
6-Port Gigabit Ethernet Switch with GMII/RGMII/MII/RMII Interface

Highlights

- One port with 10/100/1000 Ethernet MAC and configurable GMII/RGMII/MII/RMII interface
- IEEE 802.1X access control support
- Five ports with integrated 10/100/1000BASE-T PHY transceivers
- Non-blocking wire-speed Ethernet switching fabric
- Full-featured forwarding and filtering control, including Access Control List (ACL) filtering
- Full VLAN and QoS support
- EtherGreen™ power management features, including low power standby
- Flexible management interface options: SPI, I²C, MIIM, and in-band management via any port
- Commercial/Industrial temperature range support
- 128-pin TQFP-EP (14 x 14mm) RoHS compliant pkg

Target Applications

- Stand-alone 10/100/1000Mbps Ethernet switches
- VoIP infrastructure switches
- Broadband gateways/firewalls
- Wi-Fi access points
- Integrated DSL/cable modems
- Security/surveillance systems
- Industrial control/automation switches
- Networked measurement and control systems

Features

- Switch Management Capabilities
 - 10/100/1000Mbps Ethernet switch basic functions: frame buffer management, address look-up table, queue management, MIB counters
 - Non-blocking store-and-forward switch fabric assures fast packet delivery by utilizing 4096 entry forwarding table with 256kByte frame buffer
 - Jumbo packet support up to 9000 bytes
 - Port mirroring/monitoring/sniffing: ingress and/or egress traffic to any port
 - MIB counters for fully-compliant statistics gathering 34 counters per port
 - Tail tagging mode (one byte added before FCS) support at host port to inform the processor which ingress port receives the packet and its priority
 - Loopback modes for remote failure diagnostics
 - Rapid spanning tree protocol (RSTP) support for topology management and ring/linear recovery
 - Multiple spanning tree protocol (MSTP) support

- One Ext. MAC Port with GMII/RGMII/MII/RMII
 - RGMII v2.0, RMII v1.2 with 50MHz reference clock input/output option, MII in PHY/MAC mode
- Five Integrated PHY Ports
 - 1000BASE-T/100BASE-TX/10BASE-Te IEEE 802.3
 - Fast Link-up option significantly reduces link-up time
 - Auto-negotiation and Auto-MDI/MDI-X support
 - On-chip termination resistors and internal biasing for differential pairs to reduce power
 - LinkMD® cable diagnostic capabilities for determining cable opens, shorts, and length
- Advanced Switch Capabilities
 - IEEE 802.1Q VLAN support for 128 active VLAN groups and the full range of 4096 VLAN IDs
 - IEEE 802.1p/Q tag insertion/removal on per port basis
 - VLAN ID on per port or VLAN basis
 - IEEE 802.3x full-duplex flow control and half-duplex back pressure collision control
 - IEEE 802.1X access control (Port-based and MAC address based)
 - IGMP v1/v2/v3 snooping for multicast packet filtering
 - IPv6 multicast listener discovery (MLD) snooping
 - IPv4/IPv6 QoS support, QoS/CoS packet prioritization
 - 802.1p QoS packet classification with 4 priority queues
 - Programmable rate limiting at ingress/egress ports
 - Broadcast storm protection
 - Four priority queues with dynamic packet mapping for IEEE 802.1p, IPv4 DIFFSERV, IPv6 Traffic Class
 - MAC filtering function to filter or forward unknown unicast, multicast and VLAN packets
 - Self-address filtering for implementing ring topologies
- Comprehensive Configuration Registers Access
 - High-speed 4-wire SPI (up to 50MHz), I²C interfaces provide access to all internal registers
 - MII Management (MIIM, MDC/MDIO 2-wire) Interface provides access to all PHY registers
 - In-band management via any of the data ports
 - I/O pin strapping facility to set certain register bits from I/O pins at reset time
 - On-the-fly configurable control registers
- Power Management
 - Energy detect power-down mode on cable disconnect
 - Dynamic clock tree control
 - Unused ports can be individually powered down
 - Full-chip software power-down
 - Wake-on-LAN (WoL) standby power mode with PME interrupt output for system wake upon triggered events

TO OUR VALUED CUSTOMERS

It is our intention to provide our valued customers with the best documentation possible to ensure successful use of your Microchip products. To this end, we will continue to improve our publications to better suit your needs. Our publications will be refined and enhanced as new volumes and updates are introduced.

If you have any questions or comments regarding this publication, please contact the Marketing Communications Department via E-mail at docerrors@microchip.com. We welcome your feedback.

Most Current Documentation

To obtain the most up-to-date version of this documentation, please register at our Worldwide Web site at:

<http://www.microchip.com>

You can determine the version of a data sheet by examining its literature number found on the bottom outside corner of any page. The last character of the literature number is the version number, (e.g., DS30000000A is version A of document DS30000000).

Errata

An errata sheet, describing minor operational differences from the data sheet and recommended workarounds, may exist for current devices. As device/documentation issues become known to us, we will publish an errata sheet. The errata will specify the revision of silicon and revision of document to which it applies.

To determine if an errata sheet exists for a particular device, please check with one of the following:

- Microchip's Worldwide Web site; <http://www.microchip.com>
- Your local Microchip sales office (see last page)

When contacting a sales office, please specify which device, revision of silicon and data sheet (include -literature number) you are using.

Customer Notification System

Register on our web site at www.microchip.com to receive the most current information on all of our products.

Table of Contents

| | |
|---|-----|
| 1.0 Preface | 4 |
| 2.0 Introduction | 8 |
| 3.0 Pin Descriptions and Configuration | 9 |
| 4.0 Functional Description | 20 |
| 5.0 Device Registers | 64 |
| 6.0 Operational Characteristics | 169 |
| 7.0 Design Guidelines | 185 |
| 8.0 Package Information | 188 |
| Appendix A: Data Sheet Revision History | 190 |
| The Microchip Web Site | 191 |
| Customer Change Notification Service | 191 |
| Customer Support | 191 |
| Product Identification System | 192 |

KSZ9896C

1.0 PREFACE

1.1 Glossary of Terms

TABLE 1-1: GENERAL TERMS

| Term | Description |
|---------------------|---|
| 10BASE-Te | 10 Mbps Ethernet, 2.5V signaling, IEEE 802.3 compliant |
| 100BASE-TX | 100 Mbps Fast Ethernet, IEEE 802.3u compliant |
| 1000BASE-T | 1000 Mbps Gigabit Ethernet, IEEE 802.3ab compliant |
| ADC | Analog-to-Digital Converter |
| AN | Auto-Negotiation |
| BLW | Baseline Wander |
| BPDU | Bridge Protocol Data Unit. Messages which carry the Spanning Tree Protocol information. |
| Byte | 8 bits |
| CRC | Cyclic Redundancy Check. A common technique for detection data transmission errors. CRC for Ethernet is 32 bits long. |
| CSR | Control and Status Registers |
| DA | Destination Address |
| DWORD | 32 bits |
| FCS | Frame Check Sequence. The extra checksum characters added to the end of an Ethernet frame, used for error detection and correction. |
| FID | Frame or Filter ID. Specifies the frame identifier. Alternately is the filter identifier. |
| FIFO | First In First Out buffer |
| FSM | Finite State Machine |
| GPIO | General Purpose I/O |
| Host | External system (Includes processor, application software, etc.) |
| IGMP | Internet Group Management Protocol. Defined by RFC 1112, RFC 2236, and RFC 4604 to establish multicast group membership in IPv4 networks. |
| IPG | Inter-Packet Gap. A time delay between successive data packets mandated by the network standard for protocol reasons. |
| Jumbo Packet | A packet larger than the standard Ethernet packet (1518 bytes). Large packet sizes allow for more efficient use of bandwidth, lower overhead, less processing, etc. |
| Isb | Least Significant Bit |
| LSB | Least Significant Byte |
| MAC | Media Access Controller. A functional block responsible for implementing the media access control layer, which is a sublayer of the data link layer. |
| MDI | Medium Dependent Interface. An Ethernet port connection that allows network hubs or switches to connect to other hubs or switches without a null-modem, or crossover, cable. |
| MDIX | Media Independent Interface with Crossover. An Ethernet port connection that allows networked end stations (i.e., PCs or workstations) to connect to each other using a null-modem, or crossover, cable. |
| MIB | Management Information Base. The MIB comprises the management portion of network devices. This can include monitoring traffic levels and faults (statistical), and can also change operating parameters in network nodes (static forwarding addresses). |
| MII | Media Independent Interface. The MII accesses PHY registers as defined in the IEEE 802.3 specification. |
| MIIM | Media Independent Interface Management |

TABLE 1-1: GENERAL TERMS (CONTINUED)

| Term | Description |
|-----------------|--|
| MLD | Multicast Listening Discovery. This protocol is defined by RFC 3810 and RFC 4604 to establish multicast group membership in IPv6 networks. |
| MLT-3 | Multi-Level Transmission Encoding (3-Levels). A tri-level encoding method where a change in the logic level represents a code bit "1" and the logic output remaining at the same level represents a code bit "0". |
| msb | Most Significant Bit |
| MSB | Most Significant Byte |
| NRZ | Non Return to Zero. A type of signal data encoding whereby the signal does not return to a zero state in between bits. |
| NRZI | Non Return to Zero Inverted. This encoding method inverts the signal for a "1" and leaves the signal unchanged for a "0" |
| N/A | Not Applicable |
| NC | No Connect |
| OUI | Organizationally Unique Identifier |
| PHY | A device or function block which performs the physical layer interface function in a network. |
| PLL | Phase Locked Loop. A electronic circuit that controls an oscillator so that it maintains a constant phase angle (i.e., lock) on the frequency of an input, or reference, signal. |
| RESERVED | Refers to a reserved bit field or address. Unless otherwise noted, reserved bits must always be zero for write operations. Unless otherwise noted, values are not guaranteed when reading reserved bits. Unless otherwise noted, do not read or write to reserved addresses. |
| RTC | Real-Time Clock |
| SA | Source Address |
| SFD | Start of Frame Delimiter. The 8-bit value indicating the end of the preamble of an Ethernet frame. |
| SQE | Signal Quality Error (also known as "heartbeat") |
| SSD | Start of Stream Delimiter |
| TCP | Transmission Control Protocol |
| UDP | User Datagram Protocol - A connectionless protocol run on top of IP networks |
| UTP | Unshielded Twisted Pair. Commonly a cable containing 4 twisted pairs of wire. |
| UUID | Universally Unique Identifier |
| VLAN | Virtual Local Area Network |
| WORD | 16 bits |

KSZ9896C

1.2 Buffer Types

TABLE 1-2: BUFFER TYPES

| Buffer Type | Description |
|-------------|---|
| I | Input |
| IPU | Input with internal pull-up (58 k Ω \pm 30%) |
| IPU/O | Input with internal pull-up (58 k Ω \pm 30%) during power-up/reset; output pin during normal operation |
| IPD | Input with internal pull-down (58 k Ω \pm 30%) |
| IPD/O | Input with internal pull-down (58 k Ω \pm 30%) during power-up/reset; output pin during normal operation |
| O8 | Output with 8 mA sink and 8 mA source |
| O24 | Output with 24 mA sink and 24 mA source |
| OPU | Output (8mA) with internal pull-up (58 k Ω \pm 30%) |
| OPD | Output (8mA) with internal pull-down (58 k Ω \pm 30%) |
| AIO | Analog bidirectional |
| ICLK | Crystal oscillator input pin |
| OCLK | Crystal oscillator output pin |
| P | Power |
| GND | Ground |

Note: Refer to [Section 6.3, "Electrical Characteristics,"](#) on page 170 for the electrical characteristics of the various buffers.

1.3 Register Nomenclature

TABLE 1-3: REGISTER NOMENCLATURE

| Register Bit Type Notation | Register Bit Description |
|----------------------------|--|
| R | Read: A register or bit with this attribute can be read. |
| W | Write: A register or bit with this attribute can be written. |
| RO | Read only: Read only. Writes have no effect. |
| RC | Read to Clear: Contents is cleared after the read. Writes have no effect. |
| WO | Write only: If a register or bit is write-only, reads will return unspecified data. |
| WC | Write One to Clear: Writing a one clears the value. Writing a zero has no effect. |
| LL | Latch Low: Applies to certain RO status bits. If a status condition causes this bit to go low, it will maintain the low state until read, even if the status condition changes. A read clears the latch, allowing the bit to go high if dictated by the status condition. |
| LH | Latch High: Applies to certain RO status bits. If a status condition causes this bit to go high, it will maintain the high state until read, even if the status condition changes. A read clears the latch, allowing the bit to go low if dictated by the status condition. |
| SC | Self-Clearing: Contents are self-cleared after the being set. Writes of zero have no effect. Contents can be read. |
| RESERVED | Reserved Field: Reserved fields must be written with zeros, unless otherwise indicated, to ensure future compatibility. The value of reserved bits is not guaranteed on a read. |

1.4 References

- NXP I²C-Bus Specification (UM10204, April 4, 2014): www.nxp.com/documents/user_manual/UM10204.pdf

KSZ9896C

2.0 INTRODUCTION

2.1 General Description

The KSZ9896C is a highly-integrated, IEEE 802.3 compliant networking device that incorporates a layer-2 managed Gigabit Ethernet switch, five 10BASE-T/100BASE-TX/1000BASE-T physical layer transceivers (PHYs) and associated MAC units, and one configurable GMII/RGMII/MII/RMII MAC port for direct connection to a host processor/controller, another Ethernet switch, or an Ethernet PHY transceiver.

The KSZ9896C is built upon industry-leading Ethernet technology, with features designed to offload host processing and streamline the overall design:


- Non-blocking wire-speed Ethernet switch fabric
- Full-featured forwarding and filtering control, including port-based Access Control List (ACL) filtering
- Full VLAN and QoS support
- Traffic prioritization with per-port ingress/egress queues and by traffic classification
- Spanning Tree support
- IEEE 802.1X access control support

A host processor can access all KSZ9896C registers for control over all PHY, MAC, and switch functions. Full register access is available via the integrated SPI or I²C interfaces, and by in-band management via any one of the data ports. PHY register access is provided by a MIIM interface. Flexible digital I/O voltage allows the MAC port to interface directly with a 1.8/2.5/3.3V host processor/controller/FPGA.

Additionally, a robust assortment of power-management features including Wake-on-LAN (WoL) for low power standby operation, have been designed to satisfy energy-efficient system requirements.

The KSZ9896C is available in commercial (0°C to +70°C) and industrial (-40°C to +85°C) temperature ranges. An internal block diagram of the KSZ9896C is shown in [Figure 2-1](#).

FIGURE 2-1: INTERNAL BLOCK DIAGRAM




3.0 PIN DESCRIPTIONS AND CONFIGURATION

3.1 Pin Assignments

The device pin diagram for the KSZ9896C can be seen in [Figure 3-1](#). [Table 3-1](#) provides a KSZ9896C pin assignment table. Pin descriptions are provided in [Section 3.2, "Pin Descriptions"](#).

FIGURE 3-1: PIN ASSIGNMENTS (TOP VIEW)



Note: When an “_N” is used at the end of the signal name, it indicates that the signal is active low. For example, RESET_N indicates that the reset signal is active low.

The buffer type for each signal is indicated in the “Buffer Type” column of the pin description tables in [Section 3.2, "Pin Descriptions"](#). A description of the buffer types is provided in [Section 1.2, "Buffer Types"](#).

KSZ9896C

TABLE 3-1: PIN ASSIGNMENTS

| Pin | Pin Name | Pin | Pin Name | Pin | Pin Name | Pin | Pin Name |
|-----|----------|-----|------------------|-----|---------------------------------------|-----|-------------------|
| 1 | TXRX1P_A | 33 | AVDDH | 65 | TX_ER6 | 97 | SDI/SDA/MDIO |
| 2 | TXRX1M_A | 34 | TXRX4P_A | 66 | COL6 | 98 | VDDIO |
| 3 | AVDDL | 35 | TXRX4M_A | 67 | TXD6_3 | 99 | SCS_N |
| 4 | TXRX1P_B | 36 | AVDDL | 68 | TXD6_2 | 100 | SCL/MDC |
| 5 | TXRX1M_B | 37 | TXRX4P_B | 69 | TXD6_1 | 101 | GND |
| 6 | TXRX1P_C | 38 | TXRX4M_B | 70 | TXD6_0 | 102 | LED5_0 |
| 7 | TXRX1M_C | 39 | TXRX4P_C | 71 | DVDDL | 103 | LED5_1 (Note 3-1) |
| 8 | TXRX1P_D | 40 | TXRX4M_C | 72 | RX_CLK6/REFCLKO6 | 104 | DVDDL |
| 9 | TXRX1M_D | 41 | AVDDL | 73 | RX_DV6/CRS_DV6/ RX_CTL6 (Note 3-1) | 105 | LED1_0 |
| 10 | AVDDH | 42 | TXRX4P_D | 74 | VDDIO | 106 | LED1_1 (Note 3-1) |
| 11 | DVDDL | 43 | TXRX4M_D | 75 | RX_ER6/RX_CLK6 (Note 3-1) | 107 | GND |
| 12 | TXRX2P_A | 44 | AVDDH | 76 | CRS6 | 108 | NC |
| 13 | TXRX2M_A | 45 | DVDDL | 77 | RXD6_3 (Note 3-1) | 109 | GND |
| 14 | AVDDL | 46 | NC | 78 | RXD6_2 (Note 3-1) | 110 | DVDDL |
| 15 | TXRX2P_B | 47 | GND | 79 | RXD6_1 (Note 3-1) | 111 | AVDDH |
| 16 | TXRX2M_B | 48 | DVDDL | 80 | RXD6_0 (Note 3-1) | 112 | TXRX5P_A |
| 17 | TXRX2P_C | 49 | VDDIO | 81 | GND | 113 | TXRX5M_A |
| 18 | TXRX2M_C | 50 | GND | 82 | LED4_0 (Note 3-1) | 114 | AVDDL |
| 19 | AVDDL | 51 | GND | 83 | LED4_1 (Note 3-1) | 115 | TXRX5P_B |
| 20 | TXRX2P_D | 52 | TXD6_7 | 84 | DVDDL | 116 | TXRX5M_B |
| 21 | TXRX2M_D | 53 | TXD6_6 | 85 | LED3_0 | 117 | TXRX5P_C |
| 22 | AVDDH | 54 | TXD6_5 | 86 | LED3_1 (Note 3-1) | 118 | TXRX5M_C |
| 23 | DVDDL | 55 | TXD6_4 | 87 | NC | 119 | AVDDL |
| 24 | TXRX3P_A | 56 | DVDDL | 88 | NC | 120 | TXRX5P_D |
| 25 | TXRX3M_A | 57 | VDDIO | 89 | LED2_0 (Note 3-1) | 121 | TXRX5M_D |
| 26 | TXRX3P_B | 58 | RXD6_7 | 90 | LED2_1 (Note 3-1) | 122 | AVDDH |
| 27 | TXRX3M_B | 59 | RXD6_6 | 91 | PME_N | 123 | GND |
| 28 | TXRX3P_C | 60 | RXD6_5 | 92 | INTRP_N | 124 | AVDDL |
| 29 | TXRX3M_C | 61 | NC | 93 | NC | 125 | XO |
| 30 | AVDDL | 62 | RXD6_4 | 94 | CLKO_25_125 | 126 | XI |
| 31 | TXRX3P_D | 63 | TX_CLK6/REFCLKI6 | 95 | RESET_N | 127 | ISET |
| 32 | TXRX3M_D | 64 | TX_EN6/TX_CTL6 | 96 | SDO | 128 | AVDDH |

Exposed Pad Must be Connected to GND

Note 3-1 This pin provides configuration strap functions during hardware/software resets. Refer to [Section 3.2.1, "Configuration Straps"](#) for additional information.

3.2 Pin Descriptions

This sections details the functions of the various device signals.

TABLE 3-2: PIN DESCRIPTIONS

| Name | Symbol | Buffer Type | Description |
|--|------------------------------|-------------|--|
| Ports 5-1 Gigabit Ethernet Pins | | | |
| Port 5-1 Ethernet TX/RX Pair A + | TXRX[5:1]P_A | AIO | Port 5-1 1000BASE-T Differential Data Pair A (+) Note: 100BASE-TX and 10BASE-Te are also supported on the A and B pairs. |
| Port 5-1 Ethernet TX/RX Pair A - | TXRX[5:1]M_A | AIO | Port 5-1 1000BASE-T Differential Data Pair A (-) Note: 100BASE-TX and 10BASE-Te are also supported on the A and B pairs. |
| Port 5-1 Ethernet TX/RX Pair B + | TXRX[5:1]P_B | AIO | Port 5-1 1000BASE-T Differential Data Pair B (+) Note: 100BASE-TX and 10BASE-Te are also supported on the A and B pairs. |
| Port 5-1 Ethernet TX/RX Pair B - | TXRX[5:1]M_B | AIO | Port 5-1 1000BASE-T Differential Data Pair B (-) Note: 100BASE-TX and 10BASE-Te are also supported on the A and B pairs. |
| Port 5-1 Ethernet TX/RX Pair C + | TXRX[5:1]P_C | AIO | Port 5-1 1000BASE-T Differential Data Pair C (+) |
| Port 5-1 Ethernet TX/RX Pair C - | TXRX[5:1]M_C | AIO | Port 5-1 1000BASE-T Differential Data Pair C (-) |
| Port 5-1 Ethernet TX/RX Pair D + | TXRX[5:1]P_D | AIO | Port 5-1 1000BASE-T Differential Data Pair D (+) |
| Port 5-1 Ethernet TX/RX Pair D - | TXRX[5:1]M_D | AIO | Port 5-1 1000BASE-T Differential Data Pair D (-) |
| Port 6 GMII/RGMII/MII/RMII Pins | | | |
| Port 6 Transmit/Reference Clock | TX_CLK6/ REFCLKI6 | I/O8 | MII Mode: TX_CLK6 is the Port 6 25/2.5MHz Transmit Clock. In PHY mode this pin is an output, in MAC mode it is an input. RMII Mode: REFCLKI6 is the Port 6 50MHz Reference Clock input when in RMII Normal mode. This pin is unused when in RMII Clock mode. GMII Mode: GTX_CLK6 is the Port 6 125MHz Transmit Clock input. RGMII Mode: TX_CLK6 is the Port 6 125/25/2.5MHz Transmit Clock input. |
| Port 6 Transmit Enable/Control | TX_EN6/ TX_CTL6 | IPD | MII/RMII/GMII Modes: TX_EN6 is the Port 6 Transmit Enable. RGMII Mode: TX_CTL6 is the Port 6 Transmit Control. |

KSZ9896C

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|-------------------------|--------|-------------|--|
| Port 6 Transmit Error | TX_ER6 | IPD | MII/GMII Mode: Port 6 Transmit Error input. RMII/RGMII Modes: Not used. Do not connect this pin in these modes of operation. |
| Port 6 Collision Detect | COL6 | IPD/O8 | MII/GMII Mode: Port 6 Collision Detect. In PHY mode this pin is an output, in MAC mode it is an input. RMII/RGMII Modes: Not used. Do not connect this pin in these modes of operation. |
| Port 6 Transmit Data 7 | TXD6_7 | IPD | GMII Mode: Port 6 Transmit Data bus bit 7. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 6 | TXD6_6 | IPD | GMII Mode: Port 6 Transmit Data bus bit 6. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 5 | TXD6_5 | IPD | GMII Mode: Port 6 Transmit Data bus bit 5. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 4 | TXD6_4 | IPD | GMII Mode: Port 6 Transmit Data bus bit 4. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 3 | TXD6_3 | IPD | MII/GMII/RGMII Modes: Port 6 Transmit Data bus bit 3. RMII Mode: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 2 | TXD6_2 | IPD | MII/GMII/RGMII Modes: Port 6 Transmit Data bus bit 2. RMII Mode: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Transmit Data 1 | TXD6_1 | IPD | MII/RMII/GMII/RGMII Modes: Port 6 Transmit Data bus bit 1. |
| Port 6 Transmit Data 0 | TXD6_0 | IPD | MII/RMII/GMII/RGMII Modes: Port 6 Transmit Data bus bit 0. |

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|---|---|-------------|--|
| Port 6 Receive/Reference Clock | RX_CLK6/ REFCLKO6 | I/O24 | <p>MII Mode: RX_CLK6 is the Port 6 25/2.5MHz Receive Clock. In PHY mode this pin is an output, in MAC mode it is an input.</p> <p>GMII/MII 2-Wire Mode: Receive Clock. GRX_CLK6 is the Port 6 125MHz Receive Clock output for GMII (1000Mbps), and the Port 6 RX_CLK6 input for MII (10/100Mbps).</p> <p>GMII/MII 3-Wire Mode: Receive Clock. GRX_CLK6 is the Port 6 125MHz Receive Clock output for GMII (1000Mbps). Refer to the RX_ER6/RX_CLK6 pin for the MII Receive Clock in this mode.</p> <p>RMII Mode: REFCLKO6 is the Port 6 50MHz Reference Clock output when in RMII Clock mode. This pin is unused when in RMII Normal mode.</p> <p>RGMII Mode: RX_CLK6 is the Port 6 125/25/2.5MHz Receive Clock output.</p> |
| Port 6 Receive Data Valid / Carrier Sense / Control | RX_DV6/ CRS_DV6/ RX_CTL6 | IPD/O24 | <p>MII/GMII Mode: RX_DV6 is the Port 6 Received Data Valid output.</p> <p>RMII Mode: CRS_DV6 is the Carrier Sense / Receive Data Valid output.</p> <p>RGMII Mode: RX_CTL6 is the Receive Control output.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 6 Receive Error | RX_ER6/RX_CLK6 | IPD/O24 | <p>MII Mode: Port 6 Receive Error output.</p> <p>GMII/MII 2-Wire Mode: RX_ER6 Port 6 Receive Error output.</p> <p>GMII/MII 3-Wire Mode: RX_CLK6 Port 6 25/2.5MHz Receive Clock input for MII (in MAC Mode).</p> <p>RMII/RGMII Modes: Not used. Do not connect this pin in these modes of operation.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 6 Carrier Sense | CRS6 | IPD/O8 | <p>MII/GMII Mode: Port 6 Carrier Sense. In PHY mode this pin is an output, in MAC mode it is an input.</p> <p>RMII/RGMII Modes: Not used. Do not connect this pin in these modes of operation.</p> |

KSZ9896C

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|---|---------|-------------|--|
| Port 6 Receive Data 7 | RXD6_7 | IPD/O24 | GMII Mode: Port 6 Receive Data bus bit 7. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Receive Data 6 | RXD6_6 | IPD/O24 | GMII Mode: Port 6 Receive Data bus bit 6. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Receive Data 5 | RXD6_5 | IPD/O24 | GMII Mode: Port 6 Receive Data bus bit 5. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Receive Data 4 | RXD6_4 | IPD/O24 | GMII Mode: Port 6 Receive Data bus bit 4. MII/RMII/RGMII Modes: Not used. Do not connect this pin in this mode of operation. |
| Port 6 Receive Data 3 | RXD6_3 | IPD/O24 | MII/GMII/RGMII Modes: Port 6 Receive Data bus bit 3. RMII Mode: Not used. Do not connect this pin in this mode of operation. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| Port 6 Receive Data 2 | RXD6_2 | IPD/O24 | MII/GMII/RGMII Modes: Port 6 Receive Data bus bit 2. RMII Mode: Not used. Do not connect this pin in this mode of operation. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| Port 6 Receive Data 1 | RXD6_1 | IPD/O24 | MII/RMII/GMIIRGMII Modes: Port 6 Receive Data bus bit 1. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| Port 6 Receive Data 0 | RXD6_0 | IPD/O24 | MII/RMII/GMIIRGMII Modes: Port 6 Receive Data bus bit 0. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| SPI/I²C/MIIM Interface Pins | | | |
| SPI/I ² C/MIIM Serial Clock | SCL/MDC | IPU | SPI/I²C Modes: SCL serial clock. MIIM Mode: MDC serial clock. |
| SPI Data Out | SDO | O8 | SPI Mode: Data out (also known as MISO). I²C/MIIM Modes: Not used. |

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|---|--------------|-------------|---|
| SPI Data In / I ² C/MIIM Data In/Out | SDI/SDA/MDIO | IPU/O8 | <p>SPI Mode: SDI Data In (also known as MOSI).</p> <p>I²C Mode: SDA Data In/Out.</p> <p>MIIM Mode: MDIO Data In/Out.</p> <p>SDI and MDIO are open-drain signals when in the output state. An external pull-up resistor to VDDIO (1.0kΩ to 4.7kΩ) is required.</p> |
| SPI Chip Select | SCS_N | IPU | <p>SPI Mode: Chip Select (active low).</p> <p>I²C/MIIM Modes: Not used.</p> |
| LED Pins | | | |
| Port 1 LED Indicator 0 | LED1_0 | IPU/O8 | <p>Port 1 LED Indicator 0. Active low output sinks current to light an external LED.</p> |
| Port 1 LED Indicator 1 | LED1_1 | IPU/O8 | <p>Port 1 LED Indicator 1. Active low output sinks current to light an external LED.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 2 LED Indicator 0 | LED2_0 | IPU/O8 | <p>Port 2 LED Indicator 0. Active low output sinks current to light an external LED.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 2 LED Indicator 1 | LED2_1 | IPU/O8 | <p>Port 2 LED Indicator 1. Active low output sinks current to light an external LED.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 3 LED Indicator 0 | LED3_0 | IPU/O8 | <p>Port 3 LED Indicator 0. Active low output sinks current to light an external LED.</p> |
| Port 3 LED Indicator 1 | LED3_1 | IPU/O8 | <p>Port 3 LED Indicator 1. Active low output sinks current to light an external LED.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |
| Port 4 LED Indicator 0 | LED4_0 | IPU/O8 | <p>Port 4 LED Indicator 0. Active low output sinks current to light an external LED.</p> <p>Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information.</p> |

KSZ9896C

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|--------------------------------------|--------------------|-------------|---|
| Port 4 LED Indicator 1 | LED4_1 | IPU/O8 | Port 4 LED Indicator 1. Active low output sinks current to light an external LED. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| Port 5 LED Indicator 0 | LED5_0 | IPU/O8 | Port 5 LED Indicator 0. Active low output sinks current to light an external LED. |
| Port 5 LED Indicator 1 | LED5_1 | IPU/O8 | Port 5 LED Indicator 1. Active low output sinks current to light an external LED. Note: This pin also provides configuration strap functions during hardware/software resets. Refer to Section 3.2.1, "Configuration Straps" for additional information. |
| Miscellaneous Pins | | | |
| Interrupt | INTRP_N | OPU | Active low, open-drain interrupt. Note: This pin requires an external pull-up resistor. |
| Power Management Event | PME_N | O8 | Power Management Event. This output signal indicates that an energy detect event has occurred. It is intended to wake up the system from a low power mode. Note: The assertion polarity is programmable (default active low). An external pull-up resistor is required for active-low operation; an external pull-down resistor is required for active-high operation. |
| System Reset | RESET_N | IPU | Active low system reset. The device must be reset either during or after power-on. An RC circuit is suggested for power-on reset. |
| Crystal Clock / Oscillator Input | XI | ICLK | Crystal clock / oscillator input. When using a 25MHz crystal, this input is connected to one lead of the crystal. When using an oscillator, this pin is the input from the oscillator. The crystal oscillator should have a tolerance of ± 50 ppm. |
| Crystal Clock Output | XO | OCLK | Crystal clock / oscillator output. When using a 25MHz crystal, this output is connected to one lead of the crystal. When using an oscillator, this pin is left unconnected. |
| 25/125MHz Reference Clock Output | CLKO_25_125 | IPU/O24 | 25/125MHz reference clock output, derived from the crystal input. |
| Transmit Output Current Set Resistor | ISET | A | Transmit output current set resistor. This pin configures the physical transmit output current. It must be connected to GND through a 6.04k Ω 1% resistor. |
| No Connect | NC | - | No Connect. For proper operation, this pin must be left unconnected. |

TABLE 3-2: PIN DESCRIPTIONS (CONTINUED)

| Name | Symbol | Buffer Type | Description |
|--------------------------|--------------|-------------|---------------------------------|
| Power/Ground Pins | | | |
| +3.3/2.5/1.8V I/O Power | VDDIO | P | +3.3V / +2.5V / +1.8V I/O Power |
| +2.5V Analog Power | AVDDH | P | +2.5V Analog Power |
| +1.2V Analog Power | AVDDL | P | +1.2V Analog Power |
| +1.2V Digital Power | DVDDL | P | +1.2V Digital Power |
| Ground | GND | GND | Ground (pins and pad) |

3.2.1 CONFIGURATION STRAPS

The KSZ9896C utilizes configuration strap pins to configure the device for different modes. While **RESET_N** is low, these pins are hi-Z. Pull-up/down resistors are used to create high or low states on these pins, which are internally sampled at the rising edge of **RESET_N**. All of these pins have a weak internal pull-up or pull-down resistor which provides a default level for strapping. To strap an LED pin low, use a 750Ω to 1kΩ external pull-down resistor. To strap a non-LED pin high, use an external 1kΩ to 10kΩ pull-up resistor to **VDDIO**. Once **RESET_N** is high, all of these pins become driven outputs.

Because the internal pull-up/down resistors are not strong, consideration must be given to any other pull-up/down resistors which may reside on the board or inside a device connected to these pins.

When an LED pin is directly driving an LED, the effect of the LED and LED load resistor on the strapping level must be considered. This is the reason for using a small value resistor to pull an LED pin low. This is especially true when an LED is powered from a voltage that is higher than **VDDIO**.

The configuration strap pins and their associated functions are detailed in [Table 3-3](#).

TABLE 3-3: CONFIGURATION STRAP DESCRIPTIONS

| Configuration Strap Pin | Description |
|-------------------------|--|
| LED1_1 | Flow Control (All Ports) 0: Flow control disabled 1: Flow control enabled (Default) |
| LED2_1 | Link-up Mode (All PHYs) 0: Fast Link-up: Auto-negotiation and auto MDI/MDI-X are disabled 1: Normal Link-up: Auto-negotiation and auto MDI/MDI-X are enabled (Default) Note: Since Fast Link-up disables auto-negotiation and auto-crossover, it is suitable only for specialized applications. |

KSZ9896C

TABLE 3-3: CONFIGURATION STRAP DESCRIPTIONS (CONTINUED)

| Configuration Strap Pin | Description |
|----------------------------|--|
| LED2_0, LED4_0 | <p>When LED2_1 = 1 at strap-in (Normal Link-up): [LED2_0, LED4_0]: Auto-Negotiation Enable (All PHYs) / NAND Tree Test Mode 00: Reserved 01: Auto-negotiation disabled 10: NAND Tree test mode 11: Auto-negotiation enabled (Default)</p> <p>When LED2_1 = 0 at strap-in (Fast Link-up; All PHYs Full-Duplex): LED2_0: 1000BASE-T Master/Slave Mode, 10/100BASE-T MDI/MDI-X Mode (All PHYs) 0: 1000BASE-T: Slave Mode 10/100BASE-T: MDI-X 1: 1000BASE-T: Master Mode (Default) 10/100BASE-T: MDI (Default)</p> <p>LED4_0: PHY Speed Select (All PHYs) 0: 1000BASE-T 1: 100BASE-TX (Default)</p> |
| LED4_1, LED3_1 | <p>[LED4_1, LED3_1]: Management Interface Mode 00: MIIM (MDIO) 01: I²C 1x: SPI (Default)</p> |
| LED5_1 | <p>Switch Enable at Startup 0: Start Switch is disabled. The switch will not forward packets until the Start Switch bit is set in the Switch Operation Register. 1: Start Switch is enabled. The switch will forward packets immediately after reset. (Default)</p> |
| RXD6_3, RXD6_2 | <p>[RXD6_3, RXD6_2]: Port 6 Mode 00: RGMII (Default) 01: RMII 10: GMII/MII 11: MII</p> |
| RXD6_1 | <p>Port 6 MII/RMII/GMII Mode 0: MII: PHY Mode (Default) RMII: Clock Mode. RMII 50MHz reference clock is output on REFCLKO6. (Default) GMII: PHY Mode RGMII: No effect 1: MII: MAC Mode RMII: Normal Mode. RMII 50MHz reference clock is input on REFCLKI6. GMII: MAC Mode RGMII: No effect</p> |
| RXD6_0 | <p>Port 6 Speed Select 0: 1000Mbps Mode (Default) 1: 10/100Mbps Mode</p> <p>Note: If Port 6 is configured for MII or RMII, set the speed to 100Mbps.</p> |
| RX_DV6/CRS_DV6/ RX_CTL6 | <p>In-Band Management 0: Disable In-Band Management (Default) 1: Enable In-Band Management</p> |

TABLE 3-3: CONFIGURATION STRAP DESCRIPTIONS (CONTINUED)

| Configuration Strap Pin | Description |
|-------------------------|--|
| RX_ER6/RX_CLK6 | <p>GMII/MII Clock Mode 0: 2-Wire Clock Mode (Default) 1: 3-Wire Clock Mode</p> <p>These modes affect the GMII and MII clocks for outgoing data on the Port 6 RX signals. The strapped value is readable on bit 0 of the Port 6 XMII Port Control 0 Register and may be overwritten.</p> <p>In 2-Wire mode, pin 72 (RX_CLK6/REFCLKO6) is used for either the GMII output clock or the MII clock.</p> <p>In 3-Wire mode, pin 72 (RX_CLK6/REFCLKO6) is used for the GMII output clock, and pin 75 (RX_ER6/RX_CLK6) is used for the MII input clock when in MAC mode. 3-Wire mode is intended as an option, when in MAC mode, for connection to a 1000/100/10Mbps PHY which has separate pins for the GMII (1000Mbps) GTX_CLK clock and the MII (100/10Mbps) TX_-CLK clock.</p> <p>When the MAC interface does not need to be switchable between GMII and MII, or when in PHY mode, do not select 3-Wire mode.</p> |

KSZ9896C

4.0 FUNCTIONAL DESCRIPTION

This section provides functional descriptions for the following:

- [Physical Layer Transceiver \(PHY\)](#)
- [LEDs](#)
- [Media Access Controller \(MAC\)](#)
- [Switch](#)
- [NAND Tree Support](#)
- [Clocking](#)
- [Power](#)
- [Power Management](#)
- [Management Interface](#)
- [In-Band Management](#)
- [MAC Interface \(Port 6\)](#)

4.1 Physical Layer Transceiver (PHY)

Ports 1 through 5 include completely integrated triple-speed (10BASE-T_e, 100BASE-TX, 1000BASE-T) Ethernet physical layer transceivers for transmission and reception of data over standard four-pair unshielded twisted pair (UTP), CAT-5 or better Ethernet cable.

The device reduces board cost and simplifies board layout by using on-chip termination resistors for the differential pairs, eliminating the need for external termination resistors. The internal chip termination and biasing provides significant power savings when compared with using external biasing and termination resistors.

The device can automatically detect and correct for differential pair misplacements and polarity reversals, and correct for propagation delay differences between the four differential pairs, as specified in the IEEE 802.3 standard for 1000BASE-T operation.

4.1.1 1000BASE-T TRANSCEIVER

The 1000BASE-T transceiver is based on a mixed-signal/digital signal processing (DSP) architecture, which includes the analog front-end, digital channel equalizers, trellis encoders/decoders, echo cancelers, cross-talk cancelers, a precision clock recovery scheme, and power-efficient line drivers.

4.1.1.1 Analog Echo Cancellation Circuit

In 1000BASE-T mode, the analog echo cancellation circuit helps to reduce the near-end echo. This analog hybrid circuit relieves the burden of the ADC and the adaptive equalizer. This circuit is disabled in 10BASE-T_e/100BASE-TX mode.

4.1.1.2 Automatic Gain Control (AGC)

In 1000BASE-T mode, the automatic gain control circuit provides initial gain adjustment to boost up the signal level. This pre-conditioning circuit is used to improve the signal-to-noise ratio of the receive signal.

4.1.1.3 Analog-to-Digital Converter (ADC)

In 1000BASE-T mode, the analog-to-digital converter digitizes the incoming signal. ADC performance is essential to the overall performance of the transceiver. This circuit is disabled in 10BASE-T_e/100BASE-TX mode.

4.1.1.4 Timing Recovery Circuit

In 1000BASE-T mode, the mixed signal clock recovery circuit, together with the digital phase locked loop (PLL), is used to recover and track the incoming timing information from the received data. The digital PLL has very low long-term jitter to maximize the signal-to-noise ratio of the receive signal.

The 1000BASE-T slave PHY must transmit the exact receive clock frequency recovered from the received data back to the 1000BASE-T master PHY. Otherwise, the master and slave will not be synchronized after long transmission. This also helps to facilitate echo cancellation and NEXT removal.

4.1.1.5 Adaptive Equalizer

In 1000BASE-T mode, the adaptive equalizer provides the following functions:

- Detection for partial response signaling
- Removal of NEXT and ECHO noise
- Channel equalization

Signal quality is degraded by residual echo that is not removed by the analog hybrid because of impedance mismatch. The device uses a digital echo canceler to further reduce echo components on the receive signal.

In 1000BASE-T mode, data transmission and reception occurs simultaneously on all four pairs of wires (four channels). This results in high-frequency cross-talk coming from adjacent wires. The device uses three NEXT cancelers on each receive channel to minimize the cross-talk induced by the other three channels.

In 10BASE-Te/100BASE-TX mode, the adaptive equalizer needs only to remove the inter-symbol interference and recover the channel loss from the incoming data.

4.1.1.6 Trellis Encoder and Decoder

In 1000BASE-T mode, the transmitted 8-bit data is scrambled into 9-bit symbols and further encoded into 4D-PAM5 symbols. On the receiving side, the idle stream is examined first. The scrambler seed, pair skew, pair order and polarity must be resolved through the logic. The incoming 4D-PAM5 data is then converted into 9-bit symbols and de-scrambled into 8-bit data.

4.1.2 100BASE-TX TRANSCEIVER

4.1.2.1 100BASE-TX Transmit

The 100BASE-TX transmit function performs parallel-to-serial conversion, 4B/5B coding, scrambling, NRZ-to-NRZI conversion, and MLT3 encoding and transmission.

The circuitry starts with a parallel-to-serial conversion, which converts the MII data from the MAC into a 125MHz serial bit stream. The data and control stream is then converted into 4B/5B coding, followed by a scrambler. The serialized data is further converted from NRZ-to-NRZI format, and then transmitted in MLT3 current output. An external ISET resistor sets the output current for the 1:1 transformer ratio.

The output signal has a typical rise/fall time of 4ns and complies with the ANSI TP-PMD standard regarding amplitude balance, overshoot, and timing jitter. The wave-shaped 10BASE-Te output driver is also incorporated into the 100BASE-TX driver.

4.1.2.2 100BASE-TX Receive

The 100BASE-TX receiver function performs adaptive equalization, DC restoration, MLT3-to-NRZI conversion, data and clock recovery, NRZI-to-NRZ conversion, de-scrambling, 4B/5B decoding, and serial-to-parallel conversion.

The receiving side starts with the equalization filter to compensate for inter-symbol interference (ISI) over the twisted pair cable. Since the amplitude loss and phase distortion is a function of the cable length, the equalizer has to adjust its characteristics to optimize performance. In this design, the variable equalizer makes an initial estimation based on comparisons of incoming signal strength against some known cable characteristics, and then tunes itself for optimization. This is an ongoing process and self-adjusts against environmental changes such as temperature variations.

Next, the equalized signal goes through a DC restoration and data conversion block. The DC restoration circuit is used to compensate for the effect of baseline wander and to improve the dynamic range. The differential data conversion circuit converts the MLT3 format back to NRZI. The slicing threshold is also adaptive.

The clock recovery circuit extracts the 125MHz clock from the edges of the NRZI signal. This recovered clock is then used to convert the NRZI signal into the NRZ format. This signal is sent through the de-scrambler followed by the 4B/5B decoder. Finally, the NRZ serial data is converted to an MII format and provided as the input data to the MAC.

4.1.2.3 Scrambler/De-Scrambler

The purpose of the scrambler is to spread the power spectrum of the signal to reduce electromagnetic interference (EMI) and baseline wander. The scrambler is used only for 100BASE-TX.

