



SWITCHING

Link Street™ 88E6083 Datasheet

Integrated 10-Port 10/100 Ethernet Switch with
QoS, RAM, Transceivers

Doc. No. MV-S100968-00 Rev. B
April 28, 2005

MOVING FORWARD
FASTER®





Link Street™ 88E6083
Integrated 10-Port 10/100 Ethernet Switch with QoS, RAM, Transceivers

Document Status

Advanced 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. B
Preliminary	Technical Publications: 2.0

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose, without the express written permission of Marvell. Marvell retains the right to make changes to this document at any time, without notice. Marvell makes no warranty of any kind, expressed or implied, with regard to any information contained in this document, including, but not limited to, the implied warranties of merchantability or fitness for any particular purpose. Further, Marvell does not warrant the accuracy or completeness of the information, text, graphics, or other items contained within this document.

Marvell products are not designed for use in life-support equipment or applications that would cause a life-threatening situation if any such products failed. Do not use Marvell products in these types of equipment or applications.

With respect to the products described herein, the user or recipient, in the absence of appropriate U.S. government authorization, agrees:

- 1) Not to re-export or release any such information consisting of technology, software or source code controlled for national security reasons by the U.S. Export Control Regulations ("EAR"), to a national of EAR Country Groups D:1 or E:2;
- 2) Not to export the direct product of such technology or such software, to EAR Country Groups D:1 or E:2, if such technology or software and direct products thereof are controlled for national security reasons by the EAR; and,
- 3) In the case of technology controlled for national security reasons under the EAR where the direct product of the technology is a complete plant or component of a plant, not to export to EAR Country Groups D:1 or E:2 the direct product of the plant or major component thereof, if such direct product is controlled for national security reasons by the EAR, or is subject to controls under the U.S. Munitions List ("USML").

At all times hereunder, the recipient of any such information agrees that they shall be deemed to have manually signed this document in connection with their receipt of any such information. Copyright © 2005, Marvell International Ltd. All rights reserved. Marvell, the Marvell logo, Moving Forward Faster, Alaska, Fastwriter, GalNet, Libertas, Link Street, NetGX, PHYAdvantage, Pretera, Virtual Cable Tester, and Yukon are registered trademarks of Marvell. AnyVoltage, Discovery, DSP Switcher, Feroceon, GalTis, Horizon, RADLAN, Raising The Technology Bar, The Technology Within, UniMAC, and VCT are trademarks of Marvell. All other trademarks are the property of their respective owners.



Link Street™ 88E6083 Integrated 10-Port 10/100 Ethernet Switch with QoS, RAM, Transceivers

OVERVIEW

The 88E6083 device is a single chip integrated 10-port Fast Ethernet switch with Quality of Service (QoS) support, integrating eight PHY ports and two MII ports. The device contains ten independent Fast Ethernet media access controllers (MACs); eight 10BASE-T/100BASE-TX copper PHY transceivers, all eight of which can be configured as 100BASE-FX fiber; a high-speed non-blocking QoS switch fabric with four traffic classes, a high-performance address lookup engine, and a 1 Mbit frame buffer memory. The 88E6083 Fast Ethernet switch integrates special support for WAN/LAN firewall isolation, Quality of Service, dynamic VLANs, and SNMP/RMON.

The 88E6083 device supports a CPU header mode, which prepends two bytes to the frame on the CPU's port, aligning the IP header information on a 32-bit boundary, thus accelerating router performance. The header also gives the CPU wire-speed port VLAN control without needing to use 802.1Q tagged frames, so routing can be performed on all frames without exception.

The PHY transceivers are designed with Marvell® Virtual Cable Tester® (VCT™) technology for advanced cable diagnostics. VCT enables IT managers to easily pinpoint the location of cabling issues down to a meter or less, reducing network installation and support costs.

The 88E6083 device utilizes the latest Marvell® Quality of Service switch architecture with non-blocking switching performance in all traffic environments. Packets are directed into one of four traffic class priority queues based on Port priority, IEEE 802.1p, IPv4's TOS or Diff-Serv or IPv6's Traffic Class, and MAC address. The 88E6083 device also integrates IGMP snooping capability.

The 88E6083 device supports truly isolated WAN vs. LAN firewall applications. Ports 9 and 10 are Media Independent Interfaces (MIIs) that support direct connection to a Router or Management CPU. These can be configured in either RMII, MII PHY or MAC Mode, or SNI. This interface, along with BPDU handling, programmable per port VLAN configurations, and Port States, supports Spanning Tree and truly isolated WAN vs. LAN firewall.

The 88E6083 device includes complete IEEE 802.1Q VLAN ID processing per port, to process dynamic VLAN membership and VLAN tagging. The device also integrates extensive RMON network management capability, including 40 RMON statistics counters per port (twenty 32-bit counters each for ingress and egress on every port).

The 88E6083 device supports up to 64 802.1Q 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 88E6083 device supports multiple address data bases, which allows packet routing without modification of the MAC address. This allows the same MAC address to exist multiple times in the MAC Address data base with multiple port mappings, to completely isolate the WAN from the LAN data base.

The PHYs in the 88E6083 device 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 88E6083 device comply fully with the applicable sections of IEEE 802.3, IEEE 802.3u, and IEEE 802.3x standards.

The 88E6083 device is designed to work in all cable environments. True Plug-n-Play is supported with Auto-Crossover, Auto-Polarity and Auto-Negotiation in all the PHYs, along with network loop prevention. The many operating modes of the 88E6083 device can be configured using SMI (serial management interface - MDC/MDIO) and/or a low cost serial EEPROM (93C46, C56 or C66). A stand-alone QoS mode is also supported. Each PHY supports Marvell's Virtual Cable Tester™ (VCT™) feature for testing cables and connectors using TDR.

FEATURES

- Single chip integration of a 10-port Fast Ethernet QoS switch, eight PHY ports plus 2 MII ports, in a 216-pin LQFP package
- Integrates ten independent media access controllers (MACs) fully compliant with the applicable sections of IEEE802.3, IEEE802.3u and IEEE802.3x
- Integrates eight independent Fast Ethernet transceivers (PHYs) fully compliant with the applicable sections of IEEE802.3 and IEEE802.3u
- Automatic MDI/MDIX crossover for 100BASE-TX and 10BASE-T ports
- All eight PHY ports can be configured as copper ports (100BASE-TX or 10BASE-T) or fiber (100BASE-FX)
- QoS support with four traffic classes, typically supporting prioritized packet streams for management, voice, video, and data

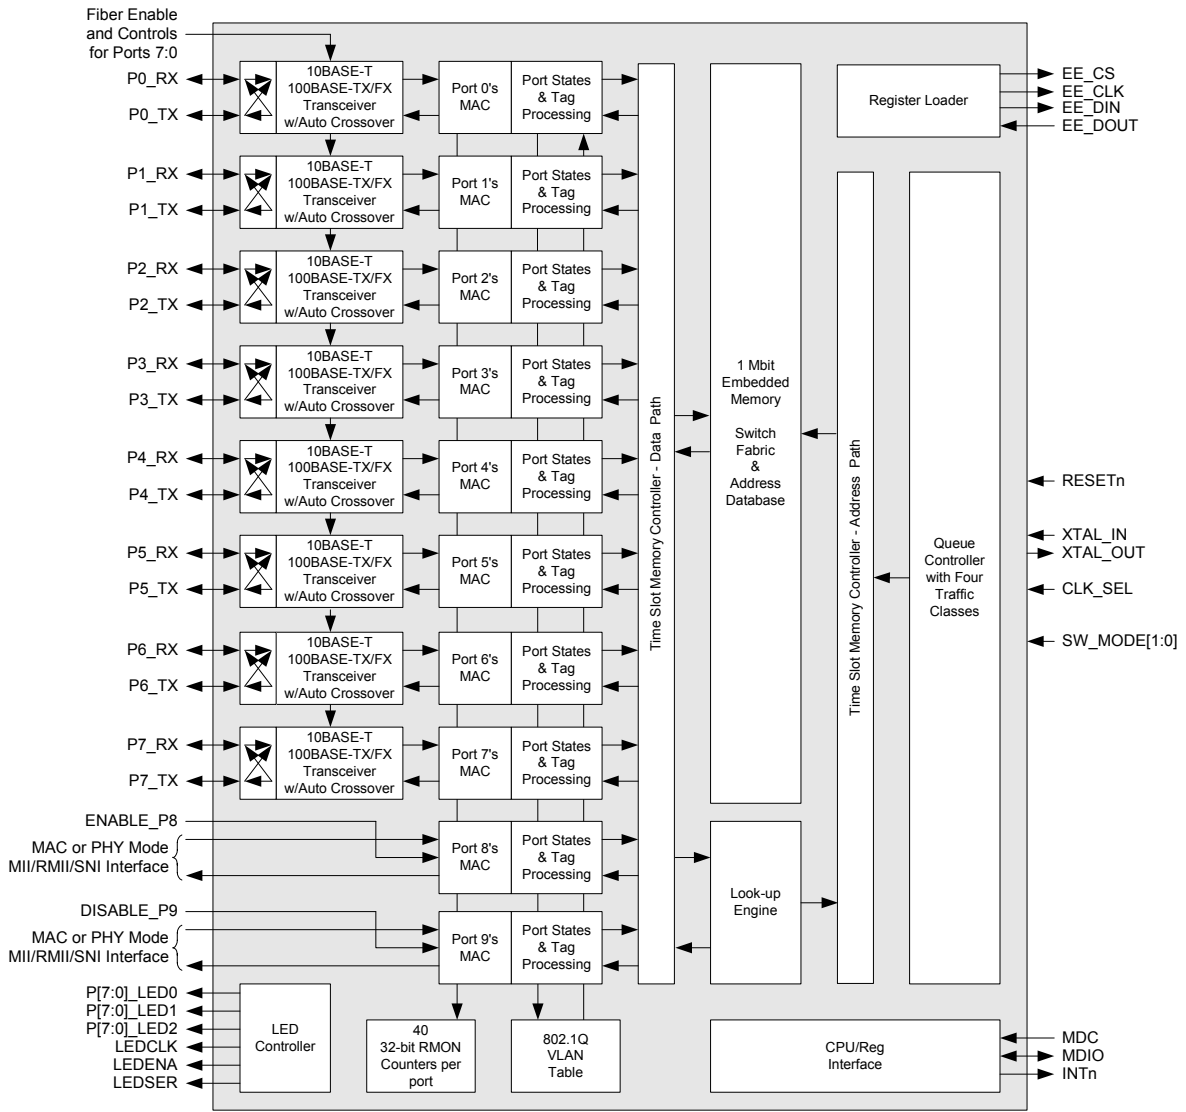


- Full IEEE 802.1Q VLAN ID processing per port, with dynamic VLAN membership, and VLAN tagging for up to 64 VLANs
- High performance lookup engine with support for 2048 MAC address entries with address lock options, automatic learning and aging
- Extensive RMON Statistics Counters, 40 Stat Counters Per Port (20 Each for Ingress and Egress)
- QoS determined by Port, IEEE 802.1p tagged frames, IPv4's Type of Service (TOS) & Differentiated Services (DS) or IPv6's Traffic Class, DA MAC address
- IGMP Snooping
- Fixed priority and weighted fair queuing
- Port based VLANs supported in any combination
- Virtual Cable Tester® (VCT™) capability
- MII ports support glueless interface to CPUs for management and firewall applications
- Flexible LED support for Link, Speed, Duplex Mode, Collision, and Tx/Rx Activities
- Supports a low cost 25 MHz XTAL clock source
- CMOS low power dissipation

- Available with Commercial or Industrial grade temperature specifications

APPLICATIONS

- Firewall gateway router with QoS support for designs with seven 10/100BASE-T LAN ports & one 10/100BASE-T WAN port with an extra MII port for an optional wireless LAN network connection
- Firewall gateway router with QoS support for designs with eight 10/100BASE-T LAN ports & an MII port for a WAN PHY
- Eight port QoS switch with Spanning Tree Support
- Firewall gateway router supporting a Fiber WAN port
- Multi-Dwelling Unit Interface gateway



88E6083 Top Level Block Diagram



Table of Contents

SECTION 1. SIGNAL DESCRIPTION	16
1.1 88E6083 216-Pin LQFP Package	16
1.2 Pin Description	17
SECTION 2. APPLICATION EXAMPLES	40
2.1 Configuration Options for the 88E6083 Devices	40
2.2 Example Configurations	40
2.3 Routing with the Marvell® Header	43
SECTION 3. FUNCTIONAL DESCRIPTION	44
3.1 Switch Data Flow	44
3.2 MII/SNI/RMII	45
3.2.1 MII PHY Mode	45
3.2.2 MII MAC Mode.....	46
3.2.3 SNI PHY Mode	47
3.2.4 RMII PHY Mode.....	48
3.2.5 MII/SNI Configuration	49
3.2.6 Enabling the MII/SNI Interfaces.....	49
3.2.7 Port Status Registers.....	49
3.2.8 MII 200 Mbps Mode.....	50
3.3 Media Access Controllers (MAC).....	50
3.3.1 Backoff.....	51
3.3.2 Half-Duplex Flow Control.....	51
3.3.3 Full-Duplex Flow Control	51
3.3.4 Forcing Flow Control	52
3.3.5 Statistics Counters.....	53
3.3.6 Extensive RMON/Statistics Counters	53
3.4 Address Management	58
3.4.1 Address Translation Unit	58
3.4.2 Address Searching or Translation	59
3.4.3 Address Learning	60
3.4.4 Address Aging	60
3.4.5 Address Translation Unit Operations.....	61
3.5 Ingress Policy	65
3.5.1 Forward Unknown/Secure Port	65
3.5.2 Quality of Service (QoS) Classification.....	66
3.5.3 VLANs	69
3.5.4 VLAN Translation Unit Operations	73

3.5.5	Switching Frames Back to their Source Port.....	78
3.5.6	IGMP Snooping	79
3.5.7	Port States.....	80
3.5.8	Ingress Double VLAN Tagging.....	81
3.5.9	Ingress Rate Limiting.....	82
3.5.10	Switch's Ingress Header.....	82
3.5.11	Switch's Ingress Trailer	84
3.6	Queue Controller.....	86
3.6.1	Frame Latencies.....	86
3.6.2	No Head-of-Line Blocking	86
3.6.3	QoS with and without Flow Control	86
3.6.4	Guaranteed Frame Delivery without Flow Control	87
3.6.5	Fixed or Weighted Priority	87
3.6.6	The Queues.....	88
3.6.7	Queue Manager	88
3.6.8	Output Queues	89
3.6.9	Multicast Handler.....	89
3.7	Egress Policy	90
3.7.1	Tagging and Untagging Frames	90
3.7.2	Priority Override	91
3.7.3	VID 0x000 and VID Override.....	91
3.7.4	Egress Double VLAN Tagging	92
3.7.5	Egress Rate Limiting	93
3.7.6	Switch's Egress Header	94
3.7.7	Switch's Egress Trailer.....	96
3.8	Spanning Tree Support	97
3.9	Embedded Memory.....	97
3.10	Interrupt Controller	98
3.11	Port Monitoring Support	98
3.12	Port Trunking Support.....	98
SECTION 4. PHYSICAL INTERFACE (PHY) FUNCTIONAL DESCRIPTION		99
4.1	Transmit PCS and PMA.....	102
4.1.1	100BASE-TX Transmitter.....	102
4.1.2	4B/5B Encoding.....	102
4.1.3	Scrambler.....	102
4.1.4	NRZ to NRZI Conversion	102
4.1.5	Pre-Driver and Transmit Clock.....	102
4.1.6	Multimode Transmit DAC	102
4.2	Receive PCS and PMA.....	103
4.2.1	10-BASE-T/100BASE-TX Receiver.....	103
4.2.2	AGC and Baseline Wander	103
4.2.3	ADC and Digital Adaptive Equalizer	103
4.2.4	Digital Phased Locked Loop (DPLL)	103



4.2.5	NRZI to NRZ Conversion.....	103
4.2.6	Descrambler	103
4.2.7	Serial-to-Parallel Conversion and 5B/4B Code-Group Alignment	104
4.2.8	5B/4B Decoder	104
4.2.9	Setting Cable Characteristics	106
4.2.10	Scrambler/Descrambler	106
4.2.11	Digital Clock Recovery/Generator	106
4.2.12	Link Monitor	107
4.2.13	Auto-Negotiation.....	107
4.2.14	Register Update.....	108
4.2.15	Next Page Support	108
4.2.16	Status Registers	108
4.3	Power Management.....	109
4.3.1	Low Power Modes	109
4.3.2	MAC Interface and PHY Configuration for Low Power Modes	109
4.3.3	IEEE Power Down Mode	110
4.4	Far End Fault Indication (FEFI)	111
4.5	Virtual Cable Tester®	112
4.6	Auto MDI/MDIX Crossover	113
4.7	LED Interface	114
4.7.1	Parallel LED Interface.....	114
4.7.2	Serial LED Interface	116
SECTION 5.	REGISTER DESCRIPTION	121
5.1	Register Types.....	122
5.2	Switch Core Registers	122
5.2.1	Switch Core Register Map.....	123
5.2.2	Switch Port Registers	125
5.2.3	Switch Global Registers	137
SECTION 6.	SERIAL MANAGEMENT INTERFACE (SMI)	155
6.1	MDC/MDIO Read and Write Operations.....	155
SECTION 7.	PHY REGISTERS.....	157
SECTION 8.	EEPROM PROGRAMMING FORMAT	188
8.1	EEPROM Programming Details.....	188
SECTION 9.	ELECTRICAL SPECIFICATIONS	190
9.1	Absolute Maximum Ratings	190

9.2	Recommended Operating Conditions	190
9.3	Package Thermal Information.....	191
9.3.1	Thermal Conditions for 216-pin LQFP Package.....	191
9.3.2	Current Consumption	192
9.3.3	Digital Operating Conditions.....	193
9.3.4	IEEE DC Transceiver Parameters.....	195
9.4	AC Electrical Specifications	196
9.4.1	Reset and Configuration Timing	196
9.4.2	Clock Timing with a 25 MHz Oscillator.....	197
9.4.3	MII Receive Timing - PHY Mode	198
9.4.4	MII Transmit Timing - PHY Mode	199
9.4.5	MAC Mode Clock Timing.....	200
9.4.6	MII Receive Timing - MAC Mode	201
9.4.7	MII Transmit Timing - MAC Mode	202
9.4.8	MAC Mode Clock Timing, 200 Mbps	203
9.4.9	MII Receive Timing - PHY Mode, 200 Mbps	204
9.4.10	MII Transmit Timing-PHY Mode, 200 Mbps	205
9.4.11	MII Receive Timing-MAC Mode, 200 Mbps.....	206
9.4.12	MII Transmit Timing—MAC Mode, 200 Mbps	207
9.4.13	SNI Falling Edge Receive Timing.....	208
9.4.14	SNI Falling Edge Transmit Timing.....	209
9.4.15	SNI Rising Edge Receive Timing	210
9.4.16	SNI Rising Edge Transmit Timing	211
9.4.17	RMII Receive Timing	212
9.4.18	RMII Transmit Timing	213
9.4.19	Serial LED Timing	214
9.4.20	Serial Management Interface Clock Timing	215
9.4.21	Serial Management Interface Timing	216
9.4.22	EEPROM Timing	217
9.4.23	IEEE AC Parameters.....	218
SECTION 10.	PACKAGE MECHANICAL DIMENSIONS	219
SECTION 11.	ORDERING PART NUMBERS AND PACKAGE MARKINGS	221



List of Tables

Table 1: Pin Type Definitions.....	17
Table 2: Network Interface	18
Table 3: PHY Configuration.....	19
Table 4: Regulator and Reference	21
Table 5: System.....	21
Table 6: Register Access Interface.....	22
Table 7: Serial EEPROM Interface.....	23
Table 8: Port 8's Enable	24
Table 9: Port 8's Input MII - If ENABLE_P8 = High	25
Table 10: Port 8's Output MII - If ENABLE_P8 = High	26
Table 11: Port 9's Enable	27
Table 12: Port 9's Input MII - If DISABLE_P9 = Low	28
Table 13: Port 9's Output MII - If DISABLE_P9 = Low	29
Table 14: Switch Configuration Interface.....	31
Table 15: Port Status LEDs	32
Table 16: Power and Ground	33
Table 17: Package Pin List—Alphabetical by Signal Name	36
Table 18: RMII/MII/SNII Configuration Options	49
Table 19: PAUSE Frame Format.....	52
Table 20: Ingress Statistics Counters.....	54
Table 21: Egress Statistics Counters	56
Table 22: ATU Operations Registers.....	61
Table 23: ATU Data Fields	62
Table 24: ATU Get Next Operation Register Usage.....	63
Table 25: ATU Load/Purge Operation Register Usage	64
Table 26: VLAN Table Settings for Figure 16	70
Table 27: VLAN Table Settings for Figure 17	71
Table 28: VTU Operations Registers.....	74
Table 29: VTU Entry Format.....	75
Table 30: VTU Get Next Operation Register Usage.....	76
Table 31: VTU Load/Purge Operation Register Usage	77
Table 32: VTU Get/Clear Violation Data Register Usage.....	78
Table 33: Port State Options	80
Table 34: Ingress Trailer Fields	85
Table 35: Egress Port Configuration	90
Table 36: Egress Header Fields	95
Table 37: Egress Trailer Fields.....	97
Table 38: 5B/4B Code Mapping	105
Table 39: Scrambler Settings	106
Table 40: Operating Mode Power Consumption	109

Table 41: FEFI Select.....	111
Table 42: MDI/MDIX Pin Functions	113
Table 43: Parallel LED Hardware Defaults.....	114
Table 44: Parallel LED Display Interpretation.....	115
Table 45: Serial LED Display Options (A = Active)	117
Table 46: Single LED Display Mode.....	118
Table 47: Dual LED Display Mode	119
Table 48: Register Types	122
Table 49: Switch Port Register (SMI Device Address) Map.....	123
Table 50: Switch Global Registers	124
Table 51: Port Status Register	125
Table 52: Switch Identifier Register.....	126
Table 53: Port Control Register	127
Table 54: Port Based VLAN Map	131
Table 55: Default Port VLAN ID & Priority.....	132
Table 56: Rate Control	133
Table 57: Port Association Vector	135
Table 58: Rx Counter	136
Table 59: Tx Counter.....	136
Table 60: Switch Global Status Register	137
Table 61: Switch MAC Address Register Bytes 0 & 1	138
Table 62: Switch MAC Address Register Bytes 2 & 3.....	138
Table 63: Switch MAC Address Register Bytes 4 & 5.....	138
Table 64: Switch Global Control Register.....	139
Table 65: VTU Operation Register	140
Table 66: VTU VID Register.....	142
Table 67: VTU Data Register Ports 0 to 3.....	142
Table 68: VTU Data Register Port 4 to 7.....	143
Table 69: VTU Data Register Port 8 to 9.....	143
Table 70: ATU Control Register	144
Table 71: ATU Operation Register	145
Table 72: ATU Data Register	146
Table 73: ATU Switch MAC Address Register Bytes 0 & 1.....	146
Table 74: ATU Switch MAC Address Register Bytes 2 & 3.....	146
Table 75: ATU Switch MAC Address Register Bytes 4 & 5.....	147
Table 76: IP-PRI Mapping Register 0.....	148
Table 77: IP-PRI Mapping Register 1.....	148
Table 78: IP-PRI Mapping Register 2.....	149
Table 79: IP-PRI Mapping Register 3.....	149
Table 80: IP-PRI Mapping Register 4.....	150
Table 81: IP-PRI Mapping Register 5.....	150
Table 82: IP-PRI Mapping Register 6.....	151
Table 83: IP-PRI Mapping Register 7.....	151
Table 84: IEEE-PRI Register.....	152



Table 85: Statistics Operation Register	152
Table 86: Stats Counter Register Bytes, 3&2	154
Table 87: Stats Counter Register Bytes, 1&0	154
Table 88: Serial Management Interface Protocol Example	156
Table 89: PHY Register Map	157
Table 90: Register Types	158
Table 91: PHY Control Register	159
Table 92: PHY Status Register	162
Table 93: PHY Identifier	163
Table 94: PHY Identifier	163
Table 95: Auto-Negotiation Advertisement Register	164
Table 96: Link Partner Ability Register (Base Page)	166
Table 97: Link Partner Ability Register (Next Page)	167
Table 98: Auto-Negotiation Expansion Register	168
Table 99: Next Page Transmit Register	169
Table 100: Link Partner Next Page Register	169
Table 101: PHY Specific Control Register	170
Table 102: PHY Specific Status Register	172
Table 103: PHY Interrupt Enable	174
Table 104: PHY Interrupt Status	175
Table 105: PHY Interrupt Port Summary (Global)	177
Table 106: Receive Error Counter	178
Table 107: LED Parallel Select Register (Global)	179
Table 108: LED Stream Select for Serial LEDs (Global)	180
Table 109: PHY LED Control Register (Global)	182
Table 110: PHY Manual LED Override	184
Table 111: VCT™ Register for TXP/N Pins	185
Table 112: VCT™ Register for RXP/N pins	186
Table 113: PHY Specific Control Register II	187
Table 114: Internal Resistor Description	194
Table 115: IEEE DC Transceiver Parameters	195
Table 116: Reset and Configuration Timing	196
Table 117: Clock Timing with a 25 MHz Oscillator	197
Table 118: MII Receive Timing-PHY Mode	198
Table 119: MII Transmit Timing-PHY Mode	199
Table 120: MAC Mode Clock Timing	200
Table 121: MII Receive Timing-MAC Mode	201
Table 122: MII Transmit Timing-MAC Mode	202
Table 123: MAC Mode Clock Timing, 200 Mbps	203
Table 124: MII Receive Timing - PHY Mode, 200 Mbps	204
Table 125: MII Transmit Timing-PHY Mode, 200 Mbps	205
Table 126: MII Receive Timing-MAC Mode, 200 Mbps	206
Table 127: MII Transmit Timing—MAC Mode, 200 Mbps	207
Table 128: SNI Falling Edge Receive Timing	208

Table 129: SNI Falling Edge Transmit Timing.....	209
Table 130: SNI Rising Edge Receive Timing	210
Table 131: SNI Rising Edge Transmit Timing	211
Table 132: RMI Receive Timing	212
Table 133: RMI Transmit Timing	213
Table 134: Serial LED Timing	214
Table 135: Serial Management Interface Clock Timing.....	215
Table 136: Serial Management Interface Timing.....	216
Table 137: EEPROM Timing	217
Table 138: IEEE AC Parameters.....	218



List of Figures

Figure 1:	88E6083 216-Pin LQFP Package (Top View)	16
Figure 2:	Firewall Router Switch supporting a Fiber WAN port	41
Figure 3:	88E6083 Firewall Router with Wireless Access Point	42
Figure 4:	88E6083 Device Switch Data Flow	44
Figure 5:	Switch Operation	45
Figure 6:	MII PHY Interface Pins	45
Figure 7:	MII MAC Interface Pins	46
Figure 8:	SNI PHY Interface Pins	47
Figure 9:	RMII PHY Interface Pins using INCLK	48
Figure 10:	ATU Size Tradeoffs	59
Figure 11:	Format of an ATU Entry	62
Figure 12:	IEEE Tag Frame Format	67
Figure 13:	IPv4 Frame Format	68
Figure 14:	IPv6 Frame Format	68
Figure 15:	Switch Operation with VLANs Disabled	69
Figure 16:	Switch Operation with a Typical Router VLAN Configuration	70
Figure 17:	Switch Operation with another Example VLAN Configuration	71
Figure 18:	Format of a VTU Entry	74
Figure 19:	IGMP Snoop Format	79
Figure 20:	Double Tag Format	81
Figure 21:	Ingress Header Format	83
Figure 22:	Ingress Trailer Format	84
Figure 23:	Switch Queues	88
Figure 24:	IEEE Tag Frame Format	91
Figure 25:	Double Tag Format	92
Figure 26:	Egress Header Format	94
Figure 27:	Egress Trailer Format	96
Figure 28:	88E6083 Device Transmit Block Diagram	100
Figure 29:	88E6083 Device Receive Block Diagram	101
Figure 30:	Serial LEDENA High Clocking with COLX in Dual Mode, Error Off, and DUPLEX in Single Mode	116
Figure 31:	Serial LED Conversion	117
Figure 32:	Serial LED Display Order	117
Figure 33:	88E6083 Register Map	121
Figure 34:	Typical MDC/MDIO Read Operation	156
Figure 35:	Typical MDC/MDIO Write Operation	156
Figure 36:	Cable Fault Distance Trend Line	186
Figure 37:	EEPROM Data Format	189
Figure 38:	Reset and Configuration Timing	196
Figure 39:	Oscillator Clock Timing	197

Figure 40: PHY Mode MII Receive Timing.....	198
Figure 41: PHY Mode MII Transmit Timing.....	199
Figure 42: MAC Clock Timing.....	200
Figure 43: MAC Mode MII Receive Timing.....	201
Figure 44: MAC Mode MII Transmit Timing.....	202
Figure 45: MAC Clock Timing, 200 Mbps.....	203
Figure 46: PHY Mode MII Receive Timing-200 Mbps.....	204
Figure 47: PHY Mode MII Transmit Timing—200 Mbps.....	205
Figure 48: MAC Mode MII Receive Timing - 200 Mbps.....	206
Figure 49: MAC Mode MII Transmit Timing, 200 Mbps.....	207
Figure 50: SNI Falling Edge Receive Timing.....	208
Figure 51: SNI Falling Edge Transmit Timing.....	209
Figure 52: SNI Rising Edge Receive Timing.....	210
Figure 53: SNI Rising Edge Transmit Timing.....	211
Figure 54: PHY Mode MII Receive Timing.....	212
Figure 55: PHY Mode MII Transmit Timing.....	213
Figure 56: Serial LED Timing.....	214
Figure 57: Serial Management Interface Clock Timing.....	215
Figure 58: Serial Management Interface Timing.....	216
Figure 59: EEPROM Timing.....	217
Figure 60: 88E6083 216-pin LQFP Package.....	219
Figure 61: 88E6083 Sample Ordering Part Number.....	221
Figure 62: 88E6083 Commercial Package Marking and Pin 1 Location.....	222
Figure 63: 88E6083 Industrial Package Marking and Pin 1 Location.....	223



Section 1. Signal Description

1.1 88E6083 216-Pin LQFP Package




Figure 1: 88E6083 216-Pin LQFP Package (Top View)

1.2 Pin Description

Table 1: Pin Type Definitions

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



Table 2: Network Interface

216-LQFP Package Pin #	Pin Name	Pin Type	Description
53 42 40 29 27 16 14 3	P7_RXP P6_RXP P5_RXP P4_RXP P3_RXP P2_RXP P1_RXP P0_RXP	Typically Input	Receiver input – Positive. P[7:0]_RXP connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode RXP connects directly to the fiber-optic receiver's positive output. For lowest power, all unused port RXP pins should be tied to VSS. These pins can become outputs if Auto MDI/MDIX Crossover is enabled (Section 4.6 "Auto MDI/MDIX Crossover" on page 113). If a PHY port's pins are not used, they should be tied to VSS.
52 43 39 30 26 17 13 4	P7_RXN P6_RXN P5_RXN P4_RXN P3_RXN P2_RXN P1_RXN P0_RXN	Typically Input	Receiver input – Negative. P[7:0]_RXN connects directly to the receiver magnetics. If the port is configured for 100BASE-FX mode RXN connects directly to the fiber-optic receiver's negative output. For lowest power, all unused port RXN pins should be tied to VSS. These pins can become outputs if Auto MDI/MDIX Crossover is enabled (Section 4.6 "Auto MDI/MDIX Crossover" on page 113). If a PHY port's pins are not used, they should be tied to VSS.
49 46 36 33 23 20 10 7	P7_TXP P6_TXP P5_TXP P4_TXP P3_TXP P2_TXP P1_TXP P0_TXP	Typically Output	Transmitter output – Positive. P[7:0]_TXP connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode 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 MDI/MDIX Crossover is enabled (Section 4.6 "Auto MDI/MDIX Crossover" on page 113). If a PHY port's pins are not used, they should be tied to VSS.
48 47 35 34 22 21 9 8	P7_TXN P6_TXN P5_TXN P4_TXN P3_TXN P2_TXN P1_TXN P0_TXN	Typically Output	Transmitter output – Negative. P[7:0]_TXN connects directly to the transmitter magnetics. If the port is configured for 100BASE-FX mode 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 MDI/MDIX Crossover is enabled (Section 4.6 "Auto MDI/MDIX Crossover" on page 113). If a PHY port's pins are not used, they should be tied to VSS.
60 59 57 56 210 209 208 207	P7_SDET P7_SDET P7_SDET P7_SDET P7_SDET P7_SDET P7_SDET P7_SDET	Input	Signal Detect input. If any PHY port 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 a PHY port is configured for 10/100BASE-T mode, Px_SDET is not used but it can not be left floating since these pins do not contain internal resistors. Px_SDET must be tied to VSS or VDDO either directly or through a 4.7 kΩ resistor.

Table 3: PHY Configuration

216-LQFP Package Pin #	Pin Name	Pin Type	Description
197 195 194 192 191 189 187 186	P7_CONFIG P6_CONFIG P5_CONFIG P4_CONFIG P3_CONFIG P2_CONFIG P1_CONFIG P0_CONFIG	Input	<p>Port Configuration. The Px_CONFIG pin is used to set the default configuration for each Port by connecting these pins to other device pins as follows:</p> <p>VSS = Auto-Negotiation enabled - default P0_LED1 = Forced 10BASE-T half-duplex P0_LED2 = Forced 10BASE-T full-duplex P1_LED0 = Forced 100BASE-TX half-duplex P1_LED1 = Forced 100BASE-TX full-duplex P1_LED2 = Forced 100BASE-FX half-duplex VDDO = Forced 100BASE-FX full-duplex</p> <p>Any port's default configuration can be modified by accessing the PHY registers by a CPU or a serial EEPROM.</p> <p>Fiber mode vs. copper mode cannot be configured in this way, however. Fiber vs. copper must be selected at Reset by using these pins.</p> <p>Px_CONFIG pins are configured after Reset and contain internal pull-down resistors so they can be left floating.</p>



Table 3: PHY Configuration (Continued)

216-LQFP Package Pin #	Pin Name	Pin Type	Description
199	CONFIG_A	Input	<p>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:</p> <p>VSS = LED Mode 0, FEFI disabled P0_LED0 = LED Mode 0, FEFI enabled P0_LED1 = LED Mode 1, FEFI disabled P0_LED2 = LED Mode 1, FEFI enabled P1_LED0 = LED Mode 2, FEFI disabled P1_LED1 = LED Mode 2, FEFI enabled P1_LED2 = LED Mode 3, FEFI disabled VDDO = LED Mode 3, FEFI enabled - default</p> <p>The LED modes are covered in section 4.7 and FEFI is covered in section 4.4.</p> <p>The CONFIG_A pin is configured after Reset and contains an internal pull-up resistor so it can be left floating.</p>
201	CONFIG_B	Input	<p>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:</p> <p>VSS = No Crossover, Class A¹ drivers, Energy Detect disabled P0_LED0 = No Crossover, Class A drivers, Energy Detect enabled P0_LED1 = No Crossover, Class B² drivers, Energy Detect disabled P0_LED2 = No Crossover, Class B drivers, Energy Detect enabled P1_LED0 = Auto Crossover, Class A drivers, Energy Detect disabled P1_LED1 = Auto Crossover, Class A drivers, Energy Detect enabled P1_LED2 = Auto Crossover, Class B drivers, Energy Detect disabled VDDO = Auto Crossover, Class B drivers, Energy Detect enabled - default</p> <p>Auto crossover is covered in section 4.6, Class B vs. Class A drivers are covered in Table 113 on page 187 and Energy Detect is covered in section 4.3.</p> <p>The CONFIG_B pin is configured after Reset contains an internal pull-up resistor so it can be left floating.</p>

1. A Class A driver is available for 100BASE-TX mode only and typically used in backplane or direct connect applications.
2. A Class B driver is typically used in CAT 5 applications.

Table 4: Regulator and Reference

216-LQFP Package Pin #	Pin Name	Pin Type	Description
215	RSET	Analog	Resistor reference. A 2 k Ω 1% resistor is placed between the RSET and VSS. This resistor is used to set an internal bias reference current.
203	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.
205	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.

Table 5: System

216-LQFP Package Pin #	Pin Name	Pin Type	Description
164	XTAL_IN	Input	25 MHz or 50 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 (25 MHz only) or an external oscillator (25 or 50 MHz). This is the only clock required as it is used for both the switch and the PHYs. Refer to section 9.4.2 for XTAL_IN timing requirements.
165	XTAL_OUT	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.
121	CLK_SEL	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.
106	RESETn	Input	Hardware reset. Active low. The 88E6083 device is configured during reset. When RESETn is low, all configuration pins become inputs and the value seen on these pins is latched on the rising edge of RESETn or some time after. Refer to section Table 116: of the AC Electrical Specifications for Reset and Configuration Timing details.



Table 6: Register Access Interface

216-LQFP Package Pin #	Pin Name	Pin Type	Description
125	MDC	Input	<p>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.</p> <p>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 14).</p> <p>MDC is internally pulled high via a resistor so it can be left floating when unused.</p>
123	MDIO	I/O	<p>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 4.7 kΩ to 10 kΩ.</p> <p>The 88E6083 device uses all of the 32 possible SMI port addresses.</p> <p>MDIO is internally pulled high via a resistor so it can be left floating when unused.</p>
122	INTn	Open Drain Output	<p>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.</p> <p>The INTn pin is asserted active low if SW_MODE[1:0] (see Table 14) bits are not 0b10 (stand-alone mode) and the EEPROM data has been completely read into the device. This EEPROM done interrupt indicates to any attached CPU that it may use the MDC/MDIO lines to access the internal registers because the EEPROM has finished using the registers. This pin also goes low when any other unmasked interrupt becomes active inside the device.</p>

Table 7: Serial EEPROM Interface

216-LQFP Package Pin #	Pin Name	Pin Type	Description
116	EE_CS /EE_1K	Typically Output	<p>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.</p> <p>EE_CS is a multi-function pin used to configure the 88E6083 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:</p> <ul style="list-style-type: none"> Low = Use 8-bit addresses (for 2K bit 93C56 & 4K bit 93C66) High = Use 6-bit addresses (for 1K bit 93C46) <p>The external EEPROM must be configured in the x16 organization.</p> <p>EE_CS is internally pulled high via a resistor so the pin can be left floating when unused or to select support for 1K bit devices. Use a 4.7 kΩ resistor to VSS for a configuration low.</p>
114	EE_CLK	Typically Input Output During EEPROM Loading	<p>Serial EEPROM clock. EE_CLK is the serial EEPROM clock reference output by the 88E6083 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.</p> <p>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.</p>



Table 7: Serial EEPROM Interface (Continued)

216-LQFP Package Pin #	Pin Name	Pin Type	Description
117	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). EE_DIN is a multifunction pin used to configure the 88E6083 during a hardware reset. When reset is asserted, EE_DIN becomes an input and the desired Half-duplex Flow Control disable is latched at the rising edge of reset as follows: Low = Enable "forced collision" flow control on all half duplex ports High = Disable flow control on all half-duplex ports EE_DIN is internally pulled high via a resistor. Use a 4.7 kΩ resistor to VSS for a configuration low. FD_FLOW_DIS is used to select flow control on full-duplex ports (see Table 14).
112	EE_DOUT	Input	Serial EEPROM data out from the EEPROM device. EE_DOUT is serial EEPROM data referenced to EE_CLK used to receive the EEPROM configuration data from the external serial EEPROM (if present). EE_DIN is internally pulled high via a resistor so it can be left floating when the pin is unused.

Table 8: Port 8's Enable

216-LQFP Package Pin #	Pin Name	Pin Type	Description
120	ENABLE_P8	Input	Enable Port 8. This pin is used to enable the Port 8 MII drivers (Table 9). A high enables the Port 8 drivers. A low disables Port 8 and its drivers (i.e., they are tri-stated). ENABLE_P8 is internally pulled high via a resistor so the pin can be left unconnected to enable Port 8.

