY or MAC mode). The Px_OUTD[3:0]/Px_MODE[3:0] pins are used to select MII/SNI/RMII modes. See Table 16 for RMII/MII/SNI Configuration options details. Only P4_OUTD0 contains meaningful data when SNI mode is selected. P4_OUTD[1:0] are used when RMII mode is selected. During RESETn these internally pulled high pins are tri-stated and used to latch in the desired operating mode for the port. These pins are tri-stated during RESETn (and by DISABLE_MII4 or by forcing the port's link down).
100	100	P4_OUTDV	Output	Output Data Valid. When P4_OUTDV is asserted high, data transmitted from the switch is decoded and presented on P4_OUTD[3:0]. P4_OUTDV is synchronous to P4_OUTCLK. This pin is tri-stated during RESETn (by DISABLE_MII4 or by forcing the port's link down).

Table 12: Port 4's Output MII Interface - If DISABLE_P4 = Low (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
91	91	P4_CRS	I/O	Carrier Sense. After reset, P4_CRS becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P4_CRS asserts (or is expected to be asserted) when the receive data path is non-idle. In half-duplex mode P4_CRS is also asserted (or is expected to be asserted) during transmission. P4_CRS is asynchronous to P4_OUTCLK and P4_INCLK. P4_CRS is tri-stated during RESETn (by DISABLE_MII4 or by forcing the port's link down) and it is internally pulled low so the pin can be left unconnected if not used. P4_CRS is tri-stated when DISABLE_P4 = High.
94	94	P4_COL	I/O	Collision. After reset, P4_COL becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P4_COL asserts (or is expected to be asserted) only in half-duplex mode when both the transmit and receive paths are non-idle. In full-duplex mode, P4_COL is always low (or it is ignored). P4_COL is asynchronous to P4_OUTCLK and P4_INCLK. P4_COL is tri-stated during RESETn (by DISABLE_MII4 or by forcing the port's link down) and it is internally pulled low so the pin can be left unconnected if not used. P4_COL is tri-stated when DISABLE_P4 = High.

Table 13: Switch Configuration Interface

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
34 33	34 33	SW_MODE1 SW_MODE0	Input	Switch Mode. These pins are used to configure the 88E6065/ 88E6035 device after reset. The Switch Mode pins work as follows: 1 0 Description 0 0 CPU attached mode – ports come up disabled[1] 0 1 Reserved 1 0 Test mode– The EEPROM is ignored 1 1 EEPROM attached mode[2] The EEPROM attached mode (when both SW_MODE pins = high) can be used with a CPU. An EEPROM can be used in CPU attached mode (when both SW_MODE pins = low). The only difference between these two modes is the initial setting of the port's Port State bits. In CPU attached mode, the ports come up disabled so that the CPU can fully configure the switch before packets are allowed to flow. This is critical in LAN to WAN routers to prevent a 'leak' between the LAN and the WAN, while the CPU is booting. In both of 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.

Table 13: Switch Configuration Interface (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
64	64	FD_FLOW_DIS	Input	Full-duplex Flow Control disable. High = Disable flow control on all full-duplex ports Low = Enable IEEE 802.3x Pause based flow control on all supported full-duplex ports Full-duplex flow control requires support from the end station. Full-duplex flow control 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 supports Pause as well (from data returned during Auto-Negotiation). If any of these requirements are not met the port will not generate full-duplex Pause frames and it will not pause when the port receives full-duplex Pause frames. FD_FLOW_DIS is latched on the rising edge of RESETn into all the port's PHY Auto-Negotiation Advertisement registers. After reset, the value on this pin can be read in the port's Port Status Register (Offset 0x00). It is internally pulled high via a resistor so the pin can be left unconnected to disable full-duplex flow control. The EE_DIN/HD_FLOW_DIS multi-function pin is used to select Flow control on half-duplex ports (Table 6).

1. The ports come up disabled in the CPU mode so the CPU can fully boot and configure the switch before packets are allowed to flow. This is required for WAN to LAN routers or for CPUs that perform bridge loop detection on link up and configure the switch before packets are allowed to flow.
2. In EEPROM attached mode the ports come up enabled unless the Port Control register (Offset 0x04) is overwritten by the EEPROM data (see Section 5.3 "EEPROM Programming Format"). This mode is for unmanaged switches without a CPU attached to the switch. It can also be used to debug a new managed switch or router designs so the switch will work without any software. If this is done for debug reasons, ensure that the production units use the CPU attached mode.

Table 14: Port Status LEDs

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
39 43 46 50 54	39 -- -- -- 54	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. P[4:0]_LED2 are driven active low whenever RESETn is active low. **NOTE:** Some of the LED pins are also used as configuration pins as documented in Table 2.
40 44 48 51 55	40 -- -- -- 55	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. P[4:0]_LED1 are driven active low whenever RESETn is active low. **NOTE:** Some of the LED pins are also used as configuration pins as documented in Table 2.
41 45 49 53 56	41 -- -- -- 56	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. P[4:0]_LED0 are driven active low whenever RESETn is active low. **NOTE:** Some of the LED pins are also used as configuration pins as documented in Table 2.
36	36	LEDSER	Output	LEDSER outputs serial status bits that can be shifted into a shift register to be displayed via LEDs. LEDSER is output synchronously to LEDCLK.
38	38	LEDENA	Output	LEDENA asserts High whenever LEDSER has valid status that is to be stored into the shift register. LEDENA is output synchronously to LEDCLK.
37	37	LEDCLK	Output	LEDCLK is the reference clock for the serial LED signals.

Table 15: Power & Ground

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
35 57 65 90 101 113	35 57 65 90 101 113	VDD	Power	1.5V Power to digital core
3 14 26	3 14 26	VDDAH	Power	2.5V Power to analog core
6 11 17 23 29	6 11 17 23 29	VDDAL	Power	1.5V Power to analog core
47 78 107 123	47 78 107 123	VDDO	Power	3.3V/2.5V Power to outputs. VDDO can be connected to 3.3V to support 3.3V I/O or 2.5V to support 2.5V I/O.

Table 15: Power & Ground (Continued)

88E6065 Pin #	88E6035 Pin #	Pin Name	Type	Description
1 20 32 42 52 59 67 80 92 103 109 115 128	1 20 32 42 52 59 67 80 92 103 109 115 128	VSS	Ground	Ground to device
62 88 119 120 121	9 10 12 13 15 16 18 19 21 22 24 25 43 44 45 46 48 49 62 88 108 119 120 121 127	NC	NC	No Connect. Do not connect these pins to anything. These pins must be left unconnected.

1.4 88E6065 128-Pin LQFP Pin Assignments (Alphabetical by Signal Name)

Pin Name	Pin #
CLK_SEL	68
CONFIG_A	106
CONFIG_B	93
CONTROL_15	124
CONTROL_25	2
DISABLE_MII4	89
EE_CLK/ADDR4	60
EE_CS/EE_4WIRE	61
EE_DIN/ HD_FLOW_DIS	69
EE_DOUT/EE_IO	58
ENABLE_MII5	66
FD_FLOW_DIS	64
INTn	63
LEDCLK	37
LEDENA	38
LEDSER	36
MDC	70
MDIO	71
NC	62
NC	88
NC	119
NC	120
NC	121
P0_CONFIG	111
P0_LED0/CONFIG_D1	56
P0_LED1/CONFIG_D2	55
P0_LED2/CONFIG_D3	54
P0_RXP	4
P0_RXN	5
P0_SDET	122

Pin Name	Pin #
P0_TXP	7
P0_TXN	8
P1_CONFIG	108
P1_LED0/CONFIG_D4	53
P1_LED1/CONFIG_D5	51
P1_LED2/CONFIG_D6	50
P1_RXP	13
P1_RXN	12
P1_SDET	127
P1_TXP	10
P1_TXN	9
P2_LED0	49
P2_LED1	48
P2_LED2	46
P2_RXP	15
P2_RXN	16
P2_TXP	18
P2_TXN	19
P3_LED0	45
P3_LED1	44
P3_LED2	43
P3_RXP	25
P3_RXN	24
P3_TXP	22
P3_TXN	21
P4_COL	94
P4_CRS	91
P4_INCLK	102
P4_IND0	117
P4_IND1	116

Pin Name	Pin #
P4_IND2	114
P4_IND3	112
P4_INDV	118
P4_LED0	41
P4_LED1	40
P4_LED2	39
P4_OUTCLK	99
P4_OUTD0/P4_MODE0	98
P4_OUTD1/P4_MODE1	97
P4_OUTD2/P4_MODE2	96
P4_OUTD3/P4_MODE3	95
P4_OUTDV	100
P4_RXP	27
P4_RXN	28
P4_TXP	30
P4_TXN	31
P5_COL	73
P5_CRS	72
P5_INCLK	82
P5_IND0	86
P5_IND1	85
P5_IND2	84
P5_IND3	83
P5_INDV	87
P5_OUTCLK	79
P5_OUTD0/P5_MODE0	77
P5_OUTD1/P5_MODE1	76
P5_OUTD2/P5_MODE2	75
P5_OUTD3/P5_MODE3	74
P5_OUTDV/WD_DIS	81

Pin Name	Pin #
RESETn	110
RSETp	126
RSETn	125
SW_MODE0	33
SW_MODE1	34
VDD	35
VDD	57
VDD	65
VDD	90
VDD	101
VDD	113
VDDAH	3
VDDAH	14

Pin Name	Pin #
VDDAH	26
VDDAL	6
VDDAL	11
VDDAL	17
VDDAL	23
VDDAL	29
VDDO	47
VDDO	78
VDDO	107
VDDO	123
VSS	1
VSS	20
VSS	32

Pin Name	Pin #
VSS	42
VSS	52
VSS	59
VSS	67
VSS	80
VSS	92
VSS	103
VSS	109
VSS	115
VSS	128
XTAL_IN	104
XTAL_OUT	105

1.5 88E6035 128-Pin LQFP Pin Assignments (Alphabetical by Signal Name)

Pin Name	Pin #
CLK_SEL	68
CONFIG_A	106
CONFIG_B	93
CONFIG_D4	53
CONFIG_D5	51
CONFIG_D6	50
CONTROL_15	124
CONTROL_25	2
DISABLE_MII4	89
EE_CLK/ADDR4	60
EE_CS/EE_4WIRE	61
EE_DIN/ HD_FLOW_DIS	69
EE_DOUT/EE_IO	58
ENABLE_MII5	66
FD_FLOW_DIS	64
INTn	63
LEDCLK	37
LEDENA	38
LEDSER	36
MDC	70
MDIO	71
NC	9
NC	10
NC	12
NC	13
NC	15
NC	16
NC	18
NC	19
NC	21

Pin Name	Pin #
NC	22
NC	24
NC	25
NC	43
NC	44
NC	45
NC	46
NC	48
NC	49
NC	62
NC	88
NC	108
NC	119
NC	120
NC	121
NC	127
P0_CONFIG	111
P0_LED0/CONFIG_D1	56
P0_LED1/CONFIG_D2	55
P0_LED2/CONFIG_D3	54
P0_RXP	4
P0_RXN	5
P0_SDET	122
P0_TXP	7
P0_TXN	8
P4_COL	94
P4_CRS	91
P4_INCLK	102
P4_IND0	117
P4_IND1	116

Pin Name	Pin #
P4_IND2	114
P4_IND3	112
P4_INDV	118
P4_LED0	41
P4_LED1	40
P4_LED2	39
P4_OUTCLK	99
P4_OUTD0/P4_MODE0	98
P4_OUTD1/P4_MODE1	97
P4_OUTD2/P4_MODE2	96
P4_OUTD3/P4_MODE3	95
P4_OUTDV	100
P4_RXP	27
P4_RXN	28
P4_TXP	30
P4_TXN	31
P5_COL	73
P5_CRS	72
P5_INCLK	82
P5_IND0	86
P5_IND1	85
P5_IND2	84
P5_IND3	83
P5_INDV	87
P5_OUTCLK	79
P5_OUTD0/P5_MODE0	77
P5_OUTD1/P5_MODE1	76
P5_OUTD2/P5_MODE2	75
P5_OUTD3/P5_MODE3	74
P5_OUTDV/WD_DIS	81

Pin Name	Pin #
RESETn	110
RSETp	126
RSETn	125
SW_MODE0	33
SW_MODE1	34
VDD	35
VDD	57
VDD	65
VDD	90
VDD	101
VDD	113
VDDAH	3
VDDAH	14

Pin Name	Pin #
VDDAH	26
VDDAL	6
VDDAL	11
VDDAL	17
VDDAL	23
VDDAL	29
VDDO	47
VDDO	78
VDDO	107
VDDO	123
VSS	1
VSS	20
VSS	32

Pin Name	Pin #
VSS	42
VSS	52
VSS	59
VSS	67
VSS	80
VSS	92
VSS	103
VSS	109
VSS	115
VSS	128
XTAL_IN	104
XTAL_OUT	105

Section 2. Switch Core Functional Description

2.1 Switch Data Flow

The 88E6065/88E6035 device accepts 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 3 and Figure 4 show the data paths inside the switch along with the major functional blocks that process the frame as it travels through the 88E6065 and the 88E6035 device. Each of these blocks along with their register-controllable options and policy are described in the sections which follow.

Figure 3: 88E6065 Switch Data Flow

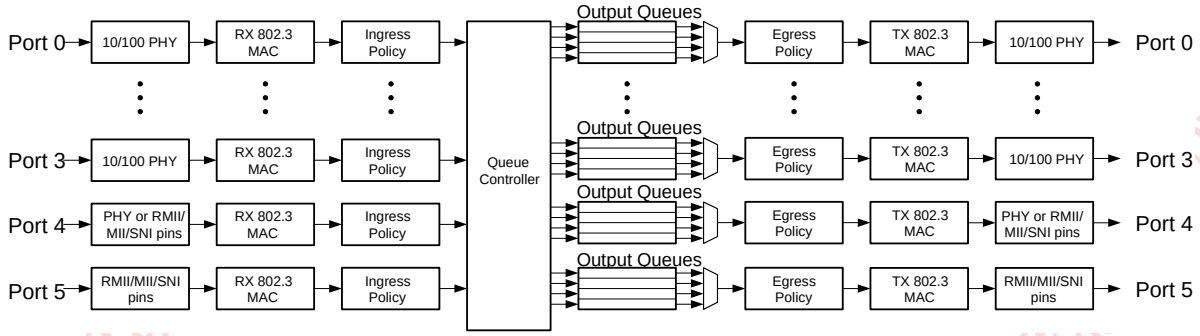

Figure 4: 88E6035 Switch Data Flow

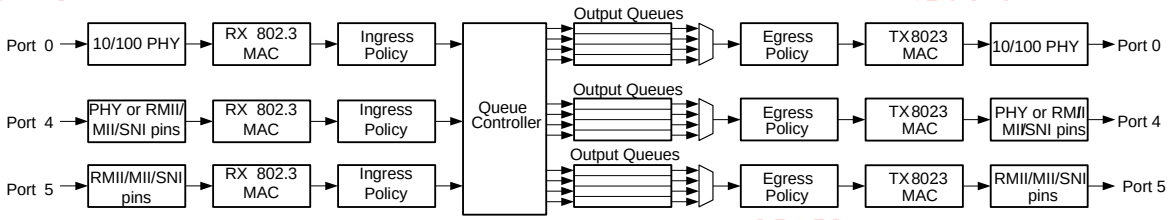

2.2 PHY or RMII/MII/SNI Pins

The PHY or physical layer interface is used to receive (and transmit) frames over long distances of Category 5 twisted pair cable or fiber-optic cables (only Port 0 and 1 support a fiber option). The 88E6065/88E6035 device contains five PHYs connected to Ports 0 to 4. The PHY block is covered in detail in Section 3. "Physical Interface (PHY) Functional Description" on page 114.

The sixth port, Port 5, and optionally the fifth port, Port 4, on the 88E6065/88E6035 device does not contain a PHY. Instead short distance industry standard digital interfaces are supported and generically called the port's MII Interface. Many interface modes and timings are supported so that a large number of external device types can be used.

The two MII interfaces on the 88E6065/88E6035 device support four major modes of operation and each MII interface can be configured independently.

2.2.1 MII PHY Mode

The MII PHY Mode, sometimes called 'Reverse MII', configures the desired MAC inside the 88E6065/88E6035 to act as a PHY so that it can be directly connected to an external MAC (like one inside a Router CPU). In this mode, the 88E6065/88E6035 device drives the interface clocks (Px_INCLK and Px_OUTCLK) so the desired frequency needs to be selected. Both full and half-duplex modes are supported and need to be selected as well to match the mode of the link partner's MAC. The MII PHY interface is compliant to the IEEE 802.3u clause 22.

Figure 5: MII PHY Interface Pins for the 88E6065/88E6035 Device

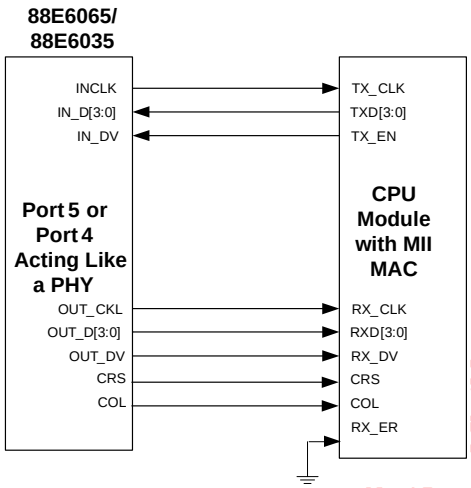

2.2.2 MII MAC Mode

The MII MAC Mode, sometimes called 'Forward MII', configures the desired MAC inside the 88E6065/88E6035 device to act as a MAC so it can be directly connected to an external PHY. In this mode, the 88E6065/88E6035 device receives the interface clocks (Px_INCLK and Px_OUTCLK) and will work at any frequency from DC to 25 MHz or 50 MHz (on Port 5 only) and 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 interface is compliant with IEEE 802.3 clause 22.

Figure 6: MII MAC Interface pins

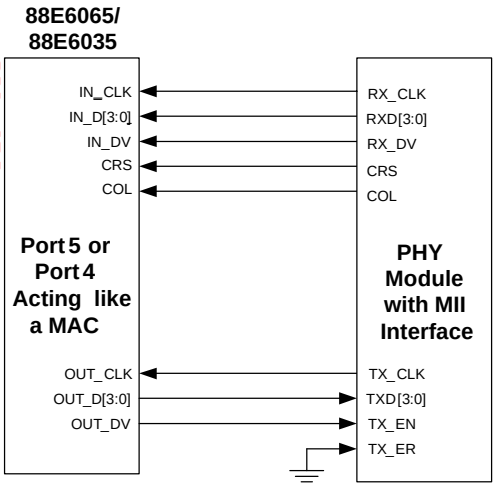

2.2.3 SNI PHY Mode

The SNI PHY Mode (7-Wire interface) configures the selected MAC inside the 88E6065/88E6035 device to act as a 10 Mbps PHY, enabling it to be directly connected to an external 10 Mbps only MAC (such as one inside a Router CPU). In this mode, only one data bit is used on each of input and output (Px_IND0 and Px_OUTD0). The interface clocks, Px_INCLK and Px_OUTCLK, are driven by the 88E6065/88E6035 device. Since SNI was never standardized, the 88E6065/88E6035 device supports various SNI modes. The active edge of the clock (either rising or falling) and the active level on the collision signal (either active high or low) can be selected.

Figure 7: SNI PHY Interface Pins

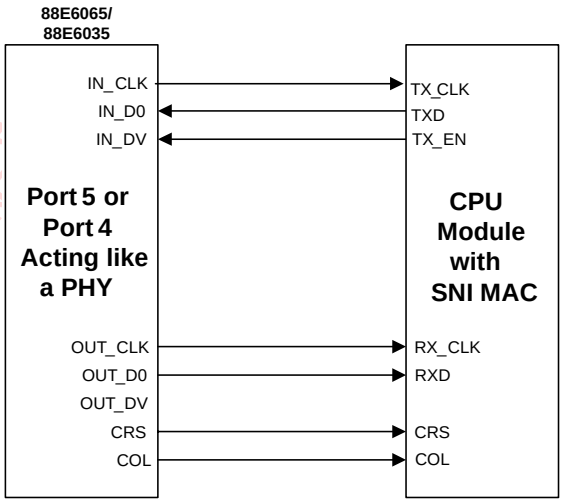

2.2.4 RMII PHY Mode

RMII PHY Mode (Reduced MII) configures the selected MAC inside the 88E6065/88E6035 device to act as a 10 or 100 Mbps PHY with a Reduced Media Independent Interface (RMII) enabling it to be directly connected to an external 10 or 100 Mbps RMII MAC (for example, one inside an ASIC or FPGA). In this mode, only two data bits are used on each of input and output (Px_IND[1:0] and Px_OUTD[1:0]). The interface clock (Px_INCLK) is driven by the 88E6065/88E6035 device at the RMII constant frequency of 50 MHz. The speed of the interface, either 10 or 100 Mbps needs to be configured the same on both sides of the interface. The 88E6065/88E6035 device's side of the interface is configured the same on both sides. The 88E6065/88E6035 device's side of the interface is configured by modifying the port's Speed bit in the Port Status register (Switch Port register - Offset 0x00). The default Speed setting is 100 Mbps.

Figure 8: RMII PHY Interface Pins using INCLK

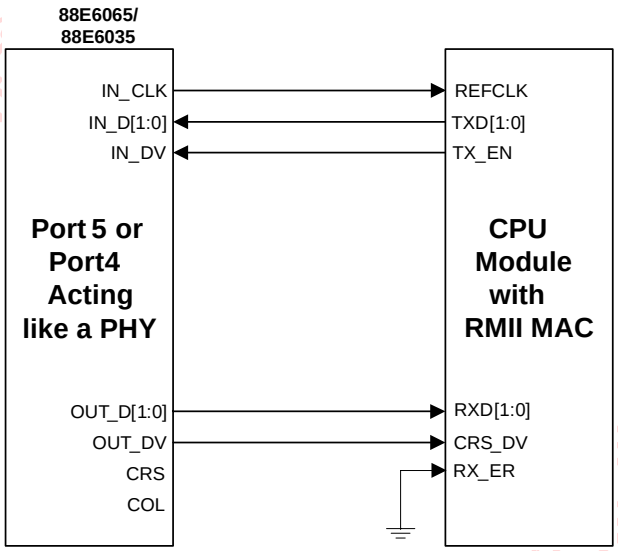

2.2.5 RMII/MII/SNI Configuration

The RMII/MII/SNI interface(s) in the 88E6065/88E6035 device are configured at the rising edges of RESETn. During reset the Px_OUTD[3:0]/Px_MODE[3:0] pins become tri-stated, and the values found on these pins becomes the interface's mode as defined in Table 16.

Table 16: RMII[1]/MII/SNI Configuration Options

Px_MODE[3:0] (at Reset)	PHY/ MAC Mode	Duplex	MII/SNI Mode	Description
0 0 0 0	PHY	Half-duplex	SNI 10 Mbps	Rising edge clock with collision active low
0 0 0 1	PHY	Half-duplex	SNI 10 Mbps	Rising edge clock with collision active high
0 0 1 0	PHY	Full-duplex	SNI 10 Mbps	Rising edge clock (collision is don't care)
0 0 1 1	MAC	Full-duplex	MII 200 Mbps	50 MHz MII Input clock mode
0 1 0 0	PHY	Half-duplex	SNI 10 Mbps	Falling edge clock with collision active low
0 1 0 1	PHY	Half-duplex	SNI 10 Mbps	Falling edge clock with collision active high
0 1 1 0	PHY	Full-duplex	SNI 10 Mbps	Falling edge clock (collision is don't care)
0 1 1 1	PHY	Full-duplex	MII 200 Mbps	50 MHz MII output clock mode
1 0 0 0	MAC	Half-duplex	MII 0 -100 Mbps	DC to 25 MHz MII input clock mode
1 0 0 1	PHY	Half-duplex	RMII 10/100 Mbps	50 MHz Reduced MII output clock mode
1 0 1 0	MAC	Full-duplex	MII 0 -100 Mbps	DC to 25 MHz MII input clock mode
1 0 1 1	PHY	Full-duplex	RMII 10/100 Mbps	50 MHz Reduced MII output clock mode
1 1 0 0	PHY	Half-duplex	MII 10 Mbps	2.5 MHz MII output clock mode
1 1 0 1	PHY	Half-duplex	MII 100 Mbps	25 MHz MII output clock mode
1 1 1 0	PHY	Full-duplex	MII 10 Mbps	2.5 MHz MII output clock mode
1 1 1 1	PHY	Full-duplex	MII 100 Mbps	25 MHz MII output clock mode

1. RMII's speed defaults to 100 Mbps. The speed can be overwritten to 10 Mbps by forcing the speed to 10 in the MAC Control Register -Offset 0x01.

2.2.6 Enabling the RMII/MII/SNI Interfaces

Port 5's MII is enabled when ENABLE_MII5 is tied high to VDDO. Port 4's MII is enabled when DISABLE_MII4 is tied to VSS.

2.2.7 Port Status Registers

Each switch port of the 88E6065/88E6035 device has a status register that reports information about that port's PHY or RMII/MII/SNI interface. See "Switch Registers" on page 146 for more information.

2.2.8 MII 200 Mbps Mode

Port 4 and Port 5 of the device's MII interfaces can run at a data rate of 200 Mbps full-duplex. Do not select this mode unless the MAC on the other end of the MII interface can also run at double speed rate. Both PHY (reverse MII) and MAC (forward MII) 200 Mbps modes are supported. When the 200 Mbps PHY mode is selected, the output MII clocks (OUTCLK and INCLK) run at a 50 MHz rate. There is no change in the format of the data. It just runs faster. When the 200 Mbps MAC mode is selected, the input MII clocks (OUTCLK and INCLK) must be 50 MHz±50 ppm. Again the format of the data is not changed.

2.3 Media Access Controllers (MAC)

The 88E6065 device contains six independent Fast Ethernet MACs. These MACs perform all of the functions in the 802.3 protocol including frame formatting, frame stripping, CRC checking/generation, CSMA/CD enforcement, and collision handling. The 88E6035 device contains three independent Fast Ethernet MACs performing the same functions as described above. Each MAC supports 100 Mbps or 10 Mbps operation in either full or half-duplex modes.

The MAC receive block checks incoming packets and discards those with CRC errors, 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 1518 bytes for un-Tagged frames and more than 1522 bytes for IEEE 802.3ac Tagged[2] frames)[3]. 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)[4] 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 2.8 for more information). The next two to sixty-six bytes are examined and may be used for QoS (Quality of Service) and/or snooping decisions made by the switch (see section 2.6 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 of the frame. Actual transmission of the frame begins immediately after the SFD.

For half-duplex mode, the 88E6065/88E6035 device also monitors 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.

2.3.1 Backoff

In half-duplex mode, the 88E6065/88E6035 device's MAC implements 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 88E6065/88E6035 device resets the progressively longer backoff time circuit after 16 consecutive re-transmit trials (if enabled by setting a 0x0 to Global Control register - Offset 0x04 bit 13). Each MAC then restarts the backoff algorithm with the shortest random backoff time and continues to retry and re-transmit the frame. A packet that successively collides with a re-transmit signal is re-transmitted until transmission is successful. This algorithm prevents packet loss in highly-congested environments. The MACS in the switch can be configured (by setting a 0x1 to Global Control register - Offset 0x04 bit 13) to meet the IEEE 802.3 specification and discard a frame after 16 consecutive collisions instead of restarting the backoff algorithm. In full-duplex mode, the collision signal and the backoff algorithm are ignored.

1. If Provider/Customer Tagging is enabled the minimum frame size for Tagged frames is 68 bytes.
2. IEEE 802.3ac VLAN Tagged frames are four byte longer than normal IEEE 802.3 frames. The extra four bytes are used to support Tagging information for IEEE 802.1p and IEEE 802.1Q.
3. A maximum frame size of 2048 bytes is supported by setting the MaxFrameSize bit in the Global Control register (Offset 0x04).
4. The 1st six bytes of a frame are processed as the frame's DA unless the Marvell Header mode is enabled on the port (Switch Port Register - Offset 0x04). If the Marvell Header mode is enabled the 1st two bytes are processed as the Marvell Header and the next 6 bytes are processed as the frame's DA.

2.3.2 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. After the EEPROM data is loaded, by driving this pin to a low, enables flow control on all half-duplex ports. HD_FLOW_DIS pin status can be read using the Port Status Register, Offset 0x00 or decimal 0 bit 2, HDFlowDis.

The duplex of a port can be forced using MAC Control Register, Offset 0x01 or decimal 1, bit3 DpxValue and bit2 ForcedDpx.

Flow control can be enabled on any particular port by setting the port's Force Flow Control mode (FCValue and ForceFC in the MAC Control Register, Offset 0x01 or decimal 1). 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 carrier which prevents further incoming packets. Only the port(s) 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. The Port Status Register, Offset 0x00 or decimal 0 bit 4 i.e., FlowControl bit being a one indicates that if the port is in half-duplex mode, the port is or will be using the back pressure. The FlowControl bit is set if the MAC determines that no more frame data should be entering the port.

Note that the duplex status of a port can be derived using the parallel LED status pins as described in Section 3.9.1.

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

2.3.3 Full-duplex Flow Control

IEEE802.3 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 – to avoid dropping packets during congestion. It is enabled on all full-duplex ports if the EE_CLK/FD_FLOW_DIS pin was low at the rising edge of RESETn, if Auto-Negotiation is enabled on the port, and if the link partner 'advertises' that it supports PAUSE during Auto-Negotiation. The status of the pin FD_FLOW_DIS can be read using the Port Status Register, Offset 0x00 or decimal 0 bit 3 – FdFlowDis.

The full-duplex flow control is automatically enabled on a port if:

- The port's PHY is advertising it supports flow control
- 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 PHY's by the PPU, if enabled. 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 to advertise flow control. This can be done by disabling the PPU and writing to the 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.

The auto negotiated value of the flow control setting for a given port can be modified using the port's Force Flow Control mode (FCValue and ForceFC in the MAC Control Register, Offset 0x01 or decimal 1). Flow control can be enabled on any particular port by setting the port's ForceFlowControl bit (Switch Port register - Offset 0x04)[1]. If the full-duplex flow control mode is not set and there is no packet buffer space available, the incoming packet is discarded.

The duplex of a port can be forced using MAC Control Register, Offset 0x01 or decimal 1, bit3 DpxValue and bit2 ForcedDpx.

In full-duplex mode, devices MAC supports the standard flow control defined in IEEE 802.3x standard. It enables stopping and restoring the transmission from 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 17.

Table 17: 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 the free buffer space is almost empty the devices send out a PAUSE frame with the maximum pause time to stop the remote node(s) from sending more frames into the switch. Only the ports that are involved in the congestion are PAUSED. When the event that caused flow control to be asserted goes away, the devices send out a PAUSE frame with pause time equal to zero, indicating that the remote node may resume transmission.

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

The Source Address of received PAUSE frames is not learned[2] 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 (the default is all zeros). Global registers 1, 2 and 3 (Switch Global Register - Offset 0x01) can be set to the desired Source Address to use. 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 Bytes 0 & 1.

The MACs will discard all IEEE 802.3x PAUSE packets they receive. This is always the case, even if full-duplex flow control is disabled or even if the port is in half-duplex mode.

2.3.4 Forcing Flow Control in the MAC

Section 2.3.3 "Full-duplex Flow Control" described 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 4 in MII 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 bits in the port's MAC Control Register, Offset 0x01 or decimal 1, bits 7 and 6). 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[3].

1. The ForceFlowControl bit works best is the FD_FLOW_DIS pin is high (i.e., disabled) at RESETn.
2. See Automatic Address Learning in section 2.5.3.
3. 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.

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.

2.4 Statistics Counters

2.4.1 MAC Based Statistics Counters

Sometimes it is necessary to debug network occurrences. For example, a technician may want to view the network remotely to solve a customer problem or a software programmer may want to trace transmitted frames. In these situations, two basic types of data are important:

- Number of good frames entering and leaving each port of the switch
- Quality of network segment performance

The 88E6065/88E6035 Statistics Counters support basic debugging needs. Each port can capture six kinds of statistics:

- A count of the number of good frames received with the number of frames transmitted, or
- A count of the number of bad frames received with the number of collisions encountered

The first statistics answers the question: "Where did all the frames go?", while the second statistic answers the question "Does the network segment have any performance problems?".

Table 18: Statistics Counters

Group 0	**Description**
RxFrames	Total good frames received
TxFrames	Total number of frames transmitted
InDiscards	Total number if good, non-filtered frames discarded due to lack of buffer memory
Group 1	**Description**
RxErrors	Total number of bad frames received
TxCollisions	Total number of collisions encountered
InFiltered	Total number of good frames discarded due to filtering rules.

The counters can be cleared and their mode chosen by the CtrlMode bit in the Switch Global Control register (Offset 0x04).

2.4.2 Policy Based RMON/Statistics Counters

The Statistics Counter logic maintains a set of thirty, 32-bit counters.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, 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 19 and Table 20.

All CPU register interfaces are slow compared with the speed of Gigabit or even Fast Ethernet frames. For this reason the 88E6065/88E6035 device supports a snapshot function to capture instantly and hold static any port's set of 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 desired 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).

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

Table 19 describes the Ingress statistics counters for each port. Note: Group 3 counters can be configured to be Ingress only. Egress only or both.

Table 19: Ingress Statistics Counters

Name	Offset Address	Description
Group 1		**Rx Octet Counters**
InGoodOctets	0x00	The sum of lengths of all good Ethernet frames received, that is frames that are not bad frames. This counter includes octets from all Group 2 frames (below).
InBadOctets	0x02	The sum of lengths of all bad Ethernet frames received not counting Preamble. This counter includes octets from all Group 3 frames (below) except for fragments due to collisions unless the collision occurred after the Start of Frame delimiter (SFD).
Group 2		**Rx Good Frame Counters**
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 as they are considered MAC Control frames nor does it include Broadcast frames counted in InBroadcasts.
InPause	0x16	The number of good Flow Control frames received.

Table 19: Ingress Statistics Counters (Continued)

Name	Offset Address	Description
Group 3		**Rx Bad Frame Counters**
InUndersize	0x18	Total frames received with a length of less than 64 octets but with a valid FCS. These frames are discarded and are not counted in Group2 above.
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.These frames are discarded and are not counted in Group2 above.
InJabber	0x1B	Total frames received with a length of more than MaxSize octets but with an invalid FCS.
InMACRcvErr	0x1C	Total frames received with an Rx_Error event seen by the receive side of the MAC.
InFCSErr	0x1D	Total frames received with a CRC error not counted in InFragments, InJabber.
Group 4		**Histogram Counters-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.
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 octets inclusive, including those with errors.

Table 20 describes the Egress counters for each port. Note: The Group 6 counters can be configures to be Ingress only, Egress only or both.

Table 20: Egress Statistics Counters

Name	Offset Address	Description
Group 5		**Tx Octet Counters**
OutOctets	0x0E	The sum of lengths of all Ethernet frames sent from this MAC. This includes Flow Control frames and partial frames transmitted due to collisions.
Group 6		**Tx Good Frame Counters**
OutUnicast	0x10	The number of frames sent that have a Unicast destination MAC address. This does not include frames dropped due to excessive collisions nor frames counted in OutFCSErr.
OutMulticasts	0x12	The number of good frames sent that had a Multicast destination MAC address. This does not include 802.3 Flow Control messages, nor frames counted in OutPause nor does it include Broadcast frames counted in OutBroadcasts.
OutBroadcasts	0x13	The number of good frames sent that have a Broadcast destination MAC address. This does not include 802.3 Flow Control frames, nor frames dropped due to excessive collisions nor frames counted in OutFCSErr.
OutPause	0x15	The number of Flow Control frames sent.
Group 7		**MiscellaneousTx Frame Counters**
OutFCSErr	0x03	The number of frames transmitted incorrectly due to an internal MAC transmit error, for example, bad CRC detected when the frame was read from memory.
Deferred	0x05	The total number of successfully transmitted frames that were delayed because the medium was busy during the first attempt and then the frame was transmitted without any collisions (this counter is valid in half-duplex mode only.
Collisions	0x1E	The number of collision events seen by the MAC not including those counted in Late or Excessive.
Late	0x1F	The number of late collision events seen by the MAC.
Excessive	0x11	The number frames dropped in the transmit MAC due to excessive collision condition. 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.
Group 8		**Histogram Counters - 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.

Table 20: Egress Statistics Counters (Continued)

Name	Offset Address	Description
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 octets inclusive, including those with errors.

2.5 Basic Switch Operation

The switch portions 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 3. These blocks modify the normal or default packet flow through the switch and are discussed in section 2.6 and section 2.8. The normal packet flow and processing is discussed first.

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 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 it came in on[3]. This ensures the packet will be received by the correct end station (if it exists) and when the end station responds back its address will be 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 a new MAC address/port number association will need to 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.

Ingress and egress policy enforcement is another basic switch operation that is performed on all incoming frames into the switch. These policies could involve VLAN based port membership, ingress rate limiting checks, egress rate limiting checks, port based security authentication etc.

Once the frames are accepted to be queued into the switch, the switch maintains quality of service differentiation of various frames by mapping them into their respective priority class queues. The switch uses various Layer1 to Layer4 header fields to determine the priority class (or it could use the port based default priorities). The act of enqueuing the frames into their respective priority queues is performed by the queue controller. The enqueued packets are read out of the packet memory using various scheduling algorithms (for e.g. Weighted Fair Queuing) described in the following sections.

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

1. The SA on switch management and Pause frames are not learned.
2. VLANs modify this operation – see section 2.6.8.
3. The devices can be configured to transmit frames out the port they came in on – Switch Port Registers - Offset 0x06.

Figure 9 shows each of the switch functions.

Figure 9: Basic MAC Address Lookup Overview

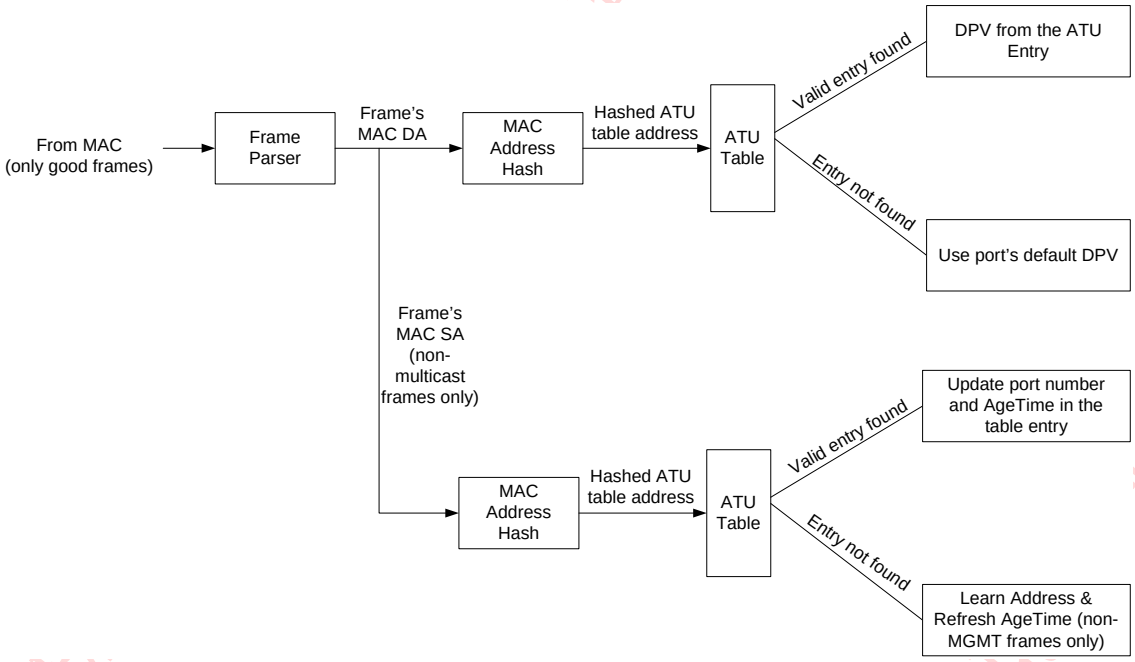

2.5.1 Lookup Engine

The devices Lookup Engine or Address Translation Unit (ATU) use the DA and SA addresses 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 will result in some MAC addresses having the same hash address.This is called a hash collision and is solved in the devices by using a four entry bin 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 the embedded SRAM and has a size of 4096 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 (or almost 64 minutes). These options are set in the ATU Control Register, Offset 0x0A or decimal 10.

2.5.2 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 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

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 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, which typically floods the frame. If the destination address in the frame is a multicast address or broadcast address, the address is searched 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 register, Offset 0x04) or the one assigned to the frame by the VTU (section 2.6.11). MAC addresses that are not members of the port's or frame's DBNum cannot be found.

2.5.3 Automatic Address Learning

The address learning engine is used to learn the source address of ingressing frames. Up to 1K MAC address/port number mappings can be stored in the address database (See section 2.5.1 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. 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). If all four address bins have the same Age time, then the first 'non-static' bin is used (see section 2.5.10.1 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 - Offset 0x00).

Multiple separate address databases are supported in the devices (see section 2.5.7). The port's database number (DBNum in the Port Based VLAN Map register, Offset 0x06, and the Port Control register, Offset 0x04) 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 2.6.15.10) 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 performing one of the following:

- Clearing the port's PAV to all zeros (see Port Association Vector, Offset 0x0B) or
- By setting the LockedPort bit in Port Association Vector, Offset 0x0B register (this enables CPU directed learning), or
- By setting the port's learn disable bit (in Global Control 1 register – ATU Control, Offset 0x0A).
- 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 2.3.3), BPDU, LAC and other management frames as long as they are determined to be MGMT frames (Section 2.5.8).

2.5.4 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 (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 (as long as the entry is not locked – see section 2.5.10.1). 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 2.5.3). 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 (See section 2.5.3) 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.

2.5.5 CPU Directed Address Learning

In some network environments, 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 2.5.10.6). 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 2.5.10.3). 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.

2.5.6 Source MAC Address Checking

The devices support the following SA filtering modes:

- 802.1X MAC authentication: This is used for MAC based port authentication as specified in IEEE 802.1X standard.

  The incoming frames would get discarded if their SA field is not in the ATU's address database or if this port's bit is not set in the PortVec bits indicating that this port is not expected to receive packets from the frame's MAC SA. CPU Directed Address Learning is required for 802.1X 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.

  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 as 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), which will drop the frame that caused the interrupt.). This configures the ingress port in a form of a Secure Port (Section 2.6.5) and will work as long as no new SA needs to be authorized on this port.