Transmitted data is scrambled through the use of an 11-bit wide linear feedback shift register (LFSR). The scrambler generates a 2047-bit non-repetitive sequence. Then the receiver de-scrambles the incoming data stream using the same sequence as at the transmitter.

KSZ9896C

4.1.3 10BASE-Te TRANSCEIVER

10BASE-Te is an energy-efficient version of 10BASE-T which is powered from a 2.5V supply. It has a reduced transmit signal amplitude and requires Cat5 cable. It is inter-operable to 100m with 10BASE-T when Cat5 cable is used.

4.1.3.1 10BASE-Te Transmit

The 10BASE-Te driver is incorporated with the 100BASE-TX driver to allow for transmission using the same magnetics. They are internally wave-shaped and pre-emphasized into outputs with typical 1.75V amplitude (compared to the typical transmit amplitude of 2.5V for 10BASE-T). The harmonic contents are at least 27dB below the fundamental frequency when driven by an all-ones Manchester-encoded signal.

4.1.3.2 10BASE-Te Receive

On the receive side, input buffers and level detecting squelch circuits are employed. A differential input receiver circuit and a phase-locked loop (PLL) perform the decoding function.

The Manchester-encoded data stream is separated into clock signal and NRZ data. A squelch circuit rejects signals with levels less than 400mV or with short pulse widths to prevent noise at the RXP1 or RXM1 input from falsely triggering the decoder. When the input exceeds the squelch limit, the PLL locks onto the incoming signal and the device decodes a data frame. The receiver clock is maintained active during idle periods in between data reception.

4.1.4 AUTO MDI/MDI-X

The automatic MDI/MDI-X feature, also known as auto crossover, eliminates the need to determine whether to use a straight cable or a crossover cable between the device and its link partner. The auto-sense function detects the MDI/MDI-X pair mapping from the link partner, and assigns the MDI/MDI-X pair mapping of the device accordingly. [Table 4-1](#) shows the device's 10/100 Mbps pin configuration assignments for MDI and MDI-X pin mapping.

TABLE 4-1: MDI/MDI-X PIN DEFINITIONS

| Pin (RJ45 pair) | MDI | | | MDI-X | | |
|------------------|------------|------------|-----------|------------|------------|-----------|
| | 1000BASE-T | 100BASE-TX | 10BASE-Te | 1000BASE-T | 100BASE-TX | 10BASE-Te |
| TXRXxP/M_A (1,2) | A+/- | TX+/- | TX+/- | B+/- | RX+/- | RX+/- |
| TXRXxP/M_B (3,6) | B+/- | RX+/- | RX+/- | A+/- | TX+/- | TX+/- |
| TXRXxP/M_C (4,5) | C+/- | Not used | Not used | D+/- | Not used | Not used |
| TXRXxP/M_D (7,8) | D+/- | Not used | Not used | C+/- | Not used | Not used |

Auto MDI/MDI-X is enabled by default. It can be disabled through the port control registers. If Auto MDI/MDI-X is disabled, the port control register can also be used to select between MDI and MDI-X settings.

An isolation transformer with symmetrical transmit and receive data paths is recommended to support Auto MDI/MDI-X.

4.1.5 PAIR-SWAP, ALIGNMENT, AND POLARITY CHECK

In 1000BASE-T mode, the device:

- Detects incorrect channel order and automatically restores the pair order for the A, B, C, D pairs (four channels).
- Supports 50±10ns difference in propagation delay between pairs of channels in accordance with the IEEE 802.3 standard, and automatically corrects the data skew so the corrected four pairs of data symbols are synchronized.

Incorrect pair polarities of the differential signals are automatically corrected for all speeds.

4.1.6 WAVE SHAPING, SLEW-RATE CONTROL, AND PARTIAL RESPONSE

In communication systems, signal transmission encoding methods are used to provide the noise-shaping feature and to minimize distortion and error in the transmission channel.

- For 1000BASE-T, a special partial-response signaling method is used to provide the bandwidth-limiting feature for the transmission path.
- For 100BASE-TX, a simple slew-rate control method is used to minimize EMI.
- For 10BASE-Te, pre-emphasis is used to extend the signal quality through the cable.

4.1.7 AUTO-NEGOTIATION

The device conforms to the auto-negotiation protocol as described by IEEE 802.3. Auto-negotiation allows each port to operate at either 10BASE-Te 100BASE-TX or 1000BASE-T by allowing link partners to select the best common mode of operation. During auto-negotiation, the link partners advertise capabilities across the link to each other and then compare their own capabilities with those they received from their link partners. The highest speed and duplex setting that is common to the two link partners is selected as the mode of operation.


The following list shows the speed and duplex operation mode from highest to lowest priority.

- Priority 1: 1000BASE-T, full-duplex
- Priority 2: 1000BASE-T, half-duplex
- Priority 3: 100BASE-TX, full-duplex
- Priority 4: 100BASE-TX, half-duplex
- Priority 5: 10BASE-Te, full-duplex
- Priority 6: 10BASE-Te, half-duplex

If the KSZ9896C link partner doesn't support auto-negotiation or is forced to bypass auto-negotiation for 10BASE-Te and 100BASE-TX modes, the KSZ9896C port sets its operating mode by observing the signal at its receiver. This is known as parallel detection, and allows the KSZ9896C to establish a link by listening for a fixed signal protocol in the absence of the auto-negotiation advertisement protocol.

The auto-negotiation link-up process is shown in [Figure 4-1](#).

FIGURE 4-1: AUTO-NEGOTIATION AND PARALLEL OPERATION



For 1000BASE-T mode, auto-negotiation is always required to establish a link. During 1000BASE-T auto-negotiation, the master and slave configuration is first resolved between link partners. Then the link is established with the highest common capabilities between link partners.

KSZ9896C

Auto-negotiation is enabled by default after power-up or hardware reset. Afterwards, auto-negotiation can be enabled or disabled via bit 12 of the [PHY Basic Control Register](#). If auto-negotiation is disabled, the speed is set by bits 6 and 13 of the [PHY Basic Control Register](#), and the duplex is set by bit 8.

If the speed is changed on the fly, the link goes down and either auto-negotiation or parallel detection initiate until a common speed between the KSZ9896C and its link partner is re-established for a link.

If link is already established and there is no change of speed on the fly, the changes (for example, duplex and pause capabilities) will not take effect unless either auto-negotiation is restarted through bit 9 of the [PHY Basic Control Register](#), or a link-down to link-up transition occurs (i.e. disconnecting and reconnecting the cable).

After auto-negotiation is completed, the link status is updated in the [PHY Basic Status Register](#), and the link partner capabilities are updated in the [PHY Auto-Negotiation Link Partner Ability Register](#), [PHY Auto-Negotiation Expansion Status Register](#), and [PHY 100BASE-T Status Register](#).

4.1.8 FAST LINK-UP

Link up time is normally determined by the time it takes to complete auto-negotiation. Additional time may be added by the auto MDI/MDI-X feature. The total link up time from power-up or cable connect is typically a second or more.

Fast Link-up mode significantly reduces 100BASE-TX link-up time by disabling both auto-negotiation and auto MDI/MDI-X, and fixing the TX and RX channels. This mode is enabled or disabled by the **LED2_1** strapping option. It is not set by registers, so fast link-up is available immediately upon power-up. Fast Link-up is available at power-up only for 100BASE-TX link speed, which is selected by strapping the **LED4_0** pin high. Fast Link-up is also available for 10BASE-T, but this link speed must first be selected via a register write.

Fast Link-up is intended for specialized applications where both link partners are known in advance. The link must also be known so that the fixed transmit channel of one device connects to the fixed receive channel of the other device, and vice versa. The TX and RX channel assignments are determined by the MDI/MDI-X strapping option on **LED2_0**.

If a device in Fast Link-up mode is connected to a normal device (auto-negotiate and auto-MDI/MDI-X), there will be no problems linking, but the speed advantage of Fast Link-up will not be realized.

For more information on configuration straps, refer to [Section 3.2.1, "Configuration Straps," on page 17](#).

4.1.9 LinkMD® CABLE DIAGNOSTICS

The LinkMD® function utilizes Time Domain Reflectometry (TDR) to analyze the cabling for common cabling problems, such as open circuits, short circuits and impedance mismatches.

LinkMD® works by sending a pulse of known amplitude and duration down the MDI or MDI-X pair, and then analyzing the shape of the reflected signal to determine the type of fault. The time duration for the reflected signal to return provides the approximate distance to the cabling fault. The LinkMD® function processes this TDR information and presents it as a numerical value that can be translated to a cable distance.


4.1.10 REMOTE PHY LOOPBACK

This loopback mode checks the line (differential pairs, transformer, RJ-45 connector, Ethernet cable) transmit and receive data paths between the KSZ9896C and its Ethernet PHY link partner, and is supported for 10/100/1000 Mbps at full-duplex.

The loopback data path is shown in [Figure 4-2](#) and functions as follows:

- The Ethernet PHY link partner transmits data to the KSZ9896C PHY port.
- Data received at the external pins of the PHY port is looped back without passing through the MAC and internal switch fabric.
- The same KSZ9896C PHY port transmits data back to the Ethernet PHY link partner.

FIGURE 4-2: REMOTE PHY LOOPBACK



The following programming steps and register settings are for remote PHY loopback mode for 1000BASE-T Master Mode, 1000BASE-T Slave Mode, 100BASE-TX Mode, and 10BASE-T Mode.

- 1000BASE-T Master Mode
 - Set Port N (1-5), [PHY 1000BASE-T Control Register](#) = 0x1F00
 - Set Port N (1-5), [PHY Remote Loopback Register](#) = 0x01F0
 - Set Port N (1-5), [PHY Basic Control Register](#) = 0x1340

- 1000BASE-T Slave Mode
 - Set Port N (1-5), [PHY 1000BASE-T Control Register](#) = 0x1300
 - Set Port N (1-5), [PHY Remote Loopback Register](#) = 0x01F0
 - Set Port N (1-5), [PHY Basic Control Register](#) = 0x1340

- 100BASE-TX Mode
 - Set Port N (1-5), [PHY Auto-Negotiation Advertisement Register](#) = 0x0181
 - Set Port N (1-5), [PHY 1000BASE-T Control Register](#) = 0x0C00
 - Set Port N (1-5), [PHY Remote Loopback Register](#) = 0x01F0
 - Set Port N (1-5), [PHY Basic Control Register](#) = 0x3300

- 10BASE-T Mode
 - Set Port N (1-5), [PHY Auto-Negotiation Advertisement Register](#) = 0x0061
 - Set Port N (1-5), [PHY 1000BASE-T Control Register](#) = 0x0C00
 - Set Port N (1-5), [PHY Remote Loopback Register](#) = 0x01F0
 - Set Port N (1-5), [PHY Basic Control Register](#) = 0x3300

KSZ9896C

4.2 LEDs

Each PHY port has two programmable LED output pins, **LEDx_0** and **LEDx_1**, to indicate the PHY link and activity status. Two different LED modes are available. The LED mode can be changed individually for each PHY port by writing to the PHY Mode bit in the PHY indirect register: MMD 2, address 0, bit 4:

- 1 = [Single-LED Mode](#)
- 0 = [Tri-Color Dual-LED Mode](#) (Default)

Each LED output pin can directly drive an LED with a series resistor (typically 220Ω to 470Ω). LED outputs are active-low.

4.2.1 SINGLE-LED MODE

In single-LED mode, the **LEDx_1** pin indicates the link status while the **LEDx_0** pin indicates the activity status, as shown in [Figure 4-2](#).

TABLE 4-2: SINGLE-LED MODE PIN DEFINITION

| LED Pin | Pin State | Pin LED Definition | Link/Activity |
|---------------|-----------|--------------------|---------------------|
| LEDx_1 | H | OFF | Link Off |
| | L | ON | Link On (any speed) |
| LEDx_0 | H | OFF | No Activity |
| | Toggle | Blinking | Activity (RX,TX) |

4.2.2 TRI-COLOR DUAL-LED MODE

In tri-color dual-LED mode, the link and activity status are indicated by the **LEDx_1** pin for 1000BASE-T; by the **LEDx_0** pin for 100BASE-TX; and by both **LEDx_1** and **LEDx_0** pins, working in conjunction, for 10BASE-T. This behavior is summarized in [Figure 4-3](#).

TABLE 4-3: TRI-COLOR DUAL-LED MODE PIN DEFINITION

| LED Pin (State) | | LED Pin (Definition) | | Link/Activity |
|-----------------|---------------|----------------------|---------------|----------------------------------|
| LEDx_1 | LEDx_0 | LEDx_1 | LEDx_0 | |
| H | H | OFF | OFF | Link off |
| L | H | ON | OFF | 1000Mbps Link / No Activity |
| Toggle | H | Blinking | OFF | 1000Mbps Link / Activity (RX,TX) |
| H | L | OFF | ON | 100Mbps Link / No Activity |
| H | Toggle | OFF | Blinking | 100Mbps Link / Activity (RX,TX) |
| L | L | ON | ON | 10Mbps Link / No Activity |
| Toggle | Toggle | Blinking | Blinking | 10Mbps Link / Activity (RX,TX) |

4.3 Media Access Controller (MAC)

4.3.1 MAC OPERATION

The device strictly abides by IEEE 802.3 standards to maximize compatibility. Additionally, there is an added MAC filtering function to filter unicast packets. The MAC filtering function is useful in applications, such as VoIP, where restricting certain packets reduces congestion and thus improves performance.

The transmit MAC takes data from the egress buffer and creates full Ethernet frames by adding the preamble and the start-of-frame delimiter ahead of the data, and generates the FCS that is appended to the end of the frame. It also sends flow control packets as needed.

The receive MAC accepts data via the integrated PHY or via the GMII/MII/RMII/RGMII interface. It decodes the data bytes, strips off the preamble and SFD of each frame. The destination and source addresses and VLAN tag are extracted for use in filtering and address/ID lookup, and the MAC also calculates the CRC of the received frame, which is compared to the FCS field. The MAC can discard frames that are the wrong size, that have an FCS error, or when the source MAC address matches the Switch MAC address.

The receive MAC also implements the Wake on LAN (WoL) feature. This system power saving feature is described in detail in the [Section 4.8, "Power Management"](#).

MIB statistics are collected in both receive and transmit directions.

4.3.2 INTER-PACKET GAP (IPG)

If a frame is successfully transmitted, then the minimum 96-bit time for IPG is specified as being between two consecutive packets. If the current packet is experiencing collisions, the minimum 96-bit time for IPG is specified as being from carrier sense (CRS) to the next transmit packet.

4.3.3 BACK-OFF ALGORITHM

The device implements the IEEE standard 802.3 binary exponential back-off algorithm in half-duplex mode. After 16 collisions, the packet is dropped.

4.3.4 LATE COLLISION

If a transmit packet experiences collisions after 512 bit times of the transmission, the packet is dropped.

4.3.5 LEGAL PACKET SIZE

On all ports, the device discards received packets smaller than 64 bytes (excluding VLAN tag, including FCS) or larger than the maximum size. The default maximum size is the IEEE standard of 1518 bytes, but can be programmed to 2000 bytes. Ports operating at 1000Mbps may be programmed to accept jumbo packets up to 9000 bytes, but for performance reasons it is recommended that no more than two ports be enabled simultaneously for jumbo packets.

4.3.6 FLOW CONTROL

The device supports standard MAC Control PAUSE (802.3x flow control) frames in both the transmit and receive directions for full-duplex connections.

In the receive direction, if a PAUSE control frame is received on any port, the device will not transmit the next normal frame on that port until the timer, specified in the PAUSE control frame, expires. If another PAUSE frame is received before the current timer expires, the timer will then update with the new value in the second PAUSE frame. During this period (while it is flow controlled), only flow control packets from the device are transmitted.

In the transmit direction, the device has intelligent and efficient ways to determine when to invoke flow control and send PAUSE frames. The flow control is based on availability of the system resources, including available buffers, available transmit queues and available receive queues.

The device issues a PAUSE frame containing the maximum pause time defined in IEEE standard 802.3x. Once the resource is freed up, the device sends out another flow control frame with zero pause time to turn off the flow control (turn on transmission to the port). A hysteresis feature is provided to prevent the flow control mechanism from being constantly activated and deactivated.

4.3.7 HALF-DUPLEX BACK PRESSURE

A half-duplex back pressure option (non-IEEE 802.3 standard) is also provided. The activation and deactivation conditions are the same as in full-duplex mode. If back pressure is required, the device sends preambles to defer the other stations' transmission (carrier sense deference).

To avoid jabber and excessive deference (as defined in the 802.3 standard), after a certain time, the device discontinues the carrier sense and then raises it again quickly. This short silent time (no carrier sense) prevents other stations from sending out packets thus keeping other stations in a carrier sense deferred state. If the port has packets to send during a back pressure situation, the carrier sense type back pressure is interrupted and those packets are transmitted instead. If there are no additional packets to send, carrier sense type back pressure is reactivated again until chip resources free up. If a collision occurs, the binary exponential back-off algorithm is skipped and carrier sense is generated immediately, thus reducing the chance of further collision and carrier sense is maintained to prevent packet reception.

KSZ9896C

To ensure no packet loss in 10BASE-T_e or 100BASE-TX half-duplex modes, the user must enable the following:

- No excessive collision drop ([Switch MAC Control 1 Register](#))
- Back pressure ([Port MAC Control 1 Register](#))

4.3.8 FLOW CONTROL AND BACK PRESSURE REGISTERS

[Table 4-4](#) provides a list of flow control and back pressure related registers.

TABLE 4-4: FLOW CONTROL AND BACK PRESSURE REGISTERS

| Registers | Description |
|---|--|
| LED Configuration Strap Register | LED configuration strap settings. (LED1_1 enables flow control and back pressure) |
| Switch MAC Address 0 Register through Switch MAC Address 5 Register | Switch's MAC address, used as source address of PAUSE control frames |
| Switch MAC Control 0 Register | "Aggressive back-off" enable |
| Switch MAC Control 1 Register | BP mode, "Fair mode" enable, "no excessive collision drop" enable |
| Switch MAC Control 4 Register | Pass PAUSE control frames |
| Port Status Register | Flow control enable (per port) |
| PHY Auto-Negotiation Advertisement Register | PHY - flow control advertisement (per port) |
| Port MAC Control 1 Register | Half-duplex back pressure enable (per port) |
| Port Ingress Rate Limit Control Register | Ingress rate limit flow control enable (per port) |
| Port Control 0 Register | Drop mode (per port) |

4.3.9 BROADCAST STORM PROTECTION

The device has an intelligent option to protect the switch system from receiving too many broadcast packets. As the broadcast packets are forwarded to all ports except the source port, an excessive number of switch resources (bandwidth and available space in transmit queues) may be utilized. The device has the option to include "multicast packets" for storm control. The broadcast storm rate parameters are programmed globally, and can be enabled or disabled on a per port basis. The rate is based on a 5ms interval for 1000BASE-T, a 50ms interval for 100BASE-TX and a 500ms interval for 10BASE-T_e. At the beginning of each interval, the counter is cleared to zero and the rate limit mechanism starts to count the number of bytes during the interval. The rate definition is described in control registers. The default setting equates to a rate of 1%.

4.3.10 SELF-ADDRESS FILTERING

Received packets can be filtered (dropped) if their source address matches the device's MAC address. This feature is useful for automatically terminating packets once they have traversed a ring network and returned to their source. It can be enabled on a per-port basis via the [Switch Lookup Engine Control 1 Register](#) and [Port Control 2 Register](#).

4.4 Switch

4.4.1 SWITCHING ENGINE

A high-performance switching engine is used to move data to and from the MAC's packet buffers. It operates in store and forward mode, while an efficient switching mechanism reduces overall latency. The switching engine has a 256KByte internal frame buffer that is shared between all the ports.

For the majority of switch functions, all of the data ports are treated equally. However, a few functions such as IGMP snooping, 802.1X, forwarding invalid VLAN packets, etc., give special recognition to the host port. Any port (but most commonly port 6) may be assigned as the host port by enabling tail tagging mode for that port. Only one port may be a host port.

When a switch receives a non-error packet, it checks the packet's destination MAC address. If the address is known, the packet is forwarded to the output port that is associated with the destination MAC address. The following paragraphs describe the key functions of destination address lookup and source address learning. These processes may be combined with VLAN support and other features, which are described in the subsequent sub-sections.

4.4.2 ADDRESS LOOKUP

Destination address lookup is performed in three separate internal address tables in the device:

1. [Address Lookup \(ALU\) Table](#): 4K dynamic + static entries
2. [Static Address Table](#): 16 static entries
3. [Reserved Multicast Address Table](#): 8 pre-configured static entries

4.4.2.1 Address Lookup (ALU) Table

The Address Lookup (ALU) Table stores MAC addresses and their associated information. This table holds both dynamic and static entries. Dynamic entries are created automatically in hardware, as described in [Section 4.4.2.4, "Learning"](#). Static entries are created by management software.

This table is a 4-way associative memory, with 1K buckets, for a total of 4K entries. A hash function translates the received packet's MAC address (and optionally the FID) into a 10-bit index for accessing the table. At each bucket are four fully-associative address entries. All four entries are simultaneously compared to the MAC address (plus optional FID) for a possible match.

Three options are available for the hashing function, as described in [Table 4-5](#). If VLAN is enabled (802.1Q VLAN Enable bit in the [Switch Lookup Engine Control 0 Register](#)), the VLAN group (FID) is included in the hashing function along with the MAC address. If VLAN is not enabled the hashing function is applied to MAC address and the FID in the default VLAN (VID=1) which is 0.

TABLE 4-5: ADDRESS LOOKUP TABLE HASHING OPTIONS

| HASH_OPTION (Switch Lookup Engine Control 0 Register) | Description |
|--|---|
| 01b (Default) | A hash algorithm based on the CRC of the MAC address plus FID. The hash algorithm uses the CRC-CCITT polynomial. The input to the hash is reduced to a 16-bit CRC hash value. Bits [9:0] of the hash value plus (binary addition) 7-bit FID (zero extended on the left) are used as an index to the table. The CRC-CCITT polynomial is: $X^{16}+X^{12}+X^5+1$. |
| 10b | An XOR algorithm based on 16 bits of the XOR of the triple-folded MAC address. Bits [9:0] of the XOR value plus 7-bit FID (left-extended) are used to index the table. |
| 00b or 11b | A direct algorithm. The 10 least significant bits of the MAC address plus 7 bit FID are used to index the table. |

4.4.2.2 Static Address Table

The 16-entry Static Address Table is typically used to hold multicast addresses, but is not limited to this. As with static entries in the ALU table, entries in the Static Address Table are created by management software. It serves the same function as static entries that are created in the ALU table, so its use is optional.

4.4.2.3 Reserved Multicast Address Table

The Reserved Multicast Address Table holds 8 pre-configured address entries, as defined in [Table 4-6](#). This table is an optional feature that is disabled at power-on. If desired, the forwarding ports may be modified.

KSZ9896C

TABLE 4-6: RESERVED MULTICAST ADDRESS TABLE

| Group | Address | MAC Group Address Function | Default PORT FORWARD Value (defines forwarding port: P6...P1) | Default Forwarding Action |
|-------|---|---|--|--|
| 0 | (01-80-C2-00)-00-00 | Bridge Group Data | 00_0000 | Forward only to the highest numbered port (default host port) |
| 1 | (01-80-C2-00)-00-01 | MAC Control Frame (typically flow control) | 00_0000 | Drop MAC flow control |
| 2 | (01-80-C2-00)-00-03 | 802.1X Access Control | 00_0000 | Forward to highest numbered port |
| 3 | (01-80-C2-00)-00-10 | Bridge Management | 11_1111 | Flood to all ports |
| 4 | (01-80-C2-00)-00-20 | GMRP | 11_1111 | Flood to all ports except highest numbered port |
| 5 | (01-80-C2-00)-00-21 | GVRP | 11_1111 | Flood to all ports except highest numbered port |
| 6 | (01-80-C2-00)-00-02, (01-80-C2-00)-00-04 – (01-80-C2-00)-00-0F | | 00_0000 | Forward to highest numbered port |
| 7 | (01-80-C2-00)-00-11 - (01-80-C2-00)-00-1F, (01-80-C2-00)-00-22 - (01-80-C2-00)-00-2F | | 11_1111 | Flood to all ports except highest numbered port |

If a match is found in one of the tables, then the destination port is read from that table entry. If a match is found in more than one table, static entries will take priority over dynamic entries.

4.4.2.4 Learning

The internal lookup engine updates the ALU table with a new dynamic entry if the following conditions are met:

- The received packet's source address (SA) does not exist in the lookup table.
- The received packet has no errors, and the packet size is of legal length.
- The received packet has a unicast SA.
- If VLAN is enabled, the received packet must belong to the indicated VLAN domain (FID).

The lookup engine inserts the qualified SA into the table, along with the port number and time stamp. If all four table entries are valid, the oldest of the (up to four) dynamic entries may be deleted to make room for the new entry. Static entries are never deleted by the learning process. If all four entries are static entries, the address is not learned but an interrupt is generated and the table index number is made available to the interrupt service routine.

4.4.2.5 Migration

The internal lookup engine also monitors whether a station has moved. If a station has moved, it updates the ALU table accordingly. Migration happens when the following conditions are met:

- The received packet's SA is in the table but the associated source port information is different.
- The received packet has no receiving errors, and the packet size is of legal length.

The lookup engine updates the existing record in the table with the new source port information.

4.4.2.6 Aging

The lookup engine updates the time stamp information of a dynamic record in the ALU table whenever the corresponding SA appears. The time stamp is used in the aging process. If a record is not updated for a period of time, the lookup engine removes the record from the table. The lookup engine constantly performs the aging process and continuously removes aging records. The aging period is about 300 seconds (± 75 seconds) and can be configured longer or shorter (1 second to 30 minutes). This feature can be enabled or disabled. Static entries are exempt from the aging process.

4.4.2.7 Forwarding


The device forwards packets using the algorithm that is depicted in Figure 4-3. Figure 4-3 shows stage one of the forwarding algorithm where the search engine looks up the VLAN ID, static table, and dynamic table for the destination address, and comes up with "port to forward 1" (PTF1). PTF1 is then further modified by spanning tree, IGMP snooping, port mirroring, and port VLAN processes.

The ACL process works in parallel with the flow outlined above. The authentication and ACL processes have the highest priority in the forwarding process, and the ACL result may override the result of the above flow. The output of the ACL process is the final "port-to-forward 2" (PTF2) destination port(s).

The device will not forward the following packets:

- Error packets: These include framing errors, frame check sequence (FCS) errors, alignment errors, and illegal size packet errors.
- MAC Control PAUSE frames: The device intercepts these packets and performs full duplex flow control accordingly.
- "Local" packets: Based on destination address (DA) lookup. If the destination port from the lookup table matches the port from which the packet originated, the packet is defined as "local".
- In-Band Management packets.

FIGURE 4-3: PACKET FORWARDING PROCESS FLOWCHART



KSZ9896C

4.4.2.8 Lookup Engine Registers

Table 4-7 provides a list of lookup engine related registers.

TABLE 4-7: LOOKUP ENGINE REGISTERS

| Registers | Description |
|--|--------------------------------|
| Global Interrupt Status Register , Global Interrupt Mask Register | Top level LUE interrupt |
| Switch Lookup Engine Control 0 Register , Switch Lookup Engine Control 1 Register , Switch Lookup Engine Control 2 Register , Switch Lookup Engine Control 3 Register | Misc. |
| Address Lookup Table Interrupt Register , Address Lookup Table Mask Register | Low level LUE interrupts |
| Address Lookup Table Entry Index 0 Register , Address Lookup Table Entry Index 1 Register | Access failure address/index |
| ALU Table Index 0 Register , ALU Table Index 1 Register , ALU Table Access Control Register , Static Address and Reserved Multicast Table Control Register , ALU / Static Address Table Entry 1 Register , ALU / Static Address / Reserved Multicast Table Entry 2 Register , ALU / Static Address Table Entry 3 Register , ALU / Static Address Table Entry 4 Register | Address table access registers |

4.4.3 IEEE 802.1Q VLAN

Virtual LAN is a means of segregating a physical network into multiple virtual networks whereby traffic may be confined to specific subsets of the greater network. IEEE 802.1Q defines a VLAN protocol using a 4-byte tag that is added to the Ethernet frame header. The device supports port-based and tag-based VLANs, including tagging, un-tagging, forwarding and filtering.

4.4.3.1 Non-Tag Port-Based VLAN

The simplest VLAN method establishes forwarding restrictions on a port-by-port basis without using VLAN tags. There is a register for each ingress port that is used to specify the allowed forwarding ports. An incoming packet is restricted from being forwarded to any egress port that is disallowed for that ingress port. The settings are made in the [Port Control 1 Register](#). This function is always enabled; it is not enabled and disabled by the 802.1Q VLAN Enable bit in the [Switch Lookup Engine Control 0 Register](#). The default setting is to allow all ingress-to-egress port paths.

4.4.3.2 Tag-Based VLAN

When 802.1Q VLAN is enabled, an internal VLAN Table with 4k entries is used to store port membership list, VLAN group ID (FID) and additional information relating to each VLAN. This table must be set up by an administrator prior to enabling 802.1Q VLAN. Enabling is done by setting the 802.1Q VLAN Enable bit in the [Switch Lookup Engine Control 0 Register](#).

In 802.1Q VLAN mode, the lookup process starts with VLAN Table lookup, using the tag's VID as the address. The first step is to determine whether the VID is valid. If the VID is not valid, the packet is dropped and its address is not learned. Alternatively, unknown VID packets may be forwarded to pre-defined ports or to the host port. If the VID is valid, the FID is retrieved for further lookup. The FID + Destination Address (hashed(DA) + FID) are used to determine the destination port. The FID + Source Address (hashed(SA) + FID) are used for address learning (see [Table 4-9](#) and [Table 4-10](#)).

The hashed(DA) + FID are hashed and used for forwarding lookup in the Address Lookup and Static Address Tables. For a successful address table lookup, the FID fields must also match. If the match fails, the packet is broadcast to all the VLAN port members defined in the VLAN Table entry. If there is a match and egress VLAN filtering is enabled, the packet is forwarded to those ports that are in both the address table port forwarding list and the VLAN table port membership list.

A similar address table lookup is performed using the hashed(SA) + FID. If the lookup fails, the FID and SA are learned.

If a non-tagged or null-VID-tagged packet is received, the ingress port default VID ([Port Default Tag 0 Register](#) and [Port Default Tag 1 Register](#)) is used for lookup.

[Table 4-8](#) details the forwarding and discarding actions that are taken for the various VLAN scenarios. The first entry in the table is explained by the fact that VLAN Table lookup is enabled even when 802.1Q VLAN is not enabled. Notice that in the [Port Default Tag 0 Register](#) and [Port Default Tag 1 Register](#), the port default VID is 1 for each port. Correspondingly, the VLAN port membership list in the VLAN Table entry for VID=1 is pre-configured at power-on to all ones. This provides the standard Ethernet switch behavior of broadcasting all packets with unknown destination address. If the VLAN table entry # 1 is changed, or if the port default VID is changed, this may affect the forwarding action for “unknown packets” even when VLAN is not enabled.

It should also be noted that the default values of the Egress VLAN Filtering bits are zero. These bits are zero only for backwards compatibility with previous “KSZ” switches. The resulting switch behavior, in the event of a successful VLAN and ALU lookups, is to forward the packet to the ports in the address table port forwarding list, without regard to the VLAN port membership list. It is suggested that the Egress VLAN Filtering bits be set to one so that the VLAN port membership list from the VLAN Table will be used to qualify the forwarding determined from the address lookup.

TABLE 4-8: VLAN FORWARDING

| VLAN Enable (Note 4-1) | VLAN Match/ Valid (Note 4-2) | Forward Option (Note 4-3) | Egress VLAN Filtering (Note 4-4) | Unknown VID Forward (Note 4-5) | Drop Invalid VID (Note 4-6) | ALU Match/ Valid (Note 4-7) | Action |
|---|--|--|---|---|--|---|--|
| 0 | X | X | X | X | X | No | Forward to port membership list of default VID in LAN table |
| 0 | X | X | X | X | X | Yes | Forward to Address Lookup port forwarding list |
| 1 | No | X | X | 0 | 0 | X | Forward to host port |
| 1 | No | X | X | 0 (def) | 1 (def) | X | Discard |
| 1 | No | X | X | 1 | X | X | Forward to Unknown VID packet forward port list |
| 1 | Yes | 0 | X | X | X | No | Broadcast: Forward to VLAN table port membership list (PORT FORWARD) Multicast: Forward to Unknown Multicast ports if UM is enabled. Else, forward to VLAN table port membership list. Unicast: Forward to Unknown Unicast ports if UU is enabled. Else forward to VLAN table port membership list. |
| 1 | Yes | 0 | 0 (def) | X | X | Yes | Forward to address table lookup port forwarding list |
| 1 | Yes | 0 | 1 | X | X | Yes | Forward to address table lookup port forwarding list & VLAN table port membership list (bitwise AND) |
| 1 | Yes | 1 | X | X | X | Yes | Forward to VLAN table port membership list |

KSZ9896C

Note: “(def)” indicates the default power-up value.

Note 4-1 **VLAN Enable** is bit 7 in the [Switch Lookup Engine Control 0 Register](#)

Note 4-2 **VLAN Match/Valid** indicates when the VLAN Table entry is valid

Note 4-3 **Forward Option** is a bit in the [VLAN Table Entry 0 Register](#)

Note 4-4 **Egress VLAN Filtering** are bits 5 and 4 in the [Switch Lookup Engine Control 2 Register](#)

Note 4-5 **Unknown VID Forwarding** is in the [Unknown VLAN ID Control Register](#)

Note 4-6 **Drop Invalid VID** is bit 6 in the [Switch Lookup Engine Control 0 Register](#)

Note 4-7 **ALU Match/Valid** indicates when the Address Lookup is a success

[Table 4-9](#) describes in more detail the address lookup process that follows the VLAN Table lookup. Lookup occurs in both the Address Lookup Table and the Static Address Table simultaneously, and the resulting action depends on the results of the two lookups.

TABLE 4-9: HASHED(DA) + FID LOOKUP IN VLAN MODE

| DA Found in Static MAC Table? | Use FID Flag? (Static MAC Table) | FID Match? | DA+FID Found in ALU Table? | Action |
|-------------------------------|----------------------------------|------------|----------------------------|---|
| No | Don't Care | Don't Care | No | Lookup has failed. Broadcast to the membership ports defined in the VLAN Table |
| No | Don't Care | Don't Care | Yes | Send to the destination port defined in the Address Lookup (ALU) Table |
| Yes | 0 | Don't Care | Don't Care | Send to the destination port(s) defined in the Static Address Table |
| Yes | 1 | No | No | Lookup has failed. Broadcast to the membership ports defined in the VLAN Table. |
| Yes | 1 | No | Yes | Send to the destination port defined in the Address Lookup (ALU) Table |
| Yes | 1 | Yes | Don't Care | Send to the destination port(s) defined in the Static Address Table |

A source address (SA) lookup is also performed in the Address Lookup Table. SA lookup also performs SA filtering and MAC priority when the address is hit. [Table 4-10](#) describes how learning is performed in the Address Lookup Table when a successful VLAN table lookup has been done and the no matching static entry is found in the Address Lookup Table or the Static Address Table.

TABLE 4-10: HASHED(SA) + FID LOOKUP IN VLAN MODE

| FID + SA Found in Address Lookup (ALU) Table? | Action |
|---|--|
| No | Learn and add FID + SA to the Address Lookup (ALU) Table |
| Yes | Update time stamp |

4.4.3.2.1 Tag Insertion and Removal

Tag insertion is enabled on all ports when the VLAN feature is enabled. At the ingress port, untagged packets are tagged with the ingress port's default tag. The default tag is separately programmable for each port. The switch does not add tags to already tagged packets unless double tagging is enabled.

At the egress port, tagged packets will have their 802.1Q VLAN tags removed if un-tagging is enabled in the VLAN table entry. This feature is controlled on a per-port basis. Untagged packets will not be modified if 802.1Q is enabled.

4.4.3.2.2 Double Tagging

The switch supports double tagging, also known as Q-in-Q or VLAN stacking. This feature can be used for service providers to append a second VLAN tag in addition to a first VLAN tag applied by the customer. VLAN support can be enabled either with or without double tagging. When double tagging is enabled, the outer tag is recognized and is used for VLAN and address lookup instead of the inner tag. The outer tag precedes the inner tag in the frame header: the outer tag is located immediately after the source address, and contains a different Tag Protocol Identifier (TPID) value than the inner tag.

Additional controls are available for full control of the VLAN function. Some of these features can be enabled on a per-port basis, while others are global:

- **Ingress VLAN Filtering:** Discard packet if VID port membership in VLAN table does not include the ingress port.
- **Discard non PVID Packet:** Discard packet if VID does not match the ingress port default VID.
- **Discard un-tagged Packet:** Discard any received packet without a tag.
- **Drop tag:** Drops the packet if it is VLAN tagged.
- **Unknown VID Forward:** Forward to a fixed set of ports if VLAN lookup fails.
- **Drop unknown VID:** Additional options for unknown VID packets: discard or forward to the host port.
- **Null VID Replacement:** Replace a null VID with the ingress port default VID.
- **PVID Replacement:** Replace a non-null VID with the ingress port default VID.
- **Double Tag Mcast Trap:** In double tag mode, trap all reserved multicast packets and forward to the host port.

4.4.3.3 VLAN Registers

Table 4-11 provides a list of VLAN related registers.

TABLE 4-11: VLAN REGISTERS

| Registers | Description |
|--|---|
| Switch Operation Register | Double tag enable |
| Switch Lookup Engine Control 0 Register | VLAN enable; Drop invalid VID frames |
| Switch Lookup Engine Control 2 Register | Trap double tagged MC frames; Dynamic & status egress VLAN filtering |
| Unknown VLAN ID Control Register | Forward unknown VID |
| Switch MAC Control 2 Register | Null VID replacement with PVID at egress |
| VLAN Table Entry 0 Register , VLAN Table Entry 1 Register , VLAN Table Entry 2 Register , VLAN Table Index Register , VLAN Table Access Control Register | Read/write access to the VLAN table |
| Port Default Tag 0 Register , Port Default Tag 1 Register | Port default tag |
| Port Ingress MAC Control Register | Drop non-VLAN frames; Tag drop |
| Port Transmit Queue PVID Register | PVID replacement at egress |
| Port Control 2 Register | VLAN table lookup for VID=0; Ingress VLAN filtering; PVID mismatch discard |

KSZ9896C

4.4.4 QUALITY-OF-SERVICE (QOS) PRIORITY SUPPORT

The device provides quality-of-service (QoS) for applications such as VoIP. There are multiple methods for assigning priority to ingress packets. Depending on the packet prioritization method, the packet priority levels are mapped to the egress queues for each port. Each port can be configured for 1, 2, and 4 egress queues, which are prioritized. The default is 1 queue per port.

When configured for 4 priority queues, Queue 3 is the highest priority queue and Queue 0 is the lowest priority. Likewise, for a 2-queue configuration, Queue 1 is the highest priority queue. If a port is not configured as 2 or 4 queues, then high priority and low priority packets have equal priority in the single transmit queue.

There is an additional option for every port to select either to always deliver packets from the highest priority queue first, or use weighted round robin queuing amongst the multiple queues. This is described later in [Section 4.4.13, "Scheduling and Rate Limiting"](#).

4.4.4.1 Port-Based Priority


With port-based priority, each ingress port is individually classified as a specific priority level. All packets received at the high-priority receiving port are marked as high priority and are sent to the high-priority transmit queue if the corresponding transmit queue is split into 2 or 4 queues.

4.4.4.2 IEEE 802.1p-Based Priority

For IEEE 802.1p-based priority, the device examines the ingress packets to determine whether they are tagged. If tagged, the 3-bit PCP priority field in the VLAN tag is retrieved and used to look up the "priority mapping" value. The "priority mapping" value is programmable.

[Figure 4-4](#) illustrates how the 802.1p priority field is embedded in the 802.1Q VLAN tag.

FIGURE 4-4: 802.P PRIORITY FIELD FORMAT



4.4.4.3 IEEE 802.1p Priority Field Re-Mapping

This is a QoS feature that allows the device to set the "User Priority Ceiling" at any ingress port. If the ingress packet's priority field has a higher priority value than the default tag's priority field of the ingress port, the packet's priority field is replaced with the default tag's priority field.

4.4.4.4 DiffServ (DSCP) Priority (IP)

DiffServ-based priority from the DSCP field in the IP header can be used to determine packet priority. The 6-bit DSCP value is used as an index to a set of registers which translate the 6-bit DSCP value to a 2-bit value that specifies one of the 4 (or 2) queues. These registers are fully programmable.

4.4.4.5 ACL Priority

The Access Control List (ACL) Filtering feature can also be used to assign priority to received packets. This is discussed in [Section 4.4.16, "Access Control List \(ACL\) Filtering"](#).

4.4.5 TRAFFIC CONDITIONING & POLICING

4.4.5.1 Two Rate Three Color Marker

The Two Rate Three Color Marker meters an IP packet stream and marks its packets green, yellow, or red. A packet is marked red if it exceeds the Peak Information Rate (PIR). Otherwise, it is marked either yellow or green depending on whether it exceeds or doesn't exceed the Committed Information Rate (CIR).

The Meter operates in one of two modes. In the Color-Blind mode, the Meter assumes that the packet stream is uncolored. In the Color-Aware mode, the Meter assumes that some preceding entity has pre-colored the incoming packet stream so that each packet is green, yellow, or red. The Marker (re)colors an IP packet according to the results of the Meter.

4.4.5.2 Weighted Random Early Detection (WRED)

The WRED feature monitors the average queue size of packet memory and ingress queue size of each traffic class, and drops packets based on memory and queue utilization. If the buffers are almost empty, all incoming traffic is accepted. As the buffer utilization increases, the probability for dropping an incoming packet also increases.

WRED is intended to avoid the problem of global synchronization. Global synchronization can occur when a switch becomes congested and begins dropping incoming packets all at once. For TCP streams, packet drops invoke the TCP congestion control mechanism, which reduce the transmission rate until there are no more packet drops. If there are many TCP streams and their congestion control mechanisms act in unison, this can cause an undesirable oscillation in traffic rates. By selectively dropping some packets early rather than waiting until the buffer is full, WRED avoids dropping large numbers of packets at once and minimizes the chances of global synchronization.

The packet drop probability is based on the minimum threshold, maximum threshold, and a probability multiplier. When the average queue depth is above the minimum threshold, packets start getting dropped. The rate of packet drop increases linearly as the average queue size increases until the average queue size reaches the maximum threshold. The probability multiplier is the fraction of packets dropped when the average queue depth is at the maximum threshold. When the average queue size is above the maximum threshold, all packets are dropped.

4.4.6 SPANNING TREE SUPPORT

To support spanning tree, one port is the designated port for the host processor, which is defined as the port for which tail tagging is enabled. Each of the other ports can be configured in one of the five spanning tree states via “transmit enable”, “receive enable” and “learning disable” register bits. [Table 4-12](#) shows the setting and software actions taken for each of the five spanning tree states.

