# Intel® Media Switch IXE5418 Gigabit Device

***Data Sheet***

## Product Features/Benefits

- Single chip, 8-port Gigabit Layer 2/3/4 switch/router
  - — High level of integration, compact footprint and low power dissipation allows high port density designs at the lowest cost-per-port
- Integrated 10/100/Gigabit Ethernet MACs
  - — Easy migration from 10/100 to Gigabit
- Wire speed performance across all ports in switching or IP/IPX routing modes
  - — Delivers congestion-free performance during periods of peak load typically seen with enterprise switches
- Hardware-assisted stacking and several Layer 2/3 protocols
  - — Reduces complexity and cost of CPU sub-system—significant in stacks or in chassis designs
- Link aggregation in any combination of up to 4 ports per group
  - — Meshed configurations with redundant paths can be built for fail-safe networks
- Advanced traffic prioritizing, QoS and bandwidth management capabilities
  - — Enables convergence of voice, video and data traffic on Ethernet/IP networks
- Fully compliant with VLAN implementation standards based on ports, tags and addresses
  - — Permits building flat, plug-and-play networks that are easy to maintain
- Advanced multicast, broadcast and filtering capabilities
  - — Enables video and voice multicasting on IP networks, protects broadcast storms and allows the building of high performance intranet firewalls

*Note:* *Key features and benefits for the IXE5418 device are given above. Please refer to Section 10.0 for a complete feature list.*



Order Number: 273363-001
May 2000

# Contents

# Figures



# Tables

# 1.0 General Description

The Intel® Media Switch IXE5418 Gigabit Device supports eight10/100/Gigabit ports, integrating the MACs, packet storage, switching, routing, and queuing logic on-chip. The IXE5418 device is capable of wire speed Layer 2 switching and wire speed IP/IPX Layer 2/3/4 switching/routing on all ports. It provides advanced filtering, quality of service, mirroring, prioritizing capabilities, and Layer 4 bandwidth management features.

A system designed using an IXE5418 device requires transceivers on the 10/100/Gigabit ports, memory for storing switching data structures and a CPU subsystem. The IXE5418 device interfaces to the CPU subsystem via a 32-bit PCI bus.

# 2.0 Applications

- Eight port 10/100/Gigabit Layer 2 and Layer 2/3/4 workgroup switches.
- Cascadable high port count Layer 2 switches and Layer 2/3/4 switch/routers when using the Gigabit ports for cascading.
- Layer 2/3/4 switch/routers and advanced bandwidth management features.