- Reverse 802.1X MAC authentication: This feature is used for blocking of known but distrusted MAC addresses. In this mode, ingressing frames would get discarded if their SA field is in the ATU's address database as a static entry with a PortVec of all zeros. This mode is like a reverse 802.1X authentication and is an enhancement to the IEEE 802.1X standard.

  Incoming frames whose MAC SA's are not in ATU's address database are learned, if the learning is not disabled as specified in Section 2.5.3. In environments where a set of MAC addresses are not trusted, those MAC addresses would get entered into the ATU database with their corresponding PortVec set to all zeros.

- Trapped 802.1X authentication: This mode is used for CPU to inspect packets from a given MAC address. In this mode, ingressing frames would get mapped to the CPUPort (Global 1, Offset 0x09) if their SA is in ATU's address database as a static entry with a PortVec of all zeros. If the SA is not in the ATU's address database then the frame would get discarded.

2.5.7 Multiple Address Database Support (DBNum)

The devices support up to 64 separate and independent address databases in the Address Translation Unit (ATU). Multiple address databases 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 1024 MAC addresses or any number in between. Any number of DBNum values can be used (from 1 to 64). 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 in the VLAN Translation Unit (VTU, Section 2.6.11). 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 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).

A frame's DBNum is generally determined by the frames VID (via the VTU, Section 2.6.11). Multiple VID's can be mapped to the same DBNum allowing for shared VLAN address databases.

2.5.8 Management/802.1D BPDU Frame Detection

The devices support two methods of management/802.1D BPDU frame detection and mapping (2.5.8.1 and 2.5.8.2). These two mechanisms support IEEE industry standards like Spanning Tree Protocol (STP), Link Aggregation Control (LAC) and Operations, Administration and Maintenance (OAM), as well as any company proprietary protocol. Both of these mechanisms:

- Detect that a frame is special by examining the frame's DA[1].
- 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 2.6.21).
- Set the priority on these MGMT frames overriding all other QoS decisions on the frame (Section 2.6.5).
- Map these MGMT frames to a port where a management CPU is directly or indirectly connected.

For all static unicast frames, their queue priority will be forced to the one specified in that particular ATU entry, if the Entry_State[0] is set to a one. For all static multicast frames, their queue priority will be forced to the one specified in the ATU entry, if the Entry_State[3] is set to a one.

All incoming frames would be considered as MGMT frames if Port Based VLAN Map register, Offset 0x06 bit ForceMap is set to a one. This feature basically bypasses ATU based destination port determination and uses the port or ports defined in the VLAN table. This forcing function is needed to get BPDU frames to egress specific ports by the CPU for the STP.

1. Frames with a DA of 01:80:C2:00:00:01 are always treated as Pause frames and are discarded and never mapped.

2.5.8.1 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 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 set to a one.

Any frame, regardless of its VLAN identifier (VID) or DBNum 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).

2.5.8.2 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[1]. 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 2.5.10).

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 2.5.7) 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.

2.5.9 Mux'ing or Ignoring 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 lookup always take place, regardless of the setting of port's MapDA bit. This is done so that DA lookups that are found to be MGMT entries in the ATU[2] are always mapped (even if the MapDA bit is configured to ignore the results). The application of SA Learning (Section 2.5.3) Default-Forward (Section 2.6.4) and ForwardUnknown (Section 2.6.5) 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 2.6.6) along with the EgressFloods (Port Control register, Offset 0x04) 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 2.6.12). This can be used for the following applications.

2.5.9.1 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 2.5.9). MGMT frames will be mapped as they normally would (Section 2.5.8).

1. 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.
2. See section 2.5.8.2 and section 2.5.10.

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.

2.5.9.2 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 2.5.9). MGMT frames, including OAM frames, will be mapped as they normally would (section 2.5.8).

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 2.6.14).

2.5.10 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 21. 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, Offset 0x04, and Global Status, global Offset 0x00).
- Read the results if appropriate.

Table 21: ATU Operation Registers

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define ATU entry configuration information like AgeTime etc.	Used to read back the ATU configuration information.
ATU Operation	0x0B (decimal 11)	Used to define the MAC's database number, the required operation and violation control bits.	Used to indicate the ATU's Busy status and its DBnum[5:0]
ATU Data	0x0C (decimal 12)	Used to further define the frame and queue priority for this MAC address.	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 desired MAC address to operate on.	Returns the resulting MAC address from the desired operation.

2.5.10.1 Format of the ATU Database

Each MAC address entry in the ATU database is 64-bits in size. The lower 48 bits contains the 48-bit MAC address and the upper 16 bits contains information about the entry as shown in Figure 10. The database is accessed 16-bits at a time via the Switch Global registers shown in Figure 10.

Figure 10: Format of an ATU Entry

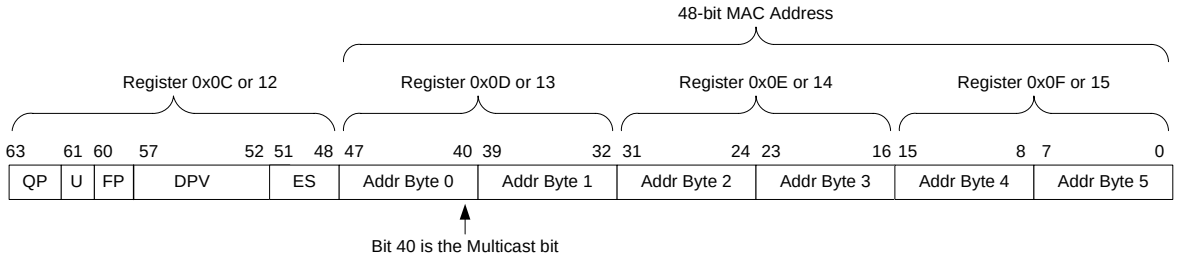

The right 48-bits in Figure 10 is the 48-bit MAC address associated with this ATU entry. The upper 16 bits is the data that is associated with the entry's MAC address. The upper 16 bits are defined as follows:

Table 22: ATU Data bits

Field	Bits	Description
QP	63:62	The MAC's Priority override value when enabled by the EntryState bits below. Used for queue priority override associated with this MAC address (if the EntryState indicates Queue priority Override and if the port's SA and/or DA QPriOverrides in Port Control2 register, Offset 0x08 are enabled.
U	61	Use MAC Frame priority override bits specified below. This bit validates the FPri bits below implying that the frame priority bits given below can override all other methods of determining a frame priority for this MAC address, provided if the port's SA and/or DA FPri overrides inPort Control2 register, Offset 0x08 are enabled.
FP	60:58	Frame priority override bits. These bits are valid only if the Frame priority override bit above is enabled. These bits are used as the frame priority (L2 only) for the frame egressing the switch by remapping the egressing frames priority bits to these bits.
DPV or Trunk ID	57:52	The Destination Port Vector or Trunk ID. These bits indicate which port or ports are associated with this MAC address (i.e., where frames should be switched). 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.

Table 22: 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. These entries would never have their frame priority or queue priority forced to some other value based on the FP and QP bits defined above. • 0x8 to 0x9 = static MGMT entry with Non rate limit Enable. Frames with this MAC address are considered static management addresses which do not age and the ingress rate limit logic would not rate limit these frames. If Entry state bit zero is a 1 i.e., for Entry state=0x9, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0xA to 0xB = Static unicast with Non rate limit Enable. Frames with this MAC address are considered as static unicast addresses which do not age and the ingress rate limit logic would not rate limit these frames. If Entry state bit zero is a 1 i.e., for Entry state=0xB, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0xC to 0xD = Static MGMT unicast. Frames with this MAC address are considered as unicast MGMT frames which do not age. If Entry state bit zero is a 1 i.e., for Entry state=0xD, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0xE to 0x0F = Static unicast. Frames with this MAC address are considered as static unicast frames which do not age. If Entry state bit zero is a 1 i.e., for Entry state=0xF, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame.

Table 22: ATU Data bits (Continued)

Field	Bits	Description
Entry State (continued.)		For multicast MAC addresses (bit 40 = 1): • 0x0 = Invalid, empty or purged entry. • 0x7 & 0xF = Static Multicast. Frames with this MAC address are considered as static multicast frames which do not age. If Entry state bit three is a 1 i.e., for Entry state=0xF, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0x6 & 0xE = Static MGMT. Frames with this MAC address are considered as static MGMT frames which do not age. If Entry state bit three is a 1 i.e., for Entry state=0xE, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0x5 & 0xD = Static Multicast with Non rate limit enable. Frames with this MAC address are considered as static multicast frames which do not age and the ingress rate limit logic would not rate limit these frame. If Entry state bit three is a 1 i.e., for Entry state=0xD, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame. • 0x4 & 0xC = Static MGMT with Non rate limit enable. Frames with this MAC address are considered as static MGMT frames which do not age and the ingress rate limit logic would not rate limit these frames. If Entry state bit three is a 1 i.e., for Entry state=0xD, the queue priority gets overridden by the value QP (bits 63:62) specified above. The frame priority for these frames can be overridden if the F bit (bit 61) is set and the FP bits (bits 60:58) become the new frame priority bits for the frame All other values are reserved for future use.

2.5.10.2 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 2.5.7), 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 23.

Table 23: ATU Get Next Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define ATU entry configuration information like AgeTime etc.	No change.
ATU Operation	0x0B (decimal 11)	Used to define the MAC's database number. the required operation and violation control bits.	Used to indicate the ATU's Busy status.
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 searched 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 searched address.

2.5.10.3 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 (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 2.5.3). The 16-bits of the ATU Data register are written into bits 63:48 of the ATU entry (see section 2.5.10.1).

A summary of how the Load operation uses the ATU's registers is shown in Table 24.

Table 24: ATU Load/Purge Operation Register Usage

Register	Offset	Before the Operation Starts	After the Operation Completes
ATU Control	0x0A (decimal 10)	Used to define the ATU entry configuration information like AgeTime etc.	No change.
ATU Operation	0x0B (decimal 11)	Used to define the operation (including the database to load or purge) 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.

2.5.10.4 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.

2.5.10.5 Moving or Removing Single Port Mappings

All MAC address port mappings associated with a specific port can be moved to another port or removed from the address database. This operation can be done for the entire set of address databases or for just a particular address database using single ATU operations. These ATU operations are:

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

The Move requires the EntryState bits in the ATU Data register (Offset 0x0B) to be 0xF. The PortVec bits in the ATU Data register are used to define the FromPort (bits [3:0] of PortVec) and the ToPort (bits [7:4] of PortVec). The ATU MAC Address registers (Offset 0x01, 0x02, and 0x03) are not used for these operations and they are left unmodified. The DBNum field of the ATU Operation register (Offset 0x0B) is used for the Flush operations that require a database number to be defined.

An ATU Move operation examines all the entries in the address database. Any valid entry that meets the requirements is processed. The requirements are:

- The address is valid (its EntryState is non-zero)
- The address is contained in the selected database (either all or the one DBNum selected)
- The address has the selected state (either all or Non-Static)
- The address has the FromPort's bit set to a one in its DPV field

Processed entries have their ATU Data register FromPort bit cleared to a zero and have their ToPort bits set to a one. If the ToPort's value is 0xF, the FromPort bits are cleared but the ToPort's bit is not set.

2.5.10.6 Servicing ATU Violations

The ATU captures ATU Full, Age Violation, SA Member Violation and SA Miss Violation data. An ATU Full violation occurs if an Automatic Address Learn (section 2.5.3) or an ATU load operation (section 2.5.10.3) 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 ATU Age violation occurs when the frame's SA is found in the address database and the entry is non-static and the age on the entry is less than 4. This interrupt can be masked on a per-port basis by setting the port's AgeIntEn bit (see the Port Association Vector register - Offset 0x0B) to a one.

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 25.

Table 25: 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 of DBNum[3:0].
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.
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

2.6 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 including reverse 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 Marvell Header mode supports faster CPU routing. Ingress storm limiting is supported for enforcing service level policies, along with a Double Tagging option (Q-in-Q) and IGMP/MLD snooping.

Figure 11: Ingress Policy Overview

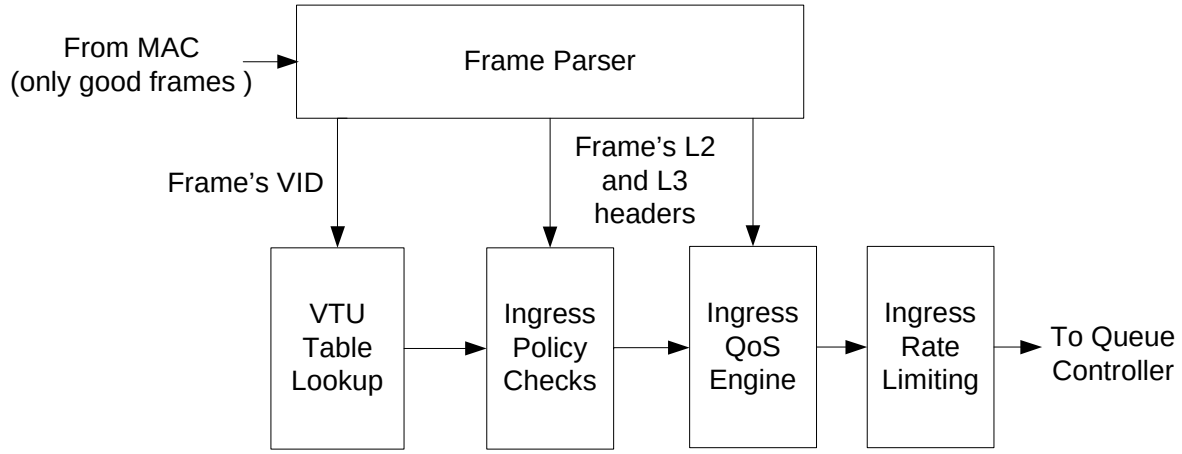

2.6.1 Port States for 802.1D Spanning Tree

The devices support four 802.1D Port States per port (802.1s per VLAN Port States are also supported - see Section 2.6.12). The 802.1D Port States are used by the Queue Controller (see Section 2.7) 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 36 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 2.5.8. Their primary purpose is to support the Spanning Tree Protocol (see Section 2.10) 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 2.5.8). 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.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 13). 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.

2.6.2 802.1X MAC Address Authentication

All non-MGMT[1] (non-management) frames received on a port with an unauthorized Source MAC Address are discarded All non-MGMT (non-management) frames received on a port with an unauthorized Source MAC Address are discarded if the port's SA_Filtering bits are set to a 0x1 (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 2.5.10). The policy of how these addresses are entered into that address database is controlled by the ATU (section 2.5.6). It is recommended that CPU Directed Address Learning (section 2.5.5) be used on ports supporting MAC based 802.1X authentication (i.e., ports with their DropOnLock bit being set to a one).

2.6.3 Reverse 802.1x MAC address Authentication

All frames that result in a MAC address hit but are known to be from malicious users are discarded if the port's SA_Filtering bits are set to 0x2 (Port Control Register, Offset 0x04). The ATU Entry needs to be a static entry for this function as higher layer management software is expected to identify these untrustworthy sources through mechanisms which are outside the scope of this discussion and program those entries into the ATU table. The policy of how these addresses are entered into that address database is controlled by the ATU. It is recommended that CPU Directed Address Learning (Section 2.5.5) be used on ports supporting MAC based 802.1X authentication (i.e., ports with their DropOnLock bit being set to a one).

2.6.4 Trapped 802.1x MAC address authentication

This mode is used if CPU wants to inspect packets coming in from certain MAC addresses. These entries are loaded into the ATU table as static and using the procedure described in the ATU operations section above and the DPV within this entry would be set to all zeros. If the frame's SA is either not found in the ATU table or if this Port's bit is not set in the PortVec bits for the frame's SA. It is recom¬mended that CPU Directed Address Learning (Section 2.5.5) be used on ports supporting MAC based 802.1X authentication.

2.6.5 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 2.5.1). 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 2.5.3), 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 2.5.10). 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.

1. MGMT frames and their Ingress Rule changes are covered in section 2.5.8.

2.6.6 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 2.5.1). 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.

2.6.7 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 2.6.21). Instead, it has the job of determining the priority of each frame for the Queue Controller. Every frame once it enters the switch gets assigned a InitialPri (both Frame priority and Queue Priority) as specified in Port Control Register (Offset 0x04) InitialPri bit setting (from either tag priority, IP priority or port's default). In the ingress policy block, this InitialPriority can get overridden by any of the items given below (see priority override list below). The ingress policy block considers the separate frame and queue priority overrides in determining the final queue priority the queue controller enqueues the packet into and also the frame priority of the outgoing frame.

Figure 12: Switch QoS Selection Overview

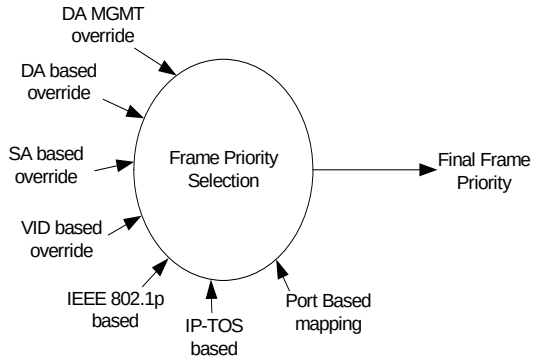

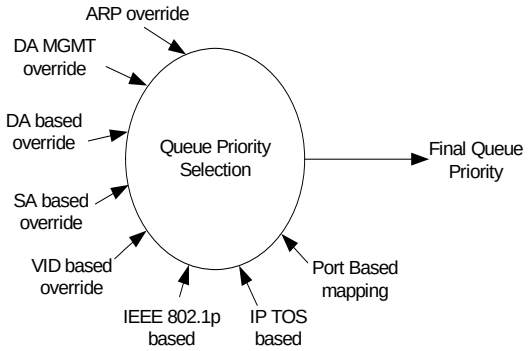

2.6.7.1 Frame priority determination

The frame priority of a frame is determined by (in priority order):

1. 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 2.5.8).
2. The frame's DA is in the address database with a priority defined and DA priority override is enabled on the port. See section 2.5.9 and DAFPriOverride in Port Control 2, Offset 0x08.
3. The frame's SA is in the address database, loaded as static, with a priority defined and with SA frame priority override enabled on the port. See section 2.5.9 and SAFPriOverride in Port Control 2, Offset 0x08.
4. The frame's VID is in the VLAN database with a priority defined and VTU frame priority override is enabled on the port. See section 2.6.13 and VTUFPriOverride in Port Control 2, Offset 0x08.
5. 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 2.6.11 and UseTag in Port Control, Offset 0x04.
6. The IPv4 Type of Service (TOS)/DiffServ field or IPv6 Traffic Class field if they are enabled on the port[0][1]. The default case is enabled on all of the ports. See section 2.6.12 and UseIP in Port Control, Offset 0x04.
7. The Port's default priority defined in DefPri (see the Default Port VLAN ID and Priority register, Offset 0x07).
8. There are enables on each of the priority classification rules giving the network administrator any combination that is needed. For instance, if a, port based higher priority port is required for a switch design, priority classifications 1 to 6, 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. Note that Port Control Register (Offset 0x04) bit TagIfBoth determines the priority, if the frame is tagged and is also a IPv4 or IPv6 and if the InitialPri bits are set to 0x3.

2.6.7.2 Queue Priority Determination

The queue priority of a frame is determined by (in priority order):

1. If the incoming frame is an ARP frame, the per-port register bit setting of ArpQPriOverride overrides the priority of these frames to the one specified in ArpQPri bits.
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 2.5.8).
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 2.5.10 and DAQPriOverride 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 frame priority over-ride enabled on the port. See section 2.5.10 and SAQPriOverride in Port Control 2, Offset 0x08.
5. The frame's VID is in the VLAN database with a priority defined and VTU frame priority override is enabled on the port. See section 2.6.11 and VTUQPriOverride 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 2.6.12 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[2]. the default case is enabled on all of the ports. See section 2.6.13 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).
9. There are enables on each of the priority classification rules giving the designer network administrator any combination that is needed. For instance, if a, port based higher priority port is required for a switch design, priority classifications 1 to 6 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.

1. IPv4/IPv6's priority classification can be configured on a per port basis to have a higher priority than the IEEE tag by the port's TagIfBoth bit in the Port Control register - Offset 0x04.

The DA priority is used to set the highest priority on BPDU's or other management frames (called MGMT[1] 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 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[2] during Ingress, will egress with the overridden priority in the frame's IEEE Tag, if the frame egresses the switch tagged.

2.6.7.3 IEEE Tagged Frames

Figure 13: IEEE Tagged Frame Format

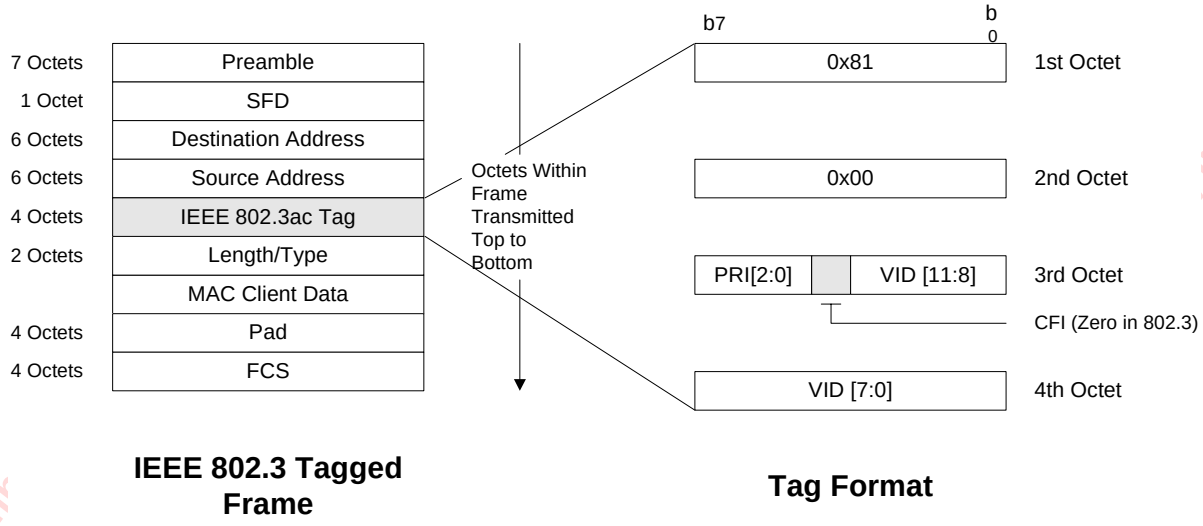

The device captures the frame's PRI[2:0] bits, remaps them, and uses them as the frame's priority. The re-mapping table is independent per port (Offset 0x18). IEEE Tagged frames that get their priority re-mapped during Ingress, will Egress with the re-mapped priority in the frame's IEEE Tag, if the frame egresses the switch tagged.

The IEEE Tag supports 8 priorities while the 88E6065/88E6035 supports 4 priorities. Since only 4 priorities are supported, 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 (Offset 0x81). The CFI bit of the IEEE Tag is ignored by the 88E6065/88E6035. The VID bits are ignored if 802.1Q is disabled on this port (see Section 2.6.9). The device's default is to capture and use IEEE Tagged frame priority data over IP priority data on all ports.

1. MGMT frames are determined by the frame's DA - section 2.5.8.
2. Each port supports an 8x3 priority re-mapping table (Offset 0x18 and 0x19) that can be used to scale a port's priorities up or down separately from the other ports.

2.6.7.4 IPv4 & IPv6 Frames

The format of an IPv4 TOS/DiffServ frame is shown in Figure 14 and an IPv6 Traffic Class frame is shown in Figure 15.

Figure 14: IPv4 Frame Format

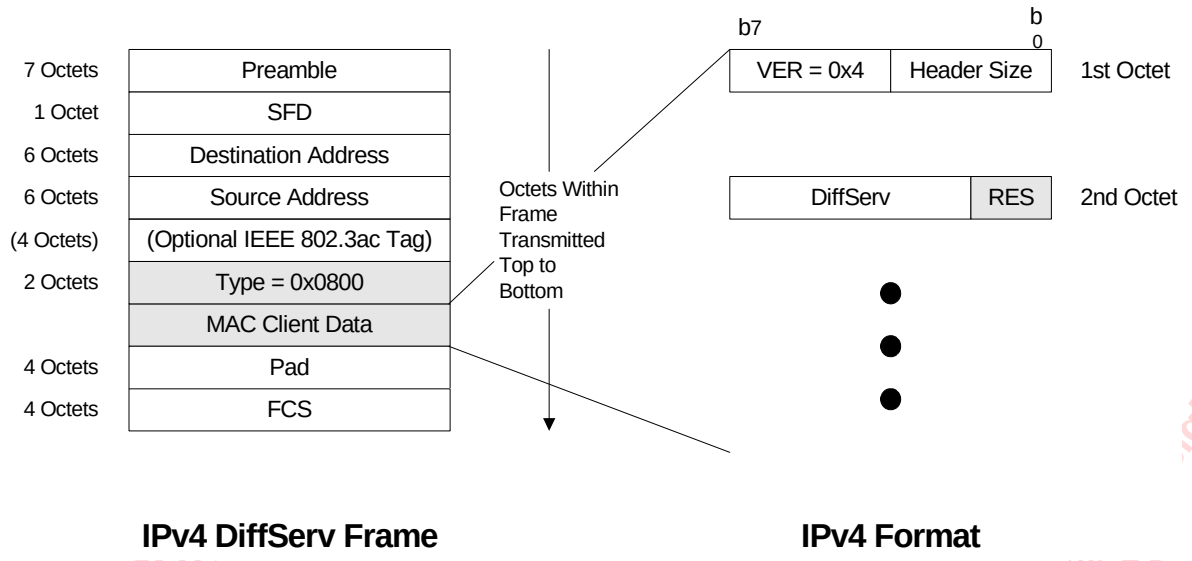

Figure 15: IPv6 Frame Format

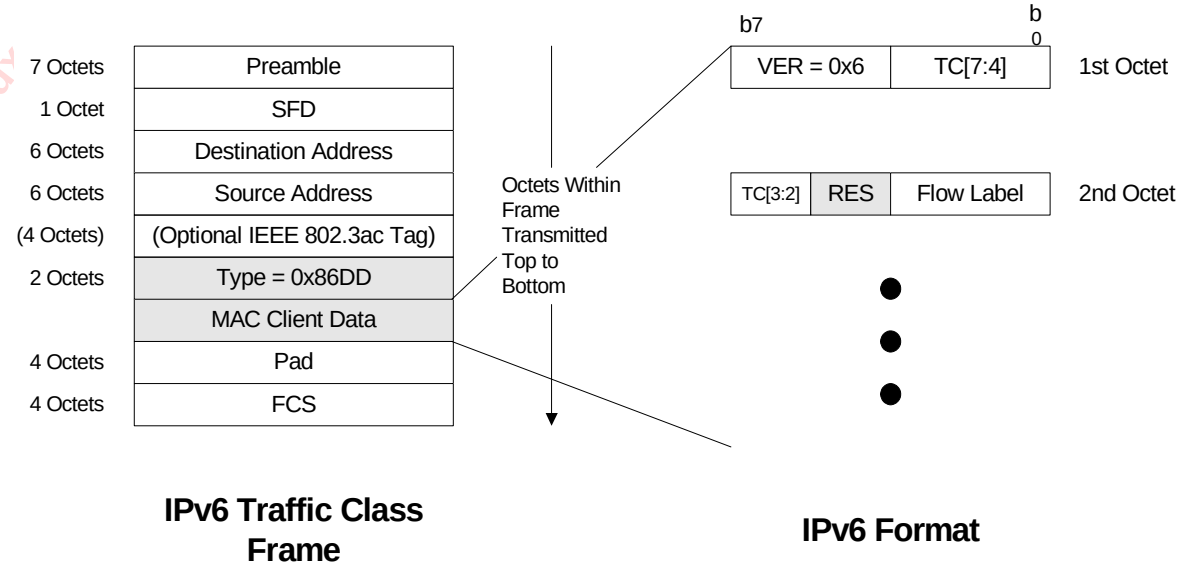

The device captures the frame's DiffServ bits (if its an IPv4 frame) or the frame's TC[7:2] bits (if its an IPv6 frame) and uses them as the frame's priority. The DiffServ/TC bits supports 64 priorities while the 88E6065/88E6035 supports 4 priorities. Since only 4 priorities are supported, the Ingress 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 (Offset 0x10). The rest of the frame's IP bits are ignored by the 88E6065/88E6035 unless IGMP/MLD snooping is enabled in the port (see Section 2.6.18). The device's default is to capture and use IP frame priority data in the absence of IEEE Tag priority data on all ports.

2.6.8 VLANs

The devices support many VLAN options including support for 802.1Q and/or port based VLANs. MGMT (management) frames bypass all port based and 802.1Q VLAN tests (section 2.13).

2.6.8.1 Port Based VLANs

The devices support a very flexible port based VLAN system. It is enabled when:

- The port's 802.1QMode is disabled (in the Port Control 2 register – Offset 0x08), or
- When the port's 802.1QMode is Fallback and the frame's VID is not contained in the VTU (Section 2.6.9), or
- When any 802.1Q mode is enabled and ProtectedPort is enabled on the port (Offset 0x08)

Each Ingress port contains a register that restricts which output (or Egress) ports it is allowed to send frames to. 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 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 16.

One exception to the above VLANTable mapping bit setting is when ForceMap (Port Based VLAN map register, Offset 0x06) bit is enabled i.e., when ForceMap bit is set to a one, then all frames entering a port are considered as MGMT frames and they get mapped to ports defined in VLANTable bits overriding the ATU based destination port identification.

Figure 16: Switch Operation with VLANs Disabled

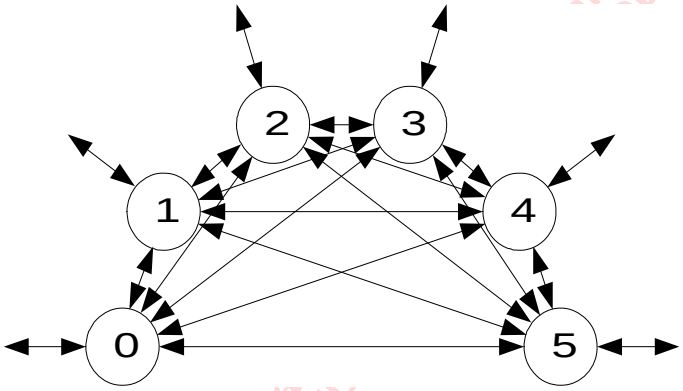

1. The 88E6065/88E6035 devices allow a port's own bit in its VLANTable to be set to a one- See section 2.6.10.

2.6.8.2 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 17 shows a typical VLAN configuration for a firewall router. Port 4 is used as the WAN port. 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 2.6.16) 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.

Figure 17: Switch Operation with a Typical Router VLAN Configuration

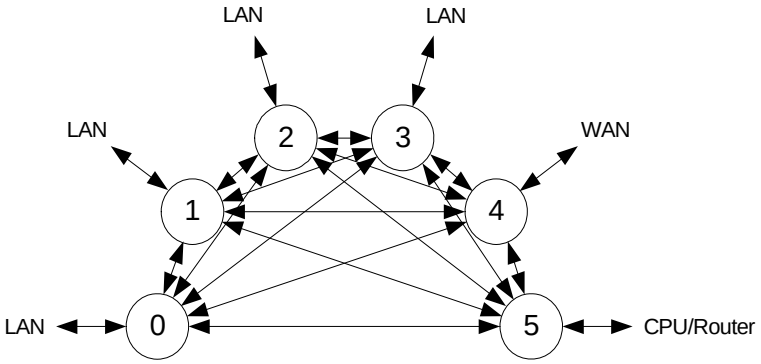

This specific VLAN configuration is accomplished by setting the port's VLANTable registers as follows:

Table 26: VLANTable Settings for Figure 17

Port #	Port Type	VLANTable Setting
0	LAN	0x3E
1	LAN	0x3D
2	LAN	0x3B
3	LAN	0x37
4	WAN	0x2F
5	CPU/Router	0x1F

To show the flexibility of the devices' VLAN Configuration options, Figure 18 shows another example. In this case, the switch is divided into two independent VLANs connected to a common router.

Figure 18: Switch Operation with example for VLAN Configuration

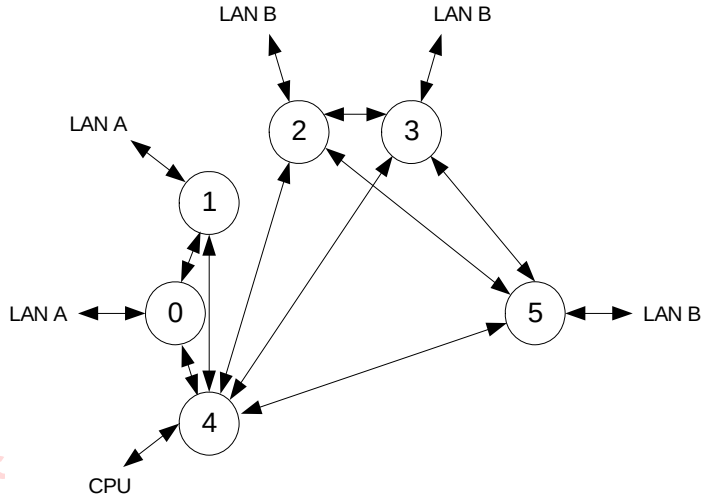

This specific VLAN configuration is accomplished by setting the port's VLANTable registers as follows:

Table 27: VLANTable Settings for Figure 18

Port #	Port Type	VLANTable Setting
0	LAN A	0x12
1	LAN A	0x11
2	LAN B	0x38
3	LAN B	0x34
4	CPU	0x2F
5	LAN B	0x1C

2.6.8.3 Tunneling Frames Through Port-Based VLANS

Normally frames cannot pass through the port based VLAN barriers. But on these devices some frames can. Before a frame can tunnel through a port based VLAN barrier its DA address must be loaded as static into the address database (section 2.5.10.1) and the VLANTunnel bit on the frame's Ingress port must be set to a one (Offset 0x04). When both of these events are true the frame is sent out 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 data is ignored in this case. This feature is enabled only on those ports that have their VLANTunnel bit set to a one and Port Based VLANs are being used on the frame because 802.1Q is disabled on the port or 802.1Q is in Fallback mode (section 2.6.9). The VLANTunnel feature is not supported in the ProtectedPort case (section 2.6.8.1).

2.6.9 802.1Q VLANs

The devices support 802.1Q with the full set of 64 different VIDs (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 2.6.8) is in effect for all 802.1Q modes described below.

2.6.9.1 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 2.6.8)
- 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:
  - ,embers of the frame's VLAN and
  - Included in the source port's VLANTable (see Section 2.6.8)
- 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 2.6.8)
- 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 2.6.8)

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).

2.6.9.2 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.

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).

2.6.9.3 Security Override of a Frame's VID

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

2.6.10 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 2.5.9.2), 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 regardless of the 802.1Q bit setting.

2.6.11 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 2.5.10.1) 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.

2.6.12 802.1s Per VLAN Spanning Tree

The devices support per VLAN Port States for 802.1s (per VLAN Spanning Tree). Each VID entry in the VTU (section 2.6.13) has two bits per port for 802.1s Port State information as shown in Table 28.

Table 28: 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 2.6.1.
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 2.6.7) or 802.1X authentication violations (section 2.6.2).
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 2.6.7) or 802.1X authentication violations (section 2.6.2).

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 2.6.1).

The 802.1s Port State options take precedence over the port's PortState bits settings (section 2.6.1), 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.

2.6.13 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 29. 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, Offset 0x04, and Global Status, Offset 0x00).
- Read the required results if appropriate.

Table 29: 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.

2.6.13.1 Format of the VTU Database

Each VID entry in the VTU database contains:

- A 1-bit valid indicator (Valid),
- A 6-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 19 and Table 30. The database is accessed 16-bits at a time via the Switch Global registers shown in Figure 19 (not all the register bits are shown). For more information about these register see the VTU Operation register (Offset 0x05) and VTU Data registers for all ports (Offset 0x06 to 0x09).

Figure 19: Format of a VTU Entry

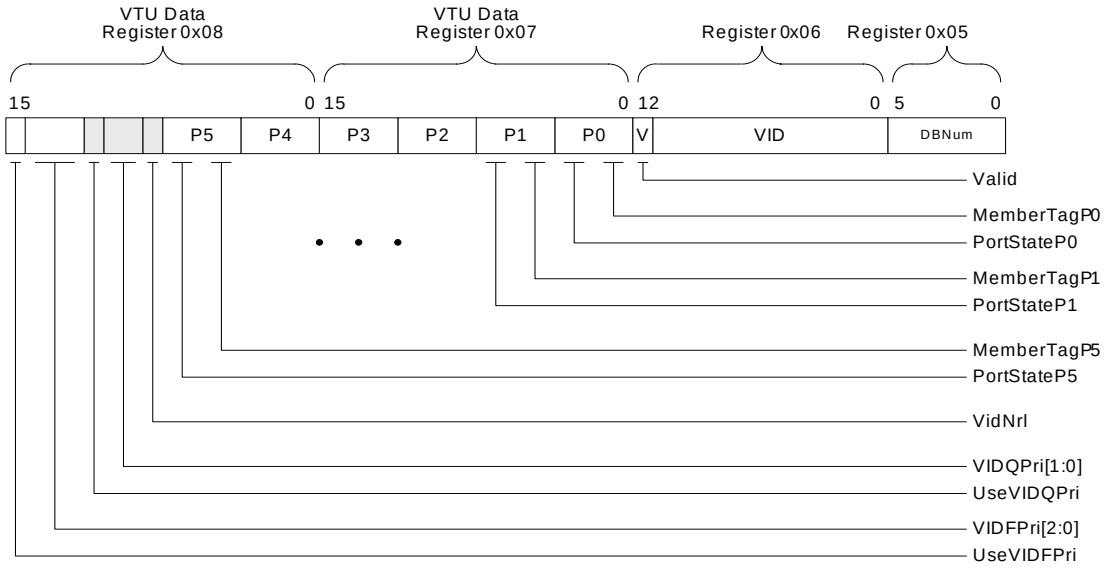

Table 30: VTU Entry Format

Field	Bits	Description
UseVIDFPRI	15 in global 1 Reg 0x08	Use VLAN identifier Frame Priority Override. This bit is used to indicate that frames assigned with this VID can have their frame priority overridden with the FPRI bits below if the port's VTUFPriOverride bit is set (see Port Control 2 register, Offset 0x08).
VIDFPRI	Bits 14:12 in global 1 Reg 0x08	VLAN identifier Frame Priority. These bits are used as a frame's priority for frames assigned with the entry's VID if the VIDFPRIOverride bit above is set to a one (and if it is enabled on the port -see Port Control 2 register, Offset 0x08).

Table 30: VTU Entry Format (Continued)

Field	Bits	Description
UseVIDQPRI	Bit 11 in global 1 Reg 0x08	Use VLAN identifier Queue Priority override. This bit is used to indicate that frames assigned with this VID can have their queue priority overridden with the QPRI bits below if the port's VTU-QPriOverride bit is set (see Port Control 2 register, Offset 0x08).
VIDQPri	Bits 10:9 in global 1 Reg 0x08	VLAN identifier Queue priority. These bits are used as a queue priority for frames assigned with the entry's VID if the VIDQPRIOverride bit above is set to a one (and if it is enabled on the port -see Port Control 2 register, Offset 0x08).
VidNrl	Bit 8 in global 1 Reg 0x08	VLAN identifier based Non-rate limit. This bit is used to indicate that frames assigned to this VID can not be ingress rate limited if the port's VidNrlEn bit is set (see Ingress Rate Control, Offset 0x09).
PortState	In global 1 Reg 0x07: Port 0 3:2 Port 1 7:6 Port 2 11:10 Port 3 15:14 In global 1 Reg 0x08: Port 4 3:2 Port 5 7:6	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 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).
MemberTag	In global 1 Reg 0x07: Port 0 1:0 Port 1 5:4 Port 2 9:8 Port 3 13:12 In global 1 Reg 0x08: Port 4 1:0 Port 5 5:4	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
Valid	Bit 12 in global 1 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 VID entries in the table are considered invalid (the Valid bit on each entry is cleared).
VID	Bits 11:0 in global 1 Reg 0x6	VLAN ID. These bits indicate the VID number that is associated with the MemberTag data, 802.1s Port State data, VTU frame/queue Priority and its override and the entry's address database number (DBNum).

Table 30: VTU Entry Format (Continued)

Field	Bits	Description
DBNum	Bits 9:8 and 3:0 in global 1 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 2.5.7). 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.

2.6.13.2 Reading the VLAN Database

The contents of the VLAN database can be dumped or searched. The dump operation is called Get Next (refer to Global1 VTU Operation register, Offset 0x05) 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 31.

Table 31: 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.

2.6.13.3 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 (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 32.

Table 32: 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.

2.6.13.4 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.

2.6.13.5 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 (Offset 0x05) and returns the VID that caused the violation in the VID field of the VTU VID register (Offset 0x06).

A summary of how the Get/Clear Violation Data operation uses the VTU's registers is shown in Table 33.

Table 33: 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

2.6.14 Switching Frames Back to its Source Port

The 88E6065/88E6035 supports the ability to send frames back out the port they came in on. While this is not a standard way to handle Ethernet frames, some applications 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 in its VLANTable register to a one. This function is valid if 802.1Q is enabled on the port or not and is valid for MGMT (management) frames (section 2.5.8).

2.6.15 IEEE Tagged Frame Handling

The format of an IEEE tagged frame is shown in Figure 20.

Figure 20: IEEE Tagged Frame Format

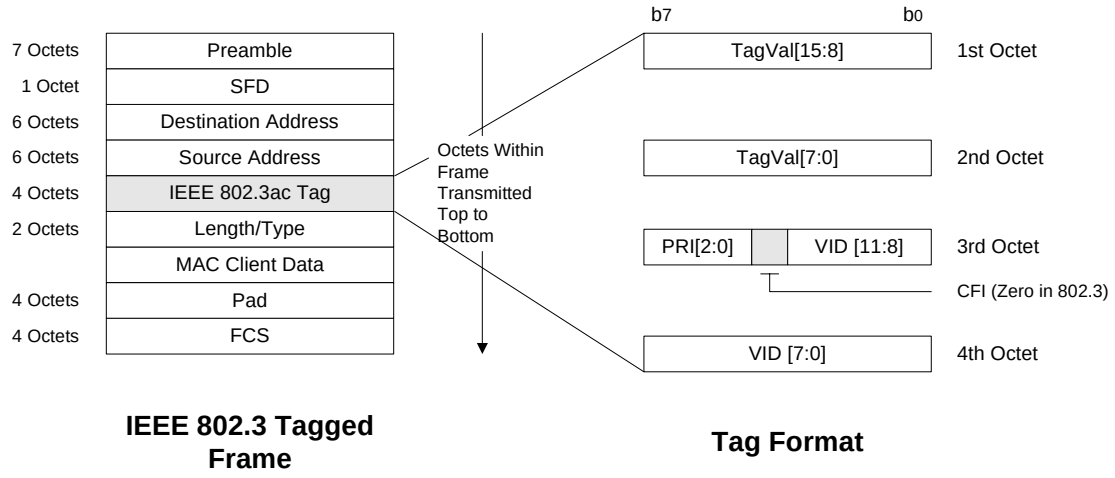

2.6.15.1 Determining if a Frame is Tagged

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, Offset 0x19).