Table 9: Port 8's Input MII - If ENABLE_P8 = High

216-LQFP Package Pin #	Pin Name	Pin Type	Description
138	P8_INCLK	I/O	<p>Input Clock. P8_INCLK is a reference for P8_INDV and P8_IND[3:0]. The direction and speed of P8_INCLK is determined by P8_MODE[3:0] (Table 10) at the end of RESETn.</p> <p>If the port is in PHY Mode, P8_INCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-X mode, and 2.5 MHz if the port is in 10BASE-T mode and 50 MHz for RMII mode.</p> <p>If the port is in MAC Mode, P8_INCLK is an input. In this mode, the frequency of the clock can be anywhere from DC to 25 MHz although it should be 25 MHz for 100BASE-X mode and 2.5 MHz for 10BASE-T mode.</p> <p>P8_INCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.</p>
140 141 142 143	P8_IND3 P8_IND2 P8_IND1 P8_IND0	Input	<p>Input Data. P8_IND[3:0] receives the data nibble to be transmitted into the switch in 100BASE-X and 10BASE-T modes. P8_IND[3:0] is synchronous to P8_INCLK. These pins are inputs regardless of the port's mode (i.e., PHY mode or MAC mode). Only P8_IND0 is used when SNI mode is selected. P8_IND[1:0] are used when RMII mode is selected.</p> <p>P8_IND[3:0] are internally pulled high via resistor so the pins can be left unconnected when they are not used.</p>
145	P8_INDV	Input	<p>Input Data Valid. When P8_INDV is asserted high, data on P8_IND[3:0] is encoded and transmitted into the switch. P8_INDV must be synchronous to P8_INCLK.</p> <p>P8_INDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.</p>



Table 10: Port 8's Output MII - If ENABLE_P8 = High

216-LQFP Package Pin #	Pin Type	Pin Name	Description
135	P8_OUTCLK	I/O	<p>Output Clock. P8_OUTCLK is a reference for P8_OUTDV and P8_OUTD[3:0]. The direction and speed of P8_OUTCLK is determined by P8_MODE[3:0] at the end of RESETn.</p> <p>If the port is in PHY Mode, P8_OUTCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-X mode, and 2.5 MHz if the port is in 10BASE-T mode.</p> <p>If the port is in MAC Mode, P8_OUTCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25 MHz although it should be 25 MHz for 100BASE-X mode and 2.5 MHz for 10BASE-T mode.</p> <p>P8_OUTCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.</p>
130 131 132 133	P8_OUTD3 /P8_MODE3 P8_OUTD3 /P8_MODE3 P8_OUTD3 /P8_MODE3 P8_OUTD3 /P8_MODE3	Normally Output Input only when RESETn is low	<p>Output Data. Data transmitted from the switch is decoded and presented on P8_OUTD[3:0] pins synchronous to P8_OUT_CLK in all modes except RMII mode where P8_OUTD[3:0] is synchronous to P8_INCLK. These pins are outputs regardless of the port's mode (i.e., PHY or MAC mode). Only P8_OUTD0 contains meaningful data when SNI mode is selected. P8_OUTD[1:0] are used when RMII mode is selected.</p> <p>During RESETn these internally pulled high pins are tri-stated and used to latch in the required operating mode for the port as described in section 3.2.5.</p>
136	P8_OUTDV	Normally Output	<p>Output Data Valid. When P8_OUTDV is asserted high, data transmitted from the switch is decoded and presented on P8_OUTD[3:0]. P8_OUTDV is synchronous to P8_OUTCLK in MII & SNI modes. It is synchronous to P8_INCLK in RMII mode.</p>

Table 10: Port 8's Output MII - If ENABLE_P8 = High (Continued)

216-LQFP Package Pin #	Pin Type	Pin Name	Description
127	P8_CRS	I/O	<p>Carrier Sense. After RESETn, P8_CRS becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P8_CRS asserts (or is expected to be asserted) when the receive data path is non-idle. In half-duplex mode P8_CRS is also asserted (or is expected to be asserted) during transmission. P8_CRS is asynchronous to P8_OUTCLK and P8_INCLK.</p> <p>P8_CRS is tri-stated during RESETn and it is internally pulled low so the pin can be left unconnected if not used.</p>
128	P8_COL	I/O	<p>Collision. After reset, P8_COL becomes an output if PHY mode is selected for this port. It remains an input if MAC mode is selected. In PHY mode, P8_COL asserts when both the transmit and receive paths are non-idle in both half and full-duplex modes. In half-duplex MAC mode, P8_COL is expected to be asserted when both the transmit and receive paths are non-idle. In full-duplex MAC mode, P8_COL is ignored. P8_COL is asynchronous with P8_OUTCLK and P8_INCLK.</p> <p>P8_COL is tri-stated during reset and it is internally pulled low so the pin can be left unconnected if not used.</p>

Table 11: Port 9's Enable

216-LQFP Package Pin #	Pin Name	Pin Type	Description
151	DISABLE_P9	Input	<p>Disable Port 9. This pin is used to disable Port 9's MII drivers (Table 9). A high disables Port 9's drivers (i.e., they are tri-stated). A low enables Port 9 and its drivers.</p> <p>DISABLE_P9 is internally pulled high via a resistor so the pin can be left unconnected to disable Port 9.</p>



Table 12: Port 9's Input MII - If DISABLE_P9 = Low

216-LQFP Package Pin #	Pin Name	Pin Type	Description
161	P9_INCLK	I/O	<p>Input Clock. P9_INCLK is a reference for P9_INDV and P9_IND[3:0]. The direction and speed of P9_INCLK is determined by P9_MODE[3:0] (Table 13) at the end of RESETn.</p> <p>If the port is in PHY Mode, P9_INCLK 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 or 200BASE mode.</p> <p>If the port is in MAC Mode, P9_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 200BASE mode.</p> <p>P9_INCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.</p>
167 169 171 173	P9_IND3 P9_IND2 P9_IND1 P9_IND0	Input	<p>Input Data. P9_IND[3:0] receives the data nibble to be transmitted into the switch in 100BASE-X and 10BASE-T modes. P9_IND[3:0] is synchronous to P9_INCLK. These pins are inputs regardless of the port's mode (i.e., PHY mode or MAC mode). Only P9_IND0 is used when SNI mode is selected. P9_IND[1:0] are used when RMII mode is selected.</p> <p>P9_IND[3:0] are internally pulled high via resistor so the pins can be left unconnected when they are not used.</p>
175	P9_INDV	Input	<p>Input Data Valid. When P9_INDV is asserted high, data on P9_IND[3:0] is encoded and transmitted into the switch. P9_INDV must be synchronous to P9_INCLK.</p> <p>P9_INDV is internally pulled low via resistor so the pin can be left unconnected when it is not used.</p>

Table 13: Port 9's Output MII - If DISABLE_P9 = Low

216-LQFP Package Pin #	Pin Name	Pin Type	Description
158	P9_OUTCLK	I/O	<p>Output Clock. P9_OUTCLK is a reference for P9_OUTDV and P9_OUTD[3:0]. The direction and speed of P9_OUTCLK is determined by P9_MODE[3:0] at the end of RESETn.</p> <p>If the port is in PHY Mode, P9_OUTCLK is an output. In this mode the frequency of the clock will be 25 MHz if the port is in 100BASE-X mode, and 2.5 MHz if the port is in 10BASE-T mode and 50 MHz for RMII mode or 200BASE mode.</p> <p>If the port is in MAC Mode, P9_OUTCLK is an input. In this mode the frequency of the clock can be anywhere from DC to 25MHz or 50 MHz although it should be 25 MHz for 100BASE-X mode and 2.5 MHz for 10BASE-T mode and 50 MHz for 200BASE mode.</p> <p>P9_OUTCLK is tri-stated during RESETn and it is internally pulled high so the pin can be left unconnected if not used.</p>
153 154 155 156	P9_OUTD3 /P9_MODE3 P9_OUTD3 /P9_MODE3 P9_OUTD3 /P9_MODE3 P9_OUTD3 /P9_MODE3	Normally Output Input only when RESETn is low	<p>Output Data. Data transmitted from the switch is decoded and presented on P9_OUTD[3:0] pins synchronous to P9_OUT_CLK. These pins are outputs regardless of the port's mode (i.e., PHY or MAC mode). Only P9_OUTD0 contains meaningful data when SNI mode is selected. P9_OUTD[1:0] are used when RMII mode is selected.</p> <p>During RESETn these internally pulled high pins are tri-stated and used to latch in the desired operating mode for the port as described in section 3.2.5.</p>
160	P9_OUTDV	Output	<p>Output Data Valid. When P9_OUTDV is asserted high, data transmitted from the switch is decoded and presented on P9_OUTD[3:0]. P9_OUTDV is synchronous to P9_OUTCLK.</p>



Table 13: Port 9's Output MII - If DISABLE_P9 = Low (Continued)

216-LQFP Package Pin #	Pin Name	Pin Type	Description
148	P9_CRS	I/O	<p>Carrier Sense. After RESETn, P9_CRS becomes an output if PHY Mode is selected for this port. It remains an input if MAC Mode is selected. P9_CRS asserts (or is expected to be asserted) when the receive data path is non-idle. In half-duplex mode P9_CRS is also asserted (or is expected to be asserted) during transmission. P9_CRS is asynchronous to P9_OUTCLK and P9_INCLK.</p> <p>P9_CRS is tri-stated during RESETn and it is internally pulled low so the pin can be left unconnected if not used.</p>
150	P9_COL	I/O	<p>Collision. After reset, P9_COL becomes an output if PHY mode is selected for this port. It remains an input if MAC mode is selected. In PHY mode, P9_COL asserts when both the transmit and receive paths are non-idle in both half and full-duplex modes. In half-duplex MAC mode, P9_COL is expected to be asserted when both the transmit and receive paths are non-idle. In full-duplex MAC mode, P9_COL is ignored. P9_COL is asynchronous with P9_OUTCLK and P9_INCLK.</p> <p>P9_COL is tri-stated during RESETn and it is internally pulled low so the pin can be left unconnected if not used.</p>

Table 14: Switch Configuration Interface

216-LQFP Package Pin #	Pin Name	Pin Type	Description
110 107	SW_MODE1 SW_MODE0	Input	<p>Switch Mode. These pins are used to configure the 88E6083 device after reset. Switch Mode pins work as follows:</p> <p><u>1 0 Description</u></p> <p>0 0 CPU attached mode – ports come up disabled¹</p> <p>0 1 Reserved</p> <p>1 0 Stand-alone mode – ports come up enabled – ignore EEPROM</p> <p>1 1 EEPROM attached mode – EEPROM defined port states²</p> <p>The EEPROM attached mode (when both SW_MODE pins = high) can be used with a CPU and the CPU attached mode can be used with an EEPROM. In all but the stand-alone mode, the INTn pin asserts active low after the EEPROM completes initializing the internal registers (i.e. a halt command has been executed)³</p> <p>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.</p>
118	FD_FLOW_DIS	Input	<p>Full-duplex Flow Control disable.</p> <p>High = disable flow control on all full-duplex ports</p> <p>Low = enable IEEE 802.3x Pause based flow control on all supported full-duplex ports</p> <p>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.</p> <p>FD_FLOW_DIS is latched on the rising edge of RESETn into all the port's PHY Auto-Negotiation Advertisement register. It is internally pulled high via a resistor so the pin can be left unconnected to disable full-duplex flow control.</p> <p>The EE_DIN/HD_FLOW_DIS multi-function pin is used to select Flow control on half-duplex ports (Table 7).</p>

1. The ports come up disabled in the CPU mode so the CPU can perform bridge loop detection on link up.
2. In EEPROM attached mode the ports come up enabled unless the Port Control register (Table 53 on page 127) is overwritten by the EEPROM data (see Section 8. "EEPROM Programming Format" for more details).
3. The INTn pin is activated so the CPU can access the registers using the MDC/MDIO pins. The CPU cannot access these registers as long as the EEPROM is still being executed.



Table 15: Port Status LEDs

216-LQFP Package Pin #	Pin Name	Pin Type	Description
62 67 72 76 81 85 90 94	P7_LED2 P6_LED2 P5_LED2 P4_LED2 P3_LED2 P2_LED2 P1_LED2 P0_LED2	Output	Parallel LED outputs – one for each port. This active low LED pin directly drives an LED in Parallel LED mode. It can be configured to display many options. P[7:0]_LED2 are driven active low whenever RESETn is active low.
64 68 73 77 82 86 91 95	P7_LED1 P6_LED1 P5_LED1 P4_LED1 P3_LED1 P2_LED1 P1_LED1 P0_LED1	Output	Parallel LED outputs – one for each port. This active low LED pin directly drives an LED in Parallel LED mode. It can be configured to display many options. P[7:0]_LED1 are driven active low whenever RESETn is active low.
65 70 74 79 83 88 92 97	P7_LED0 P6_LED0 P5_LED0 P4_LED0 P3_LED0 P2_LED0 P1_LED0 P0_LED0	Output	Parallel LED outputs – one for each port. This active low LED pin directly drives an LED in Parallel LED mode. It can be configured to display many options. P[7:0]_LED0 are driven active low whenever RESETn is active low.
99	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 (see Serial LED Interface on page 116).
103	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.
101	LEDCLK	Output	LEDCLK is the reference clock for the serial LED signals.

Table 16: Power and Ground

216-LQFP Package Pin #	Pin Name	Pin Type	Description
58 63 69 78 87 96 104 105 111 115 126 137 147 157 166 172 182 193 202	VDDO	Power	Power to outputs. VDDO must be connected to 3.3V, which supports 3.3V I/O.
2 15 28 41 54 206	VDDAH	Power	2.5V Power to analog core
5 12 18 25 31 38 44 51	VDDAL	Power	1.5V Power to analog core



Table 16: Power and Ground (Continued)

216-LQFP Package Pin #	Pin Name	Pin Type	Description
61 71 80 89 98 102 113 124 139 149 159 168 174 180 188 198	VDD	Power	1.5V Power to digital core

Table 16: Power and Ground (Continued)

216-LQFP Package Pin #	Pin Name	Pin Type	Description
1 6 11 19 24 32 37 45 50 55 66 75 84 93 100 108 109 119 129 134 144 152 162 163 170 177 184 185 190 196 200 204 211 216	VSS	Ground	Ground to device
146 176 178 179 181 183 212 213 214	NC	--	No Connect. Do not connect these pins to anything. These pins must be left unconnected.



1.3 216-Pin LQFP Pin Assignment List—by Signal Name

Table 17: Package Pin List—Alphabetical by Signal Name

Pin Number	Pin Name	Pin Number	Pin Name
121	CLK_SEL	4	P0_RXN
199	CONFIG_A	3	P0_RXP
201	CONFIG_B	207	P0_SDET
203	CONTROL_15	8	P0_TXN
205	CONTROL_25	7	P0_TXP
151	DISABLE_P9	187	P1_CONFIG
114	EE_CLK	92	P1_LED0
116	EE_CS/EE_1K	91	P1_LED1
117	EE_DIN/HD_FLOW_DIS	90	P1_LED2
112	EE_DOUT	13	P1_RXN
120	ENABLE_P8	14	P1_RXP
118	FD_FLOW_DIS	208	P1_SDET
122	INTn	9	P1_TXN
101	LEDCLK	10	P1_TXP
103	LEDENA	189	P2_CONFIG
99	LEDSER	88	P2_LED0
125	MDC	86	P2_LED1
123	MDIO	85	P2_LED2
146	NC	17	P2_RXN
176	NC	16	P2_RXP
178	NC	209	P2_SDET
179	NC	21	P2_TXN
181	NC	20	P2_TXP
183	NC	191	P3_CONFIG
212	NC	83	P3_LED0
213	NC	82	P3_LED1
214	NC	81	P3_LED2
186	P0_CONFIG	26	P3_RXN
97	P0_LED0	27	P3_RXP
95	P0_LED1	210	P3_SDET
94	P0_LED2	22	P3_TXN
		23	P3_TXP

Pin Number	Pin Name
192	P4_CONFIG
79	P4_LED0
77	P4_LED1
76	P4_LED2
30	P4_RXN
29	P4_RXP
56	P4_SDET
34	P4_TXN
33	P4_TXP
194	P5_CONFIG
74	P5_LED0
73	P5_LED1
72	P5_LED2
39	P5_RXN
40	P5_RXP
57	P5_SDET
35	P5_TXN
36	P5_TXP
195	P6_CONFIG
70	P6_LED0
68	P6_LED1
67	P6_LED2
43	P6_RXN
42	P6_RXP
59	P6_SDET
47	P6_TXN
46	P6_TXP
197	P7_CONFIG
65	P7_LED0
64	P7_LED1
62	P7_LED2
52	P7_RXN
53	P7_RXP
60	P7_SDET
48	P7_TXN

Pin Number	Pin Name
49	P7_TXP
128	P8_COL
127	P8_CRS
138	P8_INCLK
143	P8_IND0
142	P8_IND1
141	P8_IND2
140	P8_IND3
145	P8_INDV
135	P8_OUTCLK
133	P8_OUTD0/P8_MODE0
132	P8_OUTD1/P8_MODE1
131	P8_OUTD2/P8_MODE2
130	P8_OUTD3/P8_MODE3
136	P8_OUTDV
150	P9_COL
148	P9_CRS
161	P9_INCLK
173	P9_IND0
171	P9_IND1
169	P9_IND2
167	P9_IND3
175	P9_INDV
158	P9_OUTCLK
156	P9_OUTD0/P9_MODE0
155	P9_OUTD1/P9_MODE1
154	P9_OUTD2/P9_MODE2
153	P9_OUTD3/P9_MODE3
160	P9_OUTDV
106	RESETn
215	RSET
107	SW_MODE0
110	SW_MODE1
61	VDD
71	VDD



Pin Number	Pin Name
80	VDD
89	VDD
98	VDD
102	VDD
113	VDD
124	VDD
139	VDD
149	VDD
159	VDD
168	VDD
174	VDD
180	VDD
188	VDD
198	VDD
2	VDDAH
15	VDDAH
28	VDDAH
41	VDDAH
54	VDDAH
206	VDDAH
5	VDDAL
12	VDDAL
18	VDDAL
25	VDDAL
31	VDDAL
38	VDDAL
44	VDDAL
51	VDDAL
58	VDDO
63	VDDO
69	VDDO
78	VDDO
87	VDDO
96	VDDO
104	VDDO

Pin Number	Pin Name
105	VDDO
111	VDDO
115	VDDO
126	VDDO
137	VDDO
147	VDDO
157	VDDO
166	VDDO
172	VDDO
182	VDDO
193	VDDO
202	VDDO
1	VSS
6	VSS
11	VSS
19	VSS
24	VSS
32	VSS
37	VSS
45	VSS
50	VSS
55	VSS
66	VSS
75	VSS
84	VSS
93	VSS
100	VSS
108	VSS
109	VSS
119	VSS
129	VSS
134	VSS
144	VSS
152	VSS
162	VSS

Pin Number	Pin Name
163	VSS
170	VSS
177	VSS
184	VSS
185	VSS
190	VSS
196	VSS
200	VSS
204	VSS
211	VSS
216	VSS
164	XTAL_IN
165	XTAL_OUT



Section 2. Application Examples

The Marvell® 88E6083 devices support many different applications with few additional active components in a small footprint. The 88E6083 devices contain ten ports with eight PHYs. This architecture enables all PHYs to be active and available while at the same time supporting one or two independent MIIs.

The higher integration, low power, and two extra ports save BOM costs on designs requiring eight ports and one or two MIIs. The flexibility of the 88E6083 devices come from its configuration options on Port 8 and Port 9, with the options on the two MIIs. The devices also support an optional EEPROM (93C46 type) to override any required QoS default settings.

The only additional active components required to implement complete 88E6083 devices 10/100 Ethernet switch are:

- A 25 MHz crystal clock source or a 25 or 50 MHz oscillator
- A low-cost PNP transistor used to regulate 3.3V down to 2.5V
- A low-cost PNP transistor used to regulate 3.3V down to 1.5V

The 88E6083 devices are full system-on-a-chips containing the PHYs, LED drivers, voltage regulator logic, QoS switch logic, and memory.

2.1 Configuration Options for the 88E6083 Devices

The 88E6083 devices have multiple configuration options selectable through its configuration pins, which include:

- ENABLE_P8, and DISABLE_P9 ([Table 8](#) and [Table 11](#)) determine whether Port 8 and Port 9 are enabled
- Pins P[7:0]_CONFIG determine each port's mode, e.g. full or half duplex or auto negotiation enabled, and whether each port interface is configured for copper or fiber—see [Table 3](#)
- Pins Px_MODE[3:0] state is latched during RESETn and used to set the configuration of the MIIs on Port 8 and Port 9

See [Section 1. "Signal Description"](#) for a full description of the port setup options for this part.

2.2 Example Configurations

These examples describe two configurations:

- Firewall Router Switch supporting a Fiber WAN port with a double speed MII interface to the CPU - shown in [Figure 2](#)
- Firewall gateway router with QoS support for designs with one 10/100BASE-T WAN port, seven 10/100BASE-T LAN ports, and an extra MII port for an optional wireless LAN network connection—shown in [Figure 3](#)

For the firewall router examples, Port 8 is not used, Port 9 is configured in 200 BASE PHY Mode, and Ports 0 to 7 support 100BASE-TX and/or 10BASE-T operation. If required, some or all of Ports 0 to 7 could be configured as 100BASE-FX fiber interfaces for the WAN and LAN (see [Figure 2](#)). The CPU's SMI is used to access all of the PHY and Switch registers inside the 88E6083 devices.

In [Figure 3](#), Port 0 is used as a WAN interface and Port 8 is used to connect to a wireless base station PHY in the LAN. The 88E6083 devices support port-based VLANs in any programmable configuration to meet all of the sys-

tem requirements for the firewall router shown in Figure 3. In this configuration, the WAN port is isolated from the LAN ports, and vice versa. LAN frames cannot go directly to the WAN, and WAN frames cannot go directly to the LAN. The CPU's port serves as the bridge by transmitting, receiving, and translating frames between the WAN and the LAN.

Bringing the 802.3 WAN PHY into the 88E6083 devices' Port 0 not only reduces BOM costs, but it also adds the feature of a 10/100 Mbps interface with Auto Crossover cable support (supported on all copper ports). The second MII enables the economic addition of new connectivity to a Gateway Router, such as a wireless PHY. Alternatively, it can support greater integration by enabling direct connection to a WAN PHY inside the Gateway box.

Figure 2: Firewall Router Switch supporting a Fiber WAN port

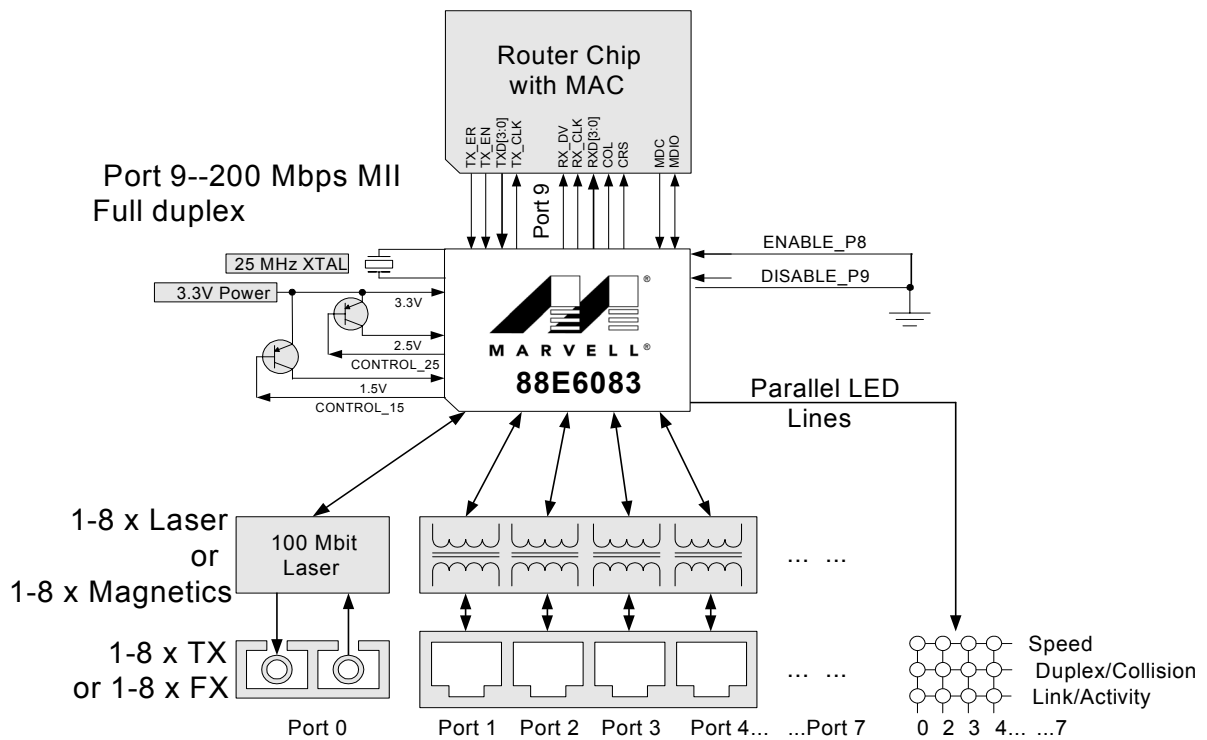
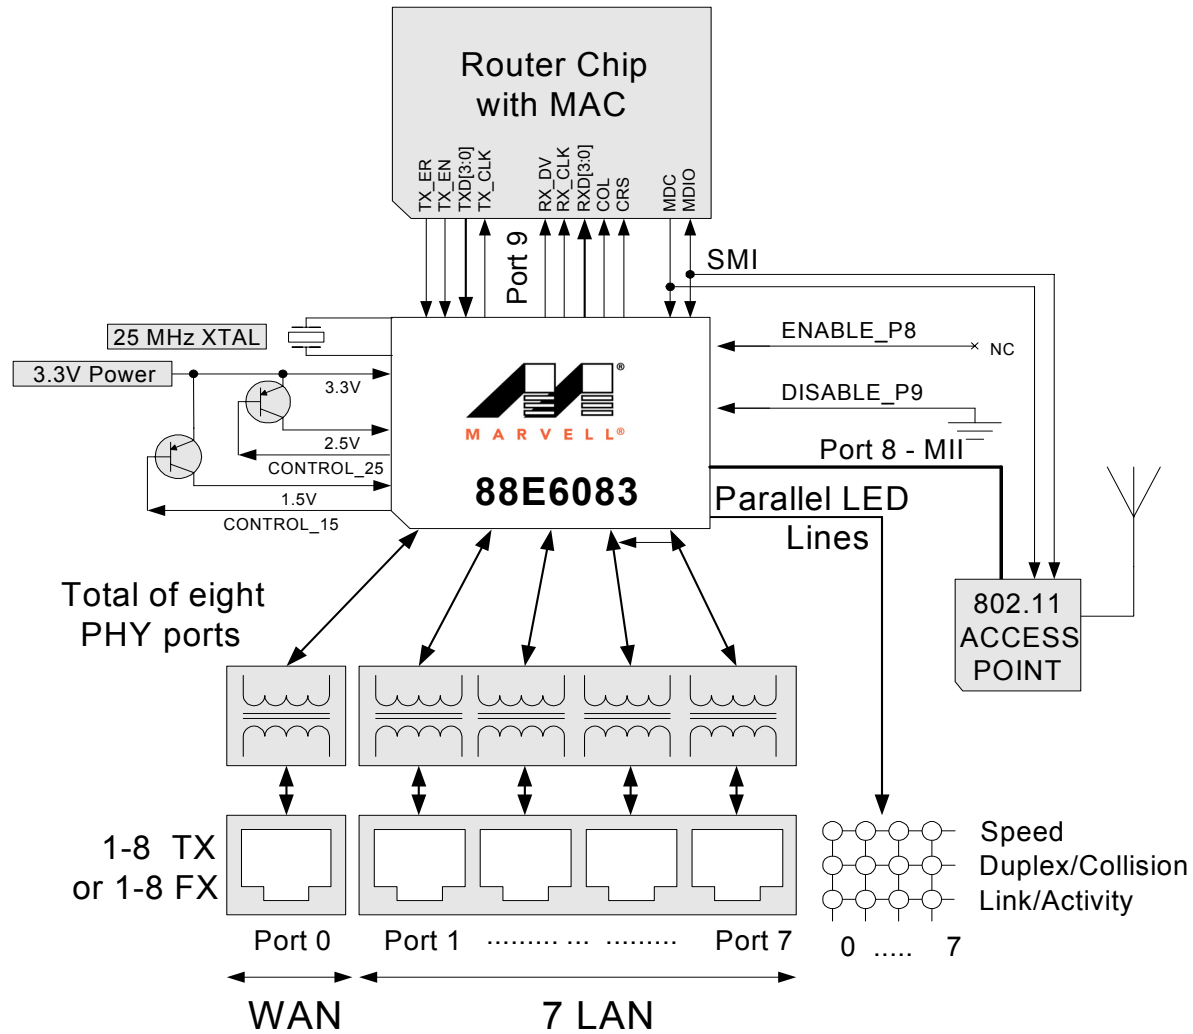


Figure 3: 88E6083 Firewall Router with Wireless Access Point



2.3 Routing with the Marvell® Header

Any port can be set to prepend two bytes of data in front of each frame sent to the external CPU and to remove these two bytes of data returning from the CPU. This Marvell Header Mode is optional, but it can increase CPU routing performance greatly by aligning the IP data portion of the frames in the CPU's memory on 32-bit boundaries. On ingress to the CPU the Marvell Header contains information that the CPU needs, such as the switch's source port of the frame and its priority. On CPU egress, the CPU uses the Header to indicate the port-based VLAN information of the frame so that frames sent to the LAN ports only go out of those LAN ports and not out of the WAN port (the header information is directly written to switch port's Port Based VLAN Map register. The Marvell Header supports the following features:

- The Marvell Header port can be any port on the switch
- Higher routing performance due to 32-bit IP data alignment
- The WAN port can be any port on the switch
- More than one WAN port can be used either as a hot backup or for link aggregation
- Any switch port can be a DMZ port or other port type

Refer to [Section 3.5.10](#) for more information on the format of the Marvell Header to and from the CPU.

The CPU can independently be set to process the Marvell Header or ignore it. If a switch port does not generate or use the Marvell Header, the CPU must not try to process it. If a switch port generates/uses a Marvell Header, the CPU can process it or ignore it.

The Marvell Header supports full wirespeed VLAN re-configuration between WAN and LAN frames coming from the CPU. It supports separate WAN and LAN VLANS without the need of using 802.1Q tagged frames. Since 802.1Q is not used inside the switch for WAN/LAN isolation, complete compatibility is insured for all traffic, including the ability to route 802.1Q tagged frames. There are no conflicts between the customer's VID's and the switch's VID's because the switch uses the Marvell Header instead.

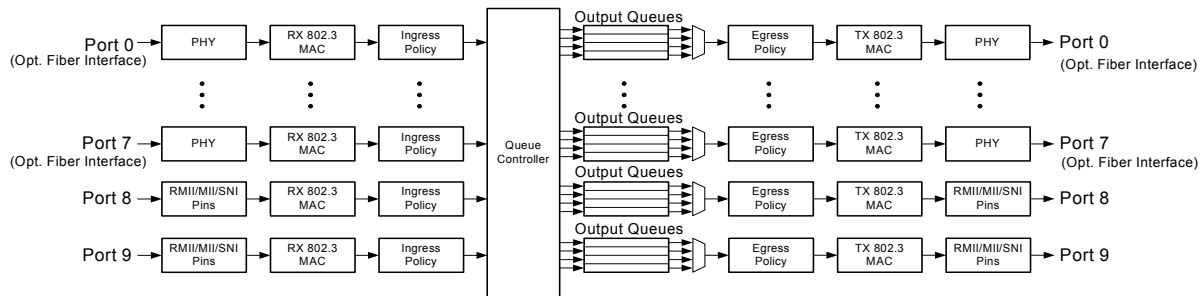
Section 3. Functional Description

This section describes the wire-speed, non-blocking 10-port 10/100-Mbps Fast Ethernet QoS (Quality-of-Service) switch core that is integrated into the 88E6083 devices. The ten ports include PHY and RMII/MII/SNI ports, as described below.

3.1 Switch Data Flow

The 88E6083 devices accept IEEE 802.3 frames and either discards them or transmits them out of one or more of the switchports. The decision on what to do with each frame is one of the many jobs handled inside the switch. Figure 4 shows the data path inside the switch along with the major functional blocks that process the frame as it travels through the 88E6083 device.

Figure 4: 88E6083 Device Switch Data Flow



The PHY, or physical layer interface, is used to receive and transmit frames over CAT 5 twisted pair cable or fiber-optic cables. The 88E6083 device contains eight PHYs connected to Port 0 to Port 7. The PHY block is covered in detail in Section 4.

Neither Port 8 nor Port 9 contain a PHY. Instead, each supports a short-distance, industry-standard digital interface, generically called the port's Media Independent Interface (MII). Many interface modes and timings are supported so that a large number of external device types can be used.

The two MIIs on the 88E6083 device support four major modes of operation (Section 3.2) and each Media Independent Interface (MII) can be configured independently:


The 88E6083 device contains ten independent MACs that perform all of the functions specified in the 802.3 protocol, which include, among others:

- Frame formatting
- CRC checking
- CSMA/CD enforcement
- Collision handling

The switch portion of the 88E6083 device receives good packets from the MACs, processes them, and forwards them to the appropriate MACs for transmission. Processing frames is the key activity, and it involves the Ingress

Policy block, the Queue Controller, the Output Queues, and the Egress Policy blocks shown in Figure 5. These blocks modify the normal or default packet flow through the switch and are discussed in Section 3.5 to Section 3.7

Figure 5: Switch Operation



3.2 MII/SNI/RMII

The Media Independent Interface (MII) supports four major modes of operation:

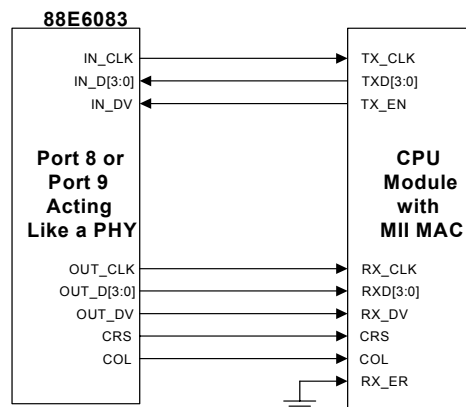
- MII PHY Mode (The 88E6083 devices drive the interface clocks. See Section 3.2.1 for details).
- MII MAC Mode (The external device drives the interface clocks. See Section 3.2.2 for details).
- SNI PHY Mode
- RMII PHY Mode

The two MII ports can be configured differently from each other.

3.2.1 MII PHY Mode

The MII PHY Mode, (Reverse MII), configures the selected MAC inside the 88E6083 device to emulate a PHY. This enables the 88E6083 device to be directly connected to an external MAC (for example, one inside a Router CPU). In this mode, the 88E6083 device drives the interface clocks (Px_INCLK and Px_OUTCLK); therefore, the required frequency needs to be selected. Both full-duplex and half-duplex modes are supported and need to be selected to match the mode of the link partner's MAC. The MII PHY interface is compliant with the IEEE 802.3u clause 22.

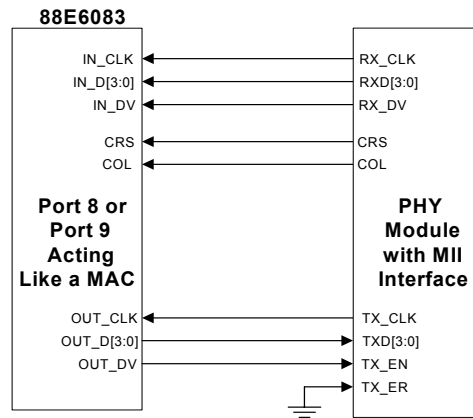
Figure 6: MII PHY Interface Pins



3.2.2 MII MAC Mode

The MII MAC Mode (Forward MII) configures the required MAC inside the 88E6083 device to act as a MAC so that it can be directly connected to an external PHY. In this mode, the 88E6083 device receives the interface clocks (Px_INCLK and Px_OUTCLK) and works at any frequency from DC to 25 MHz or 50 MHz (port 9 only - dependent upon the external source). The two clocks can be asynchronous with respect to each other. Both full-duplex 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 clause 22 of the IEEE 802.3u specification.

Figure 7: MII MAC Interface Pins

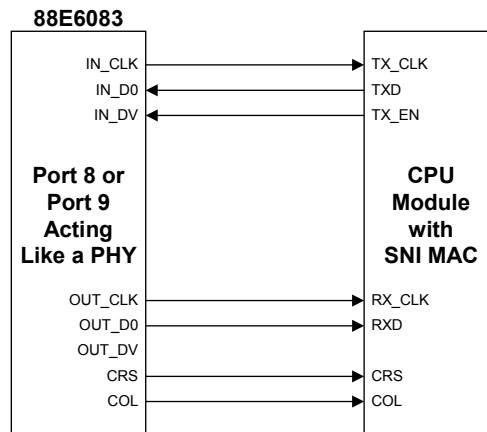


3.2.3 SNI PHY Mode

The SNI PHY Mode, 7-Wire interface, configures the selected MAC inside the 88E6083 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 88E6083 device. Since SNI was never standardized, the 88E6083 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.

In SNI mode, the output data from the switch is indicated to be valid by the CRS signal. CRS is not synchronous with OUT_CLK, however. So the receiving MAC needs to detect the beginning of the frame by looking for the Preamble and then the Start of Frame Delimiter (SFD). CRS will be deasserted asynchronously at the end of the frame so the receiving MAC needs to tolerate the presence of these extra 'dribble' bits.

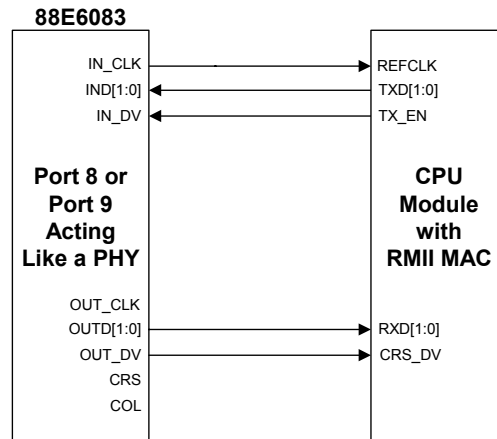
Figure 8: SNI PHY Interface Pins



3.2.4 RMII PHY Mode

RMII PHY Mode, (Reduced MII) configures the selected MAC inside the 88E6083 devices to act as a 100 Mbps PHY with a Reduced Media Independent Interface (RMII) enabling it to be directly connected to an external 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 88E6083 devices at the RMII constant frequency of 50 MHz.

Figure 9: RMII PHY Interface Pins using INCLK



3.2.5 MII/SNI Configuration

MII/SNI interfaces in the 88E6083 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 become the interface's mode as defined in [Table 18](#).

Table 18: RMII/MII/SNII 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 (supported on port 9 only) ¹
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 (supported on port 9 only)
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 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 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. If 200 Mbps mode is selected on Port 9, Port 8 is automatically disabled and cannot be used.

3.2.6 Enabling the MII/SNI Interfaces

Port 8 (the ninth port) is enabled via ENABLE_P8 (see [Table 8, "Port 8's Enable," on page 24](#)). Port 9 (the tenth port) is enabled via DISABLE_P9 (see [Table 11, "Port 9's Enable," on page 27](#)).

3.2.7 Port Status Registers

Each switch port of the 88E6083 device has a status register that reports information about that port's PHY or RMII/MII/SNI interface. See [Table 51](#) for more information.