TABLE 4-12: SPANNING TREE STATES

| Disable State | Port Setting | Software Action |
|---|---|---|
| The port should not forward or receive any packets. Learning is disabled. | transmit enable = 0 receive enable = 0 learning disable = 1 | The processor should not send any packets to the port. The switch may still send specific packets to the processor (packets that match some entries in the “Static MAC Table” with “overriding bit” set) and the processor should discard those packets. Address learning is disabled on the port in this state. |
| Blocking State | Port Setting | Software Action |
| Only packets to the processor are forwarded. Learning is disabled. | transmit enable = 0 receive enable = 0 learning disable = 1 | The processor should not send any packets to the port(s) in this state. The processor should program the “Static MAC Table” with the entries that it needs to receive (for example, BPDU packets). The “overriding” bit should also be set so that the switch will forward those specific packets to the processor. Address learning is disabled on the port in this state. |

KSZ9896C

TABLE 4-12: SPANNING TREE STATES (CONTINUED)

| Listening State | Port Setting | Software Action |
|---|---|---|
| Only packets to and from the processor are forwarded. Learning is disabled. | transmit enable = 0 receive enable = 0 learning disable = 1 | The processor should program the “Static MAC Table” with the entries that it needs to receive (for example, BPDU packets). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state. Address learning is disabled on the port in this state. |
| Learning State | Port Setting | Software Action |
| Only packets to and from the processor are forwarded. Learning is enabled. | transmit enable = 0 receive enable = 0 learning disable = 0 | The processor should program the “Static MAC Table” with the entries that it needs to receive (for example, BPDU packets). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state. Address learning is enabled on the port in this state. |
| Forwarding State | Port Setting | Software Action |
| Packets are forwarded and received normally. Learning is enabled. | transmit enable = 1 receive enable = 1 learning disable = 0 | The processor programs the “Static MAC Table” with the entries that it needs to receive (for example, BPDU packets). The “overriding” bit is set so that the switch forwards those specific packets to the processor. The processor can send packets to the port(s) in this state. Address learning is enabled on the port in this state. |

4.4.7 RAPID SPANNING TREE SUPPORT

There are three operational states assigned to each port for the Rapid Spanning Tree Protocol (RSTP):

1. [Discarding State](#)
2. [Learning State](#)
3. [Forwarding State](#)

4.4.7.1 Discarding State

Discarding ports do not participate in the active topology and do not learn MAC addresses.

- Discarding state: the state includes three states of the disable, blocking and listening of STP.
- Port setting: transmit enable = “0”, receive enable = “0”, learning disable = “1”.
- Software action: The host processor should not send any packets to the port. The switch may still send specific packets to the processor (packets that match some entries in the static table with “overriding bit” set) and the processor should discard those packets. When the port’s learning capability (learning disable = ‘1’) is disabled, port related entries in the ALU table and static MAC table can be rapidly flushed.

4.4.7.2 Learning State

Ports in “learning state” learn MAC addresses, but do not forward user traffic.

- Learning State: Only packets to and from the host processor are forwarded. Learning is enabled.
- Port setting for Learning State: transmit enable = “0”, receive enable = “0”, learning disable = “0”.
- Software action: The processor should program the Static Address Table with the entries that it needs to receive (e.g., BPDU packets). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state (see [Section 4.4.9, “Tail Tagging Mode”](#) for details). Address learning is enabled on the port in this state.

4.4.7.3 Forwarding State

Ports in “forwarding states” fully participate in both data forwarding and MAC learning.

- Forwarding state: Packets are forwarded and received normally. Learning is enabled.
- Port setting: transmit enable = “1”, receive enable = “1”, learning disable = “0”.

- Software action: The host processor should program the Static Address Table with the entries that it needs to receive (e.g., BPDUs). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state (see [Section 4.4.9, "Tail Tagging Mode"](#) for details). Address learning is enabled on the port in this state.

RSTP uses only one type of BPDU called RSTP BPDUs. They are similar to STP configuration BPDUs with the exception of a type field set to “version 2” for RSTP and “version 0” for STP, and a flag field carrying additional information.

4.4.8 MULTIPLE SPANNING TREE SUPPORT

Multiple Spanning Tree Protocol (MSTP) is an extension of RSTP that allows different VLANs to have different spanning tree configurations. The VLAN Table, Address Lookup Table, Static Address Table and Reserved Multicast Address Table all contain a 3-bit field which can be used to specify one of eight spanning trees. Each port contains state registers for specifying unique states for each of the spanning trees.


4.4.9 TAIL TAGGING MODE

Tail tagging is a method to communicate ingress and egress port information between the host processor and the switch. It is useful for spanning tree protocol, IGMP/MLD snooping, and other applications.

When the switch forwards a frame to the host port, two tail tagging bytes are added to the frame by the switch to indicate to the host processor the port that the frame was received on. In the other direction, the host processor adds the tail tagging bytes to a frame to indicate the intended egress destination port to the switch. When multiple priority queues are enabled, the tail tag is also used to indicate the priority queue. The tail tagging bytes are removed before the frame egresses the switch. Tail tagging may be enabled on any one port, and this defines the host port. tail tagging must not be enabled on multiple ports.

Tail tagging is implemented by inserting two additional bytes at the end of the packet, between the data field and the 4-byte CRC / FCS, as shown in [Figure 4-5](#).

FIGURE 4-5: TAIL TAG FRAME FORMAT



4.4.10 IGMP SUPPORT

For Internet Group Management Protocol (IGMP) support in Layer 2, the device provides two components:

- “[IGMP](#)” Snooping
- “[Multicast Address Insertion](#)” in the Static MAC Table

4.4.10.1 “IGMP” Snooping

The device traps IGMP packets and forwards them only to the processor (host port). The IGMP packets are identified as IP packets (either Ethernet IP packets, or IEEE 802.3 SNAP IP packets) with IP version = 0x4 and protocol version number = 0x2.

Note: The port for which [Tail Tagging Mode](#) is enabled is the host port.

4.4.10.2 “Multicast Address Insertion” in the Static MAC Table

Once the multicast address is programmed in the Static Address Table or Address Lookup Table, the multicast session is trimmed to the subscribed ports, instead of broadcasting to all ports.

So that the host processor knows which port the IGMP packet was received on, [Tail Tagging Mode](#) must be enabled.

4.4.11 IPV6 MLD SNOOPING

The device traps IPv6 Multicast Listener Discovery (MLD) packets and forwards them only to the processor (host port).

KSZ9896C

4.4.12 PORT MIRRORING

the device supports “port mirroring” comprehensively as:

- “Receive Only” Mirror-on-a-Port
- “Transmit Only” Mirror-on-a-Port
- “Receive and Transmit” Mirror-on-a-Port

4.4.12.1 “Receive Only” Mirror-on-a-Port

All the packets received on the port are mirrored on the sniffer port. For example, 1 is programmed to be “receive sniff” and the host port is programmed to be the “sniffer”. A packet received on port 1 is destined to port 2 after the internal lookup. The packet is forwarded to both port 2 and the host port. The device can optionally even forward “bad” received packets to the “sniffer port”.

4.4.12.2 “Transmit Only” Mirror-on-a-Port

All the packets transmitted on the port are mirrored on the sniffer port. For example, port 1 is programmed to be “transmit sniff” and the host port is programmed to be the “sniffer port”. A packet received on port 2 is destined to port 1 after the internal lookup. The device forwards the packet to both port 1 and the host port.

4.4.12.3 “Receive and Transmit” Mirror-on-a-Port

All the packets received on port A and transmitted on port B are mirrored on the sniffer port. For example, port 1 is programmed to be “receive sniff”, port 2 is programmed to be “transmit sniff”, and the host port is programmed to be the “sniffer port”. A packet received on port 1 is destined to port 2 after the internal lookup. The device forwards the packet to both port 2 and the host port.

Multiple ports can be selected as “receive sniff” or “transmit sniff”. In addition, any port can be selected as the “sniffer port”.

4.4.13 SCHEDULING AND RATE LIMITING

Each device port has two egress packet scheduling options, which can be applied when the port is configured for two or four queues. Additionally, each port has ingress and egress rate limiter features.

4.4.13.1 Strict Priority Scheduling

When an egress port is configured as two or four queues, and strict priority scheduling is selected, each queue will take absolute priority over all lower priority queues. If a packet is available to transmit from queue 3 (the highest priority queue), then it will take priority for transmission over any packet that will also be available in any of the other queues. A packet in queue 2 will be transmitted only if no packet is available in queue 3. Weighted round robin is an alternative to strict priority scheduling.

4.4.13.2 Weighted Round Robin (WRR) Scheduling

WRR scheduling is an alternative to strict priority scheduling for egress queues. It is referred to as fair queuing because it gives proportionally higher priority to the highest priority queue, but not absolute priority.

4.4.13.3 Rate Limiting

The device supports independent ingress and egress hardware rate limiting on each port. Normally these two features are considered mutually exclusive, and users are discouraged from using both on the same port.

For 10BASE-T_e, a rate setting above 10Mbps means the rate is not limited. Likewise, for 100BASE-TX, a rate setting above 100Mbps means the rate is not limited. On the receive side, the data receive rate for each priority at each port can be limited by setting up ingress rate control registers. On the transmit side, the data transmit rate for each priority queue at each port can be limited by setting up egress rate control registers. The size of each frame has options to include minimum inter-frame gap (IFG) or preamble byte, in addition to the data field (from packet DA to FCS).

For ingress rate limiting, the device provides options to selectively choose frames from all types, multicast, broadcast, and flooded unicast frames. The data rate from those selected type of frames is counted. Packets are dropped at the ingress port when the data rate exceeds the specified rate limit.

For egress rate limiting, the leaky bucket algorithm is applied to each output priority queue for shaping output traffic. Inter-frame gap is stretched on a per frame base to generate smooth, non-burst egress traffic. The throughput of each output priority queue is limited by the egress rate specified.

If any egress queue receives more traffic than the specified egress rate throughput, packets may be accumulated in the output queue and packet memory. After the memory of the queue or the port is used up, packet dropping or flow control will be triggered. As a result of congestion, the actual egress rate may be dominated by flow control/dropping at the ingress end, and may be therefore slightly less than the specified egress rate.

To reduce congestion, it is a good practice to ensure that the egress bandwidth exceeds the ingress bandwidth.

4.4.14 INGRESS MAC ADDRESS FILTERING FUNCTION

When a packet is received, the destination MAC address is looked up in both the static and dynamic MAC address tables. If the address is not found in either of these tables, then the destination MAC address is “unknown”. By default, an unknown packet is forwarded to all ports except the port at which it was received. An optional feature makes it possible to specify the port or ports to which to forward unknown packets. It is also possible to specify no ports, meaning that unknown packets will be discarded. This feature is implemented separately for unknown unicast, unknown multicast and unknown VID packets.

4.4.15 802.1X ACCESS CONTROL

IEEE 802.1X is a Port-based authentication protocol. EAPOL is the protocol normally used by the authentication process as uncontrolled Port. By receiving and extracting special EAPOL frames, the host processor can control whether the ingress and egress ports should forward packets or not. If a user port wants service from another port (authenticator), it must get approved by the authenticator. The device detects EAPOL frames by checking the destination address of the frame. The destination addresses should be either a multicast address as defined in IEEE 802.1x (01-80-C2-00-00-03) or an address used in the programmable reserved multicast address domain with offset -00-03. Once EAPOL frames are detected, the frames are forwarded to the host port so it can send the frames to the authenticator server. Eventually, the CPU determines whether the requester is qualified or not based on its source MAC address, and frames are either accepted or dropped.

When the device is configured as an authenticator, the ports of the switch must then be configured for authorization. In an authenticator-initiated port authorization, a client is powered up or plugs into the port, and the authenticator port sends an Extensible Authentication Protocol (EAP) PDU to the supplicant requesting the identification of the supplicant. At this point in the process, the port on the switch is connected from a physical standpoint; however, the 802.1X process has not authorized the port and no frames are passed from the port on the supplicant into the switching fabric. If the supplicant attached to the switch (KSZ9896C) did not understand the EAP PDU that it was receiving from the switch, it would not be able to send an ID and the port would remain unauthorized. In this state, the port would be blocked from passing any user traffic. If the supplicant is running the 802.1X EAP, it would respond to the request with its configured ID. (This could be a user name/password combination or a certificate.)

After the device receives the ID from the supplicant, it passes the ID information to an authentication server (RADIUS server) that can verify the identification information. The RADIUS server responds to the switch with either a success or failure message. If the response is a success, the port will be authorized and user traffic will be allowed to pass through the port like any switch port connected to an access device. If the response is a failure, the port will remain unauthorized and, therefore, unused. If there is no response from the server, the port will also remain unauthorized and will not pass any traffic.

Port control can be performed via the [Access Control List \(ACL\) Filtering](#) feature.

4.4.16 ACCESS CONTROL LIST (ACL) FILTERING

An Access Control List (ACL) can be created for each port to perform filtering on incoming layer 2 MAC, layer 3 IP or layer 4 TCP/UDP packets. Multicast filtering is handled in the Static Address Table and the Reserved Multicast Address Table, but the ACL provides additional capabilities for filtering routed network protocols. As shown in [Figure 4-3](#), ACL filtering may take precedence over other forwarding functions.

The ACL allows the switch to filter ingress traffic based on the following header fields:

- Source or destination MAC address and/or EtherType
- Source or destination IPv4 address with programmable mask
- IPv4 protocol
- Source or destination UDP port
- Source or destination TCP port
- TCP Flag with programmable mask

KSZ9896C

The ACL is implemented as an ordered list of up to 16 access control rules which are programmed into the ACL Table. Each entry specifies certain rules (a set of matching conditions and action rules) to control the forwarding and priority of packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the permissions required to be forwarded, based on the conditions specified in the lists. Multiple match conditions can be either AND'ed or OR'ed together.

The ACL can also implement a count function that generates an interrupt rather than a forwarding action. The counter can be either a watchdog timer or an event counter. As a watchdog timer, an interrupt is generated if a packet with a specific MAC address and EtherType is not received within a specified time interval. As an event counter, an interrupt is generated once a specified number of packets with a specific MAC address and EtherType have been received.

The ACL consists of three parts: matching rules, action rules, and processing entries. A matching rule specifies what comparison test shall be performed on the incoming packet. It can also enable a counter function. An action rule specifies the forwarding action to be taken if the matching test succeeds. Alternatively, when a count function is enabled in a matching rule, the 11-bit count value is stored in the corresponding action rule field and there is no forwarding action.

In general, the 16 matching rules are not directly linked to the 16 action rules. For example, matching entry #0 is not necessarily related to action entry #0. The exception is when the counter function is enabled in a matching rule, whereby the matching rule and action rule fields at the same ACL table entry will function together and are no longer independent.

Each of the 16 processing entries is used to link any number of matching rules (specified in RuleSet) to any one action rule (specified in FRN). When there are multiple matching rules in a RuleSet, those rules are AND'ed together. Only if all of those matching results are true will the FRN action be taken.

It is also possible to configure the ACL table so that multiple processing entries specify the same action rule. In this way, the final matching result is the OR of the matching results from each of the multiple RuleSets.

The 16 ACL rules represent an ordered list, with entry #0 having the highest priority and entry #15 having the lowest priority. All matching rules are evaluated. If there are multiple true match results and multiple corresponding actions, the highest priority (lowest numbered) of those actions will be the one taken.


4.4.16.1 Processing Entry Description

The Processing Entry consists of two parameters as described in [Table 4-13](#).

TABLE 4-13: ACL PROCESSING ENTRY PARAMETERS

| Parameter | Description |
|---------------|---|
| FRN[3:0] | First Rule Number Pointer to an Action rule entry. Possible values are 0 to 15. If all Matching rules specified in the RuleSet are evaluated true, then this is the resulting Action rule. |
| RuleSet[15:0] | Specifies a set of one or more Matching rule entries. RuleSet has one bit for each of the 16 Matching rule entries. If multiple Matching rules are selected, then all conditions will be AND'ed to produce a final match result. 0 = Matching rule not selected 1 = Matching rule selected |

FIGURE 4-6: ACL STRUCTURE AND EXAMPLE RULE VALUES



The examples in Figure 4-6 are interpreted as follows:

- Rule #0: Test the matching rule entry #0. If true, apply action rule entry #0.
- Rule #1: Test the matching rule entry #1. If true, apply action rule entry #0.
- Rule #2: Test the matching rule entry #2. If true, apply action rule entry #1.
- Matching rule entry #3 is configured for the counter function. Action entry #3 is used to hold the corresponding count value.
- Rule #4: Test the matching rule entries #4, 5 and 6. If all are true, apply action rule entry #4.
- Rule #6: Test the matching rule entry #6. If true, apply action rule entry #5.
- Rule #8: Test the matching rule entries #7 and 11. If both are true, apply action rule entry #9.

No more than one action can be taken for any packet. If the matching conditions are true for multiple RuleSets, then the corresponding FRN field with the lowest value (highest priority) determines the action to be taken.

Note that processing entries #0 and 1 produce an OR function: action #0 is taken if RuleSet #0 or RuleSet #1 is true.

Notice that processing entries #4 and 6 have overlapping RuleSets, but different FRNs. This can be summarized as:

If match #4, 5 and 6 are all true, then apply action #4,
 Else if match #6 is true, then apply action #5.

KSZ9896C

Table 4-14 summarizes the available matching options. The MD and ENB fields are used to select the desired matching option. More configuration details are given in the following section.

TABLE 4-14: MATCHING RULE OPTIONS

| MD[1:0] | ENB[1:0] | Matching Rule |
|--|----------|--|
| 00 | XX | Matching rule disabled |
| 01 (Layer 2 matching: MAC address, EtherType) | 00 | Action field is used as count value for packets matching MAC address and EtherType |
| | 01 | Compare MAC address only |
| | 10 | Compare EtherType only |
| | 11 | Compare both MAC address and EtherType |
| 10 (Layer 3 matching: IP address) | 00 | Reserved |
| | 01 | Compare IPv4 source and destination address (with mask) |
| | 10 | Compare both source and destination IPv4 addresses (without mask) |
| | 11 | Reserved |
| 11 (Layer 4 matching: TCP, UDP, IP protocol) | 00 | Compare IPv4 protocol |
| | 01 | Compare TCP source port or destination port |
| | 10 | Compare UDP source port or destination port |
| | 11 | Compare TCP sequence number |

4.4.16.2 Matching Rule Description

The Matching Rule consists of several parameters. The first two parameters, MD[1:0] and ENB[1:0], determine the organization of the remainder of each Matching Rule.

When MD = 00, the Matching Rule is disabled.

TABLE 4-15: ACL MATCHING RULE PARAMETERS FOR MD = 01

| Parameter | Description |
|-------------------|--|
| MD[1:0] | MODE 00 = Matching rule is disabled 01 = Layer 2 MAC header or counter filtering 10 = Layer 3 IP header filtering 11 = Layer 4 TCP header (and IP protocol) filtering |
| ENB[1:0] | 00 = Count Mode. Both the MAC Address and TYPE are tested. A count value (either time or packet count) is also incorporated. Details are given below this table. 01 = Comparison is performed only on the MAC Address value 10 = Comparison is performed only on the TYPE value 11 = Both the MAC Address and TYPE are tested |
| S/D | Source / Destination 0 = Destination address 1 = Source address |
| EQ | Equal / Not Equal 0 = Not Equal produces true result 1 = Equal produces true result |
| MAC ADDRESS[47:0] | 48-bit MAC address |
| TYPE[15:0] | EtherType |

Details for MD = 01, ENB = 00:

The 11 bits of the aggregated bit fields from PM, P, RPE, RP and MM in the Action rule entry specify a count value for packets matching MAC Address and TYPE in the Matching Field.

The count unit is determined by the TU bit (located in the Action rule).

- When 0, the unit is microsecond.
- When 1, the unit is millisecond.

The CA bit (located in the Action rule) determines the algorithm used to generate an interrupt when the count terminates.

- When 0, an 11-bit counter will be loaded with the count value from the list and start counting down every unit time. An interrupt will be generated when the timer expires, i.e. the next qualified packet has not been received within the period specified by the value.
- When 1, the counter is incremented with every matched packet received. An interrupt is generated when the terminal count is reached. The count resets thereafter. Time units are not used in this mode.

TABLE 4-16: ACL MATCHING RULE PARAMETERS FOR MD = 10

| Parameter | Description |
|------------------|---|
| MD[1:0] | MODE 00 = Matching rule is disabled 01 = Layer 2 MAC header or counter filtering 10 = Layer 3 IP header filtering 11 = Layer 4 TCP header (and IP protocol) filtering |
| ENB[1:0] | 00 = Reserved 01 = IPv4 source or destination address (with mask) 10 = IPv4 source and destination address (without mask) 11 = Reserved |
| S/D | Source / Destination 0 = Destination address 1 = Source address |
| EQ | Equal / Not Equal 0 = Not Equal produces true result 1 = Equal produces true result |
| IP ADDRESS[31:0] | IPv4 address Source or destination address (determined by S/D) when ENB = 01, Source address when ENB = 10 |
| IP MASK[31:0] | Mask bits for the IPv4 address when ENB = 01: 0 = This bit of the address is compared 1 = This bit of the address is not compared Destination IPv4 address when ENB = 10 |

TABLE 4-17: ACL MATCHING RULE PARAMETERS FOR MD = 11

| Parameter | Description |
|-----------|--|
| MD[1:0] | MODE 00 = Matching rule is disabled 01 = Layer 2 MAC header or counter filtering 10 = Layer 3 IP header filtering 11 = Layer 4 TCP header (and IP protocol) filtering |
| ENB[1:0] | 00 = IP Protocol comparison is enabled 01 = TCP source/destination port comparison is enabled 10 = UDP source/destination port comparison is enabled 11 = TCP sequence number is compared |
| S/D | Source / Destination 0 = Destination address 1 = Source address |
| EQ | Equal / Not Equal 0 = Not Equal produces true result 1 = Equal produces true result |

KSZ9896C

TABLE 4-17: ACL MATCHING RULE PARAMETERS FOR MD = 11 (CONTINUED)

| Parameter | Description |
|----------------------------------|---|
| MAX PORT[15:0] MIN PORT[15:0] | Max and Min Ports for TCP/UDP or TCP Sequence Number[31:0] |
| PC[1:0] | Port Comparison 00 = Port comparison is disabled 01 = Port matches either one of MAX or MIN 10 = Match if port number is in the range of MIN to MIN 11 = Match if port number is out of the range |
| PRO[7:0] | IPv4 protocol to be matched |
| FME | TCP Flag Match Enable 0 = TCP FLAG matching disabled 1 = TCP FLAG matching enabled |
| FMASK[7:0] | TCP FLAG Mask 0 = This bit of the Flag field is compared 1 = This bit of the Flag field is not compared |
| FLAG[7:0] | TCP Flag to be matched |

4.4.16.3 Action Rule Description

TABLE 4-18: ACL ACTION RULE PARAMETERS FOR NON-COUNT MODES (MD ≠ 01 OR ENB ≠ 00)


| Parameter | Description |
|----------------|--|
| PM[1:0] | Priority Mode 00 = ACL does not specify the packet priority. Priority is determined by standard QoS functions. 01 = Change packet priority to P[2:0] if it is greater than QoS result. 10 = Change packet priority to P[2:0] if it is smaller than the QoS result. 11 = Always change packet priority to P[2:0]. |
| P[2:0] | Priority value |
| RPE | Remark Priority Enable 0 = Disable priority remarking 1 = Enable priority remarking. VLAN tag priority (PCP) bits are replaced by RP[2:0]. |
| RP[2:0] | Remark Priority value |
| MM[1:0] | Map Mode 00 = No forwarding remapping 01 = The forwarding map in FORWARD is OR'ed with the forwarding map from the Address Lookup Table. 10 = The forwarding map in FORWARD is AND'ed with the forwarding map from the Address Lookup Table. 11 = The forwarding map in FORWARD replaces the forwarding map from the Address Lookup Table. |
| FORWARD[N-1:0] | Forwarding Ports Bit 0 corresponds to port 1 Bit 1 corresponds to port 2, etc. 0 = Do not forward to this port 1 = Forward to this port |

TABLE 4-19: ACL ACTION RULE PARAMETERS FOR COUNT MODE (MD = 01 OR ENB = 00)

| Parameter | Description |
|-------------|---|
| COUNT[10:0] | Count value |
| TU | Time unit for counter. 0 = Microseconds 1 = Milliseconds |
| CA | Counter Algorithm. 0 = An 11-bit counter will be loaded with the count value from the list and start counting down every unit time. An interrupt will be generated when the timer expires, i.e. the next qualified packet has not been received within the period specified by the value. 1 = The counter is incremented with every matched packet received. An interrupt is generated when the terminal count is reached. The count resets thereafter. Time units are not used in this mode. |

Figure 4-7 shows basic organization of the ACL Table. The table has 16 entries, and each entry includes a matching field, action field and process field. Although these fields are stored together in one table, it is important to note that for a given table entry, the Matching, Action and Process fields generally do not form an associated group. The one exception is when the Matching Rule is in Count Mode (MD = 01 and ENB = 00). In that case, the Matching and Action fields are used in tandem.

FIGURE 4-7: ACL TABLE FORMAT



4.4.16.4 ACL Interrupts

The ACL filtering functions do not generate interrupts. Interrupts apply only for the Count Mode (MD = 01, ENB = 00). The Matching Rule can be configured either to timeout if the interval between packets of a specific type (MAC address and EtherType), or when a set number of these packets are received. There is a separate interrupt for each port. Port specific interrupt status and masks are located in the [Port Interrupt Status Register](#) and [Port Interrupt Mask Register](#). The top level interrupt registers for each port are in the [Global Port Interrupt Status Register](#) and [Global Port Interrupt Mask Register](#).

KSZ9896C

4.4.16.5 ACL Registers

Table 4-20 provides a list of ACL related registers.

TABLE 4-20: ACL REGISTERS

| Registers | Description |
|---|-------------------------|
| Port Interrupt Status Register, Port Interrupt Mask Register | ACL interrupt |
| Port ACL Access 0 Register through Port ACL Access F Register, Port ACL Byte Enable MSB Register, Port ACL Byte Enable LSB Register, Port ACL Access Control 0 Register | ACL Table access |
| Port Priority Control Register | Priority classification |
| Port Authentication Control Register | ACL enable |

4.5 NAND Tree Support

The KSZ9896C provides parametric NAND tree support for fault detection between chip I/Os and board. The NAND tree is a chain of nested NAND gates in which each KSZ9896C digital I/O (NAND tree input) pin is an input to one NAND gate along the chain. At the end of the chain, the **INTRP_N** pin provides the output for the last NAND gate.

The NAND tree test process includes:

- Enabling NAND tree mode
- Pulling all NAND tree input pins high
- Driving low each NAND tree input pin sequentially per the NAND tree pin order, starting with the first row of Table 4-21.
- Checking the NAND tree output to ensure there is a toggle high-to-low or low-to-high for each NAND tree input driven low.

TABLE 4-21: NAND TREE TEST PIN ORDER

| NAND Tree Sequence | Pin Number | Pin Name | NAND Tree Description |
|--------------------|------------|------------------------|-----------------------|
| 1 | 52 | TXD6_7 | Input |
| 2 | 53 | TXD6_6 | Input |
| 3 | 54 | TXD6_5 | Input |
| 4 | 55 | TXD6_4 | Input |
| 5 | 58 | RXD6_7 | Input |
| 6 | 59 | RXD6_6 | Input |
| 7 | 60 | RXD6_5 | Input |
| 8 | 62 | RXD6_4 | Input |
| 9 | 63 | TXCLK6/REFCLKI6 | Input |
| 10 | 64 | TX_EN6/TXCTL6 | Input |
| 11 | 65 | TX_ER6 | Input |
| 12 | 66 | COL6 | Input |
| 13 | 67 | TXD6_3 | Input |
| 14 | 68 | TXD6_2 | Input |
| 15 | 69 | TXD6_1 | Input |
| 16 | 70 | TXD6_0 | Input |
| 17 | 72 | RX_CLK6/REFCLKO6 | Input |
| 18 | 73 | RX_DV6/CRS_DV6/RX_CTL6 | Input |

TABLE 4-21: NAND TREE TEST PIN ORDER (CONTINUED)

| NAND Tree Sequence | Pin Number | Pin Name | NAND Tree Description |
|--------------------|------------|----------------|-----------------------|
| 19 | 75 | RX_ER6/RX_CLK6 | Input |
| 20 | 76 | CRS6 | Input |
| 21 | 77 | RXD6_3 | Input |
| 22 | 78 | RXD6_2 | Input |
| 23 | 79 | RXD6_1 | Input |
| 24 | 80 | RXD6_0 | Input |
| 25 | 82 | LED4_0 | Input |
| 26 | 83 | LED4_1 | Input |
| 27 | 85 | LED3_0 | Input |
| 28 | 86 | LED3_1 | Input |
| 29 | 87 | NC | Input |
| 30 | 88 | NC | Input |
| 31 | 89 | LED2_0 | Input |
| 32 | 90 | LED2_1 | Input |
| 33 | 91 | PME_N | Input |
| 34 | 94 | CLKO_25_125 | Input |
| 35 | 96 | SDO | Input |
| 36 | 97 | SDI/SDA/MDIO | Input |
| 37 | 99 | SCS_N | Input |
| 38 | 100 | SCL/MDC | Input |
| 39 | 102 | LED5_0 | Input |
| 40 | 103 | LED5_1 | Input |
| 41 | 105 | LED1_0 | Input |
| 42 | 106 | LED1_1 | Input |
| 43 | 92 | INTRP_N | Output |

The following procedure can be used to check for faults on the KSZ9896C digital I/O pin connections to the board:

1. Enable NAND tree mode via the **LED2_1**, **LED2_0**, and **LED4_0** configuration strap pins option.
2. Use board logic to drive all KSZ9896C NAND tree input pins high and verify that the **INTRP_N** pin output is high.
3. Use board logic to drive each NAND tree input pin, per the NAND Tree pin order, as follows:
 - a) Toggle the first pin in the NAND tree sequence (**TXD6_7**) from high to low, and verify the **INTRP_N** pin switches from high to low to indicate that the first pin is connected properly.
 - b) Leave the first pin (**TXD6_7**) low.
 - c) Toggle the second pin in the NAND tree sequence (**TXD6_6**) from high to low, and verify the **INTRP_N** pin switches from low to high to indicate that the second pin is connected properly.
 - d) Leave the first pin (**TXD6_7**) and the second pin (**TXD6_6**) low.
 - e) Toggle the third pin in the NAND tree sequence (**TXD6_5**) from high to low, and verify the **INTRP_N** pin switches from high to low to indicate that the third pin is connected properly.
 - f) Continue with this sequence until all KSZ9896C NAND tree input pins have been toggled.

Each KSZ9896C NAND tree input pin must cause the **INTRP_N** output pin to toggle high-to-low or low-to-high to indicate a good connection. If the **INTRP_N** pin fails to toggle when the KSZ9896C input pin toggles from high to low, the input pin has a fault.

KSZ9896C

4.6 Clocking

4.6.1 PRIMARY CLOCK

The device requires a 25MHz reference clock input at the **XI** pin. This clock is internally multiplied up and used to clock all of the internal logic and switching functions. It is also normally used as to clock the PHY transmit paths. This clock may be supplied by connecting a crystal between the **XI** and **XO** pins (and appropriate load capacitors to ground). Alternatively, an external CMOS clock signal may drive **XI**, while **XO** is left unconnected. The **XI/XO** block is powered from **AVDDH**.

4.6.2 PORT 6 RGMII/MII/RMII CLOCKS

The MII interface is clocked asymmetrically, with the PHY device driving both the **RX_CLK6** receive clock and the **TX_CLK6** transmit clock to the MAC device. The MII port may be configured at reset by a strapping option to take the role of either the PHY or the MAC. **RX_CLK6** and **TX_CLK6** are therefore either both inputs or both outputs, depending on the MII mode.

The RMII interface uses a single 50MHz clock. This REFCLK may be sourced either from the KSZ9896C or from the connected device. A strapping option is used to select the mode. "Normal Mode" is the mode where the other device supplies the clock, and the clock is an input to the **REFCLKI6** pin of the device. "Clock Mode" is the mode where the KSZ9896C generates the 50MHz clock on the **REFCLKO6** pin.

The RGMII interface employs source synchronous clocking, so it is symmetrical and does not require a mode selection. An output clock is generated on the **RX_CLK6** pin, while an input clock is received on the **TX_CLK6** pin. The clock speed scales with the interface data rate - either 10, 100 or 1000 Mbps. A strapping option is used to select between the 100 and 1000 Mbps speeds. If the 10 Mbps rate is required, then a register setting is used to set that speed.

The Port 6 MAC interface is powered from **VDDIO**.

| |
|---|
| Note: Refer to Section 3.2.1, "Configuration Straps," on page 17 for additional information on using configuration straps. |
|---|

4.6.3 SERIAL MANAGEMENT INTERFACE CLOCK

Whether configured to be SPI, I²C or MIIM, the KSZ9896C is always a slave and receives the clock as an input. The serial management interface is powered from **VDDIO**.


4.6.4 CLKO_25_125

An output clock, derived from the local 25MHz reference at **XI**, is provided on the **CLKO_25_125** pin. The output frequency choices are 25MHz (default) and 125MHz. If not needed, this output clock can also be disabled. **CLKO_25_125** is controlled via the [Output Clock Control Register](#), and is powered from **VDDIO**.

4.7 Power

The KSZ9896C requires two to three supply voltages. The device core operates from a 1.2V supply (**DVDDL** and **AVDDL**). The PHY transceivers and XI/XO crystal/clock interface operate from a 2.5V supply (**AVDDH**). The digital I/O's can be operated from 1.8V, 2.5V or 3.3V (**VDDIO**). The digital I/Os powered from **VDDIO** include RGMII, RMII, MII, SPI, I²C, MIIM, LED, **RESET_N**, **PME_N**, **INTRP_N** and **CLKO_25_125**. An example power connection diagram can be seen in [Figure 4-8](#).

FIGURE 4-8: POWER CONNECTION DIAGRAM



4.8 Power Management

The device supports enhanced power management features in a low-power state with energy detection to ensure low-power dissipation during device idle periods. There are three operation modes under the power management function which are implemented globally (i.e., applying to all ports):

- Normal Operation Mode
- Energy Detect Mode
- Global Soft Power Down Mode

[Table 4-22](#) summarizes all internal function blocks status under the three power-management operation modes.

TABLE 4-22: MDI/MDI-X PIN DEFINITIONS

| Functional Blocks | Power Management Operation Modes | | |
|--------------------|----------------------------------|---------------------|----------------------|
| | Normal Mode | Energy Detect Mode | Soft Power Down Mode |
| Internal PLL Clock | Enabled | Disabled | Disabled |
| TX/RX PHYs | Enabled | Energy Detect at RX | Disabled |

KSZ9896C

TABLE 4-22: MDI/MDI-X PIN DEFINITIONS (CONTINUED)

| Functional Blocks | Power Management Operation Modes | | |
|-------------------|----------------------------------|--------------------|----------------------|
| | Normal Mode | Energy Detect Mode | Soft Power Down Mode |
| MACs | Enabled | Disabled | Disabled |
| Host Interface | Enabled | Disabled | Disabled |

There is one additional power saving mode that may be implemented on a per-port basis:

- Port-Based Power Down

The first three global power modes are mutually exclusive; only one mode may be selected at a time. Port-based power down may be enabled independent of the global power mode.

4.8.1 NORMAL OPERATION MODE

At power-up, the device enters into Normal operation mode. It is also selected via bits [4:3] = 00 in the [Power Down Control 0 Register](#). When the device is in normal operation mode, all PLL clocks are running, PHYs and MACs are on, and the CPU is ready to read or write the device registers through the serial interface (SPI, I²C or MIIM).

During normal operation mode, the host processor can change the power management mode bits in the [Power Down Control 0 Register](#) to transition to any of the other power management modes.

4.8.2 ENERGY-DETECT MODE

Energy-detect mode, also known as energy-detect power down (EDPD) mode, is enabled by setting bits [4:3] to 01 in the [Power Down Control 0 Register](#). Energy-detect mode provides a mechanism to save power when the device is not connected to an active link partner. Auto-negotiation must be enabled when in energy-detect mode.

Energy-detect mode consists of two states, normal-power state and low-power state. When the device is in this mode, it will monitor the cable energy. If there is no energy on the cable for a time longer than a pre-configured value, the device will go into the low-power state. While in low-power state, the device reduces power consumption by disabling all circuitry except the energy-detect circuitry of the receiver, which consumes minimal power. When the device is in the low-power state, it will transmit link pulses at long intervals, with a very low duty cycle. At the same time, it continuously monitors for energy on the cable. Once energy is detected from the cable and is present for a time longer than 100ns, the device will enter the normal-power state.

4.8.3 GLOBAL SOFT POWER-DOWN MODE

Soft power-down mode is used to power down the device when it is not in use after power-up. This mode disables all internal functions except for the serial (SPI or I²C) management interface.

When soft power-down mode is exited, all registers are reset to their default values, and all configuration strap pins are sampled to set the device settings.

4.8.4 PORT-BASED POWER DOWN

Unused ports may be powered down individually to save power.

4.8.5 WAKE ON LAN (WOL)

Wake on LAN allows a computer to be turned on or woken up by a network message. The message is usually sent by a program executed on another computer on the same local area network. Wake-up frame events are used to awaken the system whenever meaningful data is presented to the system over the network. Examples of meaningful data include the reception of a Magic Packet, a management request from a remote administrator, or simply network traffic directly targeted to the local system. The device can be programmed to notify the host of the Wake-Up frame detection with the assertion of the power management event signal (PME_N).

The device's MACs support the detection of the following Wake-Up events:

- Detection of energy signal over a pre-configured value
- Detection of a linkup in the network link state
- Receipt of a Magic Packet

There are also other types of Wake-Up events that are not listed here as manufacturers may choose to implement these in their own way.

4.8.5.1 Direction of Energy

The energy is detected from the cable and is continuously presented for a time longer than pre-configured value, especially when this energy change may impact the level at which the system should re-enter to the normal power state.

4.8.5.2 Direction of Link-up

Link status wake events are useful to indicate a linkup in the network's connectivity status.

4.8.5.3 Magic Packet™

The Magic Packet is a broadcast frame containing anywhere within its payload 6 bytes of all 1s (FF FF FF FF FF FF) followed by sixteen repetitions of the target computer's 48-bit DA MAC address. Since the magic packet is only scanned for the above string, and not actually parsed by a full protocol stack, it may be sent as any network- and transport-layer protocol.

Magic Packet technology is used to remotely wake up a sleeping or powered-off PC on a LAN. This is accomplished by sending a specific packet of information, called a Magic Packet frame, to a node on the network. When a PC capable of receiving the specific frame goes to sleep, it enables the Magic Packet RX mode in the LAN controller, and when the LAN controller receives a Magic Packet frame, it will alert the system to wake up. Once the device has been enabled for Magic Packet Detection, it scans all incoming frames addressed to the node for a specific data sequence, which indicates to the controller this is a Magic Packet frame.

A Magic Packet frame must also meet the basic requirements for the LAN technology chosen, such as Source Address (SA), Destination Address (DA), which may be the receiving station's IEEE MAC address, or a multicast or broadcast address and CRC. The specific sequence consists of 16 duplications of the MAC address of this node, with no breaks or interruptions. This sequence can be located anywhere within the packet, but must be preceded by a synchronization stream. The synchronization stream is defined as 6 bytes of 0xFF. The device will also accept a broadcast frame, as long as the 16 duplications of the IEEE address match the address of the machine to be awakened.

4.8.5.4 Interrupt Generation on Power Management Related Events

There are two ways an interrupt can be generated to the host whenever a power management related event takes place. The resulting interrupts are via the **PME_N** signal pin or via the **INTRP_N** signal pin.

4.9 Management Interface

The management interface may be used by an external host processor to read and write the device's registers. This interface has three available modes of operation: SPI, I²C or MIIM. The interface mode is selected at the deassertion of reset by a strapping option (refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional information).

Of the three interface options, SPI provides the highest performance, while MIIM performance is the lowest. Most importantly, MIIM provides access to the PHY control and status registers, but not to any of the switch registers. The vast majority of applications therefore can use SPI or I²C, but not MIIM.

Register access is also available through the high-performance in-band management interface as described in [Section 4.10, "In-Band Management," on page 57](#).

4.9.1 SPI SLAVE BUS

The KSZ9896C supports a slave mode SPI interface that provides complete access to all device registers via an SPI master device. The SPI master device supplies the clock (**SCL**), select (**SCS_N**), and serial input data (**SDI**). Serial output data (**SDO**) is driven by the KSZ9896C.

SCL is expected to stay low when SPI operation is idle. SPI operations start with the falling edge of **SCS_N** and end with the rising edge of **SCS_N**. A single read or write access consists of a 27-bit command/address phase, then a 5-bit turnaround (TA) phase, then an 8-bit data phase. For burst read or write access, **SCS_N** is held low while **SCL** continues to toggle. For every 8 cycles of **SCL**, the device will increment the address counter, and the corresponding data byte will be transferred on **SDI** or **SDO** in succession.

All commands, addresses and data are transferred most significant bit first. Input data on **SDI** is latched on the rising edge of clock **SCL**. Output data on **SDO** is clocked on the falling edge of **SCL**.

As shown in [Figure 4-23](#), there are two commands: register read and register write. [Figure 4-9](#) and [Figure 4-10](#) show the timing for these two operations.

KSZ9896C

TABLE 4-23: REGISTER ACCESS USING THE SPI INTERFACE

| SPI Operation | Command/Address Phase (SDI pin) | | TA bits (Note 4-8) | Data Phase (SDO or SDI pins) |
|----------------|---------------------------------|---|-----------------------|---------------------------------|
| | Command | Register Address | | |
| Register Read | 011 | A23 A22 A21 A20 ... A7 A6 A5 A4 A3 A2 A1 A0 | XXXXX | D7 D6 D5 D4 D3 D2 D1 D0 |
| Register Write | 010 | A23 A22 A21 A20 ... A7 A6 A5 A4 A3 A2 A1 A0 | XXXXX | D7 D6 D5 D4 D3 D2 D1 D0 |

Note 4-8 TA bits are turn-around bits. They are “don't care” bits.

Note: The actual device address space is 16 bits (A15 - A0), so the values of address bits A23 - A16 in the SPI command/address phase are “don't care”.

FIGURE 4-9: SPI REGISTER READ OPERATION





FIGURE 4-10: SPI REGISTER WRITE OPERATION



4.9.2 I²C BUS

The management interface may be configured to be an I²C slave. In this mode, an I²C master has complete programming access to the device's internal control and status registers, including all MIB counters, address lookup tables, VLAN table and ACL table.

The 7-bit device address is fixed as 1011_111. Because of the fixed address, only one KSZ9896C may be on the I²C bus at a time. The R/W control bit is then appended as the least significant bit to form these 8-bit address/control words:

1011_1110 <write>

1011_1111 <read>

The internal registers and tables of the device are accessed using 16-bit addressing and 8-bit data. The access formats are as follows:

FIGURE 4-11: SINGLE BYTE REGISTER WRITE




FIGURE 4-12: SINGLE BYTE REGISTER READ




FIGURE 4-13: BURST REGISTER WRITE





FIGURE 4-14: BURST REGISTER READ



4.9.3 MII MANAGEMENT (MIIM) INTERFACE

The device supports the IEEE 802.3 MII management interface, also known as the management data input/output (MDIO) interface. This interface allows upper-layer devices to monitor and control the states of the KSZ9896C PHY blocks, but it does not provide access to the switch registers. An external device with MDC/MDIO capability can read the PHY status or configure the PHY settings. Details on the MIIM interface can be found in Clauses 22 and 45 of the IEEE 802.3 Specification.

Use of MIIM conflicts with use of the [In-Band Management](#) interface. These interfaces cannot be used simultaneously.

The MIIM interface consists of the following:

- A physical connection that uses a data signal (**MDIO**) and a clock signal (**MDC**) for communication between an external controller and the KSZ9896C. Note that the **MDIO** signal is open-drain.
- A specific protocol that operates across the two signal physical connection that allows an external controller to communicate with the internal PHY devices.
- Access to a set of standard, vendor-specific and extended (MMD) 16-bit registers. These registers are also directly accessible via the SPI and I²C interface options.

KSZ9896C

The MIIM Interface can operate up to a maximum clock speed of 5MHz. Access is limited to only the registers in the PHY blocks of ports 1 through 5. [Table 4-24](#) summarizes the MII management interface frame format.

TABLE 4-24: MII MANAGEMENT INTERFACE FRAME FORMAT

| Operation Mode | Preamble (32-bit) | Start of Frame (2-bit) | Operation Code (2-bit) | PHY Address (5-bit) | Register Address (5-bit) | Turn Around (2-bit) | Register Data (16-bit) | Idle |
|----------------|-------------------|------------------------|------------------------|---------------------|--------------------------|---------------------|------------------------|------|
| Read | All 1s | 01 | 10 | A[4:0] | Reg[4:0] | Z0 | D[15:0] | Z |
| Write | All 1s | 01 | 01 | A[4:0] | Reg[4:0] | 10 | D[15:0] | Z |

The MIIM PHY address to PHY port mapping is as follows:

- PHY Address 1h to PHY port 1
- PHY Address 2h to PHY port 2
- PHY Address 3h to PHY port 3
- PHY Address 4h to PHY port 4
- PHY Address 5h to PHY port 5

The MIIM register address space consists of two distinct areas.