Clearing UseCoreTag to zero is for normal IEEE 802.3ac tagged frame detection. 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 2.6.15.1) or at the edge (where the Provider Tag is removed - section 2.6.20).

2.6.15.2 Provider Tag handling

These devices support a per-port ProviderTag register (Offset 0x1A). This indicates the tag value that needs to be matched to in ingress to determine if the frame is Provider tagged or not. The per-port ProviderTag register programmability allows capabilities for these devices to be connected to multiple provider networks with different ProviderTag values.

The ingress pipe would compare the incoming Tag with the ProviderTag register value and also checks the EgressTagMode (Offset 0x04, bits[13:12] equals value 0x3) for this port.

These devices have the ability to strip one or more provider tags in ingress. The ingress pipe strips provider tags till the tag value matches 0x8100. The recursive tag stripping feature is useful in applications where the frame would have traversed through several provider networks and each provider network could potentially have added a tag on top of other tags. Thus the frame at the provider network edge would have so many tags and they need to be stripped to get to the customer tag. The devices give an option to use the outer tag information for switching and/or priority determination. The stripped tag in ingress can be used as a tag for the outgoing frames in the egress.

2.6.15.3 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 2.6.15.1.

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.

2.6.15.4 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 2.6.15.1.

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.

2.6.15.5 Priority Extraction

The PRI[2:0] bits from tagged frames (section 2.6.15.1) 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 - Offset 0x18 and 0x19). IEEE tagged frames that get their priority re-mapped during ingress, will egress with the remapped 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 2.6.15.8).

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 (Offset 0x18).

2.6.15.6 Priority Override

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 2.6.7). 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 2.6.15.8).

2.6.15.7 CFI Bit

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

2.6.15.8 VID Extraction

If any of the three supported 802.1Q modes in enabled on this port (section 2.6.9) the VID read from tagged frames (as determined by section 2.6.15.1) 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 2.6.9) the VID bits from tagged frames (as determined by section 2.6.15.1) 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 2.8.1). 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).

2.6.15.9 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 2.6.9). 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).

2.6.15.10 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 (Offset 0x06) 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 34.

Table 34: 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 Force-DefaultVID = 1.
0x123	Enabled	Disabled	0x001	0x123	Use frame's VID.

2.6.16 Priority from IPv4 & IPv6 Frames

The format of an IPv4 TOS/DiffServ frame is shown in Figure 21 and an IPv6 Traffic Class frame is shown in Figure 22. 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 21: IPv4 Priority Frame Format

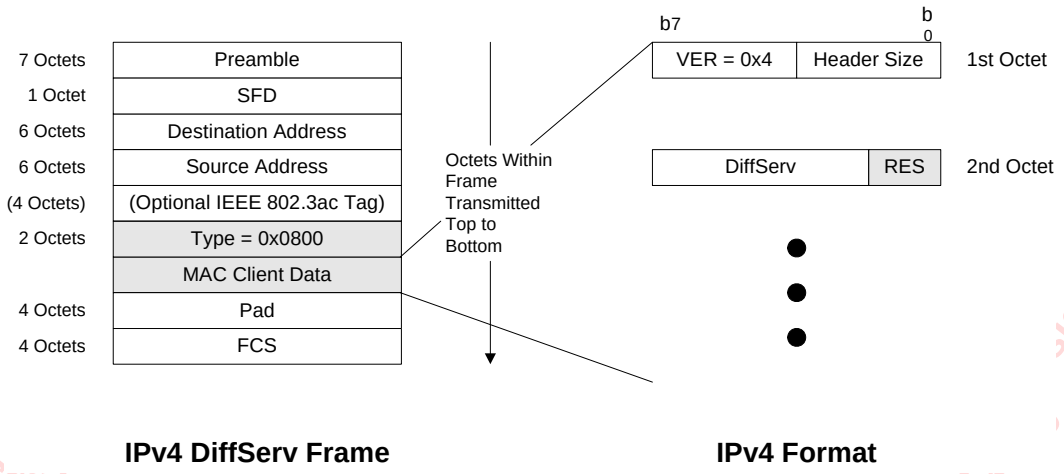

IPv4 DiffServ Frame **IPv4 Format**

Figure 22: IPv6 Priority Frame Format

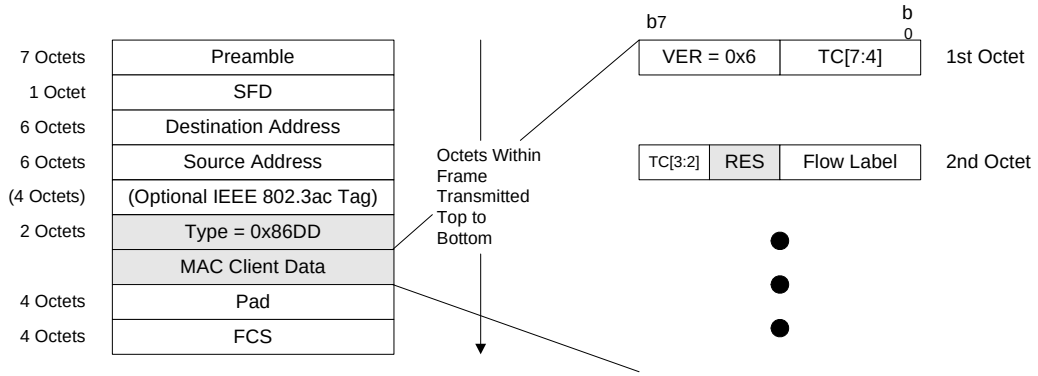

IPv6 Traffic Class Frame **IPv6 Format**

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 (Offset 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 2.6.17). 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.

2.6.17 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 23 and Figure 24. Management (MGMT) frames bypass IGMP/MLD snooping (see section 2.5.8).

Note

The gray portions of the frame in the figures are the only portions of the frame examined for this function.

Figure 23: IPv4 IGMP Snoop Format

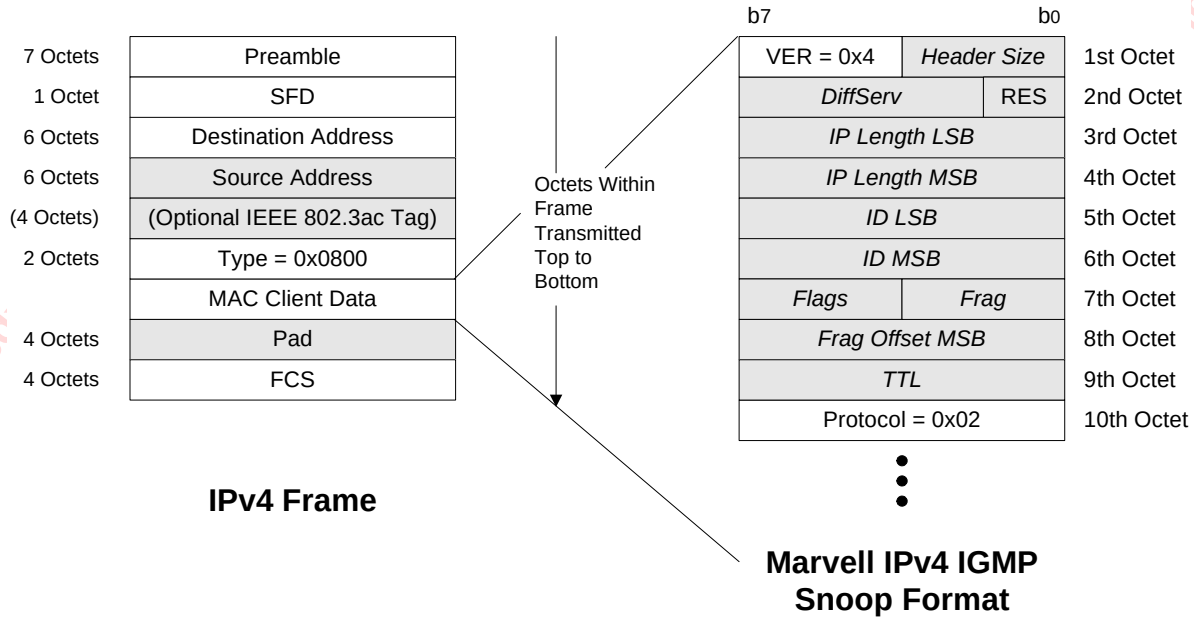

Figure 24: IPv6 MLD Snoop Format

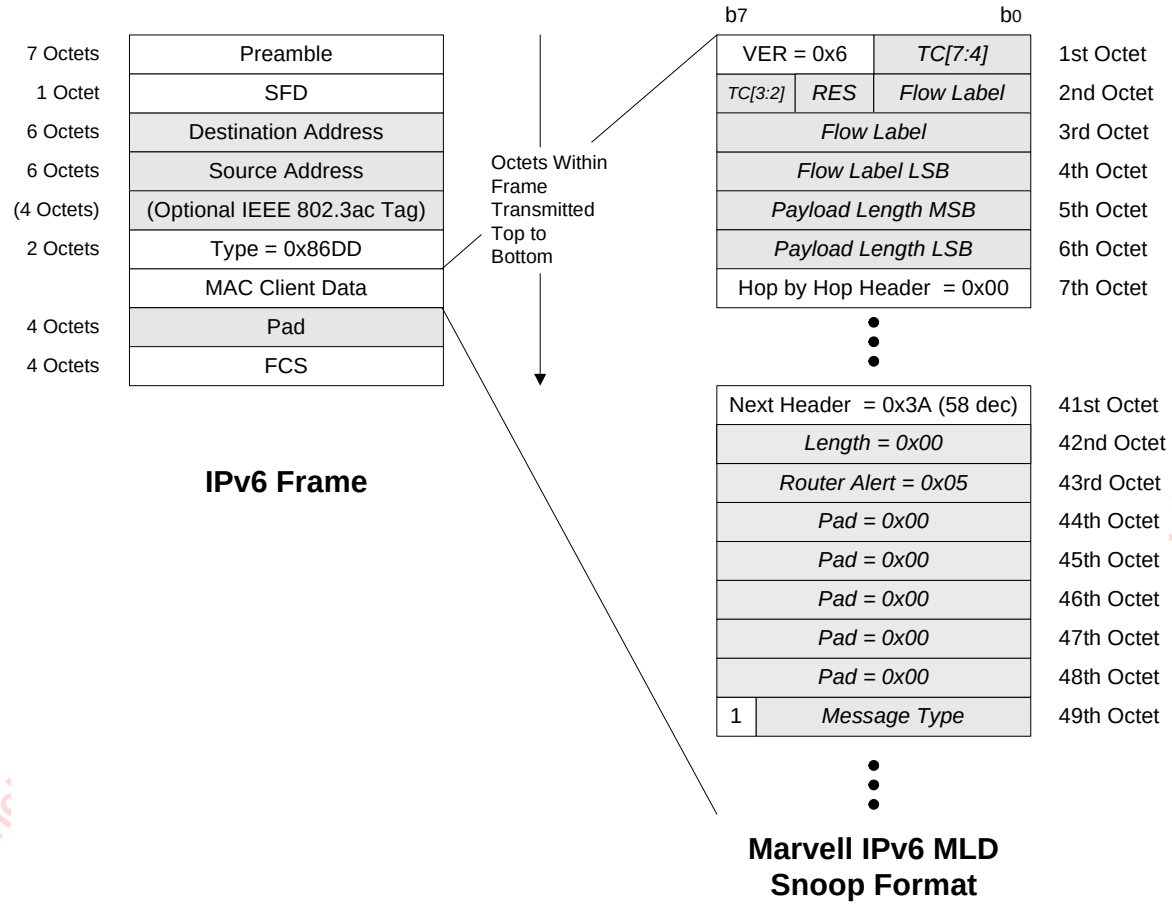

When one of these frames enters a port where IGMP/MLD snooping is enabled, the frame is sent to the CPU's port instead of where the Destination MAC Address directed it. 802.1Q VLAN membership and/or Port Based VLAN rules are still followed so the CPU port (in the local device, if cascaded devices are 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).

2.6.18 Ingress Rate Limiting

The role of ingress rate limiting in switches has been increasing as more and more applications require accurate and predictable rate limiting of traffic entering any given port. The switches may need to either limit traffic at an aggregate level from a port or limit specific streams of traffic entering a port like unicasts, multicasts, unknowns etc. 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 Bucket Configuration Memory bits using port's Ingress Rate Command and Ingress Rate Data registers (Offset 0x09 and 0x0A).

The devices support 'Best-in-Class' per port TCP/IP ingress rate limiting along with independent Storm prevention. Port based ingress rate limiting accommodates information rates from 64 Kbps to 1 Mbps in increments of 64 Kbps, from 1 Mbps to 100 Mbps in increments of 1Mbps and from 100 Mbps to 200 Mbps in increments of 10 Mbps.

One of the popular schemes for implementing rate limiting is a leaky bucket. The way a leaky bucket scheme works is that the bucket drains tokens constantly at a rate called Committed Information Rate (CIR) and bucket gets replenished with tokens whenever a frame is allowed to go through the bucket. All calculations for this bucket are done in tokens. Therefore, both bucket decrementing and incrementing is performed using tokens (i.e., frame bytes are converted into bucket tokens for calculation purposes).

These devices support a color blind leaky bucket scheme. A color blind scheme implies that the frames are not expected to be marked prior to coming into the system. In some routers it is required to have a color aware scheme of rate limiting where a frame marker would mark the frames to a lower priority than what they originally came in on if that traffic stream were to have violated any of the traffic rules. In the device's addressable applications, no such frame markers exist, thus a color blind scheme is employed.

As shown in the diagram below, the traffic below Committed Burst Size limit (CBS_Limit) is passed without any further actions. If the traffic burst were to continue and the bucket token depth approaches closer to the Excess Burst Size limit (EBS_Limit) by less than the CBS_Limit (i.e., EBS_Limit – CurrentBktDepth < CBS_Limit), then a programmable set of actions are specified.

Note

If the frame gets discarded then equivalent number of tokens for that frame will not get added to the bucket.

Figure 25: Color blind Leaky Bucket for Rate Limiting

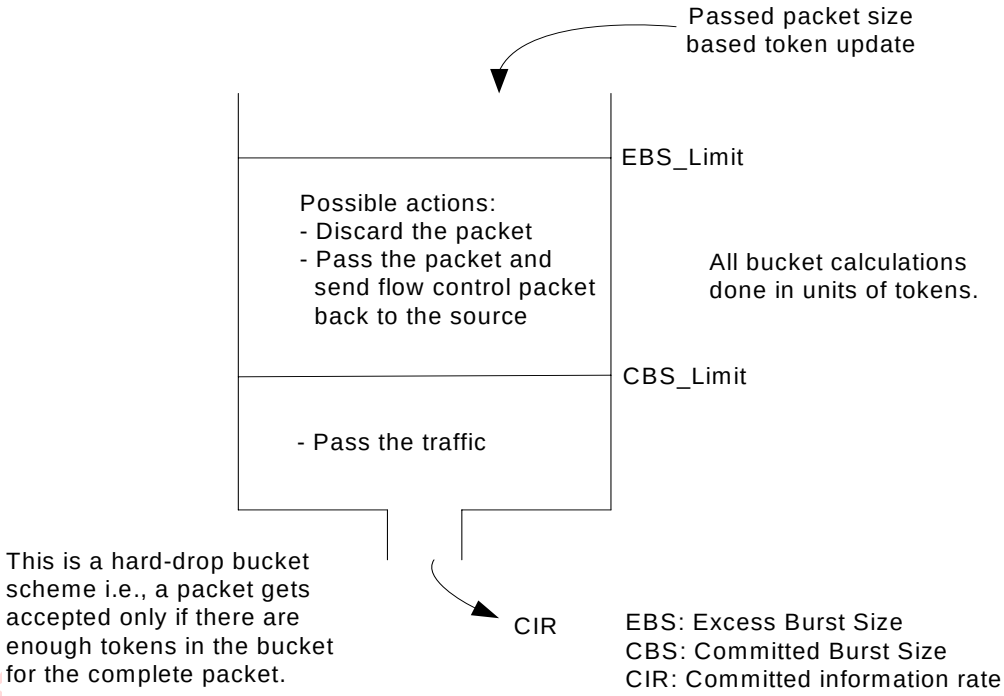

As shown in Figure 25, the traffic below CBS_Limit is passed without any further actions. If the traffic burst were to continue and the bucket token depth approaches closer to the EBSLimit by less than the CBSLimit (i.e., EBSLimit – CurrentBktDepth < CBSLimit), then a programmable set of actions are specified. Note that if the frame gets discarded then equivalent number of tokens for that frame will not get added to the bucket.

If the EBS_Limit_action were programmed to be in flow control mode, then an Ethernet flow control frame gets generated and sent to the source port but the incoming frame does get passed through the rate resource. If the port is operating in half-duplex mode then the port gets jammed.

At any point in the rate resource, if there are not enough tokens to accept a complete frame the frame wouldn't get accepted i.e., there is no concept of negative tokens for a given rate resource. In traffic management terminology this is a "strict" bucket implementation and not a "loose" one. The disadvantage of a strict implementation is that if there were to be a flow with Video/Audio streaming, as video frames are larger they may not get accepted but audio frames may go through. However, for TCP applications it is better to do a strict bucket implementation as large data frames won't get accepted till there is room for them and thus TCP ACK frames for previously transmitted frames could get accepted leading to lesser re-transmissions and better overall TCP throughput. These devices are optimized for TCP applications.

Figure 26 describes the general schema for PIRL where any one or more of the ports within the switch can be mapped onto any one or more of the buckets. Any given bucket can be programmed to be aggregate rate based or traffic type based.

Figure 26: Port based Ingress Rate Limiting Buckets

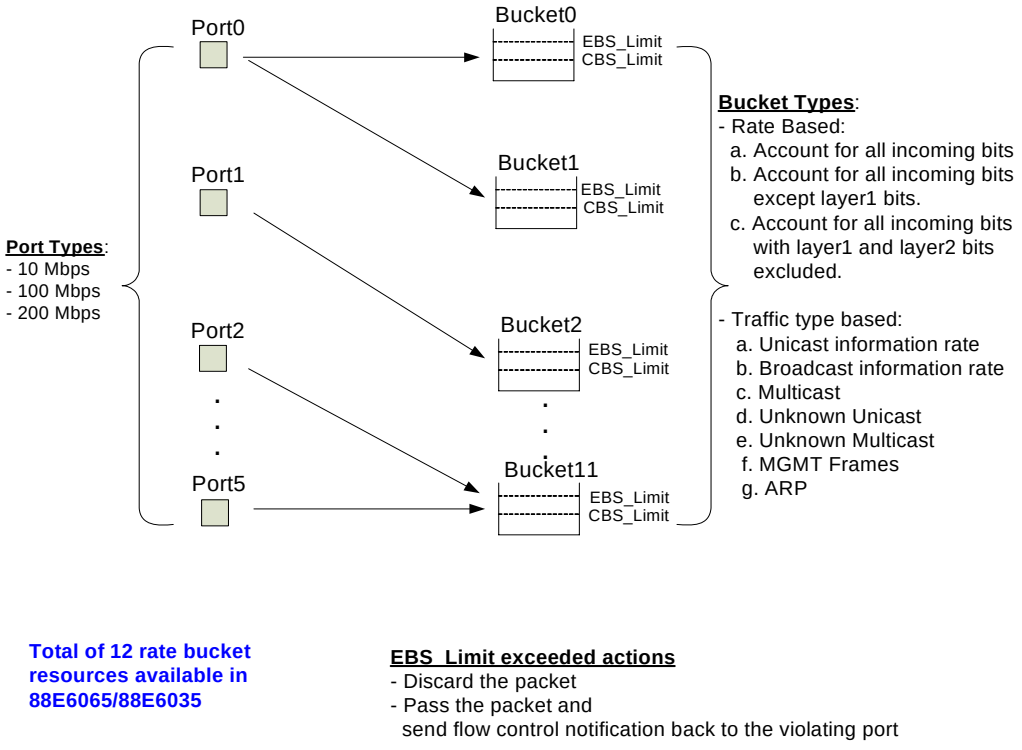

The following are the configurable parameters in this scheme.

- Any of the 12 PIRL resources can be assigned to any of the 6 ports.
- Any of the 6 ports can be assigned up to a maximum of 12 PIRL resources i.e., more than one resource can be assigned for a given port.
- Non-rate limiting (NRL) overrides can be programmed on a
  - Per-VID
  - Per-SA
  - Per-DA

For example if no rate limiting should be carried out for frames flowing through VID=391, the NRL bits for that particular VID can be set in the VTU entry.

The per-VID, per-SA or per-DA based NRL overrides are enabled by setting the corresponding bits in the register Ingress Rate Control register (Offset 0x09) bits 15 down to 13. Note that if either of the VID, SA or DA non-rate limit's are set, then both ingress and egress rate limiting logic would not account for that frame in their respective rate limiting calculations.

- Bytes to be counted
  - Account for all bytes or
  - Account for all bytes except for layer1 bytes
  - Account for all bytes except for layer2 bytes
  - Account for all bytes except for layer3 bytes
- TypeMask

  Any of the following frame types can be selected to be tracked as part of the rate resource calculations. As this is a bit vector, more than one frame type can be selected for a given rate resource.
  - Address Resolution Protocol (ARP), Management (MGMT), Multicasts, Broadcasts, Unicasts, Unknown Unicasts or Unknown Multicasts
- RateType
  - Rate based or traffic type based
- Bucket_Increment (12 bits)[1]
- BucketRateFactor (16 bits)[1]
- EBS/CBS Limits (24 bits)[1]
- EBS Limit Action
  - Discard
  - Send flow control

  Flow control mode is expected to be programmed on ports that have that have a trusted flow control mechanism available. EBSLimitAction is a per-port characteristic. That is if a port has been mapped to multiple rate resource buckets then all those buckets are expected to be programmed with the same EBSLimitAction.
- Account for Filtered discards

  This bit setting decides whether to account for frames that have been discarded because of other frame filtering reasons. To account for all frames coming into a given port(s) associated with this rate resource, this bit needs to be set.
- Account for Queue Congestion discards
- PIRLFCMode (Ingress Rate Control: Offset 0x09: bit12)

  This bit determines when the EBSLimitAction is programmed to generate a flow control message, the deassertion of flow control is controlled by the PirlFCMode bit. When this bit is programmed to a zero, flow control gets de-asserted only when the ingress rate resource has become empty and when this is programmed to a one, flow control gets de-asserted when the ingress rate resource has enough room left as specified by the CBSLimit. For example if CBSLimit is programmed to 0x60000, then if there is room for at least 0x60000 tokens available in the rate resource bucket then a frame is accepted.

2.6.19 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[2] 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

1. Refer to the Bucket Configuration register (PIRL registers - Offset 0x00) for further details.
2. The Header bit enables the Marvell Header mode for both ingress and egress.

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 27[1]. 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 frame 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 prepends 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).

Figure 27: Ingress Header Format

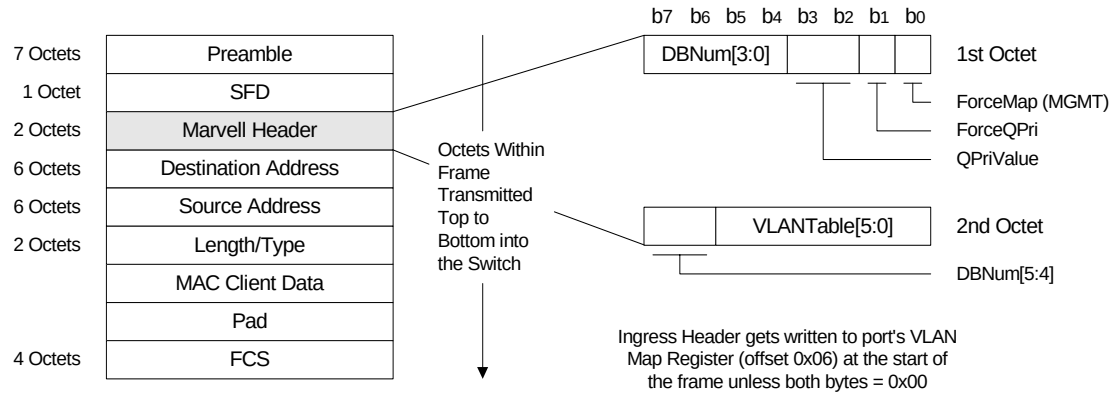

IEEE 802.3 Frame with Marvell Header In Front of DA

Marvell 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.

Note

The Marvell® Header can only modify the lower 4-bits of the port's DBNum. The upper 2-bits come from the port's DBNum[5:4] register bits from Port Control 1 register, Offset 0x05.

1. Reserved bits in the Marvell Header must always be zeros.

2.6.20 Provider/Customer Tagging

In provider network environments, it is very common to use double VLAN tagging to pass along the customer tag through the provider network by adding a provider tag on top of the customer tag. 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 Table 35.

The devices can parse the outer tag and compare that with a per-port programmable tag (Provider Tag Register, Offset 0x1A) value. If they match and if EgressTagMode is 0x3, the outer tag is identified to be a provider tag. All frames once they are identified to be provider tagged frames are stripped off of their outer tag and the next two bytes are compared with 0x8100. Till the next two bytes equal 0x8100, recursive tag stripping mechanism as specified above will continue unless RecStrDis bit (Provider Control Register, Offset 0x1C) is set to 0x1. The recursive tag stripping feature can be useful if the frame passes through multiple provider networks where each of the provider networks added a separate provider tag. The devices have the ability to remove all those provider tags recursively till the customer tag (identified by tag value of 0x8100) is surfaced.

The outer tag priority even though it gets stripped in the ingress can be used for either frame/queue priority determination and/or putting it back for the egressing frame. Similarly, VID bits from the removed tag can be used for switching, if ForceDefault (Offset 0x07) is set to a zero.The following figure illustrates the device functionality for various ingressing frame types to various egressing frame types. Note that if the incoming frames outer tag value equals 0x8100 then that frame is considered a customer port frame.

Table 35: Provider/Customer Port Supported Tagging Options

	Egress Customer Port Untagged	Egress Customer Port Tagged	Egress Customer Port UnModified	Egress Provider Port Untagged	Egress Provider Port Tagged	Egress Provider Port Unmodified
Ingress Customer Port Untagged	YES	YES	Same as Ingress Cust. untagged to Egress Cust untagged case	N/A	YES	N/A
Ingress Customer Port Tagged	YES	YES	YES	N/A	YES	N/A
Ingress Provider Port Untagged	N/A	N/A	N/A	N/A	N/A	N/A
Ingress Provider Port Tagged	YES	Supported if egress tag same as ingress customer tag	N/A	N/A	YES	YES

The EgressTagMode (Port Control register, Offset 0x04) defines how a frame ought to be transmitted out of the switch. The frame can leave the switch unmodified, tagged or untagged. If the ingress pipe determined the frame to be of MGMT type, then the egress wouldn't modify the frame for adding or removing any tags as the processing agent like a micro-processor would want to receive these MGMT frames as they arrived into the switch.

As specified in Table 35, if the egress port is a provider port then the per-port Provider Tag (Offset 0x1A) will get added to the frame before being transmitted. The frame priority for these frames would be the one selected in the ingress and passed along through the queue controller to egress. Since the frame is being modified the frame CRC will get recomputed and attached at the end of the frame.

If the egress port is to a customer port and if the frame is expected to be tagged leaving the switch, the VID selected in the ingress pipe (either port based or incoming frame based) along with the frame priority information from the ingress pipe is added to the frame with the recomputed CRC at the end of the frame.

If the egress port is to a non-customer port and if the frame is expected to be tagged leaving the switch, a tag value of 0x8100 is used. Similar to the priority port case specified above, the frame priority bits selected in ingress and passed along through the queue controller get used as the priority information for the tag.

If the EgressTagMode is to send out frames untagged and if the resulting frame is less than 60 Bytes, then the egress pipe would add a padding to the frame to make it at least 64 bytes before transmitting the frame out of the switch.

2.6.21 Port States

The 88E6065/88E6035 device supports four Port States per port shown in Table 36. The Port States are used by the Queue Controller (see section 2.6.21) in the 88E6065/88E6035 device 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 they can be modified at any time.

Table 36 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 2.5.8. Their primary purpose is to support the Spanning Tree Protocol (see section 2.9) so these frames have the ability to tunnel through blocked ports. 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. These MGMT frames are typically used for 802.1D Spanning Tree Bridge Protocol Data Units (BPDUs), but any multicast address can be used supporting new and/or proprietary protocols.

Table 36: 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 frames even if they are not MGMT frames.
Forwarding	Normal operation. All frames are allowed to enter (ingress) and leave (egress) a Forwarding port. Learning takes place on all good frames.

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 13). 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.

2.7 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 supports 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).

2.7.1 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.

2.7.2 No Head-of-Line Blocking

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

2.7.3 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 2.7.5). 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 88E6065/88E6035 device is 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 2.7.5).

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 2.7.3

2.7.4 Guaranteed Frame Delivery without Flow Control

The devices can guarantee high priority frame delivery[1], even with Flow Control disabled, due to its intelligent Queue Controller 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.

2.7.5 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 (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.

2.7.6 The Queues

The queues in the devices are shown in Figure 28.

Figure 28: Switch Queues

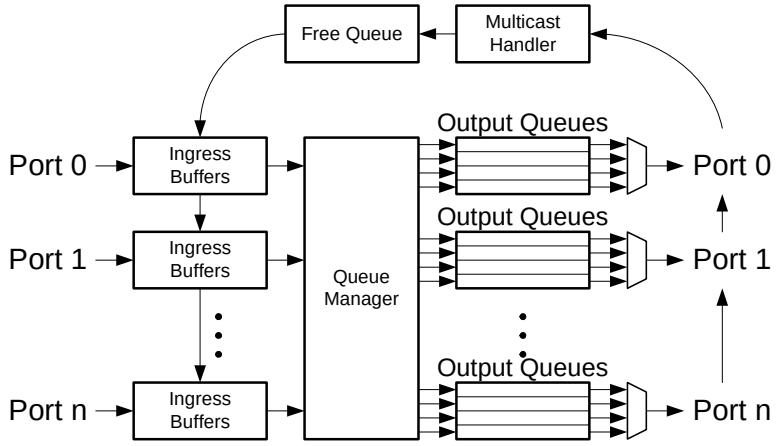

2.7.7 Queue Manager

At reset[2], 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[3] and whose link 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

1. If all the frames entering a port are all high priority at wire speed their delivery cannot be guaranteed.
2. 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 (Offset 0x04).
3. If a port is in the Disabled Port State (see section 2.6.9) its Ingress buffers are left in the Free Queue for other ports to use.

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 2.5.1) and the Ingress Policy (see section 2.6) 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 disables 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.

2.7.8 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 (Offset 0x04). A fixed or a weighted priority can be selected (see section 2.7.5).

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.

2.7.9 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.

2.8 Egress Policy

If directed, the Egress Policy block modifies frames, as they exit the switch. IEEE 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.

Note

If the frame is classified as a MGMT frame by the ingress pipe, then the frame doesn't get modified for any tagging reasons.

For all frames that get modified, Layer 2 CRC is recomputed and the original frame's CRC gets replaced by the newly computed CRC by the egress module before sending the frame out of the port. Egress module checks the CRC for the frames read out of frame memory and if the CRC happens to be not good then the frame gets sent out with the bad CRC.

For frames that get stripped off of a tag before transmitting and those that result in less than 64bytes of total frame size get padded to at least 64 byte frames.

2.8.1 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, Offset 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 2.6.13). 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 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 2.8.1.1.
- Transmit all frames Double Tagged (see section 2.9).

The format of an IEEE Tagged frame is shown in Figure 29.

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 29: IEEE Tag Frame Format

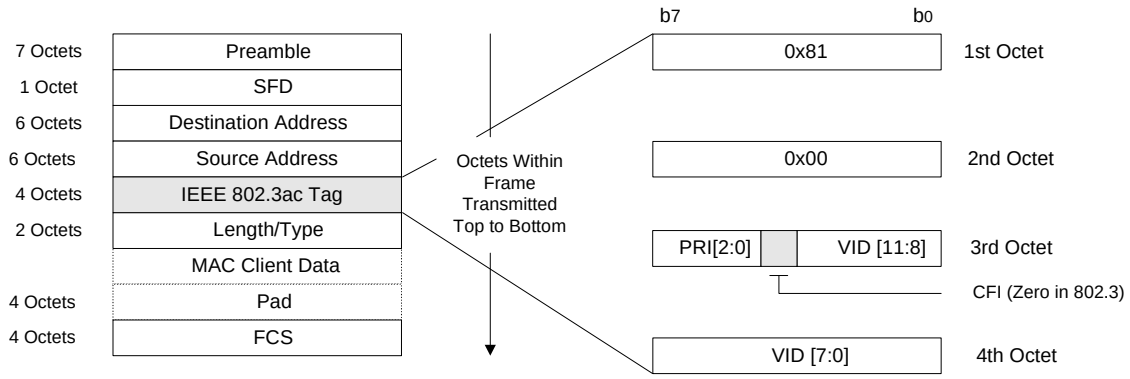

IEEE 802.3 Tagged Frame

Tag Format

2.8.1.1 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 2.6.8).
- The CFI bit is always set to a zero.
- The VID bits indicate the VID assigned to the frame during Ingress (see section 2.6.15.8).

Note

A Tag that is added due to Egress Double Tagging is done differently. See section 2.8.2.1.

2.8.1.2 Priority Re-Map and Priority Override

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 2.6.8). The PRI bits can be re-mapped by the ingress port's IEEE Priority Remapping registers (Offset 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.

2.8.1.3 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 2.6.15.8.

2.8.2 Egress Double VLAN Tagging

Double Tagging is a way to isolate one IEEE 802.1Q VLAN from other IEEE 802.1Q VLANs in a hierarchical fash¬ion 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 30.

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 Provider/Customer Tagging enabled (see section 2.6.20) 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 30: Double Tag Format

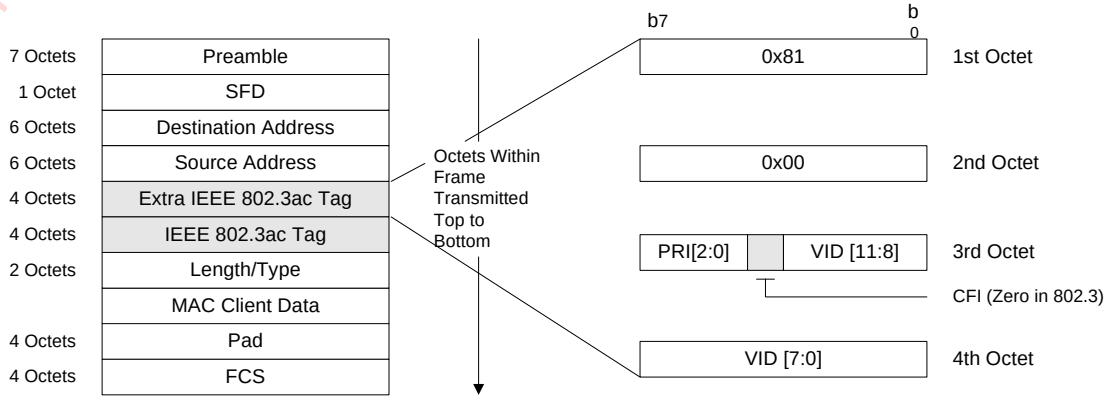

2.8.2.1 Always Adding a Tag or Double Tagging Frames

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 ProviderTag[15:8] from the per-port ProviderTag register (Offset 0x1A)
- The 2nd Octet is always ProviderTag[7:0] from the per-port ProviderTag register (Offset 0x1A).
- The PRI bits indicate the frame's priority determined during Ingress (see section 2.6.7).
- The CFI bit is always set to a zero.
- If the Ingress determines the frame to be from a provider port, the VID bit comes from the ingress pipe or else it comes from the port's default port VLAN ID.

Note

A Tag that is added due to normal IEEE Tagging is done differently. See section 2.8.1.1.

2.8.3 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 2.6.1, and if 802.1Q is enabled on the port, the VTU (section 2.6.12) which supports 802.1s Port States defined in section 2.6.11.

2.8.4 Egress Rate Limiting

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 0x0A). Egress rate limiting is performed by shaping the output load.

The egress rate limiting scheme works on the principle of adjusting the inter frame gap to achieve the desired information rate. The primary advantage with this scheme is that there are no bursts introduced by this switch onto the link partner thus reducing the burst buffering requirement of the link partner.

The required maximum rate must be selected and programmed. The 88E6065/88E6035 supports 4095 different rate speeds or shapes from 62 Kbits/sec to 200 Mbps (Egress rate shaping can be disabled by setting the Egress-Rate bits to all zero). The programmed register value uses the following equation:

EgressRate's Register Value = 8 bits/(40 ns * Egress Rate bits/sec)

For example: Egress Rate limiting needs to be set to 256 Kbps.

EgressRate = 8 bits/(40 ns * 256,000 bits/sec)
EgressRate = 8 bits/(0.04 bits)
EgressRate = 781 or 0x30D

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 are fixed to 8 bytes long. The frame's minimum IFG bytes can be fixed to a 12-byte value.

Note

The bytes transmitted on collisions on half-duplex ports are not accounted for in the egress rate limiting calculations.

2.8.5 Switch Egress Header

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 31 and its fields are defined in Table 37.

Figure 31: Egress Header Format

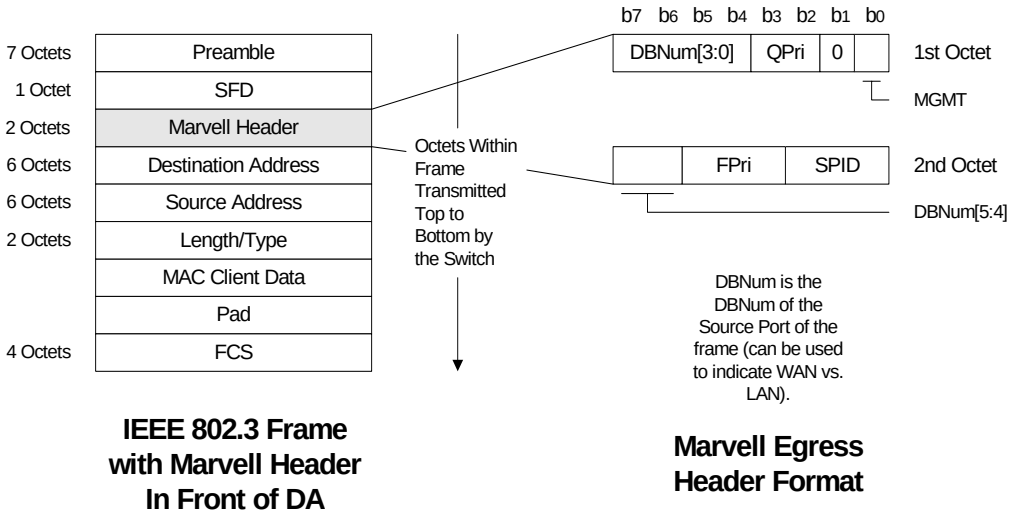

1. The Header bit enables the Header mode for both Ingress and Egress

Table 37: 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. It can be used to indicate the logical port based VLAN number of the source port.
PRI[2:0]	The frame's queue priority. This is the queue priority of the frame that is being transmitted out.This enables switches to communicate queue priority with their link partner.
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, etc.
DBNUM[5:4]	Database Number. This field represents the higher 2 bits of the address database number assigned to this frame when it Ingressed into this device. It can be used to indicate the logical port based VLAN number of the source port.

2.9 Port Trunking Support

A total of four Port based trunks are supported by the devices with any combinations of ports. 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 4 trunk groups are supported with up to 6 ports per trunk group for the devices. Load balancing is DA/SA based using the Trunk Mask Table (in global 2, Offset 0x07). If HashTrunk bit is set to a one, the hash computed for ATU lookups is used for the trunkmask selection and if the HashTrunk bit is set to a zero, the lower two bits of the frame DA and SA are XOR'ed to select the trunkmask.

2.10 Spanning Tree Support[1]

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[2] frames entering Network (external switch) ports. These frames are called MGMT (management) frames in the device (see section 2.5.8). 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 2.5.5) 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 2.6.1 and section 2.6.12). 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 thought 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 by the source port's CPUPort register (in port Control 2, Offset 0x08).

1. Rapid Spanning Tree works best if the backup port is inside the same device as the primary port. This way the flows can be transferred inside the address database quickly. If the backup port for a primary port is contained in a external cascaded device the ATU Move command can only be used to remove the flow from the primary port, but it cannot be used to transfer the ATU data to the backup port's address database.
2. BPDU = Bridge Protocol Data Unit - the frame type used to run the Spanning Tree Protocol (STP).

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.

2.11 Embedded Memory

The devices MACs interface directly to the embedded 1/2Mb (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.

2.12 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 82) 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 (Offset 0x00) 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.

2.13 MGMT Frame Handling in the Switch

MGMT (management) frames use their own set of Ingress Rules and Egress Rules when it comes to Destination Port mapping, what Priority to use, the VLAN decisions made, Ingress Rate Limiting, Port States and Egress tagging. They must follow all the other Ingress Rules like proper frame size, proper CRC, etc.

2.13.1 What is a MGMT Frame?

A frame is considered a MGMT frame when:

- On Ingress the DA address of the frame is a multicast address that is present in the address database with a MGMT EntryStates as defined in the AU Data register (Offset 0x0C), or
- If the per port register bit ForceMap (Offset 0x06) is set to a one, then all frames coming in from that port are considered MGMT frames. This forcing overrides ATU based determination.

2.13.2 DA Based MGMT Frame Rule Changes

If an ingressing frame is determined to be a MGMT frame by its DA address, the frame will be switched to the port(s) defined in the DA's ATU Entry's Destination Port Vector (DPV - section 2.5.8). This DPV will not be modified by any VLAN rules (either port based or 802.1Q based) nor will it be modified by IGMP/MLD snooping. The only potential modification that may occur to this DPV is the removal of the Source Port's bit if frames are not to be switched back to the their source port (normal switch operation – section 2.6.10).