3.2.8 MII 200 Mbps Mode

One of the 88E6083 MIIs can run at a data rate of 200 Mbps full-duplex. Only Port 9 can run at 200 Mbps and if the 200 Mbps mode is selected on Port 9, Port 8 is automatically disabled and cannot be used. Do not select this mode unless the MAC on the other end of the MII can also run at this 200 Mbps rate. Both PHY (reverse MII) and MAC (forward MII) 200 Mbps modes are supported on Port 9. When the 200 Mbps PHY mode is selected, the output MII clocks (OUTCLK and INCLK) run at a 50 MHz rate; XTAL-IN must also be 50 MHz. There is no change in the format of the data, it just moves more quickly. When the 200 Mbps MAC mode is selected, the input MII clocks (OUTCLK and INCLK) must be specified as 50 MHz \pm 50ppm. Also, the format of the data is not changed in this case.

3.3 Media Access Controllers (MAC)

The 88E6083 device contains ten independent MACs that perform all of the functions in the 802.3 protocol that include, among others, frame formatting, frame stripping, FCS checking, CSMA/CD enforcement, and collision handling.

Each MAC receive block checks incoming packets and discards packets with FCS errors, alignment errors, short packets (fewer than 64 bytes)¹, or long packets (more than 1518 bytes for un-Tagged frames and more than 1522 bytes for IEEE 802.3ac Tagged² 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), and the next six bytes are used as the packet's Source Address (SA)⁴. These two addresses are fundamental to the operation of the switch (see section 3.4). Bytes two to 16 are examined and potentially used for QoS (Quality of Service) and snooping decisions made inside the switch (see section 3.5.2 and section 3.5.9). The last four bytes of the packet contain the packet's Frame Check Sequence (FCS). The packet is discarded if the FCS does not meet the IEEE 802.3 CRC-32 requirements.

Before a packet can be sent out, the transmit block must check whether 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 when the port is in half-duplex mode. In this case, the transmitter defers its transmission until the line becomes available, when it ensures that a minimum interpacket gap of at least 96 bits occurs before transmitting a 56-bit preamble and an 8-bit Start of Frame Delimiter (SFD) ahead of the basic frame. Transmission of the frame begins immediately after the SFD.

For the half-duplex mode, the 88E6083 device also monitors the collision signal while it is transmitting. When a collision is detected (i.e., both transmitter and receiver of a half-duplex MAC are active at the same time), the MAC transmits a JAM pattern and then delays the retransmission 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.

1. When Ingress Double Tagging is enabled, the minimum frame size for Tagged frames is 68 bytes.

2. IEEE 802.3ac VLAN Tagged frames are four bytes 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 1536 bytes is supported by setting the MaxFrameSize bit in the Global Control register—see [Section 5.2.3](#).

4. If the Marvell header mode is enabled, the first two bytes after the SFD are the Header bytes (See [Section 3.5.9](#)).

3.3.1 Backoff

In half-duplex mode, the 88E6083 device's MACs implement the truncated binary exponential backoff algorithm defined in the IEEE 802.3 standard. This algorithm starts with a randomly-selected small backoff time and follows by generating progressively longer random backoff times. The random times prevent two or more MACs from always attempting re-transmission at the same time. The progressively longer backoff times give a wider random range at the expense of a longer delay, giving congested links a better chance of finding a winning transmitter. Each MAC in the 88E6083 device resets the progressively longer backoff time circuit after 16 consecutive retransmit trials. Each MAC then restarts the backoff algorithm with the shortest random backoff time and continues to retry and retransmit the frame. A packet that successively collides with a retransmit signal is re-transmitted until transmission is successful. This algorithm prevents packet loss in highly-congested environments. The MACs in the switch can be configured to meet the IEEE 802.3 specification and discard a frame after 16 consecutive collisions instead of restarting the backoff algorithm. To discard a frame after 16 consecutive collisions, set the DiscardExcessive bit to a one in the Global Control register—see [Section 5.2.3](#).

3.3.2 Half-Duplex Flow Control

Half-duplex flow control is used to throttle the end station to avoid dropping packets during network congestion. It is enabled on all half-duplex ports when the EE_DIN/HD_FLOW_DIS pin is low at the rising edge of RESETn. The 88E6083 device uses a collision-based scheme to perform half-duplex flow control. When the free buffer space is almost exhausted, the MAC forces a collision in the input port when the 88E6083 device senses an incoming packet. Only those ports that are involved in the congestion are flow controlled. When the half-duplex flow control mode is not set and no packet buffer space is available, the incoming packet is discarded.

3.3.3 Full-Duplex Flow Control

The purpose of full-duplex flow control is the same as that of flow control in the half-duplex case—to avoid dropping packets during congestion. Full-duplex flow control is enabled on all full-duplex ports when:

- FD_FLOW_DIS pin is low at the rising edge of RESETn, and
- Auto-Negotiation is enabled on the port (see [Section 4.2.13](#)), and
- The link partner 'advertises' that it supports pause during Auto-Negotiation¹

Full-duplex flow control is not automatically supported on Ports 0–7 when configured for 100BASE-FX operation, nor for Port 8 and Port 9. This is because Auto-Negotiation is not defined for 100BASE-FX or MII. Therefore, the link partners cannot advertise that they support PAUSE. It can be forced, however—see section 3.3.4. When the full-duplex flow control mode is not set and no packet buffer space is available, the incoming packet is discarded.

In full-duplex mode, the 88E6083 MACs support the standard flow control defined in the IEEE 802.3x specification. This flow control enables stopping and restoring the transmission from the remote node. The basic mechanism for performing full-duplex flow control is by means of a PAUSE frame. The format of the PAUSE frame is shown in [Table 19](#).

¹ Full-duplex flow control is not automatically supported on ports configured for 100BASE-FX operation since Auto-Negotiation is not defined for 100BASE-FX, and hence, the link partners cannot advertise that they support PAUSE. It can be forced, however – see [Section 3.3.4](#).



Table 19: PAUSE Frame Format

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 functions as follows. When the free buffer space is almost exhausted, the 88E6083 device sends out a Pause frame with the maximum Pause time (a value of all ones—0xFFFF) to stop the remote node from sending more frames into the switch. Only the node that is involved in the congestion is Paused. When the event that invoked flow control disappears, the 88E6083 device sends out a Pause frame with the Pause time equal to zero, indicating that the remote node can resume transmission. The 88E6083 device also responds to the Pause command in the MAC receiving block. When the Pause command is detected, the MAC responds within one slot time (512 bit times) to stop transmission of the new data frame for the amount of time defined in the pause time field of the Pause command packet.

The Source Address of received Pause frames is not learned (see Address Learning in [Section 3.4.3](#)) 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 88E6083 device can be configured to transmit a switch-specific Source Address on Pause frames that it transmits (the default is all zeros). Global Switch MAC Address registers 1, 2, and 3 (see [Table 61](#), [Table 62](#), and [Table 63](#)) can be set to the required 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 changing the value of the DiffAddr bit in [Table 61 on page 138](#).

The MACs always discard all IEEE 802.3x Pause frames that they receive, even when full-duplex flow control is disabled or even when the port is in half-duplex mode.

3.3.4 Forcing Flow Control

When flow control is enabled using the device's 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 have all the other ports disabled. In this case, flow control should be disabled using the appropriate device pins. Refer to the FD_FLOW_DIS and HD_FLOW_DIS pin descriptions for disabling flow control on ports by default, and to the port's ForceFlowControl bit in the Port Control Register for enabling flow control on individual ports. (see [Section 5.2.2](#)).

3.3.5 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

Frame size and the distribution of frames between multicast and unicast types are less important in this kind of debugging for an edge switch.

The 88E6083 Statistics Counter support basic debugging needs. Each port can capture two kinds of statistics:

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

The first statistic answers the question “Where did all the frames go?”, while the second statistic answers the question “Does the network segment have any performance problems?”. These counters are described in [Table 58 on page 136](#), and in [Table 59 on page 136](#). The counters can be cleared, and their mode chosen by the CtrMode bit in the Switch Global Control register ([Section 5.2.3](#)).

3.3.6 Extensive RMON/Statistics Counters

The Extensive Statistics Counters are described in [Table 20 on page 54](#) and in [Table 21 on page 56](#). The Statistics Counters’ logic maintains a total of 40 32-bit counters per port, which enable the user to monitor network performance remotely and to support Remote Monitoring (RMON) groups 1, 2, 3, and 9. These counters provide a set of Ethernet statistics for frames received on Ingress (20 counters) and transmitted on Egress (20 counters). A register interface enables the CPU to capture, read, or clear the counter values (see [Table 85, “Statistics Operation Register,” on page 152](#)).

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 which obsoletes 1229 & 1158)
- RFC 1493 - Bridge MIB (this RFC obsoletes 1286)

The complete description of each of the 40 counters is contained in [Table 20](#) and [Table 21](#).

All CPU register interfaces are slow compared to the speed of Fast Ethernet frames. For this reason the 88E6083 device supports a capture function to instantly capture and hold static any port’s set of forty 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 combined together to get the required MIB value. After capture, the CPU can read out the values of the counter or counters that it needs without the values changing between read operations. The CPU interface supports the following operations on the Statistics Counters:

- Clearing all counters for all ports
- Clearing all counters for a single port
- Capturing all counters for a single port
- Reading a captured counter (A Capture must be executed before a Read can be performed.)

See [Table 85, “Statistics Operation Register,” on page 152](#).



Table 20 describes the 20 Ingress statistics counters for each port.

Table 20: Ingress Statistics Counters

Name	Offset Address	Description
Set 1		
InUnicasts	0	Total valid ¹ frames received with a unicast Destination Address.
InBroadcasts	1	Total valid frames received with a Destination Address equal to FF:FF:FF:FF:FF:FF.
InPause	2	Total valid PAUSE frames received.
InMulticasts	3	Total valid frames received with a multicast Destination Address that are not counted in InBroadcasts or InPause.
InFCSErr	4	Total frames received with a valid length (between 64 octets and MaxSize ² octets inclusive) that have an invalid FCS and an integral number of octets.
AlignErr	5	Total frames received with a valid length (between 64 octets and MaxSize octets inclusive) that have an invalid FCS and a non-integral number of octets.
Set 2		
InGoodOctets	6	Total data octets received in frames with a valid FCS. Undersize and Oversize frames are included. The count includes the FCS but not the preamble.
InBadOctets	7	Total data octets received in frames with an invalid FCS. Fragments and Jabbers are included. The count includes the FCS but not the preamble.
Set 3		
Undersize	8	Total frames received with a length of less than 64 octets but with a valid FCS.
Fragments	9	Total frames received with a length of less than 64 octets and an invalid FCS.
In64Octets	10	Total frames received with a length of exactly 64 octets, including those with errors.
In127Octets	11	Total frames received with a length of between 65 and 127 octets inclusive, including those with errors.
In255Octets	12	Total frames received with a length of between 128 and 255 octets inclusive, including those with errors.
In511Octets	13	Total frames received with a length of between 256 and 511 octets inclusive, including those with errors.
In1023Octets	14	Total frames received with a length of between 512 and 1023 octets inclusive, including those with errors.
InMaxOctets	15	Total frames received with a length of between 1024 and MaxSize octets inclusive, including those with errors.
Jabber	16	Total frames received with a length of more than MaxSize octets but with an invalid FCS.
Oversize	17	Total frames received with a length of more than MaxSize octets but with a valid FCS.

Table 20: Ingress Statistics Counters (Continued)

Name	Offset Address	Description
Set 4		
InDiscards	18	Total valid frames received that are discarded owing to a lack of buffer space. This includes frames discarded at ingress as well as those dropped owing to priority and congestion considerations at the output queues. Frames dropped at egress owing to excessive collisions are not included but are counted in the Excessive counter.
InFiltered	19	<p>If 802.1Q is disabled on this port: Total valid frames received that are not forwarded to a destination port. These are frames for which the destination port vector is 0 or are not forwarded due to the state of the PortState bits. Valid frames discarded owing to a lack of buffer space are not included.</p> <p>If 802.1Q is enabled on this port: Total valid frames received (tagged or untagged) that were discarded owing to an unknown VID (i.e., the frame's VID was not in the VTU).</p>

1. A valid frame is one with a good FCS and whose size is 64 octets to MaxSize octets inclusive.
2. MaxSize is 1518 for non-tagged frames, 1522 for tagged frames or 1536 if MaxFrameSize = 1 (in Global Control register).



Table 21 describes the 20 Egress counters for each port.

Table 21: Egress Statistics Counters

Name	Offset Address	Description
Set 5		
OutUnicasts	20	Total frames transmitted with a unicast Destination Address.
OutBroadcasts	21	Total frames transmitted with a Destination Address equal to FF:FF:FF:FF:FF:FF.
OutPause	22	Total PAUSE frames transmitted.
OutMulticasts	23	Total frames transmitted with a multicast Destination Address that are not counted in OutBroadcasts or OutPause.
OutFCSErr	24	Total frames transmitted with an invalid FCS ¹ .
Set 6		
OutOctets	25	Total data octets transmitted from frames counted in Group 5 above. The count includes the FCS but not the preamble.
Set 7		
Out64Octets	26	Total frames transmitted with a length of exactly 64 octets, including those with errors.
Out127Octets	27	Total frames transmitted with a length of between 65 and 127 octets inclusive, including those with errors.
Out255Octets	28	Total frames transmitted with a length of between 128 and 255 octets inclusive, including those with errors.
Out511Octets	29	Total frames transmitted with a length of between 256 and 511 octets inclusive, including those with errors.
Out1023Octets	30	Total frames transmitted with a length of between 512 and 1023 octets inclusive, including those with errors.
OutMaxOctets	31	Total frames transmitted with a length of between 1024 and 1522 octets inclusive, including those with errors.
Set 8		
Collisions	32	Total number of collisions during frame transmission.
Late	33	Total number of times a collision is detected later than 512 bit-times into the transmission of a frame.
Excessive	34	Total number of frames not transmitted because the frame experienced 16 transmission attempts and it was discarded. The discard will only occur if DiscardExcessive is set to a 1 (in Global Control).
Multiple	35	Total number of successfully transmitted frames that experienced more than one collision.
Single	36	Total number of successfully transmitted frames that experienced exactly one collision.
Deferred	37	Total number of successfully transmitted frames that are delayed because the medium is busy during the first attempt.

Table 21: Egress Statistics Counters (Continued)

Name	Offset Address	Description
OutFiltered	38	Total valid frames discarded that were not transmitted owing to the following filtering rules: <ul style="list-style-type: none">• Frames that were in the port's Egress Queue that were recycled owing to Link's going down while the port is not in the Disabled Port State• Frames that were in the port's Egress Queue that are recycled owing to the Port State's not being in the Forwarding mode• Frames that were in the port's Egress Queue that are recycled owing to the 802.Q Security mode
OutDiscards	39	Total valid frames discarded that were not transmitted owing to a lack of buffer space. Always 0 in this device since all discards occur at Ingress and are counted in InDiscards.

1. Whenever a frame is modified (e.g., to add or remove a tag) and new FCS is computed for the frame. Before the new FCS is added into the frame the old FCS is inspected to ensure that the original unmodified frame is still good. If an error is detected the FCS is forcibly corrupted and the OutFCSErr counter is incremented.

3.4 Address Management

The primary function of the switch portion of the 88E6083 device is to receive good packets from the MACs, process them, and forward them to the appropriate MACs for transmission. This frame processing involves the Ingress Policy, Queue Controller, Output Queues, and Egress Policy blocks shown in [Figure 5](#). These blocks modify the normal or default packet flow through the switch and are discussed following section 3.5. The normal packet flow and processing is discussed first.

The normal packet flow involves learning how to switch packets only to the correct MACs. The switch learns which port an end station is connected to by remembering each packet's Source Address along with the port number on which the packet arrived.

When a packet is directed to a new, currently unlearned MAC address, the packet is transmitted out of all of the ports¹ except for the one on which it arrived². 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. This ensures that the packet is received by the correct end station (if it exists), and when the end station responds, its address is learned by the switch for the next series of packets. Unknown unicast MAC addresses can be prevented from egressing a port if required by clearing that port's Forward Unknown bit (see Port Control Register in [Table 53 on page 127](#)).

Owing to the limitation of physical memory, switches learn only the currently "active" MAC addresses—a small subset of the 2⁴⁸ possible MAC addresses. When an end station is moved from one port to another, a new MAC address/port number association must be learned, and the old one replaced. These issues are handled by the 'Aging' and 'Station Move Handling' functions. A MAC address/port number association is allowed to be "active" for only a limited time. This time limit is typically set to five minutes.

3.4.1 Address Translation Unit

The 88E6083 device's Lookup Engine or Address Translation Unit (ATU) gets the DA and SA 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. For example, 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 all ports.

The address database uses a hashing technique for quick storage and retrieval. Hashing a 48-bit address into fewer bits results in some MAC addresses having the same hash address. This situation is called a hash collision and is solved in the 88E6083 device by using a four entry bin per hash location that can store up to four different MAC addresses at each hash location. The four-entry bin is twice as deep as that of many competing switching devices and allows for a reduced size of address database while still holding the same number of active random value MAC addresses.

The address database is stored in the embedded SRAM and has a default size of 1024 entries with a default aging time of about 300 seconds (5 minutes). The size of the address database can be modified to a maximum of 512, 1024, or 2048 entries. Increasing the size of the address database reduces the number of buffers available for frames, however (see [Figure 10](#)). The age time can be modified in 16 second increments from 0 seconds (aging disabled) to 4080 seconds (68 minutes). These options are set in the ATU Control register ([Table 69 on page 143](#)).

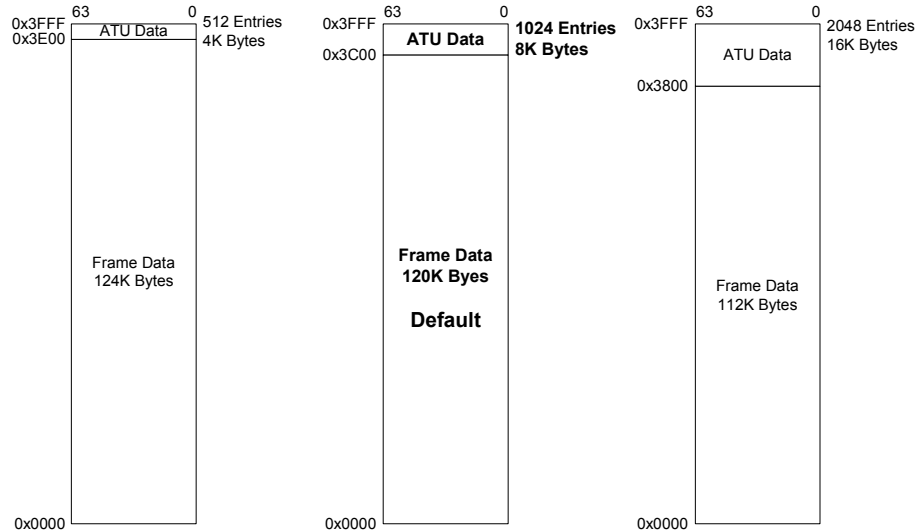
1. Port-based VLANs modify this operation ([Section 3.5.3](#)).
2. The 88E6083 device can be configured to transmit frames from the same port that they came in on—see [Table 54 on page 131](#).



Note

Changing the ATU size will result in ATU reset and SWReset. SWReset will reset the MAC and the PHY and ATU Reset will reset the entire database - See ATUSize and SWReset register descriptions in [Table 69 on page 143](#).

Figure 10: ATU Size Tradeoffs



3.4.2 Address Searching or Translation

The address search engine searches the address database to acquire the output port number, called the Destination Port Vector (DPV), for each frame's Destination Address (DA). When an address and its attendant DPV are found, the 88E6083 Switch Core can switch the frame to the appropriate port instead of flooding it. If the destination address is not found then the packet is flooded. Flooding refers to the action of sending frames out of all the ports of the switch that have link up with port state not "Disabled" except for the port on which the frame arrived. The switch arbitrates destination address lookup requests from the ports and grants one lookup at a time. The MAC address is hashed, and then data is read from the SRAM table and compared to the MAC address for a match. Four different addresses can be stored at each hash location. When a match is found, the Address Translation Unit (ATU) returns the DPV to the Ingress Policy block where it may get modified before the packet is queued to appropriate output ports. The DPV returned from the ATU may get modified by the VLANTable data or by 802.1Q processing—see section 3.5.3. When no MAC address match is found, the Ingress Policy block uses a unique default DPV for each Ingress port, which typically floods the frame. The default DPV for each port is the Port's VLANTable data — see [Table 54 on page 131](#). When 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. Multicast addresses cannot be learned, so they appear in the address database only when they are loaded into it by a CPU or the EEPROM — see section 3.5.4.3. This feature is used for multicast filtering and Bridge Protocol Data Unit (BPDU) handling. BPDU frames are special frames used for Spanning Tree or other bridge loop detection. Multiple separate address databases are supported in the 88E6083 device. The database that is searched is determined by the port's default database number, DBNum (see [Table 54 on page 131](#)) if 802.1Q is disabled on the port. If 802.1Q is enabled on the port, the database that is searched is determined by the DBNum value associated with the frame's VID (VLAN ID field of Tagged frames). MAC addresses that are not members of the port's or frame's DBNum cannot be found.

3.4.3 Address Learning

The address learning engine learns source addresses of incoming frames. Up to 2048 MAC address/port number mappings can be stored in the address database. See section 3.4.1. When the source address from an input frame cannot be found in the address database, the ATU enters the self-learning mode, places the new MAC address/port number mapping into the database, and refreshes its Age time. The Age time on a MAC Address entry is refreshed by setting its Entry_State field to 0xE. When the MAC address is 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 has moved, and the port number needs to be corrected. The entry's Age is refreshed since the MAC address is still "active". This refreshing prevents the MAC address/port number mapping from being prematurely removed as being "inactive".

When an address is added to the database it is hashed and stored in the first empty bin found at the hashed location. When all four address bins are full, a "least recently used" algorithm scans each entry's Age time in the Entry_State field. This field is described in section 3.4.5.1. If all four address bins have the same Age time, the first unlocked bin is used (see section 3.4.5.1 for more information about locked, i.e. static, addresses). If all four bins are locked, the address is not learned and an ATUFull interrupt is generated. See the Switch Global Status register—Table 60 on page 137. Multiple separate address databases are supported in the 88E6083 device.

Multiple separate address databases are supported in the 88E6083. In this device, the port's database number (DBNum—Table 54 on page 131) determines the address database into which the learned MAC address is stored 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) determines the address database into which the learned 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 for all ports by setting the LearnDis bit to a one in the ATU Control register (Table 70 on page 144). Learning is disabled on any port that has a PortState of Disabled or Blocking/Listening (see the Port Control register—Table 53 on page 127).

3.4.4 Address Aging

Address aging makes room for new active addresses by ensuring that when a node is disconnected from the network segment or when it becomes inactive, its entry is removed from the address database. An address is removed from the database after a predetermined interval from the time it last appeared as an Ingress frame's source address (SA). This interval is programmable in the 88E6083 device. The default Aging interval is about 5 minutes (282 to 304 seconds), but it can be set from 0 seconds (i.e., aging is disabled) to 4,080 seconds in 16 second increments. See AgeTime in ATU Control register – Table 70 on page 144.

The 88E6083 device runs the address aging process continuously unless disabled by clearing the AgeTime field in the ATU Control register to zero. Aging is accomplished by a periodic sweeping of the address database. The speed of these sweeps is determined by the AgeTime field. On each aging sweep of the database, the ATU reads each valid entry and updates its Age time by decrementing its Entry_State field as long as the entry is not locked. The Entry_State field is described in section 3.4.5.1. When the Entry_State field reaches zero, the entry is considered invalid and purged from the database.

The time taken to age out any entry in the MAC address database is a function of the AgeTime value in the ATU Control register and the value in the Entry_State field. A new or just-refreshed unicast MAC address has an Entry_State value of 0xE. See section 3.4.2. A purged or invalid entry has an Entry_State value of 0x0. The values from 0xD to 0x1 indicate the Age time on a valid unicast MAC address with 0x1 being the oldest. This scheme provides 14 possible age values in the Entry_State field which increases precision in the age of entries in the MAC address database. This precision is relayed to the "least recently used" algorithm that is employed by the address learning process. An address is purged from the database within 1/14th of the programmed AgeTime value in the ATU Control register.

3.4.5 Address Translation Unit Operations

The ATU in the 88E6083 device supports user commands to access and modify the contents of the MAC address database.

All ATU operations have the same user interface and protocol. Five Global registers are used and are shown in Table 22 on page 61. The protocol for an ATU operation is as follows:

1. Ensuring 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.
2. Loading the ATU Data and ATU MAC registers, if required by the required operation.
3. Starting the ATU operation by defining the required DBNum, ATUOp, and setting the ATUBusy bit to a one in the ATU Operation register. The DBNum, ATUOp and the ATUBusy bits setting can be done at the same time.
4. Waiting for the ATU operation to complete. Completion can be verified by polling the ATUBusy bit in the ATU Operation register or by receiving an ATUDone interrupt. See Switch Global Control register [Table 64 on page 139](#) and Switch Global Status register, [Table 60 on page 137](#).
5. Reading the required results, if appropriate.

Table 22: ATU Operations Registers

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
ATU Operation	0x0B or Decimal 11	Table 71 on page 145	Used to define the required operation (including which database to search, i.e., the DBNum field) and start it.	Used to indicate the ATU's Busy status.
ATU Data	0x0C or Decimal 12	Table 72 on page 146	Used further to define the required operation and used as the required ATU Data that is to be associated with the required MAC address below.	Returns the ATU Data that is associated with the resulting MAC address below.
ATU MAC (3 registers)	0x0D to 0x0F or Decimal 13 to 15	Table 73 on page 146	Used to define the required MAC address upon which to operate	Returns the resulting MAC address from the required operation.

3.4.5.1 Format of the ATU Database

Each MAC address entry in the ATU database is 64 bits in size. The lower 48 bits contain the 48-bit MAC address, and the upper 16 bits contain information about the entry as shown in Figure 11. The database is accessed 16 bits at a time via the Switch Global registers shown in the figure. For more information about these registers, see Table 73 through Table 75.

Figure 11: Format of an ATU Entry

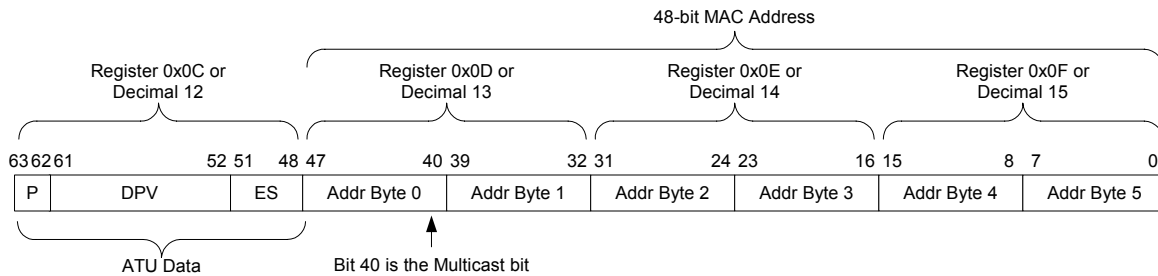


Table 23: ATU Data Fields

Field	Bits	Description
Entry_State	51:48	<p>The Entry_State field, together with the entry's Multicast bit (bit 40) is used to determine the entry's age or its type as follows:</p> <p>For unicast MAC addresses (bit 40 = 0):</p> <ul style="list-style-type: none"> 0x0 = Invalid, empty, or purged entry. 0x1 to 0xE = Valid entry where the Entry_State = the entry's age, and the DPV indicates the port or ports mapped to this MAC address. 0xF = Valid entry that is locked and does not age. The DPV indicates the port or ports mapped to this MAC address, and the Priority bits indicate the priority to use overriding any other priorities <p>For multicast MAC addresses (bit 40 = 1):</p> <ul style="list-style-type: none"> 0x0 = Invalid, empty, or purged entry. 0x5 = Valid entry that is locked and does not age. The DPV indicates the port or ports mapped to this MAC address. Frames with this MAC address are NOT limited by Ingress Rate Limiting—see Section 3.5.9. 0x7 = Valid entry that is locked does not age. The DPV indicates the port or ports mapped to this MAC address. Used for multicast filtering. 0xE = Valid entry that is locked and does not age. The DPV indicates the port or ports mapped to this MAC address. Frames with this Entry State are considered MGMT (management) frames and are allowed to tunnel through blocked ports (see Section 3.5), and the Priority bits indicate the priority to use overriding any other priorities. Used for BPDU handling. 0xF = Valid entry that is locked and does not age. The DPV indicates the port or ports mapped to this MAC address. The priority bits indicate the priority to use overriding any other priorities.

Table 23: ATU Data Fields (Continued)

Field	Bits	Description
DPV	61:52	The Destination Port Vector. These bits indicate which port or ports are associated with this MAC address (i.e., where frames should be switched to) when they are set to a one. A DPV of all zeros indicates that frames with this DA should be discarded. Bit 52 is assigned to physical Port 0, 53 to Port 1, 54 to Port 2, and so on.
Priority	63:62	These priority bits can be used as a frame's priority depending upon the value of the Entry_State (bits 51:48 above). When these bits are used, they override any other priority determined by the frame's data

3.4.5.2 Reading the Address Database - The Get Next Operation

The contents of the address database can be dumped or searched. The Get Next operation returns the active contents of the 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, the Get Next function returns all unique MAC addresses from all of the databases.

The Get Next operation starts with the MAC address contained in the ATU MAC registers and returns the next higher active MAC address currently active in the address database. Use 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 was found or that the Broadcast MAC address was found. This result indicates that the end of the database has been reached. If it were reached with a valid Broadcast address, the entry's Entry_State is returned with a non-zero value. A summary of how the Get Next operation uses the ATU's registers is shown in [Table 24](#).

Table 24: ATU Get Next Operation Register Usage

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
ATU Operation	0x0B or Decimal 11	Table 71 on page 145	Used to define the required operation (including which database to search, i.e., the DBNum field) and start it.	Used to indicate the ATU's Busy status.
ATU Data	0x0C or Decimal 12	Table 72 on page 146	Ignored.	Returns the ATU Data that is associated with the resulting MAC address below. If Entry_State = 0x0 the returned data is not a valid entry.
ATU MAC (3 registers)	0x0D to 0x0F or Decimal 13 to 15	Table 73 through Table 75	Used to define the starting MAC address to search. Use an address of all zeros to find the first 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 Entry_State ≠ 0x0).

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

3.4.5.3 Loading and 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 Entry_State in the ATU Data register ([Table 72 on page 146](#)) is non-zero. A value of zero indicates the required ATU operation is a purge.

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



Note

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

If the address is not found and if the ATU Data's Entry_State does not equal zero, the address is loaded into the address database using the same protocol as is used by automatic Address Learning. See section 3.4.3. The 16 bits of the ATU Data register are written into bits 63:48 of the ATU entry. See [Section 3.4.5.1](#).

A summary of how the Load operation uses the ATU's registers is shown in [Table 25](#).

Table 25: ATU Load/Purge Operation Register Usage

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
ATU Operation	0x0B or Decimal 11	Table 71 on page 145	Used to define the required operation (including which database to load or purge, i.e., the DBNum field) and start it.	Used to indicate the ATU's Busy status.
ATU Data	0x0C or Decimal 12	Table 72 on page 146	Used to define the associated data that is loaded with the MAC address below. When Entry_State = 0, the load becomes a purge.	No change.
ATU MAC (3 registers)	0x0D to 0x0F or Decimal 13 to 15	Table 73 through Table 75	Used to define the MAC address to load or purge.	No change.

3.4.5.4 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 Unlocked Entries
- Flush All Entries in a particular DBNum Database
- Flush all Unlocked Entries in a particular DBNum Database

The ATU Data and ATU MAC Address registers are not used for these operations and they are left unmodified.

The DBNum field of the ATU Operation register is used for the Flush operations that require a database number to be defined.

3.5 Ingress Policy

The Ingress Policy block modifies the normal packet flow through the switch. All ports have identical capabilities. The content of each frame is examined more deeply for Quality of Service (QoS) priority information, the frames are prevented from exiting from certain ports by the use of port-based VLANs or by the use of 802.1Q VLANs, and frames are prevented from entering the switch by the use of switch management Port States or by use of 802.1Q VLANs. An optional Trailer mode enables a management CPU to override the switch giving the CPU complete control over those frames that it selects. Rate limiting is also supported along with a Double Tagging option and IGMP snooping.

- An optional CPU header mode in which two bytes are prepended to the frame on the CPU's port, aligns the IP header information on a 32-bit boundary, thus accelerating router performance.
- The CPU header mode can be used to modify the port's Port based VLAN map ([Table 54](#)) at the beginning of a frame so the frame is processed with the updated setting. The ability to modify this register on a frame-by-frame basis allows the CPU to define the port of VLAN membership of each frame at wire speed.
- An optional Trailer mode enables a management CPU to override the switch giving the CPU complete control over those frames that it selects (used for Spanning Tree protocol).

3.5.1 Forward Unknown/Secure Port

The 88E6083 device 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 – [3.4.5 "Address Translation Unit Operations"](#)). The forwarding can be prevented on a per port basis (by clearing the Port Control Register's ForwardUnknown bit – [Table 53](#)) so that frames with unknown unicast addresses only go out of the port or ports where a server or router is connected.

This, together with the disabling of automatic address learning on a port ([Table 57](#)), enables a port to be configured as a Secure Port. A Secure Port is a port that allows communications to and from approved devices (by MAC address) only. In this mode, all approved devices must have their MAC addresses preloaded into the address database as 'static' ([3.4.5 "Address Translation Unit Operations"](#)) by a CPU.

When a new device tries to access the network through a Secure Port that new device's address will *not* be 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 will *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 unapproved new device and the rest of the network. Secure Port is like 802.1X without the authentication server and the associated interrupts to the CPU.



3.5.2 Quality of Service (QoS) Classification

The Ingress Policy block does not perform QoS switching policy; this task is performed by the Queue Controller. Instead, the Ingress Policy block determines the priority of each frame for the Queue Controller. The priority of a frame is determined in priority order by:

1. The CPU's Trailer when enabled on the port; See [Section 3.5.11](#).
2. The frame's DA; when that DA is in the address database with a defined priority. See section 3.4.5.
3. The IEEE 802.3ac Tag containing IEEE 802.1p priority information; this IEEE 802.1p priority information is used in determining frame priority when IEEE 802.3ac tagging is enabled on the port. See section 3.5.2.2.
4. The IPv4 Type of Service (TOS)/DiffServ field or IPv6 Traffic Class field when enabled on the port. IPv4/IPv6's priority classification can be configured on a per port basis to have a higher priority than IEEE Tag. The default state is "enabled" on all of the ports. See [Section 3.5.2.3](#).
5. The Port's default priority defined in DefPri in the Default Port VLAN ID & Priority register (see [Table 55](#), "Default Port VLAN ID & Priority," on page 132).

The designer can enable these classifications individually or in combination. For instance, if a 'hot', higher priority port is required for a switch design, priority classifications 1 to 4 can be disabled. This setting leaves only priority classification 5 active (the Port's default), which results in all ingress frames being assigned the same priority on that port.

- The priority classifications, except for #2, the DA, can be disabled separately on a per-port basis. These settings enable each port in the switch to be configured to use different rules for priority classification. See the UseTag and UseIP bits in the Port Control registers - [Table 53 on page 127](#).
- Priority classification #3 (IEEE Tag) and #4 (IP) can be reversed on a per port basis using the TagIfBoth bit in the Port Control register.

The CPU's Trailer, when enabled on the port, supersedes any other priority that may apply to the frame. Thus the CPU can have ultimate control over every frame that it transmits into the switch. The DA priority is used to set the highest priority on BPDUs or other MGMT frames since these frames should never be dropped.

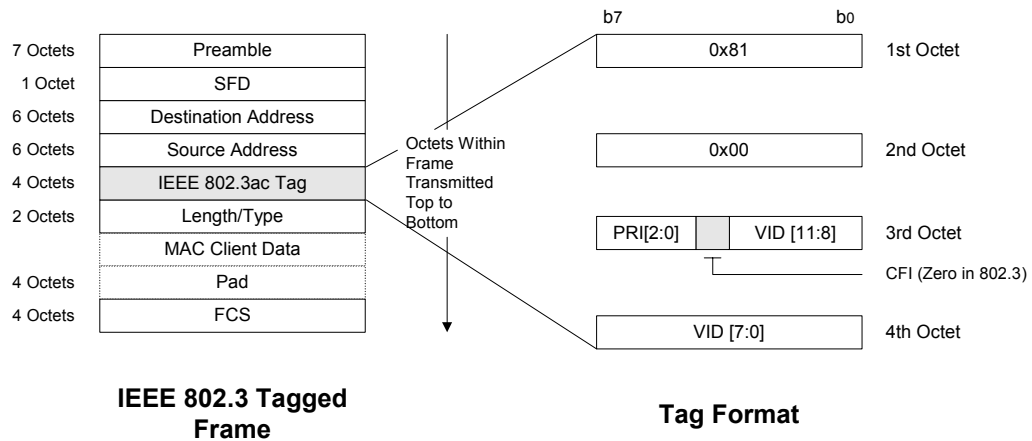
3.5.2.1 Forced Priority

All IEEE tagged frames entering a port can have their priority overridden and set to the port's default priority if none of the other priority classification rules are enabled on the port (i.e., UseIP and UseTag should be zero on the port—see Port Control Register - [Table 53 on page 127](#)). If these tagged frames egress the switch tagged, their 3-bit priority field in the tag is modified to the ingress port's default priority. This prevents an end user from using an unauthorized higher priority than that allocated.

3.5.2.2 IEEE Tagged Frames

The format of an IEEE Tagged frame is shown in Figure 12.

Figure 12: IEEE Tag Frame Format



The 88E6083 device captures the frame's PRI[2:0] bits and uses them to determine the frame's priority. The IEEE Tag supports eight priorities, while the 88E6083 device supports four. The Ingress Policy block takes PRI[2:0] and maps these bits into the two priority bits passed to the Queue Controller. This is done using the data found in the IEEE-PRI register (Table 84 on page 152). The CFI bit of the IEEE Tag is ignored by the 88E6083 device. The VID bits are ignored if 802.1Q tagging is disabled on this port (see section 3.5.3.3). The device's default condition is to capture and use IEEE Tagged frame priority data over IP priority data on all ports.

3.5.2.3 IPv4 and IPv6 Frames

The format of an IPv4 TOS/DiffServ frame is shown in [Figure 13](#), and an IPv6 Traffic Class frame is shown in [Figure 14](#).