- [Standard MIIM Registers \(Direct\)](#)
- [MDIO Manageable Device \(MMD\) Registers \(Indirect\)](#)

4.9.3.1 Standard MIIM Registers (Direct)

Standard registers provide direct read/write access to a 32-register address space, as defined in Clause 22 of the IEEE 802.3 Specification. Within this address space, the first 16 registers (Registers 0h to Fh) are defined according to the IEEE specification, while the remaining 16 registers (Registers 10h to 1Fh) are defined specific to the PHY vendor.

The KSZ9896C supports the standard registers listed in [Table 4-25](#) for each PHY port. Each 16-bit MIIM Standard Register Address maps to two corresponding 8-bit Port N Register Addresses. The register bit map and description are located at the 8-bit Port N Register Addresses.

TABLE 4-25: STANDARD MIIM REGISTERS

| MIIM Standard Register Address (hex) | Port N Register Address (hex) | Description |
|--------------------------------------|-------------------------------|--|
| IEEE-Defined Registers | | |
| 0h | 0xN100 - 0xN101 | PHY Basic Control Register |
| 1h | 0xN102 - 0xN103 | PHY Basic Status Register |
| 2h | 0xN104 - 0xN105 | PHY ID High Register |
| 3h | 0xN106 - 0xN107 | PHY ID Low Register |
| 4h | 0xN108 - 0xN109 | PHY Auto-Negotiation Advertisement Register |
| 5h | 0xN10A - 0xN10B | PHY Auto-Negotiation Link Partner Ability Register |
| 6h | 0xN10C - 0xN10D | PHY Auto-Negotiation Expansion Status Register |
| 7h | 0xN10E - 0xN10F | PHY Auto-Negotiation Next Page Register |
| 8h | 0xN110 - 0xN111 | PHY Auto-Negotiation Link Partner Next Page Ability Register |
| 9h | 0xN112 - 0xN113 | PHY 1000BASE-T Control Register |
| Ah | 0xN114 - 0xN115 | PHY 1000BASE-T Status Register |
| Bh-Ch | - | RESERVED |
| Dh | 0xN11A - 0xN11B | PHY MMD Setup Register |
| Eh | 0xN11C - 0xN11D | PHY MMD Data Register |
| Fh | 0xN11E - 0xN11F | PHY Extended Status Register |

TABLE 4-25: STANDARD MIIM REGISTERS (CONTINUED)

| MIIM Standard Register Address (hex) | Port N Register Address (hex) | Description |
|--------------------------------------|-------------------------------|--|
| Vendor-Specific Registers | | |
| 10h | - | RESERVED |
| 11h | 0xN122 - 0xN123 | PHY Remote Loopback Register |
| 12h | 0xN124 - 0xN125 | PHY LinkMD Register |
| 13h | 0xN126 - 0xN127 | PHY Digital PMA/PCS Status Register |
| 14h | - | RESERVED |
| 15h | 0xN12A - 0xN12B | Port RXER Count Register |
| 16h-1Ah | - | RESERVED |
| 1Bh | 0xN136 - 0xN137 | Port Interrupt Control / Status Register |
| 1Ch | 0xN138 - 0xN139 | PHY Auto MDI / MDI-X Register |
| 1Dh-1Eh | - | RESERVED |
| 1Fh | 0xN13E - 0xN13F | PHY Control Register |

4.9.3.2 MDIO Manageable Device (MMD) Registers (Indirect)

The MIIM interface provides indirect access to a set of MMD registers as defined in [Section 5.4, "MDIO Manageable Device \(MMD\) Registers \(Indirect\)," on page 167](#).

4.10 In-Band Management

The in-band management access (IBA) is a feature that provides full register read and write access via any one of the six data ports. Port 6 is the default IBA port. The in-band feature is enabled or disabled by a strapping option at power-up and reset. To use a different port instead of port 6 for IBA, the SPI or I²C interface or IBA must be used to write to a control register. IBA may not be used on more than one port at a time.

In-band management frames are processed differently from normal network frames. They are recognized as special frames, so address and VID lookup, VLAN tagging, source address filtering, un-tag discard, tagged frame drop, etc. are not applied to them. Received in-band management frames are never forwarded to the switch fabric or to any other port.

The [In-Band Management \(IBA\) Control Register](#) is used to enable and control the IBA feature and to specify one of the six ports as the IBA port.

The IBA frame format is shown in [Figure 4-15](#). The layer 2 portion of the IBA frame contains normal destination address (DA) and source address (SA) fields. The DA of the frames are defined to be the switch MAC address (default 00-10-A1-FF-FF-FF), and the SA is the MAC address of the source device. The DA and SA will be swapped in the response frame. A special 4-byte IBA tag follows the SA. This is then followed by the 2-byte EtherType/Length field that serves to identify this as an IBA frame.

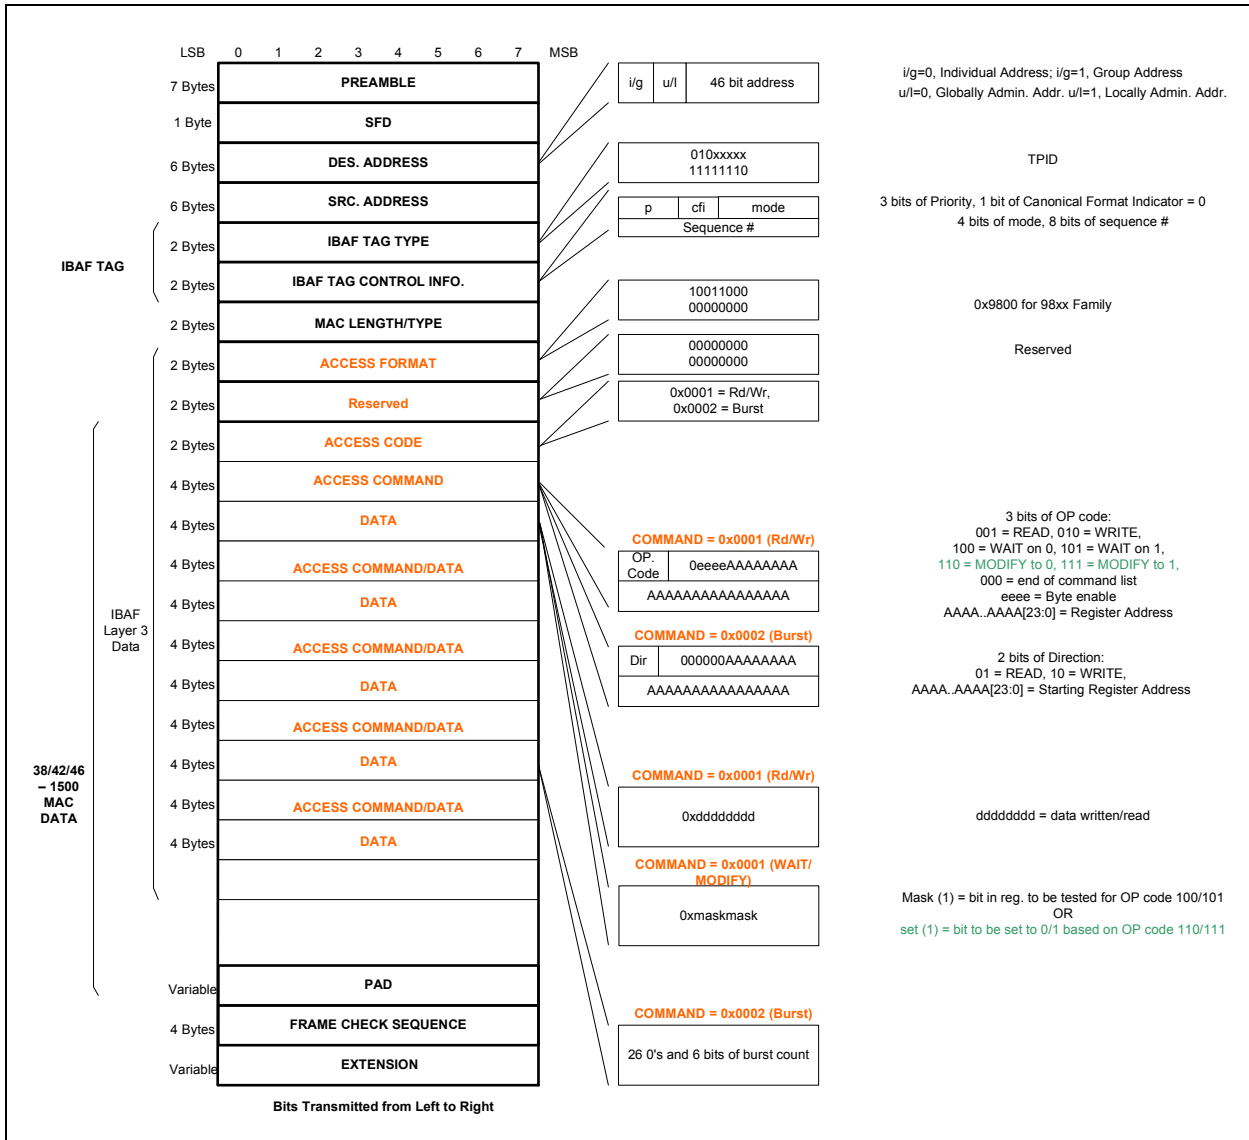
Only one IBA frame can be processed at a time. Any subsequent IBA frames received by the device will be dropped unless the most recent response frame has been fully transmitted.

There are six types of read/write commands: READ, WRITE, WAIT on 0, WAIT on 1, MODIFY to 0 and MODIFY to 1.

The burst commands offer fast and bundled data return, up to the capacity of the IBA frame buffer. There are two types of operations in burst command: READ burst and Write burst.

KSZ9896C

FIGURE 4-15: IN-BAND MANAGEMENT FRAME FORMAT



4.11 MAC Interface (Port 6)

Strapping options are used to individually select any of these MAC interface options for port 6:

- **Media Independent Interface (MII) (Port 6):** Supports 100 and 10 Mbps data rates
- **Reduced Media Independent Interface (RMII) (Port 6):** Supports 100 and 10 Mbps data rates
- **Gigabit Media Independent Interface (GMII) (Port 6):** Supports 1000Mbps data rate
- **Reduced Gigabit Media Independent Interface (RGMII) (Port 6):** Supports 1000, 100 and 10 Mbps data rates

Note that the signals on the KSZ9896C MAC interface are named as they would be for a PHY: the TX direction is into the KSZ9896C, while the RX direction is out of the KSZ9896C, as if to a host processor with integrated MAC. Signal connection to such a "MAC" device is TX-to-TX, and RX-to-RX.

An external PHY (such as the Microchip KSZ9031RNX) may be connected to port 6, but in that case the signal connection will be RX-to-TX, and TX-to-RX.

The Port 6 MII/RMII/GMII/RGMII interface is powered by the VDDIO power supply.

4.11.1 MEDIA INDEPENDENT INTERFACE (MII) (PORT 6)

The media independent interface (MII) is specified in Clause 22 of the IEEE 802.3 standard. It provides a common interface between PHY layer and MAC layer devices. The data interface is 4-bits wide and runs at one quarter the network bit rate; either 2.5MHz in 10BASE-T_e or 25MHz in 100BASE-TX (not encoded). Additional signals on the transmit side indicate when data is valid or when an error occurs during transmission. Similarly, the receive side provides signals that convey when the data is valid and without physical layer errors. For half duplex operation, the COL signal indicates if a collision has occurred during transmission.

The MII interface operates in either PHY Mode or MAC Mode. Select PHY Mode when the port is connected to a processor or other device with a MAC function; select MAC Mode when connecting to an external PHY. Note that the direction of the TX_CLK6, RX_CLK6, COL6 and CRS6 signals is affected by the PHY mode or MAC mode setting, while other MII signals do not change direction.

MII mode is selected at reset by a configuration strap option on pins RxD6_3 and RxD6_2 for port 6. The Speed strap option (on pin RxD6_0 for port 6) should be set for 100/10 Mbps Mode. PHY Mode or MAC Mode is selected by a configuration strap option on pin RxD6_1 (port 6). Refer to Section 3.2.1, "Configuration Straps," on page 17 for additional information.

Another configuration strap on the RX_ER6/RX_CLK6 pin selects either 2-Wire Mode or 3-Wire Mode for the GMII / MII clocking. This option determines which pin is used for the KSZ9896C's MII RX_CLK. The 3-Wire Mode option is intended for use when connecting to a 1000/100/10 Mbps PHY such as the Microchip KSZ9031MNX, which has separate pins for its GTX_CLK and TX_CLK signals. When 3-Wire Mode is selected, the MII (MAC Mode) RX_CLK6 input clock is on pin 75. When in 2-Wire Mode, RX_CLK6 is on pin 72. If GMII will not be used, or when in PHY Mode, the clock mode should be 2-Wire. Refer to Section 3.2.1, "Configuration Straps," on page 17 for additional information.

The interface contains two distinct groups of signals, one for transmission and the other for reception. Table 4-26 and Table 4-27 describe the signals used by the MII interface to connect to an external MAC or to an external PHY, respectively.

TABLE 4-26: MII (PHY MODE) CONNECTION TO EXTERNAL MAC

| MII Interface Signals Description | KSZ9896C Signals in PHY Mode | External MAC Device Signals |
|-----------------------------------|------------------------------|-----------------------------|
| Transmit Enable | TX_EN6 (input) | TX_EN (output) |
| Transit Error | TX_ER6 (input) | TX_ER (output) |
| Transmit Data Bits [3:0] | TXD6_[3:0] (input) | TXD[3:0] (output) |
| Transmit Clock | TX_CLK6 (output) | TX_CLK (input) |
| Collision Detection | COL6 (output) | COL (input) |
| Carrier Sense | CRS6 (output) | CRS (input) |
| Received Data Valid | RX_DV6 (output) | RX_DV (input) |
| Receive Error | RX_ER6 (output) | RX_ER (input) |
| Receive Data Bits [3:0] | RXD6[3:0] (output) | RXD[3:0] (input) |
| Receive Clock | RX_CLK6 (output) | RX_CLK (input) |

TABLE 4-27: MII (MAC MODE) CONNECTION TO EXTERNAL PHY

| MII Interface Signals Description | KSZ9896C Signals in MAC Mode | External PHY Device Signals |
|-----------------------------------|--|-----------------------------|
| Transmit Enable | RX_DV6 (output) | TX_EN (input) |
| Transit Error | RX_ER6 (output) | TX_ER (input) |
| Transmit Data Bits [3:0] | RXD6_[3:0] (output) | TXD[3:0] (input) |
| Transmit Clock | RX_CLK6 (input) (Actual pin depends on GMII/MII Clock Mode) | TX_CLK (output) |
| Collision Detection | COL6 (input) | COL (output) |
| Carrier Sense | CRS6 (input) | CRS (output) |

KSZ9896C

TABLE 4-27: MII (MAC MODE) CONNECTION TO EXTERNAL PHY (CONTINUED)

| MII Interface Signals Description | KSZ9896C Signals in MAC Mode | External PHY Device Signals |
|-----------------------------------|------------------------------|-----------------------------|
| Received Data Valid | TX_EN6 (input) | RX_DV (output) |
| Receive Error | TX_ER6 (input) | RX_ER (output) |
| Receive Data Bits [3:0] | TXD6_[3:0] (input) | RXD[3:0] (output) |
| Receive Clock | TX_CLK6 (input) | RX_CLK (output) |

4.11.2 REDUCED MEDIA INDEPENDENT INTERFACE (RMII) (PORT 6)

The reduced media independent interface (RMII) specifies a low pin count interface, which is based on MII, that provides communication with a MAC attached to the port. As with MII, RMII provides a common interface between physical layer and MAC layer devices, or between two MAC layer devices, and has the following key characteristics:

- Supports network data rates of either 10Mbps or 100Mbps.
- Uses a single 50MHz clock reference (provided internally or externally) for both transmit and receive data.
- Uses independent 2-bit wide transmit and receive data paths.
- Contains two distinct groups of signals: one for transmission and the other for reception.

The user selects one of the two RMII clocking modes by setting the appropriate strapping option.

While in RMII Normal Mode, the port will require an external 50MHz signal to be input to **TX_CLK6/REFCLKI6** from an external source. This mode is selected by strapping the **RXD6_1** pin high during reset.

While in RMII Clock Mode, the port will output a 50MHz clock on **RX_CLK6/REFCLKO6**, which is derived from the 25MHz crystal or oscillator attached to the **XI** clock input. The **TX_CLK6/REFCLKI6** input is unused in this mode. This mode is selected by strapping the **RXD6_1** pin low during reset.

Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional configuration strap information.

[Table 4-28](#) describes the signals used by the RMII interface. Refer to the RMII specification for full details on the signal descriptions.

TABLE 4-28: RMII SIGNAL DESCRIPTIONS

| RMII Signal Name (per spec) | RMII Signal (per KSZ9896C) | Pin Direction (with Respect to PHY, KSZ9896C) | Pin Direction (with Respect to MAC) | RMII Signal Description |
|-----------------------------|----------------------------|---|-------------------------------------|---|
| REF_CLK | REFCLKI6 | Input | Input or Output | Synchronous 50MHz reference clock, when port is in RMII Normal Mode |
| n/a | REFCLKO6 | Output | Input | Synchronous 50MHz reference clock, when port is in RMII Clock Mode |
| TX_EN | TX_EN6 | Input | Output | Transmit Enable |
| TXD[1:0] | TXD6_[1:0] | Input | Output | Transmit Data Bit [1:0] |
| CRS_DV | RX_DV6 | Output | Input | Carrier Sense / Receive Data Valid |
| RX_ER | RX_ER6 | Output | Input or not required | Receive Error |
| RXD[1:0] | RXD6_[1:0] | Output | Input | Receive Data Bit [1:0] |

A device port in RMII mode may connect to either an external MAC device (such as a host processor) or to an external PHY; but unlike MII, RMII does not provide separate PHY and MAC modes of operation. However, it is necessary to connect the pins properly.

TABLE 4-29: RMII CONNECTION TO EXTERNAL MAC

| RMII Interface Signals Description | KSZ9896C Signals | External MAC Device Signals |
|------------------------------------|--|------------------------------|
| Transmit Enable | TX_EN6 (input) | TX_EN (output) |
| Transmit Data Bits [1:0] | TXD6_[1:0] (input) | TXD[1:0] (output) |
| Reference Clock | REFCLKI6 (input) or REFCLKO6 (output) | REF_CLK (input or output) |
| Carrier Sense Data Valid | RX_DV6 (output) | CRS_DV (input) |
| Receive Error | RX_ER6 (output) | RX_ER (input) |
| Receive Data Bits [1:0] | RXD6_[1:0] (output) | RXD[1:0] (input) |

TABLE 4-30: RMII CONNECTION TO EXTERNAL PHY

| RMII Interface Signals Description | KSZ9896C Signals | External PHY Device Signals |
|------------------------------------|--|------------------------------|
| Transmit Enable | RX_DV6 (output) | TX_EN (input) |
| Transmit Data Bits [1:0] | RXD6_[1:0] (output) | TXD[1:0] (input) |
| Reference Clock | REFCLKI6 (input) or REFCLKO6 (output) | REF_CLK (input or output) |
| Carrier Sense Data Valid | TX_EN6 (input) | CRS_DV (output) |
| Receive Error | No connection | RX_ER (output) |
| Receive Data Bits [1:0] | TXD6_[1:0] (input) | RXD[1:0] (output) |

4.11.3 GIGABIT MEDIA INDEPENDENT INTERFACE (GMII) (PORT 6)

The Gigabit Media Independent Interface (GMII) is compliant to clause 35 of the IEEE 802.3 standard. It provides a common interface between GMII PHYs and MACs at 1000 Mbps. GMII is based heavily on the MII interface, with the key signal differences being the data bus width, clock rate, and direction of the transmit clock.

The KSZ9896C interface combines GMII mode with MII mode to form GMII/MII mode to support data transfer at 10/100/1000 Mbps. Port 6 is configured at power-up or reset to GMII/MII mode via the configuration strap option on the **RXD6_3** and **RXD6_2** pins. The Port 6 Speed Select configuration strap, on the **RXD6_0** pin, determines whether the interface is in MII mode (for 100 and 10 Mbps), or GMII mode (for 1000 Mbps) at power-up or reset. There is no method for the KSZ9896C to automatically switch speeds based on the speed of the connected device, but any of the three speeds can be selected via the port 6 [XMII Port Control 0 Register](#) and [XMII Port Control 1 Register](#). When switching from 1000 Mbps to 100 or 10 Mbps speed, the interface switches automatically from GMII to MII operation.

When this change occurs, one of the two GMII / MII clocks will change direction. If the GMII / MII interface is in PHY mode, the transmit clock (TX_CLK) will change between input (GMII) and output (MII). If the interface is in MAC mode, the receive clock (RX_CLK) will change between output (GMII) and input (MII).

Another configuration strap on the **RX_ER6/RX_CLK6** pin selects either 2-Wire Mode or 3-Wire Mode for the GMII / MII clocking. This option determines which pin is used for the KSZ9896C's MII RX_CLK. The 3-Wire Mode option is intended for use when connecting to a 1000/100/10 Mbps PHY such as the Microchip KSZ9031MNX, which has separate pins for its GTX_CLK and TX_CLK signals. When 3-Wire Mode is selected, the MII (MAC Mode) **RX_CLK6** input clock is on pin 75. When in 2-Wire Mode, **RX_CLK6** is on pin 72. If GMII will not be used, or when in PHY Mode, the clock mode should be 2-Wire. Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional information.

KSZ9896C

TABLE 4-31: GMII CONNECTION (PHY MODE) TO EXTERNAL MAC INTERFACE

| GMII Interface Signals Description | KSZ9896C Signals in PHY Mode | External MAC Device Signals |
|------------------------------------|------------------------------|-----------------------------|
| Transmit Enable | TX_EN6 (input) | TX_EN (output) |
| Transmit Error | TX_ER6 (input) | TX_ER (output) |
| Transmit Data Bits [17:0] | TXD6_[7:0] (input) | TXD[7:0] (output) |
| Transmit Clock | GTX_CLK (input) | GTX_CLK (output) |
| Collision Detection | COL6 (output) | COL (input) |
| Carrier Sense | CRS_DV6 (output) | CRS (input) |
| Receive Data Valid | RX_DV6 (output) | RX_DV (input) |
| Receive Error | RX_ER6 (output) | RX_ER (input) |
| Receive Data Bits [7:0] | RXD6_[7:0] (output) | RXD[7:0] (input) |
| Receive Clock | RX_CLK6 (output) | RX_CLK (input) |

TABLE 4-32: GMII CONNECTION (MAC MODE) TO EXTERNAL PHY

| GMII Interface Signals Description | KSZ9896C Signals in MAC Mode | External PHY Device Signals |
|------------------------------------|------------------------------|-----------------------------|
| Transmit Enable | RX_DV6 (output) | TX_EN (input) |
| Transmit Error | RX_ER6 (output) | TX_ER (input) |
| Transmit Data Bits [17:0] | RXD6_[7:0] (output) | TXD[7:0] (input) |
| Transmit Clock | GRX_CLK (output) | GTX_CLK (input) |
| Collision Detection | COL6 (input) | COL (output) |
| Carrier Sense | CRS_DV6 (input) | CRS (output) |
| Receive Data Valid | TX_EN6 (input) | RX_DV (output) |
| Receive Error | TX_ER6 (input) | RX_ER (output) |
| Receive Data Bits [7:0] | TXD6_[7:0] (output) | RXD[7:0] (output) |
| Receive Clock | GTX_CLK6 (input) | RX_CLK (output) |

4.11.4 REDUCED GIGABIT MEDIA INDEPENDENT INTERFACE (RGMII) (PORT 6)

RGMII provides a common interface between RGMII PHYs and MACs, and has the following key characteristics:

- Pin count is reduced from 24 pins for GMII to 12 pins for RGMII.
- All speeds (10Mbps, 100Mbps and 1000Mbps) are supported at both half- and full-duplex.
- Data transmission and reception are independent and belong to separate signal groups.
- Transmit data and receive data are each four bits wide - a nibble.

In RGMII operation, the RGMII pins function as follows:

- The MAC sources the transmit reference clock, **TX_CLK6**, at 125MHz for 1000Mbps, 25MHz for 100Mbps, and 2.5MHz for 10Mbps.
- The PHY recovers and sources the receive reference clock, **RX_CLK6**, at 125MHz for 1000Mbps, 25MHz for 100Mbps, and 2.5MHz for 10Mbps.
- For 1000BASE-T, the transmit data, **TXD6_[3:0]**, is presented on both edges of **TX_CLK6**, and the received data, **RXD6_[3:0]**, is clocked out on both edges of the recovered 125MHz clock, **RX_CLK6**.
- For 10BASE-T/100BASE-TX, the MAC holds **TX_CTL6** low until both the PHY and MAC operate at the same speed. During the speed transition, the receive clock is stretched on either a positive or negative pulse to ensure that no clock glitch is presented to the MAC.
- **TX_ER6** and **RX_ER6** are combined with **TX_EN6** and **RX_DV6**, respectively, to form **TX_CTL6** and **RX_CTL6**. These two RGMII control signals are valid at the falling clock edge.

After power-up or reset, the device is configured to RGMII mode if the appropriate configuration strap pins are set to one of the RGMII mode capability options. Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for available options. Note that there is no mechanism for the RGMII interface to adapt its speed automatically to the speed of the connected RGMII device. A configuration strap option sets the speed of the RGMII interface at power-up to either 1000Mbps or 100Mbps. A control register can override the configuration strap option and set the RGMII speed to either 1000, 100 or 10Mbps. If a PHY is connected to the RGMII port, it should be ensured that the PHY link speed is fixed in order to avoid a mismatch to the RGMII speed.

The device provides the option to add a minimum of 1.5ns internal delay to either **TX_CLK6** or **RX_CLK6**, via the RGMII Internal Delay control bits in the [XMII Port Control 1 Register](#). This can reduce or eliminate the need to add trace delay to the clock signals on the printed circuit board. **RGMII_ID_ig** enables delay on **TX_CLK6**, and the default is off. **RGMII_ID_eg** enables delay on **RX_CLK6**, and the default is on. Users should also be aware of any internal clock delay that may be added by the connected RGMII device.

TABLE 4-33: RGMII SIGNAL DESCRIPTIONS

| RGMII Signal Name (per spec) | RGMII Signal (per KSZ9896C) | Pin Direction (with respect to PHY, KSZ9896C) | Pin Direction (with respect to MAC) | RGMII Signal Description |
|------------------------------|-----------------------------|---|-------------------------------------|--|
| TXC | TX_CLK6 | Input | Output | Transmit Reference Clock (125MHz for 1000Mbps, 25MHz for 100Mbps, 2.5MHz for 10Mbps) |
| TX_CTL | TX_CTL6 | Input | Output | Transmit Control |
| TXD[3:0] | TXD6_[3:0] | Input | Output | Transmit Data [3:0] |
| RXC | RX_CLK6 | Output | Input | Receive Reference Clock (125MHz for 1000Mbps, 25MHz for 100Mbps, 2.5MHz for 10Mbps) |
| RX_CTL | RX_CTL6 | Output | Input | Receive Control |
| RXD[3:0] | RXD6_[3:0] | Output | Input | Receive Data [3:0] |


KSZ9896C

5.0 DEVICE REGISTERS

The KSZ9896C has a rich set of registers for device management. The registers are accessed by the SPI or I²C interfaces, or by in-band management. Alternatively, the MIIM interface can be used to access the PHY registers only. The MIIM interface cannot access the switch registers.

A 16-bit address is used to access the device registers. This address is split into three hierarchical spaces, as shown in Figure 5-1. These three spaces are used to designate the port/channel (4-bits), function (page) of the port (4-bits), and register of function (8-bits). The individual ports are numbered 1 through 6. In the port space, a value of 0 is used for global registers. Address bit 15 is always 0.


FIGURE 5-1: REGISTER ADDRESS MAPPING



Register addressing is by bytes, and the management interface (SPI, I²C or in-band) transfers data by bytes. Where registers are shown as 16-bits or 32-bits, this is for descriptive purposes only. Data can always be written and read as individual bytes and in any order.

For multi-byte registers, the data is addressed in a big-endian format, with the most significant byte at the lowest address, and the least significant byte at the highest address, as shown in [Figure 5-2](#).

FIGURE 5-2: BYTE ORDERING



The global and port register address maps are detailed in [Table 5-1](#) and [Table 5-2](#), respectively. [Table 1-3](#), “[Register Nomenclature](#),” on [page 7](#) provides a list of register bit type notations.

The remainder of this chapter is organized as follows:

- [Global Registers](#)
- [Port Registers](#)
- [Tables and MIB Counters \(Access\)](#)
- [MDIO Manageable Device \(MMD\) Registers \(Indirect\)](#)

TABLE 5-1: GLOBAL REGISTER ADDRESS MAP

| Address | Functional Group |
|-----------------|--|
| 0x0000 - 0x00FF | Global Operation Control Registers (0x0000 - 0x00FF) |
| 0x0100 - 0x01FF | Global I/O Control Registers (0x0100 - 0x01FF) |
| 0x0200 - 0x02FF | Global PHY Control and Status Registers (0x0200 - 0x02FF) |
| 0x0300 - 0x03FF | Global Switch Control Registers (0x0300 - 0x03FF) |
| 0x0400 - 0x04FF | Global Switch Look Up Engine (LUE) Control Registers (0x0400 - 0x04FF) |
| 0x0500 - 0x0FFF | RESERVED |

KSZ9896C

TABLE 5-2: PORT N (1-6) REGISTER ADDRESS MAP

| Address | Functional Group |
|-----------------|--|
| 0xN000 - 0xN0FF | Port N: Port Operation Control Registers (0xN000 - 0xN0FF) |
| 0xN100 - 0xN1FF | Port N: Port Ethernet PHY Registers (0xN100 - 0xN1FF) |
| 0xN200 - 0xN2FF | RESERVED |
| 0xN300 - 0xN3FF | Port N: Port RGMII/GMII/MII/RMII Control Registers (0xN300 - 0xN3FF) |
| 0xN400 - 0xN4FF | Port N: Port Switch MAC Control Registers (0xN400 - 0xN4FF) |
| 0xN500 - 0xN5FF | Port N: Port Switch MIB Counters Registers (0xN500 - 0xN5FF) |
| 0xN600 - 0xN6FF | Port N: Port Switch ACL Control Registers (0xN600 - 0xN6FF) |
| 0xN700 - 0xN7FF | RESERVED |
| 0xN800 - 0xN8FF | Port N: Port Switch Ingress Control Registers (0xN800 - 0xN8FF) |
| 0xN900 - 0xN9FF | Port N: Port Switch Egress Control Registers (0xN900 - 0xN9FF) |
| 0xNA00 - 0xNAFF | Port N: Port Switch Queue Management Control Registers (0xNA00 - 0xNAFF) |
| 0xNB00 - 0xNBFF | Port N: Port Switch Address Lookup Control Registers (0xNB00 - 0xNBFF) |
| 0xNC00 - 0xNFFF | RESERVED |

Note: RESERVED address space must not be written under any circumstances. Failure to heed this warning may result in untoward operation and unexpected results. If it is necessary to write to registers which contain both writable and reserved bits in the same register, the user should first read back the reserved bits (RO or R/W), “OR” the desired settable bits with the value read, and then write back the “ORed” value to the register.

5.1 Global Registers

This section details the device's global registers. For an overview of the device's entire register map, refer to [Section 5.0, "Device Registers"](#). For details on the device's port registers, refer to [Section 5.2, "Port Registers"](#).

5.1.1 GLOBAL OPERATION CONTROL REGISTERS (0x0000 - 0x00FF)

5.1.1.1 Global Chip ID 0 Register

Address: 0x0000 Size: 8 bits

| Bits | Description | Type | Default |
|------|--------------------|------|---------|
| 7:0 | Fixed Value | RO | 0x00 |

5.1.1.2 Global Chip ID 1 Register

Address: 0x0001 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------|------|---------|
| 7:0 | Chip ID (MSB) | RO | 0x98 |

5.1.1.3 Global Chip ID 2 Register

Address: 0x0002 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------|------|---------|
| 7:0 | Chip ID (LSB) | RO | 0x96 |

5.1.1.4 Global Chip ID 3 Register

Address: 0x0003 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|-----------|---------|
| 7:4 | Revision ID | RO | - |
| 3:1 | RESERVED | RO | - |
| 0 | Global Software Reset Refer to the Switch Operation Register for another reset control bit. 0 = Normal operation 1 = Resets the data path and state machines, but not register values. | R/W SC | 0b |

KSZ9896C

5.1.1.5 PME Pin Control Register

Address: 0x0006 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:2 | RESERVED | RO | - |
| 1 | PME Pin Output Enable 0 = Disabled 1 = Enabled | R/W | 0b |
| 0 | PME Pin Output Polarity 0 = PME is active low 1 = PME is active high | R/W | 0b |

5.1.1.6 Global Interrupt Status Register

Address: 0x0010 - 0x0013 Size: 32 bits

This register provides the top level interrupt status for the LUE. These interrupts are enabled in the [Global Interrupt Mask Register](#). For port specific interrupts, refer to the [Port Interrupt Status Register](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 31 | Lookup Engine (LUE) Interrupt Status Refer to the Address Lookup Table Interrupt Register for detailed LUE interrupt status bits. 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 30:0 | RESERVED | RO | - |

5.1.1.7 Global Interrupt Mask Register

Address: 0x0014 - 0x0017 Size: 32 bits

This register enables the interrupts in the [Global Interrupt Status Register](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 31 | Lookup Engine (LUE) Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 30:0 | RESERVED | RO | - |

5.1.1.8 Global Port Interrupt Status Register

Address: 0x0018 - 0x001B Size: 32 bits

This register provides the top level interrupt status for the individual ports. These interrupts are enabled in the [Global Port Interrupt Mask Register](#). Refer to the [Port Interrupt Status Register](#) for detailed port interrupt status.

| Bits | Description | Type | Default |
|------|---|------|---------|
| 31:6 | RESERVED | RO | - |
| 5 | Port 6 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 4 | Port 5 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 3 | Port 4 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 2 | Port 3 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 1 | Port 2 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |
| 0 | Port 1 Interrupt Status 0 = No interrupt 1 = Interrupt request | RO | 0b |

KSZ9896C

5.1.1.9 Global Port Interrupt Mask Register

Address: 0x001C - 0x001F Size: 32 bits

This register enables the interrupts in the [Global Port Interrupt Status Register](#).

| Bits | Description | Type | Default |
|------|---|------|---------|
| 31:6 | RESERVED | RO | - |
| 5 | Port 6 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 4 | Port 5 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 3 | Port 4 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 2 | Port 3 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 1 | Port 2 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 0 | Port 1 Interrupt Mask 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |

5.1.2 GLOBAL I/O CONTROL REGISTERS (0x0100 - 0x01FF)

5.1.2.1 Serial I/O Control Register

Address: 0x0100 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:3 | RESERVED | R/W | 0100_0b |
| 2 | <p>MIIM Preamble Suppression</p> <p>This feature affects only the MIIM (MDIO / MDC) interface. When using SPI or I²C, this bit has no effect.</p> <p>0 = Normal operation. The switch always expects the MIIM preamble.</p> <p>1 = The switch will respond to MIIM commands even in the absence of a preamble.</p> | R/W | 0b |
| 1 | <p>Automatic SPI Data Out Edge Select</p> <p>When enabled, this feature automatically determines the edge of SCL that is used to clock out the SPI data on SDO. If SCL ≥ ~25MHz, SDO data is clocked by the rising edge of SCL. If SCL < ~25 MHz, SDO data is clocked by the falling edge of SCL.</p> <p>0 = The automatic feature is disabled, and bit 0 determines the SCL clock edge used for SDO.</p> <p>1 = The automatic feature is enabled, and bit 0 is ignored.</p> | R/W | 1b |
| 0 | <p>SPI Data Out Edge Select</p> <p>When bit 1 is zero, then this bit determines the clock edge used for SPI data out. When bit 1 is set to 1, this bit is ignored.</p> <p>0 = SDO data is clocked by the falling edge of SCL</p> <p>1 = SDO data is clocked by the rising edge of SCL</p> | R/W | 0b |

5.1.2.2 Output Clock Control Register

Address: 0x0103 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:5 | RESERVED | RO | 000b |
| 4:2 | RESERVED | R/W | 000b |
| 1 | <p>CLKO_25_125 Output Pin Enable</p> <p>0 = Disabled</p> <p>1 = Enabled</p> | R/W | 1b |
| 0 | <p>CLKO_25_125 Frequency</p> <p>0 = 25 MHz</p> <p>1 = 125 MHz</p> | R/W | 0b |

KSZ9896C

5.1.2.3 In-Band Management (IBA) Control Register

Address: 0x0104 - 0x0107 Size: 32 bits

This register controls the In-Band Access (IBA) feature.

| Bits | Description | Type | Default |
|-------|--|-----------|--------------------------|
| 31 | IBA Enable The initial value is strapped in from the RX_DV6/CRS_DV6/RX_CTL6 pin. 0 = Disabled 1 = Enabled | R/W | Note 5-1 |
| 30 | IBA Destination MAC Address Match Enable Set this bit to enable checking of the destination MAC address in received IBA frames against the switch MAC address in the Switch MAC Address 0 Register through Switch MAC Address 5 Register . Non-matching frames are discarded. When not enabled, the MAC address is not checked. | R/W | 0b |
| 29 | IBA Reset Set this bit to initialize the IBA state machine. This bit is self-clearing. | R/W SC | 0b |
| 28:24 | RESERVED | RO | 0x00 |
| 23:22 | Priority Queue for IBA response Specifies the transmit priority queue for the IBA response frame. Typically this value is not changed. | R/W | 01b |
| 21:19 | RESERVED | RO | 00_0b |
| 18:16 | Port used for IBA communication 000 = Port 1 001 = Port 2 010 = Port 3 011 = Port 4 100 = Port 5 101 = Port 6 110 = Reserved 111 = Reserved | R/W | 101b |
| 15:0 | TPID (EtherType) value for IBA frame header | R/W | 0x40FE |

Note 5-1 The default value of this field is determined by the associated configuration strap value. Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional information.

5.1.2.4 I/O Drive Strength Register

Address: 0x010D Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | R/W | 0b |
| 6:4 | High Speed Drive Strength (24mA) Controls drive strength of GMII / RGMII / MII / RMII (except TX_CLK / REFCLKI, COL and CRS) and CLKO_25_125. | R/W | 110b |
| 3 | RESERVED | R/W | 0b |
| 2:0 | Low Speed Drive Strength (8mA) Controls drive strength of TX_CLK / REFCLKI, COL, CRS, LEDs, PME_N, INTRP_N, SDO and SDI/SDA/MDIO. | R/W | 10b |

5.1.2.5 In-Band Management (IBA) Operation Status 1 Register

Address: 0x0110 - 0x0113 Size: 32 bits

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31 | Good IBA Packet Detect 1 = A good IBA packet is received. | RO | 0b |
| 30 | IBA Response Packet Transmit Done 1 = An IBA response packet is sent out. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 29 | IBA Execution Done 1 = All the commands in one IBA packet are completely executed. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 28:15 | RESERVED | RO | 0x0000 |
| 14 | IBA MAC Address Mismatch Error This bit is active only when IBA_ENABLE (In-Band Management (IBA) Control Register, bit 30) is set. 1 = An IBA packet is received with an unmatched MAC address, unequal to the switch's MAC address. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 13 | IBA Access Format Error 1 = An IBA packet with a wrong access format (not equal to 0x9800) is received. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 12 | IBA Access Code Error 1 = An IBA packet with an unrecognized access code is received. (Valid access codes are 0x0001 and 0x0002.) This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |

KSZ9896C

| Bits | Description | Type | Default |
|------|--|------|---------|
| 11 | IBA Access Command Error 1 = An IBA packet with an unrecognized command code is received. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 10 | IBA Oversize Packet Error 1 = An oversized IBA packet is received. The maximum IBA packet size is 320 bytes, including 8-byte zeros before FCS and the 4-byte FCS. No response packet is sent. This bit is cleared when a packet with a matching IBA tag field is received. | RO | 0b |
| 9:7 | RESERVED | RO | 000b |
| 6:0 | IBA Access Code Error Location When IBA Access Command Error (bit 11) is set, these bits indicate the address location of the wrong command code within the IBA packet. | RO | 0x000 |

5.1.2.6 LED Override Register

Address: 0x0120 - 0x0123 Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|------------|
| 31:10 | RESERVED | RO | 0x00000 |
| 9:0 | Override LED These bits select whether each LED_{x_0} and LED_{x_1} pin will function as an LED or General Purpose Output (GPO). The LSB bit of this field represents LED1_0 , followed by LED1_1 , LED2_0 , etc. When configured as a GPO, the GPO output is controlled via the LED Output Register . 0 = LED_{x_y} pin functions as an LED 1 = LED_{x_y} pin functions as a GPO | | 000000000b |

5.1.2.7 LED Output Register

Address: 0x0124 - 0x0127 Size: 32 bits

| Bits | Description | Type | Default |
|-------|---|------|------------|
| 31:10 | RESERVED | RO | 0x00000 |
| 9:0 | GPO Output Control When configured as a GPO via the LED Override Register , the GPO output is controlled via this field. The LSB bit of this field represents LED1_0 , followed by LED1_1 , LED2_0 , etc. 0 = LED_{x_y} pin outputs low 1 = LED_{x_y} pin outputs high | R/W | 000000000b |

5.1.3 GLOBAL PHY CONTROL AND STATUS REGISTERS (0x0200 - 0x02FF)

5.1.3.1 Power Down Control 0 Register

Address: 0x0201 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:6 | RESERVED | RO | 00b |
| 5 | PLL Power Down 0 = Normal operation. 1 = Disable PLL. This may be used in combination with EDPD mode – see below. | R/W | 0b |
| 4:3 | Power Management Mode 00 = Normal operation 01 = Energy Detect Power Down (EDPD) Mode 10 = Soft Power Down Mode 11 = invalid | R/W | 00b |
| 2:0 | RESERVED | RO | 000b |

5.1.3.2 LED Configuration Strap Register

Address: 0x0210 - 0x0213 Size: 32 bits

| Bits | Description | Type | Default |
|-------|---|------|--------------------------|
| 31:10 | RESERVED | RO | 0x000000 |
| 9:0 | Configuration strap values of LED pins [LED5_1, LED5_0, LED4_1, LED4_0, LED3_1, LED3_0, LED2_1, LED2_0, LED1_1, LED1_0] | RO | Note 5-2 |

Note 5-2 The default value of this field is determined by the associated configuration strap values. Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional information.

KSZ9896C

5.1.4 GLOBAL SWITCH CONTROL REGISTERS (0x0300 - 0x03FF)

5.1.4.1 Switch Operation Register

Address: 0x0300 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|-----------|--------------------------|
| 7 | Double Tag Enable 1 = Double tagging is enabled 0 = Double tagging is disabled | R/W | 0b |
| 6:2 | RESERVED | RO | 0x00 |
| 1 | Soft Hardware Reset When set to 1, all register settings, except configuration strap options, are reset to default values. | R/W SC | 0b |
| 0 | Start Switch 1 = Switch function is enabled 0 = Switch function is disabled; no traffic will be passed until this bit is set | R/W | Note 5-3 |

Note 5-3 The default value of this field is determined by the **LED5_1** configuration strap value. Refer to [Section 3.2.1, "Configuration Straps," on page 17](#) for additional information.