The frame's priority is the priority defined in the DA's ATU Entry's Priority bits (section 2.5.8) as long as the source port's DAPriOveride bit is set to a one (See Port Control 2 register - Offset 0x08). The ATU's priority is 2 bits while the frame's priority is 3 bits. The upper two bits of the frame's priority comes from the ATU's priority with the frame's LSB of the priority being a copy of the ATU's MSB of the priority. The Queue Controller's priority is the ATU's priority. Since this is the highest priority selection for a normal frame these priorities will not be modified.

2.13.3 Ingress Rate Limiting of MGMT Frames

If the frame is determined to be a MGMT frame during ingress (by any method) the frame will be Ingress Rate Limited if the source port's Ingress Rate Limiting is enabled and LimitMGMT is set to a one (See Ingress Rate Control register - Offset 0x09).

2.13.4 Port State Handling of MGMT Frames

MGMT frames are allowed to Ingress and/or Egress any port configured in the Blocked/Listening or Learning port states (See Port Control register - Offset 0x04). They are the only frames that can pass through (or tunnel through) ports in these states.

2.13.5 Egress Rate Shaping of MGMT Frames

If a frame is determined to be a MGMT frame with non rate limiting during Ingress this information will be passed through the Queue Controller and the frame will not be Egress Rate Shaped (or delayed).

2.13.6 Egress 802.1Q Changes for MGMT Frames

Egress 802.1Q Membership checking is not performed on MGMT frames even if 802.1Q is enabled on the egress port (i.e., they will not be discarded due to the egressing port not being a member of the frame's VID).

2.13.7 Egress Tag Changes for MGMT Frames

If a frame is determined to be a MGMT frame during ingress this information will be passed through the Queue Controller and the frame will not be modified for any egress tagging purposes. For MGMT frames, the frame processor (typically an external CPU) needs to be sent out as it entered the switch. Thus these devices perform no frame modifications once it is determined to be an MGMT frame.

Egress tag rules for MGMT frames change depending upon the mode of frame's Ingress port and the mode of the frame's Egress port.

Section 3. Physical Interface (PHY) Functional Description

The 88E6065/88E6035 device contains five/two IEEE 802.3 100BASE-TX and 10BASE-T compliant media-dependent interfaces 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.

The Port 0 and Port 1[1] interfaces can be configured to support IEEE 802.3 100BASE-FX by utilizing a pseudo-ECL (PECL) interface for fiber-optics.

1. For 88E6065 only

Figure 32: 88E6065/88E6035 Device Transmit Block Diagram

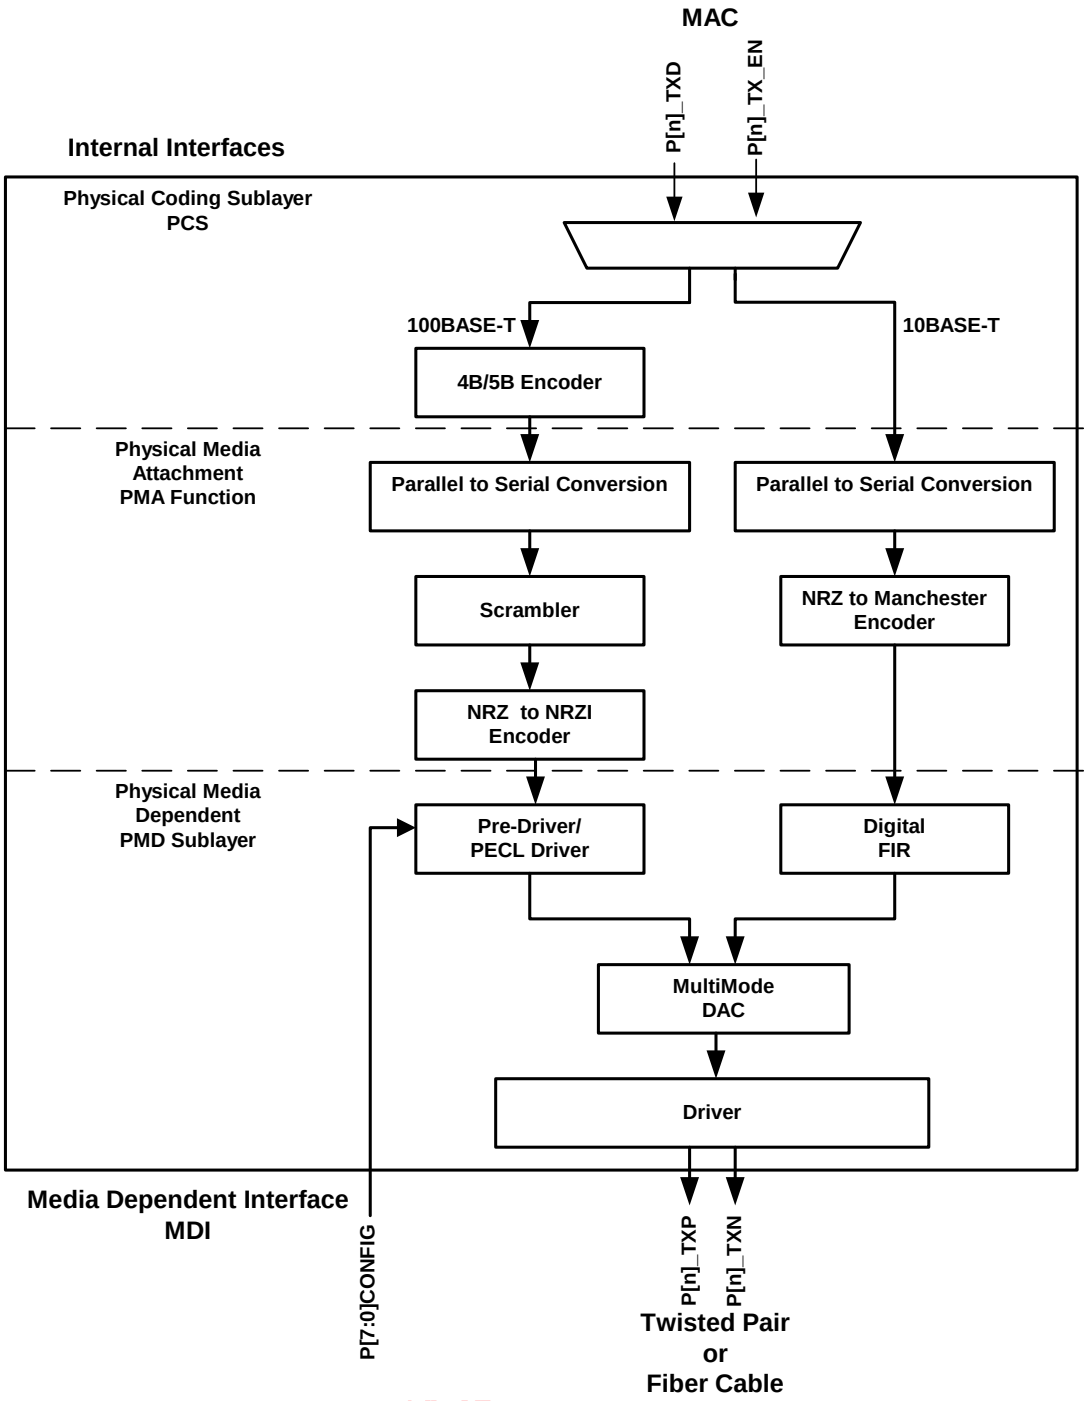

Figure 33: 88E6065/88E6035 Device Receive Block Diagram

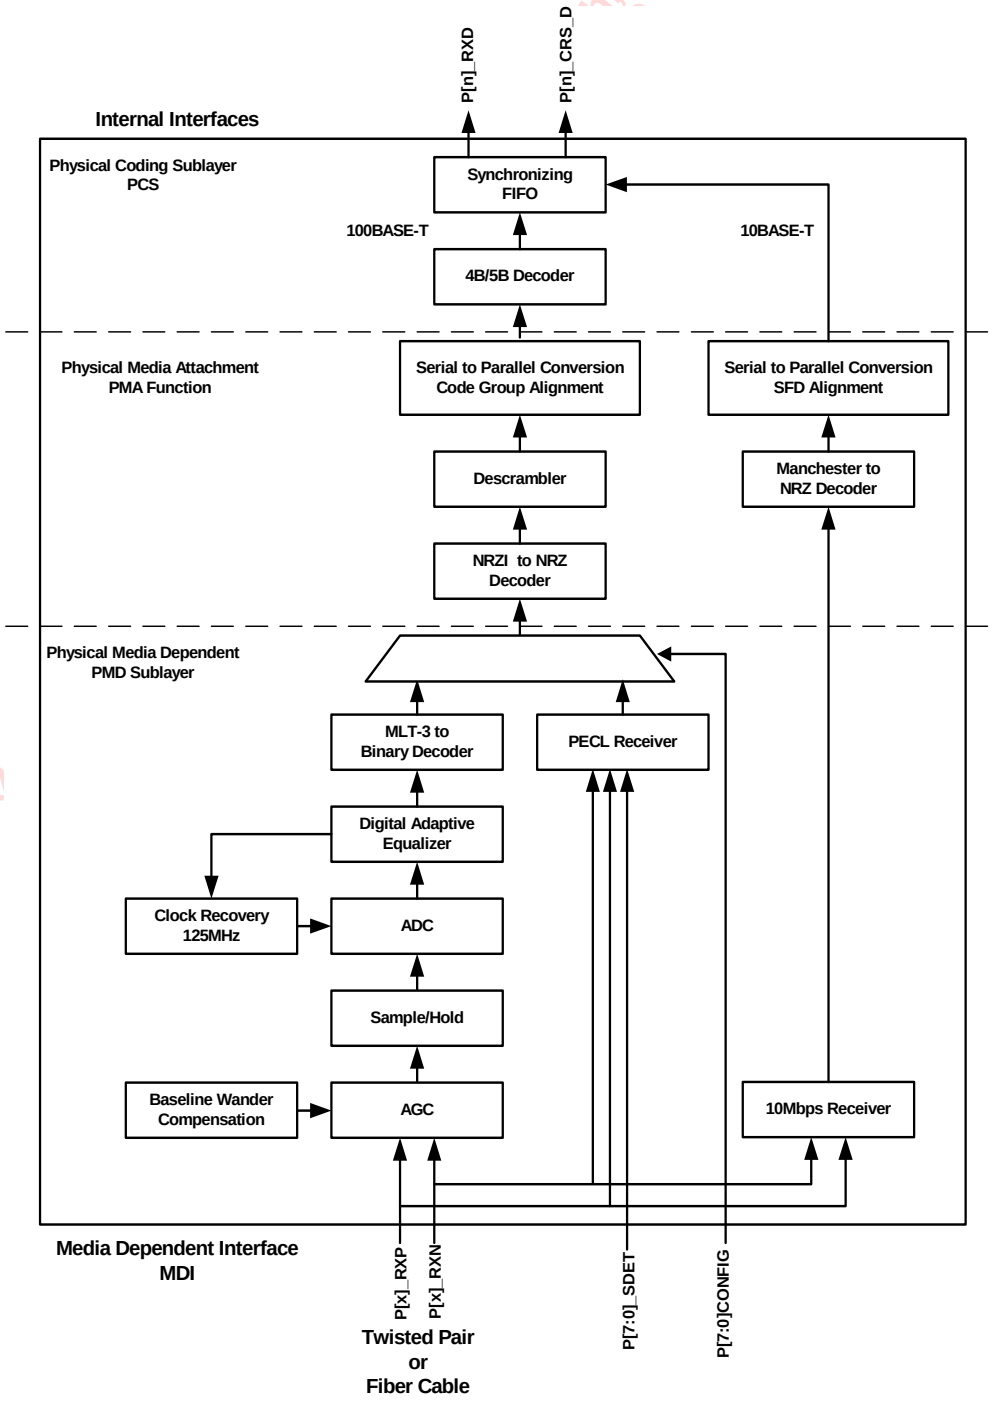

3.1 Transmit PCS and PMA

3.1.1 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.

3.1.2 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.

3.1.3 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.

3.1.4 NRZ to NRZI Conversion

The data stream is converted from NRZ to NRZI.

3.1.5 Pre-Driver and Transmit Clock

The 88E6065/88E6035 device uses 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.

3.1.6 Multimode Transmit DAC

The multimode transmit digital to analog converter (DAC) transmits MLT3-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.

3.2 Receive PCS and PMA

3.2.1 10-BASE-T/100BASE-TX Receiver

The differential RXP and RXN pins are shared by the 100BASE-TX, 100BASE-FX (supported on Port 0 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.

3.2.2 AGC and Baseline Wander

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

3.2.3 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. This equalizer removes inter-symbol interference at the receiver. The digital adaptive equalizer takes unequalized signals from the ADC output and uses a combination of feed-forward equalizer (FFE) and decision feedback equalizer (DFE) for the best optimized signal-to-noise (SNR) ratio.

3.2.4 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 88E6065/88E6035 receiver path robust in the presence of variations in process, temperature, on-chip noise, and supply voltage.

3.2.5 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.

3.2.6 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 descrambler bit (See ATU Data register - Offset 0x0C).

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.

3.2.7 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.

3.2.8 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 38.

3.2.8.1 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.

3.2.8.2 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.

3.2.8.3 Far End Fault Indication (FEFI)

When 100BASE-FX is selected and Bit 0 of CONFIG_A is low at hardware reset, the far end fault detect (FEFD) circuit is enabled. The FEFD enable state can be overridden by programming the FEFI bit (See PHY Specific Control register - Offset 0x10).

Note

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

3.2.8.4 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 ensure 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 ExtendedDistance bit in the PHY Specific Control Register (Offset 0x10).

Table 38: 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 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

3.2.9 Setting Cable Characteristics

Since cable characteristics differ 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 (Offset 0x10) for cable type.

3.2.10 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 (Offset 0x10).

The scrambler setting is also controlled by hardware configuration at the end of hardware reset. Table 39 shows the effect of various configuration settings on the scrambler.

Table 39: 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)

3.2.11 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 88E6065/88E6035 device through SDET pin for Port 0 and Port 1.

If Force Link Good is asserted (ForceLink bit is set high - PHY Specific Control Register - Offset 0x10), 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 - Offset 0x10) is set high. If Auto-Negotiation is disabled and DisNLPGen (PHY Specific Control Register - Offset 0x10) is set high, then the link pulse transmission is disabled.

3.2.12 Auto-Negotiation

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

- Power up reset
- Hardware reset
- Software reset (SWReset bit-PHY Control Register - Offset 0x00)
- Restart Auto-Negotiation (RestartAneg bit-PHY Control Register - Offset 0x00)
- Transition from power down to power up (PwrDwn bit-PHY Control Register - Offset 0x00)
- Change from the linkfail state to the link-up state

If Auto-Negotiation is enabled, the 88E6065/88E6035 device negotiates 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 88E6065/88E6035 device goes 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 - Offset 0x00). When Auto-Negotiation is disabled, the speed and duplex can be changed via the SpeedLSB and Duplex bits (PHY Control Register - Offset 0x00), respectively. The abilities that are advertised can be changed via the Auto-Negotiation Advertisement Register (Offset 0x04).

3.2.13 Register Update

Changes to the AnegEn, SpeedLSB, and Duplex bits (Offset 0x00) do not take effect unless one of the following takes place:

- Software reset (SWReset bit - Offset 0x00)
- Restart Auto-Negotiation (RestartAneg bit - Offset 0x00)
- Transition from power down to power up (PwrDwn bit - Offset 0x00)
- Loss of the link

The Auto-Negotiation Advertisement register (Offset 0x04) 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 88E6065/88E6035 device begins to transmit Fast Link Pulses (FLPs). This guarantees that a sequence of FLPs transmitted is consistent with one another.

The Next Page Transmit register (Offset 0x07) is internally latched once every time Auto-Negotiation enters the next page exchange state in the arbitration state machine.

3.2.14 Next Page Support

The 88E6065/88E6035 device supports 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 (Offset 0x05). The 88E6065/88E6035 device has an option to write the received next page into the Link Partner Next Page register - Offset 0x08 - (similar to the description provided in the IEEE 802.3ab standard) by programming the Reg8NxtPg bit (PHY Specific Control Register - Offset 0x10).

3.2.15 Status Registers

Once the 88E6065/88E6035 device completes Auto-Negotiation it updates the various status in the PHY Status (Offset 0x01), Link Partner Ability (Next Page) (Offset 0x05), and Auto-Negotiation Expansion (Offset 0x06) registers. Speed, duplex, page received, and Auto-Negotiation completed status are also available in the PHY Specific Status (Offset 0x10) and PHY Interrupt Status registers (Offset 0x13).

3.3 Power Management

The 88E6065/88E6035 supports advanced power management modes that conserve power. These features are only recommended for power-critical designs only.

3.3.1 Low Power Modes

Two low power modes are supported in the 88E6065/88E6035.

- IEEE 802.3 Clause 22.2.4.1.5 compliant power down
- Energy Detect+™

IEEE 802.3 Clause 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 88E6065/88E6035 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.

3.3.2 MAC Interface and PHY Configuration for Low Power Modes

The devices have one CONFIG bit (in CONFIG_B - Table 2) 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 (See PHY Specific Control register - Offset 0x10) enables the user to turn on Energy Detect+™. When the low power mode is not selected, the PwrDwn bit (See PHY Control register - Offset 0x00) 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 88E6065/88E6035 devices return to the mode stored in the EDet bit.

Table 40 shows how these modes are entered

Table 40: Operating Mode Power Consumption

Power Mode	Est. Power	How to Activate Mode
IEEE Power down	See Section 7.	PwrDwn bit write (See PHY Control register - Offset 0x00)
Energy Detect+™	Same as IEEE Power down mode—Section 7.	Configuration option & register EDet bit write (See PHY Specific Control register - Offset 0x10)

3.3.3 IEEE Power Down Mode

The standard IEEE power down mode is entered by setting the PwrDwn (See PHY Control register - Offset 0x00) 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.

3.3.3.1 Energy Detect +TM

In this mode, the PHY sends out a single 10 Mbps NLP (Normal Link Pulse) every second. If the 88E6065/ 88E6035 is in Energy Detect+TM mode, it can wake a connected device and it wakes upon the detection of a connected device. When ENA_EDET is 1, the mode of operation is Energy Detect+TM.

3.4 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 (Port 0 and Port 1 only).

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 (See PHY Specific Control register - Offset 0x00).

Table 41 shows the various configuration settings affecting the FEFI function on hardware reset.

Table 41: FEFI Select

EN_FEFI (CONFIG_A) Connection	FEFI	DisFEFI Bit[1]
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

1. (PHY Control register - Offset 0x00)

3.5 Virtual Cable Tester™

The device's Virtual Cable Tester™ 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 mismatches, bad connectors, termination mismatches, 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 (PHY register - Offset 0x1A and PHY register - Offset 0x1B) on the AmpRfln and DistRfln bits respectively.

Using the information from the VCT registers (PHY register - Offset 0x1A and PHY register - Offset 0x1B), 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 PHY register - Offset 0x1A and Figure 49. 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 (PHY register - Offset 0x1A), a predetermined 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 (PHY register - Offset 0x1A and PHY register - Offset 0x1B).

- 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 (PHY register - Offset 0x1A and PHY register - Offset 0x1B). The impedance mismatch at the location of the discontinuity could also be calculated knowing the magnitude of the reflection. Refer to the Application Note "Virtual Cable Tester -- How to use TDR results" for details.

3.6 Data Terminal Equipment (DTE) Detect

The devices supports 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 offers 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 (PHY Register offset 0x10 bit 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 devices. 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 (PHY Register offset 0x11 bit 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 (PHY Register offset 0x12 bit 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 (PHY Register offset 0x16 bit 15) and the DTE Detect Status bit (PHY Register offset 0x11 bit 15).

If a link partner that requires DTE power is unplugged, the DTE Detect Status bit (PHY Register offset 0x11 bit 15) will drop after a user controlled delay (default is 20 seconds - DTE Detect Status Drop bit (PHY Register offset 0x16 bit 1514) = 0x04), since the low pass filter (or similar fixture) is removed during power up.

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

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 42: Register 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 (PHY Register offset 0x11 bit 15). The wait time is 5 seconds multiplied by the value of these bits. Default at HW reset: 0x4 At SW reset: retain

3.7 Auto MDI/MDIX Crossover

The 88E6065/88E6035 device automatically determines whether or not it needs to interchange cable sense between pairs so that an external crossover cable is not required. If the 88E6065/88E6035 device interoperates with a device that cannot automatically correct for crossover, the 88E6065/88E6035 PHY makes the necessary adjustment prior to commencing Auto-Negotiation. If the 88E6065/88E6035 device interoperates 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 interoperates 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 uses signal detect to determine whether or not to crossover.

The Auto MDI/MDIX crossover function can be disabled via the AutoMDI[X] bits (PHY Register - Offset 0x10).

The 88E6065/88E6035 device is set to normal mode by default if auto MDI/MDIX crossover is disabled at hardware reset.

The pin mapping in normal and crossed modes is specified in Table 43.

Table 43: MDI/MDIX Pin Functions

Physical Pin	Normal		Crossed	
	100BASE-TX	10BASE-T	100BASE-TX	10BASE-T
TXP/TXN	Transmit	Transmit	Receive	Receive
RXP/RXN	Receive	Receive	Transmit	Transmit

3.8 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 34. This allows the link partner to receive its own frames.

Figure 34: Line Loopback Data Path

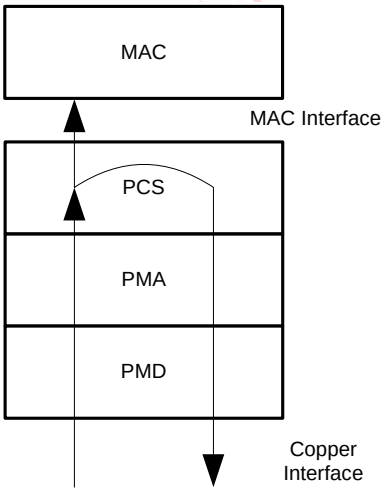

Before enabling the line loopback feature, the PHY must first establish link to another PHY link partner. If Auto-Negotiation is enabled, both link partners should advertise the same speed and full-duplex. If Auto-Negotiation is disabled, both link partners need to be forced to the same speed and full-duplex. 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)

3.9 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 (PHY register - Offset 0x19). 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 (PHY register - Offset 0x18). 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 (PHY register - Offset 0x18). The default blink period is set to 84 ms. The blink rate applies to all applicable LEDs.

3.9.1 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 the LED Parallel Select register (Offset 0x16) are based on the LED_DEF[1] and LED_DEF[0] values during hardware configuration through the CONFIG_A pin—see Table 2.

Table 44: Parallel LED Hardware Defaults

LED Mode - set by CONFIG_A at reset	P[4:0]_LED2	P[4:0]_LED1	P[4:0]_LED0
0	LINK	RX	TX
1	LINK	ACT	SPEED
2	LINK/RX	TX	SPEED
3	LINK/ACT	DUPLEX/COLX	SPEED

1. Energy Detect (section 3.3) must be disabled on ports that are configured to blink an LED but don't have a link established.

The LED Parallel Select register (Offset 0x16) shows additional display modes that can be set up by software after startup. Table 45 defines all the possible Parallel LED Display modes.

Table 45: 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.
DUPLEX/COLX	Low = Full-duplex High = Half-duplex Blink = Collision activity (blink rate is 84ms active then 84 ms inactive) The collision activity is pulse stretched to 84 ms.
ACT (blink mode)	High = No transmit or receive activity Blink = Transmit or receive activity (blink rate is 84ms active then 84 ms inactive) The transmit and receive activity is pulse stretched to 84 ms.

3.9.2 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 (Offset 0x19)).

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 35: Possible Solutions for Case One

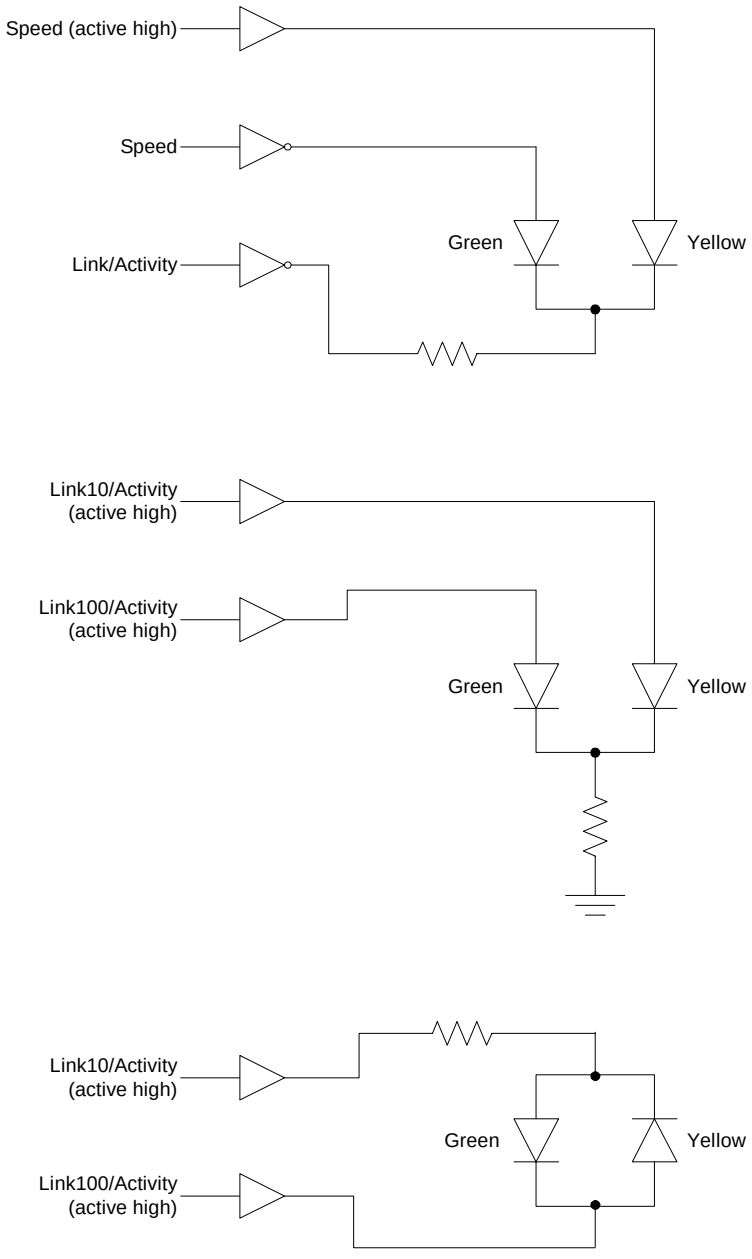

Figure 36: Possible Solutions for Case Two

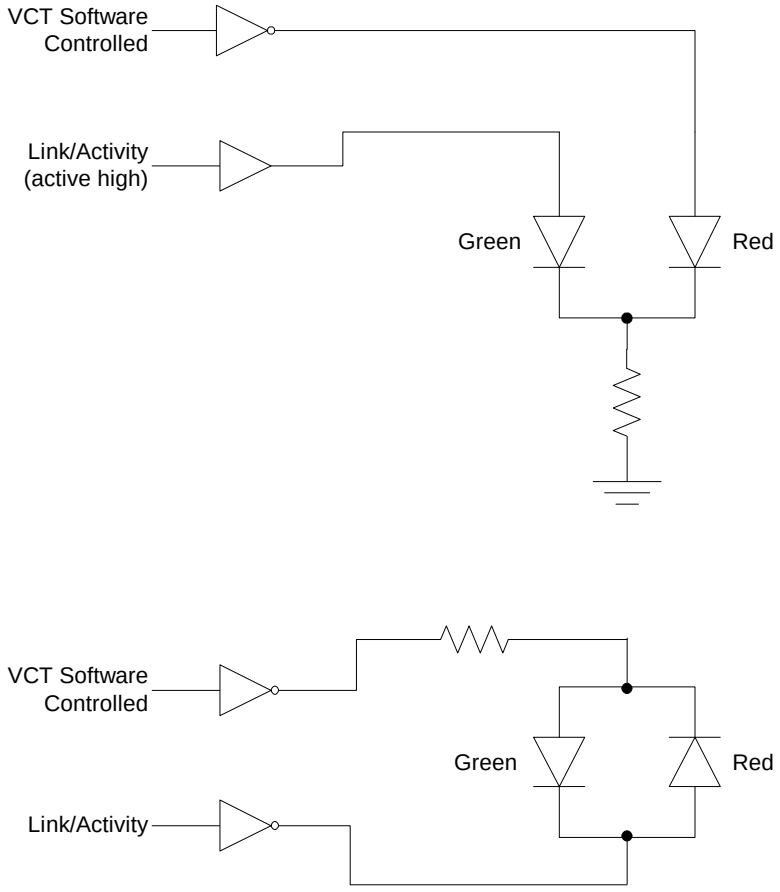

3.9.3 Serial LED Interface

LEDSER, LEDENA, and LEDCLK pins of the devices 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.

3.9.4 Single and Dual LED Modes

3.9.4.1 Single LED Display Modes

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 37.

3.9.4.2 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 37. 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 37 illustrates single and dual LED modes.

Figure 37: Serial LEDENA High Clocking with COLX in Dual Mode, Error Off, and DUPLEX in Single Mode

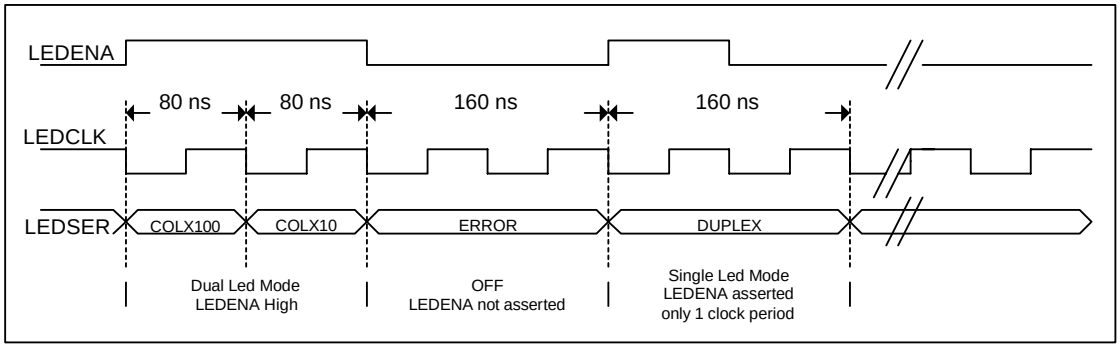

The bit stream on LEDSER can be clocked into a shift register with LEDENA as the shift enable signal as shown in Figure 38. The rate of update of the serial LED interface is controlled by programming register PHY LED Control bits 8:6. The default value is set to 42 ms.

Figure 38: Serial LED Conversion

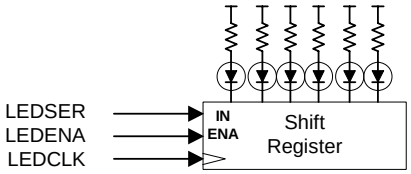

After the LED data is shifted into the correct position, the shift sequence is suspended to allow the appropriate LEDs to light or extinguish depending on status. The LED implementation used in the 88E6065/88E6035 device is self-synchronizing. The default display options are given in Table 46.

Table 46: 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 with LEDCLK. All status signals for Port 5 are sent out first followed by those for Ports 4 through 0. Each bit in the stream occupies a period of 80 ns.

Figure 39: Serial LED Display Order—(if all are selected)

<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>

Table 47 and Table 48 show the status events that can be displayed by programming the 88E6065/88E6035 device in single and dual LED display modes.

Table 47: 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 (blink rate is 84 ms active then 84 ms inactive) The collision activity is pulse stretched to 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 is pulse stretched to 170 ms.
RX	Low = Receive activity High = No receive 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.
LINK/RX	Low = Link up High = Link down Blink = Receive activity (blink rate is 84 ms active then 84 ms 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 84ms active then 84 ms inactive) The transmit and receive activity is pulse stretched to 170 ms.

Table 48: Dual LED Display Mode

Event	Description
COLX100	0 = 100 Mbps collision activity 1 = No 100 Mbps collision activity This status can be pulse stretched.
COLX10	0 = 10 Mbps collision activity 1 = No 10 Mbps collision activity This status can be pulse stretched.
ERROR100	0 = Received error, false carrier, or 100 Mbps FIFO over/underflow occurred. 1 = None of the above occurred This status can be pulse stretched.
ERROR10	0 = Jabber or 10 Mbps FIFO over/underflow occurred 1 = None of the above occurred This status can be pulse stretched.
DUPLEX100	0 = 100 Mbps full-duplex 1 = 100 Mbps half-duplex
DUPLEX10	0 = 10 Mbps full-duplex 1 = 10 Mbps half-duplex
DUPLEX/COLX100	0 = 100 Mbps full-duplex 1 = Half-duplex Blink = 100 Mbps collision activity The collision activity can be pulse stretched. The blink rate can be programmed.
DUPLEX/COLX10	0 = 10 Mbps full-duplex 1 = 10 Mbps half-duplex Blink = 10 Mbps collision activity The collision activity can be pulse stretched. The blink rate can be programmed.
SPEED100	0 = Speed is 100 Mbps 1 = Speed is 10 Mbps
SPEED10	0 = Speed is 10 Mbps 1 = Speed is 100 Mbps
LINK100	0 = 100 Mbps link up 1 = 100 Mbps link down
LINK10	0 = 10 Mbps link up 1 = 10 Mbps link down
TX100	0 = 100 Mbps transmit activity 1 = No 100 Mbps transmit activity This status can be pulse stretched.
TX10	0 = 10 Mbps transmit activity 1 = No 10 Mbps transmit activity This status can be pulse stretched.
RX100	0 = 100 Mbps receive activity 1 = No 100 Mbps receive activity This status can be pulse stretched.

Table 48: Dual LED Display Mode (Continued)

Event	Description
RX10	0 = 10 Mbps receive activity 1 = No 10 Mbps receive activity This status can be pulse stretched.
ACT100	0 = 100 Mbps transmit or 100 Mbps receive activity 1 = No 100 Mbps transmit or 100 Mbps receive activity This status can be pulse stretched.
ACT10	0 = 10 Mbps transmit or 10 Mbps receive activity 1 = No 10 Mbps transmit or 10 Mbps receive activity This status can be pulse stretched.
LINK/RX100	0 = 100 Mbps link up 1 = 100 Mbps link down Blink = 100 Mbps receive activity The receive activity can be pulse stretched. The blink rate can be programmed.
LINK/RX10	0 = 10 Mbps link up 1 = 10 Mbps link down Blink = 10 Mbps receive activity The receive activity is can be pulse stretched The blink rate can be programmed.
LINK/ACT100	0 = 100 Mbps link up 1 = 100 Mbps link down Blink = 100 Mbps transmit or 100 Mbps receive activity The transmit and receive activity can be pulse stretched. The blink rate can be programmed.
LINK/ACT10	0 = 10 Mbps link up 1 = 10 Mbps link down Blink = 10 Mbps transmit or 10 Mbps receive activity The transmit and receive activity can be pulse stretched. The blink rate can be programmed.

Section 4. Serial Management Interface (SMI)

The devices' 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.

4.1 MDC/MDIO Read and Write Operations

All of the relevant serial management registers, as well as several optional registers, are implemented in the 88E6065/88E6035 Switch Core. A description of these registers can be found in Section 5. "Register Description" on page 141.

Note

Access to the 88E6065/88E6035 device's Switch and PHY registers is not possible when the Serial EEPROM is still loading the registers. A CPU can monitor the 88E6065/88E6035 device INTn pin, which will go active (low) when the Serial EEPROM has been fully processed (i.e., a Halt instruction has been reached - see section 5.3).

Figure 40: Typical MDC/MDIO Read Operation

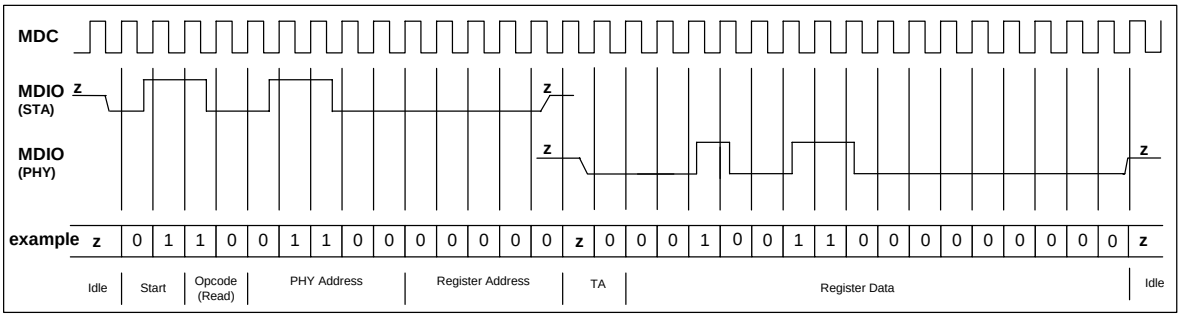

Figure 41: Typical MDC/MDIO Write Operation

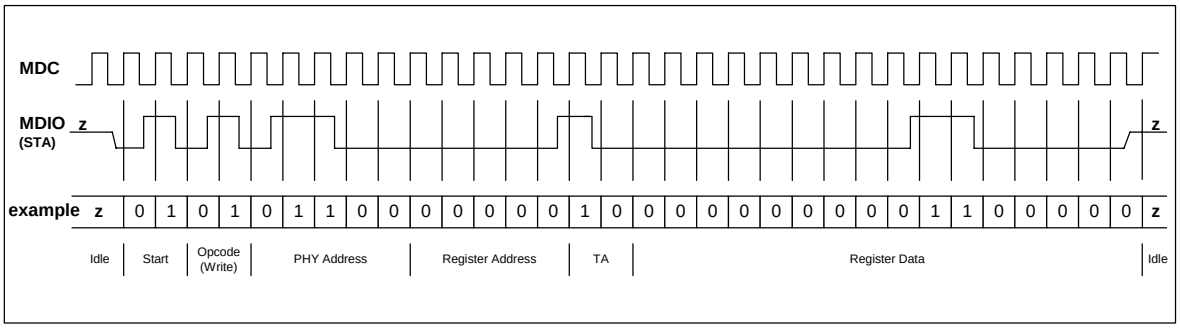

Table 49 shows an example of a read operation of PHY address 04, register 0, with data of 04C0.

Table 49: 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 5. Register Description

The 88E6065/88E6035 device registers are accessible using the IEEE Serial Management Interface (SMI) used for PHY devices (see MDC/MDIO in Table 5). The 88E6065/88E6035 device uses 16 of the 32 possible Device addresses. The 16 device addresses used are configurable at RESETn by use of the EE_CLK/ADDR4 pin (see Table 6). Figure 42 shows the register map assuming the lowest 16 SMI Device Addresses are being used.

Table 50: Register Map—Single-Chip Addressing Mode

Address	Description	Page #
Switch Per-port Registers		
00	"Port Status Register"	page 148
01	"MAC Control Register"	page 151
02	Reserved	
03	"Switch Identifier Register"	page 152
04	"Port Control Register"	page 153
05	Reserved	
06	"Port Base VLAN Map"	page 157
07	"Default Priority Register"	page 159
08	"Port Control 2 Register"	page 160
09	"Ingress Rate Control"	page 165
10	"Egress Rate Control"	page 167
11	"Port Association Vector"	page 168
12-15	Reserved	
16	"RX Counter"	page 170
17	"TX Counter"	page 170
18	"Policy Counter"	page 171
19-23	Reserved	
24	"Port IEEE Priority Remapping Registers"	page 171
25	"Port IEEE Priority Remapping Registers"	page 172
26	"Provider Tag"	page 172
27	"Queue Counter, Device Offset 0x10 to 0x19"	page 173
28–31	Reserved	
Switch Global Registers		
00	"Switch Global Status Registers"	page 175
01	"Switch MAC Address Register Bytes 0 & 1"	page 176
02	"Switch MAC Address Register Bytes 2 & 3"	page 176
03	"Switch MAC Address Register Bytes 4 & 5"	page 176
04	"Switch Global Control Register"	page 177
05	"VTU Operation Register"	page 179

Table 50: Register Map—Single-Chip Addressing Mode (Continued)

Address	Description	Page #
06	"VTU VID Register"	page 181
07	"VTU Data Register Ports 0 to 3"	page 181
08	"VTU Data Register Ports 4 to 5"	page 182
09	Reserved	
10	"ATU Control Register"	page 184
11	"ATU Operation Register"	page 185
12	"ATU Data Register for all but ATU Flush/Move"	page 187
12	"ATU Data Register for ATU Flush/Move"	page 188
13	"ATU MAC Register Bytes 0 & 1"	page 188
14	"ATU MAC Register Bytes 2 & 3"	page 189
15	"ATU MAC Register Bytes 4 & 5"	page 189
16	"IP-QPRI Mapping Register 0"	page 190
17	"IP-QPRI Mapping Register 1"	page 190
18	"IP-QPRI Mapping Register 2"	page 191
19	"IP-QPRI Mapping Register 3"	page 191
20	"IP-QPRI Mapping Register 4"	page 192
21	"IP-QPRI Mapping Register 5"	page 192
22	"IP-QPRI Mapping Register 6"	page 193
23	"IP-QPRI Mapping Register 7"	page 193
24	"IEEE-QPRI Mapping Register"	page 194
25	Reserved	
26	"Management Control"	page 194
27	"Total Free Counter"	page 195
28	"Provider Control"	page 196
29	"Stats Operation Register"	page 196
30	"Stats Counter Register"	page 198
31	"Stats Counter Register"	page 198
Switch Global 2 Registers		
01 - 06	Reserved	
07	"Trunk Mask Table Register"	page 200
08	Reserved	
09	"Ingress Rate Command Register"	page 201
10	"Ingress Rate Data Register"	page 201
Port Ingress Rate Limit Registers		
00 to 11	"Bucket Configuration Memory, IRL Unit"	page 202

Table 51 below defines the register types used in the register map.