Figure 13: IPv4 Frame Format

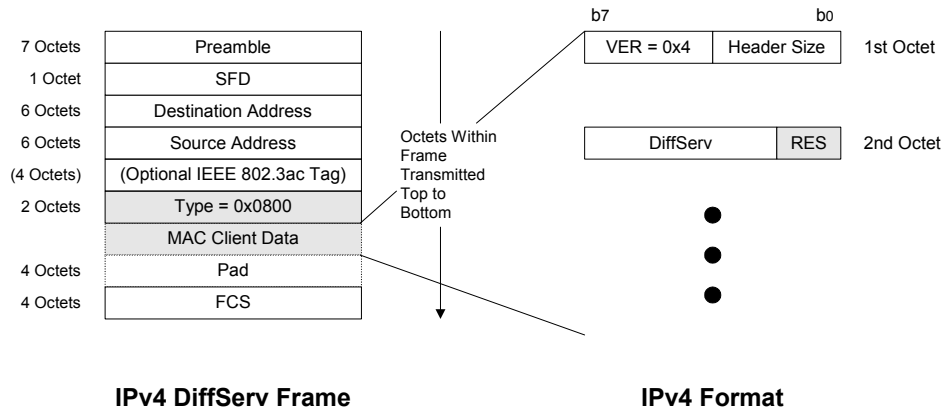
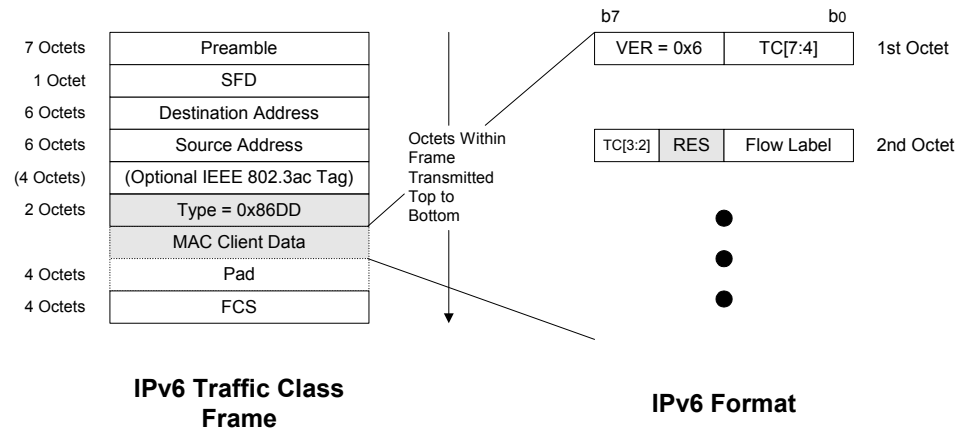


Figure 14: IPv6 Frame Format



The 88E6083 device captures the frame's DiffServ bits when it is an IPv4 frame or the frame's TC[7:2] bits when it is an IPv6 frame and uses them as the frame's priority. The DiffServ/TC bits support 64 priorities, while the 88E6083 device supports four. 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 IP-PRI registers ([Table 76 on page 148](#) and those following). The rest of the frame's IP bits are ignored by the 88E6083 device unless IGMP snooping is enabled in the port ([Section 3.5.6](#)). The 88E6083 device's default is to capture and use IP frame priority data in the absence of IEEE Tag priority data on all ports.

3.5.3 VLANs

The 88E6083 device supports many VLAN options including six 802.1Q options and one port-based VLAN option described in [Section 3.5.3.1](#).


3.5.3.1 Port Based VLANs

A very flexible port-based VLAN feature is enabled when:

- The port's 802.1QMode is disabled in the Port Based VLAN Map register ([Table 54 on page 131](#))
- Fallback is selected for the port's 802.1Q mode and the frame's VID is not contained in the VLAN Translation Unit (VTU)—see [Section 3.5.3.3](#).
- When any 802.1Q mode is enabled and ProtectedPort is enabled on the port ([Table 53](#))

Each Ingress port is associated with the VLANTable field of the Port-based VLAN Map register that restricts which egress ports its frames may use. ([Table 54 on page 131](#)). When bit 0 of a port's VLANTable field is set to a one, that port is allowed to send frames to Port 0. When bit 1 of this field is set to a one, that port is allowed to send frames to Port 1. If bit 2 equals 1, Port 2 may receive frames, etc. At reset the VLANTable value for each port is set to all ones, except for each port's own bit, which is cleared to a zero. This prevents frames from going back out of the port at which they arrived. This default VLAN configuration allows all of the ports to send frames to all of the other ports as shown in [Figure 15](#).

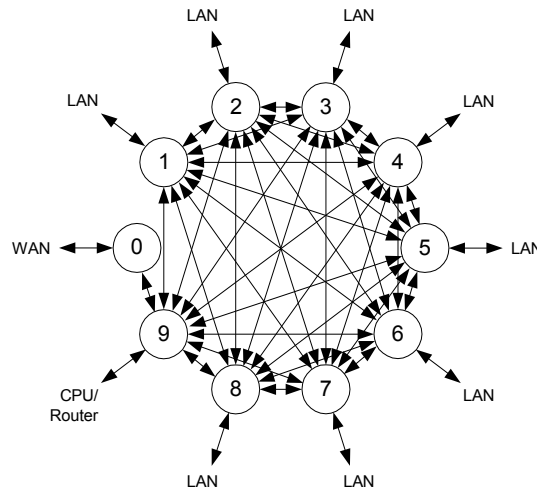
Figure 15: Switch Operation with VLANs Disabled



For unusual situations that require ingress frames to be reflected out of their ingress port, that port's own bit in the VLANTable field of the Port-based VLAN Map register can be set to a one—see [Section 3.5.5](#) and [Table 54 on page 131](#).

One reason for VLAN support in the 88E6083 device is to isolate a port for firewall router applications. Figure 16 shows a typical VLAN configuration for a firewall router. Port 0 is the WAN port. The frames arriving at this port must not go out to any of the LAN ports, but they 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 for CPU intervention but they cannot send frames directly to the WAN port. To accomplish routing, the CPU is able to send frames to all of the ports. The use of the Marvell header (Section 3.5.10) enables a CPU to define dynamically to which port or ports a frame is allowed to go. This feature is useful for purposes of WAN and LAN isolation on multicast or flooded traffic generated by the CPU.

Figure 16: 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 16

Port #	Port Type	VLANTable Setting
0	WAN	0x200
1	LAN	0x3FC
2	LAN	0x3FA
3	LAN	0x3F6
4	LAN	0x3EE
5	LAN	0x3DE
6	LAN	0x3BE
7	LAN	0x37E
8	LAN	0x2FE
9	CPU	0x1FF

To demonstrate the flexibility of the 88E6083 device VLAN configuration options, [Figure 17](#) shows another example. In this case, the switch is divided into two independent data paths with one data path containing three independent VLANs connected to a common router.

Figure 17: Switch Operation with another Example VLAN Configuration

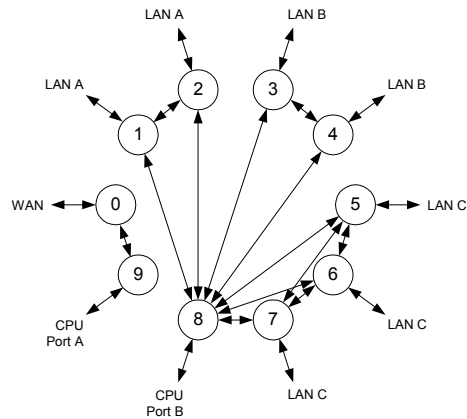


Table 27: VLANTable Settings for [Figure 17](#)

Port #	Port Type	VLANTable Setting
0	WAN	0x200
1	LAN A	0x104
2	LAN A	0x102
3	LAN B	0x110
4	LAN B	0x108
5	LAN C	0x1C0
6	LAN C	0x1A0
7	LAN C	0x160
8	CPU B	0x0FE
9	CPU A	0x001

3.5.3.2 Tunneling Frames Through Port-Based VLANs

Normally frames cannot pass through the port based VLAN barriers. However, some frames can be made to pass through the VLAN barriers on the 88E6083 device. Before a frame can tunnel through a port based VLAN barrier, its destination address (DA) must be locked into the address database (Section 3.4.5.1), and the VLANTunnel bit on the frame's Ingress port must be set to a one (Table 53 on page 127). When both of these conditions are met, the frame is sent out of the port or ports indicated in the locked address's 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. The VLANTunnel feature is not supported in the ProtectedPort case (i.e., when both the 802.1Q VLAN membership and the Port Based VLANs are being used on the frame - see Section 3.5.3.5).

3.5.3.3 802.1Q VLANs

The 88E6083 device supports 802.1Q tagging with up to 64 different VLAN identifiers (VIDs) out of the possible 4096. Any 64 VID's can be used. Since the device can contain only a subset of the possible VID's and security requirements vary, the 88E6083 device supports 802.1Q in modes called Secure, Check, and Fallback. The Secure and the Check modes completely ignore the device's Port Based VLAN features (Section 3.5.3.1), but these features are still accessible in the Fallback mode. In addition, the 88E6083 device supports three more modes, Protected Secure and Protected Check¹, in which both the frames' 802.1Q tags and the ports' VLANTable are used in making decisions about transmitting frames.

Security & Port Mapping

The 802.1Q Security features of the 88E6083 Switch Core ensure that ingress frames that meet the set criteria are only sent to the selected ports and that all other frames are discarded. One selected level of security out of six possible can be set independently on each port. The security options are processed using the ingress frame's VID or the ingress port's Default VID² 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 members of the frame's VLAN.
- Protected Secure – The VID must be contained in the VTU and the Ingress port must be a member of the VLAN or the frame will be discarded. The frame is allowed to exit only those ports that:
 - are members of the frame's VLAN *and*
 - are ports this source port can send frames to defined by the source port's VLANTable (i.e., port based VLANs - Section 3.5.3.1).
- Check - The VID must be contained in the VTU else the frame is discarded (the frame is not discarded if the ingress port is not a member of the VLAN). The frame is allowed to exit only those ports that are members of the frame's VLAN.
- Protected Check – The VID must be contained in the VTU or the frame is discarded (the frame is not discarded if the ingress port is not a member of the VLAN). The frame is allowed to exit only those ports that:
 - are members of the frame's VLAN *and*
 - are ports to which this source port can send frames defined by the source port's VLANTable (i.e., port based VLANs - Section 3.5.3.1).
- Fallback - Frames are not discarded if their VID's are 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 members of the frame's VLAN; otherwise the switch 'falls back' into Port Based VLAN mode for that frame (Section 3.5.3.1).

1. Secure, Check, and Fallback are controlled by the 802.1Q mode bits (see Table 54 on page 131). Protected Secure and Protected Check and Protected Fallback is enabled by selecting Secure or Check or Fallback as the port's 802.1Q mode (Table 54 on page 131) and by setting the port's protected port bit to a one (Table 53 on page 127)

2. The port's Default VID is used if the frame is not 802.3ac Tagged, if the frame's VID is 0x000, or if the port's ForceDefaultVID bit is set (see Table 55 on page 132).

- Protected 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 VLAN Table (Section 3.5.3.1)Otherwise the switch falls back into Port based VLAN mode for the frames (Section 3.5.3.1).

Secure, Check, Fallback or 802.1Q Disabled (i.e., port based VLANs only) modes for the port are controlled by the port's 802.1Q Mode bits (Port-based VLAN map register—see Table 54). Protected modes are enabled by selecting Secure or Check or Fallback in the port's 802.1QMode bits, and setting the port's ProtectedPort bit (Port Control registers - Table 53).

Security Violations

If a port is 802.1Q enabled, security violations are captured and an interrupt can be generated to the CPU (if unmasked by the VTUProblntEn bit—Table 63 on page 138 This is true regardless of the 802.1Q mode (interrupts are not generated if a port's 802.1Q mode is 802.1Q disabled). The interrupts (one at a time) are captured by the VLAN Translation Unit (VTU)—see Section 3.5.4). Two kinds of security violations are captured. A Miss Violation 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 the VID that caused the violation. This data is accessed by executing a Get/Clear Violation Data operation in the VTU.

3.5.3.4 Security Override of a Frame's VID

Each frame entering the switch must have a VID assigned to it. This VID is used for the frame's 802.1Q handling if the port is so enabled and is also used as the frame's VID if the frame is sent out of the switch tagged.

In the normal case, a tagged ingress frame's VID is extracted and used to process the frame through the switch. If an ingress frame is untagged, it is assigned a DefaultVID upon entry (Table 55 on page 132). However, the 88E6083 Switch Core supports a VID override function that ignores a tagged frame's VID and assigns the port's DefaultVID instead. This security feature ensures that all frames from a specific ingress port (tagged or untagged) exit the switch with the ingress port's DefaultVID. This feature prevents spoofed tags' being used to gain unauthorized access through the switch.

This feature is enabled on a per-port basis by setting the port's ForceDefaultVID bit to a one (Table 55). It is active if 802.1Q is enabled on the port or not.

3.5.3.5 Protected Port

802.1Q VLANs can be further limited by only allowing frames to leave ports that are also members of the Ingress port's port-based VLAN. The feature is enabled by setting the Protected Port bit in the port's Port Control Register (Table 53 on page 127). When this bit is set frames entering this port will not be allowed to egress from ports whose bits are not set in this port's VLANTable field of the register (Table 54).

3.5.4 VLAN Translation Unit Operations

The VLAN Translation Unit (VTU) in the 88E6083 device supports user commands to access and modify the contents of the VLAN membership database.

All VTU operations employ Global registers (shown in Table 28) and involve the same user interface and protocol. The protocol for a VTU operation comprises the following sequence:

1. Ensuring that 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.
2. Loading the VTU Data and VTU VID registers if required by the required operation.

3. Starting the VTU operation by defining the required DBNum, VTUOp, and setting the VTUBusy bit to a one in the VTU Operation register - these steps can be made simultaneously.
4. Waiting for the VTU operation to complete. This can be done by polling the VTUBusy bit in the VTU Operation register or by receiving a VTUDone interrupt. See the definitions for the Switch Global Control register [Table 63 on page 138](#)) and for the Global Status register ([Table 59 on page 136](#)).
5. Reading the required results if appropriate.

Table 28: VTU Operations Registers

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Table 65	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.
VTU VID	0x06	Table 66	Used to further define the required operation and used as the VID that is to be operated on.	Returns the VID from the required operation.
VTU Data	0x07, 0x08 and 0x09	Table 67 through Table 69	Used to define the required data that is to be associated with the required VID.	Returns the associated data from the required operation.

3.5.4.1 Format of the VTU Database

Each VID entry in the VTU database contains a 12-bit VID, a 4-bit Database number (DBNum) and 4 bits of VTU Data per port as shown in [Figure 18](#). The database is accessed 16 bits at a time via the Switch Global registers shown in the figure (not all of the register bits are shown). For more information about these registers, see [Table 64 on page 139](#) through [Table 68 on page 143](#)

Figure 18: Format of a VTU Entry

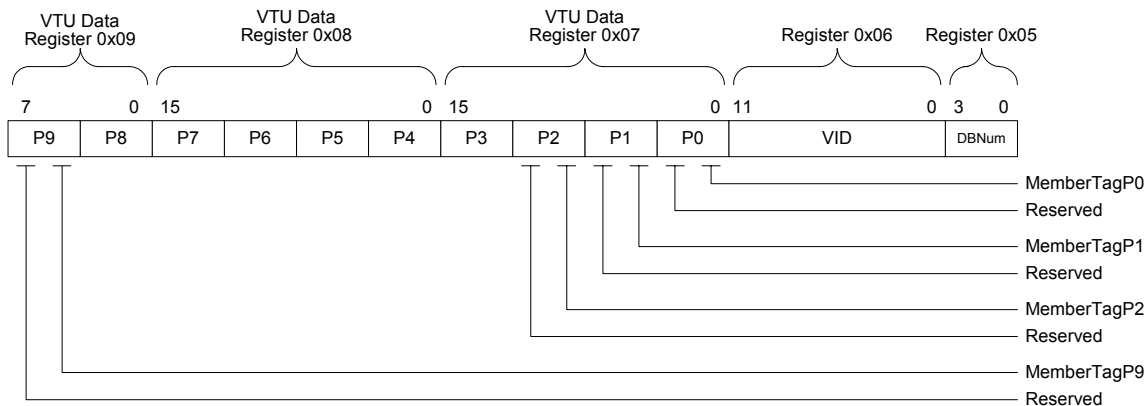


Table 29: VTU Entry Format

Field	Bits	Description
DBNum	3:0 in Register 0x5	Database Number. If separate address databases are used, these bits indicate the address database number to use for all frames assigned with the above VID (Section 3.4.2). All MAC address DA look-ups and SA learning are referred to the address database number defined by the DBNum associated with the frame's VID. Multiple VIDs can use the same DBNum. If separate address databases are not used the DBNum must be written as zeros.
VID	11:0 in Register 0x6	VLAN ID. These bits indicate the VID number that is associated with the port's MemberTag data and the address database number (DBNum).
Member TagP[2:0]	Port 0 = 1:0 Port 1 = 5:4 Port 2 = 9:8 Port 3=14:12 in Register 0x07 Port 4=1:0 Port 5=5:4 Port 6=9:8 Port 7=13:12 In Register 0x08 Port 8=1:0 Port 9=5:4 In Register 0x09	The lower 2 bits of each port's VTU data is called MemberTag. These two 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. The full definition of these bits is covered in VTU Data Register offset 7 bits 1:0. NOTE: The upper 2 bits of each port's VTU data is defined as Reserved. These bits must be written as zeros.



3.5.4.2 Reading the VLAN Database

The contents of the VLAN database can be dumped or searched. The dump operation is called Get Next since it returns the active contents of the VLAN database in ascending order. A search 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 currently active 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 registers: VTU Operation, VTU VID, and VTU Data. 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 0xFFFF was found. In either case, the end of the database has been reached. If it were reached with a valid VID of 0xFFFF, the entry's Valid bit is returned set to one (the Valid bit is contained in the VTU VID register—see [Table 65 on page 140](#)). A summary of how the Get Next operation uses the VTU's registers is shown in [Table 30](#).

Table 30: VTU Get Next Operation Register Usage

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Table 65	Used to define the required operation and start it.	Used to indicate the VTU's Busy status.
VTU VID	0x06	Table 66	Used to define the starting VID to search. Use a 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 returned indicating the end of the table has been reached. All ones is a valid entry if Valid = 0x1.
VTU Data (3 registers)	0x07, 0x08, and 0x09	Table 67 through Table 69	Ignored.	Returns the VTU Data that is associated with the VID above. If Valid = 0x0, 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 required VID that is to be searched. If the searched VID is found, it is returned in the VTU VID register along with its associated data in the VTU Data register and VTU Operations register. If the searched VID is not found active, the VID register contents do not equal the searched address.

3.5.4.3 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 operation. A VID is loaded into the database if the Valid bit in the VTU VID register ([Table 66 on page 142](#)) is a one. A value of zero indicates that the required VTU operation is a purge, and that the defined VID and its data is to be removed from the database.

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



Note

A load operation becomes a purge operation if the VTU Valid bit equals zero.

If the VID is not found and the VTU Valid bit does not equal zero, then the VID is loaded into the VLAN database assuming there is room. If the VTU is full, the load completes without any changes being made to the VLAN database, and the VTU Full Violation bit is set in the VTU Operation register. An interrupt can be generated when this event occurs ([Table 64 on page 139](#)).

A summary of how the load operation uses the VTU's registers is shown in [Table 31](#).

Table 31: VTU Load/Purge Operation Register Usage

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Table 65	Used to define the required 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	Table 66	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, 0x08, and 0x09	Table 67 through Table 69	Used to define the associated data that is loaded with the VID above.	No change.

3.5.4.4 Flushing Entries

All VIDs in the VLAN database can be purged by a single Flush All Entries VTU Operation. The VTU VID and VTU Data register are not used by the Flush command.



3.5.4.5 Servicing VTU Violations

The VTU captures VID Member Violation, VID Miss Violation, and VTU Full 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). A VTU Full Violation occurs when the CPU attempts to load a VID into the database when there is insufficient free space.

Captured VTU Violations and their associated interrupts are cleared by the Get/Clear Violation Data VTU Operation. This VTU Operation returns the ID of the source port that caused the violation in the DBNum/SPID field of the VTU Operation register (Table 64 on page 139) and returns the VID that caused the violation in the VID field of the VTU VID register (Table 65 on page 140). A SPID value of 0xF indicates the CPU port caused the violation (i.e., it was a VTU Full Violation).

A summary of how the Get/Clear Violation Data operation uses the VTU's registers is shown in Table 32.

Table 32: VTU Get/Clear Violation Data Register Usage

Register	Offset	Section	Before the Operation Starts	After the Operation Completes
VTU Operation	0x05	Table 65	Used to define the required operation (including which database) and start it.	Used to indicate the VTU's Busy status and the source of the violation.
VTU VID	0x06	Table 66	Ignored.	Used to indicate the VID that was involved in the violation.
VTU Data (3 registers)	0x07, 0x08, and 0x09	Table 67 through Table 69	Ignored.	No change.

3.5.5 Switching Frames Back to their Source Port

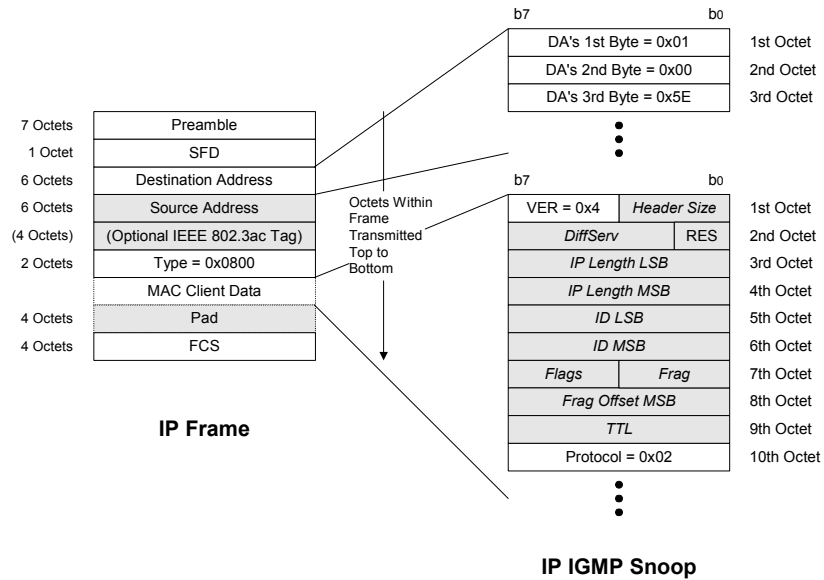
The 88E6083 device supports the ability to return frames to the port at which they arrived. While this is not a standard way to handle Ethernet frames, some applications may require this ability on some ports. This feature can be enabled on a port-by-port basis by setting the port's own bit in its VLANTable register to a one. See section 3.5.3 and Table 54 on page 131.

This function is valid whether or not 802.1Q tagging is enabled on the port.

3.5.6 IGMP Snooping

The 88E6083 device supports IGMP snooping. It is used to direct certain frames to the CPU for processing. The required format of these frames is shown in Figure 19.

Figure 19: IGMP Snoop Format



When one of these frames enters a port where IGMP snooping is enabled, the frame is sent to the CPU's port instead of where the Destination MAC Address would have sent it. VLAN membership (802.1Q) and Port Based VLAN rules, as appropriate, are still followed. The CPU port must be a member of the frame's VLAN or it will not receive the frame. IGMP snooping is enabled on a per port basis by setting the IGMPsnoop bit to a one (in the port's Port Control register—see Table 53 on page 127). The CPU's receiving port is defined by the port's Ingress-Mode settings also found in the Port Control register. Any port whose IngressMode is 01 (CPU Port with Ingress Trailer mode), or 11 (CPU Port without Ingress Trailer mode) is considered a CPU port and will receive the redirected IGMP frames. Typically only one port in the switch is configured in this way.



3.5.7 Port States

The 88E6083 device supports four Port States per port as shown in Table 33. The Port States are used by the Queue Controller (section 3.6) in the 88E6083 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 other bridge loop detection software can be supported. The PortState bits in the Port Control register (Table 53 on page 127) determine each port's Port State, and they can be modified at any time.

Table 33 below lists the Port States and their function. Two of the Port States require the detection of MGMT frames. A MGMT frame in the 88E6083 device is any multicast frame whose DA is locked into the address database with an Entry_State value of or 0x6 (section 3.4.5.1). MGMT frames can tunnel through blocked ports. MGMT frames always ignore VLANs (802.1Q and Port-Based), IGMP snooping, and Ingress Rate Limiting (i.e., they always go to the port indicated by the DA's DPV). 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 33: 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 or leave a blocked port. All other frame types are discarded. Learning is disabled on blocked ports.
Learning	Only MGMT frames are allowed to enter or leave 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 and leave 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 14). The ports should 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 after port-based VLANs are configured so a Spanning tree protocol can be run (if required).

3.5.8 Ingress Double VLAN Tagging

Double Tagging is a way to isolate one IEEE 802.1Q VLAN from other IEEE 802.1Q VLANs in a hierarchical fashion that is compatible with IEEE 802.1Q-aware switches as long as those switches support a maximum frame size of 1526 bytes or more. This method places an extra or Double Tag in front of a frame's normal Tag (assuming the frame were already tagged), increasing the frame size by 4 bytes. The Double Tag frame format is shown in Figure 20.

Ingress Double Tagging is selectable on a port-by-port basis by setting the port's IngressMode bits in its Port Control register (Table 53 on page 127). Typically any port that has Ingress Double Tagging enabled also has Egress Double Tagging enabled (see Section 3.7.4).

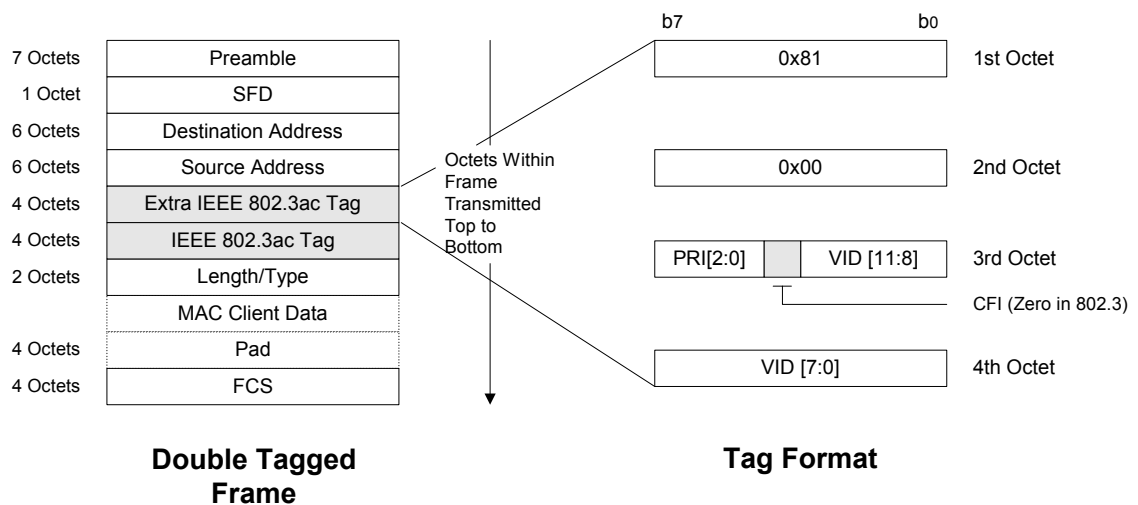
An Ingress port that has Double Tagging enabled expects all ingress frames to contain an extra Tag that needs to be removed from the frame prior to performing the port's ingress policy on the frame. In this mode, the ingress port removes the first IEEE 802.3ac Tag that appears after the Source Address in every frame. If the frame is untagged, it is not modified. When the frame has a single Tag, it is removed. If the frame has two Tags, the first Tag is removed. The CRC on modified frames is updated. The port's ingress policies of frame size checking and priority determination are performed *after* the first Tag is removed (if it were present). Thus, Tagged frames must be at least 68 bytes in size, so as not to be considered undersized frames after the Tag is removed. Frame padding is not performed on Ingress Double Tagging.



Note

All data from the removed Tags is ignored by the switch.

Figure 20: Double Tag Format



3.5.9 Ingress Rate Limiting

A switch design may need to limit the reception rate of multicast frames along with unknown or flooded unicast frames, or it may need to limit the rate for all frames but still keep QoS. The 88E6083 device supports this feature on a per port basis by setting bits in the port's Rate Control register (Table 55 on page 132). The process to arrive at such a design is shown as follows:

1. The types of frames to limit must be determined. The 88E6083 device can limit all frames, only multicast and flooded unicast frames (including broadcast frames), only multicast frames (including broadcast), or just broadcast frames. Specific multicast addresses can override this limit setting if they are locked into the address database with an Entry_State value of 0x5 (Section 3.4.5 and Table 23). MGMT (management¹) frames can also be specifically excluded by clearing the port's LimitMGMT bit to a zero in the Rate Control register. Any frame that is not limited by the above rules is ignored in the rate calculations (i.e., their size is not counted towards the limit total).
2. The required maximum rate must be selected. The 88E6083 device supports seven different rate limits from 128 kbps to 8 Mbps. Ingress rate limiting can be disabled on this port by selecting the Not Limited option.
3. The maximum rate for higher priority frames must be selected. The 88E6083 device supports four different QoS priorities and each higher priority can be limited to the same rate as the next lower priority or its limit can be twice as much. This feature supports a maximum limit of 64 Mbps for priority-3 frames.
4. The bytes to count for limiting must be determined. The default setting is to include the frame's bytes from the beginning of the Preamble field to the end of the frame check sequence (FCS) field and including a 96-bit inter frame gap (IFG). The frame's preamble bytes (the 8 bytes prior to the DA) are included if the CountPre bit is set to a one in the port's Rate Control register. The frame's minimum IFG bytes (the 12-byte inter-frame gap after the FCS) are included if the CountIFG is set to a one so either or both of these can be excluded if required.



Note

Ingress rate limiting in this device is not recommended to be used for TCP/IP traffic. UDP limiting and storm prevention are supported.

3.5.10 Switch's Ingress Header

The CPU in a router must perform many functions. One of those functions is routing IP frames, and another is bridging frames between WAN ports and LAN ports. The 88E6083 device's Ingress Header mode increases the performance of both of these functions. Any 88E6083 port can be configured to support an Ingress Header by setting the Header bit in the port's Port Control register (Table 53 on page 127), but only the CPU's port should be configured in this way².

The Ingress Header accelerates the CPU's performance when routing IP frames by aligning the IP portion of the frame with 32-bit boundaries. This is accomplished by prepending the frame with two extra bytes of zeros. Bridging between WAN and LAN ports sometimes requires the switch to support multiple address databases (one for the LAN and one for the WAN) so the same MAC address can be used on multiple ports. Since the CPU is generally a member of all VLANs, it must notify the switch which VLAN to use for 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 21³. When the Ingress Header is seen with a non-zero value, its contents are written to the port's Port Based VLAN Map register (Table 54 on page 131) 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 VLAN and address database to use on every frame at wire speed.

1. BPDU frames are MGMT frames (section 3.8).

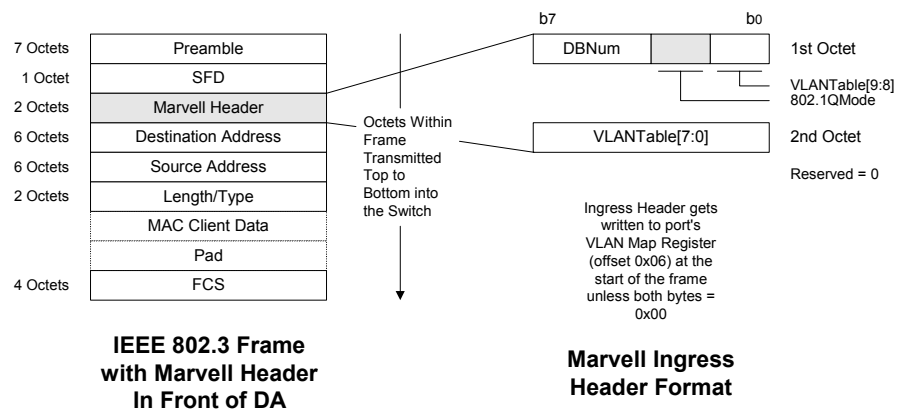
2. The Header bit enables the Header mode for both Ingress and Egress.

3. Reserved bits in the Header must always be zeros.

When the Ingress Header mode is enabled on a port, the first two bytes of the frame (just before the DA) are used to direct the switch action. The Ingress Policy block removes the Header from the frame and overwrites the frame's CRC with a new one (causing the frame to be two bytes smaller in size). 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 that the CPU must always add two bytes of data to the beginning of every frame that it sends into the switch (if the Header mode is enabled).

The Ingress Header gives the CPU the ability to control which VLAN, VLAN Mode, and address database to use on the frame it just received. However, the switch's register may not need to be changed by the CPU. If the CPU notifies the switch to process the frame based upon the switch's current ingress policy, 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 notifies the switch to ignore the header's data and process the frame normally (after the header is removed from the frame).

Figure 21: Ingress Header Format



3.5.11 Switch's Ingress Trailer

When a CPU needs to perform Spanning Tree or other bridge loop detection or when it must be able to send frames out of a special port or ports, the CPU's port must be configured for Ingress Trailer mode. Any port can be configured in this way by setting the IngressMode bits in the port's Port Control register (Table 53 on page 127), but only the CPU's port should be configured in this way

When the Ingress Trailer mode is enabled on a port, the last four bytes of the data portion of the frame are used to configure the switch—see Figure 22. The Ingress Policy block removes the Trailer from the frame and overwrites it with a new CRC, causing the frame to be four bytes smaller in size. This adjustment makes the frame 'normal' for the rest of the network since the Trailer'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 four bytes of data to the end of every frame it sends into the switch when the Trailer mode is enabled.

The Ingress Trailer gives the CPU the ability to override normal switch operation on the frame that it just transmitted but this override may not always be necessary. When the CPU requires the switch to process the frame with the switch's current ingress policy, the CPU sets the Trailer data in the frame to all zeros by inserting a four-byte pad there. This zero padding clears the Override bit in the Trailer causing the switch to ignore the Trailer's data and process the frame normally after removing the Trailer from the frame.

When the CPU needs to override all of the switch's Ingress policy, including priorities and Port Based VLANs among other functions, it sets the fields in the Trailer as shown in Figure 22 and defined in Table 34.

Figure 22: Ingress Trailer Format

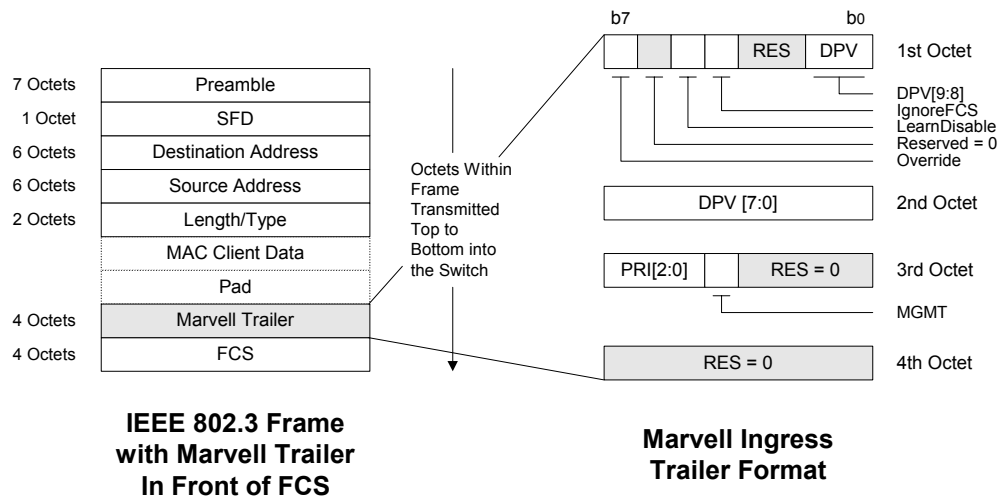


Table 34: Ingress Trailer Fields

Port State	Description
Override	When this bit is set to a one, the Trailer data is used to override the switch's operation. When this bit is cleared to a zero, the Trailer data is ignored.
LearnDisable	When this bit is set to a one (with or without the Override bit being set) the Source Address contained in this frame will not be learned in the address database.
IgnoreFCS	When this bit is set to a one (with or without the Override bit's being set) the frame's FCS is ignored (i.e., the FCS is not checked to see if the frame should be discarded owing to a CRC error). The FCS is still a required part of the frame, however, its value is not verified if this bit is set in the Trailer. A new, correct CRC is written to the last four bytes of the frame and the Frame is processed by the switch with the good CRC.
DPV[9:0]	The Destination Port Vector. These bits indicate which port or ports the frame is to be sent to if the Override bit is set to a one. A DPV of all zeros discards the frame. Bit 0 is assigned to physical Port 0 and must be set to a one for this frame to egress from that port, bit 1 for Port 1, bit 2 for Port 2, etc.
PRI[2:0]	The frame's priority. PRI[2:1] indicate the frame's priority if the Override bit is set to a one. PRI[0] is reserved for future use and it is ignored in this device.
MGMT	The frame's management bit. When this bit is set to a one (along with the Override bit) indicates the frame is a MGMT frame and is allowed to Ingress and Egress through Blocked ports (section 3.5.7). Must be set on BPDU frames.
RES = 0	These fields are reserved for future use and must be set to zeros.



3.6 Queue Controller

The 88E6083 device's queue controller uses an advanced non-blocking, four-priority output-port queue architecture with Resource Reservation. As a result, the 88E6083 device 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).

3.6.1 Frame Latencies

The 88E6083 device 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 lower priority frames are currently egressing the port. This is true regardless of the two priorities of the frames and the Scheduling¹ mode of the switch.

3.6.2 No Head-of-Line Blocking

An output port that is slow or congested never affects the transmission of frames to ports that are not congested. The 88E6083 device is designed to ensure that all flows that are not congested traverse the switch without degradation regardless of the congestion elsewhere in the switch.

3.6.3 QoS with and without Flow Control

The Queue Controller is optimized for two modes of operation – with and without Flow Control. When Flow Control is enabled, no frames are dropped, and higher priority flows receive higher bandwidth through the switch. These flows are less constrained if there is congestion between a higher and a lower priority flow. The percentage of bandwidth they receive is determined by the Scheduling mode (section 3.6.5). Flow Control prevents frames from being dropped, but it can greatly impact the available bandwidth on any network segment that is subject to flow control. The latency of higher priority data on flow-constrained network segments also increases since industry standard flow control mechanisms stop all frames from being transmitted. Therefore, flow control may not be acceptable in a QoS switch environment. For this reason, the 88E6083 device is optimized to work properly without Flow Control.

When Flow Control is disabled and congestion occurs for an extended period of time, frames are dropped. This is true in all switches. The 88E6083 device drops 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 (section 3.6.5).

1. The Scheduling mode selects either a Fixed priority or an 8-4-2-1 Weighted Fair Queuing - section 3.6.5.

3.6.4 Guaranteed Frame Delivery without Flow Control

The 88E6083 device can guarantee high-priority frame delivery¹, even with Flow Control disabled, owing to its intelligent Resource Reservation system. Having an output queue with multiple priorities is not a sufficient condition to support Quality of Service (QoS) when the higher priority frames cannot enter the switch owing to a lack of buffers. The 88E6083 device reserves buffers for higher priority frames so that they can be received and then switched. These high-priority buffers are replenished first from the Free Queue, which gets the switch ready for the next high priority frame.

3.6.5 Fixed or Weighted Priority

The 88E6083 device supports either a fixed-priority or weighted fair-queuing scheme. The selection is made by the Scheduling bit in the Switch Global Control register ([Table 63 on page 138](#)).

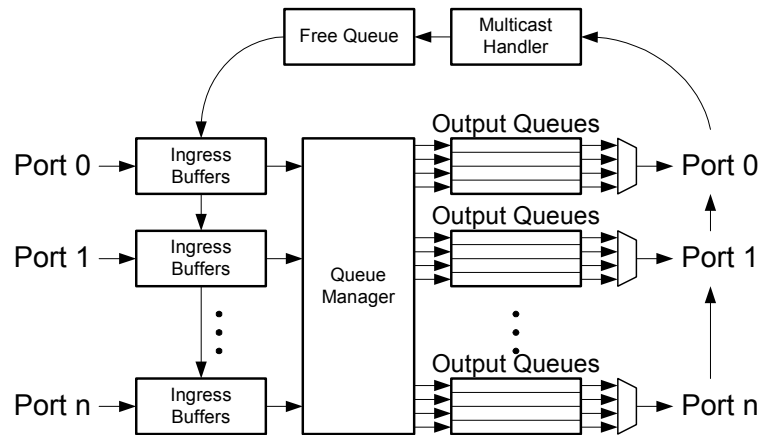
In the fixed-priority scheme, all top-priority frames egress a port until that priority's queue is empty, then the next lower priority queue's frames egress. This approach can cause the lower priorities to be starved of opportunity for transmitting any frames but ensures all high priority frames egress 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 of opportunity for transmission with only a slight delay to the higher priority frames.

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

3.6.6 The Queues

The queues in the 88E6083 device are shown in Figure 23.