5.1.4.2 Switch MAC Address 0 Register

Address: 0x0302 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:0 | MAC Address [47:40] This register, along with the Switch MAC Address 1-5 Registers, define the switch's MAC address to be used as the source address in MAC pause control frames, and for self-address filtering. | R/W | 0x00 |

5.1.4.3 Switch MAC Address 1 Register

Address: 0x0303 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [39:32] | R/W | 0x10 |

5.1.4.4 Switch MAC Address 2 Register

Address: 0x0304 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [31:24] | R/W | 0xA1 |

5.1.4.5 Switch MAC Address 3 Register

Address: 0x0305 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [23:16] | R/W | 0xFF |

5.1.4.6 Switch MAC Address 4 Register

Address: 0x0306 Size: 8 bits

| Bits | Description | Type | Default |
|------|---------------------------|------|---------|
| 7:0 | MAC Address [15:8] | R/W | 0xFF |

5.1.4.7 Switch MAC Address 5 Register

Address: 0x0307 Size: 8 bits

| Bits | Description | Type | Default |
|------|--------------------------|------|---------|
| 7:0 | MAC Address [7:0] | R/W | 0xFF |

KSZ9896C

5.1.4.8 Switch Maximum Transmit Unit Register

Address: 0x0308 - 0x0309 Size: 16 bits

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 15:14 | RESERVED | R/W | 00b |
| 13:0 | Maximum Frame Length (MTU) Specifies the maximum transmission unit (MTU), which is the maximum frame payload size. Frames which exceed this maximum are truncated. This value can be set as high as 9000 (= 0x2328) if jumbo frame support is required. Also refer to the Switch MAC Control 1 Register and Port MAC Control 0 Register . | R/W | 0x07D0 |

5.1.4.9 Switch ISP TPID Register

Address: 0x030A - 0x030B Size: 16 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 15:0 | ISP Tag TPID Default tag TPID (EtherType) for untagged incoming frames or the ISP frame tag TPID for the double tagging function. | R/W | 0x9100 |

5.1.4.10 Switch Lookup Engine Control 0 Register

Address: 0x0310 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | 802.1Q VLAN Enable This is the master enable for VLAN forwarding and filtering. Note that the VLAN Table must be set up before VLAN mode is enabled. 1 = VLAN mode enabled 0 = VLAN mode disabled | R/W | 0b |
| 6 | Drop Invalid VID 1 = All received packets with invalid VLAN ID are dropped. 0 = Received packets with invalid VLAN ID are forwarded to the host port. Note that the Unknown VID Forwarding feature (Unknown VLAN ID Control Register), if enabled, takes precedence over this bit. | R/W | 1b |
| 5:3 | Age Count This bit, in combination with the Age Period value (Switch Lookup Engine Control 3 Register), determines the aging time of dynamic entries in the address lookup table. This value is used for the Age Count field whenever a dynamic table entry is updated. | R/W | 10_0b |

| Bits | Description | Type | Default |
|------|--|------|---------|
| 2 | Reserved Multicast Lookup Enable 1 = Enable Reserved Multicast Table 0 = Disable Reserved Multicast Table | R/W | 0b |
| 1:0 | HASH_OPTION Defines the hashing option for mapping entries to the dynamic lookup table. 00, 11 = Entry is mapped directly using the 10 least significant bits of the destination address. 01 = The CRC hashing function is used. 10 = The XOR hashing function is used. Refer to Section 4.4.2.1, "Address Lookup (ALU) Table," on page 29 for additional information. | R/W | 01b |

5.1.4.11 Switch Lookup Engine Control 1 Register

Address: 0x0311 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|-----------|---------|
| 7 | Unicast Learning Disable 1 = Unicast address learning is disabled 0 = Unicast address learning is enabled | R/W | 0b |
| 6 | Self-Address Filtering – Global Enable The source address of received packets is compared to the MAC address in registers Switch MAC Address 0 Register through Switch MAC Address 5 Register , and the packet is dropped if there is a match. Self-address filtering can be enabled on a port-by-port basis by setting the port enable bit in the Port Control 2 Register in addition to setting this bit. 1 = Enable self-address filtering globally for those ports whose port enable bit (Port Control 2 Register) is set. 0 = Do not filter self-addressed packets on any port. | R/W | 0b |
| 5 | Flush Address Lookup Table The Flush Option bit in the Switch Lookup Engine Control 2 Register determines whether flushing is performed on dynamic entries, static entries, or both. 1 = Trigger a flush of the entire address lookup table. The static address table is not flushed. 0 = Normal operation | R/W SC | 0b |
| 4 | Flush MSTP Address Entries (Address Lookup Table) The Flush Option bit in the Switch Lookup Engine Control 2 Register determines whether flushing is performed on dynamic entries, static entries, or both. 1 = Trigger a flush of the matched MSTP entries 0 = Normal operation | R/W SC | 0b |
| 3 | Multicast Source Address Filtering 1 = Forward packets with a multicast source address 0 = Drop packets with a multicast source address | R/W | 1b |

KSZ9896C

| Bits | Description | Type | Default |
|------|---|------|---------|
| 2 | Aging Enable 1 = Enable address table aging 0 = Disable address table aging | R/W | 1b |
| 1 | Fast Aging 1 = Enable fast aging 0 = Disable fast aging | R/W | 0b |
| 0 | Link Down Flush 1 = Link down will cause the entries of any link down port to be flushed 0 = Link down flush is disabled | R/W | 0b |

5.1.4.12 Switch Lookup Engine Control 2 Register

Address: 0x0312 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | R/W | 0b |
| 6 | Double Tag Multicast Trap 1 = When double tagging mode is enabled, forward all reserved multicast packets to the host port only. 0 = Normal forwarding | R/W | 0b |
| 5 | Dynamic Entry Egress VLAN Filtering Egress VLAN filtering uses the forwarding port map from the VLAN table to restrict the forwarding ports determined from the address lookup. This is the recommended mode of operation when VLAN is enabled. The default value is 0 only for backwards compatibility with previous switches. 1 = Enable. For successful lookup of a dynamic entry in the address table, the forwarding ports are determined from the AND function of the address table port map and the VLAN table port map. 0 = Disable. For successful lookup of a dynamic entry in the address table, the forwarding ports are determined from the address table only. | R/W | 0b |
| 4 | Static Entry Egress VLAN Filtering Egress VLAN filtering uses the forwarding port map from the VLAN table to restrict the forwarding ports determined from the address lookup. This is the recommended mode of operation when VLAN is enabled. The default value is 0 only for backwards compatibility with previous switches. 1 = Enable. For successful lookup of a static entry in the address table, the forwarding ports are determined from the AND function of the address table port map and the VLAN table port map. 0 = Disable. For successful lookup of a static entry in the address table, the forwarding ports are determined from the address table only. | R/W | 0b |

| Bits | Description | Type | Default |
|------|---|------|---------|
| 3:2 | Flush Option Determines which address lookup table entries may be flushed by either of the flush operations in the Switch Lookup Engine Control 1 Register . 00 = No flush or flush is done 01 = Flush only dynamic table entries 10 = Flush only static table entries 11 = Flush both static and dynamic table entries | R/W | 00b |
| 1:0 | MAC Address Priority 00 = MAC Address (MACA) priority for a packet is determined from the destination address (DA) lookup 01 = MACA priority for a packet is determined from the source address (SA) lookup 10 = MACA priority for a packet is determined from the higher of the DA and SA lookups 11 = MACA priority for a packet is determined from the lower of the DA and SA lookups | R/W | 00b |

5.1.4.13 Switch Lookup Engine Control 3 Register

Address: 0x0313 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:0 | Age Period This value, multiplied by the Age Count value in the entries of the Address Lookup Table, determines the aging time of dynamic entries in that table. The unit is seconds. | R/W | 0x4B |

KSZ9896C

5.1.4.14 Address Lookup Table Interrupt Register

Address: 0x0314 Size: 8 bits

This register provides the detailed interrupt status for the Address Lookup Table. These interrupts are enabled in the [Address Lookup Table Mask Register](#). The LUE interrupt status bit in the [Global Interrupt Status Register](#) is the OR of the status bits in this register.

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:3 | RESERVED | RO | 0x00 |
| 2 | Learn Fail Interrupt Status An Address Lookup Table entry was not learned because all entries in the bucket are static | R/WC | 0b |
| 1 | Almost Full Interrupt Status Interrupt indicates that the Address Lookup Table bucket was almost full (2 or 3 valid entries) when a new static entry was written. | R/WC | 0b |
| 0 | Write Fail Interrupt Status Interrupt indicates that the Address Lookup Table bucket is full and a write failed | R/WC | 0b |

5.1.4.15 Address Lookup Table Mask Register

Address: 0x0315 Size: 8 bits

This register masks the Address Lookup Table interrupts in the [Address Lookup Table Interrupt Register](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:3 | RESERVED | RO | 0x00 |
| 2 | Learn Fail Interrupt Mask 1 = Interrupt is disabled 0 = Interrupt is enabled | R/W | 1b |
| 1 | Almost Full Interrupt Mask 1 = Interrupt is disabled 0 = Interrupt is enabled | R/W | 1b |
| 0 | Write Fail Interrupt Mask 1 = Interrupt is disabled 0 = Interrupt is enabled | R/W | 1b |

5.1.4.16 Address Lookup Table Entry Index 0 Register

Address: 0x0316 - 0x0317 Size: 16 bits

| Bits | Description | Type | Default |
|------------|--|------|---------|
| 15:12 | RESERVED | RO | 0x0 |
| 11:0 / 9:0 | <p>Almost Full Entry Index [11:0] When a static entry is successfully written into the Address Lookup Table, but the table bucket is almost full (contains 2 or 3 static entries prior to the write), the entry address is reported here.</p> <p>Fail Write Index [9:0] When a static entry write failure occurs in the Address Lookup Table, the bucket address is reported here.</p> | RO | 0x000 |

5.1.4.17 Address Lookup Table Entry Index 1 Register

Address: 0x0318 - 0x0319 Size: 16 bits

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15:10 | RESERVED | RO | 0000_00 |
| 9:0 | <p>Fail Learn Index When a destination address fails to be learned in the Address Lookup Table because the bucket contains 4 static entries, the bucket address is reported here.</p> | RO | 0x000 |

5.1.4.18 Address Lookup Table Entry Index 2 Register

Address: 0x031A - 0x031B Size: 16 bits

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15:10 | RESERVED | RO | 0000_00 |
| 9:0 | <p>CPU Access Index Whenever there is an external read or write to the Address Lookup Table, the bucket address of the access is reported here.</p> | RO | 0x000 |

KSZ9896C

5.1.4.19 Unknown Unicast Control Register

Address: 0x0320 - 0x0323 Size: 32 bits

The following three registers control forwarding of packets with 1) unknown unicast destination address, 2) unknown multicast destination address, and 3) unknown VLAN ID.

If a received packet falls into more than one of these categories, the precedence is:

1. Unknown VID
2. Unknown Unicast
3. Unknown Multicast

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31 | Unknown Unicast Packet Forward 1 = Enable forwarding of unknown unicast packets to the ports specified below 0 = Disable unknown unicast packet forwarding | R/W | 0b |
| 30:6 | RESERVED | RO | 0x000000 |
| 5:0 | Unknown Unicast Forwarding Ports Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward unknown unicast packets to that port 0 = Do not forward to that port All ones = Forwarded to all ports All zeros = Forwarded to no ports | R/W | 00_0000b |

5.1.4.20 Unknown Multicast Control Register

Address: 0x0324 - 0x0327 Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31 | Unknown Multicast Packet Forward 1 = Enable forwarding of unknown multicast packets to the ports specified below 0 = Disable unknown multicast packet forwarding | R/W | 0b |
| 30:6 | RESERVED | RO | 0x000000 |
| 5:0 | Unknown Multicast Forwarding Ports Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward unknown multicast packets to that port 0 = Do not forward to that port All ones = Forwarded to all ports All zeros = Forwarded to no ports | R/W | 00_0000b |

5.1.4.21 Unknown VLAN ID Control Register

Address: 0x0328 - 0x032B Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31 | Unknown VID Packet Forward 1 = Enable forwarding of unknown VLAN ID (VID) packets to the ports specified below 0 = Disable unknown VID packet forwarding | R/W | 0b |
| 30:6 | RESERVED | RO | 0x000000 |
| 5:0 | Unknown VID Forwarding Ports Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward unknown VID packets to that port 0 = Do not forward to that port All ones = Forwarded to all ports All zeros = Forwarded to no ports | R/W | 00_0000b |

5.1.4.22 Switch MAC Control 0 Register

Address: 0x0330 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | Alternate Back-off Mode The back-off mode applies to half-duplex only. This bit should be set if the No Excessive Collision Drop bit in the Switch MAC Control 1 Register is enabled. 1 = Enable alternate back-off mode 0 = Disable | R/W | 0b |
| 6:4 | RESERVED | R/W | 000b |
| 3 | Frame Length Field Check This applies only when the EtherType/Length field is <1500. 1 = Discard any packet if the actual packet length does not match the frame length field. 0 = Do not check the packet length. | R/W | 0b |
| 2 | RESERVED | R/W | 1b |

KSZ9896C

| Bits | Description | Type | Default |
|------|---|------|---------|
| 1 | <p>Flow Control Packet Drop Mode</p> <p>This bit controls which flow control packets may be forwarded or dropped. To enable forwarding of all flow control packets, refer to bit 0 of the Switch MAC Control 4 Register. Enabling of flow control is managed in the PHY Auto-Negotiation Advertisement Register (for PHY ports) and the XMII Port Control 0 Register (for MAC ports).</p> <p>1 = The switch will drop received packets with either EtherType = 0x8808 or destination address (DA) = 01-80-C2-00-00-01.</p> <p>0 = The switch will drop received packets with both EtherType = 0x8808 and DA = 01-80-C2-00-00-01.</p> | R/W | 0b |
| 0 | <p>Aggressive Back-off Enable</p> <p>For use with half-duplex back pressure. This is not an IEEE standard.</p> <p>1 = Enable aggressive back-off algorithm in half-duplex mode to enhance performance.</p> <p>0 = Disable</p> | R/W | 0b |

5.1.4.23 Switch MAC Control 1 Register

Address: 0x0331 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | R/W | 1b |
| 6 | <p>Multicast Storm Protection Disable</p> <p>1 = Multicast packets are not included in Broadcast Storm Protection</p> <p>0 = Multicast packets are included in Broadcast Storm Protection</p> | R/W | 1b |
| 5 | <p>Back Pressure Mode</p> <p>1 = Use carrier sense based back pressure mode</p> <p>0 = Use collision based back pressure mode. (This is the recommended mode.)</p> | R/W | 1b |
| 4 | <p>Flow Control and Back Pressure Fair Mode</p> <p>1 = Enable fair mode. If a flow controlled ingress port and a non-flow controlled ingress port forward traffic to the same egress port, packets from the non-flow controlled port may be dropped.</p> <p>0 = Disable fair mode. In this scenario the flow controlled port will be flow controlled, and the non-flow controlled port will be neither flow controlled nor packets dropped.</p> | R/W | 1b |
| 3 | <p>No Excessive Collision Drop</p> <p>If this bit is set, Alternate Back-odd Mode (bit 7 in the Switch MAC Control 0 Register) should also be set.</p> <p>1 = The switch will not drop packets when 16 or more collisions occur.</p> <p>0 = The switch will drop packets when 16 or more collisions occur</p> | R/W | 0b |

| Bits | Description | Type | Default |
|------|--|------|---------|
| 2 | <p>Jumbo Packet Support</p> <p>The programmable packet payload size limit is specified in register 0x0308 – 0x0309, up to a maximum of 9000 bytes. This bit overrides bit 1 of this register.</p> <p>1 = Enable support for jumbo packets 0 = Disable</p> | R/W | 0b |
| 1 | <p>Legal Maximum Packet Size Check Disable</p> <p>1 = Accept packets up to 2000 bytes in size. 0 = Accept only standard size packets, up to 1522 bytes for tagged packets, or 1518 bytes for untagged packets. Larger packets will be dropped.</p> | R/W | 0b |
| 0 | <p>Pass Short Packet</p> <p>1 = Accept packets between 32 and 64 bytes in size. 0 = Accept only standard size packets, at least 64 bytes. Smaller packets are dropped.</p> | R/W | 0b |

5.1.4.24 Switch MAC Control 2 Register

Address: 0x0332 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:4 | RESERVED | R/W | 0x0 |
| 3 | <p>Null VID Replacement</p> <p>Describes the behavior when a packet is received with a null (zero) VID.</p> <p>1 = Replace a null VID with the Port Default VID as defined in the Port Default Tag 0 Register and Port Default Tag 1 Register 0 = No replacement of null VID</p> | R/W | 0b |
| 2:0 | <p>Broadcast Storm Protection Rate bits [10:8]</p> <p>The remainder of this fields bits are in the Switch MAC Control 3 Register.</p> <p>Multiply this value by 64 to determine how many bytes of packet data are allowed on an input port in a preset period. The period is 5ms for a 1000Mbps port, 50ms for a 100Mbps port, or 500ms for a 10Mbps. The default is 1%.</p> | R/W | 000b |

KSZ9896C

5.1.4.25 Switch MAC Control 3 Register

Address: 0x0333 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:0 | Broadcast Storm Protection Rate bits [7:0] The remainder of this fields bits are in the Switch MAC Control 2 Register . Multiply this value by 64 to determine how many bytes of packet data are allowed on an input port in a preset period. The period is 5ms for a 1000Mbps port, 50ms for a 100Mbps port, or 500ms for a 10Mbps. The default is 1%. | R/W | 0x4A |

5.1.4.26 Switch MAC Control 4 Register

Address: 0x0334 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|-----------|
| 7:1 | RESERVED | RO | 0000_000b |
| 0 | Pass Flow Control Packets 1 = Switch will forward 802.3x PAUSE flow control frames. 0 = Switch will filter PAUSE frames. | R/W | 0b |

5.1.4.27 Switch MAC Control 5 Register

Address: 0x0335 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6 | RESERVED | R/W | 0b |
| 5:4 | Ingress Rate Limit Period 00 = 16ms 01 = 64ms 1x = 256ms | R/W | 01b |
| 3 | Queue Based Egress Rate Limit Enable 1 = Queue-based egress rate limiting 0 = Port-based egress rate limiting | R/W | 0b |
| 2:0 | RESERVED | RO | 000b |

5.1.4.28 Switch MIB Control Register

Address: 0x0336 Size: 8 bits

MIB counters are provided on a per-port basis. They are read and controlled via the [Port N: Port Switch MIB Counters Registers \(0xN500 - 0xN5FF\)](#).

| Bits | Description | Type | Default |
|------|---|-----------|----------|
| 7 | Flush MIB Counters 1 = Flush all MIB counters of enabled ports. Refer to the Port MIB Control and Status Register . 0 = Normal counter operation | R/W SC | 0b |
| 6 | Freeze MIB Counters 1 = Freeze MIB counters of enabled ports. Refer to the Port MIB Control and Status Register . 0 = Normal counter operation | R/W | 0b |
| 5:0 | RESERVED | RO | 00_0000b |

5.1.4.29 802.1p Priority Mapping 0 Register

Address: 0x0338 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When the tag PCP field is 0x1, use this value for priority | R/W | 001b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When the tag PCP field is 0x0, use this value for priority | R/W | 000b |

5.1.4.30 802.1p Priority Mapping 1 Register

Address: 0x0339 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When the tag PCP field is 0x3, use this value for priority | R/W | 011b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When the tag PCP field is 0x2, use this value for priority | R/W | 010b |

KSZ9896C

5.1.4.31 802.1p Priority Mapping 2 Register

Address: 0x033A Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When the tag PCP field is 0x5, use this value for priority | R/W | 101b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When the tag PCP field is 0x4, use this value for priority | R/W | 100b |

5.1.4.32 802.1p Priority Mapping 3 Register

Address: 0x033B Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When the tag PCP field is 0x7, use this value for priority | R/W | 111b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When the tag PCP field is 0x6, use this value for priority | R/W | 110b |

5.1.4.33 IP DiffServ Priority Enable Register

Address: 0x033E Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 7:2 | RESERVED | RO | 0000_00b |
| 1 | RESERVED | R/W | 0b |
| 0 | DiffServ Priority Remap Enable 1 = Use the following registers to remap the DSCP (DiffServ) priority to a 3-bit priority value 0 = Use DSCP bits [5:3] for priority | R/W | 0b |

5.1.4.34 IP DiffServ Priority Mapping 0 Register

Address: 0x0340 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x01, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x00, use this value for priority | R/W | 000b |

5.1.4.35 IP DiffServ Priority Mapping 1 Register

Address: 0x0341 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x03, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x02, use this value for priority | R/W | 000b |

5.1.4.36 IP DiffServ Priority Mapping 2 Register

Address: 0x0342 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x05, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x04, use this value for priority | R/W | 000b |

5.1.4.37 IP DiffServ Priority Mapping 3 Register

Address: 0x0343 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x07, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x06, use this value for priority | R/W | 000b |

KSZ9896C

5.1.4.38 IP DiffServ Priority Mapping 4 Register

Address: 0x0344 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x09, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x08, use this value for priority | R/W | 000b |

5.1.4.39 IP DiffServ Priority Mapping 5 Register

Address: 0x0345 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x0B, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x0A, use this value for priority | R/W | 000b |

5.1.4.40 IP DiffServ Priority Mapping 6 Register

Address: 0x0346 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x0D, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x0C, use this value for priority | R/W | 000b |

5.1.4.41 IP DiffServ Priority Mapping 7 Register

Address: 0x0347 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x0F, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x0E, use this value for priority | R/W | 000b |

5.1.4.42 IP DiffServ Priority Mapping 8 Register

Address: 0x0348 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x11, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x10, use this value for priority | R/W | 000b |

5.1.4.43 IP DiffServ Priority Mapping 9 Register

Address: 0x0349 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x13, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x12, use this value for priority | R/W | 000b |

5.1.4.44 IP DiffServ Priority Mapping 10 Register

Address: 0x034A Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x15, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x14, use this value for priority | R/W | 000b |

5.1.4.45 IP DiffServ Priority Mapping 11 Register

Address: 0x034B Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x17, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x16, use this value for priority | R/W | 000b |

KSZ9896C

5.1.4.46 IP DiffServ Priority Mapping 12 Register

Address: 0x034C Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x19, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x18, use this value for priority | R/W | 000b |

5.1.4.47 IP DiffServ Priority Mapping 13 Register

Address: 0x034D Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x1B, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x1A, use this value for priority | R/W | 000b |

5.1.4.48 IP DiffServ Priority Mapping 14 Register

Address: 0x034E Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x1D, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x1C, use this value for priority | R/W | 000b |

5.1.4.49 IP DiffServ Priority Mapping 15 Register

Address: 0x034F Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x1F, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x1E, use this value for priority | R/W | 000b |

5.1.4.50 IP DiffServ Priority Mapping 16 Register

Address: 0x0350 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x21, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x20, use this value for priority | R/W | 000b |

5.1.4.51 IP DiffServ Priority Mapping 17 Register

Address: 0x0351 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x23, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x22, use this value for priority | R/W | 000b |

5.1.4.52 IP DiffServ Priority Mapping 18 Register

Address: 0x0352 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x25, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x24, use this value for priority | R/W | 000b |

5.1.4.53 IP DiffServ Priority Mapping 19 Register

Address: 0x0353 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x27, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x26, use this value for priority | R/W | 000b |

KSZ9896C

5.1.4.54 IP DiffServ Priority Mapping 20 Register

Address: 0x0354 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x29, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x28, use this value for priority | R/W | 000b |

5.1.4.55 IP DiffServ Priority Mapping 21 Register

Address: 0x0355 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x2B, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x2A, use this value for priority | R/W | 000b |

5.1.4.56 IP DiffServ Priority Mapping 22 Register

Address: 0x0350 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x2D, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x2C, use this value for priority | R/W | 000b |

5.1.4.57 IP DiffServ Priority Mapping 23 Register

Address: 0x0357 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x2F, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x2E, use this value for priority | R/W | 000b |

5.1.4.58 IP DiffServ Priority Mapping 24 Register

Address: 0x0358 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x31, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x30, use this value for priority | R/W | 000b |

5.1.4.59 IP DiffServ Priority Mapping 25 Register

Address: 0x0359 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x33, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x32, use this value for priority | R/W | 000b |

5.1.4.60 IP DiffServ Priority Mapping 26 Register

Address: 0x035A Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x35, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x34, use this value for priority | R/W | 000b |

5.1.4.61 IP DiffServ Priority Mapping 27 Register

Address: 0x035B Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x37, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x36, use this value for priority | R/W | 000b |

KSZ9896C

5.1.4.62 IP DiffServ Priority Mapping 28 Register

Address: 0x035C Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x39, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x38, use this value for priority | R/W | 000b |

5.1.4.63 IP DiffServ Priority Mapping 29 Register

Address: 0x035D Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x3B, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x3A, use this value for priority | R/W | 000b |

5.1.4.64 IP DiffServ Priority Mapping 30 Register

Address: 0x035E Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x3D, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x3C, use this value for priority | R/W | 000b |

5.1.4.65 IP DiffServ Priority Mapping 31 Register

Address: 0x035F Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:4 | When IPv4 / IPv6 DSCP field is 0x3F, use this value for priority | R/W | 000b |
| 3 | RESERVED | RO | 0b |
| 2:0 | When IPv4 / IPv6 DSCP field is 0x3E, use this value for priority | R/W | 000b |

5.1.4.66 Global Port Mirroring and Snooping Control Register

Address: 0x0370 Size: 8 bits

This register contains global controls for port mirroring and IGMP and MLD snooping. Mirroring also requires additional register settings for the individual ports. Refer to the [Port N: Port Switch Ingress Control Registers \(0xN800 - 0xN8FF\)](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6 | IGMP Snooping Enable (IPv4) 1 = IGMP snooping is enabled. All IGMP packets will be forwarded to the host port. 0 = IGMP snooping is disabled. | R/W | 0b |
| 5:4 | RESERVED | RO | 00b |
| 3 | MLD Snooping Option 1 = Snooping criteria: IPv6 next header = 43, 44, 50, 51 or 60; or next header = 0 and hop-by-hop next header = 43, 44, 50, 51 or 60. 0 = Snooping criteria: IPv6 next header = 1 or 58; or next header = 0 and hop-by-hop next header = 1 or 58. | R/W | 0b |
| 2 | MLD Snooping Enable (IPv6) 1 = MLD snooping is enabled. All MLD packets will be forwarded to the host port. 0 = MLD snooping is disabled. | R/W | 0b |
| 1 | RESERVED | RO | 0b |
| 0 | Sniff Mode Select 1 = Sniffing filter is "Rx AND Tx". Both the source port and destination ports need to match. 0 = Sniffing filter is "Rx OR Tx". Either the source port or the destination port needs to match. This mode is used to implement Rx-only sniffing. | R/W | 0b |

5.1.4.67 WRED DiffServ Color Mapping Register

Address: 0x0378 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:6 | RESERVED | RO | 00b |
| 5:4 | Red This field specifies the red Differentiated Services Code Point (DSCP) value. | R/W | 11b |
| 3:2 | Yellow This field specifies the yellow DSCP value. | R/W | 10b |
| 1:0 | Green This field specifies the green DSCP value. | R/W | 01b |

KSZ9896C

5.1.4.68 Queue Management Control 0 Register

Address: 0x0390 - 0x0393 Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31:8 | RESERVED | RO | 0x000000 |
| 7:6 | Priority_2Q When the 2 queue configuration is selected, this determines how to map the 2-bit priority regeneration result from the Port Priority to Queue Mapping Register into 2 queues. 00 = Priorities 0, 1, 2 map to the Low priority queue. Priority 3 maps to the High priority queue. 01 = Not used. 10 = Priorities 0, 1 map to the Low priority queue. Priorities 2, 3 map to the High priority queue. 11 = Priority 0 maps to the Low priority queue. Priorities 1, 2, 3 map to the High priority queue. | R/W | 10b |
| 5:2 | RESERVED | R/W | 00_00b |
| 1 | Unicast Port VLAN Membership Discard This bit applies to the Port VLAN Membership function in registers 0xNA04-07. 1 = All frames are restricted to the forwarding ports as defined in the Port Control 1 Register . 0 = Frames forwarding to a single destination port are not limited to the forwarding ports defined in the Port Control 1 Register . Note that when mirroring is enabled, a single-destination frames will be dropped if it is mirrored to another port. | R/W | 1b |
| 0 | RESERVED | R/W | 0b |

5.1.5 GLOBAL SWITCH LOOK UP ENGINE (LUE) CONTROL REGISTERS (0x0400 - 0x04FF)

The following registers are used for accessing the VLAN Table, Address Lookup Table, Static Address Table and the Reserved Multicast Address Table. The organizations of these tables, and instructions for accessing them, are provided in [Section 5.3, "Tables and MIB Counters \(Access\)," on page 155](#).

5.1.5.1 VLAN Table Entry 0 Register

Address: 0x0400 - 0x403 Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|--------------|
| 31 | VALID This field in the VLAN Table specifies if the table entry is valid. 1 = Table entry is valid 0 = Table entry is invalid | R/W | 0b |
| 30:28 | RESERVED | RO | 000b |
| 27 | FORWARD OPTION This field in the VLAN Table specifies how the forwarding ports are determined. 1 = Forward to the VLAN Table port map (PORT FORWARD field). 0 = Forwarding ports are determined by other variables. It may be the ALU port map, the VLAN port map, the combination of the two, or the Unknown Unicast/Multicast feature. | R/W | 0b |
| 26:24 | PRIORITY This field in the VLAN Table specifies the priority level. | R/W | 000b |
| 23:15 | RESERVED | RO | 0000_0000_0b |
| 14:12 | MSTP INDEX This field in the VLAN Table specifies the Multiple Spanning Tree Protocol index. | R/W | 000b |
| 11:7 | RESERVED | RO | 0000_0b |
| 6:0 | FID This field in the VLAN Table specifies the Filter ID. The FID value is normally combined with the destination address and then hashed to index the address lookup table. | R/W | 000_0000b |

Note: Refer to [Section 5.3.4, "VLAN Table," on page 162](#) for additional information on VLAN tables.

KSZ9896C

5.1.5.2 VLAN Table Entry 1 Register

Address: 0x0404 - 0x407 Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|-----------|
| 31:6 | RESERVED | R/O | 0x0000000 |
| 5:0 | PORT UNTAG This field in the VLAN Table specifies the untagging policy for each egress port. Bits [5:0] correspond to ports [6:1]. 1 = Untag packets upon egress at this port 0 = Do not untag upon egress at this port | R/W | 0x00 |

Note: Refer to [Section 5.3.4, "VLAN Table," on page 162](#) for additional information on VLAN tables.

5.1.5.3 VLAN Table Entry 2 Register

Address: 0x0408 - 0x040B Size: 32 bits

| Bits | Description | Type | Default |
|------|---|------|-----------|
| 31:6 | RESERVED | R/O | 0x0000000 |
| 5:0 | PORT FORWARD This field specifies the forwarding policy to each port. The policy is applied if the FO bit is set. Bits [5:0] correspond to ports [6:1]. 1 = Forward to this port 0 = Do not forward to this port | R/W | 0x00 |

Note: Refer to [Section 5.3.4, "VLAN Table," on page 162](#) for additional information on VLAN tables.

5.1.5.4 VLAN Table Index Register

Address: 0x040C - 0x040D Size: 16 bits

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15:12 | RESERVED | RO | 0x0 |
| 11:0 | VLAN Index This value addresses the VLAN table, in combination with the VLAN table action in the following register. | R/W | 0x000 |

Note: Refer to [Section 5.3.4, "VLAN Table," on page 162](#) for additional information on VLAN tables.

5.1.5.5 VLAN Table Access Control Register

Address: 0x040E Size: 8 bits

| Bits | Description | Type | Default |
|------|--|-----------|---------|
| 7 | Start VLAN Table Action 1 = Start the action defined below 0 = Action finished | R/W SC | 0b |
| 6:2 | RESERVED | RO | 000_00b |
| 1:0 | Action Specify the action to be taken for the VLAN table entry addressed in the index register 00 = No operation 01 = Write 10 = Read 11 = Clear all entries to zero | R/W | 00b |

Note: Refer to [Section 5.3.4, "VLAN Table," on page 162](#) for additional information on VLAN tables.

5.1.5.6 ALU Table Index 0 Register

Address: 0x0410 - 0x0413 Size: 32 bits

The [ALU Table Index 0 Register](#) and [ALU Table Index 1 Register](#) contain the index values for searching and reading/writing the address lookup table as specified by the action in the [ALU Table Access Control Register](#). For additional information on ALU tables, refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#).

| Bits | Description | Type | Default |
|-------|---|------|-------------|
| 31 | RESERVED | R/W | 0b |
| 30:23 | RESERVED | RO | 000_0000_0b |
| 22:16 | FID Index [6:0] This is the FID value used to hash index the table | R/W | 000_0000b |
| 15:0 | MAC Index [47:32] These are the upper 16 bits of the MAC addressed used to hash index the table | R/W | 0x0000 |

KSZ9896C

5.1.5.7 ALU Table Index 1 Register

Address: 0x0414 - 0x0417 Size: 32 bits

The [ALU Table Index 0 Register](#) and [ALU Table Index 1 Register](#) contain the index values for searching and reading/writing the address lookup table as specified by the action in the [ALU Table Access Control Register](#). For additional information on ALU tables, refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#).

| Bits | Description | Type | Default |
|------|---|------|------------|
| 31:0 | MAC Index [31:0] These are the lower 32 bits of the MAC addressed used to hash index the address lookup table. If direct addressing is enabled, then bits [11:0] are used to directly index the address lookup table. | R/W | 0x00000000 |

5.1.5.8 ALU Table Access Control Register

Address: 0x0418 - 0x041B Size: 32 bits

This register provides control and status for searching and reading or writing the ALU Table. The [ALU Table Index 0 Register](#) and [ALU Table Index 1 Register](#) contain the index values, while the [ALU / Static Address Table Entry 1 Register](#), [ALU / Static Address / Reserved Multicast Table Entry 2 Register](#), [ALU / Static Address Table Entry 3 Register](#), and [ALU / Static Address Table Entry 4 Register](#) are used for the entry values. Refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#) for additional information on ALU tables.

| Bits | Description | Type | Default |
|-------|--|---------|---------|
| 31:30 | RESERVED | RO | 00b |
| 29:16 | VALID_COUNT Indicates the total number of valid entries in the table after the search finishes | RO | 0x0000 |
| 15:8 | RESERVED | RO | 0x00 |
| 7 | START_FINISH 1 = Start the action defined below 0 = Action finished | R/W, SC | 0b |
| 6 | VALID For search operation. It is cleared when register 0x042F (ALU / Static Address Table Entry 4 Register) is read. 1 = Indicates that the next valid entry is ready 0 = Next valid entry is not ready | RO | 0b |
| 5 | VALID_ENTRY_OR_SEARCH_END For search operation. It is intended for added convenience when accessing registers by in-band management (IBA) rather than SPI or I ² C. It combines bits 6 and 7. 1 = Indicates either that the next valid entry is ready, or that the search has ended. 0 = Neither next valid entry ready nor search ended. | RO | 0b |
| 4:3 | RESERVED | R/W | 00b |

| Bits | Description | Type | Default |
|------|--|------|---------|
| 2 | DIRECT 1 = Access by direct addressing the ALU Table. This method is used only for debugging, if at all. 0 = Hashing function is used to index the table. This is the normal method. Refer to the Switch Lookup Engine Control 0 Register . | R/W | 0b |
| 1:0 | ACTION Specifies the action to be taken for the ALU Table entry access 00 = No operation 01 = Write 10 = Read 11 = Search | R/W | 00b |

5.1.5.9 Static Address and Reserved Multicast Table Control Register

Address: 0x041C - 0x041F Size: 32 bits

This register provides control and index for reading and writing the Static Address Table and the Reserved Multicast Table. The [ALU / Static Address Table Entry 1 Register](#), [ALU / Static Address / Reserved Multicast Table Entry 2 Register](#), [ALU / Static Address Table Entry 3 Register](#), and [ALU / Static Address Table Entry 4 Register](#) are used for the Static Address Table entry values, while only the [ALU / Static Address / Reserved Multicast Table Entry 2 Register](#) is used for the Reserved Multicast Table entry value. Refer to [Section 5.3.2, "Static Address Table," on page 159](#) and [Section 5.3.3, "Reserved Multicast Address Table," on page 161](#) for additional information on these tables.

| Bits | Description | Type | Default |
|-------|--|---------|----------|
| 31:22 | RESERVED | RO | 0x000 |
| 21:16 | TABLE_INDEX Bits [21:16] used to index the Reserved Multicast Table Bits [19:16] used to index the Static Address Table | R/W | 00_0000b |
| 15:8 | RESERVED | RO | 0x00 |
| 7 | START_FINISH 1 = Start access 0 = Access is finished | R/W, SC | 0b |
| 6:2 | RESERVED | R/W | 000_00b |
| 1 | TABLE_SELECT Specifies which table is being accessed 1 = Access Reserved Multicast Table 0 = Access Static Address Table | R/W | 0b |
| 0 | ACTION Specifies the action to be taken for the table 1 = Read 0 = Write | R/W | 0b |

KSZ9896C

5.1.5.10 ALU / Static Address Table Entry 1 Register

Address: 0x0420 - 0x0423 Size: 32 bits

This register contains the table entry values for read and write operations to the Address Lookup Table and Static Address Table. The field definitions of this register differ dependent on the table type used, as defined in the following sections:

- [ALU Table Entry 1 Register on page 158](#)
- [Static Address Table Entry 1 Register on page 160](#)

Refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#) and [Section 5.3.2, "Static Address Table," on page 159](#) for additional information.

5.1.5.11 ALU / Static Address / Reserved Multicast Table Entry 2 Register

Address: 0x0424 - 0x0427 Size: 32 bits

The field definitions of this register differ dependent on the table type used, as defined in the following sections:

- [ALU Table Entry 2 Register on page 158](#)
- [Static Address Table Entry 2 Register on page 160](#)
- [Reserved Multicast Address Table Entry 2 Register on page 162](#)

Refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#), [Section 5.3.2, "Static Address Table," on page 159](#), and [Section 5.3.3, "Reserved Multicast Address Table," on page 161](#) for additional information.

5.1.5.12 ALU / Static Address Table Entry 3 Register

Address: 0x0428 - 0x042B Size: 32 bits

The field definitions of this register differ dependent on the table type used, as defined in the following sections:

- [ALU Table Entry 3 Register on page 159](#)
- [Static Address Table Entry 3 Register on page 161](#)

Refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#) and [Section 5.3.2, "Static Address Table," on page 159](#) for additional information.

5.1.5.13 ALU / Static Address Table Entry 4 Register

Address: 0x042C - 0x042F Size: 32 bits

This register contains the table entry values for read and write operations to the Address Lookup Table and Static Address Table. The field definitions of this register differ dependent on the table type used, as defined in the following sections:

- [ALU Table Entry 4 Register on page 159](#)
- [Static Address Table Entry 4 Register on page 161](#)

Refer to [Section 5.3.1, "Address Lookup \(ALU\) Table," on page 155](#) and [Section 5.3.2, "Static Address Table," on page 159](#) for additional information.

5.2 Port Registers

This section details the device's port registers. All settings are on a per-port basis. Address field "N" specifies the port number. Valid values for "N" are 1 to 6 for some registers, 6 for MAC port specific registers, and 0 to 5 for PHY specific registers.

For an overview of the device's entire register map, refer to [Section 5.0, "Device Registers"](#). For details on the device's global registers, refer to [Section 5.1, "Global Registers"](#).

5.2.1 PORT N: PORT OPERATION CONTROL REGISTERS (0xN000 - 0xN0FF)

5.2.1.1 Port Default Tag 0 Register

Address: 0xN000 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:0 | Port Default 802.1Q Tag [15:8] Bits [7:5]: Priority Code Point (PCP) Bit [4]: Drop Eligible Indicator (DEI) Bits [3:0]: VLAN Identifier (VID) [11:8] | R/W | 0x00 |

5.2.1.2 Port Default Tag 1 Register

Address: 0xN001 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:0 | Port Default 802.1Q tag [7:0] VLAN Identifier (VID) [7:0] | R/W | 0x01 |

5.2.1.3 Port PME_WoL Event Register

Address: 0xN013 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|-------|---------|
| 7:3 | RESERVED | RO | 0000_0b |
| 2 | Magic Packet Detect To be detected, the destination address of the Magic Packet must match the Global Switch MAC Address in the Switch MAC Address 0 Register through Switch MAC Address 5 Register . 1 = Magic Packet is detected at this port 0 = Not detected | RO/WC | 0b |

KSZ9896C

| Bits | Description | Type | Default |
|------|--|-------|---------|
| 1 | Link Up Detect Applicable only to ports with integrated PHY. 1 = Link up is detected at this port 0 = Not detected | RO/WC | 0b |
| 0 | Energy Detect Applicable only to ports with integrated PHY. 1 = Cable energy is detected at this port 0 = Not detected | RO/WC | 0b |

5.2.1.4 Port PME_WoL Enable Register

Address: 0xN017 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|-------|---------|
| 7:3 | RESERVED | RO | 0000_0b |
| 2 | Magic Packet Detect Enable To be detected, the destination address of the Magic Packet must match the Global Switch MAC Address in the Switch MAC Address 0 Register through Switch MAC Address 5 Register . 1 = The PME pin will be asserted when a Magic Packet is detected at this port 0 = The PME pin will not be asserted by Magic Packet detection at this port | RO/WC | 0b |
| 1 | Link Up Detect Enable Applicable only to ports with integrated PHY. 1 = The PME pin will be asserted when link up is detected at this port 0 = The PME pin will not be asserted by link up detection at this port | RO/WC | 0b |
| 0 | Energy Detect Enable Applicable only to ports with integrated PHY. 1 = The PME pin will be asserted when cable energy is detected at this port 0 = The PME pin will not be asserted by cable energy detection at this port | RO/WC | 0b |

5.2.1.5 Port Interrupt Status Register

Address: 0xN01B Size: 8 bits
Port N: 1-6

These registers provide interrupt status for the individual ports. these interrupts are enabled in the [Port Interrupt Mask Register](#). For non-port specific interrupt status, refer to the [Global Interrupt Status Register](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:2 | RESERVED | RO | 0000b |
| 1 | PHY Interrupt Status Applies only to PHY ports, not MAC port(s). Refer to the Port Interrupt Control / Status Register for individual PHY interrupts. 0 = No interrupt 1 = Interrupt detected | RO | 0b |
| 0 | ACL Interrupt Status ACL interrupts can be generated by configuring an entry in the ACL Table for counter mode. All ports have an ACL Table. To clear this bit, toggle the ACL Interrupt Mask bit in the Port Interrupt Mask Register . 0 = No interrupt 1 = Interrupt detected | RO | 0b |

5.2.1.6 Port Interrupt Mask Register

Address: 0xN01F Size: 8 bits
Port N: 1-6

This register enables the interrupts in the [Port Interrupt Status Register](#).