Table 51: Register Types

Type	Description
LH	Register field with latching high function. If status is high, then the register 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 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.
Retain	The register value is retained after software reset is executed.
RO	Read only.
ROC	Read only clear. After read, register field is cleared.
RW	Read and Write with initial value indicated.
RWC	Read/Write clear on read. All bits are readable and writable. After reset or after the register field is read, register field is cleared to zero.
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 doesn't take effect until soft reset is executed.
WO	Write only. Reads from this type of register field return undefined data.
NR	Non-Rollover Register

Figure 42: 88E6065 Register Map

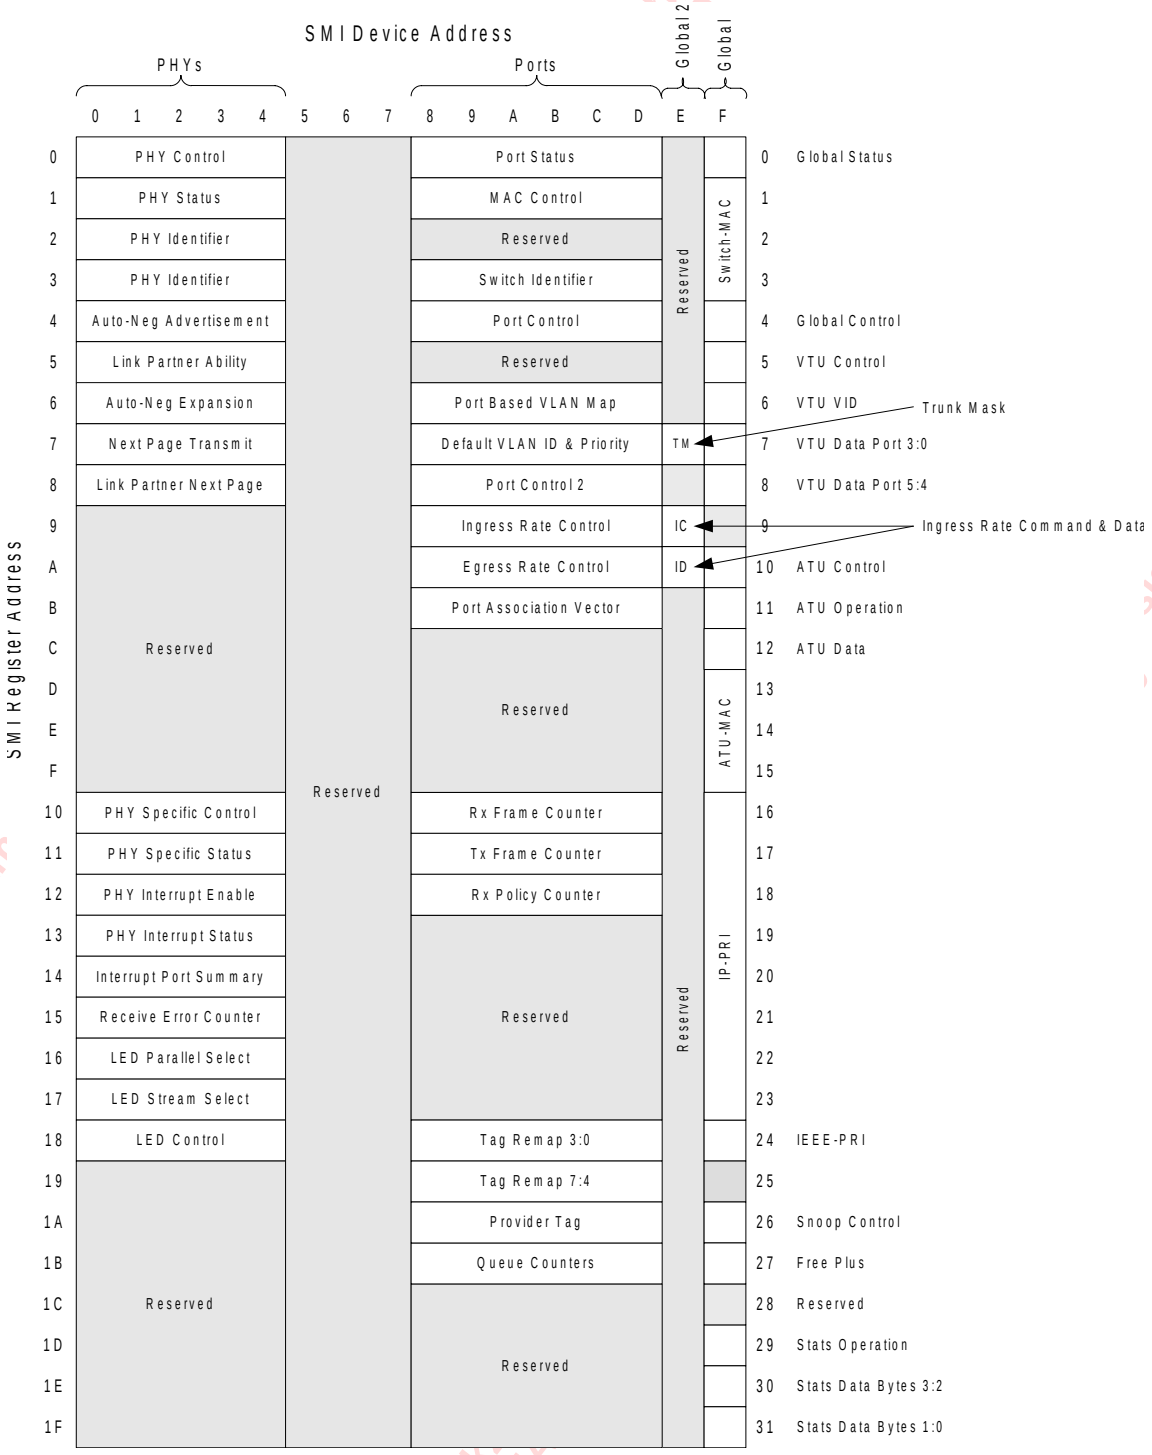

Figure 43: 88E6035 Register Map

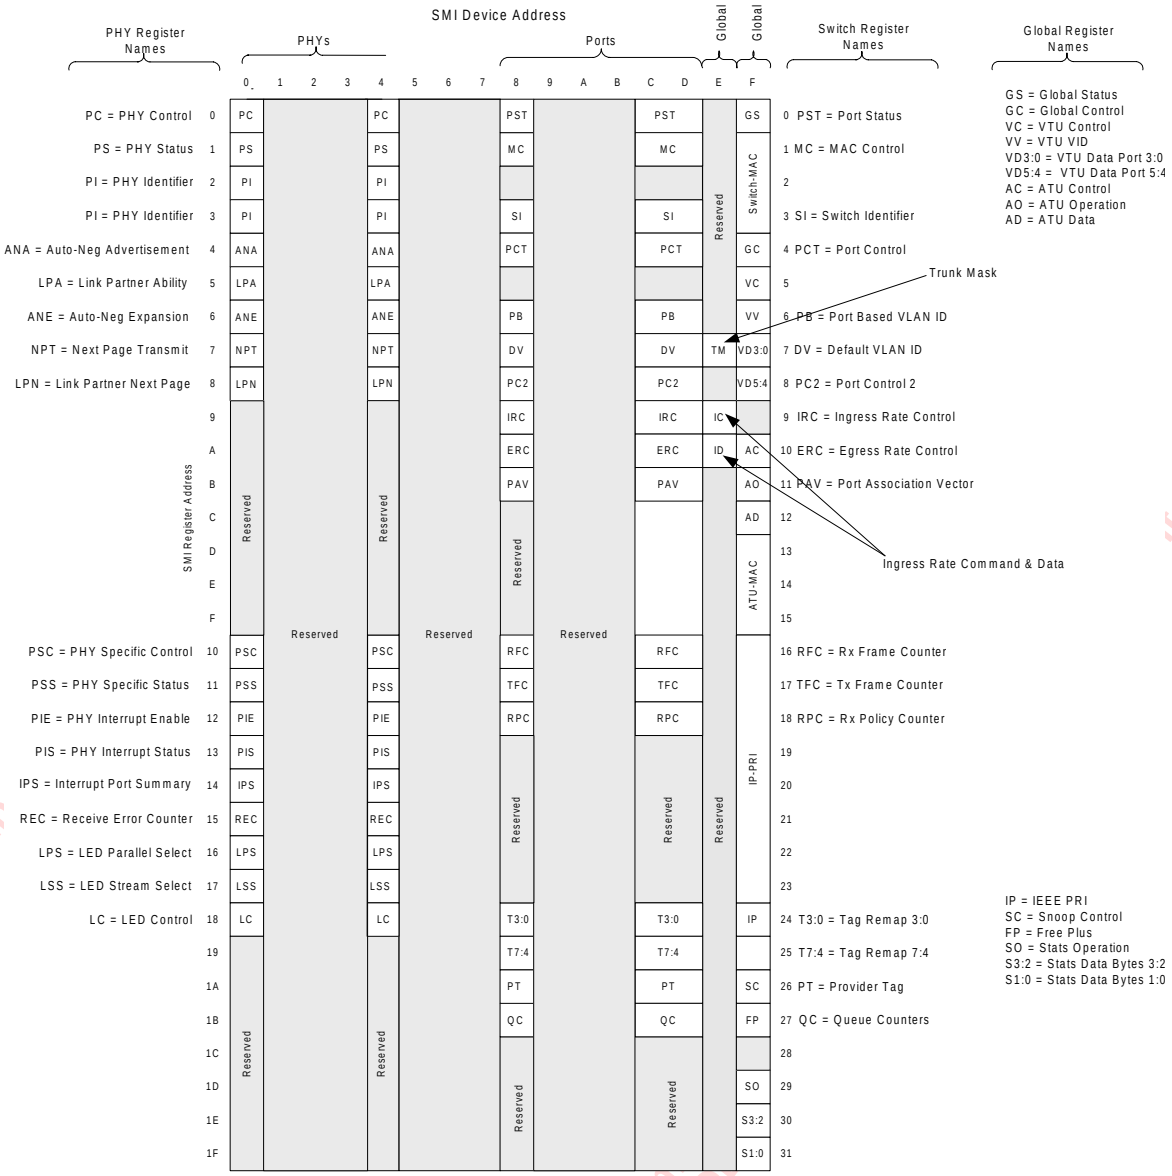

5.1 PHY Registers

The 88E6065/88E6035 device contains five/two physical layer devices (PHYs). In the 88E6065, these devices are accessible using SMI device addresses 0x00 to 0x04 (or 0x10 to 0x14) depending upon the value of the EE_CLK/ADDR4 pin at reset – Table 6. In the 88E6035, these devices are accessible using SMI device addresses 0x00 and 0x04 (or 0x10 and 0x14) depending upon the value of the EE_CLK/ADDR4 pin at reset – Table 6. The PHYs are fully IEEE 802.3 compliant including their register interface.

The PHYs in the 88E6065/88E6035 device are identical to the Marvell® 88E3082 Octal Transceiver except there are only five/two transceivers (transceivers 5 to 7/ 2 to3, 5 to 7 do not exist and are not accessible).

Refer to the Marvell® 88E3082 datasheet for a description on how to program the 88E6065/88E6035 device's PHYs.

5.2 Switch Registers

The 88E6065/88E6035 device contains six/three ports (MACs). In the 88E6065, the supported ports are accessible using SMI device addresses 0x08 to 0x0D (or 0x18 to 0x1D) depending upon the value of the EE_CLK/ADDR4 pin at reset – Table 6. In the 88E6035, the supported ports are accessible using SMI device addresses 0x08, 0x0C to 0x0D (or 0x14, 0x1C to 0x1D) depending upon the value of the EE_CLK/ADDR4 pin at reset – Table 6. The MACs are fully IEEE 802.3 compliant. Since there is no IEEE standard covering required MAC registers, these registers are 88E6065/88E6035 device specific.

The switch contains many global registers that are used to control features and functions that are common to all ports in the switch. The global registers are accessible using two SMI device addresses. Device address 0x0F (or 0x1F depending upon the value of the EE_CLK/ADDR4 pin at reset – Table 6) is referred to as Global or Global 1 space. The registers is this address space are defined starting in 5.2.2 "Switch Global 1 Registers" on page 174. Device address 0x0E (0x1E) is referred to as Global 2 space and its registers are defined starting in 5.2.3 "Switch Global 2 Registers" on page 199.

5.2.1 Switch Port Registers

Each Ethernet port in the 88E6065/88E6035 device contains its own per port registers. Each per port register is 16-bits wide and their bit assignments are shown in Figure 44.

Figure 44: Switch Port Register Map

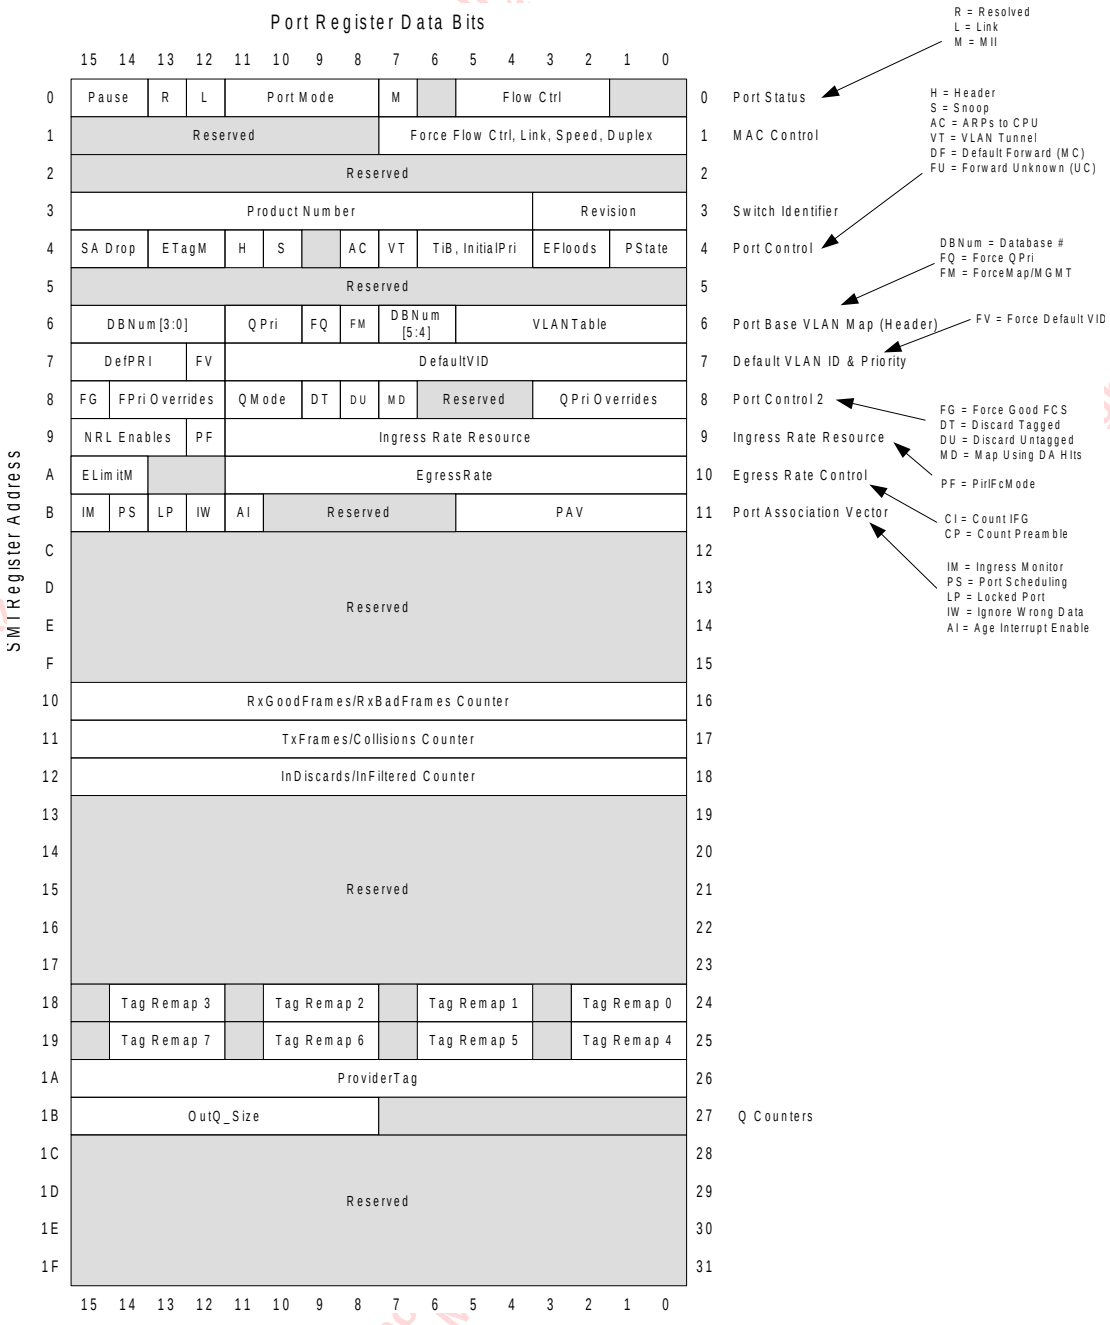

Table 52: Port Status Register
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
15	LinkPause	RO	Link Partner's Pause bit from the PHY, returned from the link partner through Auto-Negotiation. This bit is valid for Ports 0 through 4 only, and when the Resolved bit, below, is set to a one. 0 = MAC Pause not implemented in the link partner 1 = MAC Pause is implemented in the link partner **NOTE:** LinkPause information is directly reflected from the PHY register status even if the ForcedFC (Offset 0x01) bit is set in the MAC control register as ForcedFC (Offset 0x01) bit gives an ability to change the MAC configuration for only flow control enabling/disabling.
14	MyPause	RO	My Pause bit from the PHY, sent to the link partner during Auto-Negotiation. This bit is valid for Ports 0 through 4 only. It is set high if FD_FLOW_DIS (Table 13) is low during RESETn. 0 = MAC Pause is not implemented in the link partner 1 = MAC Pause is implemented in the link partner **NOTE:** MyPause information is directly reflected from the PHY register status even if the ForcedFC (Offset 0x01) bit is set in the MAC control register as ForcedFC (Offset 0x01) bit gives an ability to change the MAC configuration only for flow control enabling/disabling.
13	Resolved	RO	Link Mode is resolved. 0 = Link is undergoing auto-negotiation or the port is disabled 1 = Link has determined its Speed, Duplex and LinkPause settings
12	Link	RO	Link Status in real time (i.e., it is not latched) unless it is being forced up or down by ForcedLink (Offset 0x01). 0 = Link is down 1 = Link is up **NOTE:** The expected way to bring the link up and down is to control through the PHY registers. If ForcedLink (Offset 0x01) bit is set in the MAC control register, this bits' status will no longer reflect the link status coming from the PHY, it will reflect the link value bit instead.

Table 52: Port Status Register (Continued)
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
11:8	PortMode	RO	Port mode. These bits represent the mode the port is in. The meaning of these bits for PHY ports is listed 1st followed by the meaning of these bits for MII ports. **NOTE:** Port 4 can be either a PHY port or an MII port. The MII bit (bit 7 below) returns the port's interface type. For PHY ports: 1000 = 10 Mbps half-duplex 1001 = 100 Mbps half-duplex 1010 = 10 Mbps full-duplex 1011 = 100 Mbps full-duplex All other values are reserved. For MII Ports: 0000 = 10 Mbps half-duplex SNI with rising edge clock & low collision 0001 = 10 Mbps half-duplex SNI with rising edge clock & high collision 0010 = 10 Mbps full-duplex SNI with rising edge clock 0011 = 200 Mbps full-duplex MAC mode MII with 50 MHz input clock 0100 = 10 Mbps half-duplex SNI with falling edge clock & low collision 0101 = 10 Mbps half-duplex SNI with falling edge clock & high collision 0110 = 10 Mbps full-duplex SNI[1] with falling edge clock 0111 = 200 Mbps full-duplex PHY mode MII with 50 MHz output clock 1000 = DC to 25 MHz half-duplex MII MAC mode with input MII clocks 1001 = 10/100[2] half-duplex RMII PHY mode with 50 MHz output clock 1010 = DC to 25 MHz full-duplex MII MAC mode with input MII clocks 1011 = 10/100[2] full-duplex RMII PHY mode with 50 MHz output clock 1100 = 10 Mbps half-duplex MII PHY mode with 2.5 MHz output MII clocks 1101 = 100 Mbps half-duplex MII PHY mode with 25 MHz output MII clocks 1110 = 10 Mbps full-duplex MII PHY mode with 2.5 MHz output MII clocks 1111 = 100 Mbps full-duplex MII PHY mode with 25 MHz output MII clocks For MII ports the PortMode bits reflect the Px_MODE[3:0] configuration pin values latched in at reset (assuming they are not being forced by the MAC Control register's forcing bits, Offset 0x1). **NOTE:** Forcing either the speed or Duplex using MAC control register (Offset 0x01) will change these bits.
7	MII	RO	MII Interface Mode. This bit is set to a one on ports 4 or 5 if the port's MII interface is enabled. This bit is cleared to a zero for all ports using the PHY interface pins.
6	Reserved	RES	Reserved for future use.

Table 52: Port Status Register (Continued)
Offset: 0x00 or Decimal 0

Bits	Field	Type	Description
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 will be used by the Tx MAC (i.e., the transmitter will be paused off). If the port is in half-duplex mode, this bit will never be a one since all Rx Pause frames will be ignored. If the port is in full-duplex mode, this bit will never be 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 being a one indicates that the port is or will be using back pressure. If the port is in full-duplex mode, this bit 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. **NOTE:** This bit gets updated only if flow control is enabled.
3	FdFlowDis	RO	Full-duplex Flow Disable. This bit reads back the real time value that currently appears on the FD_FLOW_DIS pin.
2	HdFlowDis	RO	Half-duplex Flow Disable. This bit reads back the real time value that currently appears on the HD_FLOW_DIS pin.
1:0	Reserved	RES	Reserved for future use.

1. If an SNI full-duplex mode is forced to half-duplex (by ForcedDpx, Offset 0x01) then the low collision mode is chosen.
2. The initial speed of the RMII interface is 100 Mbps but the speed can be overridden to 10 Mbps by forcing the speed to 10 in the MAC Control register - Offset 0x01.

Table 53: MAC Control Register
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
15:8	Reserved	RES	Reserved for future use.
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 will be forced enabled when this bit is set to a one. It will be 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	RWR	Force Flow Control. When this bit is set to a one, flow control (if full-duplex) or backpressure (if half-duplex) for this port will be 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. **NOTE:** This bit setting will change the MAC level flow control setting only (not the PHYs). Setting this bit will not change the Port Status Register (Offset 0x00) bits LinkPause and MyPause as those bits reflect the information from the PHY directly.
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 will be 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 down or up. When this bit is cleared to a zero, normal link detection occurs. This bit forces the MAC to consider Link to be up or down. It does not force Link up or down in the PHY but it will on the MII pins. If software needs to stop the flow of frames on a port it is better to use the Disabled port state in the Port State bits in the Port Control registers (Offset 0x04). **NOTE:** Setting this bit will change the value reflected in the Port Status Register (Offset 0x00) bit "Link".
3	DpxValue	RWR	Duplex's Forced value. This bit is used to force the duplex to full or half when the ForcedDpx bit (below) is set to a one. The duplex will be 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 (200BASE mode does not support half-duplex as undefined results will occur – do not try to force this mode). If an SNI full-duplex port is forced to half duplex the low collision mode will be used. If the ForcedDpx bit (below) is cleared to a zero this bit has no effect.

Table 53: MAC Control Register (Continued)
Offset: 0x01 or Decimal 1

Bits	Field	Type	Description
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. This bit changes the MAC's duplex mode only (not the PHYs). It is intended to be used on MII ports where the initial duplex set on the Px_MODE[3:0] pins needs to be changed after a hardware reset. **NOTE:** Setting this bit will change the value reflected in the Port Status Register (Offset 0x00) bits "PortMode".
1:0	ForceSpd	RWS to 0x03	Force Speed. These bits are used to force the speed on this port as follows: 00 = 10 Mbps 01 = 100 Mbps 10 = 200 Mbps (selects 100 Mbps on Ports 0 to 3) 11 = Speed is not forced. Normal speed detection occurs. The 200 Mbps speed can only be selected on the MII ports and only if they are in a full-duplex MII mode. Selecting this mode in other cases will cause undefined results. If an MII port is configured by the Px_MODE[3:0] pins to use an RMII interface the speed of the port can be forced to 10 or 100 by using this register. **NOTE:** Setting this bit will change the value reflected in the Port Status Register (Offset 0x00) bits "PortMode".

Warning

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

Table 54: Switch Identifier Register
Offset: 0x03 or Decimal 3

Bits	Field	Type	Description
15:4	ProductNum	RO	Product Number or identifier. The 88E6065 is identified by the value 0x065. The 88E6035 is identified by the value 0x035.
3:0	Rev	RO	Revision Identifier. This Rev field may change at any time. Contact the Marvell® FAE's for current information on the device revision identifier.

Table 55: Port Control Register
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
15:14	SA Filtering	RWR	Source Address Filtering controls. These bits select the SA (Source Address) filtering method to be used on the port as follows: 00 = SA Filtering Disabled. No frame will be filtered (i.e., discarded) due to the contents of its Source Address field 01 = Drop On Lock. Ingressing frames will be discarded if their SA field is not in the ATU's address database (i.e., it's 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). Used for MAC based 802.1X. 10 = Drop On Unlock. Ingressing frames will be discarded if their SA field is in the ATU's address database as a Static entry with a PortVec of all zeros. Used to discard frames from known distrusted sources. 11 = Drop to CPU. Ingressing frames will be mapped to the CPUPort (Global 1, Offset 0x09) if their SA field is in the ATU's address database as a Static entry with a PortVec of all zeros. Otherwise, the frames will be discarded if their SA field is not in the ATU's address database (i.e., it's 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). This mode is a form of MAC based 802.1X where some frames can be forced to the CPU for further authentication prior to full authorization.
13:12	Egress TagMode	RWR	Egress Tagging Mode. The first three Egress Tag Modes are used as the default Egress mode for frames whose VID[1] is not contained in the VTU (Switch Port register - Offset 0x05). The 4th Egress Tag Mode is applied to all frames all the time (if selected). 00 = Default to Normal mode – frames are transmitted unmodified 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 = Provider Tag mode (or always add a Tag). This setting is not a default setting. It ignores the MemberTag data in the VTU for this port.
11	Header	RWR	Ingress & Egress Header Mode. When this bit is set to a one all frames Egressing out this port are prepended with the Marvell 2 byte Egress Header just before the frame's DA field. Also, all frames Ingressing into this port are expected to be prepended with the Mavell® 2 byte Ingress Header just before the frame's DA field. On Ingress the first 2 bytes after the SFD are removed from the frame and the frame's CRC and size is recomputed. The frame's Ingress Header is used to update the port's VLAN Map register value (at register Offset 0x06) and if the ForceMap bit is set the frame will be directed out the port or ports defined by the VLANTable (and it will be considered an MGMT frame). When the Header bit is cleared to a zero normal Ethernet frames egress the switch and are expected to ingress the switch.

Table 55: 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[2] overriding the destination port(s) determined by the DA mapping[3]. When this bit is cleared to a zero IGMP/MLD frames are not treated specially.
9	Reserved	RES	Reserved for future use.
8	ARPtoCPU	RWR	ARPs to CPU. When this bit is cleared to a zero, ARP frames are not mapped differently. When this bit is set to a one, ARP frames are also mapped to the port pointed to by the CPUPort register[4] (Global 1, Offset 0x1A). This additional mapping takes place after all VLAN membership masking has occurred (both Port based VLAN masking using the port's VLANTable, Offset 0x06, and 802.1Q based membership masking). This allows ARP frames to get to the CPU even if the CPU is not a member of the port's or frame's VLAN.
7	VLAN Tunnel	RWR	VLAN Tunnel. When this bit is cleared to a zero the port based VLANs defined in the VLANTable (Switch Port register - Offset 0x06) and 802.1Q based VLAN membership defined by the frame's VID stored in the VTU (Switch Port register - Offset 0x05), are enforced for ALL frames. When this bit is set to a one, the port based VLANTable masking & 802.1Q based membership masking is bypassed for any frame entering this port with a DA address that is currently 'static' in the ATU. This includes unicast as well as multicast frames.
6	TagIfBoth	RWS	Use Tag information for the initial QPri assignment if the frame is both tagged and its also IPv4 or IPv6 (and if InitialPri, below, = 0x03, Use Tag & IP Priority). The initial QPri is assigned as follows: 0 = QPri is frame's DiffServ bits (for IPv4) or the frame's Traffic Class bits (for IPv6) mapped using the IP PRI Mapping registers (Global 1 Offset 0x10 to 0x17). 1 = QPri is the determined FPri mapped using the IEEE PRI Mapping register (Global 1, Offset 0x18).

Table 55: Port Control Register (Continued)
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
5:4	InitialPri	RWS to 0x03	Initial Priority assignment. Each frame entering a port is assigned a Frame Priority (FPri) and a Queue Priority (QPri). The FPri determined during ingress is written to the frame's IEEE 802.3ac tag PRI bits if the frame egresses tagged or if the frame egresses Provider Tagged (see EgressTagMode bits in Port Control 1, Offset 0x04)[5]. The QPri is used internally to determine which egress priority queue the frame is mapped into. This register is used to select the frame's initial FPri & QPri depending upon the frame's type and content. This initial FPri & QPri assignments can be overridden by various overrides if enabled on the port (see Port Control 2 register, Offset 0x08). If a frame does not meet the condition listed in the following table the defaults are assigned to frame. The default FPri is the port's DefFPri bit in the Default VLAN ID and Priority register at Offset 0x07. The default QPri is obtained by mapping the frame's FPri value using the IEEE PRI Mapping register (Global 1, Offset 0x18). The initial FPri and QPri are assigned as follows: 00 = Use Port defaults for FPri and QPri. 01 = Use Tag Priority. If the frame is tagged[6], FPri is set to the frame's tag PRI bits re-mapped by the port's Tag Remap registers (Offset 0x17 & 0x18) and the QPri is the determined FPri mapped by the IEEE PRI Mapping register (Global 1, Offset 0x18). If the frame is untagged the port defaults are used for FPri and QPri. 10 = Use IP Priority. If the frame is IPv4 or IPv6, QPri is the frame's Diff-Serv bits (for IPv4) or the frame's Traffic Class bits (for IPv6) mapped by the IP PRI Mapping registers (Global 1, Offset 0x10 to 0x17) and FPri[2:1] is the frame's QPri and FPri[0] is the port's DefFPri[0]. If the frame is not IPv4 nor IPv6 the port defaults are used for FPri and QPri. 11 = Use Tag & IP Priority. If the frame is tagged, FPri is the frame's tag PRI bits re-mapped by the port's Tag Remap registers (Offset 0x17 & 0x18). If the frame is also IPv4 or IPv6 QPri's value will be determined by the TagIfBoth bit above. If the frame is untagged but it is IPv4 or IPv6, FPri and QPri are set according to the Use IP Priority setting above. If the frame is neither tagged nor IPv4 nor IPv6 the port defaults are used for FPri and QPri.

Table 55: Port Control Register (Continued)
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
3:2	Egress Floods	RWS to 0x3	Egress Flooding mode. The DA of every frame, unicast and multicast, is searched in the ATU. If the DA is found in the address database it is considered known. If it not found it is considered unknown. Frames with known DA's are not effected by this register. Frames with unknown DA's generally flood out all the ports (except for the port it originally came in on). This register can be used to prevent frames with unknown DA's from egressing this port as follows: 00 = Do not egress any frame with an unknown DA (unicast or multicast) 01 = Do not egress any frame with an unknown multicast DA 10 = Do not egress any frame with an unknown unicast DA 11 = Egress all frames with an unknown DA (unicast and multicast) The broadcast DA is considered a multicast DA.
1:0	PortState	RWR or RWS to 0x3[7]	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 with the exception of entering the Disabled port state. Disabling a port will cause frame transmission to stop immediately, potentially causing a CRC to be received by the link partner. To prevent any CRC's, move the port to the Blocking port state and wait until the port's egress queue empties (OutQSize, Offset 0x1B) before moving the port to the Disabled port state. 00 = Disabled. The switch port is completely disabled and it will not receive or transmit any frames. The Queue Controller (QC) will return any pre-allocated ingress queue buffers when the port is in this mode. The Disabled port state overrides all other PortStates as no buffers are made available for a Disabled port. 01 = Blocking/Listening. The switch will examine all frames without learning any SA addresses, and discard all but MGMT[8] frames. It will allow MGMT frames only to exit the port. This mode is used for BPDU handling for bridge loop detection and Spanning Tree support. 10 = Learning. The switch will examine all frames, learning all SA addresses (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 will examine all frames, learning all SA addresses (except those from MGMT frames), and receive and transmit all frames like 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. The CPU port is determined by the CPUPort bits in the Global 1 register at Offset 0x09.
3. IGMP/MLD frames get to the CPUPort even if the CPUPort is VLANTable isolated from this port.
4. The CPUPort register must map to a valid physical port number.
5. Frames that egress unmodified do not get modified so the frame's tag PRI bits will not be modified.
6. Priority Only tagged frames (tagged frames with a VID=0x000) are considered tagged for priority decisions.
7. The power-on-reset values for the PortState bits are latched from the SW_MODE[1:0]. If the SW_MODE[1:0] pins are set to 0x0, then the ports are disabled. If SW_MODE[1:0] pins are not set to 0x0, the ports will be in forwarding state. If an EEPROM is present, it can override the PortState values. Note that no traffic is accepted by the switch until the InitReady bit is set (Offset 0x00).
8. MGMT (management) frames are the only kind of frame that can be tunneled through Blocked ports. A MGMT frame is any frame whose DA address appears in the ATU Database with the MGMT Entry State.

Table 56: Port Base VLAN Map[1]
Offset: 0x06 or Decimal 6

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 should be zero if it is not being used. It needs to be a unique number for each independent, non-overlapping, VLAN if used. The upper 2 bits of the port's default DBNum are contained in bits 7:6 of this register.
11:10	QPriValue	RWR	Queue priority value to use when forced. When the ForceQPri bit below is set to a one, all frames entering this port are mapped to the priority queue defined in these bits, overriding the frame's initial QPri, unless a VTU, SA, DA, or ARP priority override occurs (see Port Control 2, Offset 0x08). This is an internal queue priority override only. The frame's priority (FPri) is not affected by these bits.
9	ForceQPri	RWR	Force Queue Priority. When this bit is cleared to zero, normal priority queue mapping is used on all ingressing frames entering this port. When this bit is set to a one all frames ingressing this port are mapped to the priority queue defined in the QPri bits above, overriding the frame's initial QPri, unless a VTU, SA, DA or ARP priority override occurs (see Port Control 2, Offset 0x08). This is an internal queue priority override only. The frame's priority (FPri) is not affected by these bits.
8	ForceMap	RWR	Force Mapping. When this bit is cleared to a zero, normal frame processing occurs. When this bit is set to a one, all received frames will be considered MGMT[2] and they are mapped to the port or ports defined in the VLANTable bits below overriding the mapping from the address database. The forcing function is needed to get BPDU frames to egress specific ports by the CPU for the Spanning Tree Protocol. ForceMapped frames will egress ports that are not in the Disabled port state (i.e., they are MGMT frames and will egress out Blocked ports). This bit is accessible by the CPU's Ingress Header so the CPU can enable and disable MGMT and forcing on a frame by frame basis.
7:6	DBNum[5:4]	RWR	Port's Default Database VLAN Number bits 5:4. 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 should be zero if it is not being used. It needs to be a unique number for each independent, non-overlapping, VLAN if used. The lower 4 bits of the port's default DBNum are contained in bits 15:12 of this register.

Table 56: Port Base VLAN Map[1] (Continued)
Offset: 0x06 or Decimal 6

Bits	Field	Type	Description
5:0	VLANTable	RWS to all ones except for this port's bit	Port based VLAN Table & Forced Mapping bits. When the ForceMap bit above is cleared to a zero, 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 non-snooped frames. In this mode, these bits restrict where a port can send frames to (unless a VLANTunnel frame is being received – Switch Port register - Offset 0x04). If ForceMap bit is set to a one, these bits indicate which port or ports all frames that ingress into this port are sent to overriding the mapping from the address database. 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 and the ForceMap bit is zero. This Port's bit will always be zero after Reset. This prevents frames going out the port they came in on. The 88E6065/88E6035 allows this Port's bit to be set to a one allowing frames to be switched back to the port they came in on so software needs to be careful with these bits. This register is normally reset to 0x3E for Port0 (SMI Device Address 0x08 or 0x18), and it resets to 0x3D for Port1 (Addr 0x09 or 0x19), to 0x3B for Port2 (Addr 0x0A or 0x1A).

1. The contents of this register can be modified on a frame by frame basis if the port's Header Mode is enabled (section 2.6.19).
2. MGMT = Management frames, frames that can tunnel through Blocked ports (see PortStates in Port Control 1, Offset 0x04).

Table 57: Default Priority Register
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15:13	DefPri	RWR	Default Frame Priority. The bits of this register are used as the default frame priority (FPri) to use when no other priority information is available (either the frame is not IEEE Tagged, nor is it an IPv4 nor an IPv6 frame – or the frame is a priority type that is currently disabled (see InitialPri, Offset 0x04). The initial FPri assignment can be overridden by various overrides if enabled on the port (see Port Control 2 register, Offset 0x08).
12	Force DefaultVID	RWR	Force to use Default VID. When this bit is set to a one, all frames received with an IEEE 802.3ac Tag have the Tag's VID ignored and the DefaultVID below is assigned to the frame instead. This can be used to force all frames entering a port to use a known VID even if the frames are already tagged. When this bit is cleared to a zero, all IEEE 802.3ac Tagged frames with a non-zero VID have the frame's VID assigned to the frame unmodified. The VID that is assigned to the frame during ingress is written to the frame's IEEE 802.3ac tag VID bits if the frame egresses tagged[1]. If the frame egresses Provider Tagged (see EgressTagMode bits in Port Control 1, Offset 0x04) it will use the VID assigned to the frame during ingress if the frame came from a Provider Port (i.e., the frame's EtherType equaled the port's Provider Tag register, port – Offset 0x1A), otherwise it will use the DefaultVID of the source port.
11:0	DefaultVID	RWS to 0x001	Default VLAN Identifier. The DefaultVID field is used as the VID assigned to untagged frames received on 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. The VID that is assigned to the frame during ingress is written to the frame's IEEE 802.3ac tag VID bits if the frame egresses tagged[1]. If the frame egresses Provider Tagged (see EgressTagMode bits in Port Control 1, Offset 0x04) it will use the VID assigned to the frame during ingress if the frame came from a Provider Port (i.e., the frame's EtherType equaled the port's Provider Tag register, port – Offset 0x1A), otherwise it will use the DefaultVID of the source port.

1. Frames that egress unmodified do not get modified so the frame's tag VID bits will not be modified.

Table 58: Port Control 2 Register
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
15	ForceGood FCS	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 they will be 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 will be processed by the rest of the switch policy decisions.
14	VID FPri Override	RWR	VID Frame Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one, then VID frame priority overrides can occur on this port. A VID frame priority override occurs when the determined VID of a frame[1] results in a VID whose UseVIDFPri override bit in the VLAN database is set to a one. When this occurs the VIDFPri value assigned to the frame's VID (in the VLAN database) is used to overwrite the frame's previously determined frame priority. If the frame egresses tagged[2], the priority in the frame will be this new VIDFPri value. This function does not effect the egress queue priority (QPri) the frame is switched into (see the VID QPri Override, bit 3, below). The VID Frame Priority Override has higher priority than the port's Default Priority and the frame's IEEE and/or IP priorities. The FPri determined by the frame's VID can be overridden, however, by the frame's SA and/or DA Frame Priority Overrides (see bits 13:12 below) or if the frame is identified as a MGMT frame.
13	SA FPri Override	RWR	SA Frame 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 frame priority overrides can occur on this port. An SA ATU frame priority override occurs when the source address of a frame results in an ATU hit where the SA's MAC address entry contains the UseATUFPri bit = 1. When this occurs the ATUFPri value assigned to the frame's SA (in the ATU database) is used to overwrite the frame's previously determined frame priority. If the frame egresses tagged[2] the priority in the frame will be this new ATUFPri value. This function does not effect the egress queue priority (QPri) the frame is switched into (see the SA QPri Override, bit 2, below). The SA Frame Priority Override has higher priority than the port's Default Priority, and the frame's IEEE and/or IP priorities, and the VID Frame Priority Override. The priority determined by the frame's SA can be overridden, however, by the frame's DA Frame Priority Override (see bit 12 below) or if the frame is identified as a MGMT frame.

Table 58: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
12	DA FPri Override	RWR	DA Frame 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 frame priority overrides can occur on this port. A DA ATU frame priority override occurs when the destination address of a frame results in an ATU hit where the DA's MAC address entry contains the UseATUFPri bit = 1. When this occurs the ATUFPri value assigned to the frame's DA (in the ATU database) is used to overwrite the frame's previously determined frame priority. If the frame egresses tagged, the priority in the frame will be this new ATUFPri value. This function does not effect the egress queue priority (QPri) the frame is switched into (see the SA QPri Override, bit 2, below). The DA ATU Frame Priority Override has highest priority over the port's Default Priority, the frame's IEEE and/or IP priorities, the VID Frame Priority Override, and the SA Frame Priority Override. Note: If a frame's DA is contained in the ATU with a MGMT Entry State, the frame's FPri will be overridden regardless of the state of this bit (as long as the DA's ATU entry has the UseATUFPri bits set to a 1).
11:10	802.1Q Mode	RWR	IEEE 802.1Q Mode for this port. These bits determine if 802.1Q based VLANs are used along with port based VLANs for this Ingress port (port based VLANs defined by the VLANTable, Offset 0x06, are always used). It also determines the action to be taken if an 802.1Q VLAN Violation is detected. VLAN barriers (both port based and 802.1Q based) can be bypassed by VLANTunnel (Offset 0x04) and MGMT (section 2.5.8) frames. These bits work as follows: 00 = 802.1Q Disabled. Use Port Based VLANs only. The VLANTable bits (Offset 0x06) determine which Egress ports this Ingress port is allowed to switch frames to for all frames. VID violations are not generated for this port. 01 = Fallback. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violations and use the VLANTable bits below if the VID assigned to the frame during ingress is not contained in the VTU (both errors are logged – Switch Port register - Offset 0x05). 10 = Check. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violation but discard the frame if the VID assigned to it during ingress is not contained in the VTU (both errors are logged – Switch Port register - Offset 0x05). 11 = Secure. Enable 802.1Q for this Ingress port. Discard Ingress Membership violations and discard frames whose VID assigned to it during ingress is not contained in the VTU (both errors are logged – Switch Port register - Offset 0x05).