Figure 23: Switch Queues



3.6.7 Queue Manager

At reset¹, the Queue Manager initializes the Free Queue by setting all of the buffer pointers to point into it and ensuring that all of the other queues are empty. The Queue Manager then takes the first available free buffer pointers from the Free Queue and assigns them to any Ingress port that is not disabled² and whose link³ is up. When these conditions are met, the switch is ready to accept and switch packets. Whenever any port's link goes down or 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 feature prevents "stale buffers" or "lost buffers" conditions and maximizes the size of the Free Queue so that of momentary congestion can be handled. When an enabled port's link comes back up, the port gets its Ingress Buffers back and it can start receiving frames again.

When a MAC receives a packet, it places it into the embedded memory at the address indicated by the input pointers that the MAC received from the Queue Manager. When packet reception is complete, the MAC transfers the pointers to the Queue Manager and requests new buffers from the Free Queue. When the Free Queue is empty, the MAC is not allocated any pointers until they become available. If the MAC starts to receive a packet when it has no pointers allocated, the packet is dropped. If flow control is enabled, it prevents this condition from occurring.

The Queue Manager uses the data returned from the Lookup Engine (section 3.4.1) and the Ingress Policy block (section 3.5) to determine to which output queues the packet's pointers should point and at what priority. At this point, the Queue Manager modifies the required 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. One port can have Flow Control enabled, while another has it disabled.

1. The Queue Manager is reset either by toggling the hardware RESETn pin, a software reset by the SWReset bit, or by an ATUSize change (both in the ATU Control register—Table 69 on page 143).
 2. When a port is in the Disabled Port State (Section 3.2.5), its Ingress buffers are left in the Free Queue for other ports to use.
 3. PHY based ports need to get a Link Up signal from the PHY. MII based ports have a link up state if they are enabled.

When Flow Control is enabled on an ingress port, the frame is switched to the required output queue without modification. This operation is done so that frames are not dropped. The Queue Manager monitors which output queues are congested and enables or disables flow control on the ingress ports that are causing the congestion. This approach allows flows that are not congested to progress through the switch without degradation.

When Flow Control is disabled on an ingress port, the frame can be discarded instead of being switched to the required Output Queue. If a frame is destined for more than one output queue, it can be switched to some queues and not to others. The decisions are quite complex because the Queue Manager takes many pieces of information into account before the decision is made.

The Queue Manager monitors the priority of the current frame, the current level of congestion in the output queues to which the frame is being switched, and the current number of free buffers in the Free Queue. As a result, uncongested flows traverse the switch unimpeded, and higher priority frames traverse the switch more quickly even when there is congestion elsewhere in the switch.

3.6.8 Output Queues

The Output Queues receive and transmit packets in the order received for any given priority. This is very important for some forms of Ethernet traffic. The Output Queues are emptied as fast as possible, but they can empty at different rates, possibly owing to a port's being configured for a slower speed, or because of network congestion (collisions or Flow Control).

Each 88E6083 port contains four independent Output Queues, one for each priority. The order in which the frames are transmitted out of each port is controlled by the Scheduling bit in the Switch Global Control registers ([Table 64 on page 139](#)). A fixed or a weighted priority can be selected ([Section 3.6.5](#)).

After a packet has been completely transmitted to the MAC, the Output Queue passes the transmitted packet's pointers to the Multicast Handler for processing, after which the MAC begins transmitting the next packet.

3.6.9 Multicast Handler

The Multicast Handler receives the pointers from all of the packets that are transmitted. It looks up each pointer to determine whether each packet was directed to more than one output queue. If not, the pointer is returned to the Free Queue where it can be used again. When the frame is switched to multiple output queues, the Multicast Handler ensures that the frame has exited all of the ports to which it was switched before returning the pointers to the Free Queue.

3.7 Egress Policy

The Egress Policy block is used to modify frames, if so directed, as they exit the switch. IEEE Tags can be added or removed or specific switch information can be added to the frame for the switch's CPU.

3.7.1 Tagging and Untagging Frames

Egress tagging and untagging is supported dynamically using 802.1Q VLANs or statically using port-based VLANs. The mode that is used on a port is determined by the port's 802.1Q Mode bits ([Table 54 on page 131](#)) as follows:

- Secure or Check - The MemberTag bits ([Table 67 on page 142](#)) associated with the VID assigned to the frame during ingress determine if the frame should egress unmodified, tagged, or untagged.
- Fallback - The MemberTag bits associated with the VID assigned to the frame during ingress determine whether the frame should egress unmodified, tagged, or untagged if the VID is contained in the VLAN database (section 3.5.4). If the VID is not found in the VLAN database, the frame will egress unmodified, tagged, or untagged determined by the port's EgressMode bits ([Table 53 on page 127](#)).
- 802.1Q Disabled - The port's EgressMode bits determine if the frame will egress unmodified, tagged, or untagged.



Note

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

The 88E6083 device performs the following on the egressing frame depending upon the egress tagging mode that was determined above. The format of an IEEE Tagged frame is shown in [Figure 24](#).

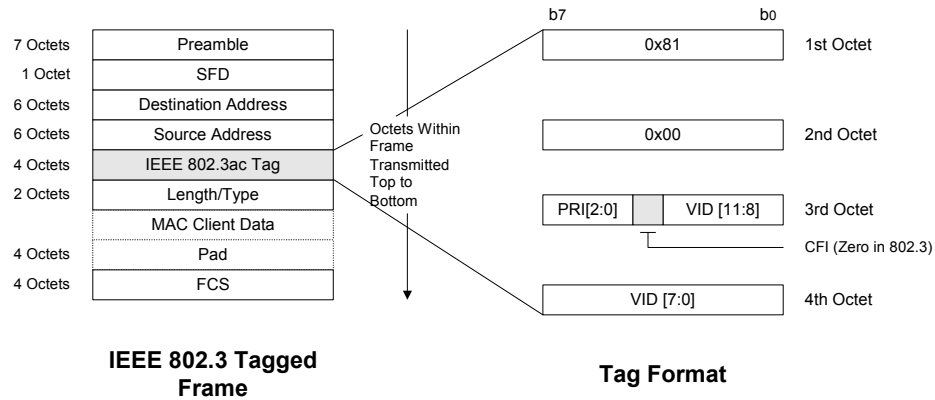
Table 35: Egress Port Configuration

Configuration	Result	
Leave the frames unmodified. This is the default setting so the switch acts as a transparent switch.	UnTagged frames egress the port UnTagged.	Tagged frames egress the port Tagged.
Transmit all frames UnTagged. Needed when switching frames to end stations that do not recognize Tags.	UnTagged frames egress the port unmodified.	The IEEE Tag on Tagged frames is removed, the frame is zero padded if needed, and a new CRC is computed for the frame.
Transmit all frames Tagged ¹ . Typically used when switching frames into the core or up to a server.	An IEEE Tag is added to UnTagged frames and a new CRC is computed.	Tagged frames egress the port unmodified.
Transmit all frames Double Tagged.	See section 3.7.4	

1. Tagged frames egressing tagged can be forced to be modified to use the source port's Default VID and/or the source port's Default Priority (see [Table 55 on page 132](#)).

The format of an IEEE Tagged frame is shown in [Figure 24](#).

Figure 24: IEEE Tag Frame Format



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

- The first octet is always 0x81.
- The second octet is always 0x00.
- PRI[2:1] indicate the Egress Priority Queue that the frame was switched through (section 3.6).
- The CFI bit is always set to a zero.
- PRI0 and VID come from the frame's source port's Default Port VLAN ID & Priority register ([Table 55 on page 132](#)). The frame's source port determines the Tag data so that when multiple ports are switching to one port, different Tags can be added.

3.7.2 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 ([Section 3.5.2](#)). The frame's PRI bits can be ignored and the ingress port's default priority used instead (DefPri—[Table 55](#) controlled by UseTag—[Table 53](#)).

3.7.3 VID 0x000 and VID Override

A Tagged frame egressing a port Tagged may have its VID bits modified. If the Ingressing frame's VID was 0x000 the Ingress port's DefaultVID ([Table 55](#)) is assigned to the frame instead. The Ingress port's DefaultVID can be forced as the VID to use on all frames entering a port. See [Section 3.5.3.4](#).

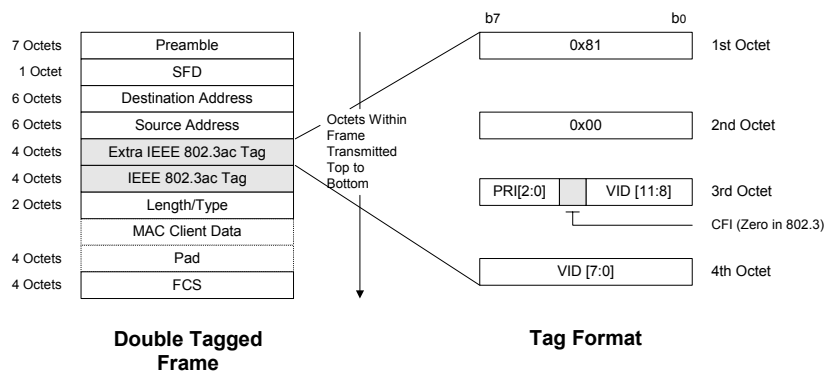
3.7.4 Egress Double VLAN Tagging

Double Tagging is a way to isolate one IEEE 802.1Q VLAN from other IEEE 802.1Q VLANs in a hierarchical fashion that is compatible with IEEE 802.1Q-aware switches as long as those switches support a maximum frame size of 1526 bytes or more. This method places an extra or Double Tag in front of a frame's normal Tag (assuming the frame were already Tagged) increasing the frame size by 4 bytes. The Double Tag frame format is shown in Figure 25.

Egress Double Tagging is selectable on a port-by-port basis by setting the port's EgressMode bits in its Port Control register (Table 53 on page 127). Typically, any port that has Egress Double Tagging enabled also has Ingress Double Tagging enabled (see Section 3.5.8).

An Egress port that has Double Tagging enabled transmits all Egress frames with an extra Tag. When a frame is UnTagged, it egresses Tagged. If a frame is Tagged, it egresses Double Tagged. The extra or Double Tag is inserted just after the frame's Source Address. This new Tag becomes the frame's first Tag. The Frame's Tag data comes from the same sources as a normal tag insertion described in Section 3.7.1.

Figure 25: Double Tag Format



3.7.5 Egress Rate Limiting

A switch design may need to limit the transmission rate of frames but still keep QoS. The 88E6083 device supports this capability on a per-port basis by setting bits in the port's Rate Control register ([Table 55 on page 132](#)). Egress rate limiting is performed by shaping the output load.

1. The type of frame that is to be limited must be determined. The 88E6083 device can limit all frames, or all but MGMT (management¹) frames. MGMT frames can be specifically excluded by clearing the port's LimitMGMT bit to a zero in the Rate Control register. Any frame that is not limited by the above rules is ignored in the rate calculations (i.e., its size is not counted toward the limit total).
2. The required maximum rate must be selected. The 88E6083 device supports seven different rate limits from 128 kbps to 8 Mbps. The selected rate will not be exceeded. Egress rate limiting can be disabled on this port by selecting the Not Limited option.
3. The bytes to count for limiting must be determined. The default setting is to include the frame's bytes from the beginning of the Preamble field to the end of the frame check sequence (FCS) field and including a 96-bit Inter Frame Gap (IFG). The frame's preamble bytes (the eight bytes prior to the DA) are included if the CountPre bit is set to a one in the port's Rate Control Register. The frame's minimum inter-frame gap (IFG, the 12 byte inter frame gap after the FCS) are included if the CountIFG bit is set to a one so either or both of these can be excluded if required.

1. BPDU frames are MGMT frames ([Section 3.5.7](#) and [Section 3.8](#)).

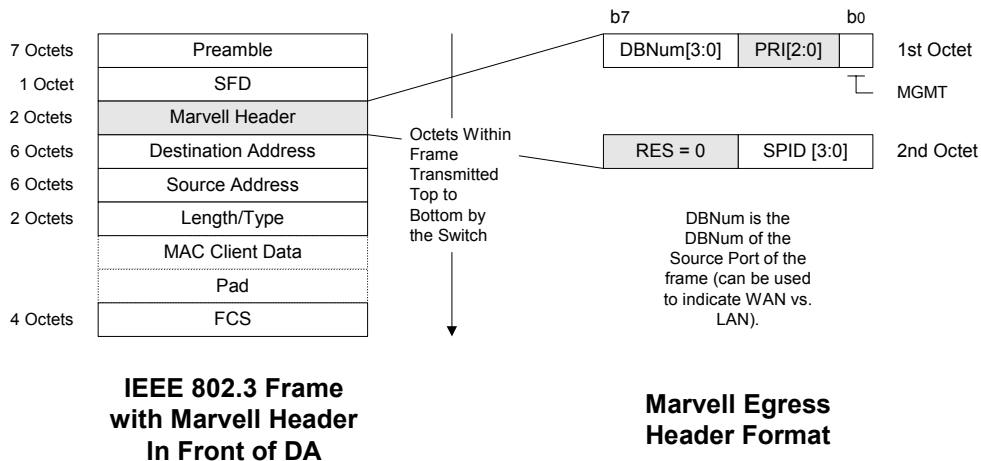
3.7.6 Switch's Egress Header

A CPU can have the IP frame data of the Ethernet frames aligned to 32-bit boundaries for faster routing. The switch supports a Marvell® Header that inserts two bytes into the frame just before the frame's Destination Address. Any port in the 88E6083 device can be configured in this way by setting the Header bit in the port's Port Control register (see [Table 53 on page 127](#)). In practice, only the CPU's port should be configured in this way¹.

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 occupies the first two bytes of the frame. If the CPU's MAC must 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. Routers generally do not filter frames, so the shifting of the frame's data will not create problems. These routers set their MACs in promiscuous mode allowing all frames to enter the CPU for processing.

The format of the Egress Header is shown in [Figure 26](#) and its fields are defined in [Table 36](#).

Figure 26: Egress Header Format



1. The Header bit enables the Header mode for both ingress and egress.

Table 36: Egress Header Fields

Field	Description
DBNum[3:0]	Database Number. This field represents the address database assigned to this frame at ingress into the switch (i.e., it is assigned by the source port). DBNum can be used to indicate the port based VLAN number of the source port.
PRI[2:0]	The frame's priority. PRI[2:1] indicate the frame's Egress Priority Queue the frame was switched through (section 3.6). PRI0 comes from the frame's source port's Default Port VLAN ID & Priority register (Table 55 on page 132).
MGMT	The frame management bit. When this bit is set to a one, it indicates that the frame is a MGMT frame which is allowed to ingress and egress through blocked ports (section 3.5.7).
RES	Reserved for future use. Currently set to 0x0.
SPID[3:0]	The Source Port ID. These bits indicate the physical port of entry of the frame. An SPID of all zeros indicates Port 0. An SPID of 0x1 indicates Port 1. 0x2 indicates Port 2, etc.

3.7.7 Switch's Egress Trailer

If a CPU needs to perform Spanning Tree or other bridge loop detection, the CPU must have information about the originating physical source port of its received frames, since those frames' SA information cannot be relied upon. To get this physical source port data, the CPU's port needs to be configured for Egress Trailer mode. Any port can be configured in this way by setting the TrailerMode bit in the port's Port Control register (Table 53 on page 127).

When the Egress Trailer mode is enabled on a port, four extra bytes are added to the end of the frame before the frame's CRC or FCS, and a new CRC is appended to the end of the frame. When the frame is received by the CPU, the MAC removes the CRC, so the Trailer occupies the last four bytes of the frame. The CPU's software can examine portions of the frame to determine if it needs the frame's source port information. Generally it is needed only on MGMT/Spanning Tree frames. If the information is not needed, the CPU reduces the frame size variable by four and passes the frame to the appropriate routines for processing. Removing the Trailer from a frame is essentially a subtraction operation without the need to move any data. Consequently CPU overhead is kept to a minimum.

The format of the Egress Trailer is shown in Figure 27 and its fields are defined in Table 37.

Figure 27: Egress Trailer Format

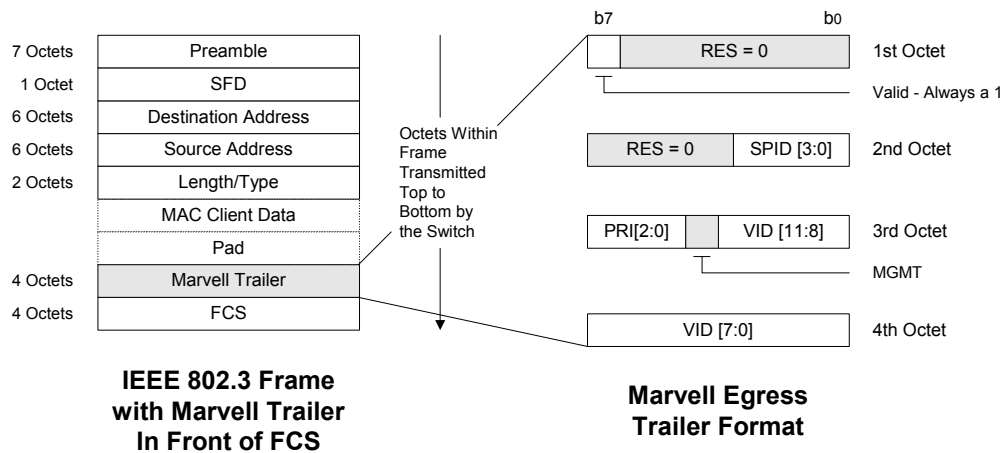


Table 37: Egress Trailer Fields

Field	Description
Valid	Always set to a one, indicating that Trailer data is present.
SPID[3:0]	Source Port ID. These bits indicate the physical port of entry of the frame. SPID[3:0] = 0 indicates Port 0. SPID[3:0] = 0x1 indicates Port 1, SPID[3:0] = 0x2 indicates Port 2, etc.
PRI[2:0]	The frame priority. PRI[2:1] indicate the frame's Egress Priority Queue that the frame was switched through (section 3.6). PRI[0] comes from the frame's source port Default Port VLAN ID & Priority register (Table 55).
MGMT	The frame management bit. When this bit is set to a one, it indicates that the frame is a MGMT frame which is allowed to ingress and egress through blocked ports (Section 3.5.7).
VID[11:0]	The VID field, comes from the frame's source port Default Port VLAN ID & Priority register (Table 55).
RES = 0	These fields are reserved for future use and must be set to zeros.

3.8 Spanning Tree Support

IEEE 802.1D Spanning Tree is supported in the 88E6083 device with the help of an external CPU that runs the Spanning Tree algorithm. The 88E6083 device supports Spanning Tree by:

- Detection of Bridge Protocol Data Unit (BPDU) frames. These frames are called MGMT (management) frames in the 88E6083 device. They are detected by loading the BPDU's multicast address (01:80:C2:00:00:00) into the address database with a MGMT Entry_State indicator (see Section 3.4.5.1).
- Tunnelling of BPDU frames through Blocked ports. Blocked ports are controlled by the Port's PortState bits (Section 3.2). When a port is in the Blocked state, all frames are discarded except for multicast frames with a DA that is contained in the address database with a MGMT indicator (see above).
- Redirection of BPDU frames. BPDU frames need to go to the CPU only, even though they are multicast frames. This is handled in the detection of BPDU frames above by mapping the BPDU's multicast address to the CPU port. (The value of the DPV bits when the address is loaded).
- Source Port information. The CPU needs information of the physical source port of origin of the BPDU frame. The source port is supplied in the frame's Egress Trailer that is sent to the CPU (Section 3.7.7).
- CPU transmission of BPDU frames. The CPU needs to be able to transmit BPDU frames out of any physical port of the switch. This is supported in the Ingress Trailer data that is supplied by the CPU (Section 3.5.11).

The 88E6083 device can support 802.1D Spanning Tree, or it can be used to perform simpler bridge loop detection on new link up. These different options are accommodated by running appropriate software on the attached CPU.

Any vendor's proprietary protocol can be handled with the same mechanism.

3.9 Embedded Memory

The 88E6083 device contains an embedded 1 Mb (16K x 64) Synchronous SRAM (SSRAM) that runs at 50 MHz, on a 64-bit wide data bus. The memory interface provides up to 3.2 Gbps bandwidth for packet reception and transmission and address mapping of data accesses. This memory bandwidth is enough for all of the ports running at full wire speed in full-duplex mode with minimum size frames.



3.10 Interrupt Controller

The 88E6083 device contains a switch Interrupt Controller that is used to merge various interrupts on to the CPU's interrupt signal. Each switch interrupt can be individually masked by an enable bit contained in the Switch Global Control register ([Table 63 on page 138](#)). When an unmasked interrupt occurs and the interrupt pin goes active low, the CPU needs to read the Switch Global Status register ([Table 59 on page 136](#)) to determine the source of the interrupt. When the interrupt comes from the switch core (from ATUFull, ATUDone, EEInt, etc.), the switch's interrupt pin goes inactive after the Switch Global Status register is read. The interrupt status bits are cleared on read. If the interrupt comes from the PHY (PHYInt), then the switch's interrupt pin goes inactive only after the PHY's interrupt is cleared by reading the appropriate registers in the PHY. See [Table 89](#) for details.

3.11 Port Monitoring Support

Port monitoring is supported by the 88E6083 device with Egress only monitoring or Egress and Ingress monitoring. Egress monitoring duplicates egress frames from a particular port to a selected monitor port. Ingress monitoring duplicates any good ingress frames for a particular port to a selected monitor port (such frames are processed normally through the switch). Port monitoring is enabled by modifying the Port Association Vector ([Table 56](#)) for the particular port that is to be monitored.

3.12 Port Trunking Support

Port Trunking is supported by the 88E6083 device with any combinations of ports. The ports that are to be associated with the trunk need to have all the port member's bits set in each of their Port Association Vectors ([Table 56](#)). Port based VLANs are then used to prevent loops and to load balance the trunk.

Section 4. Physical Interface (PHY) Functional Description

The 88E6083 device contains eight 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 to 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.

All PHY-port interfaces can be configured to support IEEE 802.3 100BASE-FX by utilizing a pseudo-ECL (PECL) interface for fiber-optics.

Figure 28: 88E6083 Device Transmit Block Diagram

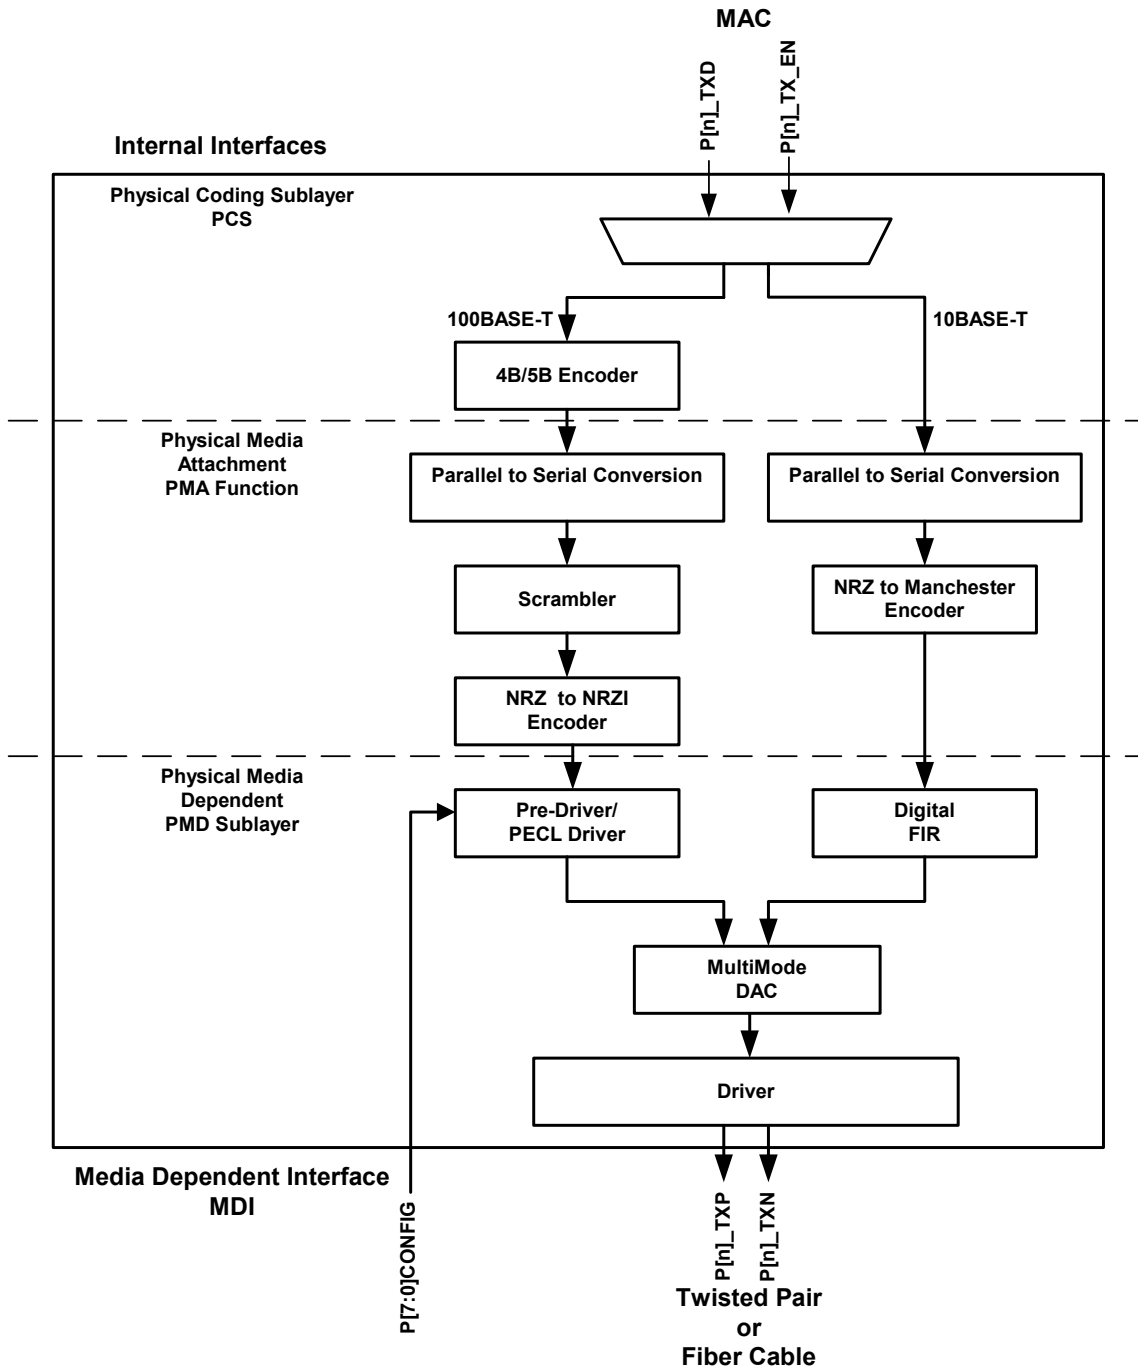
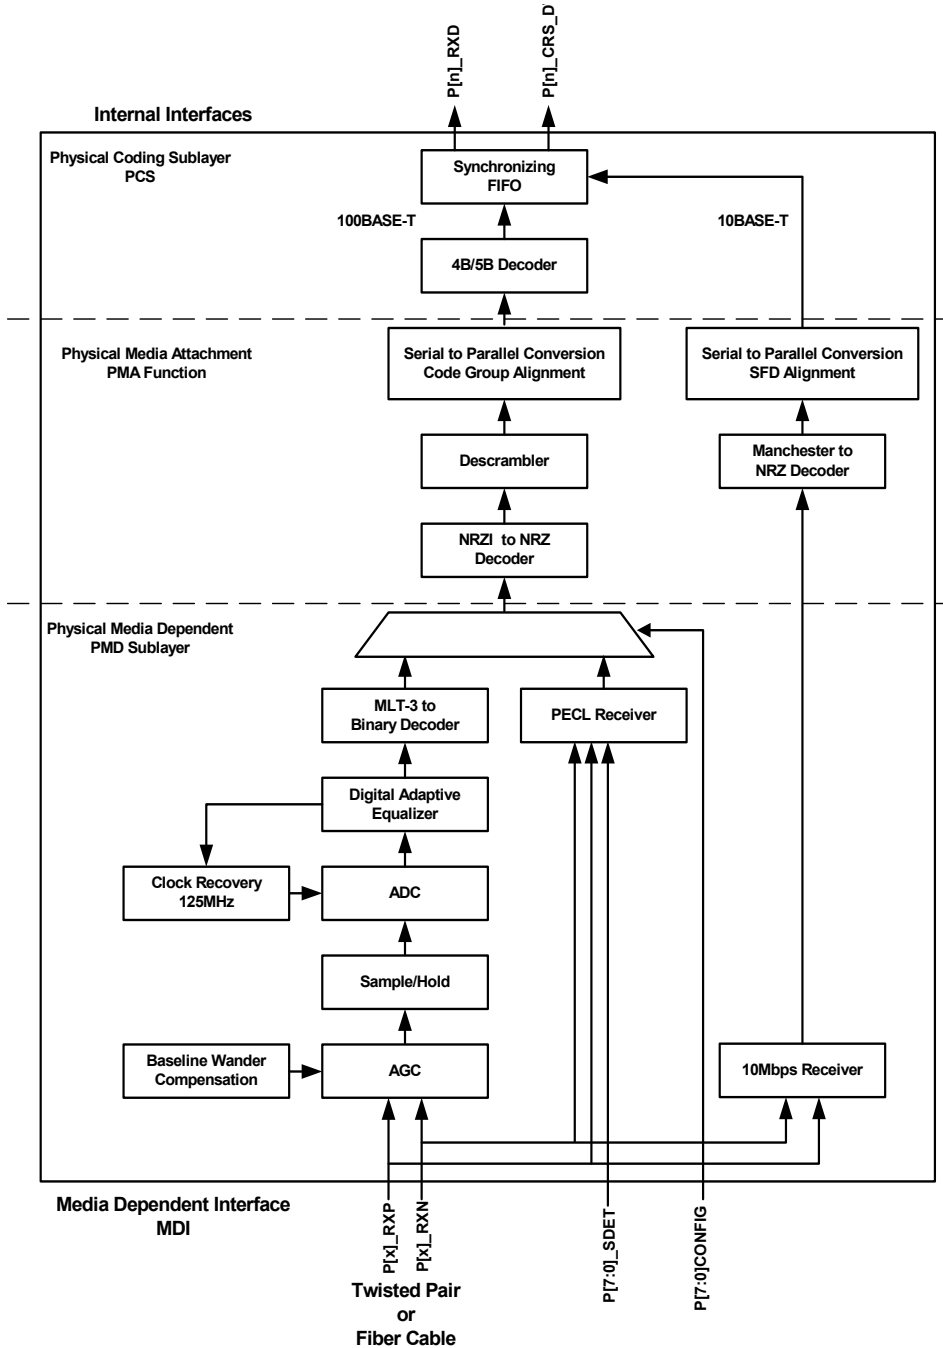


Figure 29: 88E6083 Device Receive Block Diagram





4.1 Transmit PCS and PMA

4.1.1 100BASE-TX Transmitter

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

4.1.2 4B/5B Encoding

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

4.1.3 Scrambler

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



Note

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

4.1.4 NRZ to NRZI Conversion

The data stream is converted from NRZ to NRZI.

4.1.5 Pre-Driver and Transmit Clock

The 88E6083 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 MLT3-coded symbols. The digital time base generator (TBG) produces the locked 125 MHz transmit clock.

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

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

4.1.6 Multimode Transmit DAC

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

4.2 Receive PCS and PMA

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

The differential RXP and RXN pins are shared by the 100BASE-TX, 100BASE-FX, and 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 RMII's.

4.2.2 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.

4.2.3 ADC and Digital Adaptive Equalizer

In 100BASE-T mode, an analog to digital converter (ADC) samples and quantizes the input analog signal and sends the result into the digital adaptive equalizer. 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.

4.2.4 Digital Phased Locked Loop (DPLL)

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

Digital interpolator clock recovery circuits are optimized for MLT-3, NRZI, and Manchester modes. A digital approach makes the 88E6083 receiver path robust in the presence of variations in process, temperature, on-chip noise, and supply voltage.

4.2.5 NRZI to NRZ Conversion

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

4.2.6 Descrambler

The descrambler is initially enabled upon hardware reset if 100BASE-TX is selected. The scrambler can be enabled or disabled via software by setting the descrambler bit ([Table 101](#)).

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

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

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

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

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

4.2.8 5B/4B Decoder

The 5B/4B decoder translates 5B code-groups into 4B nibbles to be presented to the MAC interfaces. The 5B/4B code mapping is shown in Table 38 on page 105.

4.2.8.1 FIFO

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

4.2.8.2 100BASE-FX Receiver

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

4.2.8.3 Far End Fault Indication (FEFI)

When 100BASE-FX is selected and the CONFIG_A pin enables FEFI at hardware reset, the far end fault detect (FEFD) circuit is enabled. The FEFD enable state can be overridden by programming the FEFI bit ([Table 101](#)).

**Note**

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

4.2.8.4 10BASE-T Receiver

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

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

In this mode, the recovered 10BASE-T signal is decoded from Manchester to NRZ data. The data stream is converted from serial to parallel format and aligned. The alignment is necessary to 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 ([Table 101](#)).

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 0 1 1	5	0 1 0 1	Data 5
0 1 1 1 0	6	0 1 1 0	Data 6
0 1 1 1 1	7	0 1 1 1	Data 7
1 0 0 1 0	8	1 0 0 0	Data 8
1 0 0 1 1	9	1 0 0 1	Data 9
1 0 1 1 0	A	1 0 1 0	Data A
1 0 1 1 1	B	1 0 1 1	Data B
1 1 0 1 0	C	1 1 0 0	Data C
1 1 0 1 1	D	1 1 0 1	Data D
1 1 1 0 0	E	1 1 1 0	Data E
1 1 1 0 1	F	1 1 1 1	Data F
1 1 1 1 1	I	Undefined	IDLE; used as inter-stream fill code
1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
0 1 1 0 1	T	Undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
0 0 1 1 1	R	Undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
0 0 0 0 0	V	Undefined	Invalid code
0 0 0 0 1	V	Undefined	Invalid code
0 0 0 1 0	V	Undefined	Invalid code



Table 38: 5B/4B Code Mapping (Continued)

PCS Code-Group [4:0] 4 3 2 1 0	Name	TXD/RXD <3:0> 3 2 1 0	Interpretation
0 0 0 1 1	V	Undefined	Invalid code
0 0 1 0 1	V	Undefined	Invalid code
0 0 1 1 0	V	Undefined	Invalid code
0 1 0 0 0	V	Undefined	Invalid code
0 1 1 0 0	V	Undefined	Invalid code
1 0 0 0 0	V	Undefined	Invalid code
1 1 0 0 1	V	Undefined	Invalid code

4.2.9 Setting Cable Characteristics

Since cable characteristics 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 (Table 101) for cable type.

4.2.10 Scrambler/Descrambler

The scrambler block is initially enabled upon hardware reset if 100BASE-TX is selected. If 100BASE-FX or 10BASE-T is selected, the scrambler is disabled by default. The scrambler is controlled by programming the Descrambler bit in the PHY Specific Control Register (Table 101).

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)	Descramble 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

4.2.11 Digital Clock Recovery/Generator

The 88E6083 device uses an all-digital clock recovery and generator circuit to create all of the needed receive and transmit clocks. The digital time base generator (TBG) takes the 25 MHz or 50 MHz reference input clock (XTAL_IN) and produces the locked 25 MHz transmit clock for the MAC in 100BASE-TX mode. It produces a 2.5 MHz transmit clock for the MAC in 10BASE-T mode as well as producing the over-sample clock for 10BASE-T waveshaping.

4.2.12 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 88E6083 device through SDETP/N pins.

If Force Link Good is asserted (ForceLink bit is set high - PHY Specific Control Register, [Table 101](#)), the link is forced to be good, and the link monitor is bypassed. Pulse checking is disabled if Auto-Negotiation is disabled, and DisNLPCheck (PHY Specific Control Register, [Table 101](#)) is set high. If Auto-Negotiation is disabled and DisNLPGen (PHY Specific Control Register, [Table 101](#)) is set high, then the link pulse transmission is disabled.

4.2.13 Auto-Negotiation

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

- Power up reset
- Hardware reset
- Software reset
- Restart Auto-Negotiation
- Transition from power down to power up
- Change from the *linkfail* state to the *link-up* state

If Auto-Negotiation is enabled, the 88E6083 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 88E6083 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 - [Table 91](#)). When Auto-Negotiation is disabled, the speed and duplex can be changed via the SpeedLSB and Duplex bits (PHY Control Register - [Table 91](#)), respectively. The abilities that are advertised can be changed via the Auto-Negotiation Advertisement Register ([Table 95](#)).



4.2.14 Register Update

Changes to the AnegEn, SpeedLSB, and Duplex bits ([Table 95](#)) do not take effect unless one of the following takes place:

- Software reset (SWReset bit - [Table 95](#))
- Restart Auto-Negotiation (RestartAneg bit - [Table 95](#))
- Transition from power down to power up (PwrDwn bit - [Table 95](#))
- Loss of the link

The Auto-Negotiation Advertisement register ([Table 95](#)) 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 88E6083 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 ([Table 99](#)) is internally latched once every time Auto-Negotiation enters the *next page exchange* state in the arbitration state machine.

4.2.15 Next Page Support

The 88E6083 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 ([Table 96](#)). The 88E6083 device has an option to write the received next page into the Link Partner Next Page register - [Table 100](#) - (similar to the description provided in the IEEE 802.3ab standard) by programming the Reg8NxtPg bit (PHY Specific Control Register - [Table 101](#)).

4.2.16 Status Registers

Once the 88E6083 device completes Auto-Negotiation it updates the various status in the PHY Status ([Table 102](#)), Link Partner Ability (Next Page) ([Table 97](#)), and Auto-Negotiation Expansion ([Table 98](#)) registers. Speed, duplex, page received, and Auto-Negotiation completed status are also available in the PHY Specific Status ([Table 102](#)) and PHY Interrupt Status registers ([Table 104](#)).

4.3 Power Management

The 88E6083 supports advanced power management modes that conserve power.

4.3.1 Low Power Modes

Two low power modes are supported in the 88E6083.

- 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 88E6083 to wake up when energy is detected on the wire with the additional capability to wake up a link partner. The 10BASE-T link pulses are sent once every second while listening for energy on the line.

4.3.2 MAC Interface and PHY Configuration for Low Power Modes

The 88E6083 has one CONFIG bit (in CONFIG_B - [Table 40](#)) dedicated to support the low power modes.

Low power modes are also register programmable. The EDet bit ([Table 101](#)) enables the user to turn on Energy Detect+™. When the low power mode is not selected, the PwrDwn bit ([Table 91](#)) 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 88E6083 returns 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 8 .	PwrDwn bit write (Table 91)
Energy Detect+™	Same as IEEE Power down mode—see Section 8 .	Configuration option & register EDet bit write (Table 101)



4.3.3 IEEE Power Down Mode

The standard IEEE power down mode is entered by setting the PwrDwn ([Table 91](#)) bit equal to one. In this mode, the PHY does not respond to any MAC interface signals except the MDC/MDIO. It also does not respond to any activity on the CAT 5 cable.

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

4.3.3.1 Energy Detect +™

In this mode, the PHY sends out a single 10 Mbps NLP (Normal Link Pulse) every second. If the 88E6083 is in Energy Detect+™ 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+™.

4.4 Far End Fault Indication (FEFI)

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

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

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

The FEFI function is controlled by CONFIG_A connection and the DisFEFI bit (Table 101).

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

Table 41: FEFI Select

Pn_CONFIG ¹	EN_FEFI (CONFIG_A)	FEFI	DisFEFI Bit (Table 101)
Auto-Negotiation	X	Disabled	HW reset set to 1
10BASE-T	X	Disabled	HW reset set to 1
100BASE-TX	X	Disabled	HW reset set to 1
100BASE-FX	0	Disabled	HW reset set to 1
100BASE-FX	1	Enabled	HW reset set to 0

1. "n" in Pn_CONFIG may only take the value "0" or "1".