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:2 | RESERVED | RO | 0000b |
| 1 | PHY Interrupt Enable Applies only to PHY ports. 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |
| 0 | ACL Interrupt Enable Applies to all ports. 0 = Interrupt enabled 1 = Interrupt disabled | R/W | 0b |

KSZ9896C

5.2.1.7 Port Operation Control 0 Register

Address: 0xN020 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | Local Loopback Data going from the internal switch fabric to an egress port is looped back at that egress port and returned to the internal switch fabric. 1 = Local loopback 0 = Normal operation | R/W | 0b |
| 6 | Remote Loopback Data received at the external pins of a port is looped back and sent out the same port, without passing through the internal switch fabric. For PHY ports, data is received on the RX+/- pins and transmitted out the TX+/- pins. For xMII ports, data is received on the TXD pins and transmitted on the RXD pins. 1 = Remote loopback 0 = Normal operation | R/W | 0b |
| 5:3 | RESERVED | RO | 000b |
| 2 | Tail Tag Enable When tail tagging is enabled for a port, it designates that port to be the “host” or “CPU” port. Do not enable tail tagging for more than one port. See the tail tagging description for details. 1 = Enable tail tagging on this port 0 = Disable tail tagging for this port | R/W | 0b |
| 1:0 | Egress Queue Split Enable 11 = Reserved 10 = Four egress queues. Packets are assigned to a queue based on priority. 01 = Two egress queues. Packets are assigned to a queue based on priority. 00 = Single egress queue. There is no priority differentiation. | R/W | 00b |

5.2.1.8 Port Status Register

Address: 0xN030 Size: 8 bits
 Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:5 | RESERVED | RO | 000 |
| 4:3 | <p>Port Speed Status</p> <p>For non-PHY ports, these bits duplicate the speed setting in the XMII Port Control 1 Register bit 6 and XMII Port Control 0 Register bit 4.</p> <p>For PHY ports, these bits indicate the actual link speed, which is also available in the PHY Control Register.</p> <p>00 = 10 Mb/s 01 = 100 Mb/s 10 = 1000 Mb/s</p> | RO | - |
| 2 | <p>Port Duplex Status</p> <p>For non-PHY ports, this bit duplicates the duplex setting in the XMII Port Control 0 Register bit 6.</p> <p>For PHY ports, this bit indicates the actual link duplex, which is also available in the PHY Control Register.</p> <p>1 = Full duplex 0 = Half duplex</p> | RO | 1 or - |
| 1 | <p>Transmit Flow Control Enabled Status</p> <p>For non-PHY ports, this bit duplicates the Tx FC enable bit 5 in the XMII Port Control 0 Register.</p> <p>For PHY ports, this bit is set only when FC is enabled (PHY Auto-Negotiation Advertisement Register, bits 11:10), link is up and FC is established via auto-negotiation.</p> <p>1 = TX flow control is enabled 0 = Disabled</p> | RO | - |
| 0 | <p>Receive Flow Control Enabled Status</p> <p>For non-PHY ports, this bit duplicates the Rx FC enable bit 3 in the XMII Port Control 0 Register.</p> <p>For PHY ports, this bit is set only when FC is enabled (PHY Auto-Negotiation Advertisement Register, bits 11:10), link is up and FC is established via auto-negotiation.</p> <p>1 = RX flow control is enabled 0 = Disabled</p> | RO | - |

KSZ9896C

5.2.2 PORT N: PORT ETHERNET PHY REGISTERS (0xN100 - 0xN1FF)

The registers in this section are for PHY ports only. Refer to the IEEE802.3 clause 22.3.4 for additional details.

5.2.2.1 PHY Basic Control Register

Address: 0xN100 - 0xN101 Size: 16 bits
Port N: 1-5

PHY register 0x00.

| Bits | Description | Type | Default |
|------|---|-----------|--------------------------|
| 15 | PHY Software Reset Set this bit to reset this PHY. Registers are not reset. This bit is self-clearing. | R/W SC | 0b |
| 14 | Local Loopback Mode Data going from the internal switch fabric to an egress port is looped back at that egress port and returned to the internal switch fabric. 1 = Local Loopback mode 0 = Normal operation | R/W | 0b |
| 13 | Speed Select (LSB) This bit is ignored if auto-negotiation is enabled (bit 12 in this register). Bits [6, 13] 11 = Reserved 10 = 1000 Mb/s 01 = 100 Mb/s 00 = 10 Mb/s | R/W | Note 5-4 |
| 12 | Auto-Negotiation Enable 1 = Auto-negotiation is enabled 0 = Auto-negotiation is disabled The initial value of this bit is determined by a strapping option, but it may be overwritten. | R/W | Note 5-4 |
| 11 | Power Down 1 = Power-down mode 0 = Normal operation When this bit is set to '1', the link-down status might not get updated in the PHY status register. Software should note link is down and should not rely on the PHY status register link status. After this bit is changed from '1' to '0', an internal reset is performed. Wait a minimum of 1ms before read/write access to this PHY's registers. | R/W | 0b |
| 10 | Isolate 1 = Logical isolation of the PHY from the switch core 0 = Normal operation | R/W | 0b |
| 9 | Re-start Auto-Negotiation Set this bit to re-start auto-negotiation. This bit is self-clearing. | R/W, SC | 0b |
| 8 | Duplex Mode This bit is ignored if auto-negotiation is enabled (bit 12 in this register). 1 = Full duplex 0 = Half duplex | R/W | 1b |

| Bits | Description | Type | Default |
|------|--|------|--------------------------|
| 7 | Collision Test 1 = Enable COL test 0 = Disable COL test | R/W | 0b |
| 6 | Speed Select (MSB) This bit is ignored if auto-negotiation is enabled (bit 12 in this register). Bits [6, 13] 11 = Reserved 10 = 1000 Mb/s 01 = 100 Mb/s 00 = 10 Mb/s | R/W | Note 5-4 |
| 5:0 | RESERVED | RO | 0x00 |

Note 5-4 The default value of this field is determined by the associated configuration strap value. Refer to [Section 3.2.1, "Configuration Straps,"](#) on page 17 for additional information.

5.2.2.2 PHY Basic Status Register

Address: 0xN102 - 0xN103 Size: 16 bits
Port N: 1-5

PHY register 0x01.

| Bits | Description | Type | Default |
|------|--|----------|---------|
| 15 | 100BASE-T4 0 = PHY is not 100BASE-T4 capable | RO | 0b |
| 14 | 100BASE-TX Full Duplex 1 = PHY is 100BASE-TX full-duplex capable | RO | 1b |
| 13 | 100BASE-TX Half Duplex 1 = PHY is 100BASE-TX half-duplex capable | RO | 1b |
| 12 | 10 Mb/s Full Duplex 1 = PHY is 10 Mb/s full-duplex capable | RO | 1b |
| 11 | 10 Mb/s Half Duplex 1 = PHY is 10 Mb/s half-duplex capable | RO | 1b |
| 10:9 | RESERVED | RO | 00b |
| 8 | Extended Status 1 = PHY Extended Status Register contains extended status information. | RO | 1b |
| 7 | RESERVED | RO | 0b |
| 6 | MF Preamble Suppression 1 = PHY will accept management frames with preamble suppressed | RO | 1b |
| 5 | Auto-negotiation Complete 1 = Auto-negotiation process is complete 0 = Auto-negotiation process is not complete | RO | - |
| 4 | Remote (Far End) Fault 1 = Remote fault condition detected 0 = No remote fault detected | RO LH | - |

KSZ9896C

| Bits | Description | Type | Default |
|------|---|----------|---------|
| 3 | Auto-Negotiation Ability 1 = PHY is able to perform auto-negotiation | RO | 1b |
| 2 | Link Status 1 = Link is up 0 = Link is down | RO LL | - |
| 1 | Jabber Detect 1 = Jabber condition detected 0 = No jabber condition detected | RO LH | - |
| 0 | Extended Capability 1 = Supports extended capabilities register 0 = Basic register set capabilities only | RO | 1b |

5.2.2.3 PHY ID High Register

Address: 0xN104 - 0xN105 Size: 16 bits
Port N: 1-5

PHY register 0x02.

| Bits | Description | Type | Default |
|------|---------------------------------|------|---------|
| 15:0 | PHY Identifier High Word | RO | 0x0022 |

5.2.2.4 PHY ID Low Register

Address: 0xN106 - 0xN107 Size: 16 bits
Port N: 1-5

PHY register 0x03.

| Bits | Description | Type | Default |
|------|--------------------------------|------|---------|
| 15:0 | PHY Identifier Low Word | RO | 0x1631 |

5.2.2.5 PHY Auto-Negotiation Advertisement Register

Address: 0xN108 - 0xN109 Size: 16 bits
Port N: 1-5

PHY register 0x04.

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15 | Next Page 1 = Next page capable 0 = No next page capability | R/W | 0b |
| 14 | Reserved | RO | 0b |
| 13 | Remote Fault 1 = Remote fault capable 0 = No remote fault capability | R/W | 0b |
| 12 | RESERVED | RO | 0b |
| 11:10 | Pause (Flow Control) Capability 11 = Both Symmetric PAUSE and Asymmetric PAUSE toward local device 10 = Asymmetric PAUSE toward link partner 01 = Symmetric PAUSE 00 = No PAUSE | R/W | 01b |
| 9 | 100BASE-T4 1 = 100BASE-T4 capable 0 = No 100BASE-T4 capability | RO | 0b |
| 8 | 100BASE-TX Full Duplex 1 = 100BASE-TX full duplex capable 0 = No 100BASE-TX full duplex capability | R/W | 1b |
| 7 | 100BASE-TX Half Duplex 1 = 100BASE-TX half duplex capable 0 = No 100BASE-TX half duplex capability | R/W | 1b |
| 6 | 10BASE-T Full Duplex 1 = 10BASE-T full duplex capable 0 = No 10BASE-T full duplex capability | R/W | 1b |
| 5 | 10BASE-T Half Duplex 1 = 10BASE-T half duplex capable 0 = No 10BASE-T half duplex capability | R/W | 1b |
| 4:0 | Selector Field 00001 = IEEE 802.3 | R/W | 0x01 |

KSZ9896C

5.2.2.6 PHY Auto-Negotiation Link Partner Ability Register

Address: 0xN10A - 0xN10B Size: 16 bits
Port N: 1-5

PHY register 0x05.

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15 | Next Page 1 = Next page capable 0 = No next page capability | RO | - |
| 14 | Acknowledge 1 = Link code word received from partner 0 = Link code word not yet received | RO | - |
| 13 | Remote Fault 1 = Remote fault detected 0 = No remote fault | RO | - |
| 12 | RESERVED | RO | - |
| 11:10 | Pause (Flow Control) Capability 11 = Both Symmetric PAUSE and Asymmetric PAUSE toward local device 10 = Asymmetric PAUSE toward link partner 01 = Symmetric PAUSE 00 = No PAUSE | RO | - |
| 9 | 100BASE-T4 1 = 100BASE-T4 capable 0 = No 100BASE-T4 capability | RO | - |
| 8 | 100BASE-TX Full Duplex 1 = 100BASE-TX full duplex capable 0 = No 100BASE-TX full duplex capability | RO | - |
| 7 | 100BASE-TX Half Duplex 1 = 100BASE-TX half duplex capable 0 = No 100BASE-TX half duplex capability | RO | - |
| 6 | 10BASE-T Full Duplex 1 = 10BASE-T full duplex capable 0 = No 10BASE-T full duplex capability | RO | - |
| 5 | 10BASE-T Half Duplex 1 = 10BASE-T half duplex capable 0 = No 10BASE-T half duplex capability | RO | - |
| 4:0 | Selector Field 00001 = IEEE 802.3 | RO | - |

5.2.2.7 PHY Auto-Negotiation Expansion Status Register

Address: 0xN10C - 0xN10D Size: 16 bits
Port N: 1-5

PHY register 0x06.

| Bits | Description | Type | Default |
|------|---|----------|---------|
| 15:5 | RESERVED | RO | 0x000 |
| 4 | Parallel Detection Fault 1 = Fault detected by parallel detection 0 = No fault detected by parallel detection | RO LH | - |
| 3 | Link Partner Next Page Able 1 = Link partner has next page capability 0 = Link partner does not have next page capability | RO | - |
| 2 | Next Page Able 1 = Local device has next page capability 0 = Local device does not have next page capability | RO | 1b |
| 1 | Page Received 1 = New page received 0 = New page not received | RO LH | - |
| 0 | Link Partner Auto-Negotiation Able 1 = Link partner has auto-negotiation capability 0 = Link partner does not have auto-negotiation capability | RO | - |

5.2.2.8 PHY Auto-Negotiation Next Page Register

Address: 0xN10E - 0xN10F Size: 16 bits
Port N: 1-5

PHY register 0x07.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 15 | Next Page 1 = Additional next page(s) will follow 0 = Last page | R/W | 0b |
| 14 | Reserved | RO | 0b |
| 13 | Message Page 1 = Message page 0 = Unformatted page | R/W | 1b |
| 12 | Acknowledge 2 1 = Will comply with message 0 = Cannot comply with message | R/W | 0b |

KSZ9896C

| Bits | Description | Type | Default |
|------|--|------|---------|
| 11 | Toggle 1 = Previous value of transmitted link code word equaled logic one 0 = Previous value of transmitted link code word equaled logic zero | RO | 0b |
| 10:0 | Message Field | R/W | 0x001 |

5.2.2.9 PHY Auto-Negotiation Link Partner Next Page Ability Register

Address: 0xN110 - 0xN111 Size: 16 bits
Port N: 1-5

PHY register 0x08.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 15 | Next Page 1 = Additional next page(s) will follow 0 = Last page | R/W | - |
| 14 | Acknowledge 1 = Successful receipt of link word 0 = No successful receipt of link word | RO | - |
| 13 | Message Page 1 = Message page 0 = Unformatted page | RO | - |
| 12 | Acknowledge 2 1 = Able to act on the information 0 = Not able to act on the information | RO | - |
| 11 | Toggle 1 = Previous value of transmitted link code word equaled logic zero 0 = Previous value of transmitted link code word equaled logic one | RO | - |
| 10:0 | Message Field | RO | - |

5.2.2.10 PHY 1000BASE-T Control Register

Address: 0xN112 - 0xN113 Size: 16 bits
Port N: 1-5

PHY register 0x09.

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15:13 | Test Mode Bits | R/W | 000b |
| 12 | Master-Slave Manual Configuration Enable 1 = Enable master-slave manual configuration 0 = Disable master-slave manual configuration | R/W | 0b |
| 11 | Master-Slave Manual Configuration Value This bit is ignored if master-slave manual configuration is disabled (bit 12) 1 = Configure PHY as master during master-slave negotiation 0 = Configure PHY as slave during master-slave negotiation | R/W | 0b |
| 10 | Port Type This bit is valid only if master-slave manual configuration is disabled (bit 12) 1 = Indicate the preference to operate as multi-port device (master) 0 = Indicate the preference to operate as single-port device (slave) | R/W | 0b |
| 9 | 1000BASE-T Full Duplex 1 = Advertise that PHY is 1000BASE-T full duplex capable 0 = Advertise that PHY is not 1000BASE-T full duplex capable | R/W | 1b |
| 8 | 1000BASE-T Half Duplex 1 = Advertise that PHY is 1000BASE-T half duplex capable 0 = Advertise that PHY is not 1000BASE-T half duplex capable | R/W | 1b |
| 7:0 | RESERVED | RO | - |

5.2.2.11 PHY 1000BASE-T Status Register

Address: 0xN114 - 0xN115 Size: 16 bits
Port N: 1-5

PHY register 0x0A.

| Bits | Description | Type | Default |
|------|---|----------------|---------|
| 15 | Master-Slave Configuration Fault 1 = Master-slave configuration fault detected 0 = No master-slave configuration fault detected | RO LH SC | 0b |
| 14 | Master-Slave Configuration Resolution 1 = Local PHY configuration resolved to master 0 = Local PHY configuration resolved to slave | RO | - |
| 13 | Local Receiver Status 1 = Local receiver OK 0 = Local receiver not OK | RO | - |

KSZ9896C

| Bits | Description | Type | Default |
|------|--|----------|---------|
| 12 | Remote Receiver Status 1 = Remote receiver OK 0 = Remote receiver not OK | RO | - |
| 11 | Link Partner 1000BASE-T Full Duplex Capability 1 = Link partner is capable of 1000BASE-T full duplex 0 = Link partner is not capable of 1000BASE-T full duplex | RO | - |
| 10 | Link Partner 1000BASE-T Half Duplex Capability 1 = Link partner is capable of 1000BASE-T half duplex 0 = Link partner is not capable of 1000BASE-T half duplex | RO | - |
| 9:8 | RESERVED | RO | 00b |
| 7:0 | Idle Error Count Cumulative count of errors detected when receiver is receiving idles and PMA_TXMODE.indicate = Send_N. The counter is incremented every symbol period that rxerror_status = ERROR. | RO SC | 0x00 |

5.2.2.12 PHY MMD Setup Register

Address: 0xN11A - 0xN11B Size: 16 bits
Port N: 1-5

PHY register 0x0D.

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 15:14 | MMD Operation Mode For the selected MMD device address (bits [4:0] of this register), these two bits select one of the following register or data operations and the usage for the PHY MMD Data Register . 00 = Register 01 = Data, no post increment 10 = Data, post increment on reads and writes 11 = Data, post increment on writes only | R/W | 00b |
| 13:5 | RESERVED | R/W | 0x000 |
| 4:0 | MMD Device Address | R/W | 0x00 |

5.2.2.13 PHY MMD Data Register

Address: 0xN11C - 0xN11D Size: 16 bits
Port N: 1-5

PHY register 0x0E.

| Bits | Description | Type | Default |
|------|---|------|---------|
| 15:0 | <p>MMD R/W Index/Data</p> <p>For the selected MMD device address (in the PHY MMD Setup Register, bits [4:0]):</p> <p>When the PHY MMD Setup Register bits [15:14] = 00, this register contains the read/write register access for the MDD device address.</p> <p>Otherwise, this register contains the read/write data value for the MMD device address and its selected register address</p> <p>Refer to the PHY MMD Setup Register bits [15:14] for descriptions of post increment reads and writes of this register for data operation.</p> | R/W | 0x0000 |

5.2.2.14 PHY Extended Status Register

Address: 0xN11E - 0xN11F Size: 16 bits
Port N: 1-5

PHY register 0x0F.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 15 | <p>1000BASE-X Full-Duplex</p> <p>1 = PHY able to perform 1000BASE-T full duplex 0 = PHY cannot perform 1000BASE-T full-duplex</p> | RO | 0b |
| 14 | <p>1000BASE-X Half-Duplex</p> <p>1 = PHY able to perform 1000BASE-T half duplex 0 = PHY cannot perform 1000BASE-T half-duplex</p> | RO | 0b |
| 13 | <p>1000BASE-T Full-Duplex</p> <p>1 = PHY able to perform 1000BASE-T full duplex 0 = PHY cannot perform 1000BASE-T full-duplex</p> | RO | 1b |
| 12 | <p>1000BASE-T Half-Duplex</p> <p>1 = PHY able to perform 1000BASE-T half duplex 0 = PHY cannot perform 1000BASE-T half-duplex</p> | RO | 1b |
| 11:0 | RESERVED | RO | – |

KSZ9896C

5.2.2.15 PHY Remote Loopback Register

Address: 0xN122 - 0xN123 Size: 16 bits
Port N: 1-5

PHY register 0x11.

| Bits | Description | Type | Default |
|------|--|-----------|----------|
| 15:9 | RESERVED | R/W | 0x00 |
| 8 | Remote Loopback Data received at the external pins of a port is looped back and sent out the same port, without passing through the internal switch fabric. For PHY ports, data is received on the RX+/- pins and transmitted out the TX+/- pins. For xMII ports, data is received on the TXD pins and transmitted on the RXD pins. 1 = Remote loopback 0 = Normal operation | R/W | 0b |
| 7:2 | RESERVED | R/W | 1111_01b |
| 1 | RESERVED | R/W RC | 0b |
| 0 | RESERVED | RO | 0b |

5.2.2.16 PHY LinkMD Register

Address: 0xN124 - 0xN125 Size: 16 bits
Port N: 1-5

PHY register 0x12.

| Bits | Description | Type | Default |
|-------|--|---------|---------|
| 15 | Cable Diagnostic Test Enable Write value: 1 = Enable cable diagnostic test. After test has completed, this bit is self-cleared. 0 = Disable cable diagnostic test. Read value: 1 = Cable diagnostic test is in progress. 0 = Indicates cable diagnostic test (if enabled) has completed and the status information is valid for read. | R/W, SC | 0b |
| 14 | RESERVED | R/W | 0b |
| 13:12 | Cable Diagnostic Test Pair These two bits select the differential pair for testing: 00 = Differential pair A 01 = Differential pair B 10 = Differential pair C 11 = Differential pair D | R/W | 00b |

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 11:10 | RESERVED | R/W | 00b |
| 9:8 | Cable Diagnostic Status 00 = Normal cable condition (no fault detected) 01 = Open cable fault detected 10 = Short cable fault detected 11 = Reserved | RO | 00b |
| 7:0 | RESERVED | RO | 0x00 |

5.2.2.17 PHY Digital PMA/PCS Status Register

Address: 0xN126 - 0xN127 Size: 16 bits
Port N: 1-5

PHY register 0x13.

| Bits | Description | Type | Default |
|------|---|----------|---------|
| 15:2 | RESERVED | RO LH | 0x0000 |
| 1 | 1000BASE-T Link Status 1 = Link is up 0 = Link is down | RO | - |
| 0 | 100BASE-TX Link Status 1 = Link is up 0 = Link is down | RO | - |

5.2.2.18 Port RXER Count Register

Address: 0xN12A - 0xN12B Size: 16 bits
Port N: 1-5

PHY register 0x15.

| Bits | Description | Type | Default |
|------|--|----------|---------|
| 15:0 | RXER Counter Count of receive frames with one or more symbol errors. | RO RC | 0x0000 |

KSZ9896C

5.2.2.19 Port Interrupt Control / Status Register

Address: 0xN136 - 0xN137 Size: 16 bits
Port N: 1-5

PHY register 0x1B.

| Bits | Description | Type | Default |
|------|---|----------|---------|
| 15 | Jabber Interrupt Enable 1 = Enable jabber interrupt 0 = Disable jabber interrupt | R/W | 0b |
| 14 | Receive Error Interrupt Enable 1 = Enable receive error interrupt 0 = Disable receive error interrupt | R/W | 0b |
| 13 | Page Received Interrupt Enable 1 = Enable page received interrupt 0 = Disable page received interrupt | R/W | 0b |
| 12 | Parallel Detect Fault Interrupt Enable 1 = Enable parallel detect fault interrupt 0 = Disable parallel detect fault interrupt | R/W | 0b |
| 11 | Link Partner Acknowledge Interrupt Enable 1 = Enable link partner acknowledge interrupt 0 = Disable link partner acknowledge interrupt | R/W | 0b |
| 10 | Link Down Interrupt Enable 1 = Enable link down interrupt 0 = Disable link down interrupt | R/W | 0b |
| 9 | Remote Fault Interrupt Enable 1 = Enable remote fault interrupt 0 = Disable remote fault interrupt | R/W | 0b |
| 8 | Link Up Interrupt Enable 1 = Enable link up interrupt 0 = Disable link up interrupt | R/W | 0b |
| 7 | Jabber Interrupt 1 = Jabber occurred 0 = Jabber did not occur | RO RC | 0b |
| 6 | Receive Error Interrupt 1 = Receive error occurred 0 = Receive error did not occur | RO RC | 0b |
| 5 | Page Received Interrupt 1 = Page receive occurred 0 = Page receive did not occur | RO RC | 0b |
| 4 | Parallel Detect Fault Interrupt 1 = Parallel detect fault occurred 0 = Parallel detect fault did not occur | RO RC | 0b |

| Bits | Description | Type | Default |
|------|--|----------|---------|
| 3 | Link Partner Acknowledge Interrupt 1 = Link partner acknowledge occurred 0 = Link partner acknowledge did not occur | RO RC | 0b |
| 2 | Link Down Interrupt 1 = Link down occurred 0 = Link down did not occur | RO RC | 0b |
| 1 | Remote Fault Interrupt 1 = Remote fault occurred 0 = Remote fault did not occur | RO RC | 0b |
| 0 | Link Up Interrupt 1 = Link up occurred 0 = Link up did not occur | RO RC | 0b |

5.2.2.20 PHY Auto MDI / MDI-X Register

Address: 0xN138 - 0xN139 Size: 16 bits
Port N: 1-5

PHY register 0x1C.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 15:8 | RESERVED | R/W | 0x00 |
| 7 | MDI Set When Swap-Off bit is one, 1 = PHY is set to operate in MDI mode 0 = PHY is set to operate in MDI-X mode This bit has no function when Swap-Off is zero. | R/W | 0b |
| 6 | Swap-Off 1 = Disable Auto MDI / MDI-X function 0 = Enable Auto MDI / MDI-X function (normal operation) | R/W | 0b |
| 5:0 | RESERVED | R/W | 0x00 |

KSZ9896C

5.2.2.21 PHY Control Register

Address: 0xN13E - 0xN13F Size: 16 bits
Port N: 1-5

PHY register 0x1F.

| Bits | Description | Type | Default |
|-------|--|----------------|---------|
| 15:12 | RESERVED | R/W | 0x0 |
| 11:10 | RESERVED | RO LH RC | 00b |
| 9 | Enable Jabber 1 = Enable jabber counter 0 = Disable jabber counter | R/W | 1b |
| 8:7 | RESERVED | R/W | 10b |
| 6 | Speed Status 1000BASE-T 1 = PHY final speed status is 1000BASE-T | RO | - |
| 5 | Speed Status 100BASE-TX 1 = PHY final speed status is 100BASE-TX | RO | - |
| 4 | Speed Status 10BASE-T 1 = PHY final speed status is 10BASE-T | RO | - |
| 3 | Duplex Status 1 = Full duplex 0 = Half duplex | RO | - |
| 2 | 1000BASE-T Master / Slave Status 1 = 1000BASE-T master mode 0 = 1000BASE-T slave mode | RO | - |
| 1 | RESERVED | R/W RC | 0b |
| 0 | RESERVED | RO RC | - |

5.2.3 PORT N: PORT RGMII/GMII/MII/RMII CONTROL REGISTERS (0xN300 - 0xN3FF)

5.2.3.1 XMII Port Control 0 Register

Address: 0xN300 Size: 8 bits
Port N: 6

| Bits | Description | Type | Default |
|------|---|------|--------------------------|
| 7 | RESERVED | RO | 0b |
| 6 | MAC Port Duplex 1 = Port MAC interface operates in full-duplex mode 0 = Port MAC interface operates in half-duplex mode | R/W | 1b |
| 5 | MAC Port Transmit Flow Control Enable 1 = Enable transmit flow control on this port 0 = Disable transmit flow control on this port | R/W | Note 5-5 |
| 4 | MAC Port Speed 10/100 This bit is ignored if bit 6 in the XMII Port Control 1 Register selects 1000 Mb/s. If bit 6 in the XMII Port Control 1 Register is one, then this bit determines the port speed: 1 = Port operates at 100 Mb/s 0 = Port operates at 10 Mb/s | R/W | 1b |
| 3 | MAC Port Receive Flow Control Enable 1 = Enable receive flow control on this port 0 = Disable receive flow control on this port | R/W | Note 5-5 |
| 2:1 | RESERVED | R/W | 00b |
| 0 | RESERVED GMII/MII Clock Mode These modes affect the GMII and MII clocks for outgoing data on the Port 6 RX signals. 1 = 3-Wire mode. Pin 72 (RX_CLK6/REFCLK06) is used for the GMII output clock, and pin 75 (RX_ER6/RX_CLK6) is used for the MII input clock when in MAC mode. 3-Wire mode is intended as an option, when in MAC mode, for connection to a 1000/100/10Mbps PHY which has separate pins for the GMII (1000Mbps) GTX_CLK clock and the MII (100/10Mbps) TX_CLK clock. 0 = 2-Wire mode. Pin 72 (RX_CLK6/REFCLK06) is used for either the GMII output clock or the MII clock. Note: When the MAC interface does not need to be switchable between GMII and MII, or when in PHY mode, do not select 3-Wire mode. | R/W | Note 5-6 |

Note 5-5 The default value of this field is determined by the **LED1_1** configuration strap value. Refer to [Section 3.2.1, "Configuration Straps,"](#) on page 17 for additional information.

Note 5-6 The default value of this field is determined by the **RX_ER6/RX_CLK6** configuration strap value. Refer to [Section 3.2.1, "Configuration Straps,"](#) on page 17 for additional information.

KSZ9896C

5.2.3.2 XMII Port Control 1 Register

Address: 0xN301 Size: 8 bits
Port N: 6

| Bits | Description | Type | Default |
|------|---|------|--|
| 7 | RESERVED | R/W | 0b |
| 6 | Port Speed 1000 1 = RGMII operates at 10 or 100 Mb/s, as determined by bit 4 in the XMII Port Control 0 Register 0 = RGMII operates at 1000 Mb/s | R/W | Config. strap value from RXD6_0 |
| 5 | RESERVED | R/W | 0b |
| 4 | RGMII Ingress Internal Delay (RGMII_ID_ig) 1 = Minimum 1.5 ns delay is added to ingress RGMII clock 0 = No delay is added | R/W | 0b |
| 3 | RGMII Egress Internal Delay (RGMII_ID_eg) 1 = Minimum 1.5 ns delay is added to egress RGMII clock 0 = No delay is added | R/W | 1b |
| 2 | MII / RMII Modes For MII interface: 1 = The MII interface operates as a MAC device (receives clocks, etc.) 0 = The MII interface operates as a PHY device (drives clocks, etc.) For RMII interface: 1 = The 50 MHz RMII REFCLK is received at the RXC pin 0 = A 50 MHz RMII REFCLK is generated from the RXC pin | R/W | Config. strap value from RXD6_1 |
| 1:0 | Port Interface Type Select 00 = Interface is RGMII 01 = Interface is RMII 10 = Interface is GMII (if bit 6 = 0) or MII (if bit 6 = 1) 11 = Interface is MII | R/W | Config. strap value from RXD6_[3:2] |

5.2.4 PORT N: PORT SWITCH MAC CONTROL REGISTERS (0xN400 - 0xN4FF)

5.2.4.1 Port MAC Control 0 Register

Address: 0xN400 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:4 | RESERVED | RO | 0x0 |
| 3 | RESERVED | R/W | 0b |
| 2 | RESERVED | RO | 0b |
| 1 | Broadcast Storm Protection Enable 1 = Enable broadcast storm protection for ingress traffic 0 = Disable broadcast storm protection for ingress traffic | R/W | 0b |
| 0 | Jumbo Packet Enable 1 = Accept packets sizes up to 9000 bytes payload (excludes header and CRC) 0 = Standard payload limit of 1500 bytes applies | R/W | 0b |

5.2.4.2 Port MAC Control 1 Register

Address: 0xN401 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | R/W | 0b |
| 6 | RESERVED | RO | 0b |
| 5:4 | RESERVED | R/W | 00b |
| 3 | Back pressure Enable Back pressure is used only for half duplex. Full duplex uses PAUSE frame flow control which is enabled separately. 1 = Enable port back pressure 0 = Disable port back pressure | R/W | 0b |
| 2:1 | Reserved | R/W | 00b |
| 0 | Pass All Frames Bad frames (CRC error, too large, too small) are normally dropped. This feature allows them to be forwarded for mirroring purposes only. This bit does not affect filtering of flow control frames. To disable filtering of flow control frames, refer to the Switch MAC Control 4 Register . 1 = Enable 0 = Disable | R/W | 0b |

KSZ9896C

5.2.4.3 Port Ingress Rate Limit Control Register

Address: 0xN403 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6 | Port or Priority Based Ingress Rate Limiting 1 = Ingress rate limiting is port based 0 = Ingress rate limiting is priority based | R/W | 0b |
| 5 | Ingress PPS Mode 1 = Ingress rate limiting is based on number of packets (PPS mode enabled) 0 = Ingress rate limiting is based on number of bits (PPS mode disabled) | R/W | 0b |
| 4 | Ingress Rate Limit Flow Control Enable To use this feature, flow control must also be enabled for this port. Refer to the XMII Port Control 0 Register for non-PHY ports, and the PHY Auto-Negotiation Advertisement Register for PHY ports. 1 = Flow control is asserted if the port's receive rate is exceeded 0 = Flow control is not asserted based on ingress rate limits | R/W | 0b |
| 3:2 | Ingress Limit Mode These bits determine what kind of frames are limited and counted against ingress rate limiting. 00 = Count and limit all frames 01 = Count and limit broadcast, multicast and flooded unicast frames only 10 = Count and limit broadcast and multicast frames only 11 = Count and limit broadcast frames only | R/W | 00b |
| 1 | Count IFG Bytes 1 = Each frame's minimum inter-frame gap (IFG) bytes (12 per frame) are included in ingress rate limiting calculations 0 = IFG byte count is not included | R/W | 0b |
| 0 | Count Preamble Bytes Not valid when PPS mode is enabled (bit 5) 1 = Each frame's preamble bytes (8 per frame) are included in ingress rate limiting calculations 0 = Preamble byte count is not included | R/W | 0b |

5.2.4.4 Port Priority 0 Ingress Limit Control Register

Address: 0xN410 Size: 8 bits
 Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | Reserved | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 0 Frames Reference Table 5-3 for code values. Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

TABLE 5-3: DATA RATE SELECTION TABLE FOR INGRESS AND EGRESS RATE LIMITING

| Code | 10Mbps | | 100Mbps | | 1000Mbps | |
|----------------|--------------|--------------|--------------|--------------|--------------|---------------|
| | PPS | BPS | PPS | BPS | PPS | BPS |
| 7d'0 | 19.2k | 10Mbps | 192k | 100Mbps | 1.92M | 1000Mbps |
| 7d'1 – 7d'10 | 1.92k * code | 1Mbps * code | 1.92k * code | 1Mbps * code | 19.2k * code | 1Mbps * code |
| 7d'11 – 7d'100 | – | 10Mbps | 1.92k * code | 1Mbps * code | 19.2k * code | 10Mbps * code |
| 7d'101 | 64 | 64kbps | 64 | 640kbps | 640 | 640kbps |
| 7d'102 | 128 | 128kbps | 128 | 1280kbps | 1280 | 1280kbps |
| 7d'103 | 256 | 192kbps | 256 | 1920kbps | 2560 | 1920kbps |
| 7d'104 | 384 | 256kbps | 384 | 256kbps | 3840 | 2560kbps |
| 7d'105 | 512 | 320kbps | 512 | 320kbps | 5120 | 3200kbps |
| 7d'106 | 640 | 384kbps | 640 | 384kbps | 6400 | 3840kbps |
| 7d'107 | 768 | 448kbps | 768 | 448kbps | 7680 | 4480kbps |
| 7d'108 | 896 | 512kbps | 896 | 512kbps | 8960 | 5120kbps |
| 7d'109 | 1024 | 576kbps | 1024 | 576kbps | 10,240 | 5760kbps |
| 7d'110 | 1152 | 640kbps | 1152 | 640kbps | 11,520 | 6400kbps |
| 7d'111 | 1280 | 704kbps | 1280 | 704kbps | 12,800 | 7040kbps |
| 7d'112 | 1408 | 768kbps | 1408 | 768kbps | 14,080 | 7680kbps |
| 7d'113 | 1536 | 832kbps | 1536 | 832kbps | 15,360 | 8320kbps |
| 7d'114 | 1664 | 896kbps | 1664 | 896kbps | 16,640 | 8960kbps |
| 7d'115 | 1792 | 960kbps | 1792 | 960kbps | 17,920 | 9600kbps |

Note: PPS = Packets Per Second, BPS = Bits Per Second.

KSZ9896C

5.2.4.5 Port Priority 1 Ingress Limit Control Register

Address: 0xN411 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 1 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.6 Port Priority 2 Ingress Limit Control Register

Address: 0xN412 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 2 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.7 Port Priority 3 Ingress Limit Control Register

Address: 0xN413 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 3 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.8 Port Priority 4 Ingress Limit Control Register

Address: 0xN414 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 4 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.9 Port Priority 5 Ingress Limit Control Register

Address: 0xN415 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 5 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.10 Port Priority 6 Ingress Limit Control Register

Address: 0xN416 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 6 Frames Note that the update will not take effect until the Port Priority 7 Ingress Limit Control Register is written. | R/W | 0x00 |

5.2.4.11 Port Priority 7 Ingress Limit Control Register

Address: 0xN417 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Ingress Data Rate Limit for Priority 7 Frames | R/W | 0x00 |

KSZ9896C

5.2.4.12 Port Queue 0 Egress Limit Control Register

Address: 0xN420 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Egress Data Rate Limit for Queue 0 Frames When egress rate limiting is configured to “port-based” via bit 3 of the Switch MAC Control 5 Register , only this register is used for setting. The Port Queue 1-3 Egress Limit Control Registers are used only for queue-based rate limiting. Note that the update will not take effect until the Port Queue 3 Egress Limit Control Register is written. | R/W | 0x00 |

5.2.4.13 Port Queue 1 Egress Limit Control Register

Address: 0xN421 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Egress Data Rate Limit for Queue 1 Frames Note that the update will not take effect until the Port Queue 3 Egress Limit Control Register is written. | R/W | 0x00 |

5.2.4.14 Port Queue 2 Egress Limit Control Register

Address: 0xN422 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Egress Data Rate Limit for Queue 2 Frames Note that the update will not take effect until the Port Queue 3 Egress Limit Control Register is written. | R/W | 0x00 |

5.2.4.15 Port Queue 3 Egress Limit Control Register

Address: 0xN423 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Egress Data Rate Limit for Queue 3 Frames | R/W | 0x00 |

5.2.5 PORT N: PORT SWITCH MIB COUNTERS REGISTERS (0xN500 - 0xN5FF)

Thirty MIB counters are provided for each port. They are accessed via the following two indirect registers. MIB Counters are read-clear. Refer to [Section 5.3.6, "Management Information Base \(MIB\) Counters"](#) for additional details.

Note: The [Switch MIB Control Register](#) contains two additional bits for global control of MIB counters.

5.2.5.1 Port MIB Control and Status Register

Address: 0xN500 - 0xN503 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|-----------|---------|
| 31 | MIB Counter Overflow Indication 1 = Counter has overflowed 0 = Counter has not overflowed | RO | 0b |
| 30:26 | RESERVED | RO | 0x00 |
| 25 | MIB Read Enable / Count Valid Set this bit to '1' to initiate a counter read. When the counter read is complete and the count value is available in the MIB Counter Value, this bit will automatically clear to '0'. | R/W SC | 0b |
| 24 | MIB Flush and Freeze Enable Refer to the Switch MIB Control Register for the associated global MIB control bits for flushing and freezing port counters. 1 = Enable MIB counter flush and freeze function for this port 0 = Disable MIB counter flush and freeze function for this port | R/W | 0b |
| 23:16 | MIB Index | R/W | 0x00 |
| 15:4 | RESERVED | RO | 0x000 |
| 3:0 | MIB Counter Value [35:32] | RO | 0x0 |

KSZ9896C

5.2.5.2 Port MIB Data Register

Address: 0xN504 - 0xN507 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---------------------------------|------|---------|
| 31:0 | MIB Counter Value [31:0] | RO | 0x0000 |

5.2.6 PORT N: PORT SWITCH ACL CONTROL REGISTERS (0xN600 - 0xN6FF)

An Access Control List (ACL) Table is provided for each port. It is accessed via the following two indirect registers. Refer to [Section 5.3.5, "Access Control List \(ACL\) Table"](#) for additional details.