Table 58: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
9	Discard Tagged	RWR	Discard Tagged Frames. When this bit is set to a one, all non-MGMT frames that are processed as tagged will be discarded as they enter this switch port. Priority only tagged frames (those that Ingressed with a VID of 0x000 even if the VID assigned to the frame is non-zero) are considered untagged.
8	Discard Untagged	RWR	Discard Untagged Frames. When this bit is set to a one, all non-MGMT frames that are processed as untagged will be discarded as they enter this switch port. Priority only tagged frames (those that ingressed with a VID of 0x000 even if the VID assigned to the frame is non-zero) are considered untagged.
7	MapDA	RWS	Map using DA hits. When this bit is set to a one, normal switch operation will occur where a frame's DA address is used to direct the frame out the correct port(s). When this bit is cleared to a zero, the frame will be sent out the port(s) defined by the EgressFloods bits (Switch Port register - Offset 0x04) 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).
6:4	Reserved	RES	Reserved for future use.
3	VID QPri Override	RWR	VID Queue Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one then VID queue priority overrides can occur on this port. A VID queue priority override occurs when the determined VID of a frame[3] results in a VID whose UseVIDQPri override bit in the VLAN database is set to a one. When this occurs the VIDQPri value assigned to the frame's VID (in the VLAN database) is used to overwrite the frame's previously determined queue priority. If the frame egresses tagged the priority in the frame will not be modified by this new VIDQPri value (see the VID FPri Override, bit 14, above). This function effects the egress queue priority (QPri) the frame is switched into. The VID Priority Queue Override has higher priority than the initial QPri determined by the port's Default Priority, the frame's IEEE and/or IP priorities and the ForceQPri (Offset 0x06). The QPri determined by the frame's VID can be overridden, however, by the frame's SA, DA or ARP Queue Priority Overrides (see bits 2:0 below).

Table 58: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
2	SA QPri Override	RWR	SA Queue 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 queue priority overrides can occur on this port. An SA ATU queue 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 Queue Priority Override[4]. When this occurs the QPRI value assigned to the frame's SA (in the ATU database) is used to overwrite the frame's previously determined egress queue priority. If the frame egresses tagged, the priority in the frame will not be modified. When used, the two bits of the QPRI priority determine the egress queue the frame is switched into. The SA ATU Queue Priority Override has higher priority than the initial QPri determined by the port's Default Priority, the frame's IEEE and/or IP priorities, the ForceQPri (Offset 0x06) and the VID Queue Priority Override. The QPri determined by the frame's SA can be overridden, however, by the frame's DA and ARP Queue Priority Override (see bits 1:0 below).
1	DA QPri Override	RWR	DA Queue 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 queue priority overrides can occur on this port. A DA ATU queue 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 Queue Priority Override. When this occurs the QPRI value assigned to the frame's DA (in the ATU database) is used to overwrite the frame's previously determined egress queue priority. If the frame egresses tagged the priority in the frame will not be modified. When used the two bits of the QPRI priority determine the egress queue the frame is switched into. The DA ATU Queue Priority Override has higher priority than the initial QPri determined by the port's Default Priority, the frame's IEEE and/or IP priorities, the ForceQPri (Offset 0x06) and the VID and SA Queue Priority Overrides. The QPri determined by the frame's DA can be overridden, however, by the frame's ARP Queue Priority Override (see bit 0 below). **NOTE:** If a frame's DA is contained in the ATU with a MGMT Entry State the frame's QPri will be overridden regardless of the state of this bit.

Table 58: Port Control 2 Register (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
0	ARP QPri Override	RWR	ARP Queue Priority Override. When this bit is cleared to a zero, normal frame priority processing occurs. When this bit is set to a one, then ARP queue priority overrides can occur on this port. An ARP queue priority override occurs for all ARP frames. When this occurs the frame's previously determined egress queue priority will be overwritten with ArpQPri (Global 1, Offset 0x1A). If the frame egresses tagged the priority in the frame will not be modified. When used, the two bits of the ArpQPri priority determine the egress queue the frame is switched into. The ARP Queue Priority Override has highest priority over the frame's initial QPri determined by the port's Default Priority, the frame's IEEE and/or IP priorities, and the ForceQPri (Offset 0x06), VID, SA and DA Queue Priority Overrides.

1. The VID of a frame could be a tagged frame's VID or the port's DefaultVID.
2. Frames that egress unmodified do not get modified so the frame's tag PRI bits will not be modified.
3. The VID of a frame could be a tagged frame's VID or the port's DefaultVID.
4. SA Queue Priority Override can only occur on MAC addresses that are Static or where the Port is Locked, and where the port is the mapped source port for the MAC address.

Table 59: Ingress Rate Control
Offset: 0x09 or Decimal 9

Bits	Field	Type	Description
15	VidNrlEn	RWR	VID Non Rate Limit enable. When this bit is cleared to a zero, normal frame processing occurs. When this bit is set to a one, then VID non rate limiting overrides can occur on this port. A VID non rate limiting override occurs when the determined VID of a frame[1] results in a VID whose VIDNonRateLimit bit in the VLAN database is set to a one. When this occurs the frame will not be ingress nor egress rate limited.
14	SaNrlEn	RWR	SA Non Rate Limit enable. When this bit is cleared to a zero, normal frame processing occurs. When this bit is set to a one, then SA ATU non rate limiting overrides can occur on this port. An SA ATU non rate limiting 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 Non Rate Limited[2]. When this occurs the frame will not be ingress nor egress rate limited.
13	DaNrlEn	RWR	DA Non Rate Limit enable. When this bit is cleared to a zero, normal frame processing occurs. When this bit is set to a one, DA ATU non rate limiting overrides can occur on this port. A DA ATU non rate limiting 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 Non Rate Limited. When this occurs the frame will not be ingress nor egress rate limited. **NOTE:** If the frame's DA is contained in the ATU with a MGMT Entry State the frame will be non rate limited based upon the ATU entry's EntryState value regardless of the state of this bit.
12	PirlFc Mode	RWR	Port Ingress Rate Limit Flow Control Mode. This bit is used to determine when Flow Control (asserted due to ingress rate limiting threshold exceeding reasons) gets deasserted. 0 = De-Assert flow control when ingress rate resource has become empty i.e., the ingress rate resource can accept more packets as it is empty 1 = De-assert flow control when ingress rate resource has enough room to accept at least one packet of size determined by the value programmed in the CBS_Limit field as specified in Switch Global register - Offset 0x00 to 0x0B. For example, if the CBS_Limit for the ingress rate resource is programmed to be 2 Kbytes, then the flow control will get de-asserted if there is at least 2 Kbytes worth of room available in the ingress rate resource.

Table 59: Ingress Rate Control (Continued)
Offset: 0x09 or Decimal 9

Bits	Field	Type	Description
11:0	Ingress Rate Resource	RWR	Each bit in this field identifies an ingress rate resource in the device assigned to this port. If all the bits of this field are zeros then none of the ingress rate resources are assigned to this port thus this port's ingress traffic would not get rate limited. More than one rate resource can be assigned to a given port. For example if TCP/IP rate limiting and limiting of broadcasts needs to be enabled for a given port, then two rate resources can be assigned to this port by setting the resource bucket's bit in this field. If rate resources 0 and 1 need to be assigned to this port, then bits zero and one of this field need to be set to a one. Multiple ports can share the rate limiting resources. For example if ports 3 & 5 need to share a rate resource bucket 2, then bit2 of this field need to enabled for both port 3 and port 5. This feature is especially useful for port link aggregation cases where multiple ports form logically one port and thus rate limiting needs to be controlled across multiple ports.

1. The VID of a frame could be a tagged frame's VID or the port's DefaultVID.
2. SA Non Rate Limiting Override can only occur on MAC addresses that are Static or where the Port is Locked, and where the port is the mapped source port for the MAC address.

Table 60: Egress Rate Control
Offset: 0x0A or Decimal 10

Bits	Field	Type	Description
15:14	ELimit Mode	RWS to 0x02	Egress limit counting mode. These bits are used to indicate which bytes in the transmitted frames are counted for egress rate limiting as follows: 00 = Reserved 01 = Count all Layer 1 bytes 10 = Count all Layer 2 bytes 11 = Count all Layer 3 bytes Layer 1 = Preamble (8 bytes) + Frame's DA to CRC + IFG (12 bytes) Layer 2 = Frames's DA to CRC Layer 3 = Frames's DA to CRC – 18[1] – 4 (if the frame is tagged[2]) A frame is considered tagged if it was either Customer tagged or Provider tagged during ingress.
13:12	Reserved	RES	Reserved for future use.
11:0	Egress Rate	RWR	Egress data rate limit. The EgressRate bits modify this port's effective transmission rate (i.e., rate shaping). When this register is cleared to zero, egress rate shaping is disabled. The 88E6065/88E6035 uses the following formula to limit the Egress data rate: EgressRate = 8 bits/(40 ns * Egress Rate bits/sec) Where Egress Rate is in bits per second. For example for 256 kbps: Pri0Rate = 8 bits/(40 ns * 256000 bits/sec) = Pri0Rate = 8 bits/(0.01024 bits) <- (0.000040 * 256) = 781 or 0x30D <- The value programmed into this register. This register supports a range of 200 Mbps to 50 kbps but a port will not be able to go faster than its current speed. Egress Rate Shaping transmits a frame at wire speed counting the transmitted bytes determined in ELimitMode above. The value in this register determines the time it takes for the transmitted byte count to reach zero. When it reaches zero, the next frame is allowed to be transmitted and process repeats. This burstless rate shaping is the best method for supporting the minimal amount of buffering required in the link partner this device is connected to.

1. The 18 bytes are: 6 for DA, 6 for SA, 2 for EtherType and 4 for CRC.
2. Only one tag is counted even if the frame contains more that one tag (i.e. it is Provider Tagged).

Table 61: Port Association Vector
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
15	Ingress Monitor	RWR	Ingress Monitor enable. When this bit is set to a one Ingress port monitoring is enabled on the same ports that Egress port monitoring is enabled on as defined by the PAV bits below. When this bit is set to a one the PAV bits below are used to modify the DPV bits returned to a port on a DA hit. 1st the port's PAV bits below are modified internally by masking out this port's bits (forcing this port's bit to a zero). Then the modified PAV bits from below are OR'ed together with the ATU's DPV bits read from memory on an ATU hit for this port.
14	PortSched	RWR	Port Scheduling mode. When the UsePortSched bit is set (in Global Control, Global 1, Offset 0x04), this bit is used to select the Queue Controller's scheduling mode on this port as follows: 0 = Use an 8, 4, 2, 1 weighted fair queuing scheme 1 = Use a strict priority scheme
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 (for 802.1X MAC authentication) is enabled on this port. In this mode, an ATU Miss Violation interrupt will occur when a new SA address is received in a frame on this port. Automatic SA learning and refreshing is disabled for this port 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 2.5.10.1), or if the port is 'locked' (see bit 13 above), the source port's bit is checked to insure 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 will be masked and ignored.
11	AgeIntEn	RWR	Age Interrupt Enable. When this bit is set to a one, ATU Age Violation interrupts from this port are enabled. An Age Violation will occur anytime a port is Locked (bit 13 above is a one) and the ingressing frame's SA is contained in the ATU as a non-Static entry with a EntryState less than 0x04. This interrupt can be used as part of CPU directed learning to inform the CPU MAC addresses are still being used and need to be refreshed by the CPU.
10:6	Reserved	RES	Reserved for future use.

Table 61: Port Association Vector (Continued)
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
5:0	PAV	RWS to all zeroes except for this port's bit	Port Association Vector for ATU learning. The value in these bits are used as the port's DPV on automatic ATU Learning or EntryState refresh whenever these bits contain a non-zero value. When these bits are all zero, automatic Learning and EntryState 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 – Switch Port register - Offset 0x04). The PAV bits can be used to set up an Egress port monitor. To configure Port 0 to get a copy of the frames that exit Port 1, set bit 0 in Port 1's PAV register (along with Port 1's bit). To configure Port 0 to get a copy of the frames that exit Port 2, set bit 0 in Port 2's PAV register, etc. The PAV bits can be used to set up an Ingress port monitor along with an Egress port monitor. To configure Port 0 to get a copy of the frames that enter Port 1 as well as exit Port 1, set bit 0 in Port 1's PAV register (along with Port 1's bit – bit 1) and set the IngressMonitor bit above. The PAV bits can be used to set up port trunking (along with the VLANTable bits (Switch Port register - Offset 0x06). 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 or the Trunk Mask Table (Global 2, Offset 0x07) 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. The PAV bits can be used to disable learning on a port by clearing them to 0x00 for that port.

Table 62: RX Counter
Offset: 0x10 or Decimal 16

Bits	Field	Type	Description
15:0	RxGood Frames/ RxBad Frames	RO	Receive Counter. When CtrMode is cleared to a zero (Global Control – Switch Global register - Offset 0x04), this counter increments each time a good frame enters this port. It does not matter if the frame is switched or discarded. When CtrMode is set to a one, this counter increments each time an error frame enters this port. An error frame is one that is 64 bytes or greater with a bad CRC (including alignment errors but not dribbles). Fragments and properly formed frames are not counted. The counter will wrap around back to zero. The only time this counter will not increment is when this port is Disabled (see PortState - Offset 0x04). This register can be cleared by changing the state of the CtrMode bit in Global Control (Switch Port register - Offset 0x04).

Table 63: TX Counter
Offset: 0x11 or Decimal 17

Bits	Field	Type	Description
15:0	TxFrames/ Collisions	RO	Transmit Counter. When CtrMode is cleared to a zero (Global Control – Switch Port register - Offset 0x04), this counter increments each time a frame successfully exits this port. When CtrMode is set to a one, this counter increments each time a collision occurs during an attempted transmission. It no longer counts all transmitted frames – but only those transmission attempts that resulted in a collision. The counter will wrap around back to zero. The only time this counter will not increment is when this port is Disabled (see PortState - Offset 0x04). This register can be cleared by changing the state of the CtrMode bit in Global Control (Switch Global register - Offset 0x04).

Table 64: Policy Counter
Offset: 0x12 or Decimal 18

Bits	Field	Type	Description
15:0	InDiscards/ InFiltered	RO	Policy Counter. When CtrMode is cleared to a zero (Global Control – Switch Global register - Offset 0x04), this counter increments each time a good, non-filtered, frame is discarded due to a lack of buffer memory. This counter also gets incremented if a good frame gets discarded because it arrived with less than minimum inter-packet gap. When CtrMode is set to a one, this counter increments each time a good frame is discarded due to filtering rules. The counter will wrap around back to zero. The only time this counter will not increment is when this port is Disabled (see PortState - Offset 0x04). This register can be cleared by changing the state of the CtrMode bit in Global Control (Switch Global register - Offset 0x04).

Table 65: Port IEEE Priority Remapping Registers
Offset: 0x18 or Decimal 24

Bits	Field	Type	Description
15	Reserved	RES	Reserved for future use.
14:12	TagRemap3	RWS to 0x03	Tag Remap 3. All tagged frames with a priority of 3 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
11	Reserved	RES	Reserved for future use.
10:8	TagRemap2	RWS to 0x02	Tag Remap 2. All tagged frames with a priority of 2 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
7	Reserved	RES	Reserved for future use.
6:4	TagRemap1	RWS to 0x01	Tag Remap 1. All tagged frames with a priority of 1 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
3	Reserved	RES	Reserved for future use.
2:0	TagRemap0	RWR	Tag Remap 0. All tagged frames with a priority of 0 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.

Table 66: Port IEEE Priority Remapping Registers
Offset: 0x19 or Decimal 25

Bits	Field	Type	Description
15	Reserved	RES	Reserved for future use.
14:12	TagRemap7	RWS to 0x07	Tag Remap 7. All tagged frames with a priority of 7 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
11	Reserved	RES	Reserved for future use.
10:8	TagRemap6	RWS to 0x06	Tag Remap 6. All tagged frames with a priority of 6 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
7	Reserved	RES	Reserved for future use.
6:4	TagRemap5	RWS to 0x05	Tag Remap 5. All tagged frames with a priority of 5 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.
3	Reserved	RES	Reserved for future use.
2:0	TagRemap4	RWS to 0x04	Tag Remap 4. All tagged frames with a priority of 4 get this register's value as the frame's new FPri (if InitialPri[0] = 1). If a tagged frame egresses the switch tagged this new priority will be written to the frame's tag.

Table 67: Provider Tag
Offset: 0x1A or Decimal 26

Bits	Field	Type	Description
15:0	ProviderTag	RWS to 0x9100	Provider Tag. These bits indicate the EtherType value that needs to be matched to ingress to determine if a frame is Provider Tagged or not. If incoming tag matches this register value and the EgressTagMode (Port control register - Offset 0x04) is 0x3 then the frame is considered to be provider tagged. **NOTE:** If the programmed Provider Tag value is 0x8100 then the first 4 bytes of provider tag information is stripped for all provider tagged packets. If this register is programmed to a value other than 0x8100, multiple provider tags get stripped as long as RecStrDis (Switch Global registers- Offset 0x1C) is 0x0. If the Egress Tag Mode (Port control register - Offset 0x04) is 0x3, this register's value is inserted as a tag (EtherType) for outgoing packets.

Table 68: Queue Counter, Device Offset 0x10 to 0x19
Offset: 0x1B or Decimal 27

Bits	Field	Type	Description
15:8	OutQ_Size	RO	Egress Queue Size Counter. This counter reflects the current number of Egress buffers switched to this port. This is the total number of buffers across all four priority queues.
7:0	Reserved	RES	Reserved for future use.

5.2.2 Switch Global 1 Registers

The 88E6065/88E6035 device contains global registers that effect all the Ethernet ports in the device. This section documents the Global 1 registers. Global 2 registers are documented in Section 5.2.3 "Switch Global 2 Registers" on page 199. Each global register is 16-bit wide and their bit assignments are shown in Figure 45.

Figure 45: Switch Global 1 Register Map

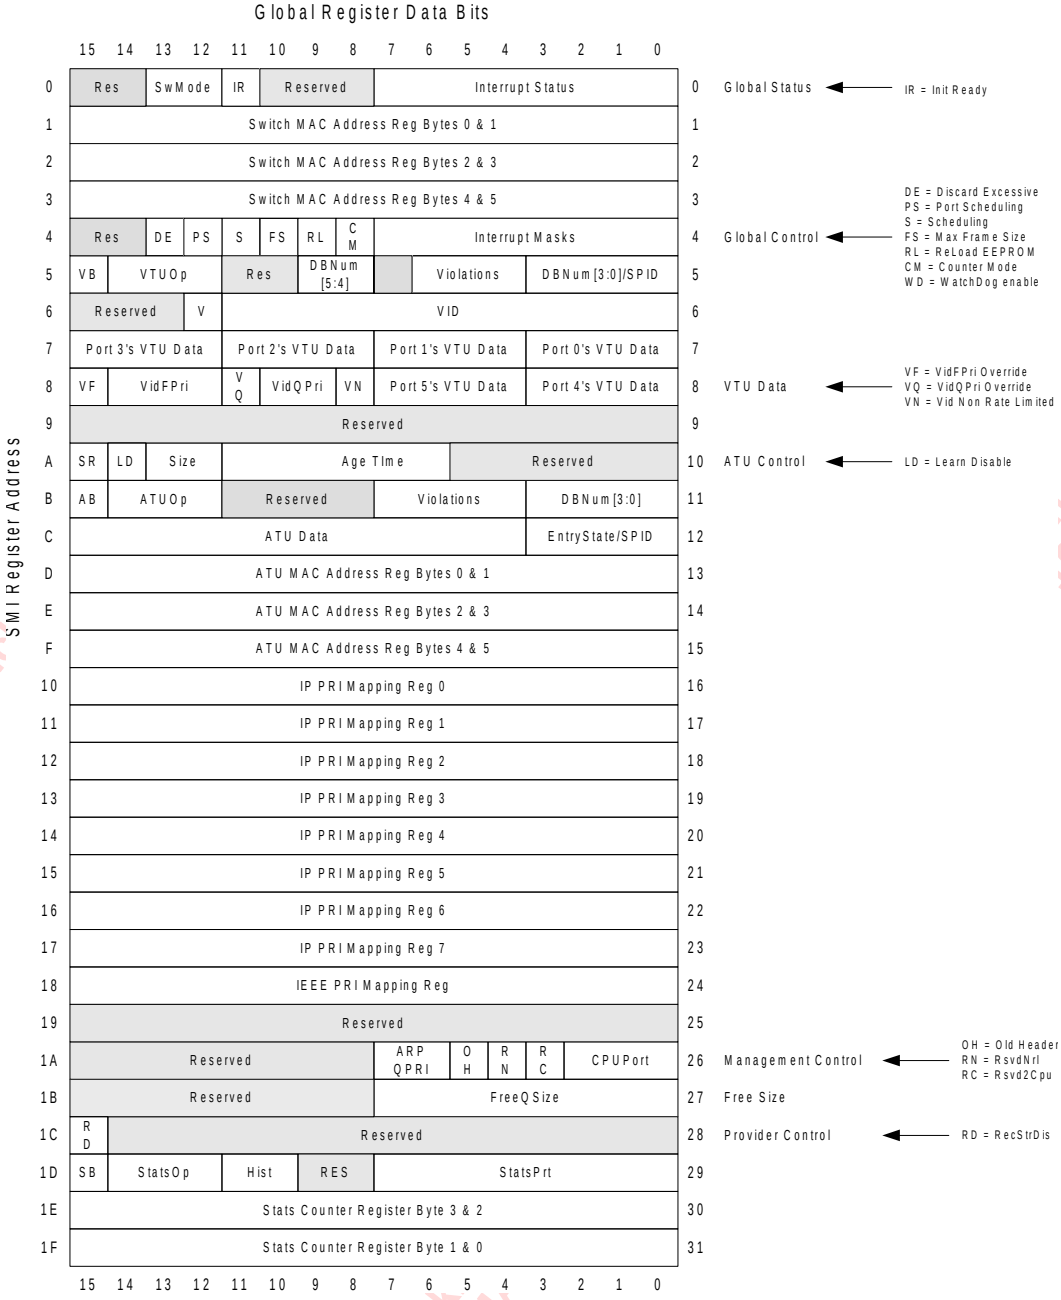

Table 69: Switch Global Status Registers
Offset: 0x00 or Decimal 0

Bits	Field	Mode	Description
15:14	Reserved	RES	Reserved for future use.
13:12	SW_Mode	RO	Switch Mode. These bits return the value of the SW_MODE[1:0] pins.
11	InitReady	RO	Switch Ready. This bit is set to a one when the Address Translation Unit, the VLAN Translation Unit, the Queue Controller and the Statistics Controller are done with their initialization and are ready to accept frames.
10:7	Reserved	RES	Reserved for future use.
6	StatsDone	ROC	Statistics Done Interrupt. This bit is set to a one whenever the STATBusy bit (Switch Global register - Offset 0x1D) transitions from a one to a zero. It is automatically cleared when read. This bit being high will cause the 88E6065/88E6035 device's INTn pin to go low if the STATDoneIntEn bit in Global Control (Switch Global register - Offset 0x04) 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 the pending VTU Violations have been serviced by the VTU Get/Clear Violation Data operation (Switch Port register - 0x05). This bit being high will cause the 88E6065/88E6035 device's INTn pin to go low if the VTUProbIntEn bit in Global Control (Switch Global register - Offset 0x04) is set to a one.
4	VTUDone	ROC	VTU Done Interrupt. This bit is set to a one whenever the VTUBusy bit (Switch Port register - 0x05) transitions from a one to a zero. It is automatically cleared when read. This bit being high will cause the 88E6065/88E6035 device's INTn pin to go low if the VTUDoneIntEn bit in Global Control (Switch Global register - Offset 0x04) is 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 (Switch Global register - 0x0B). This bit being high will cause the 88E6065/88E6035 device's INTn pin to go low if the ATUFullIntEn bit in Global Control (Switch Global register - Offset 0x04) is set to a one.
2	ATUDone	ROC	ATU Done Interrupt. This bit is set to a one whenever the ATUBusy bit (Switch Global register - 0x0B) transitions from a one to a zero. It is automatically cleared when read. This bit being high will cause the 88E6065/88E6035 device's INTn pin to go low if the ATUDoneIntEn bit in Global Control (Switch Global register - Offset 0x04) 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 will cause the 88E6065/88E6035 device's INTn pin to go low if the PHYIntEn bit in Global Control (Switch Global register - Offset 0x04) is set to a one.
0	EEInt	ROC	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 will cause the 88E6065/88E6035 device's INTn pin to go low if the EEIntEn bit in Global Control (Switch Global register - Offset 0x04) is set to a one.

Table 70: 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 1st 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 2:0 are the port number (Port 0 = 0, Port 1 = 1, 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 71: 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 72: 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 2:0 of this register are ignored if DiffAddr, above, is set to a one.

Table 73: Switch Global Control Register
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
15:14	Reserved	RES	Reserved for future use.
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	UsePort Schedule	RWR	Use per port scheduling mode as follows: 0 = Use Scheduling (bit 11 below) as the scheduling mode on all ports 1 = Use PortSched bit (Switch Port register - Offset 0x0B) as the scheduling mode on a per port basis
11	Scheduling	RWR	Scheduling mode. This bit is used to select the Queue Controller's scheduling mode (when UsePortSchedule is zero) as follows: 0 = Use an 8, 4, 2, 1 weighted fair queuing scheme 1 = Use a strict priority scheme
10	MaxFrame Size	SC	Maximum Frame Size allowed. The Ingress block will discard all frames less than 64 bytes in size. It will also discard all frames greater than a certain size regardless if the frame as follows: 0 = Max size is 1522 if IEEE 802.3ac tagged, or 1518 if not tagged 1 = Max size is 2048
9	ReLoad	RWR	Reload the registers using the EEPROM. When this bit is set to a one the contents of the external EEPROM will be used to load the registers as if a reset had occurred. When the reload operation is complete, this bit will be cleared to a zero automatically and the EEInt interrupt will be set. In customer configurations which shut down the MDC (clock) if the MDIO bus is IDLE, in order to guarantee proper clock operation, it is recommended that a dummy MDIO read be carried out before a switch over between MDIO mode and EEPROM mode i.e., after setting the Reload bit, a dummy MDIO read operation should be performed.
8	CtrMode	RWR	Counter Mode. When this bit is set to a one the Rx Counters for all ports (Switch Port register - Offset 0x10) count Rx Errors, the Tx Counters for all ports (Switch Port register - Offset 0x11) count Tx Collisions and the Policy Counters for all ports (Switch Port register - Offset 0x12) count Filtered frames. When this bit is cleared to a zero the Rx Counters for all ports count Rx Frames, the Tx Counters for all ports count Tx Frames and the Policy Counters count Discarded frames. The Rx Counters, the Tx Counters and the Policy Counters for all ports are cleared to a zero whenever this bit changes state (i.e., it transitions from a one to a zero or from a zero to a one).
7	Reserved	RES	Reserved for future use.

Table 73: Switch Global Control Register (Continued)
Offset: 0x04 or Decimal 4

Bits	Field	Type	Description
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 88E6065/88E6035 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 VT Problem interrupt to drive the 88E6065/88E6035 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 VT Done interrupt to drive the 88E6065/88E6035 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 88E6065/88E6035 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 88E6065/88E6035 device's INTn pin low.
1	PHYIntEn		PHY Interrupt Enable. This bit must be set to a one to allow active interrupts enabled in PHY registers 0x12 to drive the 88E6065/88E6035 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 88E6065/88E6035 device's INTn pin low.

Table 74: 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 (Switch Global register - Offset 0x04).
14:12	VTUOp	RWR	VLAN Table Unit Opcode. The 88E6065/88E6035 device supports 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]
11:10	Reserved	RES	Reserved for future use
9:8	DBNum[5:4]	RWR	VTU MAC Address Database Number bits 5:4. On all VTUOps except for Get Violation Data, this field is DBNum[5:4] 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 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 lower 4 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 Port 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 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 1st 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 1st Miss Violation or Member Violation (above) will be saved until cleared.

Table 74: VTU Operation Register (Continued)
Offset: 0x05 or Decimal 5

Bits	Field	Type	Description
4	VTUFull Violation	RO	VLAN Translation Table Full Violation. This bit is set to a one if the Load VTUOp could not store the desired entry. This bit causes the VTUProbInt (in Global Status – Offset 0x00) to be set causing an interrupt if enabled. This bit will be cleared after a Get/Clear Violation Data VTUOp that returns an SPID (bits 3:0 below) of 0x0F.
3:0	DBNum[3:0] /SPID	RWR	VTU MAC Address Database Number bits 3:0 or Source Port ID. On all VTUOps except for Get Violation Data, 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 desired 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 2 bits of the VTU's DBNum are in bits 9:8 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 = 0x0F the source of the violation was the CPU's registers interface (i.e., the VTU was full during a CPU Load operation).

1. An Entry is Loaded if the Valid bit (Switch Global register - Offset 0x06) 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 if it exists and the Valid bit (Switch Global register - Offset 0x06) is cleared to a zero.
3. A Get Next operation finds the next higher VID currently in the VTU's database. The VID value (Switch Global register - Offset 0x06) 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 a 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 - Offset 0x00) 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 (Switch Global register - Offset 0x06). When all Violations currently pending in the VTU have been serviced the VTUProb bit in Global Status will be cleared to a zero.

Table 75: 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 one's 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 (Switch Global registers - Offset 0x07) or the VID that caused the VTU Violation.

Table 76: VTU Data Register Ports 0 to 3
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15:14	PortState P3	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	PortState P2	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	PortState P1	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 if 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	PortState P0	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 the port's PortState bit (port Offset 0x04) 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.

Table 76: VTU Data Register Ports 0 to 3 (Continued)
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
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[2] and are not allowed to Egress this port. **NOTE:** These bits follow a different order compared to previous, pin compatible) devices.

1. The VID used comes from the VID in Tagged frames or the default VID assigned to Untagged frames.
2. The Ingress Discard occurs if the port's 802.Q mode is secure (port Offset 0x08).

Table 77: VTU Data Register Ports 4 to 5
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
15	UseVID FPri	RWR	Use VID Frame Priority Override. When this bit is set to a one the VTUFPri bits (below) can be used to override the frame priority (FPri) on any frame associated with this VID. The FPri assigned to a frame is used to as the tag's PRI bits if the frame egresses the switch tagged. The egress queue the frame is switched into is not modified by an FPri override.
14:12	VIDFPri	RWR	VID Frame Priority bits. These bits are used to override the frame priority on any frame associated with this VID value, if the UseVIDFPri bit (above) is set to a one and the port's VID FPriOverride is enabled (in Port Control 2 – Switch Port registers - Offset 0x0A).
11	UseVID QPri	RWR	Use VID Queue Priority Override. When this bit is set to a one the VTUQPri bits (below) can be used to override the queue priority (QPri) on any frame associated with this VID. The QPri of a frame is used to determine the egress queue the frame is switched into. If the frame egresses tagged the priority in the frame will not be modified by a QPri override.
10:9	VIDQPri	RWR	VID Queue Priority bits. These bits are used to override the queue priority on any frames associated with this VID value, if the UseVIDQPri bit (above) is set to a one and the port's VID QPriOverride is enabled (in Port Control 2 – Switch Port registers - Offset 0x08).
8	VidNrl	RWR	VID Non Rate Limit. This bit is used to indicate any frames associated with this VID value are to bypass ingress and egress rate limiting, if the ingress port's VidNrl bit (Switch Port registers - Offset 0x09) is also set to a one.

Table 77: VTU Data Register Ports 4 to 5 (Continued)
Offset: 0x08 or Decimal 8

Bits	Field	Type	Description
7:6	PortState P5	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	PortState P4	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 the port's PortState bit (port Offset 0x04) 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 are discarded at Ingress and are not allowed to Egress this port. **NOTE:** These bits follow a different order compared to previous, pin compatible) devices.

Table 78: ATU Control Register
Offset: 0x0A or Decimal 10

Bits	Field	Type	Description
15	SWReset	SC	Switch Software Reset. Writing a one to this bit causes the QC and the MAC state machines in the switch to be reset (i.e., the datapath). Register values are not modified and the EEPROM will not be re-read. The PHYs, VTU, Stats and ATU are not effected by this bit. When the reset operation is done, this bit will be cleared to a zero automatically. The reset will occur immediately. So to prevent transmission of CRC frames, all the ports should be set to the Disabled state (Switch Port registers - Offset 0x04), and wait for 2 ms. (i.e., the time for a maximum frame to be transmitted at 10 Mbps) before the SWReset bit is set to a one. Refer to Figure 46 to see the portions of the device that are reset by this bit.
14	LearnDis	RWR	ATU Learn Disable. 0 = Normal operation, learning is determined by the PortState (Switch Port registers - Offset 0x04) 1 = Automatic learning is disabled on all ports – CPU ATU Loads still work
13:12	ATUSize	RWS to 0x02	Address Translation Unit Table Size. The initial size of the ATU database is 1,024 entries. The size of the ATU database can be modified at any time, but an ATU reset (causing a full flush of the existing address database) and a Switch Software Reset (bit 15 above) occurs automatically if the new ATUSize is different from the old ATUSize. 00 = 256 Entry Address Database 01 = 512 Entry Address Database 10 = 1024 Entry Address Database (Default) 11 = Reserved
11:4	AgeTime	RWS to 0x16	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 0x16 is 22 decimal. 22 x 15 = 330 seconds or just over 5.5 minutes. The minimum age time is 0x01 or 15 seconds. The maximum age time is 0xFF or 3825 seconds or almost 64 minutes. If the AgeTime is set to 0x00 the Aging function is disabled and all learned addresses will remain in the database forever.
3:0	Reserved	RES	Reserved for future use.

Table 79: 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 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 (Switch Global registers - Offset 0x04).
14:12	ATUOp	RWR	Address Translation Unit Opcode. The 88E6065/88E6035 supports the following ATU operations (all of these operations can be executed while frames are transiting through the switch): 000 = No Operation 001 = Flush[1] /Move[2] All Entries 010 = Flush[3] /Move all Non-Static Entries 011 = Load[4] or Purge[5] an Entry in a particular DBNum Database 100 = Get Next[6] from a particular DBNum Database 101 = Flush/Move All Entries in a particular DBNum Database 110 = Flush/Move all Non-Static Entries in a particular DBNum Database 111 = Get/Clear Violation Data[7]
11:10	Reserved	RES	Reserved for future use.
9:8	DBNum[5:4]	RWR	ATU MAC Address Database Number bits 5:4. 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 on the Database supported commands (ATUOps 0x03, 0x04, 0x05 and 0x06 above). On Get/Clear Violation Data ATUOps these bits return the DBNum[5:4] value associated with the ATU violation that was just serviced. The lower 4 bits of the ATU's DBNum are in bits 3:0 of this register.
7	Age Violation	RO	Age Violation. On Get/Clear Violation Data ATUOps, this bit is returned set to a one if the Violation being serviced was due to a Source Address look-up that resulted in a Hit but where the age on the non-Static entry was less than 4 (i.e., the entry needs to be refreshed). This violation will occur only on Locked Ports (Switch Port registers - Offset 0x0B). Only the 1st Age, Member, Miss or Full Violation will be saved per port until cleared.
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 was due to a Source Address look-up that resulted in a Hit but where the PortVec bits did not contain the frame's Ingress port bit set to a one (i.e., a station move occurred). This violation will occur if the ATU entry was Static or on Locked Ports (Switch Port registers - Offset 0x0B) with non-Static entries. This violation can be masked on a per port basis by setting the port's IgnoreWrongData bit. Only the 1st Age, Member, Miss or Full Violation will be saved per port until cleared.

Table 79: ATU Operation Register (Continued)
Offset: 0x0B or Decimal 11

Bits	Field	Type	Description
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 was due to a Source Address look-up that resulted in a Miss. This violation will occur only on Locked Ports (Switch Port registers - Offset 0x0B - i.e., CPU directed learning is enabled on the port). Only the 1st Age, Member, Miss or Full Violation will be saved per port until cleared.
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 was due to a Load ATUOp or automatic learn that could not store the desired entry. This will only occur if all available locations for the desired address contain other MAC addresses that are loaded Static. Only the 1st Age, Member, Miss or Full Violation will be 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 0x03, 0x04, 0x05 and 0x06 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 2 bits of the ATU's DBNum are in bits 9:8 of this registers.

1. A Flush occurs if the EntryState (Switch Port registers - Offset 0x0B) is zero.
2. Move is used for 802.1w (rapid spanning tree) to reassign all valid entries associated with one port (the FromPort - Switch Port registers - Offset 0x0B) and move the association to another port (the ToPort - Switch Global registers - Offset 0x0C). It can also be used to completely disassociate a port from the database (if the ToPort = 0x0F). The Move occurs if the EntryState (Switch Port registers - Offset 0x0B) is 0x0F.
3. A Non-Static entry is any unicast address with an EntryState less than 0x08. All unicast frames will flood until new addresses are learned.
4. An Entry is Loaded if the EntryState (Switch Global registers - Offset 0x0C) is non-zero.
5. An Entry is Purged if it exists and if the EntryState (Switch Global registers - Offset 0x0C) is zero.
6. 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 (Switch Global registers - Offset 0x0C) are used as the address to start from. 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 Entry_State of 0x00, no higher MAC address was found. If ATUByte[5:0] is returned set to all one's with a non-zero Entry_State, 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.
7. When the ATUProb bit is set to a one (Global Status - Switch Global registers - Offset 0x00), 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 will be cleared to a zero.

Table 80: ATU Data Register for all but ATU Flush/Move
Offset: 0x0C or Decimal 12

Bits	Field	Type	Description
15:14	MacQPri	RWR	MAC Queue Priority Data. These bits are used to override the queue priority on any frame associated with this MAC value, if the EntryState indicates Queue Priority Override and if the port's SA and/or DA QPriOverrides are enabled (in Port Control 2 – Switch Port registers - Offset 0x08). SA Frame Priority Override can only occur on MAC addresses that are Static or where the Port is Locked, and where the port is the mapped source port for the MAC address. The QPri of a frame is used to determine the egress queue the frame is switched into. If the frame egresses tagged the priority in the frame will not be modified by a QPri override.
13	UseMacFPri	RWR	Use MAC Frame Priority Override. When this bit is set to a one, the MacFPri bits (below) can be used to override the frame priority (FPri) on any frame associated with this MAC. The FPri assigned to a frame is used to as the tag's PRI bits if the frame egresses the switch tagged. The egress queue the frame is switch into is not modified by an FPri override.
12:10	MacFPri	RWR	MAC Frame Priority Data. These bits are used to override the frame priority on any frame associated with this MAC value, if the UseMacFPri bit (above) is set to a one and the port's SA and/or DA FPriOverrides are enabled (in Port Control 2 – Switch Port registers - Offset 0x08). SA Frame Priority Overrides can only occur on MAC addresses that are Static or if the Port is Locked, and if the port is the mapped source port for the MAC address.
9:4	PortVec	RWR	Port Vector. These bits are used as the Port Vector for ATU Load operations and it's the resulting Port Vector from ATU Get Next operations. On Get/Clear Violation Data ATUOps these bits return the PortVec value associated with the ATU violation that was just serviced.
3:0	EntryState/ SPID	RWR	ATU Entry State. These bits are used as the Entry State for ATU Load/ Purge operations and it's the resulting Entry State from ATU Get Next operations. If these bits equal 0x0 then the ATUOp will be a Purge. If these bits are not 0x0 then the ATUOp will be a Load. 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 81: ATU Data Register for ATU Flush/Move
Offset: 0x0C or Decimal 12

Bits	Field	Type	Description
15:12	Reserved	RES	Reserved for future use.
11:8	ToPort		To Port. These bits are used as the ToPort during ATU Move operations. If the ToPort = 0xF, the operation becomes a Remove Port (i.e., the From-Port is removed from the database and the entry is purged if the resulting PortVec equals zeros).
7:4	FromPort		From Port. These bits are used as the FromPort during Move operations.
3:0	EntryState		ATU Entry State. These bits are used as the Entry State for ATU Flush/ Move operations. If these bits equal 0x0 then the ATUOp will be a Flush. If these bits are 0xF then the ATUOp will be a Move. All other values are reserved for future use.

Table 82: ATU MAC 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 input MAC address for ATU Load, Purge or Get Next operations and it's the resulting MAC address from ATU Get Next operations. Bit 0 of byte 0 (bit 40) is the multi-cast bit (it is the 1st 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's 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 83: ATU MAC 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's 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's 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 84: ATU MAC 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's 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's 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 85: IP-QPRI 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_0x16	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 86: IP-QPRI 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 87: IP-QPRI 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	RW0S 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 88: IP-QPRI 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 89: IP-QPRI 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 90: IP-QPRI 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 91: IP-QPRI 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 92: IP-QPRI 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 93: IEEE-QPRI Mapping 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 94: Management Control
Offset: 0x1A or Decimal 26

Bits	Field	Type	Description
15:8	Reserved	RES	Reserved for future use.
7:6	ArpQPri	RWS to 0x1	ARP Queue Priority is used for ARP QPri Overridden frames. When a ARP frame is received on a port that has its ARP QPriOverride bit set to a one (Port Control 2 register, Offset 0x08), the QPri assigned to the frame comes from these bits.
5	OldHeader	RWR	Egress Old Headers. When this bit is set to a one, and frames are egressed with a Marvell Header, the format of the Header is slightly modified to be backwards compatible with previous devices that used the original Header. Specifically, bit 3 of the Header's 2nd octet is cleared to a zero such that only FPri[2:1] is available in the Header.
4	RsvdNrl	RWR	Reserved Non Rate Limit. When this bit is set to a one, frames that match the requirements of the Rsvd2Cpu bit below will also be considered to be ingress and egress non rate limited.

Table 94: Management Control (Continued)
Offset: 0x1A or Decimal 26

Bits	Field	Type	Description
3	Rsvd2Cpu	RWR	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 EntryState, mapping it toward the CPU's port. This allows proprietary protocols to be supported while also supporting standard protocols. If multiple address databases are used, each multicast address needs to be added 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 CPUPort (bits 2:0 below). These frames are considered to have a QPri of 0x3.
2:0	CPUPort	RWS to 0x7	CPU Port. These bits indicate the destination port for Rsvd2Cpu frames (bit 3 above) and IGMP/MLD Snooped frames (Switch Port registers - Offset 0x04). The CPUPort bits need to indicate the port number on this device where the CPU is connected. **NOTE:** If the ATU is used for MGMT or BPDU frame detection, these frames are directed to the correct port where the CPU is connected by insuring that the ATU entry's PortVec bit points to the CPU's port. This register is ignored in this case.

Table 95: Total Free Counter
Offset: 0x1B or Decimal 27

Bits	Field	Type	Description
15:9	Reserved	RES	Reserved for future use.
8:0	FreeQSize	RO	Free Queue Size Counter. This counter reflects the current number of unallocated buffers available for all the ports.