4.5 Virtual Cable Tester®

The 88E6083 PHY 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 88E6083 Switch Core conducts a cable diagnostic test by transmitting a signal of known amplitude (+1V) sequentially along each of the TX and RX pairs of an attached cable. The transmitted signal continues along the cable until it is reflected from a cable imperfection. The magnitude of this echo signal and its return time are shown in the VCT™ registers (Table 111 and Table 112) on the AmpRfln and DistRfln bits respectively.

Using the information from these registers, the VCT registers (Table 111 and Table 112), the distance to the problem location and the type of problem can be determined. For example, the echo time can be converted to distance using Table 111 and Figure 36. The polarity and magnitude of the reflection together with the distance indicates the type of discontinuity. For example, a +1V reflection indicates an open close to the PHY and a -1V reflection indicates a short close to the PHY.

When the cable diagnostic feature is activated by setting the ENVCT bit to one (Table 111), a pre-determined amount of time elapses before a test pulse is transmitted. This is to ensure that the link partner loses link, so that it stops sending 10BASE-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 88E6083 device receives a continuous signal for 125 ms, it declares a test failure because it cannot start the TDR test. In the test failure case, the received data is not valid. The results of the test are also summarized in the VCTTst bits (Table 111 and Table 112).

- 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 can be user-defined using the information in the VCT registers. The impedance mismatch at the location of the discontinuity can also be calculated knowing the magnitude of the echo signal. Refer to the Application Note "Virtual Cable Tester -- How to use TDR results" for details.

4.6 Auto MDI/MDIX Crossover

The 88E6083 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 88E6083 device interoperates with a device that cannot automatically correct for crossover, the 88E6083 PHY makes the necessary adjustment prior to commencing Auto-Negotiation. If the 88E6083 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 88E6083 device interoperates with legacy 10BASE-T devices that do not implement Auto-Negotiation, it follows 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 88E6083 device uses signal detect to determine whether to implement crossover.

The Auto MDI/MDIX crossover function can be disabled via the AutoMDI[X] bits ([Table 101](#)).

The 88E6083 device is set to MDIX mode by default if auto MDI/MDIX crossover is disabled at hardware reset.

The pin mapping in MDI and MDIX modes is specified in [Table 42](#).

Table 42: 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



4.7 LED Interface

The LED interface pins can either be controlled by a port's PHY or controlled directly, independently of the state of the PHY. Four display options are available and both interfaces (PHY or direct control) can be used together to provide even more LED combinations.

The LEDs can be controlled through a parallel interface or via a serial interface that supports up to seven LEDs per port.

Direct control is achieved by writing to the PHY Manual LED Override register (Table 110). Any of the LEDs can be turned on, off, or made to blink at variable rates independent of the state of the PHY.

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

The 88E6083 device can be programmed to display serial LED status in single or dual LED modes. See Section 4.7.2.1 "Single and Dual LED Modes" on page 116. Some of the status indications 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.

Some port PHY events configured to drive LED pins are too short to result in an observable change in LED state. For these events, pulse stretching is provided to translate a short PHY event into an observable LED response. The duration of a pulse stretch can be programmed via the PulseStretch bits (Table 109). The default pulse stretch duration is set to 170 to 340 ms. The pulse stretch duration applies to all applicable LED interface pins.

Some of the status indicators signal multiple events by toggling LED interface pins resulting in a blinking LED. The blink period can be programmed via the BlinkRate bits (Table 109). The default blink period is set to 84 ms. The blink rate information is true for all applicable LEDs.

4.7.1 Parallel LED Interface

These LED interface pins can be used to display port status information. The LED interface has three different status indicators for each port with four different display options, using the Px_LED2, Px_LED1, and Px_LED0 pins. The LED Parallel Select Register (Table 107) specifies which single LED mode status to display on the LED interface pins. The default display for each mode is shown in Table 43 and the default mode that applies depends upon the hardware configuration established using the CONFIG_A pin - see Table 3.

Table 43: 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

Table 107 shows additional display modes that can be set up by software after startup. Table 44 includes these extra modes that cannot be set up by hardware configuration pins at reset. Refer to Table 107 for details.

Table 44: 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
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 84 ms active then 84 ms inactive). The transmit and receive activity is pulse stretched to 84 ms.
ACT (Blink mode)	High = No transmit or receive activity Blink = Transmit or receive activity (blink rate is 84 ms active then 84 ms inactive) The transmit and receive activity is pulse stretched to 84 ms

4.7.2 Serial LED Interface

The LEDSER, LEDENA, and LEDCLK pins are used for the serial interface.

However, the default displays in this case are different. The serial LED interface can display a variety of different status indications in 100BASE-TX and 10BASE-T modes (see Table 46 and Table 47). Status to display, pulse stretching, and blink speed can be programmed via the LED Stream Select for Serial LEDs register (Table 108)

4.7.2.1 Single and Dual LED Modes

Single LED mode is initiated at power up according to the CONFIG_A pin settings. Dual LED mode can be entered (or single mode can be reentered) by setting the appropriate register fields according to Table 108. These registers enable LED functions to be programmed individually in dual LED or single LED mode. The data states are sent out on the LEDSER data stream and may be extracted using the LEDCLK and LEDENA signals as shown in the example of Figure 30.

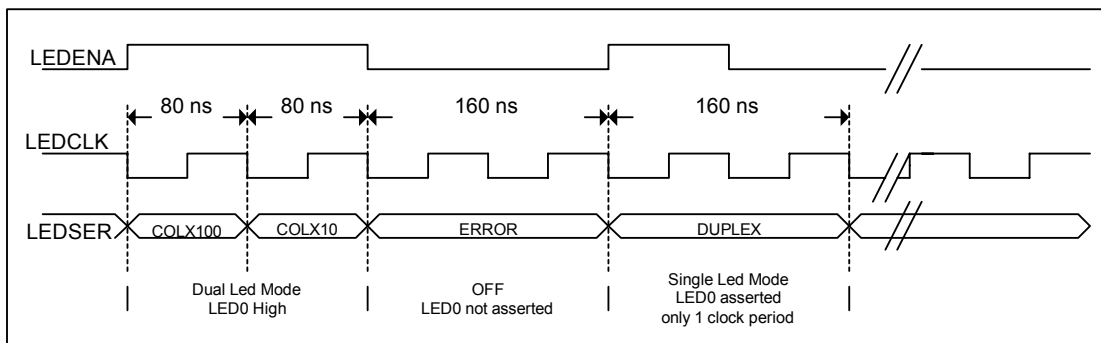
4.7.2.2 Single LED Modes

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, the LEDENA signal is high over Duplex100 position only for one clock period. Refer to Figure 30 for more details.

4.7.2.3 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 this stream. For example, the LEDENA signal asserts over COLX100 and COLX10 in Figure 30. and remains high for two clock periods. If one of these status bits (shown in Table 47) is turned off, LEDENA is not asserted in both positions.

Figure 30: 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 31. The rate of update of the serial LED interface is controlled by programming register 24.8:6. The default value is set to 42 ms.

Figure 31: 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 88E6083 device is self-synchronizing. The default display options are given in Table 45 on page 117.

Table 45: Serial LED Display Options (A = Active)

LED Mode	COLX	ERROR	DUPLEX	DUPLEX /COLX	SPEED	LINK	TX	RX	ACT	LINK /RX	LINK /ACT
0	-	A	-	A	A	-	-	-	-	-	A
1	A	A	A	-	A	-	A	-	-	A	-
2	A	A	A	-	A	A	-	-	A	-	-
3	A	A	A	-	A	A	A	A	-	-	-

The LED status bits are output in the order shown on the LEDSER pin synchronously with LEDCLK. All status signals for Port 7 are sent out first followed by those for Ports 6 through 0. Each bit in the stream occupies a period of 80 ns.

Figure 32: Serial LED Display Order

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

Table 46 and Table 47 show the status events that can be displayed by programming the 88E6083 device in single and dual LED display modes.



4.7.2.4 Single LED Display Mode

Table 46: 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 84ms inactive) The transmit and receive activity is pulse stretched to 170ms.

4.7.2.5 Dual LED Display Mode

Table 47: Dual LED Display Mode

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




Table 47: Dual LED Display Mode (Continued)

Event	Description
TX10	1 = No 10 Mbps transmit activity 0 = 10 Mbps transmit activity This status can be pulse stretched.
RX100	1 = No 100 Mbps receive activity 0 = 100 Mbps receive activity This status can be pulse stretched.
RX10	1 = No 10 Mbps receive activity 0 = 10 Mbps receive activity This status can be pulse stretched.
ACT100	1 = No 100 Mbps transmit or 100 Mbps receive activity 0 = 100 Mbps transmit or 100 Mbps receive activity This status can be pulse stretched.
ACT10	1 = No 10 Mbps transmit or 10 Mbps receive activity 0 = 10 Mbps transmit or 10 Mbps receive activity This status can be pulse stretched.
LINK/RX100	1 = 100 Mbps link down 0 = 100 Mbps link up Blink = 100 Mbps receive activity The receive activity can be pulse stretched. The blink rate can be programmed.
LINK/RX10	1 = 10 Mbps link down 0 = 10 Mbps link up Blink = 10 Mbps receive activity The receive activity is can be pulse stretched The blink rate can be programmed.
LINK/ACT100	1 = 100 Mbps link down 0 = 100 Mbps link up Blink = 100 Mbps transmit or 100 Mbps receive activity The transmit and receive activity can be pulse stretched. The blink rate can be programmed.
LINK/ACT10	1 = 10 Mbps link down 0 = 10 Mbps link up Blink = 10 Mbps transmit or 10 Mbps receive activity The transmit and receive activity can be pulse stretched. The blink rate can be programmed.

Section 5. Register Description

The 88E6083 registers are accessible using the IEEE Serial Management Interface (SMI) used for PHY devices (see MDC/MDIO in [Table 6 on page 22](#)). The 88E6083 device uses 19 of the 32 possible Device addresses. The allocation of these 19 device addresses is fixed and cannot be changed. The 88E6083 device does not respond to the 13 unused addresses so these addresses can be used by other devices that are sharing the Serial Management Interface. [Figure 33 on page 121](#) shows the register map.

Figure 33: 88E6083 Register Map





5.1 Register Types

The registers in the 88E6083 device are made up of one or more fields. The way in which each of these fields operate is defined by the field's Type. The function of each Type is described in [Table 48](#).

Table 48: Register Types

Type	Description
LH	Register field with latching high function. If status is high, then the register is set to a 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 a zero and remains cleared until a read operation is performed through the management interface or until a reset occurs.
RES	Reserved for future use. All reserved bits are read as zero unless otherwise noted.
RO	Read only.
ROC	Read only clear. After read, register field is cleared to zero.
R/W	Read and Write with no definable initial value.
RWR	Read/Write reset. All bits are readable and writable. After reset the register field is cleared to zero.
RWS	Read/Write set. All bits are readable and writable. After reset the register field is set to a non-zero value specified in the text.
SC	Self-Clear. Writing a one to this register causes the required function to be immediately executed, then the register field is cleared to zero when the function is complete.
Update	Value written to the register field does not take effect until a soft reset is executed.
WO	Write only. Reads to this type of register field return undefined data.

5.2 Switch Core Registers

Switch Core Registers are of two types:

- Switch Port Registers—are separately configured for each port comprised in the switch. Each port has a unique associated SMI device address and this address is used to access that port's switch register set.
- Switch Global Registers—are configured using a single SMI device address. A global register's setting affects all ports of the switch.

The 88E6083 device contains ten ports (MACs). The supported ports are accessible using SMI device addresses 0x10 to 0x19. The MACs are fully IEEE 802.3 compliant. Since there is no IEEE standard covering required MAC registers, these registers are 88E6083 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 SMI device address 0x1B.

5.2.1 Switch Core Register Map

Table 49: Switch Port Register (SMI Device Address) Map

Description	Offset Hex	Offset Decimal	Page Number
Port Status Register	0x00	0	page 125
Reserved Registers	0x01 - 0x02	1 - 2	page 126
Switch Identifier Register	0x03	3	page 126
Port Control Register	0x04	4	page 127
Reserved Registers	0x05	5	page 130
Port Based VLAN Map	0x06	6	page 131
Default Port VLAN ID & Priority	0x07	7	page 132
Reserved Registers	0x08 - 0x09	8 - 9	page 132
Rate Control	0x0A	10	page 133
Port Association Vector	0x0B	11	page 135
Reserved Registers	0x0C - 0x0F	12 - 15	page 135
Rx Counter	0x10	16	page 136
Tx Counter	0x11	17	page 136
Reserved Registers	0x12 - 0x1F	18 - 31	page 136



Table 50: Switch Global Registers

Description	Offset Hex	Offset Decimal	Page Number
Switch Global Status Register	0x00	0	page 137
Switch MAC Address Register Bytes 0 & 1	0x01	1	page 138
Switch MAC Address Register Bytes 2 & 3	0x02	2	page 138
Switch MAC Address Register Bytes 4 & 5	0x03	3	page 138
Switch Global Control Register	0x04	4	page 139
VTU Operation Register	0x05	5	page 140
VTU VID Register	0x06	6	page 142
VTU Data Register Ports 0 to 3	0x07	7	page 142
VTU Data Register Ports 4 to 7	0x08	8	page 143
VTU Data Register Ports 8 to 9	0x09	9	page 143
ATU Control Register	0x0A	10	page 144
ATU Operation Register	0x0B	11	page 145
ATU Data Register	0x0C	12	page 146
ATU MAC Address Register Bytes 0 & 1	0x0D	13	page 146
ATU Switch MAC Address Register Bytes 2 & 3	0x0E	14	page 146
ATU Switch MAC Address Register Bytes 4 & 5	0x0F	15	page 147
IP-PRI Mapping Register 0	0x10	16	page 148
IP-PRI Mapping Register 1	0x11	17	page 148
IP-PRI Mapping Register 2	0x12	18	page 149
IP-PRI Mapping Register 3	0x13	19	page 149
IP-PRI Mapping Register 4	0x14	20	page 150
IP-PRI Mapping Register 5	0x15	21	page 150
IP-PRI Mapping Register 6	0x16	22	page 151
IP-PRI Mapping Register 7	0x17	23	page 151
IEEE-PRI Register	0x18	24	page 152
Reserved Registers	0x19 - 0x1C	25 - 28	page 152
Stats Operation Register	0x1D	29	page 152
Stats Counter Register Bytes, 3 & 2	0x1E	30	page 154
Stats Counter Register Bytes 1 & 0	0x1F	31	page 154

5.2.2 Switch Port Registers

Table 51: Port Status Register
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Type	Description
15	LinkPause	RO	Link Partner's Pause bit, returned from the link partner through Auto-Negotiation. This bit is valid for Ports 0 to 7 only and when the Resolved bit is set to a one. 0 = MAC Pause not implemented in the link partner 1 = MAC Pause is implemented in the link partner
14	MyPause	RO	My Pause bit, sent to the link partner during Auto-Negotiation. This bit is valid for Ports 0 to 7 only. It is set high if FD_FLOW_DIS (Table 14) is low during RESETn. 0 = MAC Pause not implemented in the link partner 1 = MAC Pause is implemented in the link partner
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). 0 = Link is down 1 = Link is up
11	PortMode	RO	Port mode. The value of this bit is always a one for Ports 0 thru 7 and it comes from Px_MODE3 during configuration for Port 8 and Port 9. 0 = SNI mode or MII 200 ¹ mode is enabled 1 = MII 10/100 or RMII 100 Mbps mode is enabled
10	PHYMode	RO	PHY mode. This bit is valid for all ports but it is meaningful for Port 8 and Port 9 only and only when the PortMode (bit 11 above) indicates an MII mode of operation (i.e., PortMode is a one). This bit may be either value for SNI configured ports – but all SNI port configurations are PHY mode. The value of this bit is always zero for Ports 0 to 7 and it comes from Px_MODE2 during configuration for Port 8 and Port 9. 0 = Pins are in MII MAC Mode (i.e., INCLK and OUTCLK are inputs) or the pins are in RMII PHY Mode (all RMII Modes are PHY mode). 1 = Pins are in MII PHY Mode (i.e., INCLK and OUTCLK are output)

Table 51: Port Status Register (Continued)
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Type	Description
9	Duplex	RO R/W on Port 8 and Port 9 Only	Duplex mode. This bit is valid when the Resolved bit is set to a one. 0 = Half-duplex 1 = Full-duplex This bit can be written to on Port 8 and Port 9 only so that the Port's duplex can be modified after Reset occurs. Care must be taken to ensure this Port's duplex matches the duplex of its link partner.
8	Speed	RO	Speed mode. This bit is valid when the Resolved bit is set to a one and the PortMode (bit 11 above) is also a one. If the PortMode bit is a zero then the port's speed is 10 Mbps unless both this bit and the Duplex (bit 9 above) are both ones, then the speed is 200 Mbps. When the port is configured in MII MAC mode, the Speed bit is undefined. 0 = 10 Mbps 1 = 100 Mbps
7:0	Reserved	RES	Reserved for future use.

1. MII 200 mode is supported on Port 9 only. Enabling MII 200 mode on Port 9 forces Port 8 to be disabled.

The PortMode, PHYMode, Duplex, and Speed bits for Port 8 and Port 9 come from the Px_MODE[3:0] pins at Reset. When the PortMode bit is a one, the PHYMode, Duplex, and Speed bits map the same as the other 10/100 ports that contain PHYs. When the PortMode is a zero bit the mapping is easier to see by looking at [Table 18](#).



Note

Registers 0x01 and 0x02 (Hex), or 1 and 2 (Decimal) are reserved.

Table 52: Switch Identifier Register
Offset: 0x03 (Hex), or 3 (Decimal)

Bits	Field	Type	Description
15:4	DeviceID	RO	Device Identifier. The 88E6083 device is identified by the value 0x083.
3:0	RevID	RO	Revision Identifier. The initial version of the 88E6083 device has a RevID of 0x0. This RevID field may change at any time. Contact a Marvell FAE for current information on the device revision identifier.

Table 53: Port Control Register
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Type	Description
15	ForceFlow Control	RWR	Force Flow Control. This bit is used to enable full-duplex flow control or half-duplex back pressure on this port depending upon the port's current duplex. This bit can only be changed when the port is Disabled. 0 = Use configuration pin settings for flow control/back pressure 1 = Force flow control/back pressure to be enabled on this port
14	Trailer	RWR	Egress Trailer Mode 0 = Normal mode – frames are transmitted unmodified 1 = Add a Trailer to all frames transmitted out of the port. NOTE: Egress Trailer mode is intended for management CPU ports for source port information on BPDU frame.
13:12	Egress Mode	RWR	Egress Mode. The first three Egress Modes are used as the default Egress Mode for frames whose VID is not contained in the VTU (Table 66, on page 142). The fourth Egress 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 = Always add a Tag (or Egress Double Tag). This setting is not a default setting. It ignores the MemberTag data in the VTU for this port. NOTE: 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 the switch.
11	Header	RWR	Ingress & Egress Header Mode. When this bit is set to a one all frames Egressing out this port are pre-pended 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 pre-pended with the Marvell 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. If the frame's Ingress Header is non-zero it is used to update the port's VLAN Map register value (Table 54).



Table 53: Port Control Register (Continued)
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Type	Description
10	IGMP Snoop	RWR	<p>IGMP Snooping. When this bit is set to a one and this port receives an IPv4 IGMP frame, the frame is switched to the CPU port overriding the destination port(s) determined by the DA mapping. When this bit is cleared to a zero, IGMP frames are not treated specially.</p> <p>NOTES:</p> <ul style="list-style-type: none"> The CPU port is determined by the Ingress Mode bits. A port is considered the CPU port if its Ingress Mode (bits 9:8 in this register) are either 0b01 or 0b11. VLAN masking is performed on IGMP frames.
9:8	Ingress Mode	RWR	<p>Ingress Mode.</p> <p>00 = Normal mode – frames are received unmodified 01 = CPU Port with Ingress Trailer mode – frames must be received with a Trailer and the Trailer is checked, potentially used and then removed from the frame. Also used to identify the CPU port for IGMP Snooping. 10 = Remove IEEE 802.3ac tag if one is present (or Ingress Double Tag) 11 = CPU Port without Ingress Trailer mode - used to identify the CPU port for IGMP Snooping.</p> <p>NOTE: Trailer mode is intended for Ingress data control from a management CPU port so that BPDU frames can be directed.</p>
7	VLAN Tunnel	RWR	<p>VLAN Tunnel. When this bit is cleared to a zero, the VLANs defined in the VLANTable (Table 54) are enforced for ALL frames. When this bit is set to a one, the VLANTable masking is bypassed for any frame entering this port with a DA that is currently 'locked' in the ATU. This includes unicast as well as multicast frames.</p> <p>VLANs defined by the frame's VID membership are always enforced regardless of the value of the VLANTunnel bit.</p>
6	TagIfBoth	RWS	<p>Use IEEE 802.1p tag fields over IP fields. If the current frame is both IEEE 802.3ac tagged and an IPv4 or an IPv6 frame, and UseIP and UseTag bits below are both enabled, a priority selection between the two needs to be made.</p> <p>0 = Use IP fields for priority mapping when both field types are present 1 = Use Tag fields for priority mapping when both field types are present</p>
5	UseIP	RWS	<p>Use IP for Priority. Use IPv4 TOS and/or DiffServ fields and/or IPv6 Traffic Class fields, if present, for priority data.</p> <p>0 = Ignore IPv4 and IPv6 priority fields 1 = Use IPv4 fields when the frame is IPv4, and use IPv6 fields when the frame is IPv6 for priority data.</p>

Table 53: Port Control Register (Continued)
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Type	Description
4	UseTag	RWS	<p>Use IEEE Tags. Use IEEE 802.1p Traffic Class field for priority mapping when the frame is an IEEE 802.3ac tagged frame.</p> <p>0 = Ignore IEEE 802.1p tag fields even if the frame is tagged. In this mode tagged frames that enter this port and egress the switch tagged will get this port's default priority bits written to the tag, if the frame's IP priority is not used (i.e., the frame was not an IP frame or the UseIP bit above is a zero). This works as a Force Default Priority function on tagged frames.</p> <p>1 = Use IEEE 802.1p tag Traffic Class for priority mapping if the frame is tagged</p>
3	Protected Port	RWR	<p>Protected Port. When this bit is cleared to a zero normal switch operation occurs (i.e., frames are allowed to Egress ports defined by the 802.1Q VLAN membership for the frame's VID <i>or</i> by the port's VLANTable depending upon the port's 802.1Q Mode). When this bit is set to a one, frames are allowed to Egress ports defined by the 802.1Q VLAN membership for the frame's VID <i>and</i> by the port's VLANTable if 802.1Q is enabled on the port. Both must allow the frame to egress.</p>
2	Forward Unknown	RWS	<p>Forward Unknown. When this bit is set to a one, normal switch operation occurs and unicast frames with unknown DAs are allowed to egress this port (assuming the VLAN settings also allow the frame to Egress this port). When this bit is cleared to a zero, unicast frames with unknown DAs do not egress this port.</p>

Table 53: Port Control Register (Continued)
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Type	Description
1:0	PortState	RWR Or RWS to 0b11	<p>Port State. These bits are used to manage a port to determine what kind of frames, if any, are allowed to enter or leave a port for simple bridge loop detection or 802.1D Spanning Tree. The state of these bits can be changed at any time without disrupting frames currently in transit. The Port States are:</p> <p>00 = Disabled. The switch port is completely disabled and it will not receive or transmit any frames. The QC will return any pre-allocated ingress queue buffers when the port is in this mode.</p> <p>01 = Blocking/Listening. The switch examines all frames without learning any SA addresses, and discard all but MGMT 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.</p> <p>10 = Learning. The switch will examine all frames, learning all SA addresses, and still discard all but MGMT frames. It will allow MGMT frames only to exit the port.</p> <p>11 = Forwarding. The switch will examine all frames, learning all SA addresses, and receive and transmit all frames as a normal switch.</p> <p>When Port 9 is configured for 200 Mbits/s mode (see section 3.2.5 "MII/SNI Configuration") Port 8 is automatically set to Disabled. Software Reset, (SWReset—see Table 69) causes the PortState bits to revert back to their reset state (i.e., either disable or forwarding).</p> <p>NOTES:</p> <ul style="list-style-type: none"> The PortState bits for all ports come up in the Forwarding state unless the SW_MODE[1:0] pins are set to 0b00, the CPU attached mode. MGMT (management) frames are the only kind of frame that can be tunneled through Blocked ports. A MGMT frame is any frame whose multicast DA address appears in the ATU Database with the MGMT Entry State.



Note

Register 0x05 (Hex), or 5 (Decimal) is reserved.

Table 54: Port Based VLAN Map
Offset: 0x06 (Hex), or 6 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:12	DBNum	RWR	Port's Default DataBase VLAN Number. This field can be used with non-overlapping VLANs to keep each VLAN's MAC address mapping database separate from the other VLANs. This allows the same MAC address to appear multiple times in the address database (at most one time per VLAN) with a different port mapping per entry. This field is overridden by the DBNum returned from a VTU hit and it should be zero if not being used. It needs to be a unique number for each independent, non-overlapping, VLAN if used.
11:10	802.1Q Mode	RWR	IEEE 802.1Q Mode for this port. These bits determine if port based VLANs or 802.1Q based VLANs are used for this Ingress port. It also determines the action to be taken if an 802.1Q VLAN Violation is detected. These bits work as follows: 00 = 802.1Q Disabled. Use Port Based VLANs only. The VLANTable bits below determine which Egress ports this Ingress port is allowed to switch frames to for all frames. 01 = Fallback. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violations and use the VLANTable bits below if the frame's VID is not contained in the VTU (both errors are logged - Table 66). 10 = Check. Enable 802.1Q for this Ingress port. Do not discard Ingress Membership violation but discard the frame if its VID is not contained in the VTU (both errors are logged - Table 66). 11 = Secure. Enable 802.1Q for this Ingress port. Discard Ingress Membership violations and discard frames whose VID is not contained in the VTU (both errors are logged - Table 65)
9:0	VLANTable	RWS to all ones except for this port's bit	Port-based VLAN Table. The bits in this table are used to restrict the output ports to which this input port can send frames. The VLANTable bits are used if 802.1Q is disabled on this port or if it cannot determine membership. When used, these bits restrict where a port can send frames (unless a VLANTunnel frame is being received - Table 54). To send frames to Port 0, bit 0 of this register must be a one. To send frames to Port 1, bit 1 of this register must be a one, etc. After Reset, all ports are accessible since all the other port number bits are set to a one. This Port's bit will always be zero at reset. This prevents frames exiting the port on which they arrived. This port's bit can be set to a one, enabling frames to be returned on their ingress port; consequently care should be taken in writing code to control these bits. This register is reset to 0x3FE for Port 0 (SMI Device Address 0x10), and it resets to 0x3FD for Port 1 (Addr 0x11) and to 0x3B for Port 2 (Addr 0x12), etc.

Table 55: Default Port VLAN ID & Priority
Offset: 0x07 (Hex), or 7 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:14	DefPri	RWR	Default Priority. The bits of this register are used as the default Ingress priority to use when no other priority information is available (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 UseIP and UseTag in). The Ingress priority determines the Egress Queue the frame is switched into.
13	TagPriLSB	RWR	Tag Priority Least Significant Bit. This bit is used as the lower bit of the IEEE Tagged PRI added during egress to untagged frames that ingress this port. The upper two bits of the IEEE Tagged PRI added to untagged frames come from the Egress Queue into which the frame was switched.
12	Force Default-VID	RWR	Force to use Default VID. When this bit is set to a one all Ingress frames with IEEE802.ac Tags are treated as if the Tag's VID contains a value of 0x000 (i.e., it is a priority tagged frame). In this case, the frame's VID is ignored and the DefaultVID below is used and replaced into the frame. When this bit is cleared to a zero all IEEE 802.3ac Tagged frames with a non-zero VID use the frame's VID unmodified.
11:0	DefaultVID	RWS to 0x001	VLAN Identifier. The VID field is used as the IEEE Tagged VID added during egress to untagged frames that arrived at this port.



Note

Registers 0x08 and 0x09 (Hex), or 8 and 9 (Decimal) are reserved.

Table 56: Rate Control
Offset: 0x0A (Hex), or 10 (Decimal)

Bits	Field	Type	Description
15:14	Limit Mode	RWR	Ingress Limit Mode. These bits determine what kinds of frames are limited and counted against Ingress limiting as follows: 00 = Limit and count all frames 01 = Limit and count Broadcast, Multicast and flooded unicast frames 10 = Limit and count Broadcast and Multicast frames only 11 = Limit and count Broadcast frames only If a frame is not limited by the above setting its size is not counted against the limit for the other frames.
13	Pri3Rate	RWR	Ingress data rate limit for priority-3 frames. Priority 3 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri2Rate 1 = Use twice the rate as Pri2Rate (i.e., it can be up to 64 Mbps)
12	Pri2Rate	RWR	Ingress data rate limit for priority-2 frames. Priority 2 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri1Rate 1 = Use twice the rate as Pri1Rate (i.e., it can be up to 32 Mbps)
11	Pri1Rate	RWR	Ingress data rate limit for priority 1 frames. Priority 1 frames are discarded after the ingress rate selected below is reached or exceeded: 0 = Use the same rate as Pri0Rate 1 = Use twice the rate as Pri0Rate (i.e., it can be up to 16 Mbps)
10:8	Pri0Rate	RWR	Ingress data rate limit for priority 0 frames. Priority 0 frames are discarded after the ingress rate selected below is reached or exceeded: 000 = Not Limited = Default 001 = 128 Kbps 010 = 256 Kbps 011 = 512 Kbps 100 = 1 Mbps 101 = 2 Mbps 110 = 4 Mbps 111 = 8 Mbps
7	Reserved	RES	Reserved for future use
6	Limit MGMT	RWR	Limit and count MGMT frame bytes. When this bit is set to a one MGMT frames are included in ingress and egress rate limiting calculations and can be limited. When this bit is cleared to a zero, MGMT bytes are not counted and are not limited



Table 56: Rate Control (Continued)
Offset: 0x0A (Hex), or 10 (Decimal)

Bits	Field	Type	Description
5	Count IFG	RWS	Count IFG bytes. When this bit is set to a one each frame's minimum inter frame gap (IFG) bytes (12 per frame) are included in ingress and egress rate limiting calculations. When this bit is cleared to a zero, IFG bytes are not counted.
4	Count Pre	RWS	Count Preamble bytes. When this bit is set to a one each frame's preamble bytes (8 per frame) are included in ingress and egress rate limiting calculations. When this bit is cleared to a zero, preamble bytes are not counted.
3	Reserved	RES	Reserved for future use
2:0	Egress Rate	RWR	Egress data rate limit. The EgressRate bits modify this port's effective transmission rate as follows: 000 = Not Limited = Default 001 = 128 Kbps 010 = 256 Kbps 011 = 512 Kbps 100 = 1 Mbps 101 = 2 Mbps 110 = 4 Mbps 111 = 8 Mbps

Table 57: Port Association Vector
Offset: 0x0B (Hex), or 11 (Decimal)

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 upon which egress port monitoring is enabled, as defined by the PAV bits below. This is the recommended operation mode for port monitoring.
14:10	Reserved	RES	Reserved for future use.
9:0	PAV	RWS to all zeros except for this port's bit	<p>Port Association Vector for ATU learning. The value of these bits is used as the port's DPV on automatic ATU Learning or Entry_State refresh whenever these bits contain a non-zero value. When these bits are all zero automatic Learning and Entry_State refresh is disabled on this port.</p> <p>For normal switch operation this port's bit should be the only bit set in the vector. These bits must only be changed when frames are not entering the port (see PortState bits in Port Control – Table 53).</p> <p>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. Egress port monitoring is performed using the ATU's address database, therefore it will not take effect until source addresses are learned/updated on the port where the PAV are modified.</p> <p>The PAV bits can be used to set up an ingress port monitor along with an egress port monitor. To configure Port 0 to receive 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.</p> <p>The PAV bits can be used to set up port trunking (along with the VLANTable bits Table 54). For the two ports that form a trunk, set both of the port's bits in both port's PAV registers. Then use the VLANTable to isolate the two ports from each other and to steer the traffic from the other ports down the required trunk line of the pair.</p>



Note

Registers 0x0C - 0x0F (Hex), or 12 - 15 (Decimal) are reserved.

Table 58: Rx Counter
 Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:0	RxCtr	RO	<p>Received Counter. When CtrMode is cleared to a zero (Global Control – Table 64) this counter increments each time a good frame enters this port. It does not matter if the frame is switched or discarded by the switch.</p> <p>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.</p> <p>The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (See PortState field of Port Control Register - Table 53). This register can be cleared by changing the state of the CtrMode bit in the Switch Global Control register.</p>

Table 59: Tx Counter
 Offset: 0x11 (Hex), or 17 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:0	TxCtr	RO	<p>Transmit Counter. When CtrMode is cleared to a zero (see Global Control) this counter increments each time a frame successfully exits this port.</p> <p>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.</p> <p>The counter wraps back to zero. The only time this counter does not increment is when this port is Disabled (See PortState field of Port Control Register - Table 53). This register can be cleared by changing the state of the CtrMode bit in Global Control (Table 64).</p>



Note

Registers 0x12 - 0x1F (Hex), or 18 - 31 (Decimal) are reserved.

5.2.3 Switch Global Registers

Table 60: Switch Global Status Register
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Type	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 Queue Controller and the Statistics Controller have finished 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 (Table 85) transitions from a one to a zero. It is automatically cleared when read. This bit being high will cause the 88E6083 device INTn pin to go low if the STATDoneIntEn bit in Global Control (Table 64) is set to a one.
5	VTUProb	ROC	VLAN Table Problem/Violation Interrupt. This bit is set to a one of the VTU cannot load a new VLAN entry owing to all of the available locations' being already in use or if a VLAN Violation is detected. It is automatically cleared when read. This bit's being high causes the 88E6083 device INTn pin to go low if the VTUProbIntEn bit in Global Control (Table 64) is set to a one.
4	VTUDone	ROC	VTU Done Interrupt. This bit is set to a one whenever the VTUBusy bit (Table 65) transitions from a one to a zero. It is automatically cleared when read. This bit's being high causes the 88E6083 device INTn pin to go low if the VTUDoneIntEn bit in Global Control (Table 64) is set to a one.
3	ATUFull	ROC	ATU Full Interrupt. This bit is set to a one if the ATU cannot load or learn a new mapping owing to all the available locations for an address being locked. It is automatically cleared when read. This bit's being high causes the 88E6083 device's INTn pin to go low if the ATUFullIntEn bit in GlobalControl (Table 64) is set to a one.
2	ATUDone	ROC	ATU Done Interrupt. This bit is set to a one whenever the ATUBusy bit (Table 71) changes from a one to a zero. It is automatically cleared when read. This bit's being high causes the 88E6083 device INTn pin to go low if the ATUDoneIntEn bit in Global Control (Table 64) is set to a one.
1	PHYInt	RO	PHY Interrupt. This bit is set to a one when the PHYs interrupt logic has at least one active interrupt. This bit being high causes the 88E6083 device INTn pin to go low if the PHYIntEn bit in Global Control (Table 64) is set to a one.
0	EEInt	ROC	EEPROM Done Interrupt. This bit is set to a one after the EEPROM has finished loading registers, and it is automatically cleared when read. This bit's being high causes the 88E6083 device INTn pin to go low when the EEIntEn bit in Global Control (Table 64) is set to a one.



Table 61: Switch MAC Address Register Bytes 0 & 1
Offset: 0x01 (Hex), or 1 (Decimal)

Bits	Field	Type	Description
15:9	MACByte0	RWR	MAC Address Byte 0 (bits 47:41) is 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 1 st 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 47:3 of the MAC address are the same, but bits 3:0 are the port number (Port 0 = 0, Port 1 = 1, and so on.)
7:0	MACByte1	RWR	MAC Address Byte 1 (bits 39:32) is used as the switch's source address (SA) in transmitted full-duplex Pause frames.

Table 62: Switch MAC Address Register Bytes 2 & 3
Offset: 0x02 (Hex), or 2 (Decimal)

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

Table 63: Switch MAC Address Register Bytes 4 & 5
Offset: 0x03 (Hex), or 3 (Decimal)

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

Table 64: Switch Global Control Register
Offset: 0x04 (Hex), or 4 (Decimal)

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 single frame.
12	Reserved	RES	Reserved for future use.
11	Scheduling	RWR	Scheduling mode. This bit is used to select the Queue Controller's scheduling mode as follows: 0 = Use an 8, 4, 2, 1 weighted fair queuing scheme 1 = Use a strict priority scheme
10	MaxFrame Size	RWR	Maximum Frame Size allowed. The Ingress block discards all frames that are less than 64 bytes in size. It also discards all frames that are greater than a certain size as follows: 0 = Max size is 1522 if IEEE 802.3ac tagged, or 1518 if not tagged 1 = Max size is 1536
9	ReLoad	SC	Reload the registers using the EEPROM. When this bit is set to a one the contents of the external EEPROM are used to load the registers just as if a reset had occurred. When the reload operation is done, this bit is cleared to a zero automatically, and the EEInt interrupt is set.
8	CtrMode	RWR	Counter Mode. When this bit is set to a one, the Rx counters for all ports (Table 58) count Rx errors and the Tx counters for all ports (Table 59) count Tx collisions. When this bit is cleared to a zero, the Rx counters for all ports count Rx frames and the Tx counters for all ports count Tx frames. The Rx counters and the Tx 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 further use
6	StatsDoneIntEn	RWR	Statistics Operation Done Interrupt Enable. This bit must be set to a one to allow the Stat Done interrupt to drive the 88E6083 device INTn pin low.
5	VTUProblntEn	RWR	VLAN Problem/Violation Interrupt Enable. This bit must be set to a one to allow the VT Problem Interrupt to drive the 88E6083 device 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 88E6083 device INTn pin low.
3	ATUFull IntEn	RWR	ATU Full Interrupt Enable. This bit must be set to a one to allow the ATU Full interrupt to drive the 88E6083 device INTn pin low.



Table 64: Switch Global Control Register (Continued)
 Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Type	Description
2	ATUDone IntEn	RWR	ATU Done Interrupt Enable. This bit must be set to a one to allow the ATU Done interrupt to drive the 88E6083 device INTn pin low.
1	PHYIntEn	RWR	PHY Interrupt Enable. This bit must be set to a one to allow active interrupts enabled in PHY registers 0x12 to drive the 88E6083 device 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 88E6083 device INTn pin low.

Table 65: VTU Operation Register
 Offset: 0x05 (Hex), or 5 (Decimal) Register

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 is automatically cleared to a zero. The transition of this bit from a one to a zero can be used to generate an interrupt (Table 64).
14:12	VTUOp	RWR	VLAN Table Unit Table Opcode. The 88E6083 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
11:7	Reserved	RES	Reserved for future use.
6	Member Violation	RO	Source port violation. This bit is set to a one if any 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. This bit causes the VTU-ProbInt (in GlobalStatus - Table 60) to be set causing an interrupt if enabled. This bit is cleared after a Get/Clear Violation Data VTUOp that returns a SPID (bits (3:0 below) ≠ 0xF. Only the first Member Violation or Miss Violation (below) is saved until cleared.
5	Miss Violation	RO	VTU Miss violation. This bit is set to a one if any 802.1Q enabled ingress port accesses the VTU with a VID that is not contained in the VTU. This bit causes the VTUProbInt (in GlobalStatus - Table 60) to be set causing an interrupt if enabled. This bit will be cleared after a Get/Clear Violation Data VTUOp that returns a SPID (bits 3:0 below) ≠ 0xF. Only the first Miss violation or Member violation (above) is saved until cleared.