5.2.6.1 Port ACL Access 0 Register

Address: 0xN600 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--------------------------------|------|---------|
| 7:4 | RESERVED | RO | 0x0 |
| 3:0 | First Rule Number (FRN) | R/W | 0x0 |

5.2.6.2 Port ACL Access 1 Register

Address: 0xN601 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:6 | RESERVED | R/W | 00b |
| 5:4 | Mode 00 = No action taken 01 = Layer 2 MAC header filtering 10 = Layer 3 IP address filtering 11 = Layer 4 TCP port number / IP protocol filtering | R/W | 00b |
| 3:2 | Enable | R/W | 00b |
| 1 | Source / Destination 1 = Source 0 = Destination | R/W | 0b |
| 0 | Compare Equal 1 = Match if the compared values are equal 0 = Match if the compared values are not equal | R/W | 0b |

5.2.6.3 Port ACL Access 2 Register

Address: 0xN602 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [47:40] | R/W | 0x00 |

5.2.6.4 Port ACL Access 3 Register

Address: 0xN603 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [39:32] | R/W | 0x00 |

5.2.6.5 Port ACL Access 4 Register

Address: 0xN604 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [31:24] | R/W | 0x00 |

5.2.6.6 Port ACL Access 5 Register

Address: 0xN605 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|----------------------------|------|---------|
| 7:0 | MAC Address [23:16] | R/W | 0x00 |

5.2.6.7 Port ACL Access 6 Register

Address: 0xN606 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---------------------------|------|---------|
| 7:0 | MAC Address [15:8] | R/W | 0x00 |

KSZ9896C

5.2.6.8 Port ACL Access 7 Register

Address: 0xN607 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--------------------------|------|---------|
| 7:0 | MAC Address [7:0] | R/W | 0x00 |

5.2.6.9 Port ACL Access 8 Register

Address: 0xN608 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|-------------------------|------|---------|
| 7:0 | EtherType [15:8] | R/W | 0x00 |

5.2.6.10 Port ACL Access 9 Register

Address: 0xN609 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|------------------------|------|---------|
| 7:0 | EtherType [7:0] | R/W | 0x00 |

5.2.6.11 Port ACL Access A Register

Address: 0xN60A Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:6 | Priority Mode (PM) | R/W | 00b |
| 5:3 | Priority | R/W | 00_0b |
| 2 | Remark Priority Enable (RPE) | R/W | 0b |
| 1:0 | Remark Priority [2:1] The two MSB of the 3-bit Remark Priority field. Refer to the Port ACL Access B Register for the LSB. | R/W | 00b |

5.2.6.12 Port ACL Access B Register

Address: 0xN60B Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | Remark Priority [0] The LSB of the 3-bit Remark Priority field. Refer to the Port ACL Access B Register for the MSB. | R/W | 0b |
| 6:5 | Map Mode (MM) | R/W | 00b |
| 4:0 | RESERVED | R/W | 0b |

5.2.6.13 Port ACL Access C Register

Address: 0xN60C Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|-----------------|------|---------|
| 7:0 | RESERVED | R/W | 0x00 |

5.2.6.14 Port ACL Access D Register

Address: 0xN60D Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|----------|
| 7:6 | RESERVED | R/W | 00b |
| 5:0 | Port Forward Map Each bit corresponds to a forwarding port. Bit 0 corresponds to port 1, Bit 1 corresponds to port 2, etc. 1 = enable forwarding to that port 0 = do not forward to that port | R/W | 00_0000b |

5.2.6.15 Port ACL Access E Register

Address: 0xN60E Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|-----------------------|------|---------|
| 7:0 | Ruleset [15:8] | R/W | 0x00 |

KSZ9896C

5.2.6.16 Port ACL Access F Register

Address: 0xN60F
Port N: 1-6 Size: 8 bits

| Bits | Description | Type | Default |
|------|----------------------|------|---------|
| 7:0 | Ruleset [7:0] | R/W | 0x00 |

5.2.6.17 Port ACL Byte Enable MSB Register

Address: 0xN610
Port N: 1-6 Size: 8 bits

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:0 | Byte Enable [15:8] Each bit enables accessing one of the ACL bytes when a read or write is initiated by writing to the Port ACL Byte Enable LSB Register . Bit 0 applies to the Port ACL Access 7 Register Bit 1 applies to the Port ACL Access 6 Register , etc. Bit 7 applies to the Port ACL Access 0 Register 1 = Byte is selected for read/write 0 = Byte is not selected | R/W | 0x00 |

5.2.6.18 Port ACL Byte Enable LSB Register

Address: 0xN611
Port N: 1-6 Size: 8 bits

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:0 | Byte Enable [7:0] Each bit enables accessing one of the ACL bytes when a read or write is initiated by writing to the Port ACL Byte Enable LSB Register . Bit 0 applies to the Port ACL Access F Register Bit 1 applies to the Port ACL Access E Register , etc. Bit 7 applies to the Port ACL Access 8 Register 1 = Byte is selected for read/write 0 = Byte is not selected | R/W | 0x00 |

5.2.6.19 Port ACL Access Control 0 Register

Address: 0xN612 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | RESERVED | R/O | 0b |
| 6 | Write Status 1 = Write operation is complete 0 = Write operation is not complete | RO | 1b |
| 5 | Read Status 1 = Read operation is complete 0 = Read operation is not complete; continue polling until this bit is set before reading results registers | RO | 1b |
| 4 | Write / Read 1 = Write 0 = Read | R/W | 0b |
| 3:0 | ACL Index | R/W | 0x0 |

KSZ9896C

5.2.7 PORT N: PORT SWITCH INGRESS CONTROL REGISTERS (0xN800 - 0xN8FF)

5.2.7.1 Port Mirroring Control Register

Address: 0xN800 Size: 8 bits
Port N: 1-6

This register contains the port controls for port mirroring. The [Global Port Mirroring and Snooping Control Register](#) must also be properly configured.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6 | Receive Sniff 1 = All packets received on this port are designated as “monitored packets” and will be forwarded to the designated “sniffer port”. 0 = No receive monitoring. | R/W | 0b |
| 5 | Transmit Sniff 1 = All packets transmitted on this port are designated as “monitored packets” and will be forwarded to the designated “sniffer port”. 0 = No transmit monitoring. | R/W | 0b |
| 4:2 | RESERVED | RO | 0_00b |
| 1 | Sniffer Port 1 = This port is designated as the sniffer port and will transmit monitored packets. 0 = Not a sniffer port. Normal operation. | R/W | 0b |
| 0 | RESERVED | RO | 0b |

5.2.7.2 Port Priority Control Register

Address: 0xN801 Size: 8 bits
Port N: 1-6

To achieve the desired functionality, do not set more than one bit at a time in this register.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | Highest Priority 1 = Highest priority selected 0 = Highest priority not selected | R/W | 0b |
| 6 | OR'ed Priority 1 = All available priority OR'ed 0 = All available priority not OR'ed | R/W | 0b |
| 5 | RESERVED | RO | 0b |
| 4 | MAC Address Priority Classification 1 = Enable MAC address priority classification for ingress packets on port 0 = Disable MAC address classification | R/W | 0b |

| Bits | Description | Type | Default |
|------|--|------|---------|
| 3 | VLAN Priority Classification 1 = Enable VLAN priority classification for ingress packets on port 0 = Disabled VLAN classification | R/W | 0b |
| 2 | 802.1p Priority Classification 1 = Enable 802.1p priority classification for ingress packets on port 0 = Disable 802.1p priority classification | R/W | 0b |
| 1 | Diffserv Priority Classification 1 = Enable Diffserv priority classification for ingress packets on port 0 = Disable Diffserv priority classification | R/W | 0b |
| 0 | ACL Priority Classification 1 = Enable ACL priority classification for ingress packets on port 0 = Disable ACL priority classification | R/W | 0b |

5.2.7.3 Port Ingress MAC Control Register

Address: 0xN802 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | User Priority Ceiling 1 = If a packet's "user priority field" is greater than the "user priority field" in the port default tag register, replace the packet's user priority field with the user priority field in the port default tag from the Port Default Tag 0 Register . 0 = Do not replace the user priority field at this port. | R/W | 0b |
| 6:5 | RESERVED | RO | 0b |
| 4 | Discard Untagged Packets This bit should be left as 0 for the host port. 1 = Any ingress packet at this port without an IEEE 802.1Q tag is discarded. 0 = Do not discard. | R/W | 0b |
| 3 | Discard Tagged Packets 1 = Any ingress packet at this port with an IEEE 802.1Q tag is discarded. 0 = Do not discard. | R/W | 0b |
| 2:0 | Port Default Priority Classification If Diffserv classification, 802.1p classification and VLAN classification are not enabled or fail to classify, ingress packets at this port will be classified with the priority specified here. Possible values are 0 to 7. | R/W | 000b |

KSZ9896C

5.2.7.4 Port Authentication Control Register

Address: 0xN803 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:3 | RESERVED | RO | 0000_0b |
| 2 | Access Control List (ACL) Enable 1 = enable 0 = enable | R/O | 0b |
| 1:0 | Authentication Mode 00 = Pass Mode. Authentication is disabled. When ACL is enabled, all traffic that misses the ACL rules is forwarded; otherwise ACL actions apply. 01 = Block Mode. Authentication is enabled. When ACL is enabled, all traffic that misses the ACL rules is blocked; otherwise ACL actions apply. 10 = Trap Mode. Authentication is enabled. All traffic is forwarded to the host port. When ACL is enabled, all traffic that misses the ACL rules is blocked; otherwise ACL actions apply. 11 = Reserved | R/W | 00b |

5.2.7.5 Port Pointer Register

Address: 0xN804 - 0xN807 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 31:19 | RESERVED | RO | 0x0000 |
| 18:16 | Port Index Used to point to the indirect mapping locations for the 6 ports | R/W | 000b |
| 15:2 | RESERVED | RO | 0x0000 |
| 1:0 | Queue Pointer Used to point to the indirect mapping locations for the 4 queues | R/W | 00b |

5.2.7.6 Port Priority to Queue Mapping Register

Address: 0xN808 - 0xN80B Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31:28 | Regenerated priority (queue) value for priority 7. The 2 most significant bits are reserved. | R/W | xx11b |
| 27:24 | Regenerated priority (queue) value for priority 6. The 2 most significant bits are reserved. | R/W | xx11b |
| 23:20 | Regenerated priority (queue) value for priority 5. The 2 most significant bits are reserved. | R/W | xx10b |
| 19:16 | Regenerated priority (queue) value for priority 4. The 2 most significant bits are reserved. | R/W | xx10b |
| 15:12 | Regenerated priority (queue) value for priority 3. The 2 most significant bits are reserved. | R/W | xx01b |
| 11:8 | Regenerated priority (queue) value for priority 2. The 2 most significant bits are reserved. | R/W | xx01b |
| 7:4 | Regenerated priority (queue) value for priority 1. The 2 most significant bits are reserved. | R/W | xx00b |
| 3:0 | Regenerated priority (queue) value for priority 0. The 2 most significant bits are reserved. | R/W | xx00b |

KSZ9896C

5.2.7.7 Port Police Control Register

Address: 0xN80C - 0xN80F Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 31:12 | RESERVED | RO | 0x00000 |
| 11 | Dropped Color 1 = Color packet PMON holds dropped packets of that color 0 = Color packet PMON holds all packets of that color | R/W | 0b |
| 10 | Drop All 1 = All packets are dropped while max threshold is exceeded in PM WRED 0 = Drop packet based on WRED_PM_PROB_MULTIPLIER | R/W | 0b |
| 9:8 | Packet Type PMON packet type to be read for the connection from Queue Pointer of the port pointed by Port Index 11 = WRED_PMON holds the number of RED packets while read 10 = WRED_PMON holds the number of YELLOW packets while read 01 = WRED_PMON holds the number of GREEN packets while read 00 = WRED_PMON holds the number of dropped packets while read | R/W | 00b |
| 7 | Port Based Policing 1 = Policing is per-port per-queue based 0 = Policing is per-queue based only, Port Index = 0 is used to set aggregated CIR, PIR for each queue | R/W | 0b |
| 6:5 | NONDSCP_COLOR Color of non-IP frame for color aware | R/W | 01b |
| 4 | Color Mark Enable 1 = DSCP color mark enable 0 = DSCP color mark disable | R/W | 0b |
| 3 | Color Remap Enable 1 = DSCP color remap enable for color aware 0 = DSCP color remap disable, use DSCP color | R/W | 0b |
| 2 | Drop SRP 1 = Allow drop SRP packets while WRED is enabled 0 = Don't allow drop SRP packet while WRED is enabled | R/W | 0b |
| 1 | Police Mode This bit sets the policing modes for the queue if policing is enabled 1 = Color blind 0 = Color aware mode | R/W | 0b |
| 0 | Police Enable 1 = Enable policing and WRED 0 = Disable policing and WRED | R/W | 0b |

5.2.7.8 Port Police Queue Rate Register

Address: 0xN820 - 0xN823 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 31:16 | Committed Information Rate Specifies the committed information rate for the connection from the Queue Pointer of the port pointed by the Port index. | R/W | 0x1000 |
| 15:0 | Peak Information Rate Specifies the queue's peak information rate for the connection from the Queue Pointer of the port pointed by the Port index. | R/W | 0x2000 |

5.2.7.9 Port Police Queue Burst Size Register

Address: 0xN824 - 0xN827 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31:16 | Committed Burst Size Specifies the queue's committed burst size in bytes that is supported for the connection from the Queue Pointer of the port pointed by the Port index. | R/W | 0x1000 |
| 15:0 | Peak Information Rate Specifies the queue's peak burst size in bytes that is supported for the connection from the Queue Pointer of the port pointed by the Port index. | R/W | 0x3000 |

5.2.7.10 Port WRED Packet Memory Control Register 0

Address: 0xN830 - 0xN833 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31:27 | RESERVED | RO | 00000b |
| 26:16 | WRED Packet Memory Maximum Threshold | R/W | 0x400 |
| 15:11 | RESERVED | RO | 00000b |
| 10:0 | WRED Packet Memory Minimum Threshold | R/W | 0x080 |

KSZ9896C

5.2.7.11 Port WRED Packet Memory Control Register 1

Address: 0xN834 - 0xN837 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31:27 | RESERVED | RO | 00000b |
| 26:16 | WRED Packet Memory Probability Multiplier | R/W | 0x020 |
| 15:11 | RESERVED | RO | 00000b |
| 10:0 | WRED Packet Memory Average Queue Size | R | - |

5.2.7.12 Port WRED Queue Control Register 0

Address: 0xN840 - 0xN843 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|--|------|---------|
| 31:27 | RESERVED | RO | 00000b |
| 26:16 | WRED Maximum Queue Threshold WRED maximum threshold for the connection from the Queue Pointer of the port pointed to by the Port index. | R/W | 0x080 |
| 15:11 | RESERVED | RO | 00000b |
| 10:0 | WRED Minimum Queue Threshold WRED minimum threshold for the connection from the Queue Pointer of the port pointed to by the Port index. | R/W | 0x009 |

5.2.7.13 Port WRED Queue Control Register 1

Address: 0xN844 - 0xN847 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31:27 | RESERVED | RO | 00000b |
| 26:16 | WRED Queue Probability Multiplier Probability multiplier for the connection from the Queue Pointer of the port pointed to by the Port index. | R/W | 0x010 |
| 15:11 | RESERVED | RO | 00000b |
| 10:0 | WRED Packet Memory Average Queue Size Average queue size for the connection from the Queue Pointer of the port pointed to by the Port index. | R | - |

5.2.7.14 Port WRED Queue Performance Monitor Control Register

Address: 0xN848 - 0xN84B Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|-------|---|------|---------|
| 31 | Random Drop Enable 1 = Random Drop enabled 0 = Random Drop disabled | R/W | 0b |
| 30 | PMON Flush 1 = PMON counters are flushed. This is a low to high edge triggered clear. 0 = No activity | R/W | 0x010 |
| 29 | GYR Drop Disable 1 = Drop GREEN/YELLOW/RED is disabled 0 = Drop GREEN/YELLOW/RED is enabled | R/W | 0b |
| 28 | YR Drop Disable 1 = Drop YELLOW/RED is disabled 0 = Drop YELLOW/RED is enabled | R/W | 0b |
| 27 | R Drop Disable 1 = Drop RED is disabled 0 = Drop RED is enabled | R/W | 0b |
| 26 | Drop All 1 = Drop all packets wile the max threshold is exceeded 0 = Drop based on WRED Queue Probability Multiplier | R/W | 0b |
| 25:24 | RESERVED | RO | 00b |
| 23:0 | Packet Event Counter Packer event count for the connection from the Queue Pointer of the port pointed to by the Port index. | R | 0x00000 |

KSZ9896C

5.2.8 PORT N: PORT SWITCH EGRESS CONTROL REGISTERS (0xN900 - 0xN9FF)

5.2.8.1 Port Transmit Queue Index Register

Address: 0xN900 - 0xN903 Size: 32 bits
Port N: 1-6

This register holds an index value that is used when accessing several registers that follow. This register may be accessed as an 8-bit register, with bits [7:0] located at address 0xN903.

| Bits | Description | Type | Default |
|------|--|------|------------|
| 31:2 | RESERVED | RO | 0x00000000 |
| 1:0 | Queue Index Points to the queue number for subsequent queue configuration registers. | R/W | 00 |

5.2.8.2 Port Transmit Queue PVID Register

Address: 0xN904 - 0x907 Size: 32 bits
Port N: 1-6

This register may be accessed as an 8-bit register, with bits [7:0] located at address 0xN903.

| Bits | Description | Type | Default |
|------|---|------|------------|
| 31:1 | RESERVED | RO | 0x00000000 |
| 0 | Port VID Replacement 1 = For any egress packet with a non-zero VLAN ID, replace the VID with the port default VID from the Port Default Tag 0 Register and Port Default Tag 1 Register . If double tagging is enabled, it is the ISP tag ID that will be replaced. 0 = Do not replace the VID. | R/W | 0b |

5.2.8.3 Port Transmit Queue Control 0 Register

Address: 0xN914 Size: 8 bits
Port N: 1-6

This register is indexed. Settings are applied on a per-queue and per-port basis. Specify the target queue number in the [Port Transmit Queue Index Register](#) before accessing this register.

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:6 | Scheduler Mode Determines the egress scheduling policy when 2 or 4 transmit queues are enabled. 00 = Strict Priority. All packets in queue 3 are transmitted before any packets from any lower number queue will be transmitted. All packets in queue 2 are transmitted before any packets from any lower number queue, etc. 10 = Weighted Round Robin (WRR). As determined by the weight values in the (queue indexed) Port Transmit Queue Control 1 Register , a limited number of packets are transmitted from a queue before the next queue is serviced. All queues are serviced in turn. 01 and 11 = RESERVED | R/W | 10b |
| 5:0 | RESERVED | RO | 000000b |

5.2.8.4 Port Transmit Queue Control 1 Register

Address: 0xN915 Size: 8 bits
Port N: 1-6

This register is indexed. Settings are applied on a per-queue and per-port basis. Specify the target queue number in the [Port Transmit Queue Index Register](#) before accessing this register.

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7 | RESERVED | RO | 0b |
| 6:0 | Queue Weight for WRR Scheduling During weighted round robin (WRR) scheduling, this value specifies the number of packets that may be transmitted from the particular queue before the next queue is serviced. These values are not relevant when the port is configured for a single egress queue. Do not assign 0 to any queue. Example values for a 4-queue configuration are 8 (for queue 3), 4 (for queue 2), 2 (for queue1) and 1 (for queue 0). | R/W | 0000001 |

KSZ9896C

5.2.9 PORT N: PORT SWITCH QUEUE MANAGEMENT CONTROL REGISTERS (0xNA00 - 0xNAFF)

5.2.9.1 Port Control 0 Register

Address: 0xNA00 - 0xNA03 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|-----------|
| 32:2 | RESERVED | RO | 0x0000000 |
| 1:0 | Drop Mode 00 = No drop, issue flow control to this port when destination is in congestion state. 01 = Drop packets with priority 0 sourced from this port when destination is in congestion state and disable the flow control to this port. 10 = Drop packets with priority 0, 1 sourced from this port when destination is in congestion state and disable the flow control to this port. 11 = Drop packets with priority 0, 1, 2 sourced from this port when destination is in congestion state and disable the flow control to this port. | R/W | 00b |

5.2.9.2 Port Control 1 Register

Address: 0xNA04 - 0xNA07 Size: 32 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|-----------|
| 31:7 | RESERVED | RO | 0x0000000 |
| 6:0 | Port VLAN Membership Each bit corresponds to a device port. This feature does not utilize VLAN tags or the VLAN Table, and is unrelated to tag-based VLAN functions. Also refer to bit 1 in the Queue Management Control 0 Register . Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Frames may be forwarded to the corresponding port 0 = Frames are blocked from being forwarded to corresponding port | R/W | 0x7F |

5.2.10 PORT N: PORT SWITCH ADDRESS LOOKUP CONTROL REGISTERS (0xNB00 - 0xNBFF)

5.2.10.1 Port Control 2 Register

Address: 0xNB00 Size: 8 bits
 Port N: 1-6

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7 | Null VID Lookup Enable How to handle tagged packets with VID = 0. 1 = Non-standard operation. VLAN Table lookup using VID = 0. 0 = Standard operation. VLAN Table lookup using the port default VID (Port Default Tag 0 Register and Port Default Tag 1 Register) | R/W | 0b |
| 6 | Ingress VLAN Filtering 1 = Discard packets whose VID port membership in the VLAN Table does not include the ingress port 0 = No ingress filtering | R/W | 0b |
| 5 | Discard Non-PVID Packet 1 = Discard packets whose VID does not match the ingress port default VID 0 = Do not compare VID to the port default VID | R/W | 0b |
| 4 | MAC Based 802.1X Enable 1 = Enable MAC based 802.1X authentication in lookup engine 0 = Only use ACL (if enabled) to perform MAC based authentication | R/W | 0b |
| 3 | Self-Address Filtering – Port Enable The source address of received packets is compared to the MAC address in the Switch MAC Address 0 Register through Switch MAC Address 5 Register , and the packet is dropped if there is a match. Both this port enable bit and the global enable bit in the Switch Lookup Engine Control 1 Register must be set to enable self-address filtering. 1 = Enable self-address filtering for this port. 0 = Disable self-address filtering for this port. | R/W | 0b |
| 2 | RESERVED | RO | 0b |
| 1 | RESERVED | R/W | 0b |
| 0 | RESERVED | RO | 0b |

KSZ9896C

5.2.10.2 Port MSTP Pointer Register

Address: 0xNB01 Size: 8 bits
Port N: 1-6

| Bits | Description | Type | Default |
|------|--|------|---------|
| 7:3 | RESERVED | R/W | 0000_0b |
| 2:0 | MSTP Pointer Points to one of the 8 MSTPs. It is used as an index when reading or writing the Port MSTP State Register . | R/W | 000b |

5.2.10.3 Port MSTP State Register

Address: 0xNB04 Size: 8 bits
Port N: 1-6

There are eight MSTPs, and the MSTP Pointer in the [Port MSTP Pointer Register](#) is used as an index to select a particular MSTP when reading or writing this register.

| Bits | Description | Type | Default |
|------|---|------|---------|
| 7:3 | RESERVED | RO | 0000_0b |
| 2 | Port Transmit Enable 1 = Enable packet transmission on the port 0 = Disable packet transmission on the port | R/W | 1b |
| 1 | Port Receive Enable 1 = Enable packet reception on the port 0 = Disable packet reception on the port | R/W | 1b |
| 0 | Port Learning Disable 1 = Disable switch address learning on the port 0 = Enable switch address learning on the port | R/W | 0b |

5.3 Tables and MIB Counters (Access)

Indirect address and data registers are used to access the various tables and counters:

- [Address Lookup \(ALU\) Table](#)
- [Static Address Table](#)
- [Reserved Multicast Address Table](#)
- [VLAN Table](#)
- [Access Control List \(ACL\) Table](#)
- [Management Information Base \(MIB\) Counters](#)

5.3.1 ADDRESS LOOKUP (ALU) TABLE

The Address Lookup Table is the largest of three tables used for MAC address lookup. It supports both dynamic and static MAC address entries. In response to a destination address (DA) lookup, all tables are searched to make a packet forwarding decision. In response to a source address (SA) lookup, only this table is searched for aging, migration and learning of the dynamic entries.

It is suggested that static address entries be programmed to the Static Address Table. When that table is full, additional static address entries may be programmed into this table. Static entries will not be aged out.

A static DA lookup result (in either this table or the [Static Address Table](#)) takes precedence over the dynamic DA lookup result.

This table is a 4-way associative memory, with 1K buckets, for a total of 4K entries. In normal operation, the MAC address (and optionally the FID) are hashed to generate a 10-bit index. The 10-bit index specifies a bucket of up to 4 entries, but the entries within each bucket are not individually addressable. A new entry can be added to a bucket if the bucket contains 3 or fewer valid entries. Once a bucket is full with 4 valid entries, any new entry will overwrite the least-recently-used dynamic entry.

If a bucket contains 2 or 3 static entries, adding an additional static entry to that bucket will generate an Almost Full interrupt. (Refer to the [Address Lookup Table Interrupt Register](#) and [Global Interrupt Status Register](#)). Once the interrupt is generated, the 12-bit absolute address of the new static entry will be available for reading in the [Address Lookup Table Entry Index 0 Register](#) bits [11:0].

If a bucket is full with 4 static entries, attempting to write an additional static entry will fail and will result in a Write Fail Interrupt. Once the interrupt is generated, the 10-bit index of the full bucket will be available for reading in the [Address Lookup Table Entry Index 0 Register](#) bits [9:0].

A bucket with 4 static entries will also block any dynamic entries from being learned. Such a failure will generate a Learn Fail Interrupt. Once the interrupt is generated, the 10-bit index of the full bucket will be available for reading in the [Address Lookup Table Entry Index 1 Register](#) bits [9:0].


This table is accessed indirectly via the ALU Table registers:

- [ALU Table Index 0 Register](#)
- [ALU Table Index 1 Register](#)
- [ALU Table Access Control Register](#)
- [ALU Table Entry 1 Register](#)
- [ALU Table Entry 2 Register](#)
- [ALU Table Entry 3 Register](#)
- [ALU Table Entry 4 Register](#)

All static and dynamic entries may be read, but only static entries may be written. Because most of the table entries are created automatically by hardware address learning, and because the table is 4-way associative, three different methods are available for accessing the memory:

1. Lookup by MAC address. This utilizes the same hash tag lookup circuitry that is used for ingress packet forwarding lookup. The hash tag algorithm is specified by the Hash Option bits in the [Switch Lookup Engine Control 0 Register](#). This is normally how static entries should be written (Read or write).
2. Direct addressing (Read or write).
3. Sequential search, returning all valid entries (Read only).

FIGURE 5-3: ADDRESS LOOKUP TABLE CONFIGURATION



5.3.1.1 Address Lookup Table Read Operation

1. Write the index to the [ALU Table Index 0 Register](#) and [ALU Table Index 1 Register](#).
 - a) A MAC address is used for the hashing indexing function, which is the normal access mechanism for this table.
 - b) A FID is also required if VLAN is enabled.
 - c) If directly addressing the table, a 12-bit index is written to bits [11:0] of the MAC address field.
2. Write to the [ALU Table Access Control Register](#).
 - a) Set the ACTION field to 10 to indicate a read operation.
 - b) Select the addressing method via the DIRECT bit.
 - c) Set the START_FINISH bit to initiate the operation.
3. Read / poll the [ALU Table Access Control Register](#).
 - a) The START_FINISH bit transitions to 0 to indicate that the operation is complete.
 - b) Once START_FINISH is 0, the VALID bit indicates whether or not a valid entry was found. If VALID is false, there is no need to proceed to step 4.
 - c) The VALID_ENTRY_OR_SEARCH_END bit is an aggregate of the START_FINISH bit and the VALID bit. It is intended for use when accessing registers by in-band management (IBA), where polling multiple bits is less practical.
4. Read the contents of the returned table entry from the [ALU Table Entry 1 Register](#), [ALU Table Entry 2 Register](#), [ALU Table Entry 3 Register](#), and [ALU Table Entry 4 Register](#). If no VALID entry is returned, these registers will contain all zeros.

5.3.1.2 Address Lookup Table Search Operation

The second method to access the Address Lookup Table is through the search operation. The entire table is searched sequentially, revealing each valid entry. Invalid address entries are skipped, providing an efficient way to search the entire table. Setting the START_FINISH bit in the [ALU Table Access Control Register](#) begins the search from the top of the table. This bit is cleared when the search is complete. During the table search, the VALID bit in the [ALU Table Access Control Register](#) is cleared when the search starts. The VALID bit indicates when a found valid entry is available in the ALU Table Entry registers (0x0420 - 0x042F). After reading the last ALU Table Entry register ([ALU Table Entry 4 Register](#)), the search process automatically continues to seek the next valid entry in the address table. The START_FINISH bit remains set until all the valid entries in the table have been returned. The search can be stopped any time by setting the [ALU Table Access Control Register](#) START_FINISH bit to 0.

1. Write to the [ALU Table Access Control Register](#).
 - a) Set the ACTION field to 11 to indicate a search operation.
 - b) Set the START_FINISH bit to initiate the operation.
2. Poll the VALID bit until it is set.
 - a) If register access is by in-band management (IBA) rather than SPI or I²C, poll VALID_ENTRY_OR_SEARCH_END instead. This bit goes high to indicate either a new valid entry is returned or the search is complete.
3. Read the entry from the [ALU Table Entry 1 Register](#), [ALU Table Entry 2 Register](#), [ALU Table Entry 3 Register](#), and [ALU Table Entry 4 Register](#), in that order.
 - a) If the search has ended and there are no more valid entries to read, these registers will return all zeros. In this case, go to step 5.
4. If START_FINISH = 0, go to step 5; else go to step 2.
5. Read VALID_COUNT to verify the number of valid entries.

5.3.1.3 Address Lookup Table Write Operation

1. Perform a read operation to get the contents of the current entry. The values are kept in the [ALU Table Entry 1 Register](#), [ALU Table Entry 2 Register](#), [ALU Table Entry 3 Register](#), and [ALU Table Entry 4 Register](#).
2. Modify the correct entry as necessary. Set the STATIC bit so that the entry is not aged out.
3. Write to the [ALU Table Access Control Register](#).
 - a) Set the ACTION field to 01 to indicate a write operation.
 - b) Select the addressing method via the DIRECT bit.
 - c) Set the START_FINISH bit to initiate the operation.

KSZ9896C

5.3.1.4 ALU Table Entry 1 Register

Address: 0x0420 - 0x0423 Size: 32 bits

| Bits | Description | Type | Default |
|-------|---|------|----------|
| 31 | STATIC 1 = Entry is static and updated by a host processor; will not be aged out 0 = Entry is dynamically learned and aged | R/W | 0b |
| 30 | SRC FILTER 1 = Drop packet if source address match during source learning 0 = Don't drop if source address match | R/W | 0b |
| 29 | DES FILTER 1 = Drop packet if destination address match during lookup 0 = Don't drop if destination address match | R/W | 0b |
| 28:26 | PRIORITY (for static entries) AGE COUNT (for dynamic entries) >0 = Entry has been accessed or learned since last aging process. A default value is reloaded every time the entry is learned or accessed. It is decremented during aging process. 0 = Entry has not been accessed or learned since last aging process. Entry is not valid if it's not static. | R/W | 0_00b |
| 25:3 | RESERVED | RO | 0x000000 |
| 2:0 | MSTP Multiple Spanning Tree Protocol group ID for matching | R/W | 000b |

5.3.1.5 ALU Table Entry 2 Register

Address: 0x0424 - 0x0427 Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31 | OVERRIDE 1 = Enable overriding of port state 0 = Do not enable | R/W | 0b |
| 30:6 | RESERVED | RO | 0x000000 |
| 5:0 | PORT FORWARD Each bit corresponds to a device port. Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward to the corresponding port 0 = Do not forward to the corresponding port | R/W | 0x00 |

5.3.1.6 ALU Table Entry 3 Register

Address: 0x0428 - 0x042B Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|----------|
| 31:23 | RESERVED | RO | 0x000 |
| 22:16 | FID VLAN group ID for matching | R/W | 000_0000 |
| 15:0 | MAC Address [47:32] | R/W | 0x0000 |

5.3.1.7 ALU Table Entry 4 Register

Address: 0x042C - 0x042F Size: 32 bits

| Bits | Description | Type | Default |
|------|---------------------------|------|------------|
| 31:0 | MAC Address [31:0] | R/W | 0x00000000 |

5.3.2 STATIC ADDRESS TABLE

The Static Address Table is one of three tables used for MAC address lookup. It can hold up to 16 static address entries, thereby minimizing the number of static entries that may need to be programmed into the Address Lookup Table, which is used primarily for dynamic entries. In response to a destination address (DA) lookup, all tables are searched to make a packet forwarding decision. Entries in this table are programmed by the host processor, and are never aged.

A static DA lookup result (in either this table or the Address Lookup Table) takes precedence over the dynamic DA lookup result.

The Static Address Table has 16 entries and is accessed indirectly. The [Static Address and Reserved Multicast Table Control Register](#) is used for indexing and read/write control. The following registers are used for the data fields:

- [Static Address Table Entry 1 Register](#)
- [Static Address Table Entry 2 Register](#)
- [Static Address Table Entry 3 Register](#)
- [Static Address Table Entry 4 Register](#)

5.3.2.1 Static Address Table Write Operation

1. Write the content of the table entry to the [Static Address Table Entry 1 Register](#), [Static Address Table Entry 2 Register](#), [Static Address Table Entry 3 Register](#), and [Static Address Table Entry 4 Register](#).
2. Write to the [Static Address and Reserved Multicast Table Control Register](#).
 - a) Write the TABLE_INDEX field with the 4-bit index value.
 - b) Set the TABLE_SELECT bit to 0 to select the Static Address Table.
 - c) Set the ACTION bit to 0 to indicate a write operation.
 - d) Set the START_FINISH bit to 1 to initiate the operation.
3. When the operation is complete, the START_FINISH bit will be cleared automatically.

5.3.2.2 Static Address Table Read Operation

1. Write to the [Static Address and Reserved Multicast Table Control Register](#).
 - a) Write the TABLE_INDEX field with the 4-bit index value.
 - b) Set the TABLE_SELECT bit to 0 to select the Static Address Table.

KSZ9896C

- c) Set the ACTION bit to 1 to indicate a read operation.
- d) Set the START_FINISH bit to 1 to initiate the operation.
2. When the operation is complete, the START_FINISH bit will be cleared automatically.
 - a) Read the contents of the indexed entry from the [Static Address Table Entry 1 Register](#), [Static Address Table Entry 2 Register](#), [Static Address Table Entry 3 Register](#), and [Static Address Table Entry 4 Register](#).

5.3.2.3 Static Address Table Entry 1 Register

Address: 0x0420 - 0x0423 Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|----------|
| 31 | VALID 1 = Entry is valid 0 = Entry is not valid | R/W | 0b |
| 30 | SRC FILTER 1 = Drop packet if source address match during source learning 0 = Don't drop if source address match | R/W | 0b |
| 29 | DES FILTER 1 = Drop packet if destination address match during lookup 0 = Don't drop if destination address match | R/W | 0b |
| 28:26 | PRIORITY | R/W | 0_00b |
| 25:3 | RESERVED | RO | 0x000000 |
| 2:0 | MSTP Multiple Spanning Tree Protocol group ID for matching | R/W | 000b |

5.3.2.4 Static Address Table Entry 2 Register

Address: 0x0424 - 0x0427 Size: 32 bits

| Bits | Description | Type | Default |
|------|--|------|----------|
| 31 | OVERRIDE 1 = Enable overriding of port state 0 = Do not enable | R/W | 0b |
| 30 | USE FID Use FID on multicast packets for matching | R/W | 0b |
| 29:6 | RESERVED | RO | 0x000000 |
| 5:0 | PORT FORWARD Each bit corresponds to a device port. Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward to the corresponding port 0 = Do not forward to the corresponding port | R/W | 0x00 |

5.3.2.5 Static Address Table Entry 3 Register

Address: 0x0428 - 0x042B Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|-----------|
| 31:23 | RESERVED | RO | 0x000 |
| 22:16 | FID VLAN group ID for matching | R/W | 000_0000b |
| 15:0 | MAC Address [47:32] | R/W | 0x0000 |

5.3.2.6 Static Address Table Entry 4 Register

Address: 0x042C - 0x042F Size: 32 bits

| Bits | Description | Type | Default |
|------|---------------------------|------|------------|
| 31:0 | MAC Address [31:0] | R/W | 0x00000000 |

5.3.3 RESERVED MULTICAST ADDRESS TABLE

The Reserved Multicast Address Table determines the forwarding ports for 48 specific multicast addresses. The table is addressed by the least significant 6 bits of the multicast address, and the table contents are the bits (the PORT_FORWARD field) that represent each possible forwarding port of the device. It is not addressed by the group number in the first column of [Table 4-6](#). Note that the 48 addresses are organized into 8 fixed groups, and changing a forwarding port for one address also makes the same change for all other addresses in the same group.

The Reserved Multicast Table is accessed in the same manner as the [Static Address Table](#), using the same indirect access registers. The [Static Address and Reserved Multicast Table Control Register](#) is used for indexing and read/write control, while the [Reserved Multicast Address Table Entry 2 Register](#) is used for the data fields.

5.3.3.1 Reserved Multicast Table Write Operation

- Write the PORT_FORWARD value to the [Reserved Multicast Address Table Entry 2 Register](#).
- Write to the [Static Address and Reserved Multicast Table Control Register](#).
 - Write the TABLE_INDEX field with the 6-bit index value.
 - Set the TABLE_SELECT bit to 1 to select the Reserved Multicast Table.
 - Set the ACTION bit to 0 to indicate a write operation.
 - Set the START_FINISH bit to 1 to initiate the operation.
- When the operation is complete, the START_FINISH bit will be cleared automatically.

5.3.3.2 Reserved Multicast Table Read Operation

- Write to the [Static Address and Reserved Multicast Table Control Register](#).
 - Write the TABLE_INDEX field with the 6-bit index value.
 - Set the TABLE_SELECT bit to 1 to select the Reserved Multicast Table.
 - Set the ACTION bit to 1 to indicate a read operation.
 - Set the START_FINISH bit to 1 to initiate the operation.
- When the operation is complete, the START_FINISH bit will be cleared automatically.
 - Read the contents of the indexed entry from the [Reserved Multicast Address Table Entry 2 Register](#).

KSZ9896C

5.3.3.3 Reserved Multicast Address Table Entry 2 Register

Address: 0x0424 - 0x0427 Size: 32 bits

| Bits | Description | Type | Default |
|-------|--|------|----------|
| 31:30 | RESERVED | R/W | 00b |
| 29:6 | RESERVED | RO | 0x000000 |
| 5:0 | PORT FORWARD Each bit corresponds to a device port. Bit 0 is for port 1 Bit 1 is for port 2, etc. 1 = Forward to the corresponding port 0 = Do not forward to the corresponding port | R/W | 0x00 |

5.3.4 VLAN TABLE

An internal VLAN Table is used for VLAN lookup. If 802.1Q VLAN mode is enabled ([Switch Lookup Engine Control 0 Register](#)), this table will be used to retrieve the VLAN information that is associated with the ingress packet. The table holds 4096 entries - one for each possible VLAN. The table must be set up before 802.1Q VLAN is enabled. The VLAN table is accessed one entry at a time using the following indirect registers:

- [VLAN Table Entry 0 Register](#)
- [VLAN Table Entry 1 Register](#)
- [VLAN Table Entry 2 Register](#)
- [VLAN Table Index Register](#)
- [VLAN Table Access Control Register](#)

The table data fields are described in [Figure 5-4](#) and [Table 5-4](#).

FIGURE 5-4: VLAN TABLE STRUCTURE




TABLE 5-4: VLAN TABLE DATA FIELDS

| Field | Size (bits) | Description |
|----------------|-------------------|--|
| VALID | 1 | When 1, indicates that the table entry is valid. |
| FORWARD OPTION | 1 | When 1, forward to VLAN port table (PORT FORWARD field). When 0, see Table 4-8, "VLAN Forwarding" for details. |
| PRIORITY | 3 | Priority value for this VID. |
| MSTP INDEX | 3 | Multiple Spanning Tree Protocol index. |
| FID | 7 | Filter ID value. It is combined with destination address and hashed to index the Address Lookup Table. |
| PORT UNTAG | 7 (1 per port) | When 1, untag at that egress port. When 0, don't untag. |
| PORT FORWARD | 7 (1 per port) | VLAN port membership list. There is one bit per port, starting with the LSB which corresponds to port 1. A bit value of 1 indicates the associated port is included in the port membership list for that VID. When 0, that port is excluded. |

5.3.4.1 VLAN Table Write Operation

1. Write to the [VLAN Table Entry 0 Register](#), [VLAN Table Entry 1 Register](#), and [VLAN Table Entry 2 Register](#) to set up the data fields as described in [Figure 5-4](#) and [Table 5-4](#).
2. Write the VLAN Index value in the [VLAN Table Index Register](#). This is the 12-bit index (address) to select the table entry. It is equivalent to the VID which indexes the table during lookup.
3. Write the [VLAN Table Access Control Register](#) to specify a write operation, and set START (bit 7). When the operation is complete, bit 7 will be cleared automatically.

5.3.4.2 VLAN Table Read Operation

1. Write the VLAN Index value in the [VLAN Table Index Register](#) to select one of the 4k table entries.
2. Write the [VLAN Table Access Control Register](#) to specify a read operation and set START (bit 7). When the operation is complete, bit 7 will be cleared automatically.
3. Read the [VLAN Table Entry 0 Register](#), [VLAN Table Entry 1 Register](#), and [VLAN Table Entry 2 Register](#) to retrieve the read results from the VLAN table.

5.3.5 ACCESS CONTROL LIST (ACL) TABLE

ACL filtering is implemented individually per-port. The ACL tables are accessed using the [Port N: Port Switch ACL Control Registers \(0xN600 - 0xN6FF\)](#). The 16 entries in each ACL table are addressed indirectly by an index register.

[Table 5-5](#) shows how the various fields of the ACL Table entries are mapped to data registers. The [Port ACL Byte Enable MSB Register](#) and [Port ACL Byte Enable LSB Register](#) make it possible to write or read any combination of bytes. This is useful for writing the Matching rule, Action rule and Process field separately. There are 16 bits in these byte enable registers, corresponding to the 16 data registers [Port ACL Access 0 Register](#) through [Port ACL Access F Register](#). Note that the enable bits are applied in reverse order:

Bit 0 for the [Port ACL Access F Register](#)

Bit 1 for the [Port ACL Access E Register](#)

...

Bit 14 for the [Port ACL Access 1 Register](#)

Bit 15 for the [Port ACL Access 0 Register](#)

Also note that the [Port ACL Access C Register](#) is not used, so byte enable bit 3 is a don't care.

KSZ9896C

TABLE 5-5: ACL FIELD REGISTER MAPPING

| Register | Bits | MD = 01 ENB = 00 Count Mode | MD = 01 ENB ≠ 00 | MD = 10 | MD = 11 | | | |
|----------|------|-----------------------------------|-----------------------------|-----------------|-------------|--|------------|--|
| 0xN600 | 7:4 | RESERVED | | | | | | |
| | 3:0 | Process Field: FRN [3:0] | | | | | | |
| 0xN601 | 7:6 | RESERVED | | | | | | |
| | 5:4 | MD [1:0] | | | | | | |
| | 3:2 | ENB [1:0] | | | | | | |
| | 1 | S / D | | | | | | |
| | 0 | EQ | | | | | | |
| 0xN602 | 7:0 | MAC ADDRESS [47:0] | IP Address [31:0] | MAX PORT [15:0] | | | | |
| 0xN603 | 7:0 | | | MIN PORT [15:0] | | | | |
| 0xN604 | 7:0 | | | | | | | |
| 0xN605 | 7:0 | | | | | | | |
| 0xN606 | 7:3 | | | IP MASK [31:0] | RESERVED | | | |
| | 2:1 | | | | PC [1:0] | | | |
| | 0 | | | | PRO [7:0] | | | |
| 0xN607 | 7:1 | | | | TYPE [15:0] | | FME | |
| | 0 | | | | | | | |
| 0xN608 | 7:0 | | | | | | FMSK [7:0] | |
| 0xN609 | 7:0 | FLAG [7:0] | | | | | | |
| 0xN60A | 7:6 | COUNT [10:3] | Action Rule: PM [1:0] | | | | | |
| | 5:3 | | Action Rule: P [2:0] | | | | | |
| | 2 | | Action Rule: RPE | | | | | |
| | 1:0 | | Action Rule: RP [2:1] | | | | | |
| 0xN60B | 7 | COUNT [2:0] | Action Field: RP [0] | | | | | |
| | 6:5 | | Action Field: MM [1:0] | | | | | |
| | 4:0 | | RESERVED | | | | | |
| 0xN60C | 7:0 | RESERVED | | | | | | |
| 0xN60D | 7 | RESERVED | | | | | | |
| | 6 | TU | Action Field: FORWARD [6:0] | | | | | |
| | 5 | CA | | | | | | |
| | 4:0 | RESERVED | | | | | | |
| 0xN60E | 7:0 | Process Field: RuleSet [15:8] | | | | | | |
| 0xN60F | 7:0 | Process Field: RuleSet [7:0] | | | | | | |

5.3.5.1 ACL Table Read

1. Write to the [Port ACL Access Control 0 Register](#) with the table entry number (0 to 15) in the ACL Index field, and the Write/Read bit 4 cleared to zero. This one write to this register initiates the read operation.
2. Poll the Read Status bit in the [Port ACL Access Control 0 Register](#) to determine when the read operation is complete.
3. When the operation is complete, data may be retrieved from the [Port ACL Access 0 Register](#) through [Port ACL Access F Register](#).

5.3.5.2 ACL Table Write

1. Write the ACL table entry values to the [Port ACL Access 0 Register](#) through [Port ACL Access F Register](#).
2. Write the [Port ACL Byte Enable MSB Register](#) and [Port ACL Byte Enable LSB Register](#) to select which registers ([Port ACL Access 0 Register](#) through [Port ACL Access F Register](#)) are to be written into the ACL table.
3. Write to the [Port ACL Access Control 0 Register](#) with the table entry number in the ACL Index field, and the Write/Read bit 4 set to one. This one write to this register initiates the write operation.
4. The Write Status bit in the [Port ACL Access Control 0 Register](#) may be polled to determine when the operation is complete.

5.3.6 MANAGEMENT INFORMATION BASE (MIB) COUNTERS

There are 36 MIB counters per port. These counters accumulate a variety of statistics on ingress and egress traffic and events for network management. They are accessed indirectly using the [Port MIB Control and Status Register](#) and [Port MIB Data Register](#). The [Switch MIB Control Register](#) provides global flush and freeze control of the MIB counters.

TABLE 5-6: MIB COUNTERS

| MIB Index | MIB Counter | Size (bits) | Description |
|-----------|--------------------|-------------|--|
| 0x00 | RxHiPriorityByte | 30 | RX high priority octet count, including bad packets. |
| 0x01 | RxUndersizePkt | 30 | RX undersize packets with good CRC. |
| 0x02 | RxFragments | 30 | RX fragment packets with bad CRC, symbol errors or alignment errors. |
| 0x03 | RxOversize | 30 | RX oversize packets w/ good CRC (max: 1536 or 1522 bytes). |
| 0x04 | RxJabbers | 30 | RX packets longer than 1522 bytes with either CRC errors, alignment errors or symbol errors (depends on max packet size setting); or RX packets longer than 1916 bytes only. |
| 0x05 | RxSymbolError | 30 | RX packets with invalid data symbol; and legal preamble and packet size. |
| 0x06 | RxCRCError | 30 | RX packets between 64 and 1522 bytes in size, with an integral number of bytes and a bad CRC. (Upper limit depends on max packet size setting.) |
| 0x07 | RxAlignmentError | 30 | RX packets between 64 and 1522 bytes in size, with a non-integral number of bytes and a bad CRC. (Upper limit depends on max packet size setting.) |
| 0x08 | RxControl8808Pkts | 30 | MAC control frames received with 0x8808 in the EtherType field. |
| 0x09 | RxPausePkts | 30 | PAUSE frames received. PAUSE is defined as EtherType (0x8808), DA, control opcode (0x0001), minimum 64 byte data length, and a valid CRC. |
| 0x0A | RxBroadcast | 30 | RX good broadcast packets. Does not include erred broadcast packets or valid multicast packets. |
| 0x0B | RXMulticast | 30 | RX good multicast packets. Does not include MAC control frames, erred multicast packets, or valid broadcast packets. |
| 0x0C | RxUnicast | 30 | RX good unicast packets. |
| 0x0D | Rx64Octets | 30 | RX packets (bad packets included) that are 64 bytes in length. |
| 0x0E | Rx65to127Octets | 30 | RX packets (bad packets included) that are 65 to 127 bytes in length. |
| 0x0F | Rx128to255Octets | 30 | RX packets (bad packets included) that are 128 to 255 bytes in length. |
| 0x10 | Rx256to511Octets | 30 | RX packets (bad packets included) that are 256 to 511 bytes in length. |
| 0x11 | Rx512to1023Octets | 30 | RX packets (bad packets included) that are 512 to 1023 bytes in length. |
| 0x12 | Rx1024to1522Octets | 30 | RX packets (bad packets included) that are 1024 to 1522 bytes in length. |
| 0x13 | Rx1523to2000Octets | 30 | RX packets (bad packets included) that are 1523 to 2000 bytes in length. |
| 0x14 | Rx2001+Octets | 30 | RX packets (bad packets included) that are between 2001 bytes and the upper limit in length. |

TABLE 5-6: MIB COUNTERS (CONTINUED)

| MIB Index | MIB Counter | Size (bits) | Description |
|-----------|----------------------|-------------|---|
| 0x15 | TxHiPriorityByte | 30 | TX high priority good octet count, including PAUSE packets. |
| 0x16 | TxLateCollision | 30 | Collision is detected later than 512 bit times into the transmission of a packet. |
| 0x17 | TxPausePkts | 30 | PAUSE frames transmitted. PAUSE is EtherType (0x8808), DA, control opcode (0x0001), minimum 64 byte data length, and a valid CRC. |
| 0x18 | TxBroadcastPkts | 30 | TX good broadcast packets. Does not include erred broadcast packets or valid multicast packets. |
| 0x19 | TxMulticastPkts | 30 | TX good multicast packets. Does not include MAC control frames, erred multicast packets, or valid broadcast packets. |
| 0x1A | TxUnicastPkts | 30 | TX good unicast packets. |
| 0x1B | TxDeferred | 30 | TX packets where the first transmit attempt is delayed due to the busy medium. |
| 0x1C | TxTotalCollision | 30 | TX total collisions. Half duplex only. |
| 0x1D | TxExcessiveCollision | 30 | TX fails due to excessive collisions. |
| 0x1E | TxSingleCollision | 30 | Successfully transmitted frames where transmission is inhibited by exactly one collision. |
| 0x1F | TxMultipleCollision | 30 | Successfully transmitted frames where transmission is inhibited by more than one collision. |
| 0x80 | RxByteCnt | 36 | RX byte count. |
| 0x81 | TxByteCnt | 36 | TX byte count. |
| 0x82 | RxDropPackets | 30 | RX packets dropped due to lack of resources. |
| 0x83 | TXDropPackets | 30 | TX packets dropped due to lack of resources. |

5.3.6.1 MIB Counter Read Operation

Indirect access registers are used to read the MIB counters. Separate access registers are provided for each port via the [Port MIB Control and Status Register](#) and [Port MIB Data Register](#). All MIB Counters are read-clear. The steps for reading a counter are as follows:

1. Write the MIB Index to bits [23:16] of the [Port MIB Control and Status Register](#).
2. Set the MIB Read Enable in bit 25 of the [Port MIB Control and Status Register](#). This step and the previous step may be done together.
3. Read the MIB Read Enable / Count Valid in bit 25 of the [Port MIB Control and Status Register](#). A '0' value indicates that the read is complete and the count is valid.
4. Read the count value from the [Port MIB Data Register](#). For 36-bit counters, counter bits [35:32] are read from the [Port MIB Control and Status Register](#). The Counter Overflow bit is also found in the [Port MIB Control and Status Register](#).