Table 96: Provider Control
Offset: 0x1C or Decimal 28

Bits	Field	Type	Description
15	RecStrDis	RWR	Recursive Strip Disable. If the ProviderTag (Offset 0x1A) is programmed to something other than 0x8100 and if the incoming packet is provider tagged (refer to ProviderTag register, Offset 0x1A for further details), a value of zero for this bit ensures that multiple provider tags get stripped (note that multiple provider tags only get stripped from the packet if the incoming tag matches the Provider Tag register, Offset 0x1A). When this bit is set to a one, recursive tag stripping feature is disabled and only one provider tag will be removed during ingress on provider ports.
14:0	Reserved	RES	Reserved for future use.

Table 97: 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 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 (Switch Global registers - Offset 0x04).
14:12	StatsOp	RWR	Statistics Unit Opcode. The 88E6065/88E6035 supports 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 or Direct Counter 101 = Capture All Counters for a Port 11x = Reserved
11:10	Histogram Mode	RWS 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 received and transmitted frames
9:8	Reserved	RES	Reserved for future use.

Table 97: Stats Operation Register (Continued)
Offset: 0x1D or Decimal 29

Bits	Field	Type	Description
7:5	StatsPort	RWR	Access Statistics Counters directly for a Port or the Capture area. These bits can be used to directly access a ports counters without doing a capture first. Use bits 7:5 = 000 to access the captured counters. Use bits 7:5 = 001 to access the counters for Port 0. Use bits 7:5 = 010 to access the counters for Port 1, etc. These bits represents the port number for the following StatsOps: • Flush (clear) All Counters for a Port • Read a Captured or Direct Counter • Capture all Counters for a Port
4:0	StatsPtr	RWR	Statistics Pointer. This field is used as a parameter for the StatsOp commands. StatsPtr must be set to the desired counter to Read a Captured or Direct Counter (0x4) StatsOp (valid range is 0x00 to 0x1F). A Capture All Counters for a Port StatsOp should be carried out prior to using the Read A Captured Counter StatsOp with StatsPort = 0x0. The counter that is read is defined as follows: Ingress Counters[1] 0x00 – InGoodOctets 0x01 – Reserved 0x02 – InBadOctets 0x04 – InUnicast 0x06 – InBroadcasts 0x07 – InMulticasts 0x16 – InPause 0x18 – InUndersize 0x19 – InFragments 0x1A – InOversize 0x1B – InJabber 0x1C – InRxErr 0x1D – InFCSErr Egress Counters 0x0E – OutOctets[2] 0x0F – Reserved 0x10 – OutUnicast 0x12 – OutMulticasts 0x13 – OutBroadcasts 0x15 – OutPause 0x03 – OutFCSErr 0x05 – Deferred 0x11 – Excessive 0x14 – Single 0x17 – Multiple 0x1E – Collisions[3] 0x1F – Late Histogram Counters[4] 0x08 – 64Octets 0x09 – 65to127Octets 0x0A – 128to255Octets 0x0B – 256to511Octets 0x0C – 512to1023Octets 0x0D – 1024toMaxOctets

1. If Marvell Header mode is used on Ports 4 to 5 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. The Collisions, Deferred, Single, Multiple, Excessive and Late counters will not increment on full-duplex ports.
4. If Marvell Header mode is used on Ports 4 to 5, the extra two bytes in the frame are not included in the count before determining which Histogram Counter to increment.

Table 98: Stats Counter Register
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 – Switch Global registers - Offset 0x1D).
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 – Switch Global registers - Offset 0x1D).

Table 99: Stats Counter Register
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 – Switch Global registers - Offset 0x1D).
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 – Switch Global registers - Offset 0x1D).

5.2.3 Switch Global 2 Registers

The 88E6065/88E6035 device contains global registers that affect all the Ethernet ports in the device. This section documents the Global 2 registers. Global 1 registers are documented in Section 5.2.2 "Switch Global 1 Registers" on page 174. Each global register is 16-bit wide and their bit assignments are shown in Figure 46.

Figure 46: Switch Global 2 Registers

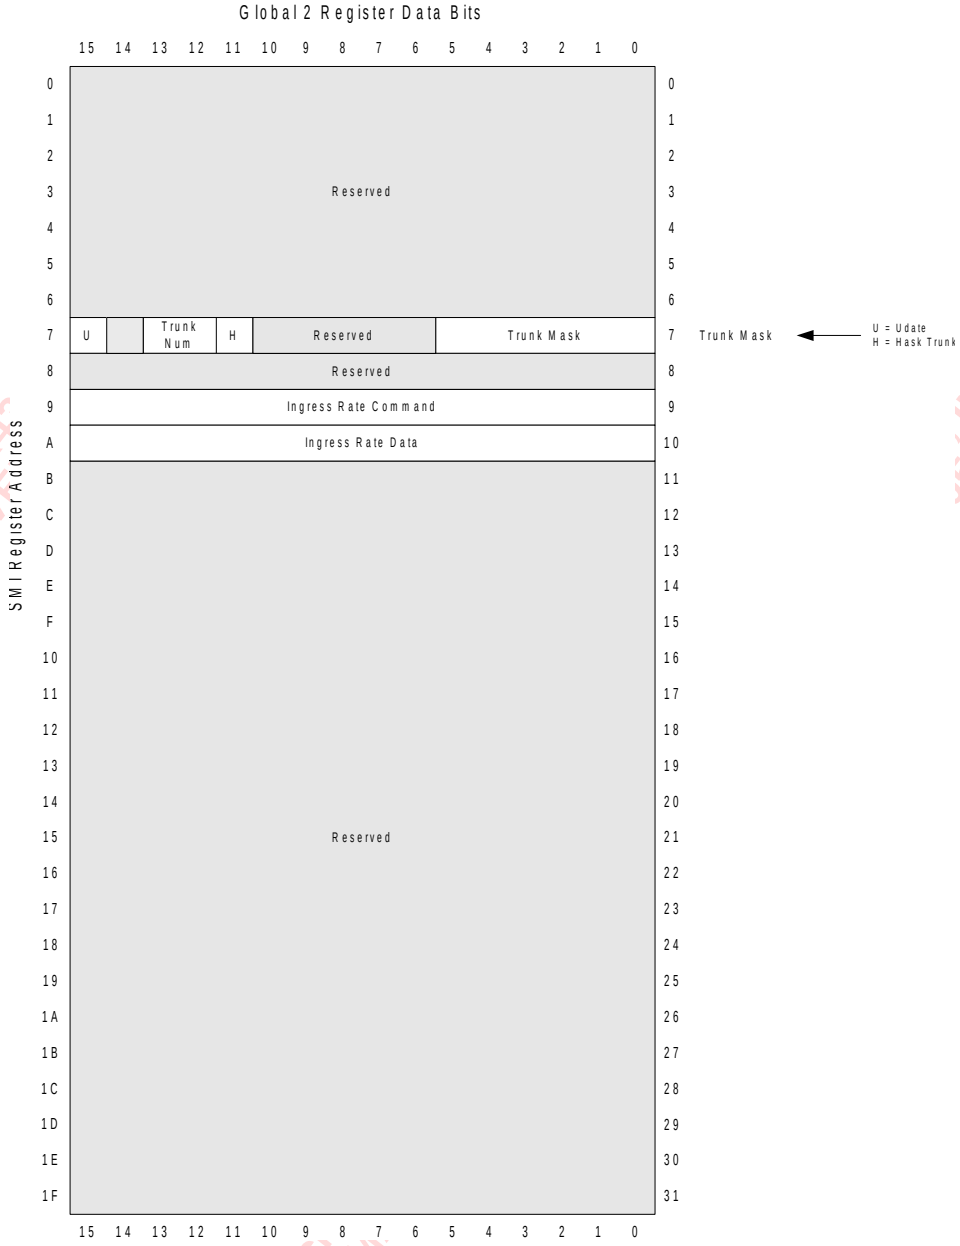

Table 100: Trunk Mask Table Register
Offset: 0x07 or Decimal 7

Bits	Field	Type	Description
15	WriteTMD	SC	Write Trunk Mask data. When this bit is set to a one, the data written to bits 5:0 will be written to the Trunk Mask selected by the MaskNum bits below. After the write has taken place this bit self clears to zero.
14	Reserved	RES	Reserved for future use.
13:12	MaskNum	RWR	Mask Number. These bits select one of four possible Trunk Mask vectors for both read and write operations. A write operation occurs if the WriteTMD bit is a one (the Trunk Mask Table can be written by writing to this register once). Otherwise, a read of the current MaskNum occurs and the data found there is placed in the TrunkMask bits below (the Trunk Mask Table can be read by 1st writing to this register – with WriteTMD = 0, and then reading this register).
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 2 bits of the frame's DA and SA are XOR'ed together to select the TrunkMask to use.
10:6	Reserved	RES	Reserved for future use.
5:0	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 WriteTMD 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 WriteTMD bit being a zero, these bits are not written anywhere (this allow the TrunkNum 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 TrunkNum bits. The TrunkMask is reset to all ones for all TrunkNum entries.

Table 101: Ingress Rate Command Register[1]
Offset: 0x09 or Decimal 9

Bits	Field	Type	Description
15	IRLBusy	SC	Ingress Rate Limit unit Busy. This bit must be set to a one to start an IRL operation (see IRLOp below). Only one IRL operation can be executing at one time so this bit must be zero before setting it to a one. When the requested IRL operation completes this bit will automatically be cleared to a zero.
14:12	IRLOp	RWR	Ingress Rate Limit unit Opcode. The 88E6065/88E6035 supports the following IRL operations (all of these operations can be executed while frames are transiting through the switch): 000 = No operation 001 = Init all resources to the initial state 010 = Init the selected resource (pointed to by IRLUnit) to the initial state. This initializes internal rate limiting related counters. 011 = Write to the selected resource/register (IRLUnit/IRLReg)[2] 100 = Read the selected resource/register (IRLUnit/IRLReg)[3] 101 = Reserved 110 = Reserved 111 = Reserved
11:8	Reserved	RES	Reserved for future use.
7:4	IRLUnit	RWR	Ingress Rate Limit Unit. These bits are used to define one of twelve Ingress Rate Limiting resources that are being written or read or initialized. Use a value of 0x00 to access resource 0, a value of 0x01 to access resource 1, etc.
3:0	IRLReg	RWR	Ingress Rate Limit register. These bits are used to define the controlling register being written or read on the resource defined in IRLUnit above. Use a value of 0x00 to access register 0, a value of 0x01 to access register 1, etc.

1. This register is used to access PIRL registers.
2. 1st write the data to the IRLData register (Global 2, Offset 0x0A), and then perform an IRLWrite.
3. 1st perform an IRLRead and then the requested data will be available in the IRLData register (Global 2, Offset 0x0A).

Table 102: Ingress Rate Data Register[1]
Offset: 0x0A or Decimal 10

Bits	Field	Type	Description
15:0	IRLData	RWR	Ingress Rate Limit Data. The actual fields within the Bucket Configuration Memory (BCM) register are described in PIRL registers - Offset 0x00.

1. This register is used to access PIRL registers.

5.2.4 Port Ingress Limit (PIRL) Registers

The 88E6065/88E6035 device contains global registers that effect all the ingress rate limiting resources. These registers are accessed using IngressRateCommand and IngressRateData registers.

Table 103: Bucket Configuration Memory, IRL Unit
Offset: 0x00 to 0x0B; IRL Reg 0 to 7

Bits	Field	Type	Description
111	Account_for_ QCong	RWR	When enabled, this bit indicates that even if the incoming packets are discarded in the chip because of queue congestion, account for the incoming bytes in the rate bucket calculations.
110	Account_for_ filtered	RWR	When enabled, this bit indicates that even if the incoming packets are filtered in the chip, account for the incoming bytes in the rate bucket calculations.
109	Reserved	RES	Reserved for future use.
108	EBS_Limit_ Action	RWR	If the incoming port information rate exceeds the EBS_Limit, this field specifies the action that needs to be taken for the violating traffic. = 1'd0 indicates Discard the packet = 1'd1 indicates send flow control packet back to the source.
107-104	Reserved	RES	Reserved for future use.
103-80	EBS_Limit	RWR	Excess Burst Size limit; This effectively indicates the size of the bucket.
79-76	Reserved	RES	Reserved for future use.
75-52	CBS_Limit	RWR	Committed Burst Size limit; This indicates the committed burst amount. It is expected that this value be set to a fixed 2 Kbytes for 88E6065/ 88E6035.
51-48	Reserved	RES	Reserved for future use.
47-32	Bucket_Rate Factor	RWR	Bucket_decrement indicates the amount of tokens that need to be removed per each bucket decrement. This parameter, along with the amount of time since this particular rate resource has been updated, is taken into consideration for computing the total number of tokens that need to be decremented for a given configured information rate.
31-28	Reserved	RES	Reserved for future use.
27-16	Bucket_ Increment	RWR	This indicates the amount of tokens that need to be added for each byte of packet information.
15	Bucket_ RateTypeR	WR	1 = Indicates the bucket is rate based. 0 = Indicates the bucket is traffic type based.

Table 103: Bucket Configuration Memory, IRL Unit (Continued)
Offset: 0x00 to 0x0B; IRL Reg 0 to 7

Bits	Field	Type	Description
14-4	Bucket_TypeMask	RWR	This field has the following definition if Bucket_RateType = 0x0; [4] – Unknown Unicast [5] – Unknown Multicast [6] – Broadcast [7] – Multicast [8] – Unicast [9] – MGMT Frames [10] – ARP [14:11] – Reserved.
3:0	Bytes_tobe_counted	RWR	[0] – Count all bytes; [1] – Count all bytes except for layer1 bits; [2] – Count all bytes except for layer2 bits; [3] - Reserved All bytes: Preamble + SFD + IFG + SIZE L1: SIZE L2: SIZE – MRVLHDR - DA – SA – Len/Type – Tag where SIZE: Byte count reported by Ingress block and * is used to indicate that MRVLHDR and Tag may not be present always. The power on reset value for this is 0x0 For logic simplification it is assumed that the Preamble is always a constant 7 bytes long, SFD is 1 byte long and IFG is assumed to be a constant 12 bytes long. For example, if bits 1 and 2 are set then account for all bytes except for layer1 and layer2 bytes.

Figure 47: Bucket Configuration Data structure CPU view

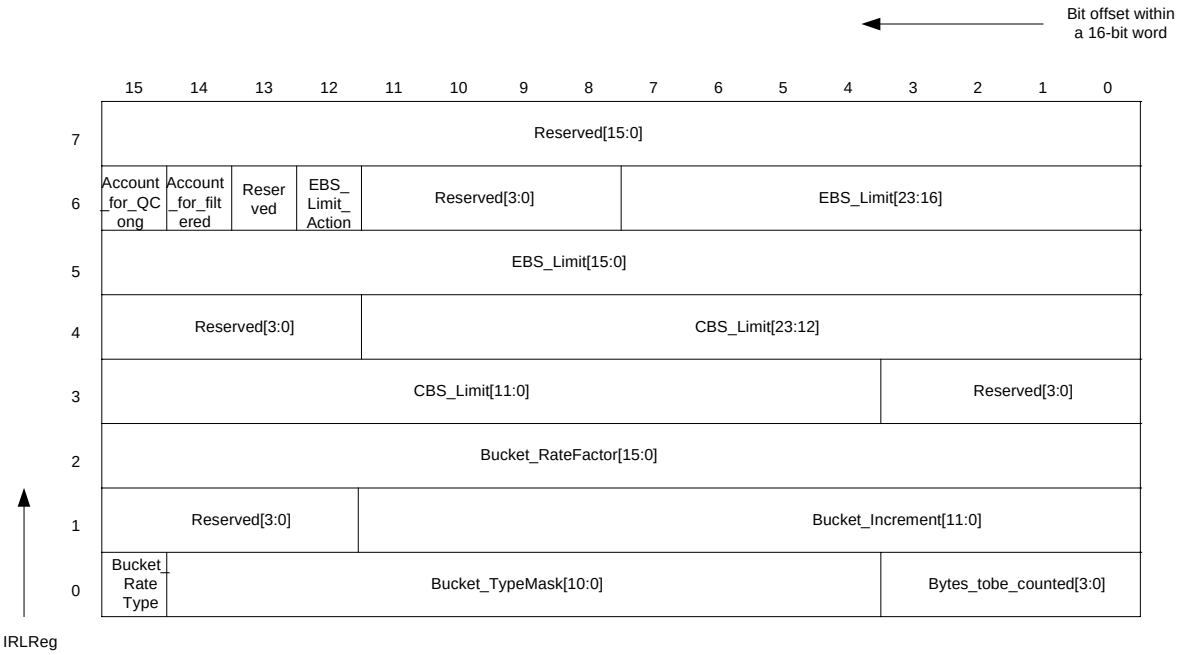

5.3 EEPROM Programming Format

The 88E6065/88E6035 supports an optional external serial EEPROM device for programming its internal registers. The EEPROM data will be read in once after Reset is deasserted unless the Stand Alone Switch Mode is selected (SW_MODE[1:0] = 0x2 – Table 13).

The 88E6065/88E6035 supports 2K bit (93C56) or 4K bit (93C66) 4-wire EEPROM devices as well as 1K bit (24C01), 2K bit (24C02) or 4K bit (24C04) 2-wire EEPROM devices (the type of the device, 2- or 4-wire, is selected by the EE_CS/EE_2WIRE pin at reset – Table 6). 4-Wire external EEPROM devices must be configured in x16 data organization mode. 2-Wire external EEPROM devices are treated as if they are in a x16 data organization by always reading both an odd address (as the upper 8-bits) and an even address (as the lower 8-bits).

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, this is called the Command
3. If the just read in Command is all one's, terminate the serial EEPROM reading process, go to 8
4. Increment the address to the next 16-bit address
5. Read in the 16-bits of data from the new address, this is called RegData
6. Write RegData to the location or locations defined by the previous Command
7. Go to 1
8. Set the EEInt bit in Global Status to a one (Switch Global registers - Offset 0x00) generating an Interrupt (if enabled)
9. Done

The 16-bit Command determines which register or registers inside the 88E6065/88E6035 are updated as follows (refer to Figure 48):

1. Bit 15 determines which set of registers can be written. If bit 15 = 0 the five external PHY device's registers can be written to or the two Global register spaces can be written to (SMI Device Addresses 0x0F & 0x0E). If bit 15 = 1 the Switch registers can be written to (SMI Device Addresses 0x08 to 0x0D). See section Section 5. and Figure 42 for more information on the registers and their addresses.
2. Bits 10:5 (or 9:5), the Device Vector, determines which Device Address or Addresses are written. Each bit of the Device Vector that is set to a one causes one Device Address to be written. Bit 5 controls writes for Port 0 (either PHY address 0x00 or Switch address 0x08). Bit 6 controls writes for Port 1 (either PHY address 0x01 or Switch address 0x09). Bits 7 controls writes for Switch Port 2. The Switch Global register set that is written to is determined by bit 6 the G bit. When G=1 the Global space at device address 0x0F is written to. When G=0 the Global 2 space at device address 0x0E is written to. Care is needed to insure Reserved registers are not written.
3. Bits 4:0, the Register Address, determine which SMI Register Address is written.
4. Two special commands are supported. Halt (end of EEPROM processing) and Release Internal Reset. Release Internal Reset is used to allow packets to start flowing (if Ports are in the Forwarding PortState) and to allow the EEPROM to be able to load the ATU and/or VTU as these blocks need internal reset released before they can be written to.

The format of the EEPROM Commands support writing the same RegData to all the PHY's or all the Switch's MAC's with one Command/RegData pair. The Command/RegData list can be as short or as long as needed. This makes optimum use of the limited number of Command/RegData pairs that can fit in a given size EEPROM (31 [1] Command/RegData pairs in the 24C01, 63 in the 93C56 or 24C02 and 127 in the 93C66 or 24C04).

1. The maximum number of Command/RegData pairs is one less than expected because the last entry must be the End of List Indicator of all one's.

Figure 48: EEPROM Data Form

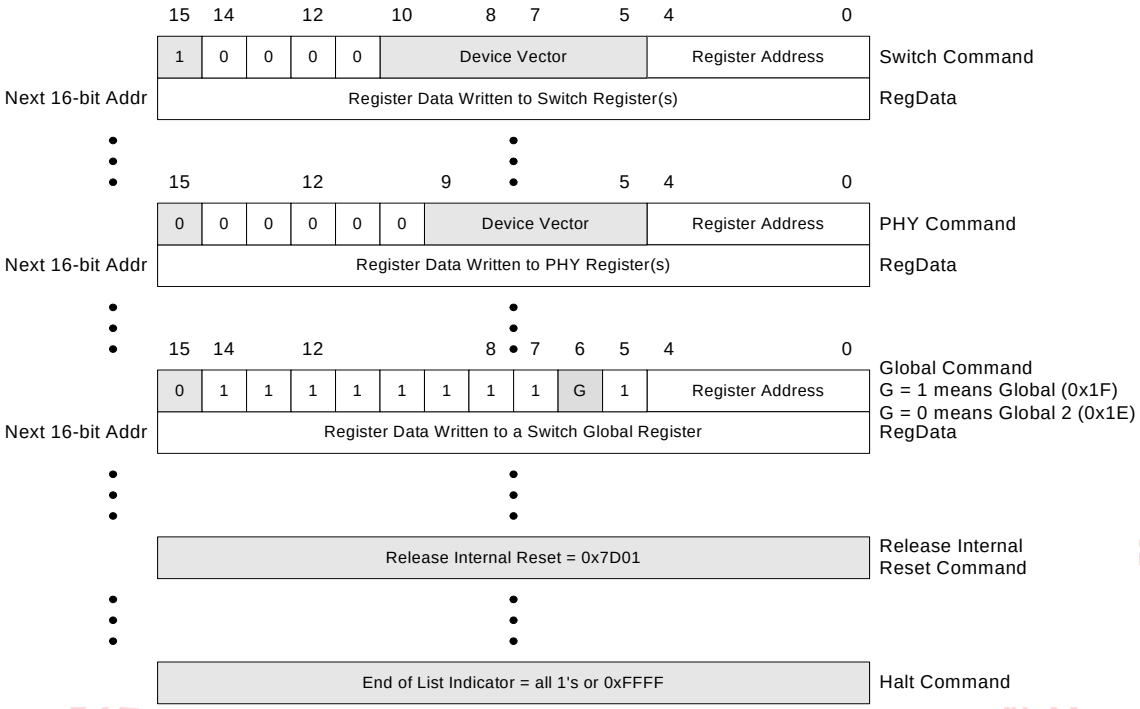

Section 6. PHY Register Description

The device contains six/three physical layer devices (PHYs). These devices are accessible using SMI device addresses 0x00 to 0x04 (or 0x10 to 0x12) depending upon the value of the P5_OUTD[3:0]/P5_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 104: PHY Register Map

Description	Offset Hex	Offset Decimal	Page Number
PHY Control Register	0x00	0	page 208
PHY Status Register	0x01	1	page 211
PHY Identifier	0x02	2	page 213
PHY Identifier	0x03	3	page 213
Auto-Negotiation Advertisement Register	0x04	4	page 214
Link Partner Ability Register (Base Page)	0x05	5	page 216
Link Partner Ability Register (Next Page)	0x05	5	page 217
Auto-Negotiation Expansion Register	0x06	6	page 218
Next Page Transmit Register	0x07	7	page 219
Link Partner Next Page Register	0x08	8	page 220
Reserved Registers	0x09-0x0F	9 - 15	page 220
PHY Specific Control Register I	0x10	16	page 221
PHY Specific Status Register	0x11	17	page 223
PHY Interrupt Enable	0x12	18	page 225
PHY Interrupt Status	0x13	19	page 226
PHY Interrupt Port Summary	0x14	20	page 228
PHY Receive Error Counter	0x15	21	page 229
LED Parallel Select Register	0x16	22	page 229
LED Stream Select for Serial LEDs (Global Register)	0x17	23	page 231
PHY LED Control Register	0x18	24	page 233
PHY Manual LED Override Register	0x19	25	page 234
VCT™ Control Register	0x1A	26	page 235
VCT™ Status Register	0x1B	27	page 236
PHY Specific Control Register II	0x1C	28	page 237
Reserved Registers	0x1D to 0x1F	29 - 31	page 237

Table 105: PHY Control Register
Offset: 0x00 or Decimal 0

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: 88E6065/88E6035 device = 1

Table 105: PHY Control Register (Continued)
Offset: 0x00 or Decimal 0

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: 88E6065/88E6035 device 0.12 = 1
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 105: PHY Control Register (Continued)
Offset: 0x00 or Decimal 0

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 106: PHY Status Register
Offset: 0x01 or Decimal 1

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 106: PHY Status Register (Continued)
Offset: 0x01 or Decimal 1

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 PHY registers - Offset 0x11. 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 107: PHY Identifier
Offset: 0x02 or Decimal 2

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 108: PHY Identifier
Offset: 0x03 or Decimal 3

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 109: Auto-Negotiation Advertisement Register
Offset: 0x04 or Decimal 4

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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) or link goes down.

Table 109: Auto-Negotiation Advertisement Register (Continued)
Offset: 0x04 or Decimal 4

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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) 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 - PHY registers - Offset 0x00) or link goes down.
4:0	AnegAd Selector	R/W	Always 0x01	Always 0x01	Selector Field Mode 00001 = 802.3

Table 110: Link Partner Ability Register (Base Page)
Offset: 0x05 or Decimal 5

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 (PHY registers - Offset 0x11) is disabled. When Reg8NxtPg (PHY registers - Offset 0x11) is enabled, then next page is stored in the Link Partner Next Page register (PHY registers - Offset 0x05), and the Link Partner Ability Register (PHY registers - Offset 0x05) 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 111: Link Partner Ability Register (Next Page)
Offset: 0x05 or Decimal 5

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 (PHY registers - Offset 0x10) is disabled. When Reg8NxtPg (PHY registers - Offset 0x11) is enabled, then next page is stored in the Link Partner Next Page register (PHY registers - Offset 0x05), and Link Partner Ability Register (PHY registers - Offset 0x05) 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 112: Auto-Negotiation Expansion Register
Offset: 0x06 or Decimal 6

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 (PHY registers - Offset 0x01) 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 (PHY registers - Offset 0x01). 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 113: Next Page Transmit Register
Offset: 0x07 or Decimal 7

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 114: Link Partner Next Page Register
Offset: 0x08 or Decimal 8

Bits	Field	Mode	HW Rst	SW Rst	Description
15	RxNxtPage	RO	0x0	0x0	If Reg8NxtPg (PHY registers - Offset 0x10) is enabled, then next page is stored in the Link Partner Next Page register (PHY registers - Offset 0x05); otherwise, the Link Partner Next Page register (PHY registers - Offset 0x05) 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 115: PHY Specific Control Register
Offset: 0x10 or Decimal 16

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 2 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 (PHY registers - Offset 0x08) to store Next Page. If set to store next page in the Link Partner Next Page register (PHY registers - Offset 0x08), then 802.3u is violated to emulate 802.3ab. 0 = Store next page in the Link Partner Ability Register (Base Page) register (PHY registers - Offset 0x05). 1 = Store next page in the Link Partner Next Page register (PHY registers - Offset 0x08)
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 115: PHY Specific Control Register (Continued)
Offset: 0x10 or Decimal 16

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.
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

Table 115: PHY Specific Control Register (Continued)
Offset: 0x10 or Decimal 16

Bits	Field	Mode	HW Rst	SW Rst	Description
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 116: PHY Specific Status Register
Offset: 0x11 or Decimal 17

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.

Table 116: PHY Specific Status Register (Continued)
Offset: 0x11 or Decimal 17

Bits	Field	Mode	HW Rst	SW Rst	Description
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
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 117: PHY Interrupt Enable
Offset: 0x12 or Decimal 18

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 117: PHY Interrupt Enable (Continued)
Offset: 0x12 or Decimal 18

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 118: PHY Interrupt Status
Offset: 0x13 or Decimal 19

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

Table 118: PHY Interrupt Status (Continued)
Offset: 0x13 or Decimal 19

Bits	Field	Mode	HW Rst	SW Rst	Description
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
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 119: PHY Interrupt Port Summary (Global[1])
Offset: 0x14 or Decimal 20

Bits	Field	Mode	HW Rst	SW Rst	Description
15:6	Reserved	RO	0x0	0x0	Must be 00000000.
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
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 120: Receive Error Counter
Offset: 0x15 or Decimal 21

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 121: LED Parallel Select Register (bits 11:0 are Global[1] bits)
Offset: 0x16 or Decimal 22

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)

Table 121: LED Parallel Select Register (Continued) (bits 11:0 are Global[1] bits)
Offset: 0x16 or Decimal 22

Bits	Field	Mode	HW Rst	SW Rst	Description
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)
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 122: LED Stream Select for Serial LEDs (Global Register)
Offset: 0x17 or Decimal 23

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 122: LED Stream Select for Serial LEDs (Global Register) (Continued)
Offset: 0x17 or Decimal 23

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 123: PHY LED Control Register (bits 14:0 are Global[1] bits)
Offset: 0X18 or Decimal 24

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 124: PHY Manual LED Override
Offset: 0x19 or Decimal 25

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 124: PHY Manual LED Override
Offset: 0x19 or Decimal 25

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 must be disabled on ports that are configured to blink an LED but don't have a link established.

Table 125: VCT™ Register for TXP/N Pins
Offset: 0x1A[1] or Decimal 26

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 125.
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.

Table 125: VCT™ Register for TXP/N Pins (Continued)
Offset: 0x1A[1] or Decimal 26

Bits	Field	Mode	HW Rst	SW Rst	Description
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 126: VCT™ Register for RXP/N pins
Offset: 0x1B[1] or Decimal 27

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 49: Cable Fault Distance Trend Line

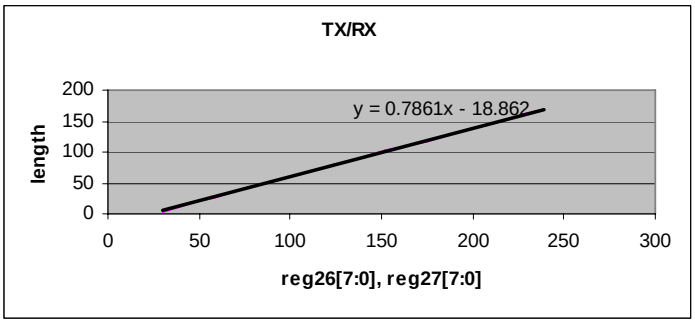

Table 127: PHY Specific Control Register II
Offset: 0x1C or Decimal 28

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 88E6065/88E6035 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 7. Electrical Specifications

7.1 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 128: Absolute Maximum Ratings

Symbol	Parameter	Min	Typ	Max	Units
$V_{DD(3.3)}$	Power Supply Voltage on V_{DDO} with respect to V_{SS}	-0.5	3.3	+3.6	V
$V_{DD(2.5)}$	Power Supply Voltage on V_{DDAH}, V_{DDO} with respect to V_{SS}	-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 V_{DD}, or V_{DDAL} with respect to V_{SS}	-0.5	1.5	+3.6 or $V_{DD(2.5)}$ +0.5[2] whichever is less	V
V_{PIN}	Voltage applied to any input pin with respect to V_{SS}	-0.5		+3.6 or V_{DDO} +0.5[3] whichever is less	V
$T_{STORAGE}$	Storage temperature	-55		+125[4]	°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. VPIN must never be more than 0.5V greater than VDDO or damage will result.
4. 125°C is the re-bake temperature. For extended storage time greater than 24 hours, +85°C should be the maximum.

7.2 Recommended Operating Conditions

Table 129: Recommended Operating Conditions

Symbol	Parameter	Condition	Min	Typ	Max	Units
$V_{DD(3.3)}$	3.3V power supply	For pins V_{DDO}	3.135	3.3	3.465	V
$V_{DD(2.5)}$	2.5V power supply	For pins V_{DDAH}, V_{DDO}	2.375	2.5	2.625	V
$V_{DD(1.5)}$	1.5V power supply	For pins V_{DD}, V_{DDAL}	1.425	1.5	1.575	V
T_A	Ambient operating temperature	Commercial Parts[1]	0		70	°C
		Industrial Parts[2,3]	-40		85	°C
T_J	Maximum junction temperature				125[4]	°C
RSET	Internal bias reference	Resistor value placed between RSET- and RSET+ pins	1980	2000	2020	Ω

1. The important parameter is the junction temperature. As long as the junction temperature is between 0 °C and 120 °C, the device can be operated at any ambient temperature.
2. The important parameter is the junction temperature. As long as the junction temperature is between -40 °C and 120 °C, the device can be operated at any ambient temperature.
3. Industrial part numbers have an “I” following the commercial part numbers. 9.1 "Ordering Part Numbers and Package Markings" on page 272.
4. Refer to white paper on TJ Thermal Calculations for more Information.

7.3 Package Thermal Information

7.3.1 Thermal Conditions for 128-pin LQFP Package

Table 130: Thermal Conditions for 128-pin PQFP Package

Symbol	Parameter	Condition	Min	Typ	Max	Units
θ_{JA}	Thermal resistance[1] - junction to ambient of the 88E6065/88E6035 device 128-Pin LQFP package $\theta_{JA} = (T_J - T_A)/ P$ P = Total Power Dissipation	JEDEC 4 in. x 4.5 in. 4-layer PCB with no air flow		33.1		°C/W
		JEDEC 4 in. x 4.5 in. 4-layer PCB with 1 meter/sec air flow		30.3		°C/W
		JEDEC 4 in. x 4.5 in. 4-layer PCB with 2 meter/sec air flow		29.1		°C/W
		JEDEC 4 in. x 4.5 in. 4-layer PCB with 3 meter/sec air flow		28.3		°C/W
ψ_{JT}	Thermal characteristic parameter[1] - junction to top center of the 88E6065/88E6035 device 128-Pin LQFP package $\psi_{JT} = (T_J - T_C)/P$. P = Total Power Dissipation	JEDEC 4 in. x 4.5 in. 4-layer PCB with no air flow		2.3		°C/W
θ_{JC}	Thermal resistance[1] - junction to case of the 88E6065/88E6035 device 128-Pin LQFP package $\theta_{JC} = (T_J - T_C)/ P_{Top}$ P_{Top} = Power Dissipation from the top of the package Thermal resistance[1] - junction to board of the 88E6065/88E6035 device 128-Pin LQFP package	JEDEC with no air flow		16.7		°C/W
θ_{JB}	$\theta_{JB} = (T_J - T_B)/ P_{bottom}$ P_{bottom} = power dissipation from the bottom of the package to the PCB surface.	JEDEC with no air flow		25.7		°C/W

1. Refer to white paper on TJ Thermal Calculations for more information.

7.4 DC Electrical Characteristics

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 131: DC Electrical Characteristics

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
$I_{DD(3.3)}$	3.3 volt Power to outputs	V_{DDO}	No link on any port		15		mA
			All ports 10 Mbps		17		mA
			All ports 100 Mbps		22		mA
			Power Down Mode		18		mA
			Sleep Mode		18		mA
$I_{DD(2.5)}$	2.5 volt[1] Power to analog core and digital I/O (If applicable)	V_{DDAH}, V_{DDO}	No link on any port		40		mA
			All ports 10 Mbps and active		126		mA
			All ports 100 Mbps		111		mA
			Power Down Mode		32		mA
			Sleep Mode		36		mA
	2.5 volt Power to Center Tap	External Magnetics Center Tap Pin	No link on any port		10		mA
			All ports 10 Mbps and active		301		mA
			All ports 100 Mbps		103		mA
			Power Down Mode		0		mA
			Sleep Mode		0		mA
$I_{DD(1.5)}$	1.5 volt Power to analog core	V_{DDAL}	No link on any port		0		mA
			All ports 10 Mbps and active		1		mA
			All ports 100 Mbps		29		mA
			Power Down Mode		1		mA
			Sleep Mode		0		mA
	1.5 volt Power to digital core	V_{DD}	No link on any port		40		mA
			All ports 10 Mbps and active		54		mA
			All ports 100 Mbps and active		119		mA
			Power Down Mode		39		mA
			Sleep Mode		36		mA

1. If class A mode of operation is selected than there is 20 mA additional current in the center tap for each port that is in 100 Mb/s mode. For 10M mode there is no change in current since it is always class B. Class A mode can be selected by setting PHY register 0x1C.0 = 1 for backplane applications.

7.4.1 Digital Operating Conditions

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 132: Digital Operating Conditions

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
V_{IH}	High level input voltage	All pins		2.0		V_{DDO} +0.5	V
		XTAL_IN		1.4			V
V_{IL}	Low level input voltage	All pins		-0.5		0.8	V
		XTAL_IN				0.6	V
V_{OH}	High level output voltage	LED pins[1]	I_{OH} = -8 mA V_{DDO} = Min	2.4			V
		XTAL_OUT	I_{OH} = -1 mA		$V_{IH(XTAL\_IN)}$ +0.2		V
		All others	I_{OH} = -4 mA V_{DDO} = Min	2.4			V
V_{OL}	Low level output voltage	LED pins[1]	I_{OL}= 8 mA			0.4	V
		XTAL_OUT	I_{OL} = 1 mA		$V_{IL(XTAL\_IN)}$ -0.2		V
		INTn pin[2]	I_{OL}= 8 mA				V
		All others	I_{OL}= 4 mA			0.4	V
I_{ILK}	Input leakage current	With pull-up resistor (typ. 120 kΩ)	$0<V_{IN}<V_{DDO}$			+ 10 - 50	µA
		With pull-down resistor (typ. 120 kΩ)	$0<V_{IN}<V_{DDO}$			+ 50 - 10	µA
		All others	$0<V_{IN}<V_{DDO}$			±10	µA
C_{IN}	Input capacitance	All pins				5	pF

1. The LED pins are as follows: P4_LED2[2:0], P3_LED1[2:0], P2_LED2[2:0], P1_LED1[2:0], P0_LED1[2:0]

Table 133: Internal Resistor Description

Pin #	Pin Name	Resistor	Pin #	Pin Name	Resistor
34, 33	SW_MODE[1:0]	Internal pull-up	82	P5_INCLK	Internal pull-high
58	EE_DOUT/ EE_IO	Internal pull-up	83, 84, 85, 86	P5_IND[3:0]	Internal pull-up
60	EE_CLK/ ADDR4	Internal pull-up	87	P5_INDV	Internal pull-down
61	EE_CS/ EE_4WIRE	Internal pull-up	89	DISABLE_MII4	Internal pull-up
64	FD_FLOW_DIS	Internal pull-up	91	P4_CRS	Internal pull-down
66	ENABLE_MII5	Internal pull-up	93	CONFIG_B	Internal pull-up
68	CLK_SEL	Internal pull-up	94	P4_COL	Internal pull-down
69	EE_DIN/ HD_FLOW_DIS	Internal pull-up	95, 96, 97, 98	P4_OUTD[3:0]/ P4_MODE[3:0]	Internal pull-up
70	MDC	Internal pull-up	99	P4_OUTCLK	Internal pull-up
71	MDIO	Internal pull-up	102	P4_INCLK	Internal pull-up
72	P5_CRS	Internal pull-down	106	CONFIG_A	Internal pull-up
73	P5_COL	Internal pull-down	108	P1_CONFIG	Internal pull-down
74, 75, 76, 77	P5_OUTD[3:0]/ P5_MODE[3:0]	Internal pull-up	111	P0_CONFIG	Internal pull-down
79	P5_OUTCLK	Internal pull-up	112, 114, 116, 117	P4_IND[3:0]	Internal pull-up
81	P5_OUTDV/ WD_DIS	Internal pull-up	118	P4_INDV	Internal pull-down

7.4.2 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 134: IEEE DC Transceiver Parameters

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
V_{ODIFF}	Absolute peak differential output voltage	TXP/N[1:0]	10BASE-T no cable	2.2	2.5	2.8	V
		TXP/N[1:0]	10BASE-T cable model	585[1]			mV
		TXP/N[0]	100BASE-FX mode	0.4	0.8	1.2	V
		TXP/N[1:0]	100BASE-TX mode	0.950	1.0	1.05	V
	Overshoot[2]	TXP/N[4:0]	100BASE-TX mode	0		5%	V
	Amplitude Symmetry (positive/ negative)	TXP/N[1:0]	100BASE-TX mode	0.98x		1.02x	V+/V-
V_{IDIFF}	Peak Differential Input Voltage accept level	RXP/N[1:0]	10BASE-T mode	585[3]			mV
		RXP/N[0] P[1:0]_SDET P/N	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 88E6065/88E6035 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 88E6065/88E6035 will reject signals typically with peak-to-peak differential amplitude less than 360 mV.

7.5 AC Electrical Specifications

7.5.1 Reset and Configuration Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified.)

Table 135: Reset and Configuration Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
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 XTAL_IN cycles prior to RESETn de-asserted		10			Cycles	
T_{SU}	Configuration data valid prior to RESETn de-asserted		200			ns	1
T_{HD}	Configuration data valid after RESETn de-asserted		0			ns	
T_{CO}	Configuration output driven after RESETn de-asserted		40			ns	2

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.
2. P[x]_OUTD[3:0]/P[x]_MODE[3:0] are normally outputs that are also used to configure the 88E6065/88E6035 device during hardware reset. When reset is asserted, these pins become inputs and the required LED configuration is latched at the rising edge of RESETn.

Figure 50: Reset and Configuration Timing

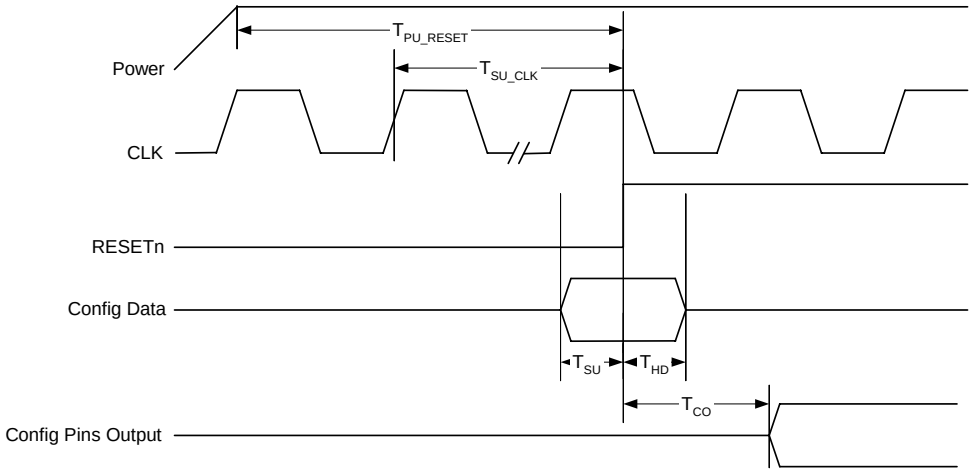

7.5.2 Clock Timing when using a 25 MHz Oscillator

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified.)