Table 65: VTU Operation Register (Continued)
Offset: 0x05 (Hex), or 5 (Decimal) Register

Bits	Field	Type	Description
4	VTUFull Violation	RO	VLAN Translation Table Full Violation. This bit is set to a one if the Load VTUOp cannot store the required entry. This bit causes the VTUProblnt (in GlobalStatus - Table 60) to be set causing an interrupt if enabled. This bit is cleared after a Get/Clear Violation Data VTUOp that returns an SPID (bits 3:0 below) of 0xF.
3:0	DBNum/ SPID	RWR	VTU MAC Address Database Number or Source Port ID. On all VTUOps except for Get Violation Data, this field is DBNum and it is used to separate MAC address databases by a frame's VID. If multiple address databases are not being used these bits must remain zero. If multiple address databases are being used these bits are used to set the required address database number that is associated with a VID value on Load operations (or used to read the currently assigned DBNum on Get Next operations). On the Get Violation Data VTUOp this field returns the Source Port ID of the port that caused the violation (an SPID of 0xF indicates a VTUOp violation).

1. An Entry is loaded if the Valid bit (Table 66) 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 required VID.
2. An Entry is Purged if it exists and the Valid bit (Table 66) is cleared to a zero.
3. A Get Next operation finds the next higher VID currently in the VTU's database. The VID value (Table 66) is used as the VID from which to start. To find the lowest VID set the VID field to ones. When the operation completes, 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 ones with a Valid bit cleared to zero, no higher VIDs were found. To Search for a particular VID, perform a Get Next operation using a VID field with a value one less than the object of the search.
4. The Get/Clear Violation VTUOp returns the source port of the violation in the SPID field of this register (bits 3:0) and it returns the VID of the violation in the VID field of the VTU VID register (Table 66).



Table 66: VTU VID Register
Offset: 0x06 (Hex), or 6 (Decimal)

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 that the VID value below is valid. If this bit is cleared to a zero and the VID is all ones it indicates that the end of the VID list was reached with no new valid entries found. On Load or Purge operations this bit indicates the required operation of a Load (when set to a one) or a Purge (when cleared to a zero).
11:0	VID	RWR	VLAN Identifier. This VID is used in all the VTUOp commands (except Get/ Clear Violation Data) and it is the VID that is associated with the VTU data below (Table 67) or the VID that caused the VTU Violation.

Table 67: VTU Data Register Ports 0 to 3
Offset: 0x07 (Hex) or 7 (Decimal)

Bits	Field	Type	Description
15:14	Reserved P3	RES	Reserved for future use.
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	Reserved P2	RES	Reserved for future use.
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	Reserved P1	RES	Reserved for future use.
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	Reserved P0	RES	Reserved for future use.
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 not a member of this VLAN. Any frames with this VID ¹ are discarded at Ingress and are not allowed to egress this port. 10 = Port is a member of this VLAN and frames are to egress Untagged. 11 = Port is a member of this VLAN and frames are to egress Tagged.

1. The VID used comes from the VID in Tagged frames or the default VID assigned to Untagged frames.

Table 68: VTU Data Register Port 4 to 7
Offset: 0x08 (Hex), or 8 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:14	Reserved P7	RES	Reserved for future use.
13:12	Member TagP7	RWR	Membership and Egress Tagging for Port 7. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP4 below.
11:10	Reserved P6	RES	Reserved for future use.
9:8	Member TagP6	RWR	Membership and Egress Tagging for Port 6. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP4 below.
7:6	Reserved P5	RES	Reserved for future use.
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	Reserved P4	RES	Reserved for future use.
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 not a member of this VLAN. Any frames with this VID ¹ are discarded at Ingress and are not allowed to egress this port. 10 = Port is a member of this VLAN and frames are to egress Untagged. 11 = Port is a member of this VLAN and frames are to egress Tagged.

1. The VID used comes from the VID in Tagged frames or the default VID assigned to Untagged frames.

Table 69: VTU Data Register Port 8 to 9
Offset: 0x09 (Hex), or 9 (Decimal)

Bits	Field	Type/ Reset Value	Description
15:8	Reserved	RES	Reserved for future use.
7:6	Reserved P9	RES	Reserved for future use.
5:4	Member TagP9	RWR	Membership and Egress Tagging for Port 9. These bits are used to support 802.1Q membership and Egress Tagging. See MemberTagP8 below.
3:2	Reserved P8	RES	Reserved for future use.
1:0	Member TagP8	RWR	Membership and Egress Tagging for Port 8. These bits are used to support 802.1Q membership and Egress Tagging as follows: 00 = Port is a member of this VLAN and frames are to Egress unmodified. 01 = Port is not a member of this VLAN. Any frames with this VID ¹ are discarded at Ingress and are not allowed to egress this port. 10 = Port is a member of this VLAN and frames are to egress Untagged. 11 = Port is a member of this VLAN and frames are to egress Tagged.

1. The VID used comes from the VID in Tagged frames or the default VID assigned to Untagged frames.

Table 70: ATU Control Register
Offset: 0x0A (Hex), or 10 (Decimal)

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. Register values are not modified, except for the PortState bits (Table 53) and the EEPROM is not re-read. The PHYs are not affected by this bit, but the ATU and the VTU are. When the reset operation is done, this bit is cleared to a zero automatically. The reset occurs immediately. To prevent transmission of CRC frames, use the PortState bits to set each port to the Disabled state and wait for 2 ms (i.e., the time for a maximum frame to be transmitted at 10 Mbps) before setting the SWReset bit to a one.
14	LearnDis	RWR	ATU Learn Disable. 0 = Normal operation, learning is determined by the PortState (Table 53) 1 = Automatic learning is disabled on all ports – CPU ATU loads still work
13:12	ATUSize	RWS to 0x1	Address Translation Unit Table Size. The initial size of the ATU database is 1024 entries. The size of the ATU database can be modified at any time, but an ATU reset and a Switch Software Reset (bit 15 above) occurs automatically if the new ATUSize is different from the old ATUSize. 00 = 512 Entry Address Database 01 = 1024 Entry Address Database - default 10 = 2048 Entry Address Database 11 = Reserved
11:4	AgeTime	RWS to 0x13	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 16 is the age time in seconds. For example: The default value of 0x13 is 19 decimal. 19 x 16 = 304 seconds or just over 5 minutes. The minimum age time is 0x1 or 16 seconds. The maximum age time is 0xFF or 4080 seconds or 68 minutes. When the AgeTime is set to 0x0 the Aging function is disabled, and all learned addresses will remain in the database forever.
3:0	Reserved	RES	Reserved for future use.



Caution

SWRESET and ATU size changes will cause the contents of the ATU & the VTU to be flushed.

Table 71: ATU Operation Register
Offset: 0x0B (Hex), or 11 (Decimal)

Bits	Field	Type/ Reset Value	Description
15	ATUBusy	SC	Address Translation Unit Busy. This bit must be set to a one to start an ATU operation (see ATUOp below). Since only one ATU operation can be executing at one time, this bit must be zero before setting it to a one. When the requested ATU operation completes, this bit is automatically cleared to a zero. The transition of this bit from a one to a zero can be used to generate an interrupt (Table 60).
14:12	ATUOp	RWR	Address Translation Unit Table Opcode. The 88E6083 device supports the following ATU operations. All of these operations can be executed while frames are transiting through the switch: 000 = No Operation 001 = Flush All Entries ¹ 010 = Flush all Unlocked Entries ² 011 = Load or Purge an Entry ³ in a particular DBNum Database 100 = Get Next from a particular DBNum Database ⁴ 101 = Flush All Entries in a particular DBNum Database 110 = Flush All Unlocked Entries in a particular DBNum Database 111 = Reserved
11:4	Reserved	RES	Reserved for future use.
3:0	DBNum	RWR	ATU MAC Address Database Number. If multiple address databases are not being used, these bits must remain zero. If multiple address databases are being used these bits are used to set the required address database number that is to be used on the Database supported commands (ATUOps 0x3, 0x4, 0x5, and 0x6 above).

1. All frames flood until new addresses are learned.
2. An Unlocked entry is any unicast address with an EntryState less than 0xF. All unicast frames will flood until new addresses are learned
3. An Entry is Loaded if the EntryState (Table 72) is non-zero. An Entry is Purged if it exists and if the EntryState is zero.
4. A Get Next operation finds the next higher MAC address currently in a particular ATU database (defined by the DBNum field). The ATUByte[5:0] values (Table 73) 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 found. To find the next address simply issue the Get Next opcode again. If ATUByte[5:0] is returned set to all ones, no higher MAC address was found (if EntryState = 0) or the Broadcast Address was found (if EntryState ≠ 0) In either case, 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.



Table 72: ATU Data Register
Offset: 0x0C (Hex), or 12 (Decimal)

Bits	Field	Type	Description
15:14	EntryPri	RWR	ATU Entry Priority Data. These bits are used as the input Entry Priority data for ATU load operation. It is the resulting Entry Priority from the ATU Get Next operation.
13:4	PortVec	RWR	Port Vector. These bits are used as the input Port Vector for ATU load operation. It is the resulting Port Vector from the ATU Get Next operation.
3:0	EntryState	RWR	ATU Entry State. These bits are used as the input Entry State for ATU Load or purge operations. It is the resulting Entry State from the ATU Get Next operation.

Table 73: ATU Switch MAC Address Register Bytes 0 & 1
Offset: 0x0D (Hex), or 13 (Decimal)

Bits	Field	Type	Description
15:8	ATUByte0	RWR	ATU MAC Address Byte 0 (bits 47:40) is used as the input MAC address for ATU load, purge, or Get Next operations. It is the resulting MAC address from ATU Get Next operation. Bit 0 of byte 0 (bit 40) is the multicast bit (it is the 1 st bit down the wire). Any MAC address with the multicast bit set to a one is considered locked by the ATU.
7:0	ATUByte1	RWR	ATU MAC Address Byte 1 (bits 39:32) is used as the input MAC address for ATU load, purge, or Get Next operations. It is the resulting MAC address from the ATU Get Next operation.

Table 74: ATU Switch MAC Address Register Bytes 2 & 3
Offset: 0x0E (Hex), or 14 (Decimal)

Bits	Field	Type	Description
15:8	ATUByte2	RWR	ATU MAC Address Byte 2 (bits 31:24) is used as the input MAC address for ATU load, purge or Get Next operations. It is the resulting MAC address from ATU Get Next operations.
7:0	ATUByte3	RWR	ATU MAC Address Byte 3 (bits 23:16) that are used as the input MAC address for ATU load, purge or Get Next operations. It is the resulting MAC address from ATU Get Next operations.

Table 75: ATU Switch MAC Address Register Bytes 4 & 5
Offset: 0x0F (Hex), or 15 (Decimal)

Bits	Field	Type	Description
15:8	ATUByte4	RWR	ATU MAC Address Byte 4 (bits 15:8) is used as the input MAC address for ATU load, purge or Get Next operations. It is the resulting MAC address from ATU Get Next operations.
7:0	ATUByte5	RWR	ATU MAC Address Byte 5 (bits 7:0) is used as the input MAC address for ATU load, purge or Get Next operations. It is the resulting MAC address from ATU Get Next operations.



Table 76: IP-PRI Mapping Register 0
 Offset: 0x10 (Hex), or 16 (Decimal)

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

Table 77: IP-PRI Mapping Register 1
 Offset: 0x11 (Hex), or 17 (Decimal)

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

Table 78: IP-PRI Mapping Register 2
Offset: 0x12 (Hex), or 18 (Decimal)

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

Table 79: IP-PRI Mapping Register 3
Offset: 0x13 (Hex), or 19 (Decimal)

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



Table 80: IP-PRI Mapping Register 4
Offset: 0x14 (Hex), or 20 (Decimal)

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

Table 81: IP-PRI Mapping Register 5
Offset: 0x15 (Hex), or 21 (Decimal)

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

Table 82: IP-PRI Mapping Register 6
Offset: 0x16 (Hex), or 22 (Decimal)

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

Table 83: IP-PRI Mapping Register 7
Offset: 0x17 (Hex), or 23 (Decimal)

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



Note

Registers 0x19 - 0x1C (Hex), or 25 - 28 (Decimal) are reserved.

Table 84: IEEE-PRI Register
Offset: 0x18 (Hex), or 24 (Decimal)

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 when 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 when 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 when 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 when 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 when 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 when 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 when 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 when its IEEE Tag has a value of 0.

Table 85: Statistics Operation Register
Offset: 0x1D (Hex) or 29 (Decimal)

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 (Table 64).

Table 85: Statistics Operation Register (Continued)
Offset: 0x1D (Hex) or 29 (Decimal)

Bits	Field	Type	Description	
14:12	StatsOp	RWR	<p>Statistics Unit Opcode. The 88E6083 supports the following Stats operations (all of these operations can be executed while frames are in transit through the switch):</p> <p>000 = No Operation 001 = Flush (clear) All Counters for all Ports 010 = Flush (clear) All Counters for a Port 011 = Reserved 100 = Read a Captured Counter 101 = Capture All Counters for a Port 11x = Reserved</p>	
11:6	Reserved	RES	Reserved for future use.	
5:0	StatsPr	RWR	<p>Statistics Pointer. This field is used as a parameter for the above StatsOp commands. It must be set to the required Port number for the Capture All Counters for a Port (0x5) StatsOp and Flush All Counters for a Port (0x2) StatsOp. Use 0x00 for Port 0, 0x01 for Port 1, etc.</p> <p>StatsPr must be set to the required counter to read for the Read a Captured Counter (0x4) StatsOp (valid range is 0x00 to 0x27). A Capture All Counters for a Port StatsOp must be done prior to using the Read A Captured Counter StatsOp. The counter that is read is defined as follows:</p>	
			Ingress Counters	Egress Counters
			0x00 - InUnicasts 0x01 - InBroadcasts 0x02 - InPause 0x03 - In Multicasts 0x04 - InFCSErr 0x05 - AlignErr 0x06 - InGoodOctets 0x07 - InBadOctets 0x08 - Undersize 0x09 - Fragments 0x0A - In64Octets 0x0B - In127Octets 0x0C - In255Octets 0x0D - In511Octets 0x0E - In1023Octets 0x0F - InMaxOctets 0x10 - Jabber 0x11 - Oversize 0x13 - InFiltered 0x12 - InDiscards	0x14 - OutUnicasts 0x15 - OutBroadcasts 0x16 - OutPause 0x17 - OutMulticasts 0x18 - OutFCSErr 0x19 - OutOctets 0x1A - Out64Octets 0x1B - Out127Octets 0x1C - Out255Octets 0x1D - Out511Octets 0x1E - Out1023Octets 0x1F - OutMaxOctets 0x20 - Collisions 0x21 - Late 0x22 - Excessive 0x23 - Multiple 0x24 - Single 0x25 - Deferred 0x26 - OutFiltered 0x27 - OutDiscards



Table 86: Stats Counter Register Bytes, 3&2
Offset: 0x1E (Hex), or 30 (Decimal)

Bits	Field	Type	Description
15:8	StatsByte3	RO	Statistics Counter Byte 3. These bits contain bits 31:24 of the last statistics counter requested to be read by the CPU (by using the Read a Counter-StatsOp - Table 85)
7:0	StatusByte2	RO	Statistics Counter Byte 2. These bits contain bits 23:16 of the last stat counter requested to be read by the CPU (by using the Read a Counter-StatsOp - Table 85)

Table 87: Stats Counter Register Bytes, 1&0
Offset: 0x1F (Hex), or 31 (Decimal)

Bits	Field	Type	Description
15:8	StatsByte1	RO	Statistics Counter Byte 1. These bits contain bits 31:24 of the last statistics counter requested to be read by the CPU (by using the Read a Counter-StatsOp - Table 85)
7:0	StatusByte0	RO	Statistics Counter Byte 0. These bits contain bits 23:16 of the last statistics counter requested to be read by the CPU (by using the Read a Counter-StatsOp - Table 85)

Section 6. Serial Management Interface (SMI)

The 88E6083 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.

6.1 MDC/MDIO Read and Write Operations

All of the PHY management registers, as well as the switch core registers, are accessed through this interface. A description of these registers can be found in [Section 5, Register Description](#) and [Section 7, PHY Registers](#).



Note

Access to the 88E6083 device's Switch and PHY registers is not possible when the Serial EEPROM is still loading the registers. A CPU can monitor the 88E6083 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 6](#).)

Figure 34: Typical MDC/MDIO Read Operation

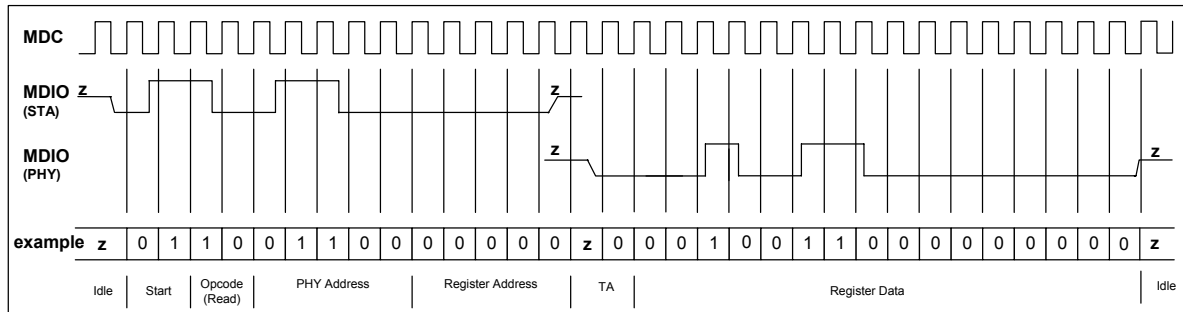


Figure 35: Typical MDC/MDIO Write Operation

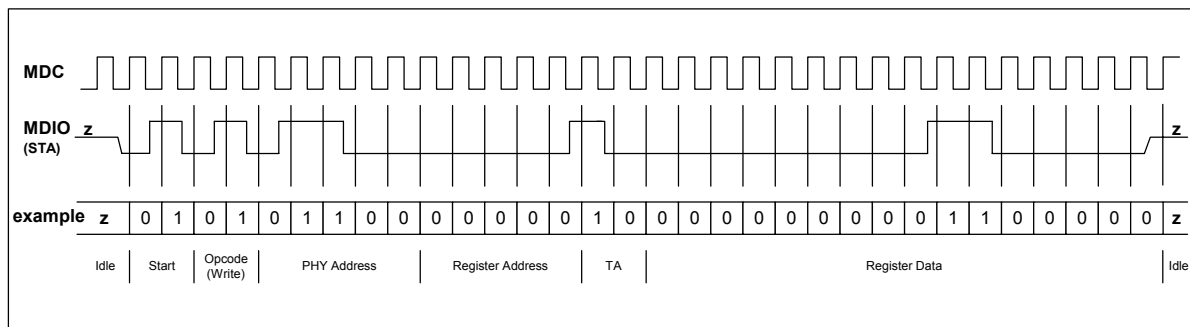


Table 88 shows an example of a read operation.

Table 88: 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	01100	00000	z0	00010011000000	11111111

Section 7. PHY Registers

The 88E6083 device contains eight physical layer devices (PHYs). These devices are accessible using SMI device addresses 0x00 to 0x07. The PHYs are fully IEEE 802.3 compliant including their register interface.

The PHYs in the 88E6083 device are identical to those in the Marvell® 88E3082 Octal Transceiver.

Table 89: PHY Register Map

Description	Offset Hex	Offset Decimal	Page Number
PHY Control Register	0x00	0	page 159
PHY Status Register	0x01	1	page 162
PHY Identifier	0x02	2	page 163
PHY Identifier	0x03	3	page 163
Auto-Negotiation Advertisement Register	0x04	4	page 164
Link Partner Ability Register (Base Page)	0x05	5	page 166
Link Partner Ability Register (Next Page)	0x05	5	page 167
Auto-Negotiation Expansion Register	0x06	6	page 168
Next Page Transmit Register	0x07	7	page 169
Link Partner Next Page Register	0x08	8	page 169
Reserved Registers	0x09-0x0F	9 - 15	page 170
PHY Specific Control Register I	0x10	16	page 170
PHY Specific Status Register	0x11	17	page 172
PHY Interrupt Enable	0x12	18	page 174
PHY Interrupt Status	0x13	19	page 175
PHY Interrupt Port Summary (Common register to all ports)	0x14	20	page 177
PHY Receive Error Counter	0x15	21	page 178
LED Parallel Select Register (Common register to all ports)	0x16	22	page 179
LED Stream Select Register (Common register to all ports)	0x17	23	page 180
PHY LED Control Register (Common register to all ports)	0x18	24	page 182
PHY Manual LED Override Register	0x19	25	page 184
VCT™ Control Register	0x1A	26	page 185
VCT™ Status Register	0x1B	27	page 186
PHY Specific Control Register II	0x1C	28	page 187
Reserved Registers	0x1D to 0x1F	29 - 31	page 187



7.1 Register Types

The registers in the 88E6083 device are made up of one or more fields. The way in which each of these fields operate is defined by the field's Type. The function of each Type is described in [Table 90](#).

Table 90: Register Types

Type	Description
LH	Register field with latching high function. If status is high, then the register is set to a 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 a zero and remains cleared until a read operation is performed through the management interface or until a reset occurs.
RES	Reserved for future use. All reserved bits are read as zero unless otherwise noted.
RO	Read only.
ROC	Read only clear. After read, register field is cleared to zero.
R/W	Read and write with no definable initial value.
RWR	Read/Write reset. All bits are readable and writable. After reset the register field is cleared to zero.
RWS	Read/Write set. All bits are readable and writable. After reset the register field is set to a non-zero value specified in the text.
SC	Self-Clear. Writing a one to this register causes the required function to be immediately executed, then the register field is cleared to zero when the function is complete.
Retain	The register value is retained after software reset is executed.
Update	Value written to the register field does not take effect until a soft reset is executed.
WO	Write only. Reads to this type of register field return undefined data.
SW Rst	Descriptions listed under the SW Reset column specify the condition of the bit(s) after a soft reset is asserted.
HW Rst	Descriptions listed under the HW Reset column specify the condition of the bit(s) after a hard reset is asserted.

Table 91: PHY Control Register
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	SWReset	SC	0x0	Self Clear	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. 1 = PHY reset 0 = Normal operation
14	Loopback	R/W	Retain	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. Enable or Disable of loopback bit must be followed by PHY software Reset (bit 15 above). 1 = Enable loopback 0 = Disable loopback NOTE: See Customer Information Section in the Release Notes for enable/disable of loopback
13	SpeedLSB	R/W	ANEG [2:0]	Update	Speed Selection (LSB) When a speed change occurs, the PHY drops link and tries to determine speed when Auto-Negotiation is on. Speed, Auto-Negotiation enable, and duplex enable take on the values set by ANEG[2:0] on hardware reset. A write to these registers 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. 1 = 100 Mbps 0 = 10 Mbps



Table 91: PHY Control Register (Continued)
 Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
12	AnegEn	R/W	ANEG [2:0]	Update	<p>Auto-Negotiation Enable Speed, Auto-Negotiation enable, and duplex enable take on the values set by ANEG[2:0] on hardware reset. A write to these registers has no effect unless any one of the following also occurs: Software reset is asserted (bit 15, above), Power down (bit 11, below), or transitions from power down to normal operation. If the AnegEn bit is set to 0, the speed and duplex bits of the PHY Control Register (Offset 0x00) take effect. If the AnegEn bit is set to 1, speed and duplex advertisement is found in the Auto-Negotiation Advertisement Register (Offset 0x04).</p> <p>1 = Enable Auto-Negotiation Process 0 = Disable Auto-Negotiation Process</p>
11	PwrDwn	R/W	0x0	0x0	<p>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.</p> <p>1 = Power down 0 = Normal operation</p>
10	Isolate	RO	Always 0	Always 0	<p>Isolate Mode Will always be 0. The Isolate function is not available, since full MII is not implemented.</p> <p>0 = Normal operation</p>
9	RestartAneg	SC	0x0	Self Clear	<p>Restart Auto-Negotiation Auto-Negotiation automatically restarts after hardware or software reset regardless of whether or not the restart bit is set.</p> <p>1 = Restart Auto-Negotiation Process 0 = Normal operation</p>

Table 91: PHY Control Register (Continued)
Offset: 0x00 (Hex), or 0 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
8	Duplex	R/W	ANEG [2:0]	Update	<p>Duplex Mode Selection Speed, Auto-Negotiation enable, and duplex enable take on the values set by ANEG[2:0] on hardware reset. 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.</p> <p>1 = Full-duplex 0 = Half-duplex</p>
7	ColTest	RO	Always 0	Always 0	<p>Collision Test Mode Will always be 0. The Collision test is not available, since full MII is not implemented.</p> <p>0 = Disable COL signal test</p>
6	SpeedMSB	RO	Always 0	Always 0	<p>Speed Selection Mode (MSB) Will always be 0.</p> <p>0 = 100 Mbps or 10 Mbps</p>
5:0	Reserved	RO	Always 0	Always 0	Will always be 0.



Table 92: PHY Status Register
Offset: 0x01 (Hex), or 1 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	100T4	RO	Always 0	Always 0	100BASE-T4 This protocol is not available. 0 = PHY not able to perform 100BASE-T4
14	100FDX	RO	Always 1	Always 1	100BASE-T and 100BASE-X full-duplex 1 = PHY able to perform full-duplex
13	100HDX	RO	Always 1	Always 1	100BASE-T and 100BASE-X half-duplex 1 = PHY able to perform half-duplex
12	10FDX	RO	Always 1	Always 1	10BASE-T full-duplex 1 = PHY able to perform full-duplex
11	10HPX	RO	Always 1	Always 1	10BASE-T half-duplex 1 = PHY able to perform half-duplex
10	100T2FDX	RO	Always 0	Always 0	100BASE-T2 full-duplex. This protocol is not available. 0 = PHY not able to perform full-duplex
9	100T2HDX	RO	Always 0	Always 0	100BASE-T2 half-duplex This protocol is not available. 0 = PHY not able to perform half-duplex
8	ExtdStatus	RO	Always 0	Always 0	Extended Status 0 = No extended status information in Register 15
7	Reserved	RO	Always 0	Always 0	Must always be 0.
6	MFPPreSup	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 1 = Auto-Negotiation process completed 0 = Auto-Negotiation process not completed

Table 92: PHY Status Register (Continued)
Offset: 0x01 (Hex), or 1 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
4	RemoteFault	RO, LH	0x0	0	Remote Fault Mode 1 = Remote fault condition detected 0 = Remote fault condition not detected
3	AnegAble	RO	Always 1	Always 1	Auto-Negotiation Ability Mode 1 = PHY able to perform Auto-Negotiation
2	Link	RO, LL	0x0	0	Link Status Mode This register indicates when link was lost since the last read. For the current link status, either read this register back-to-back or read RTLink in Table 102 . 1 = Link is up 0 = Link is down
1	JabberDet	RO, LH	0x0	0	Jabber Detect 1 = Jabber condition detected 0 = Jabber condition not detected
0	ExtdReg	RO	Always 1	Always 1	Extended capability mode. 1 = Extended register capabilities

Table 93: PHY Identifier
Offset: 0x02 (Hex), or 2 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	OUI MSb	RO	0x0141	0x0141	Organizationally Unique Identifier bits 3:18 0000000101000001 Marvell® OUI is 0x005043

Table 94: PHY Identifier
Offset: 0x03 (Hex), or 3 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:10	OUI LSb	RO	Always 000011	Always 000011	Organizationally Unique Identifier bits 19:24
9:4	ModelNum	RO	Varies	Varies	Model Number = 001000



Table 94: PHY Identifier
 Offset: 0x03 (Hex), or 3 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
3:0	RevNum	RO	Varies	Varies	Revision Number Contact Marvell® FAEs for information on the device revision number.

Table 95: Auto-Negotiation Advertisement Register
 Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	AnegAd NxtPage	R/W	0x0	Retain	Next Page 1 = Advertise 0 = Not advertised Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
14	Ack	RO	Always 0	Always 0	Must be 0.
13	AnegAd ReFault	R/W	0x0	Retain	Remote Fault Mode 1 = Set Remote Fault bit 0 = Do not set Remote Fault bit Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
12:11	Reserved	R/W	0x0	Retain	Must be 00. Reserved bits are R/W to allow for forward compatibility with future IEEE standards. Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
10	AnegAd Pause	R/W	0x0	Retain	Pause Mode 1 = MAC Pause implemented 0 = MAC Pause not implemented Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.

Table 95: Auto-Negotiation Advertisement Register (Continued)
Offset: 0x04 (Hex), or 4 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
9	AnegAd 100T4	RO	Always 0	Always 0	100BASE-T4 mode 0 = Not capable of 100BASE-T4
8	AnegAd 100FDX	R/W	0x1	Retain	100BASE-TX full-duplex Mode 1 = Advertise 0 = Not advertised Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
7	AnegAd 100HDX	R/W	0x1	Retain	100BASE-TX half-duplex Mode 1 = Advertise 0 = Not advertised Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
6	AnegAd 10FDX	R/W	0x1	Retain	10BASE-TX full-duplex Mode 1 = Advertise 0 = Not advertised Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
5	AnegAd 10HDX	R/W	0x1	Retain	10BASE-TX half-duplex Mode 1 = Advertise 0 = Not advertised Values programmed into the Auto-Negotiation Advertisement Register have no effect unless Auto-Negotiation is restarted (RestartAneg - Table 91) or link goes down.
4:0	AnegAd Selector	RO	Always 0x01	Always 0x01	Selector Field Mode 00001 = 802.3



Table 96: Link Partner Ability Register (Base Page)
Offset: 0x05 (Hex), or 5 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	LPNxt Page	RO	0x0	0	Next Page Mode Base page will be overwritten if next page is received and if Reg8NxtPg (Table 101) is disabled. When Reg8NxtPg (Table 101) is enabled, then next page is stored in the Link Partner Next Page register (Table 100), and the Link Partner Ability Register (Table 96) holds the base page. Received Code Word Bit 15
14	LPAck	RO	0x0	0	Acknowledge Received Code Word Bit 14
13	LPRemote Fault	RO	0x0	0	Remote Fault Received Code Word Bit 13
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 97: Link Partner Ability Register (Next Page)
Offset: 0x05 (Hex), or 5 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	LPNxtPage	RO	--	--	Next Page Mode Base page will be overwritten if next page is received and if Reg8NxtPg (Table 101) is disabled. When Reg8NxtPg (Table 101) is enabled, then next page is stored in the Link Partner Next Page register (Table 100), and Link Partner Ability Register (Table 96) 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 98: Auto-Negotiation Expansion Register
 Offset: 0x06 (Hex), or 6 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:5	Reserved	RO	Always 0x000	Always 0x000	Reserved. Must be 00000000000. The Auto-Negotiation Expansion Register is not valid until the AnegDone (Table 92) indicates completed.
4	ParFaultDet	LH/ ROC	0x0	0x0	Parallel Detection Level 1 = A fault has been detected via the Parallel Detection function 0 = A fault has not been detected via the Parallel Detection function
3	LPNxtPg Able	RO	0x0	0x0	Link Partner Next Page Able 1 = Link Partner is Next Page able 0 = Link Partner is not Next Page able
2	LocalNxtPg Able	RO	Always 0x1	Always 0x1	Local Next Page Able This bit is equivalent to AnegAble (Table 92). 1 = Local Device is Next Page able
1	RxNewPage	RO/LH	0x0	0x0	Page Received 1 = A New Page has been received 0 = A New Page has not been received
0	LPAnegAble	RO	0x0	0x0	Link Partner Auto-Negotiation Able 1 = Link Partner is Auto-Negotiation able 0 = Link Partner is not Auto-Negotiation able

Table 99: Next Page Transmit Register
Offset: 0x07 (Hex), or 7 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	TxNxtPage	R/W	0x0	0x0	Next Page 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 100: Link Partner Next Page Register
Offset: 0x08 (Hex), or 8 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	RxNxtPage	RO	0x0	0x0	Next Page If Reg8NxtPg (Table 101) is enabled, then next page is stored in the Link Partner Next Page register (Table 100); otherwise, the Link Partner Next Page register (Table 100) is cleared to all 0's. Received Code Word Bit 15
14	RxAck	RO	0x0	0x0	Acknowledge Received Code Word Bit 14
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	0x001	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 101: PHY Specific Control Register
Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RES	--	--	Reserved.
14	EDet	R/W	ENA_ EDET	Retain	Energy Detect 1 = Enable with sense and pulse 0 = Disable Enable with sense only is not supported Enable Energy Detect takes on the appropriate value defined by the CONFIG9 pin at hardware reset.
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. 1 = Disable linkpulse check 0 = Enable linkpulse check
12	Reg8NxtPg	R/W	0x0	0x0	Enable the Link Partner Next Page register (Table 97) to store Next Page. If set to store next page in the Link Partner Next Page register (Table 97), then 802.3u is violated to emulate 802.3ab. 1 = Store next page in the Link Partner Next Page register (Table 97) 0 = Store next page in the Link Partner Ability Register (Base Page) register (Table 96).
11	DisNLPGen	R/W	0x0	0x0	Disable Linkpulse Generation. Linkpulse check and generation disable have no effect, when Auto-Negotiation is enabled locally. 1 = Disable linkpulse generation 0 = Enable 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. 1 = Force link good 0 = Normal operation

Table 101: PHY Specific Control Register (Continued)
Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
9	DisScrambler	R/W	ANEG [2:0]	Retain	Disable Scrambler If fiber mode is selected then the scrambler is disabled at hardware reset. 1 = Disable scrambler 0 = Enable scrambler
8	DisFEFI	R/W	DIS_ FEFI ANEG [2:0]	Retain	Disable FEFI In 100BASE-FX mode, Disable FEFI takes on the appropriate value defined by the CONFIG8 pin at hardware reset. FEFI is automatically disabled regardless of the state of this bit if copper mode is selected. 1 = Disable FEFI 0 = Enable FEFI For the 88E3083 device, this bit applies to Port 7 only.
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. 1 = Lower 10BASE-T receive threshold 0 = Normal 10BASE-T receive threshold
6	TPSelect	R/W	ANEG [2:0]	Update	(Un)Shielded Twisted Pair This setting can be changed by writing to this bit followed by software reset. 1 = Shielded Twisted Pair 0 = Unshielded Twisted Pair - default
5:4	MDI/MDIX Crossover	R/W	ENA_XC, 1	Update	MDI/MDIX Crossover This setting can be changed by writing to these bits followed by software reset. If ENA_XC is 1 at hardware reset then bits 5:4 = 11 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
3:2	RxFIFO Depth	R/W	0x0	0x0	Receive FIFO Depth 00 = 4 Bytes 01 = 6 Bytes 10 = 8 Bytes 11 = Reserved



Table 101: PHY Specific Control Register (Continued)
 Offset: 0x10 (Hex), or 16 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
1	AutoPol	R/W	0x0	00	Polarity Reversal If polarity is disabled, then the polarity is forced to be normal in 10BASE-T mode. Polarity reversal has no effect in 100BASE-TX mode. 1 = Disable automatic polarity reversal 0 = Enable automatic polarity reversal
0	DisJabber	R/W	0x0	00	Disable Jabber Jabber has no effect in full-duplex or in 100BASE-X mode. 1 = Disable jabber function 0 = Enable jabber function

Table 102: PHY Specific Status Register
 Offset: 0x11 (Hex), or 17 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RES	--	--	Reserved
14	ResSpeed	RO	ANEG [2:0]	Retain	Resolved Speed Speed and duplex takes on the values set by ANEG[2:0] pins on hardware reset only. The values are updated after the completion of Auto-Negotiation. The registers retain their values during software reset. 1 = 100 Mbps 0 = 10 Mbps
13	ResDuplex	RO	ANEG [2:0]	Retain	Resolved Duplex Mode Speed and duplex takes on the values set by ANEG[2:0] pins on hardware reset only. The values are updated after the completion of Auto-Negotiation. The registers retain their values during software reset. This bit is valid only after the resolved bit 11 is set. 1 = Full-duplex 0 = Half-duplex
12	RcvPage	RO, LH	0x0	0x0	Page Receive Mode 1 = Page received 0 = Page not received

Table 102: PHY Specific Status Register (Continued)
Offset: 0x11 (Hex), or 17 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
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 when Auto-Negotiation is completed or Auto-Negotiation is disabled. 1 = Resolved 0 = Not resolved
10	RTLink	RO	0x0	0x0	Link (real time) 1 = Link up 0 = Link down
9:7	Reserved	RES	0x0	0x0	Must be 000.
6	MDI/MDIX Crossover Status	RO	0x1	0x1	MDI/MDIX Crossover Status 1 = Transmit on pins TXP/TXN, Receive on pins RXP/RXN 0 = Transmit on pins RXP/RXN, Receive on pins TXP/TXN
5	Reserved	RES	0x4	0x4	Must be 0.
4	Sleep	RO	0	0x0	Energy Detect Status 1 = Chip is in sleep mode (No wire activity) 0 = Chip is not in sleep mode (Active)
3:2	Reserved	RES	0x0	0x0	Must be 00.
1	RTPolarity	RO	0x0	0x0	Polarity (real time) 1 = Reversed 0 = Normal
0	RTJabber	RO	Retain	0x0	Jabber (real time) 1 = Jabber 0 = No Jabber



Table 103: PHY Interrupt Enable
 Offset: 0x12 (Hex), or 18 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RES	--	--	Reserved
14	SpeedIntEn	R/W	0x0	Retain	Speed Changed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
13	DuplexIntEn	R/W	0x0	Retain	Duplex Changed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
12	RcvPageIntEn	R/W	0x0	Retain	Page Received Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
11	AnegDone IntEn	R/W	0x0	Retain	Auto-Negotiation Completed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
10	LinkIntEn	R/W	0x0	Retain	Link Status Changed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
9	SymErrIntEn	R/W	0x0	Retain	Symbol Error Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
8	FlsCrsIntEn	R/W	0x0	Retain	False Carrier Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
7	FIFOErrInt	R/W	0x0	Retain	FIFO Over/Underflow Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
6	MDI[x]IntEn	R/W	0x0	0x0	MDI/MDIX Crossover Changed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
5	Reserved	RES	0x0	Retain	Must be 0.

Table 103: PHY Interrupt Enable (Continued)
Offset: 0x12 (Hex), or 18 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
4	EDetIntEn	R/W	0x0	Retain	Energy Detect Interrupt Enable 1 = Enable 0 = Disable
3:2	Reserved	RES	0x0	Retain	Must be 00.
1	PolarityIntEn	R/W	0x0	Retain	Polarity Changed Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable
0	JabberIntEn	R/W	0x0	Retain	Jabber Interrupt Enable 1 = Interrupt enable 0 = Interrupt disable

Table 104: PHY Interrupt Status
Offset: 0x13 (Hex), or 19 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RES	--	--	Reserved
14	SpeedInt	RO, LH	0x0	0x0	Speed Changed 1 = Speed changed 0 = Speed not changed
13	DuplexInt	RO, LH	0x0	0x0	Duplex Changed 1 = Duplex changed 0 = Duplex not changed
12	RxPageInt	RO, LH	0x0	0x0	1 = Page received 0 = Page not received
11	AnegDoneInt	RO, LH	0x0	0x0	Auto-Negotiation Completed 1 = Auto-Negotiation completed 0 = Auto-Negotiation not completed
10	LinkInt	RO, LH	0x0	0x0	Link Status Changed 1 = Link status changed 0 = Link status not changed



Table 104: PHY Interrupt Status (Continued)
 Offset: 0x13 (Hex), or 19 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
9	SymErrInt	RO, LH	0x0	0x0	Symbol Error 1 = Symbol error 0 = No symbol error
8	FisCrslnt	RO, LH	0x0	0x0	False Carrier 1 = False carrier 0 = No false carrier
7	FIFOErrInt	RO, LH	0x0	0x0	FIFO Over /Underflow Error 1 = Over/underflow error 0 = No over/underflow error
6	MDIMDIXInt	RO, LH	0x0	0x0	MDI/MDIX Crossover Changed 1 = MDI/MDIX crossover changed 0 = MDI/MDIX crossover not changed
5	Reserved	RO	Always 0	Always 0	Must be 0
4	EDetChg	RO, LH	0x0	0x0	Energy Detect Changed 1 = Changed 0 = No Change
3:2	Reserved	RO	0x0	0x0	Must be 00
1	PolarityInt	RO	0x0	0x0	Polarity Changed 1 = Polarity changed 0 = Polarity not changed
0	JabberInt	RO, LH	0x0	0x0	Jabber Mode 1 = Jabber 0 = No Jabber

Table 105: PHY Interrupt Port Summary (Global¹)
Offset: 0x14 (Hex), or 20 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:8	Reserved	RO	0x0	0x0	Must be 00000000.
7	Port7Int Active	RO	0x0	0x0	Port 7 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 1 = Port has active interrupt 0 = Port does not have active interrupt
6	Port6Int Active	RO	0x0	0x0	Port 6 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 1 = Port has active interrupt 0 = Port does not have active interrupt
5	Port5Int Active	RO	0x0	0x0	Port 5 Interrupt Active Bit is set high, if any enabled interrupt is active for the port. Bit is cleared only when all bits in register 19 are cleared. 1 = Port has active interrupt 0 = Port does not have 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. 1 = Port has active interrupt 0 = Port does not have 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. 1 = Port has active interrupt 0 = Port does not have active interrupt



Table 105: PHY Interrupt Port Summary (Global¹) (Continued)
Offset: 0x14 (Hex), or 20 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
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. 1 = Port has active interrupt 0 = Port does not have 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. 1 = Port has active interrupt 0 = Port does not have 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. 1 = Port has active interrupt 0 = Port does not have active interrupt

1. A Global register is accessible for writing from any switch port.

Table 106: Receive Error Counter
Offset: 0x15 (Hex), or 21 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:0	RxErrCnt	RO	0x0000	0x0000	Receive Error Count This register counts receive errors on the media interface. When the maximum receive error count reaches 0xFFFF, the counter will not roll over.