**Figure 1. IXE5418 Block Diagram**

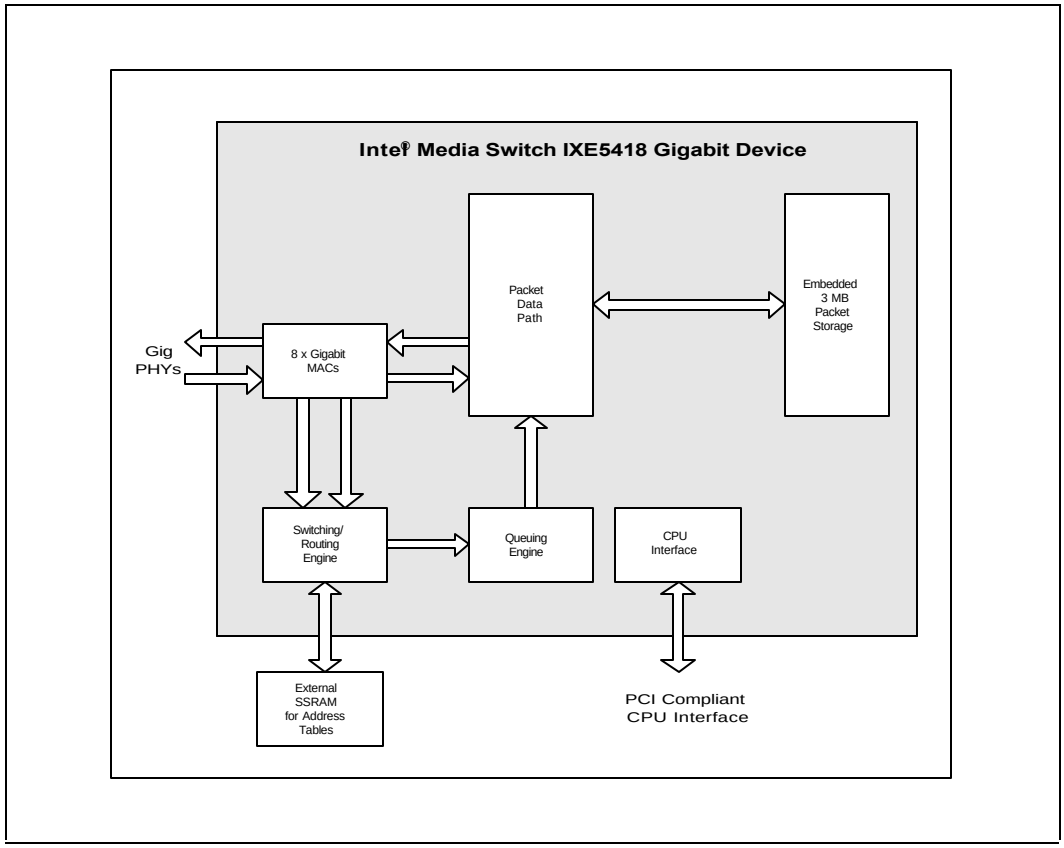

# 3.0 Functional Description

## 3.1 Introduction

The IXE5418 device is a highly integrated Layer 2 switch, as well as a Layer 2/3/4 switch/router. It supports eight 10/100/Gigabit ports with on-chip MACs. It also supports integrated switching and routing logic, on-chip packet queuing memory, and on-chip packet storage memory. The IXE5418 is capable of switching and routing packets at wire speed on all ports regardless of packet size.

In Layer 2 mode, the IXE5418 supports:

- IEEE 802.1D, 1998 Edition standard, including 802.1P, 802.1Q

It also supports the IEEE 802.3, 1998 Edition standard, which includes:

- 802.3x flow control
- 802.3z Gigabit Ethernet standards
- 802.3ad Link Aggregation standards

It provides on-chip 32-bit statistics counters to support the MIB-II, the Etherstats group of RMON, and SMON standards. In addition, it provides counters for monitoring flow based statistics. Details of the statistics counters are provided in future sections.

In addition to the IEEE standards support, the IXE5418 provides many advanced features all in a single device such as Filtering, Mirroring, Bandwidth Management, and Broadcast and Multicast Storm Control. Details of these features are provided in future sections.

The IXE5418 connects to other devices using standards based interfaces such as GPCS (SERDES)/GMII, CPU, and SSRAM. This simplifies the system design and provides flexibility when used with other devices. The following sections describe the external interfaces to the IXE5418 device.

Figure 2 illustrates a system block diagram of the IXE5418 device.