Table 136: Clock Timing with a 25 MHz Oscillator

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	XTAL_IN period		40 -50 ppm	40	40 +50 ppm	ns	25 MHz
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	

Figure 51: Oscillator Clock Timing

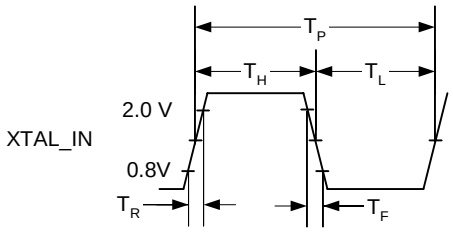

7.5.3 MII Receive Timing—PHY Mode

In PHY mode, the P[x]_INCLK pins are outputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 137: MII Receive Timing—PHY Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_TX\_CLK}$	P[x]_INCLK period	10BASE mode		400		ns	1
		100BASE mode		40		ns	1
$T_{H\_TX\_CLK}$	P[x]_INCLK high	10BASE mode	160	200	240	ns	
		100BASE mode	16	20	24	ns	
$T_{L\_TX\_CLK}$	P[x]_INCLK low	10BASE mode	160	200	240	ns	
		100BASE mode	16	20	24	ns	
T_{SU_TX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid prior to P[x]_INCLK going high.		15			ns	
T_{HD_TX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high.		0			ns	

1. 2.5 MHz for 10 Mbps or 25 MHz for 100 Mbps.

Figure 52: PHY Mode MII Receive Timing

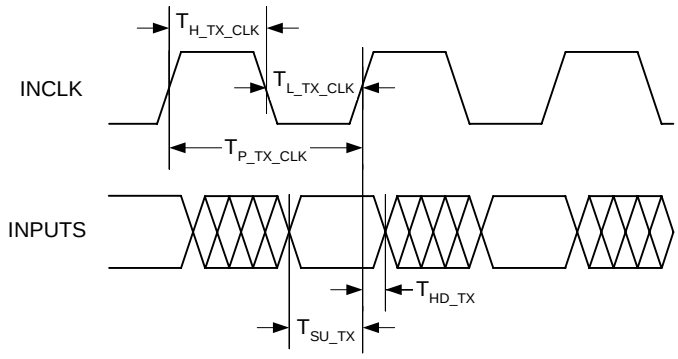

NOTE: INCLK is the clock used to clock the input data.
It is an output in this mode.

7.5.4 MII Transmit Timing—PHY Mode

In PHY mode, the P[x]_OUTCLK pins are outputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 138: MII Transmit Timing—PHY Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_RX\_CLK}$	P[x]_OUTCLK period	10BASE mode		400		ns	1
		100BASE mode		40		ns	1
$T_{H\_RX\_CLK}$	P[x]_OUTCLK high	10BASE mode	160	200	240	ns	
		100BASE mode	16	20	24	ns	
$T_{L\_RX\_CLK}$	P[x]_OUTCLK low	10BASE mode	160	200	240	ns	
		100BASE mode	16	20	24	ns	
T_{CQ_MAX}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV) valid				25	ns	
T_{CQ_MIN}	P[x]_OUTCLK to outputs P[x]_OUTD[3:0], P[x]_OUTDV) invalid		10			ns	

1. 2.5 MHz for 10 Mbps or 25 MHz for 100 Mbps.

Figure 53: PHY Mode MII Transmit Timing

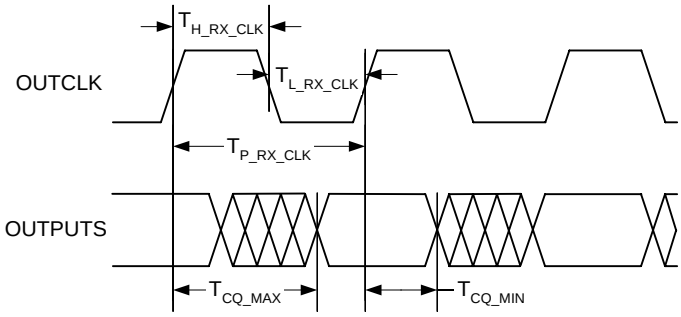

NOTE: OUTCLK is the clock used to clock the output data.
It is an output in this mode.

7.5.5 MAC Mode Clock Timing

In MAC mode, INCLK and OUTCLK are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 139: MAC Mode Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	MACCLK_IN period		0	4 or 40	40 +50 ppm	ns	DC to 25 MHz
T_H	MACCLK_IN high time		16			ns	
T_L	MACCLK_IN low time		16			ns	
T_R	MACCLK_IN rise				3	ns	
T_F	MACCLK_IN fall				3	ns	

Figure 54: MAC Clock Timing

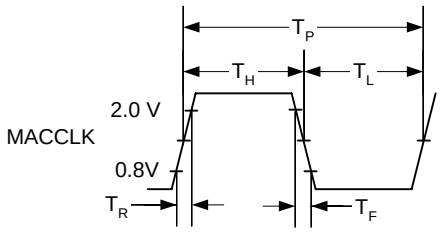

7.5.6 MII Receive Timing—MAC Mode

In MAC mode, the P[x]_INCLK pins are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 140: MII Receive Timing—MAC Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{SU_RX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid prior to P[x]_INCLK going high	With 10 pF load	10			ns	
T_{HD_RX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high	With 10 pF load	10			ns	

Figure 55: MAC Mode MII Receive Timing

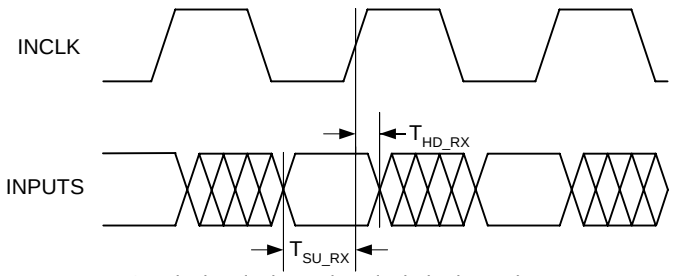

NOTE: INCLK is the clock used to clock the input data.
It is an input in this mode.

7.5.7 MII Transmit Timing—MAC Mode

In MAC mode, the P[x]_OUTCLK pins are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 141: MII Transmit Timing—MAC Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{CQ_MAX}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV) valid	With 10 pF load			25	ns	
T_{CQ_MIN}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV) invalid	With 10 pF load	0			ns	

Figure 56: MAC Mode MII Transmit Timing

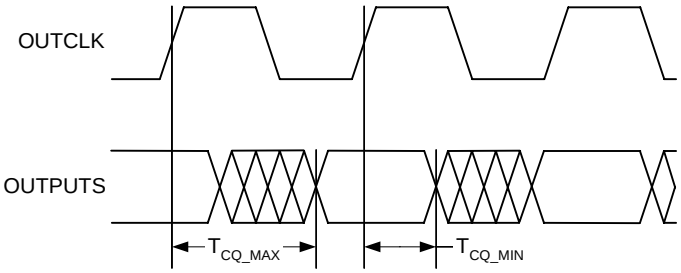

NOTE: OUTCLK is the clock used to clock the output data. It is an input in this mode.

7.5.8 MAC Mode Clock Timing, 200 Mbps

In MAC mode, INCLK and OUTCLK are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 142: MAC Mode Clock Timing, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	MACCLK_IN period		0	20	20 +50 ppm	ns	DC to 25 MHz
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	

Figure 57: MAC Clock Timing, 200 Mbps

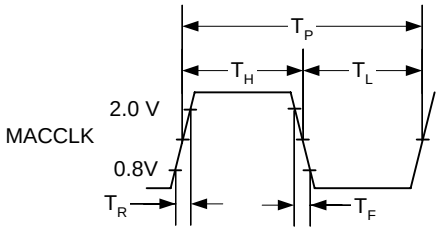

7.5.9 MII Receive Timing-PHY Mode, 200 Mbps

In PHY mode, the P[x]_INCLK pins are outputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 143: MII Receive Timing—PHY Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_TX\_CLK}$	P[x]_INCLK period	200BASE mode		20		ns	1
$T_{H\_TX\_CLK}$	P[x]_INCLK high	200BASE mode	8	10	12	ns	
$T_{L\_TX\_CLK}$	P[x]_INCLK low	200BASE mode	8	10	12	ns	
T_{SU_TX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid prior to P[x]_INCLK going high.		8.592			ns	
T_{HD_TX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high.		0			ns	

Figure 58: PHY Mode MII Receive Timing—200 Mbps

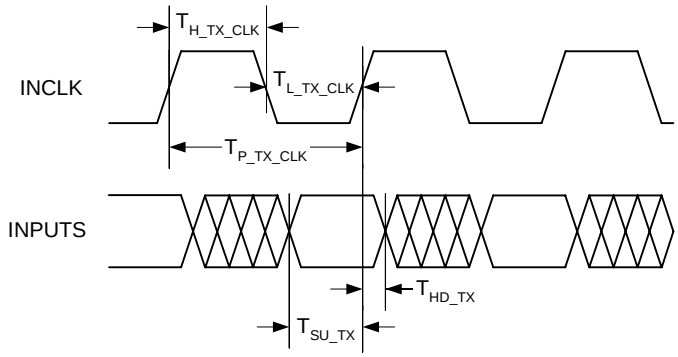

NOTE: INCLK is the clock used to clock the input data.
It is an output in this mode.

7.5.10 MII Transmit Timing-PHY Mode, 200 Mbps

In PHY mode, the P[x]_OUTCLK pins are outputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 144: MII Transmit Timing—PHY Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_RX\_CLK}$	P[x]_OUTCLK period	200BASE mode		20		ns	1
$T_{H\_RX\_CLK}$	P[x]_OUTCLK high	200BASE mode	8	10	12	ns	
$T_{L\_RX\_CLK}$	P[x]_OUTCLK low	200BASE mode	8	10	12	ns	
T_{CQ_MAX}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV) valid				15	ns	
T_{CQ_MIN}	P[x]_OUTCLK to outputs P[x]_OUTD[3:0], P[x]_OUTDV) invalid		3			ns	

1. 2.5 MHz for 10 Mbps or 25 MHz for 100 Mbps.

Figure 59: PHY Mode MII Transmit Timing—200 Mbps

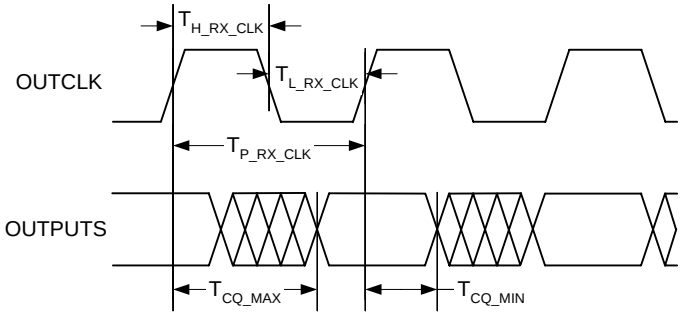

NOTE: OUTCLK is the clock used to clock the output data.
It is an output in this mode.

7.5.11 MII Receive Timing-MAC Mode, 200 Mbps

In MAC mode, the P[x]_INCLK pins are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 145: MII Receive Timing—MAC Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{SU_RX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid prior to P[x]_INCLK going high	With 10 pF load	5			ns	
T_{HD_RX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high	With 10pF load	3			ns	

Figure 60: MAC Mode MII Receive Timing—200 Mbps

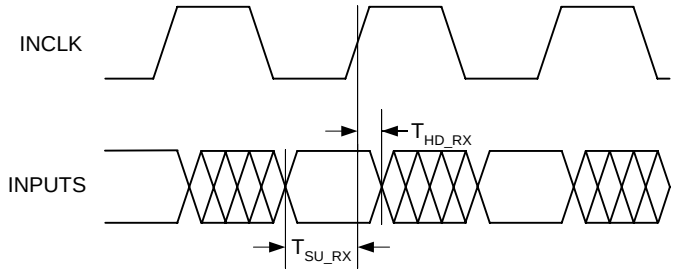

NOTE: INCLK is the clock used to clock the input data.
It is an input in this mode.

7.5.12 MII Transmit Timing-MAC Mode, 200 Mbps

In MAC mode, the P[x]_OUTCLK pins are inputs.
(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 146: MII Transmit Timing—MAC Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{CQ_MAX}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV valid	With 10 pF load			8.358	ns	
T_{CQ_MIN}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV invalid	With 10pF load	0			ns	

Figure 61: MAC Mode MII Transmit Timing, 200 Mbps

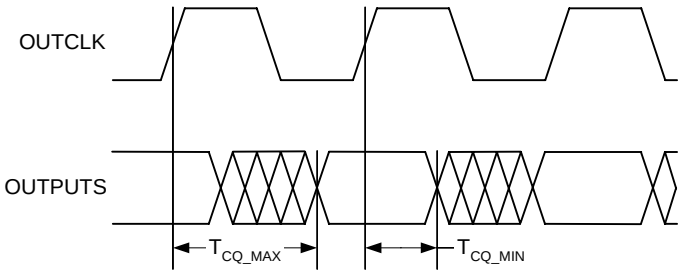

NOTE: OUTCLK is the clock used to clock the output data.
It is an input in this mode.

7.5.13 SNI Falling Edge Receive Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 147: SNI Falling Edge Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_FRX\_CLK}$	SNI falling edge INCLK period	10BASE-T Mode		100		ns	
$T_{H\_FRX\_CLK}$	SNI falling edge INCLK high		35	50	65	ns	
$T_{L\_FRX\_CLK}$	SNI falling edge INCLK low		35	50	65	ns	
T_{SU_FRX}	SNI receive data valid prior to INCLK going low		20			ns	
T_{HD_FRX}	SNI receive data valid after INCLK going low		10			ns	

Figure 62: SNI Falling Edge Receive Timing

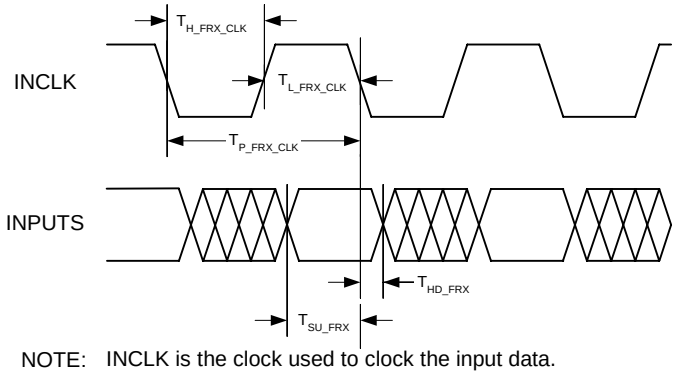

NOTE: INCLK is the clock used to clock the input data. It is an output in this mode.

7.5.14 SNI Falling Edge Transmit Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 148: SNI Falling Edge Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_FTX\_CLK}$	SNI falling edge OUTCLK period	10BASE-T Mode		100		ns	
$T_{H\_FTX\_CLK}$	SNI falling edge OUTCLK high		35	50	65	ns	
$T_{L\_FTX\_CLK}$	SNI falling edge OUTCLK low		35	50	65	ns	
T_{CQ_MXFTX}	SNI falling edge OUTCLK to output valid				65	ns	
T_{CQ_MNFTX}	SNI falling edge OUTCLK to output invalid		35			ns	

Figure 63: SNI Falling Edge Transmit Timing

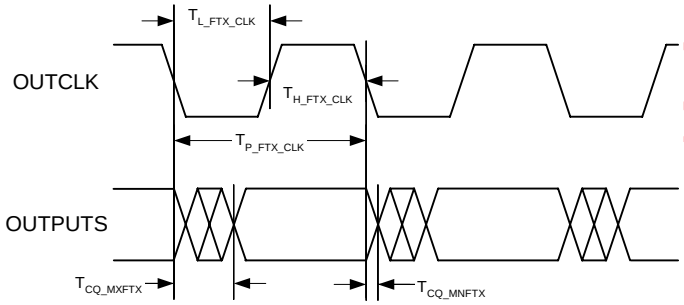

NOTE: OUTCLK is the clock used to clock the output data. It is an output in this mode.

7.5.15 SNI Rising Edge Receive Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 149: SNI Rising Edge Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_RRX\_CLK}$	SNI rising edge INCLK period	10BASE-T Mode		100		ns	
$T_{H\_RRX\_CLK}$	SNI rising edge INCLK high		35	50	65	ns	
$T_{L\_RRX\_CLK}$	SNI rising edge INCLK low		35	50	65	ns	
T_{SU_RRX}	SNI receive data valid prior to INCLK going high		20			ns	
T_{HD_RRX}	SNI receive data valid after INCLK going high		10			ns	

Figure 64: SNI Rising Edge Receive Timing

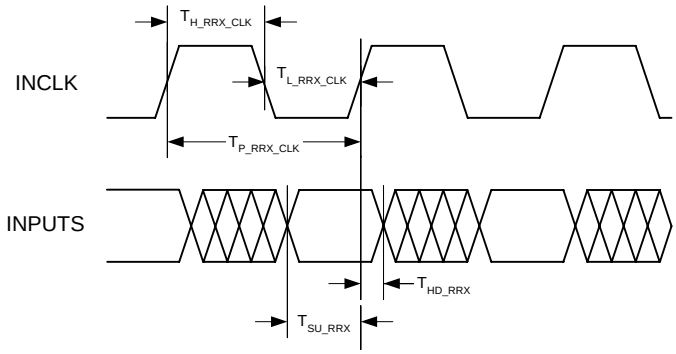

NOTE: INCLK is the clock used to clock the input data. It is an output in this mode.

7.5.16 SNI Rising Edge Transmit Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 150: SNI Rising Edge Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_RTX\_CLK}$	SNI rising edge OUTCLK period	10BASE-T Mode		100		ns	
$T_{H\_RTX\_CLK}$	SNI rising edge OUTCLK high		35	50	65	ns	
$T_{L\_RTX\_CLK}$	SNI rising edge OUTCLK low		35	50	65	ns	
T_{CQ_MXRTX}	OUTCLK to output valid				65	ns	
T_{CQ_MNRTX}	OUTCLK to output invalid		35			ns	

Figure 65: SNI Rising Edge Transmit Timing

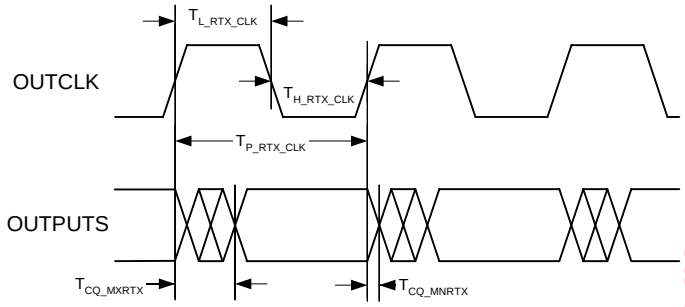

NOTE: OUTCLK is the clock used to clock the output data. It is an output in this mode.

7.5.17 RMII Receive Timing using INCLK

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 151: RMII Receive Timing using INCLK

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_TX\_CLK}$	P[x]_INCLK period	100BASE mode		20		ns	1
$T_{H\_TX\_CLK}$	P[x]_INCLK high	100BASE mode	8	10	12	ns	
$T_{L\_TX\_CLK}$	P[x]_INCLK low	100BASE mode	8	10	12	ns	
T_{SU_TX}	MII inputs (P[x]_IND[1:0], P[x]_INDV) valid prior to P[x]_INCLK going high.		9.5			ns	
T_{HD_TX}	MII inputs (P[x]_IND[1:0], P[x]_INDV) valid after P[x]_INCLK going high.		0			ns	

1. 50 MHz for 100 Mbps.

Figure 66: PHY Mode RMII Receive Timing using INCLK

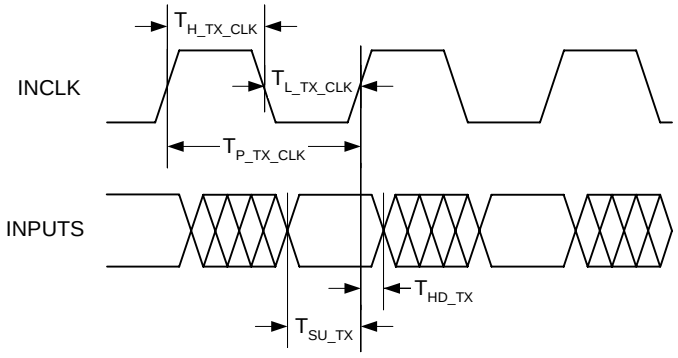

NOTE: INCLK is the clock used to clock the input data. It is an output in this mode.

7.5.18 RMII Transmit Timing using INCLK

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 152: RMII Transmit Timing using INCLK

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
$T_{P\_RX\_CLK}$	P[x]_INCLK period	100BASE mode		20		ns	1
$T_{H\_RX\_CLK}$	P[x]_INCLK high	100BASE mode	8	10	12	ns	
$T_{L\_RX\_CLK}$	P[x]_INCLK low	100BASE mode	8	10	12	ns	
T_{CQ_MAX}	P[x]_INCLK to outputs (P[x]_OUTD[1:0], P[x]_OUTDV) valid				8.0	ns	
T_{CQ_MIN}	P[x]_INCLK to outputs P[x]_OUTD[1:0], P[x]_OUTDV) invalid		0			ns	

1. 50 MHz for 100 Mbps.

Figure 67: PHY Mode RMII Transmit Timing using INCLK

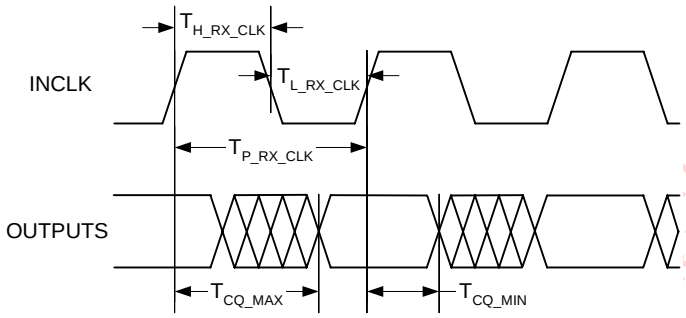

NOTE: INCLK is the clock used to clock the output data.
It is an output in this mode.

7.5.19 Serial LED Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 153: Serial LED Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_{P_LEDCLK}	LEDCLK period			80		ns	
T_{H_LEDCLK}	LEDCLK high		32	40	48	ns	
T_{L_LEDCLK}	LEDCLK low		32	40	48	ns	
T_{CQ_MAX}	LEDCLK falling edge to link outputs (LEDSER, LEDENA) valid	With 10 pF load			20	ns	
T_{CQ_MIN}	LEDCLK falling edge to link outputs (LEDSER, LEDENA) invalid	With 10 pF load	0			ns	

Figure 68: Serial LED Timing

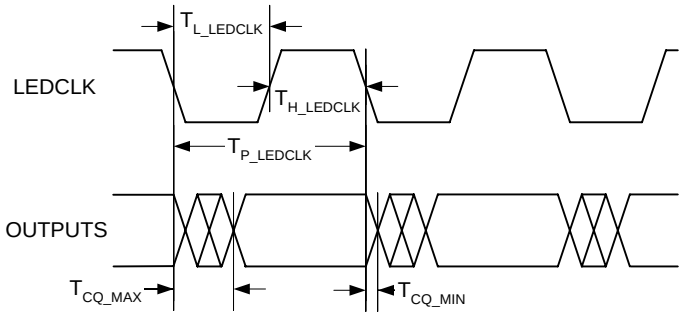

7.5.20 Serial Management Interface Clock Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 154: Serial Management Interface Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	MDCCLK_IN period		120			ns	8.33 MHz
T_H	MDCCLK_IN high time		48			ns	
T_L	MDCCLK_IN low time		48			ns	
T_R	MDCCLK_IN rise				6	ns	
T_F	MDCCLK_IN fall				6	ns	

Figure 69: Serial Management Interface Clock Timing

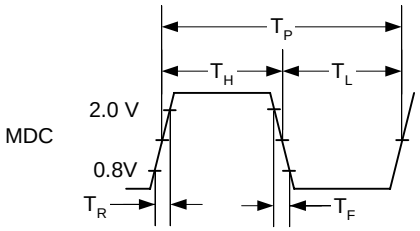

7.5.21 Serial Management Interface Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 155: Serial Management Interface Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{CQ_MDIO}	DC to MDIO (Output data valid time				20	ns
T_{DLY_MDIO}	MDC to MDIO (Output) delay time		0			ns
T_{SU}	MDIO (Input) to MDC setup time		10			ns
T_{HD}	MDIO (Input) to MDC hold time		10			ns

Figure 70: Serial Management Interface Timing

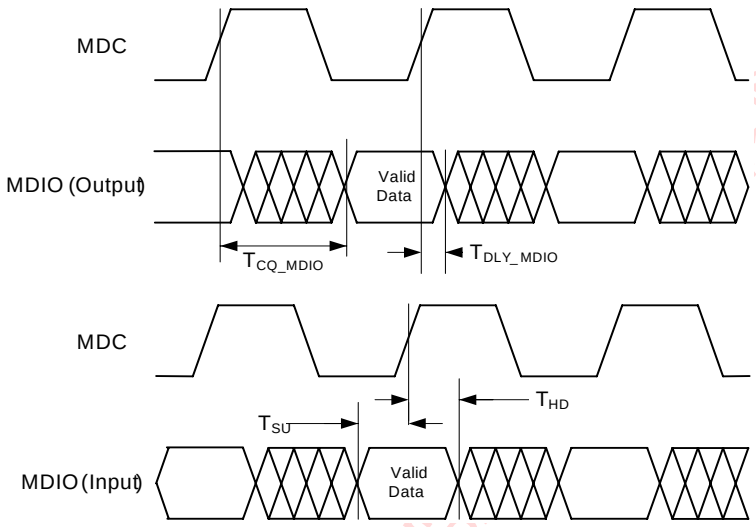

7.5.22 2-Wire EEPROM Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 156: 2-Wire EEPROM Input Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	EE_CLK period			20000		ns	
T_H	EE_CLK high time			10000		ns	
T_L	EE_CLK low time			10000		ns	
T_{IN}	EE_CLK input time		50		4500	ns	

Figure 71: 2-Wire Input Timing

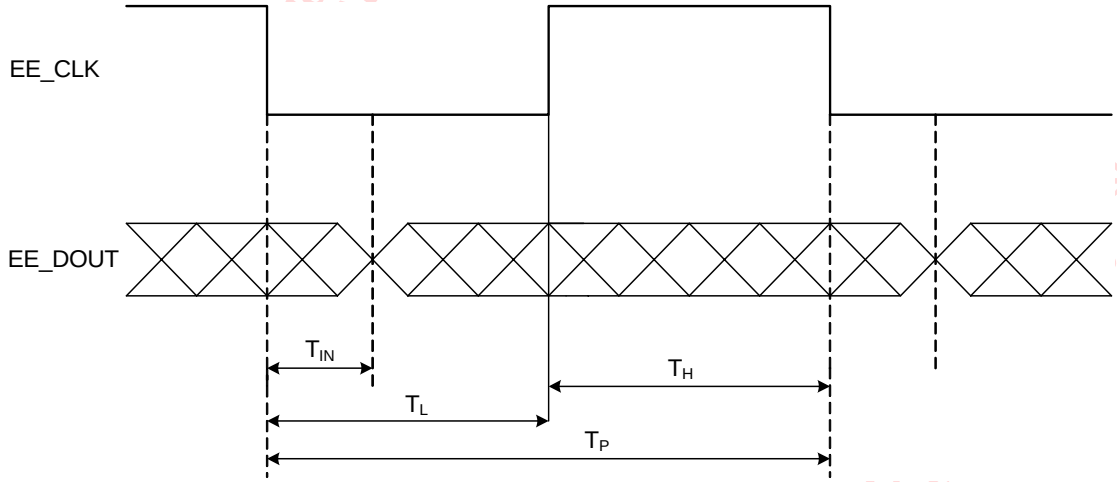

Table 157: 2-Wire EEPROM Output Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	EE_CLK period			20000		ns	
T_H	EE_CLK high time			10000		ns	
T_L	EE_CLK low time			10000		ns	
T_{IN}	EE_CLK output time		0		9800	ns	

Figure 72: 2-Wire Output Timing

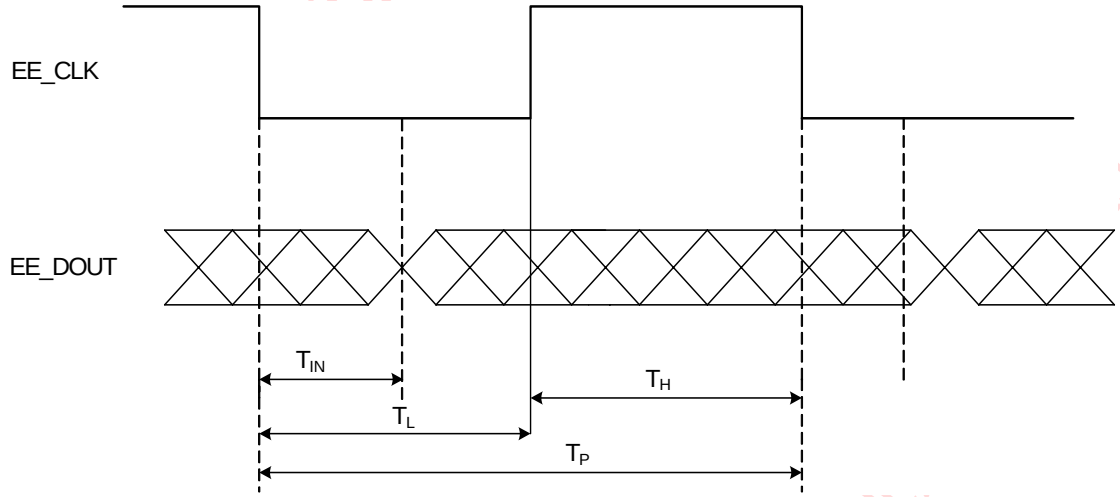

7.5.23 4-Wire EEPROM Timing

(Over full range of values listed in the Recommended Operating Conditions unless otherwise specified)

Table 158: 4-Wire EEPROM Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units	Notes
T_P	EE_CLK period			10000		ns	
T_H	EE_CLK high time			5000		ns	
T_L	EE_CLK low time			5000		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 73: EEPROM Timing

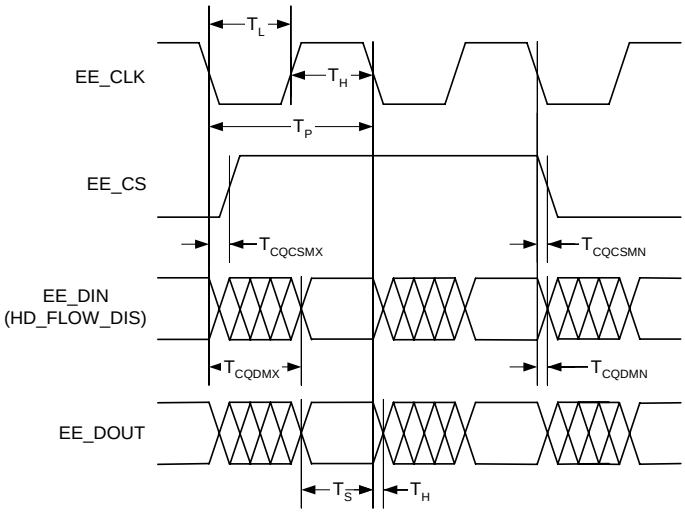

7.5.24 IEEE AC Parameters

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 159: IEEE AC Parameters

Symbol	Parameter	Pins	Condition	Min	Typ	Max	Units
T_{RISE}	Rise time	TXP/N[4:0]	100BASE-TX	3.0	4.0	5.0	ns
T_{FALL}	Fall time	TXP/N[4:0]	100BASE-TX	3.0	4.0	5.0	ns
T_{RISE}/ T_{FALL} Symmetry		TXP/N[4:0]	100BASE-TX	0		0.5	ns
DCD	Duty cycle distortion	TXP/N[4:0]	100BASE-TX	0		0.5[1]	ns, peak-peak
Transmit Jitter		TXP/N[4:0]	100BASE-TX	0		1.4	ns, peak-peak

1. ANSI X3.263-1995 Figure 9-3.

Section 8. Mechanical Drawings

8.1 88E6065/88E6035 128-Pin LQFP Package Mechanical Drawing

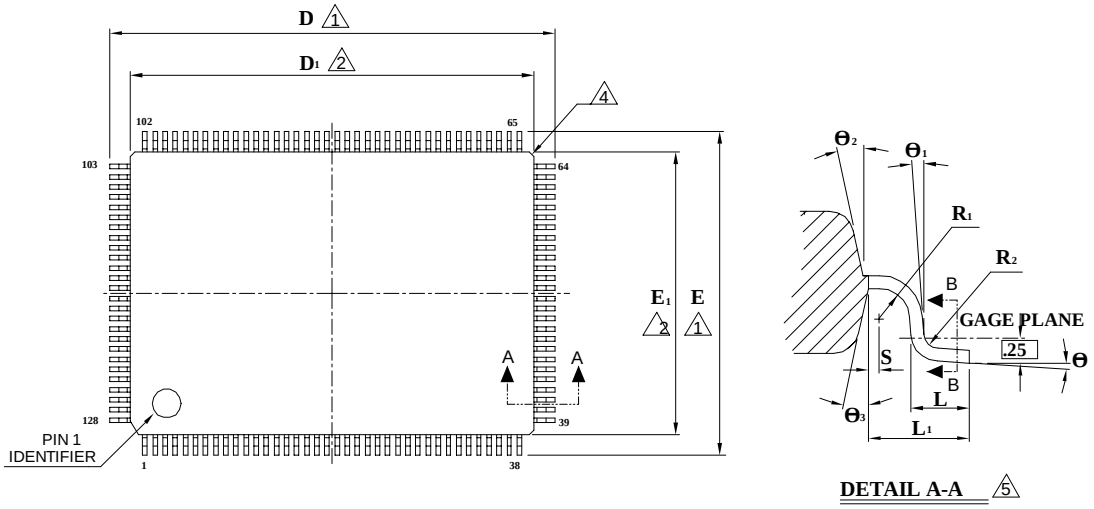

DETAIL A-A 5

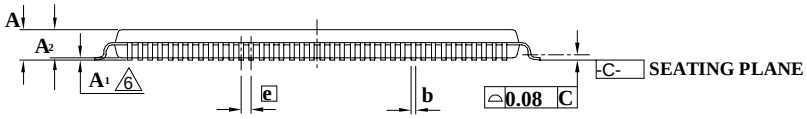

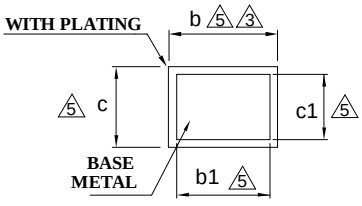

DETAIL B-B

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 MISMACH.
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.

8.2 Dimensions in mm.

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	0.17	0.22	0.27
b_1	0.17	0.20	0.23
c	0.09	—	0.20
c_1	0.09	—	0.16
D	21.90	22.00	22.10
D_1	19.90	20.00	20.10
E	15.90	16.00	16.10
E_1	13.90	14.00	14.10
e	0.50 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	—	—
θ	0°	3.5°	7°
$θ_1$	4° TYP		
$θ_2$	12°TYP		
$θ_3$	12°TYP		

Section 9. Ordering Information

9.1 Ordering Part Numbers and Package Markings

Figure 74 shows the ordering part numbering scheme for the devices. Contact Marvell® FAEs or sales representatives for complete ordering information.

Figure 74: Sample Part Number

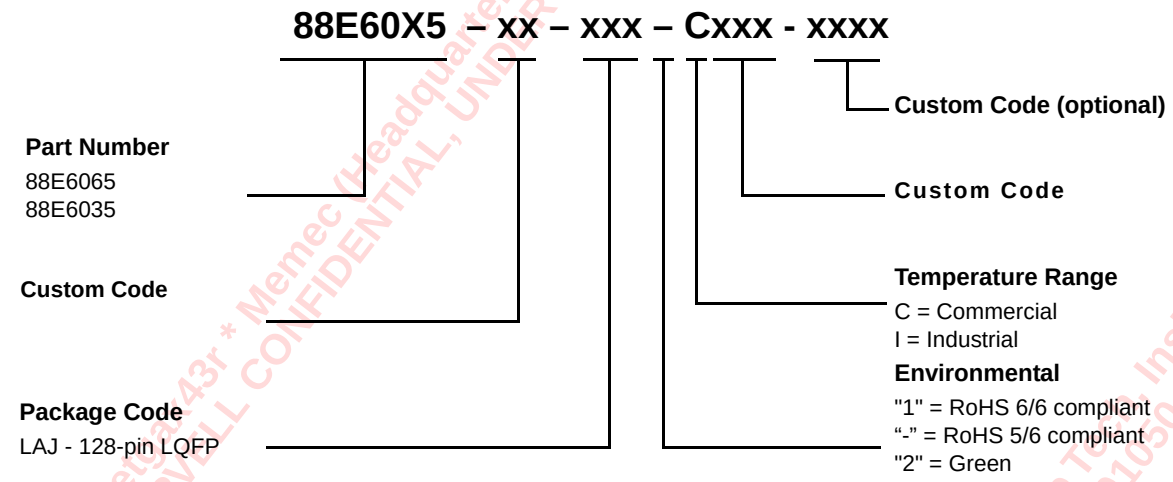

Table 160: Part Order Options - RoHS 6/6 Compliant (Lead Free)

Package Type	Part Order Number
88E6065 128-pin LQFP - Commercial Temp	88E6065-XX-LAJ1C000
88E6065 128-pin LQFP - Industrial Temp	88E6065-XX1LAJ1I000
88E6035 128-pin LQFP - Commercial Temp	88E6035-XX-LAJ1C000

Table 161: Part Order Options - RoHS 5/6 Compliant (Non Lead Free)

Package Type	Part Order Number
88E6065 128-pin LQFP - Commercial Temp	88E6065-XX-LAJ-C000
88E6065 128-pin LQFP - Industrial Temp	88E6065-XX1LAJ-I000
88E6035 128-pin LQFP - Commercial Temp	88E6035-XX-LAJ-C000

Table 162: Part Order Options - Green

Package Type	Part Order Number
88E6065 128-pin LQFP - Commercial Temp	88E6065-XX-LAJ2C000
88E6065 128-pin LQFP - Industrial Temp	88E6065-XX1LAJ2I000
88E6035 128-pin LQFP - Commercial Temp	88E6035-XX-LAJ2C000

9.1.1 Package Marking Examples

Figure 75 is an example of the package marking and pin 1 location for the 88E6065 128-pin LQFP Commercial package. Markings for the other variants are similar.

Figure 75: 88E6065 128-pin LQFP Commercial Package Marking and Pin 1 Location

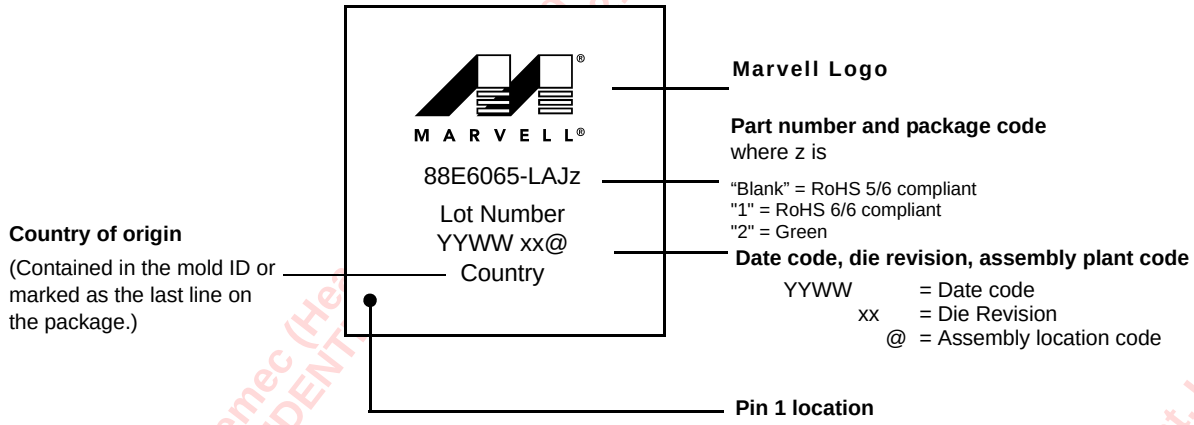

Note: The above example is not drawn to scale. Location of markings is approximate.

Figure 76 is an example of the package marking and pin 1 location for the 88E6065 128-pin LQFP Industrial package. Markings for the other variants are similar.

Figure 76: 88E6065 128-pin LQFP Industrial Package Marking and Pin 1 Location

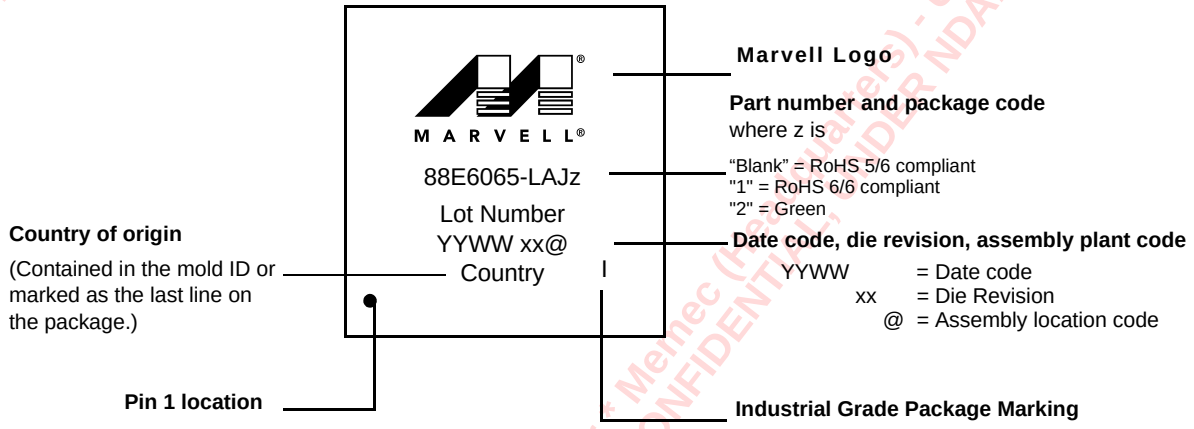

Note: The above example is not drawn to scale. Location of markings is approximate.

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