Table 107: LED Parallel Select Register (Global^{1,2})
Offset: 0x16 (Hex), or 22 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:12	Reserved	RES	0000	Retain	Must be 0000
11:8	LED2	R/W	LED[1:0] 00 = LINK 01 = LINK 10 = LINK/RX 11 = LINK ACT	Retain	LED2 Control 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 to 1111 = Force to 1
7:4	LED1	R/W	LED[1:0] 00 = RX 01 = ACT 10 = TX 11 = DUPLEX /COLX	Retain	LED1 Control 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 to 1111 = Force to 1



Table 107: LED Parallel Select Register (Global^{1,2}) (Continued)
 Offset: 0x16 (Hex), or 22 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
3:0	LED0	R/W	LED[1:0] 00 = TX 01 = SPEED 10 = SPEED 11= SPEED	Retain	LED0 Control 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 to 1111 = Force to 1

1. See [Table 43 on page 114](#) and [Table 44 on page 115](#) for parallel LED display modes that can be selected by hardware pin configuration at reset.
2. This global register is accessible for writing or reading from any PHY port (i.e., only one physical register exists for all the PHY Ports).

Table 108: LED Stream Select for Serial LEDs (Global¹)
 Offset: 0x17 (Hex), or 23 (Decimal)

Bits	Function	Mode	HW Rst	SW Rst	Description
15:14	LEDLnkActy	R/W	LED[1:0]	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]	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

Table 108: LED Stream Select for Serial LEDs (Global¹) (Continued)
Offset: 0x17 (Hex), or 23 (Decimal)

Bits	Function	Mode	HW Rst	SW Rst	Description
11:10	LEDActy	R/W	LED[1:0]	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]	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]	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]	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
3:2	LEDSpd	R/W	LED[1:0]	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



Table 108: LED Stream Select for Serial LEDs (Global¹) (Continued)
 Offset: 0x17 (Hex), or 23 (Decimal)

Bits	Function	Mode	HW Rst	SW Rst	Description
1:0	LEDDx/ COLX	R/W	LED[1:0]	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

1. This global register is accessible for writing or reading from any PHY port (i.e., only one physical register exists for all the PHY Ports).

Table 109: PHY LED Control Register (Global¹)
 Offset: 0x18 (Hex), 0r 24 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RO	Always 0	Always 0	Must be 0.
14:12	PulseStretch	R/W	0x4	Retain	Pulse stretch duration. 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

Table 109: PHY LED Control Register (Global¹)
Offset: 0X18 (Hex), 0r 24 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
8:6	SrStrUpdate	R/W	0x2	Retain	Serial Stream Update 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. This global register is accessible for writing or reading from any PHY port (i.e., only one physical register exists for all the PHY Ports).



Table 110: PHY Manual LED Override
Offset: 0x19 (Hex), or 25 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:6	Reserved	R/W	0x0	Retain	0000000000
5:4	LED2	R/W	0x0	Retain	00 = Normal 01 = Blink 10 = LED Off 11 = LED On
3:2	LED1	R/W	0x0	Retain	00 = Normal 01 = Blink 10 = LED Off 11 = LED On
1:0	LED0	R/W	0x0	Retain	00 = Normal 01 = Blink 10 = LED Off 11 = LED On

Table 111: VCT™ Register for TXP/N Pins
Offset: 0x1A¹ (Hex), or 26 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	EnVCT	R/W, SC	0x0	0x0	<p>Enable VCT The 88E6083 device must be in forced 100 Mbps mode before enabling this bit.</p> <p>1 = Run VCT™ 0 = VCT completed</p> <p>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 in Section 4.5.</p>
14:13	VCTTst	RO	0x0	Retain	<p>VCT Test Status These VCT test status bits are valid after completion of VCT.</p> <p>11 = Test fail 00 = valid test, normal cable (no short or open in cable) 10 = valid test, open in cable (Impedance > 333 ohm) 01 = valid test, short in cable (Impedance < 33 ohm)</p>
12:8	AmpRfln	RO	0x0	Retain	<p>Amplitude of Reflection The amplitude of reflection is stored in these register bits. These amplitude bits range from 0x07 to 0x1F.</p> <p>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.</p>
7:0	DistRfln	RO	0x0	Retain	<p>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). See Figure 36. 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.</p>

1. The results stored in this register apply to the TX pin pair.

Table 112: VCT™ Register for RXP/N pins
 Offset: 0x1B¹ (Hex), or 27 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15	Reserved	RO	Always 0	Always 0	Reserved
14:13	VCTTst	RO	0	Retain	VCT Test Status The VCT test status bits are valid after completion of VCT. 11 = Test fail 00 = valid test, normal cable (no short or open in cable) 10 = valid test, open in cable (Impedance > 333 ohm) 01 = valid test, short in cable (Impedance < 33 ohm)
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). See Figure 36 . 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 36: Cable Fault Distance Trend Line




Table 113: PHY Specific Control Register II
Offset: 0x1C (Hex), or 28 (Decimal)

Bits	Field	Mode	HW Rst	SW Rst	Description
15:1	Reserved	R/W	0x0	0x0	Must be 0000000000000000
0	SelClsA	R/W	SEL_ CLASS A/B ¹	Update	1 = Select Class A driver; available for 100BASE-TX mode only - typically used in backplane or direct connect applications 0 = Select Class B driver - typically used in CAT 5 applications

1. This is the PHY SEL_CLASSA/B bit Hardware Reset Name. The CONFIGB Pin of the 88E6083 device contains an internal pull-up resistor, setting the default SEL_CLASSA/B PHY bit Hardware Reset to class B drivers.



Note

Registers 0x1D through 0x1F (hexadecimal (29 through 31 decimal)) are reserved. Do not read or write to these registers.

Section 8. EEPROM Programming Format

8.1 EEPROM Programming Details

The 88E6083 device supports an optional external serial EEPROM device for programming its internal registers. The EEPROM data is read in once after RESETn is deasserted if the EEPROM attached Switch Mode is selected (both SW_MODE[1:0] = high—see [Table 14 on page 31](#)).

The 88E6083 supports 1K bit (93C46), 2K bit (93C56), or 4K bit (93C66) EEPROM devices. The size of the device is selected by the EE_DIN/EE_1K pin at reset—see [Table 7 on page 23](#). The external EEPROM device must be configured in x16 data organization mode.

Regardless of which particular device is attached, the EEPROM is read and processed in the same way:

1. Start at EEPROM address 0x00.
2. Read in the 16 bits of data from the even address This step is called the Command.
3. Terminate the serial EEPROM reading process and go to step 8—If the just read in Command is all ones.
4. Increment the address by 1—to the next odd address.
5. Read in the 16 bits of data from the odd address—This is called the RegData.
6. Write RegData to the locations defined by the previous Command.
7. Go to step 2.
8. Set the EEInt bit in Global Status to a one ([Table 60 on page 137](#)) generating an Interrupt (if enabled).
9. Done.

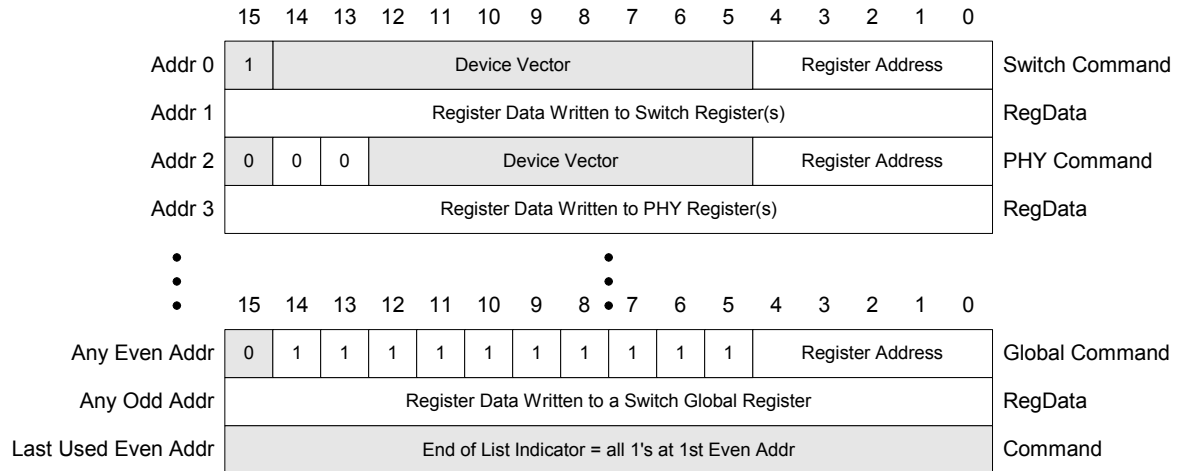
The 16-bit Command determines which register or registers inside the 88E6083 are updated as follows. Refer to [Figure 37](#) below.

- Bit 15 determines which set of registers can be written.
If bit 15 = 0 the PHY registers can be written (SMI Device Addresses 0x00 to 0x07) or the Switch Global registers can be written (SMI Device Address 0x7).
If bit 15 = 1 the Switch registers can be written (SMI Device Addresses 0x10 to 0x19). See [Section 5](#) and [Figure 33](#) for more information on the registers and their addresses.
- Bits 14:5 (or 12:5), the Device Vector, determines which Device 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 0x10).
- Bit 6 controls writes for Port 1 (either PHY address 0x01 or Switch address 0x11).
- Bit 7 controls writes for Port 2
- The Switch Global registers (Switch address 0x1B) are written when bit 15 is a zero and bits 14:5 are ones. Care is needed to ensure that Reserved registers are not written.
- Bits 4:0, the Register Address, determine which SMI Register Address is written.

The format of the EEPROM Commands supports writing the same RegData to all the PHYs or all the Switch's MACs 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 EEPROM¹. The maximum number of Command/RegData pairs is one fewer than might be expected from a simple calculation because the last entry must be the End of List Indicator of all ones.

1. 31 in the 93C46, 63 in the 93C56, and 127 in the 93C66 EEPROM.

Figure 37: EEPROM Data Format



Section 9. Electrical Specifications

9.1 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.

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} 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

- V_{DD(2.5)} must never be more than 0.5V greater than V_{DD(3.3)} or damage will result. This implies that power must be applied to V_{DD(3.3)} before or at the same time as V_{DD(2.5)}.
- V_{DD(1.5)} must never be more than 0.5V greater than V_{DD(2.5)} or damage will result. This implies that power must be applied to V_{DD(2.5)} before or at the same time as V_{DD(1.5)}.
- V_{PIN} must never be more than 0.5V greater than V_{DDO} or damage will result.
- 125°C is the re-bake temperature. For extended storage time greater than 24 hours, +85°C should be the maximum.

9.2 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}	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	0		70	°C
		Industrial Parts ¹	-40		85	°C
T _J	Maximum junction temperature				125 ²	°C
RSET	Internal bias reference	Resistor value placed between RSET- and RSET+ pins	1980	2000	2020	Ω

- Industrial part numbers have an "I" following the commercial part numbers. [Section Section 11. "Ordering Part Numbers and Package Markings" on page 221](#)
- Refer to white paper on T_J Thermal Calculations for more Information.

9.3 Package Thermal Information

9.3.1 Thermal Conditions for 216-pin LQFP Package

Symbol	Parameter	Condition	Min	Typ	Max	Units
θ_{JA}	Thermal resistance ¹ - junction to ambient of the 88E6083 device 216-Pin LQFP package $\theta_{JA} = (T_J - T_A) / P$ P = Total Power Dissipation	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		35.8		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/sec air flow		32.2		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/sec air flow		30.8		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/sec air flow		30		°C/W
ψ_{JT}	Thermal characteristic parameter ¹ - junction to top center of the 88E6083 device 216-Pin LQFP package $\psi_{JT} = (T_J - T_{TOP}) / P$ T_{TOP} = Temperature on the top center of the package	JEDEC 3 in. x 4.5 in. 4-layer PCB with no air flow		0.39		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 1 meter/sec air flow		-		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 2 meter/sec air flow		-		°C/W
		JEDEC 3 in. x 4.5 in. 4-layer PCB with 3 meter/sec air flow		-		°C/W
θ_{JC}	Thermal resistance ¹ - junction to case of the 88E6083 device 216-Pin LQFP package $\theta_{JC} = (T_J - T_C) / P_{TOP}$ P_{TOP} = Power Dissipation from the top of the package	JEDEC with no air flow		7.9		°C/W
θ_{JB}	Thermal resistance ¹ - junction to board of the 88E6083 device 216-Pin LQFP package $\theta_{JB} = (T_J - T_B) / P_{bottom}$ P_{bottom} = power dissipation from the bottom of the package to the PCB surface.	JEDEC with no air flow		29.9		°C/W

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

9.3.2 Current Consumption

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

Symbol	Parameter	Pins	Condition	Min	Typ ¹	Max	Units
I _{DD(3.3)}	3.3 volt Power to outputs	V _{DDO}	Both MII ports disabled		20		mA
			Both MII ports 10 Mbps		25		mA
			Both MII ports 100 Mbps		30		mA
I _{DD(2.5)}	2.5 volt Power to analog core	V _{DDAH}	No link on any port		65		mA
			All ports 10 Mbps linked but idle		55		mA
			All ports 10 Mbps and active		200		mA
			All ports 100 Mbps active or idle		180		mA
			All ports powered down		55		mA
	2.5V volt Power to Center Tap	External magnetics Center Tap	No link on any port		15		mA
			All ports 10 Mbps linked but idle		1		mA
			All ports 10 Mbps and active		480		mA
			All ports 100 Mbps active or idle		160 ²		mA
			All ports powered down		0		mA
I _{DD(1.5)}	1.5 volt Power to analog core	V _{DDAL}	No link on any port		2		mA
			All ports 10 Mbps linked but idle		2		mA
			All ports 10 Mbps and active		2		mA
			All ports 100 Mbps active or idle		50		mA
			All ports powered down		2		mA
	1.5 volt Power to digital core	V _{DD}	No link on any port		100		mA
			All ports 10 Mbps linked but idle		120		mA
			All ports 10 Mbps and active		125		mA
			All ports 100 Mbps linked but idle		200		mA
			All ports 100 Mbps and active		230		mA
			All ports powered down		40		mA

1. The values listed are typical values with LED mode 3, 3 LEDs per port, Auto-Negotiation on, and Energy Detect enabled.
2. The 2.5V power source must supply an additional 20 mA per port for a Class A operation vs. Class B operation listed.

9.3.3 Digital Operating Conditions

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

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
		INTn ²	I _{OH} = -8 mA V _{DDO} = Min	2.4			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	0<V _{IN} <V _{DD}			+ 10 - 50	μA
		With pull-down resistor	0<V _{IN} <V _{DD}			+ 50 - 10	μA
		All others	0<V _{IN} <V _{DD}			±10	μA
C _{IN}	Input capacitance	All pins				5	pF

1. The 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]

2. The INTn is an active low, open drain pin. See INTn description in the Signal Description.



Table 114: Internal Resistor Description

Pin #	Pin Name	Resistor	Pin #	Pin Name	Resistor
199	CONFIG_A	Internal pull-up	130,131,132, 133	P8_OUTD[3:0]/ P8_MODE[3:0]	Internal pull-up
201	CONFIG_B	Internal pull-up	136	P8_OUTDV	Internal pull-up
121	CLK_SEL	Internal pull-up	127	P8_CRS	Internal pull-down
125	MDC	Internal pull-up	128	P8_COL	Internal pull-down
123	MDIO	Internal pull-up	151	DISABLE_P9	Internal pull-up
116	EE_CS/EE_1K	Internal pull-up	161	P9_INCLK	Internal pull-up
114	EE_CLK	Internal pull-up	167,169,171, 173	P9_IND[3:0]	Internal pull-up
117	EE_DIN/ HD_FLOW_DIS	Internal pull-up	175	P9_INDV	Internal pull-down
112	EE_DOUT	Internal pull-up	158	P9_OUTCLK	Internal pull-up
120	ENABLE_P8	Internal pull-up	153,154,155 156	P9_OUTD[3:0]/ P9_MODE[3:0]	Internal pull-up
138	P8_INCLK	Internal pull-up	148	P9-CRS	Internal pull-down
140,141,142, 143	P5_IND[3:0]	Internal pull-up	150	P9_COL	Internal pull-down
145	P8_INDV	Internal pull-down	110, 107	SW_MODE[1:0]	Internal pull-up
135	P8_OUTCLK	Internal pull-up	118	FD_FLOW_DIS	Internal pull-up

9.3.4 IEEE DC Transceiver Parameters

IEEE tests are typically based on templates and cannot simply be specified by a number. For an exact description of the template and the test conditions, refer to the IEEE specifications:

-10BASE-T IEEE 802.3 Clause 14

-100BASE-TX ANSI X3.263-1995

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

Table 115: 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 88E6083 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 88E6083 will reject signals typically with peak-to-peak differential amplitude less than 360 mV.

9.4 AC Electrical Specifications

9.4.1 Reset and Configuration Timing


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

Table 116: Reset and Configuration Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{PU_RESET}	Valid power to RESETn de-asserted		10			ms
T_{SU_CLK}	Number of valid REFCLK cycles prior to RESETn de-asserted		10			Clks
T_{SU}	Configuration data valid prior to RESETn ¹ de-asserted		200			ns
T_{HD}	Configuration data valid after RESETn de-asserted		0			ns
T_{CO}	Configuration output driven after RESETn ² de-asserted		40			ns

1. When RESETn is low all configuration pins become inputs, and the value seen on these pins is latched on the rising edge of RESETn.
2. P[x]_OUTD[3:0]/P[x]_MODE[3:0], and P[8]_OUTDV/WD_DIS (Table 10) are normally outputs that are also used to configure the 88E6083 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 38: Reset and Configuration Timing



9.4.2 Clock Timing with a 25 MHz Oscillator


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

Table 117: Clock Timing with a 25 MHz Oscillator

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P^1	XTAL_IN period		40 -50 ppm	40	40 +50 ppm	ns
T_H	XTAL_IN high time		16			ns
T_L	XTAL_IN low time		16			ns
T_R	XTAL_IN rise				3	ns
T_F	XTAL_IN fall				3	ns

1. 25 MHz

Figure 39: Oscillator Clock Timing



9.4.3 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 118: MII Receive Timing-PHY Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_TX_CLK}^1$	P[x]_INCLK period	10BASE mode		400		ns
		100BASE mode		40		ns
$T_{H_TX_CLK}^1$	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 40: PHY Mode MII Receive Timing



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

9.4.4 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 119: MII Transmit Timing-PHY Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_RX_CLK}^1$	P[x]_OUTCLK period	10BASE mode		400		ns
		100BASE mode		40		ns
$T_{H_RX_CLK}^1$	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				27	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 41: PHY Mode MII Transmit Timing



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

9.4.5 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 120: MAC Mode Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P ¹	MACCLK_IN period		0	40	40 +50 ppm	ns
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

1. DC to 25 MHz

Figure 42: MAC Clock Timing



9.4.6 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 121: MII Receive Timing-MAC Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units
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 43: MAC Mode MII Receive Timing



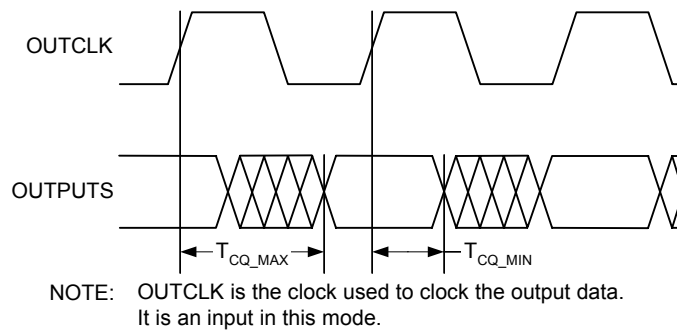
9.4.7 MII Transmit Timing - MAC Mode

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 122: MII Transmit Timing-MAC Mode

Symbol	Parameter	Condition	Min	Typ	Max	Units
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 44: MAC Mode MII Transmit Timing



9.4.8 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 123: MAC Mode Clock Timing, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P^1	MACCLK_IN period		0	20	20 +50 ppm	ns
T_H	MACCLK_IN high time		8			ns
T_L	MACCLK_IN low time		8			ns
T_R	MACCLK_IN rise				3	ns
T_F	MACCLK_IN fall				3	ns

1. DC to 25 MHz

Figure 45: MAC Clock Timing, 200 Mbps




9.4.9 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 124: MII Receive Timing - PHY Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_TX_CLK}$	P[x]_INCLK period	200BASE mode		20		ns
$T_{H_TX_CLK}$	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.		13.5			ns
T_{HD_TX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high.		-1.5			ns

Figure 46: 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.

9.4.10 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 125: MII Transmit Timing-PHY Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_RX_CLK}^1$	P[x]_OUTCLK period	200BASE mode		20		ns
$T_{H_RX_CLK}$	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		7.2			ns

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

Figure 47: 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.

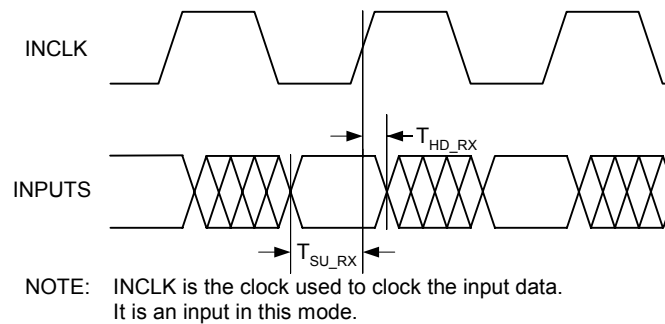
9.4.11 MII Receive Timing-MAC Mode, 200 Mbps

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 126: MII Receive Timing-MAC Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units
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	7			ns
T_{HD_RX}	MII inputs (P[x]_IND[3:0], P[x]_INDV) valid after P[x]_INCLK going high	With 10pF load	2			ns

Figure 48: MAC Mode MII Receive Timing - 200 Mbps




9.4.12 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 127: MII Transmit Timing—MAC Mode, 200 Mbps

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_{CQ_MAX}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV valid)	With 10 pF load			10.5	ns
T_{CQ_MIN}	P[x]_OUTCLK to outputs (P[x]_OUTD[3:0], P[x]_OUTDV invalid)	With 10pF load	3			ns

Figure 49: 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.


9.4.13 SNI Falling Edge Receive Timing

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

Table 128: SNI Falling Edge Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_FRX_CLK}$	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 50: SNI Falling Edge Receive Timing



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


9.4.14 SNI Falling Edge Transmit Timing

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

Table 129: SNI Falling Edge Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_FTX_CLK}$	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 51: SNI Falling Edge Transmit Timing



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


9.4.15 SNI Rising Edge Receive Timing

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

Table 130: SNI Rising Edge Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_RRX_CLK}$	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 52: SNI Rising Edge Receive Timing



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


9.4.16 SNI Rising Edge Transmit Timing

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

Table 131: SNI Rising Edge Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_RTX_CLK}$	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 53: SNI Rising Edge Transmit Timing



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

9.4.17 RMI Receive Timing


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

Table 132: RMI Receive Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_TX_CLK}^1$	P[x]_INCLK period	100BASE mode		20		ns
$T_{H_TX_CLK}$	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.		10			ns
T_{HD_TX}	MII inputs (P[x]_IND[1:0], P[x]_INDV) valid after P[x]_INCLK going high.		0			ns

1. 25 MHz for 100 Mbps.

Figure 54: PHY Mode MII Receive Timing



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

9.4.18 RMI Transmit Timing


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

Table 133: RMI Transmit Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
$T_{P_RX_CLK}^1$	P[x]_INCLK period	100BASE mode		20		ns
$T_{H_RX_CLK}$	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				10	ns
T_{CQ_MIN}	P[x]_INCLK to outputs (P[x]_OUTD[1:0], P[x]_OUTDV) invalid		0			ns

1. 25 MHz for 100 Mbps.

Figure 55: PHY Mode MII Transmit Timing



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


9.4.19 Serial LED Timing

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

Table 134: Serial LED Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
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 (LEDSE, LEDENA) valid	With 10 pF load			20	ns
T_{CQ_MIN}	LEDCLK falling edge to link outputs (LEDSE, LEDENA) invalid	With 10 pF load	0			ns

Figure 56: Serial LED Timing



9.4.20 Serial Management Interface Clock Timing


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

Table 135: Serial Management Interface Clock Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P	MDCCLK_IN period		120 ¹			ns
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

1. 8.33 MHz

Figure 57: Serial Management Interface Clock Timing



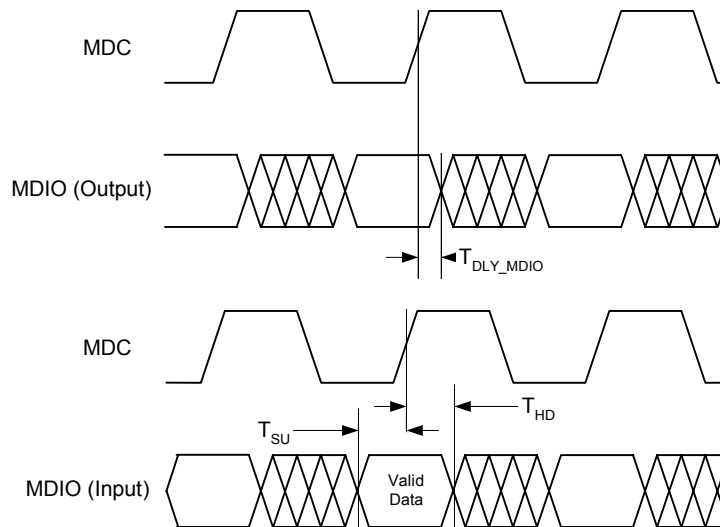
9.4.21 Serial Management Interface Timing

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

Table 136: Serial Management Interface Timing

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

Figure 58: Serial Management Interface Timing




9.4.22 EEPROM Timing

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

Table 137: EEPROM Timing

Symbol	Parameter	Condition	Min	Typ	Max	Units
T_P	EE_CLK period			5120		ns
T_H	EE_CLK high time			2560		ns
T_L	EE_CLK low time			2560		ns
T_{CQCSMX}	Serial EEPROM chip select valid	Referenced to EE_CLK			5	ns
T_{CQCSMN}	Serial EEPROM chip select invalid				5	ns
T_{CQDMX}	Serial EEPROM data transmitted to EEPROM valid				10	ns
T_{CQDMN}	Serial EEPROM data transmitted to EEPROM invalid		3			ns
T_S	Setup time for data received from EEPROM		10			ns
T_H	Hold time for data received from EEPROM		10			ns

Figure 59: EEPROM Timing



9.4.23 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 138: 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 10. Package Mechanical Dimensions




Figure 60: 88E6083 216-pin LQFP Package



Symbol	Dimension in mm		
	Min	Nom	Max
A	--	--	1.60
A ₁	0.05	--	0.15
A ₂	1.35	1.40	1.45
b	0.13	0.18	0.23
b ₁	0.13	0.16	0.19
c	0.09	0.14	0.20
c ₁	0.09	0.12	0.16
D	25.60	26.00	26.40
D ₁	--	24.00	--
E	25.60	26.00	26.40
E ₁	--	24.00	--
e	0.40 BSC		
L	0.45	0.60	0.75
L ₁	1.00 REF		
R ₁	0.08	--	--
R ₂	0.08	--	--
S	0.20	--	--
q	0°	3.5°	7°
θ ₁	0°	--	--
θ ₂	11°	12°	13°
θ ₃	11°	12°	13°

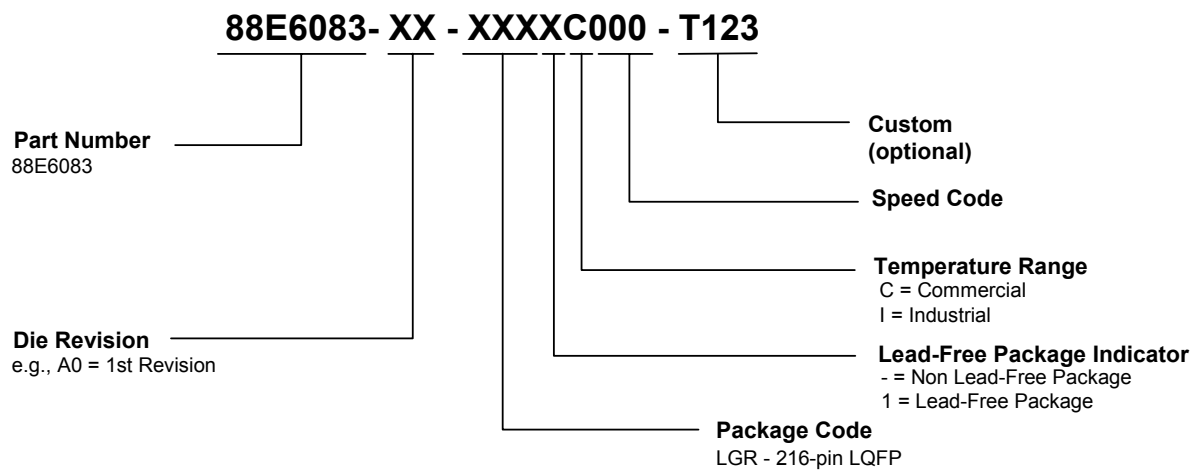
Notes:

10. TO BE DETERMINED AT SEATING PLANE -C-.
11. DIMENSIONS D1 AND E1 DO NOT INCLUDE MOLD PROTRUSION. D1 AND E1 ARE MAXIMUM PLASTIC BODY SIZE DIMENSIONS INCLUDING MOLD MISMATCH.
12. DIMENSION b DOES NOT INCLUDE DAMBAR PROTRUSION. DAMBAR CAN NOT BE LOCATED ON THE LOWER RADIUS OR THE FOOT.
13. EXACT SHAPE OF EACH CORNER IS OPTIONAL
14. THESE DIMENSIONS APPLY TO THE FLAT SECTION OF THE LEAD BETWEEN 0.10 mm AND 0.25 mm FROM THE LEAD TIP.
15. A1 IS DEFINED AS THE DISTANCE FROM THE SEATING PLANE TO THE LOWEST POINT OF THE PACKAGE BODY.
16. CONTROLLING DIMENSION: MILLIMETER.

Section 11. Ordering Part Numbers and Package Markings

Figure 61 shows the ordering part numbering scheme for the 88E6083. Contact Marvell® FAEs or sales representatives for complete ordering information.

Figure 61: 88E6083 Sample Ordering Part Number



The standard ordering part numbers for the available alternatives are:

Commercial Temp:

Non Lead-Free Package: 88E6083-XX-LGR-C000

Lead-Free Package: 88E6083-XX-LGR1C000

Industrial Temp:

Non Lead-Free Package: 88E6083-XX-LGR-I000

Lead-Free Package: 88E6083-XX-LGR1I000



Figure 62 shows a typical package marking and pin 1 location for an 88E6083 part.

Figure 62: 88E6083 Commercial Package Marking and Pin 1 Location

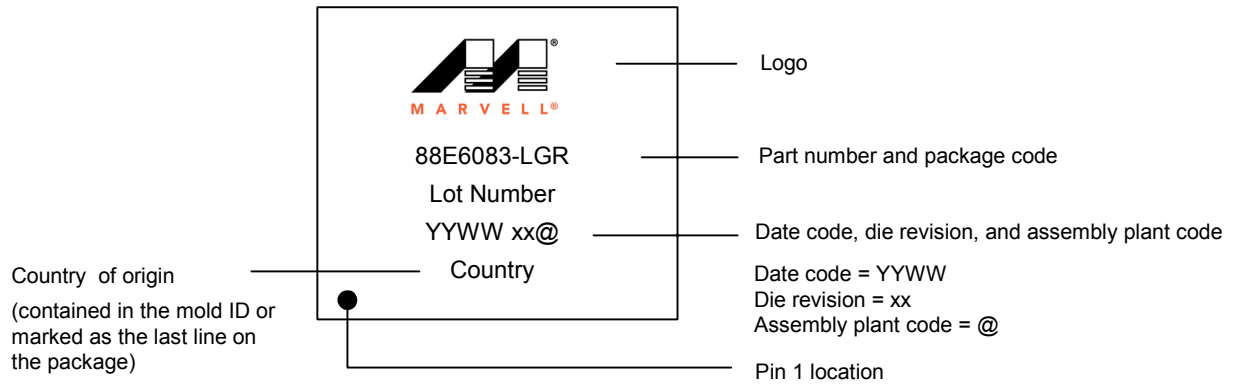
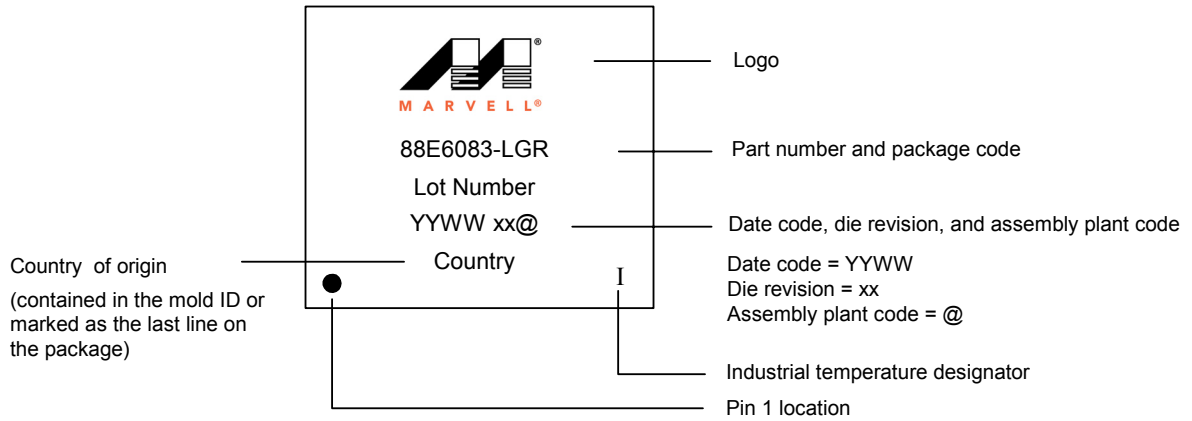


Figure 63: 88E6083 Industrial Package Marking and Pin 1 Location





MOVING FORWARD
FASTER®

Marvell Semiconductor, Inc.

700 First Avenue
Sunnyvale, CA 94089

Phone 408.222.2500
Fax 408.752.9028

www.marvell.com

US and Worldwide Offices

Marvell Semiconductor, Inc.

700 First Avenue
Sunnyvale, CA 94089
Tel: 1.408.222.2500
Fax: 1.408.752.9028

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. 50F
1-25-1, Nishi-Shinjuku, Shinjuku-ku
Tokyo 163-0650
Tel: 81.(0).3.5324.0355
Fax: 81.(0).3.5324.0354

Marvell Semiconductor Israel, Ltd.

Moshav Manof
D.N. Misgav 20184
Israel
Tel: 972.4.995.1000
Fax: 972.4.995.1001

Worldwide Sales Offices

Western US Sales Office

Marvell
700 First Avenue
Sunnyvale, CA 94089
Tel: 1.408.222.2500
Fax: 1.408.752.9028
Sales Fax: 1.408.752.9029

Central US Sales Office

Marvell
11709 Boulder Lane, Ste. #220
Austin, TX 78726
Tel: 1.512.336.1551
Fax: 1.512.336.1552

Eastern US/Canada Sales Office

Marvell
Parlee Office Park
1 Meeting House Road, Suite 1
Chelmsford, MA 01824
Tel: 978 250-0588
Fax: 978 250-0589

Europe Sales Office

Marvell
3 Clifton Court
Corner Hall
Hemel Hempstead
Hertfordshire, HP3 9XY
United Kingdom
Tel: 44.(0).1442.211668
Fax: 44.(0).1442.211543

Marvell

Fagerstagatan 4
163 08 Spanga
Stockholm, Sweden
Tel: 46.16.146348
Fax: 46.16.482425

Marvell

5 Rue Poincare
56400 Le Bono
France
Tel: 33.297.579697
Fax: 33.297.578933

Israel Sales Office

Marvell
Ofek Center Bldg. 2, Floor 2
Northern Industrial Zone
LOD 71293
Israel
Tel: 972.8.914.1300
Fax: 972.8.914.1301

China Sales Office

Marvell
5J, 1800 Zhong Shan West Road
Shanghai, China 200233
Tel: 86.21.6440.1350
Fax: 86.21.6440.0799

Japan Sales Office

Marvell
Helios Kannai Bldg. 12F
3-21-2 Motohama-cho, Naka-ku
Yokohama, Kanagawa
Japan 231-0004
Tel: 81.45.222.8811
Fax: 81.45.222.8812

Taiwan Sales Office

Marvell
2Fl., No. 1, Alley 20, Lane 407
Ti-Ding Blvd., Nei Hu District
Taipei, Taiwan 114, R. O. C
Tel: (886-2).7720.5700
FAX: (886-2).7720.5707

For more information, visit our website at: www.marvell.com

Copyright © 2005, Marvell International Ltd. All rights reserved. Marvell, the Marvell logo, Moving Forward Faster, Alaska, Fastwriter, GalNet, Libertas, Link Street, NetGX, PHYAdvantage, Prestera, Virtual Cable Tester, and Yukon are registered trademarks of Marvell. AnyVoltage, Discovery, DSP Switcher, Feroceon, GalTis, Horizon, RADLAN, Raising The Technology Bar, The Technology Within, UniMAC, and VCT are trademarks of Marvell. All other trademarks are the property of their respective owners.

X-ON Electronics

Largest Supplier of Electrical and Electronic Components

Click to view similar products for [Marvell](#) manufacturer:

Other Similar products are found below :

[88E1112-C2-NNC1C000](#) [88E1512-A0-NNP2C000](#) [88E1512-A0-NNP2I000](#) [88E1518-A0-NNB2C000](#) [88E1543-A1-LKJ2C000](#) [88E2110-A0-BWS4C000](#) [88E6083-B0-LGR1I000](#) [88E6085-A2-TAH1I000](#) [88E6185-A2-LKJ1C000](#) [88E6320-A0-NAZ2C000](#) [88Q2112-A2-NYD2A000](#)
[88Q2112-A2-NYD2A000-P123](#) [88SE9170A2-NNX2I000](#) [88X3310-A1-BUS4C000](#) [98DX4122C0-BIH2I000](#)