5.3.6.2 MIB Counter Freeze and Flush Functions

Counter freeze and flush functions are available on a port-by-port basis. Freezing or flushing counters is initiated by setting the appropriate bit in the [Switch MIB Control Register](#). The freeze or flush function will be applied to all ports for which the flush and freeze functions have been enabled. To enable flush and freeze for a port, set bit 24 in the [Port MIB Control and Status Register](#).

The following steps show an example of how flush and freeze are used to collect MIB statistics for all ports for a period of 1 second:

1. Set the MIB Flush and Freeze Enable bit 24 in the [Port MIB Control and Status Register](#) for all ports N.
2. Write 0x40 to the [Switch MIB Control Register](#) to freeze the MIB counters for all enabled ports.
3. Write 0xC0 to the [Switch MIB Control Register](#) to clear the MIB counters for all enabled ports (while continuing to also freeze the counters).
4. At the beginning of the 1 second period, write 0x00 to the [Switch MIB Control Register](#) to enable the counters.
5. At the end of the 1 second period, write 0x40 to the [Switch MIB Control Register](#) to freeze the counters.
6. Read each counter for each port.

5.4 MDIO Manageable Device (MMD) Registers (Indirect)

MMD registers provide indirect read/write access to up to 32 MMD device addresses with each device supporting up to 65,536 16-bit registers, as defined in Clause 22 of the IEEE 802.3 Specification. However, the KSZ9896C uses only a small fraction of the available registers. Refer to [Table 5-7, "MMD Register Map"](#) for a list of accessible MMD device addresses and their associated register addresses. Detailed descriptions of the supported MMD registers are provided in the following subsections. Additional information on the MIIM interface is provided in [Section 4.9.3, "MIIM Management \(MIIM\) Interface,"](#) on page 55.

The following two standard port registers serve as the portal registers to access the indirect MMD registers.

- [PHY MMD Setup Register](#)
- [PHY MMD Data Register](#)

TABLE 5-7: MMD REGISTER MAP

| Device Address (hex) | Register Address (hex) | Description |
|----------------------|------------------------|--|
| 2h | 00h | MMD LED Mode Register |
| 7h | 3Ch | MMD EEE Advertisement Register |

Example: MMD Register Write

Write MMD - Device Address 2h, Register 00h = 0010h to enable single-LED mode.

1. Write the [PHY MMD Setup Register](#) with 0002h // Set up register address for MMD – Device Address 2h.
2. Write the [PHY MMD Data Register](#) with 0000h // Select Register 00h of MMD – Device Address 2h.
3. Write the [PHY MMD Setup Register](#) with 4002h // Select register data for MMD – Device Address 2h, Reg. 00h.
4. Write the [PHY MMD Data Register](#) with 0010h // Write value 0010h to MMD – Device Address 2h, Reg. 00h.

Example: MMD Register Read

Read MMD - Device Address 2h, Register 11h – 13h for the LED mode status.

1. Write the [PHY MMD Setup Register](#) with 0002h // Set up register address for MMD – Device Address 2h.
2. Write the [PHY MMD Data Register](#) with 0000h // Select Register 00h of MMD – Device Address 2h.
3. Write the [PHY MMD Setup Register](#) with 4002h // Select register data for MMD – Device Address 2h, Reg. 00h.
4. Read the [PHY MMD Data Register](#) // Read data in MMD – Device Address 2h, Reg. 00h.

KSZ9896C

5.4.1 MMD LED MODE REGISTER

MMD Address: 0x02 Size: 16 bits
Register: 0x00

| Bits | Description | Type | Default |
|------|---|------|---------|
| 15:5 | RESERVED | RO | 0x000 |
| 4 | LED Mode 1 = Single-LED Mode 0 = Tri-color Dual-LED Mode | R/W | 0b |
| 3:0 | RESERVED | RO | 0001b |

5.4.2 MMD EEE ADVERTISEMENT REGISTER

MMD Address: 0x07 Size: 16 bits
Register: 0x3C

| Bits | Description | Type | Default |
|------|--|------|---------|
| 15:3 | RESERVED | RO | 0x000 |
| 2 | 1000BASE-T EEE Enable 1 = 1000 Mbps EEE capable 0 = No 1000 Mbps EEE capability | R/W | 1b |
| 1 | 100BASE-T EEE Enable 1 = 100 Mbps EEE capable 0 = No 100 Mbps EEE capability | R/W | 1b |
| 0 | RESERVED | RO | 0b |

6.0 OPERATIONAL CHARACTERISTICS

6.1 Absolute Maximum Ratings*

| | |
|--|------------------|
| Supply Voltage (AVDDL, DVDDL) | -0.5 V to +1.8 V |
| Supply Voltage (AVDDH, VDDIO) | -0.5 V to +5.0 V |
| Input Voltage (all inputs) | -0.5 V to +5.0 V |
| Output Voltage (all outputs) | -0.5 V to +5.0 V |
| Lead Temperature (soldering, 20 sec.) | +260°C |
| Storage Temperature (T _S) | -65°C to +150°C |
| Maximum Junction Temperature (T _J) | +125°C |
| HBM ESD Performance | +/-6 kV |

*Stresses exceeding those listed in this section could cause permanent damage to the device. This is a stress rating only. Exposure to absolute maximum rating conditions for extended periods may affect device reliability. Functional operation of the device at any condition exceeding those indicated in [Section 6.2, "Operating Conditions**"](#), [Section 6.3, "Electrical Characteristics"](#), or any other applicable section of this specification is not implied.

6.2 Operating Conditions**

| | |
|---|--------------------------|
| Supply Voltage (AVDDL, DVDDL) | +1.14V to +1.26 V |
| Supply Voltage (AVDDH) | +2.375 V to +2.625 V |
| Supply Voltage (VDDIO @ 3.3V) | +3.135 V to +3.465 V |
| Supply Voltage (VDDIO @ 2.5V) | +2.375 V to +2.625 V |
| Supply Voltage (VDDIO @ 1.8V) | +1.71 V to +1.89 V |
| Ambient Operating Temperature in Still Air (T _A) | Note 6-1 |
| Junction to Ambient Resistance (Θ _{JA}) (Note 6-2) | Note 6-3 |
| Junction to Case Characterization (Ψ _{JT}) (Note 6-2) | 0.04°C/W |
| Junction to Case Resistance (Θ _{JC}) (Note 6-2) | Note 6-4 |

Note 6-1 0°C to +70°C for commercial version, -40°C to +85°C for industrial version.

Note 6-2 Ψ_{JT} and Θ_{JA} are under a 0 m/s air velocity. A 6-layer PCB is required for industrial applications.

Note 6-3 11.3°C/W on a 6-layer PCB per JESD51, 14.4°C/W on a 4-layer PCB per JESD51.

Note 6-4 1.5°C/W on a 6-layer PCB per JESD51, 1.21°C/W on a 4-layer PCB per JESD51.

**Proper operation of the device is ensured only within the ranges specified in this section.

KSZ9896C

6.3 Electrical Characteristics

T_A = 25°C.

TABLE 6-1: ELECTRICAL CHARACTERISTICS

| Symbol | Parameter | Conditions | Min | Typ | Max | Units |
|--|-----------------------------|--|-------------|------|-------------|-------|
| Supply Current - Full 1000 Mbps Operation | | | | | | |
| I _{DD_AH} | AVDDH supply current | VDDIO @ 3.3V Ports 1-5 in 1000BASE-T Port 6 @ 1000 Mbps All ports 100% utilization | | 330 | | mA |
| I _{DD_IO} | VDDIO supply current | | | 80 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 460 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 750 | | mA |
| Supply Current - Mixed 1000/100 Mbps Operation | | | | | | |
| I _{DD_AH} | AVDDH supply current | VDDIO @ 3.3V Ports 1-5 in 100BASE-TX Port 6 @ 1000 Mbps All ports 100% utilization | | 140 | | mA |
| I _{DD_IO} | VDDIO supply current | | | 80 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 140 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 350 | | mA |
| Supply Current - Full 100 Mbps Operation | | | | | | |
| I _{DD_AH} | AVDDH supply current | VDDIO @ 3.3V Ports 1-5 in 100BASE-TX Port 6 @ 100 Mbps All ports 100% utilization | | 140 | | mA |
| I _{DD_IO} | VDDIO supply current | | | 35 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 140 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 350 | | mA |
| Supply Current - Full 10 Mbps Operation | | | | | | |
| I _{DD_AH} | AVDDH supply current | VDDIO @ 3.3V Ports 1-5 in 10BASE-T Port 6 @ 10Mbps All ports 100% utilization | | 100 | | mA |
| I _{DD_IO} | VDDIO supply current (3.3V) | | | 30 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 30 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 150 | | mA |
| Supply Current - Power Management - Energy Detect Mode | | | | | | |
| I _{DD_AH} | AVDDH supply current | | | 20 | | mA |
| I _{DD_IO} | VDDIO supply current (3.3V) | | | 30 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 30 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 150 | | mA |
| Supply Current - Power Management - Global Soft Power Down Mode | | | | | | |
| I _{DD_AH} | AVDDH supply current | | | 2 | | mA |
| I _{DD_IO} | VDDIO supply current (3.3V) | | | 6 | | mA |
| I _{DD_CA} | AVDDL supply current | | | 0.01 | | mA |
| I _{DD_CD} | DVDDL supply current | | | 5 | | mA |
| I Type CMOS Input Buffers (VDDIO = 3.3/2.5/1.8V) | | | | | | |
| V _{IH} | Input High Voltage | | 2.1/1.7/1.3 | | | V |
| V _{IL} | Input Low Voltage | | | | 0.9/0.9/0.6 | V |
| I _{IN} | Input Current | V _{IN} = GND ~ VDDIO | -10 | | 10 | μA |

TABLE 6-1: ELECTRICAL CHARACTERISTICS (CONTINUED)

| Symbol | Parameter | Conditions | Min | Typ | Max | Units |
|--|--|---|-------------|------|-------------|-------|
| O8 Type CMOS Output Buffers (VDDIO = 3.3/2.5/1.8V) | | | | | | |
| V _{OH} | Output High Voltage | I _{OH} = 8/8/6 mA | 2.4/1.9/1.5 | | | V |
| V _{OL} | Output Low Voltage | I _{OL} = 8/8/6 mA | | | 0.4/0.4/0.2 | V |
| I _{OZ} | Output Tri-State Leakage | V _{IN} = GND ~ VDDIO | | | 10 | μA |
| O24 Type CMOS Output Buffers (VDDIO = 3.3/2.5/1.8V) | | | | | | |
| V _{OH} | Output High Voltage | I _{OH} = 24/24/20 mA | 2.4/1.9/1.5 | | | V |
| V _{OL} | Output Low Voltage | I _{OL} = 24/24/20 mA | | | 0.4/0.4/0.2 | V |
| I _{OZ} | Output Tri-State Leakage | V _{IN} = GND ~ VDDIO | | | 10 | μA |
| I/O Pin Internal Pull-Up and Pull-Down Effective Resistance | | | | | | |
| R1.8PU | I/O Pin Effective Pull-Up Resistance | VDDIO = 1.8V | | 125 | | kΩ |
| R1.8PD | I/O Pin Effective Pull-Down Resistance | | | 97 | | kΩ |
| R2.5PU | I/O Pin Effective Pull-Up Resistance | VDDIO = 2.5V | | 58 | | kΩ |
| R2.5PD | I/O Pin Effective Pull-Down Resistance | | | 51 | | kΩ |
| R3.3PU | I/O Pin Effective Pull-Up Resistance | VDDIO = 3.3V | | 38 | | kΩ |
| R3.3PD | I/O Pin Effective Pull-Down Resistance | | | 39 | | kΩ |
| 100BASE-TX Transmit (Measured Differentially After 1:1 Transformer) | | | | | | |
| V _O | Peak Differential Output | 100Ω termination on the differential output | ±0.95 | | ±1.05 | V |
| V _{imb} | Output Voltage Imbalance | 100Ω termination on the differential output | | | 2 | % |
| t _r , t _f | Rise/Fall Time | | 3 | | 5 | ns |
| | Rise/Fall Time Imbalance | | 0 | | 0.5 | ns |
| | Duty cycle Distortion | | | | ±0.25 | ns |
| | Overshoot | | | | 5 | % |
| V _{SET} | Reference Voltage of ISET (using 6.04kΩ - 1% resistor) | | | 1.21 | | V |
| | Output Jitter | Peak-to-Peak | | 0.7 | 1.4 | ns |
| 10BASE-Te Receive | | | | | | |
| V _{sq} | Squelch Threshold | 5MHz Square Wave | | 400 | | mV |

KSZ9896C

TABLE 6-1: ELECTRICAL CHARACTERISTICS (CONTINUED)

| Symbol | Parameter | Conditions | Min | Typ | Max | Units |
|---|----------------------------------|--|------|------|------|-------|
| 10BASE-Te Transmit (Measured Differentially After 1:1 Transformer) | | | | | | |
| V_p | Peak Differential Output Voltage | 100 Ω termination on the differential output | 1.54 | 1.75 | 1.96 | V |
| | Jitter Added | 100 Ω termination on the differential output (peak-to-peak) | | | 3.5 | ns |
| t_r, t_f | Rise/Fall Time | | | 25 | | ns |

6.4 Timing Specifications

This section details the various timing specifications of the device.

Note: The I²C interface timing adheres to the NXP I²C-Bus Specification (UM10204, Rev. 6) (high-speed mode and slower). Refer to the I²C-Bus Specification for additional information.

6.4.1 GMII TIMING

Figure 6-1 illustrates the GMII timing requirements.

FIGURE 6-1: GMII TIMING




TABLE 6-2: GMII TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|------------------|--|-----|-----|-----|-------|
| t _{CYC} | Clock Period | 7.5 | 8.0 | 8.5 | ns |
| t _{SU} | TX_DV, TX_ER, TXD[7:0] setup time to rising edge of GTX_CLK | 2.0 | | | ns |
| t _{HD} | TX_DV, TX_ER, TXD[7:0] hold time from rising edge of GTX_CLK | 1.2 | | | ns |
| t _{OD} | Output skew between GTX_CLK and RX_DV, RX_ER, RXD[7:0] | | | 1 | ns |

KSZ9896C

6.4.2 RGMII TIMING

Figure 6-2 illustrates the RGMII timing requirements.

FIGURE 6-2: RGMII TIMING




TABLE 6-3: RGMII TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|---------------------------------|--------------------------------------|-----|-----|----------|-------|
| T _{SETUP} | Data to clock input setup (Note 6-5) | 1.3 | | | ns |
| T _{HOLD} | Data to clock input hold (Note 6-5) | 0.7 | | | ns |
| T _{SKEW} | Data to clock output skew (Note 6-6) | 1.2 | 2.0 | | ns |
| T _{CYC} | Clock cycle duration (Note 6-7) | 7.2 | 8 | 8.8 | ns |
| Duty_G | 1000Mbps duty cycle | 45 | 50 | 55 | % |
| Duty_T | 10/100Mbps duty cycle | 40 | 50 | 60 | % |
| T _r / T _f | Rise / Fall time (20-80%) | | | Note 6-8 | ns |

Note 6-5 For cases where there is no (or insufficient) skew between the input data and input clock, it is possible to add internal delay to the TX_CLK pinout by setting the RGMII Ingress Internal Delay bit in the [XMII Port Control 1 Register](#) register. This feature reduces the setup time requirement and increases the hold time requirement nominally by 1.3ns.

Note 6-6 The RGMII interface adheres to the RGMII Specification Version 2.0, which specified that the driving device delay the output clock relative to the output data. This is the T_{SKEW} parameter. This skew can be disabled by clearing the RGMII Egress Internal Delay bit in the [XMII Port Control 1 Register](#) register. Generally this is not recommended.

Note 6-7 For 10Mbps and 100Mbps, T_{CYC} will scale to 400ns +/- 40ns and 40ns +/- 4 ns, respectively.

Note 6-8 0.75ns for VDDIO = 3.3V/2.5V, 1.0ns for VDDIO = 1.8V

6.4.3 MII TIMING

6.4.3.1 MII Transmit Timing in MAC Mode

Figure 6-3 illustrates a write operation from the KSZ9896C to a PHY or other device while operating the KSZ9896C in MAC Mode.

FIGURE 6-3: MII TRANSMIT TIMING IN MAC MODE




TABLE 6-4: MII TRANSMIT TIMING IN MAC MODE VALUES

| Symbol | Description | Min | Typ | Max | Units |
|--|--|-----|--------|-----|-------|
| t_p (100BASE-TX/ 10BASE-Te) | RX_CLK period | | 40/400 | | ns |
| t_{WL} (100BASE-TX/ 10BASE-Te) | RX_CLK pulse width low | | 20/200 | | ns |
| t_{WH} (100BASE-TX/ 10BASE-Te) | RX_CLK pulse width high | | 20/200 | | ns |
| t_{OD} | RX_DV, RXD_[3:0] output delay from rising edge of RX_CLK | | 16 | | ns |

KSZ9896C

6.4.3.2 MII Receive Timing in MAC Mode

Figure 6-4 illustrates a read operation by the KSZ9896C from a PHY or other device while operating the KSZ9896C in MAC Mode.

FIGURE 6-4: MII RECEIVE TIMING IN MAC MODE




TABLE 6-5: MII RECEIVE TIMING IN MAC MODE VALUES

| Symbol | Description | Min | Typ | Max | Units |
|--|---|-----|--------|-----|-------|
| t_p (100BASE-TX/ 10BASE-Te) | TX_CLK period | | 40/400 | | ns |
| t_{WL} (100BASE-TX/ 10BASE-Te) | TX_CLK pulse width low | | 20/200 | | ns |
| t_{WH} (100BASE-TX/ 10BASE-Te) | TX_CLK pulse width high | | 20/200 | | ns |
| t_{SU1} | TXD_[3:0] setup time to rising edge of TX_CLK | 10 | | | ns |
| t_{SU2} | TX_EN, TX_ER setup time to rising edge of TX_CLK | 10 | | | ns |
| t_{HD1} | TXD_[3:0] hold time from rising edge of TX_CLK | 10 | | | ns |
| t_{HD2} | TX_EN, TX_ER hold time from rising edge of TX_CLK | 10 | | | ns |

6.4.3.3 MII Receive Timing in PHY Mode

FIGURE 6-5: MII RECEIVE TIMING IN PHY MODE




TABLE 6-6: MII RECEIVE TIMING IN PHY MODE VALUES

| Symbol | Description | Min | Typ | Max | Units |
|--|--|-----|--------|-----|-------|
| t_p (100BASE-TX/ 10BASE-Te) | RX_CLK period | | 40/400 | | ns |
| t_{WL} (100BASE-TX/ 10BASE-Te) | RX_CLK pulse width low | | 20/200 | | ns |
| t_{WH} (100BASE-TX/ 10BASE-Te) | RX_CLK pulse width high | | 20/200 | | ns |
| t_{OD} | RX_DV, RXD_[3:0] output delay from rising edge of RX_CLK | | 20 | | ns |

KSZ9896C

6.4.3.4 MII Transmit Timing in PHY Mode

FIGURE 6-6: MII TRANSMIT TIMING IN PHY MODE




TABLE 6-7: MII TRANSMIT TIMING IN PHY MODE VALUES

| Symbol | Description | Min | Typ | Max | Units |
|--|---|-----|--------|-----|-------|
| t_p (100BASE-TX/ 10BASE-Te) | TX_CLK period | | 40/400 | | ns |
| t_{WL} (100BASE-TX/ 10BASE-Te) | TX_CLK pulse width low | | 20/200 | | ns |
| t_{WH} (100BASE-TX/ 10BASE-Te) | TX_CLK pulse width high | | 20/200 | | ns |
| t_{SU1} | TXD_[3:0] setup time to rising edge of TX_CLK | 10 | | | ns |
| t_{SU2} | TX_EN, TX_ER setup time to rising edge of TX_CLK | 10 | | | ns |
| t_{HD1} | TXD_[3:0] hold time from rising edge of TX_CLK | 0 | | | ns |
| t_{HD2} | TX_EN, TX_ER hold time from rising edge of TX_CLK | 0 | | | ns |

6.4.4 RMII TIMING

Figure 6-7 and Figure 6-8 illustrate the RMII timing requirements.

FIGURE 6-7: RMII TRANSMIT TIMING




FIGURE 6-8: RMII RECEIVE TIMING




TABLE 6-8: RMII TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|-----------|--------------|-----|-----|-----|-------|
| t_{cyc} | Clock cycle | | 20 | | ns |
| t_1 | Setup time | 4 | | | ns |
| t_2 | Hold time | 2 | | | ns |
| t_{od} | Output delay | 7 | 9 | 13 | ns |

KSZ9896C

6.4.5 MIIM TIMING

Figure 6-9 illustrates the MIIM timing requirements.

FIGURE 6-9: MIIM TIMING




TABLE 6-9: MIIM TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|----------|--|-----|-----|-----|-------|
| t_p | MDC period | | 400 | | ns |
| t_{OD} | Output delay | | 200 | | ns |
| t_{SU} | MDIO setup time to rising edge of MDC | 10 | | | ns |
| t_{HD} | MDIO hold time from rising edge of MDC | 5 | | | ns |

6.4.6 SPI TIMING

Figure 6-10 and Figure 6-11 illustrate the SPI timing requirements.

FIGURE 6-10: SPI DATA INPUT TIMING




FIGURE 6-11: SPI DATA OUTPUT TIMING




TABLE 6-10: SPI TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|------------|---|-----|-----|-----|-------|
| f_{SCLK} | SCL clock frequency | | | 50 | MHz |
| t_1 | SCS_N active setup time | 8 | | | ns |
| t_2 | SDI data input setup time | 3 | | | ns |
| t_3 | SDI data input hold time | 3 | | | ns |
| t_4 | SCS_N active hold time | 8 | | | ns |
| t_5 | SCS_N disable high time | 8 | | | ns |
| t_6 | SCL falling edge to SDO data output valid | 2 | | 9 | ns |
| t_7 | SCS_N inactive to SDO data input invalid | 1 | | | ns |

KSZ9896C

6.4.7 AUTO-NEGOTIATION TIMING

Figure 6-12 illustrates the Auto-Negotiation timing requirements.

FIGURE 6-12: AUTO-NEGOTIATION TIMING




TABLE 6-11: AUTO-NEGOTIATION TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|------------|---------------------------------------|------|-----|------|---------|
| t_{BTB} | FLP burst to FLP burst | 8 | 16 | 24 | ms |
| t_{FLPW} | FLP burst width | | 2 | | ms |
| t_{PW} | Clock/Data pulse width | | 100 | | ns |
| t_{CTD} | Clock pulse to data pulse | 55.5 | 64 | 69.5 | μ s |
| t_{CTC} | Clock pulse to clock pulse | 111 | 128 | 139 | μ s |
| | Number of clock/data pulses per burst | 17 | | 33 | |

6.4.8 POWER-UP AND RESET TIMING

Figure 6-13 illustrates the power-up and reset timing requirements.

FIGURE 6-13: POWER-UP AND RESET TIMING




TABLE 6-12: POWER-UP AND RESET TIMING VALUES

| Symbol | Description | Min | Typ | Max | Units |
|--------|---|-----|-----|-----|---------|
| tvr | Supply voltage rise time (must be monotonic) | 200 | | | μ s |
| tsr | Stable supply voltages to de-assertion of reset | 10 | | | ms |
| tcs | Configuration strap input setup time | 5 | | | ns |
| tch | Configuration strap input hold time | 5 | | | ns |
| trc | De-assertion of reset to configuration strap pin output | 6 | | | ns |
| tpc | Supply voltages cycle off-to-on time | 150 | | | ms |

Note 1: The recommended powering sequence is to bring up all voltages at the same time. However, if that cannot be attained, then the recommended power-up sequence is to power-up the transceiver (AVDDH) and digital I/Os (VDDIO) voltages before the low voltage core (AVDDL and DVDDL). There is no power sequence requirement between transceiver (AVDDH) and digital I/Os (VDDIO) power rails. The power-up waveforms should be monotonic for all supply voltages.

Note 2: After the de-assertion of reset, it is recommended to wait a minimum of 100 μ s before starting to program the device through any interface.

Note 3: The recommended power down sequence is to power down the low voltage core before powering down the transceiver and digital I/O voltages, or to have all supplies power down in unison. Before the next power-up cycle, all supply voltages to the device should reach less than 0.4V and there should be a minimum wait time of 150ms from power-off to power-on.

KSZ9896C

6.5 Clock Specifications

A crystal or external clock source, such as an oscillator, is used to provide a 25MHz reference clock for the KSZ9896C. If an external clock source is used, the **XO** pin must be left floating. Since the **XI/XO** circuit is powered from **AVDDH**, the external clock source should also be powered from the same power rail. [Figure 6-14](#) details the available connection methods. [Table 6-13](#) details the recommended crystal specifications.

FIGURE 6-14: INPUT REFERENCE CLOCK CONNECTION OPTIONS




TABLE 6-13: REFERENCE CRYSTAL CHARACTERISTICS

| Characteristic | Min | Typ | Max | Units |
|------------------------------------|-------------|-----|-----|-------|
| Oscillation Mode | Fundamental | | | |
| Frequency | | 25 | | MHz |
| Frequency tolerance | | | ±50 | ppm |
| Effective Series Resistance (ESR) | | | 50 | Ω |
| Total period jitter (peak-to-peak) | | | 100 | ps |
| Drive level | | | 100 | uW |

7.0 DESIGN GUIDELINES

This section provides general design guidelines for the following:

- [Reset Circuit Guidelines](#)
- [Magnetics Connection and Selection Guidelines](#)

7.1 Reset Circuit Guidelines

Figure 7-1 illustrates the recommended reset circuit for powering up the KSZ9896C if reset is triggered by the power supply.

FIGURE 7-1: SIMPLE RESET CIRCUIT





Figure 7-2 illustrates a reset circuit recommended for applications where reset is driven by another device, such as a CPU. At power-on reset, R, C and D1 provide the necessary ramp rise time to reset the KSZ9896C. The RST_OUT_N from the CPU provides a warm reset after power-up.

FIGURE 7-2: RESET CIRCUIT FOR CPU RESET INTERFACE



KSZ9896C

7.2 Magnetics Connection and Selection Guidelines

A 1:1 isolation transformer is required at the line interface. For designs exceeding FCC requirements, utilize one with integrated common-mode chokes. An optional auto-transformer stage following the chokes provides additional common-mode noise and signal attenuation.

The KSZ9896C PHY port design incorporates voltage-mode transmit drivers and on-chip terminations. With the voltage-mode implementation, the transmit drivers supply the common-mode voltages to the four differential pairs. Therefore, the four transformer center tap pins on the KSZ9896C chip side should not be connected to any power supply source on the board; rather, the center tap pins should be separated from one another and connected through separate 0.1µF common-mode capacitors to ground. Separation is required because the common-mode voltage could be different between the differential pairs, depending on the connected speed mode.

Figure 7-3 details a typical magnetic interface circuit for the KSZ9896C PHY port.

FIGURE 7-3: TYPICAL MAGNETIC INTERFACE CIRCUIT




Table 7-1 provides a list of recommended magnetic characteristics.

TABLE 7-1: MAGNETICS SELECTION CRITERIA

| Parameter | Value | Test Condition |
|--------------------------------|-------------|--------------------|
| Turns ratio | 1 CT : 1 CT | |
| Open-circuit inductance (min.) | 350µH | 100mV, 100KHz, 8mA |
| Insertion loss (typ.) | 1.0dB | 100KHz to 100MHz |
| HIPOT (min.) | 1500vrms | |

Table 7-2 provides a list of KSZ9896C compatible single-port magnetics with separated transformer center tap pins on the Gigabit PHY chip side.


TABLE 7-2: COMPATIBLE SINGLE-PORT 10/100/1000 MAGNETICS

| Manufacturer | Part Number | Auto-Transformer | Temperature Range | Magnetic + RJ-45 |
|--------------|------------------|------------------|-------------------|------------------|
| Bel Fuse | 0826-1G1T-23-F | Yes | 0°C to 70°C | Yes |
| HALO | TG1G-E001NZRL | No | -40°C to 85°C | No |
| HALO | TG1G-S001NZRL | No | 0°C to 70°C | No |
| HALO | TG1G-S002NZRL | Yes | 0°C to 70°C | No |
| Pulse | H5007NL | Yes | 0°C to 70°C | No |
| Pulse | H5062NL | Yes | 0°C to 70°C | No |
| Pulse | HX5008NL | Yes | -40°C to 85°C | No |
| Pulse | JK0654219NL | Yes | 0°C to 70°C | Yes |
| Pulse | JK0-0136NL | No | 0°C to 70°C | Yes |
| TDK | TLA-7T101LF | No | 0°C to 70°C | No |
| Würth/Midcom | 000-7093-37R-LF1 | Yes | 0°C to 70°C | No |

KSZ9896C

8.0 PACKAGE INFORMATION

8.1 Package Marking Information



| | | |
|----------------|-----|---|
| Legend: | t | Temperature range designator (C = commercial, I = industrial) |
| | R | Product revision |
| | nnn | Internal code |
| | e3 | Pb-free JEDEC® designator for Matte Tin (Sn) |
| | VVV | Plant assembly |
| | CO | Country of origin |
| | YY | Year code (last two digits of calendar year) |
| | WW | Week code (week of January 1 is week '01') |
| | NNN | Alphanumeric traceability code |


Note: In the event the full Microchip part number cannot be marked on one line, it will be carried over to the next line, thus limiting the number of available characters for customer-specific information.

* Standard device marking consists of Microchip part number, year code, week code and traceability code. For device marking beyond this, certain price adders apply. Please check with your Microchip Sales Office. For QTP devices, any special marking adders are included in QTP price.

8.2 Package Drawings

Note: For the most current package drawings, please see the Microchip Packaging Specification located at <http://www.microchip.com/packaging>

FIGURE 8-1: PACKAGE (DRAWING & DIMENSIONS)



KSZ9896C

APPENDIX A: DATA SHEET REVISION HISTORY

TABLE A-1: REVISION HISTORY

| Revision | Section/Figure/Entry | Correction |
|------------------------|--------------------------|------------|
| DS00002390A (03-06-17) | Initial Document Release | |

THE MICROCHIP WEB SITE

Microchip provides online support via our WWW site at www.microchip.com. This web site is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the web site contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, Microchip consultant program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

CUSTOMER CHANGE NOTIFICATION SERVICE

Microchip's customer notification service helps keep customers current on Microchip products. Subscribers will receive e-mail notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, access the Microchip web site at www.microchip.com. Under "Support", click on "Customer Change Notification" and follow the registration instructions.

CUSTOMER SUPPORT

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or field application engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the web site at: <http://www.microchip.com/support>

KSZ9896C

PRODUCT IDENTIFICATION SYSTEM

To order or obtain information, e.g., on pricing or delivery, refer to the factory or the listed sales office.

| <u>PART NO.</u> | <u>XX</u> | <u>X</u> | <u>[XX]</u> ⁽¹⁾ |
|------------------------------|---|--|----------------------------|
| Device | Package | Temp. Range | Tape & Reel Option |
| Device: | KSZ9896C = 7-Port Switch with 1 GMII/RGMII/MII/RMII | | |
| Package: | TX | = 128-pin TQFP-EP | |
| Temperature Range: | C | = 0°C to +70°C (Commercial) | |
| | I | = -40°C to +85°C (Industrial) | |
| Tape and Reel Option: | Blank | = Standard packaging (tray) | |
| | -TR | = Tape and Reel (Note 1) | |

Examples:

a) KSZ9896CTXC
128-pin TQFP-EP package,
Commercial temperature,
Standard packaging

b) KSZ9896CTXI-TR
128-pin TQFP-EP package,
Industrial temperature,
Tape and reel

Note 1: Tape and Reel identifier only appears in the catalog part number description. This identifier is used for ordering purposes and is not printed on the device package. Check with your Microchip Sales Office for package availability with the Tape and Reel option.

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specification contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is one of the most secure families of its kind on the market today, when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods used to breach the code protection feature. All of these methods, to our knowledge, require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Most likely, the person doing so is engaged in theft of intellectual property.
- Microchip is willing to work with the customer who is concerned about the integrity of their code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of their code. Code protection does not mean that we are guaranteeing the product as "unbreakable."

Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Information contained in this publication regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications. MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION, INCLUDING BUT NOT LIMITED TO ITS CONDITION, QUALITY, PERFORMANCE, MERCHANTABILITY OR FITNESS FOR PURPOSE. Microchip disclaims all liability arising from this information and its use. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, AnyRate, AVR, AVR logo, AVR Freaks, BeaconThings, BitCloud, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, Heldo, JukeBlox, KEELoQ, KEELoQ logo, Kleer, LANCheck, LINK MD, maXStylus, maXTouch, MediaLB, megaAVR, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, Prochip Designer, QTouch, RightTouch, SAM-BA, SpyNIC, SST, SST Logo, SuperFlash, tinyAVR, UNI/O, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

ClockWorks, The Embedded Control Solutions Company, EtherSynch, Hyper Speed Control, HyperLight Load, IntelliMOS, mTouch, Precision Edge, and Quiet-Wire are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, BodyCom, chipKIT, chipKIT logo, CodeGuard, CryptoAuthentication, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, EtherGREEN, In-Circuit Serial Programming, ICSP, Inter-Chip Connectivity, JitterBlocker, KleerNet, KleerNet logo, Mindi, MiWi, motorBench, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PureSilicon, QMatrix, RightTouch logo, REAL ICE, Ripple Blocker, SAM-ICE, Serial Quad I/O, SMART-I.S., SQI, SuperSwitcher, SuperSwitcher II, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

Silicon Storage Technology is a registered trademark of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2017, Microchip Technology Incorporated, All Rights Reserved.

ISBN: 9781522414520

QUALITY MANAGEMENT SYSTEM
CERTIFIED BY DNV
== ISO/TS 16949 ==

Microchip received ISO/TS-16949:2009 certification for its worldwide headquarters, design and wafer fabrication facilities in Chandler and Tempe, Arizona; Gresham, Oregon and design centers in California and India. The Company's quality system processes and procedures are for its PIC® MCUs and dsPIC® DSCs, KEELoQ® code hopping devices, Serial EEPROMs, microperipherals, nonvolatile memory and analog products. In addition, Microchip's quality system for the design and manufacture of development systems is ISO 9001:2000 certified.



MICROCHIP

Worldwide Sales and Service

AMERICAS

Corporate Office
2355 West Chandler Blvd.
Chandler, AZ 85224-6199
Tel: 480-792-7200
Fax: 480-792-7277
Technical Support:
<http://www.microchip.com/support>
Web Address:
www.microchip.com

Atlanta
Duluth, GA
Tel: 678-957-9614
Fax: 678-957-1455

Austin, TX
Tel: 512-257-3370

Boston
Westborough, MA
Tel: 774-760-0087
Fax: 774-760-0088

Chicago
Itasca, IL
Tel: 630-285-0071
Fax: 630-285-0075

Dallas
Addison, TX
Tel: 972-818-7423
Fax: 972-818-2924

Detroit
Novi, MI
Tel: 248-848-4000

Houston, TX
Tel: 281-894-5983

Indianapolis
Noblesville, IN
Tel: 317-773-8323
Fax: 317-773-5453
Tel: 317-536-2380

Los Angeles
Mission Viejo, CA
Tel: 949-462-9523
Fax: 949-462-9608
Tel: 951-273-7800

Raleigh, NC
Tel: 919-844-7510

New York, NY
Tel: 631-435-6000

San Jose, CA
Tel: 408-735-9110
Tel: 408-436-4270

Canada - Toronto
Tel: 905-695-1980
Fax: 905-695-2078

ASIA/PACIFIC

Asia Pacific Office
Suites 3707-14, 37th Floor
Tower 6, The Gateway
Harbour City, Kowloon

Hong Kong
Tel: 852-2943-5100
Fax: 852-2401-3431

Australia - Sydney
Tel: 61-2-9868-6733
Fax: 61-2-9868-6755

China - Beijing
Tel: 86-10-8569-7000
Fax: 86-10-8528-2104

China - Chengdu
Tel: 86-28-8665-5511
Fax: 86-28-8665-7889

China - Chongqing
Tel: 86-23-8980-9588
Fax: 86-23-8980-9500

China - Dongguan
Tel: 86-769-8702-9880

China - Guangzhou
Tel: 86-20-8755-8029

China - Hangzhou
Tel: 86-571-8792-8115
Fax: 86-571-8792-8116

China - Hong Kong SAR
Tel: 852-2943-5100
Fax: 852-2401-3431

China - Nanjing
Tel: 86-25-8473-2460
Fax: 86-25-8473-2470

China - Qingdao
Tel: 86-532-8502-7355
Fax: 86-532-8502-7205

China - Shanghai
Tel: 86-21-3326-8000
Fax: 86-21-3326-8021

China - Shenyang
Tel: 86-24-2334-2829
Fax: 86-24-2334-2393

China - Shenzhen
Tel: 86-755-8864-2200
Fax: 86-755-8203-1760

China - Wuhan
Tel: 86-27-5980-5300
Fax: 86-27-5980-5118

China - Xian
Tel: 86-29-8833-7252
Fax: 86-29-8833-7256

ASIA/PACIFIC

China - Xiamen
Tel: 86-592-2388138
Fax: 86-592-2388130

China - Zhuhai
Tel: 86-756-3210040
Fax: 86-756-3210049

India - Bangalore
Tel: 91-80-3090-4444
Fax: 91-80-3090-4123

India - New Delhi
Tel: 91-11-4160-8631
Fax: 91-11-4160-8632

India - Pune
Tel: 91-20-3019-1500

Japan - Osaka
Tel: 81-6-6152-7160
Fax: 81-6-6152-9310

Japan - Tokyo
Tel: 81-3-6880-3770
Fax: 81-3-6880-3771

Korea - Daegu
Tel: 82-53-744-4301
Fax: 82-53-744-4302

Korea - Seoul
Tel: 82-2-554-7200
Fax: 82-2-558-5932 or
82-2-558-5934

Malaysia - Kuala Lumpur
Tel: 60-3-6201-9857
Fax: 60-3-6201-9859

Malaysia - Penang
Tel: 60-4-227-8870
Fax: 60-4-227-4068

Philippines - Manila
Tel: 63-2-634-9065
Fax: 63-2-634-9069

Singapore
Tel: 65-6334-8870
Fax: 65-6334-8850

Taiwan - Hsin Chu
Tel: 886-3-5778-366
Fax: 886-3-5770-955

Taiwan - Kaohsiung
Tel: 886-7-213-7830

Taiwan - Taipei
Tel: 886-2-2508-8600
Fax: 886-2-2508-0102

Thailand - Bangkok
Tel: 66-2-694-1351
Fax: 66-2-694-1350

EUROPE

Austria - Wels
Tel: 43-7242-2244-39
Fax: 43-7242-2244-393

Denmark - Copenhagen
Tel: 45-4450-2828
Fax: 45-4485-2829

Finland - Espoo
Tel: 358-9-4520-820

France - Paris
Tel: 33-1-69-53-63-20
Fax: 33-1-69-30-90-79

France - Saint Cloud
Tel: 33-1-30-60-70-00

Germany - Garching
Tel: 49-8931-9700

Germany - Haan
Tel: 49-2129-3766400

Germany - Heilbronn
Tel: 49-7131-67-3636

Germany - Karlsruhe
Tel: 49-721-625370

Germany - Munich
Tel: 49-89-627-144-0
Fax: 49-89-627-144-44

Germany - Rosenheim
Tel: 49-8031-354-560

Israel - Ra'anana
Tel: 972-9-744-7705

Italy - Milan
Tel: 39-0331-742611
Fax: 39-0331-466781

Italy - Padova
Tel: 39-049-7625286

Netherlands - Drunen
Tel: 31-416-690399
Fax: 31-416-690340

Norway - Trondheim
Tel: 47-7289-7561

Poland - Warsaw
Tel: 48-22-3325737

Romania - Bucharest
Tel: 40-21-407-87-50

Spain - Madrid
Tel: 34-91-708-08-90
Fax: 34-91-708-08-91

Sweden - Gothenberg
Tel: 46-31-704-60-40

Sweden - Stockholm
Tel: 46-8-5090-4654

UK - Wokingham
Tel: 44-118-921-5800
Fax: 44-118-921-5820

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Microchip:](#)

[KSZ9896CTXC](#) [KSZ9896CTXC-TR](#) [KSZ9896CTXI](#) [KSZ9896CTXI-TR](#)



Стандарт Электрон Связь

Мы молодая и активно развивающаяся компания в области поставок электронных компонентов. Мы поставляем электронные компоненты отечественного и импортного производства напрямую от производителей и с крупнейших складов мира.

Благодаря сотрудничеству с мировыми поставщиками мы осуществляем комплексные и плановые поставки широчайшего спектра электронных компонентов.

Собственная эффективная логистика и склад в обеспечивает надежную поставку продукции в точно указанные сроки по всей России.

Мы осуществляем техническую поддержку нашим клиентам и предпродажную проверку качества продукции. На все поставляемые продукты мы предоставляем гарантию .

Осуществляем поставки продукции под контролем ВП МО РФ на предприятия военно-промышленного комплекса России , а также работаем в рамках 275 ФЗ с открытием отдельных счетов в уполномоченном банке. Система менеджмента качества компании соответствует требованиям ГОСТ ISO 9001.

Минимальные сроки поставки, гибкие цены, неограниченный ассортимент и индивидуальный подход к клиентам являются основой для выстраивания долгосрочного и эффективного сотрудничества с предприятиями радиоэлектронной промышленности, предприятиями ВПК и научно-исследовательскими институтами России.

С нами вы становитесь еще успешнее!

Наши контакты:

Телефон: +7 812 627 14 35

Электронная почта: sales@st-electron.ru

Адрес: 198099, Санкт-Петербург,
Промышленная ул, дом № 19, литера Н,
помещение 100-Н Офис 331