**Figure 2. IXE5418 System Block Diagram**

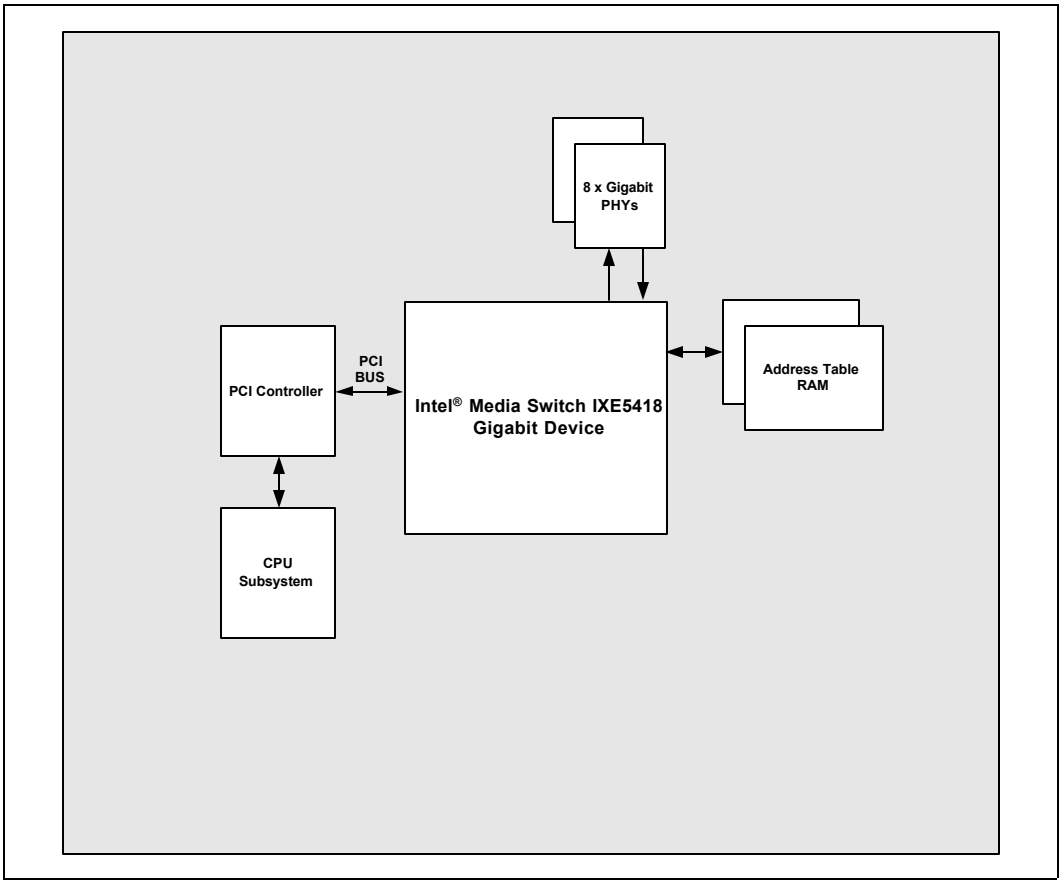

# 4.0 Interface Descriptions

Table 1 provides a list of IXE5418 interfaces.

**Table 1. IXE5418 Interfaces**

| Interface | Description |
|---|---|
| CPU | PCI compliant interface with bus mastering capability for packet and unresolved entry transfers to CPU. |
| GMII/GPCS | Configurable GMII or 8B/10B encoded GPCS interface used to link the IXE5418 to the eight Gigabit ports. |
| Address Table SRAM | Dual external 64-bit wide Synchronous SRAMs that are used for address storage. The size of each device supported on the interfaces are programmable from 64K x 32 to 256K x 32, with two 32-bit devices per bank. |
| MDIO PHY Scanning | Phy Scan provides a parallel to serial data converter that can be used to send and receive data to the external transceivers over the MDIO/MDC interface. The parallel to serial data converter also is equipped with an automatic polling machine that monitors the condition of the transceivers and interrupts the CPU on a state change. |
| LED | Can be attached to four external HC595-type shift registers. Drives 4 bits of status per port onto the interface. |

# 5.0 Data Structures

## 5.1 Switching Block Entries

The Switching Engine block of the IXE5418 device uses a group of data structures when making its forwarding decisions. These data structures called “entries” are linked to each other via pointers and are usually stored in an external SSRAM (some of them can also be stored on-chip), as shown below in Figure 3.

Each type of “entry” contains different information that allows the IXE5418 to support its diverse set of features. Multiple addresses can point to the same entry type unless they use different features. This results in compact and efficient memory usage, with memory requirements increasing only as more of these features are enabled.

**Figure 3. IXE5418 Data Structures**

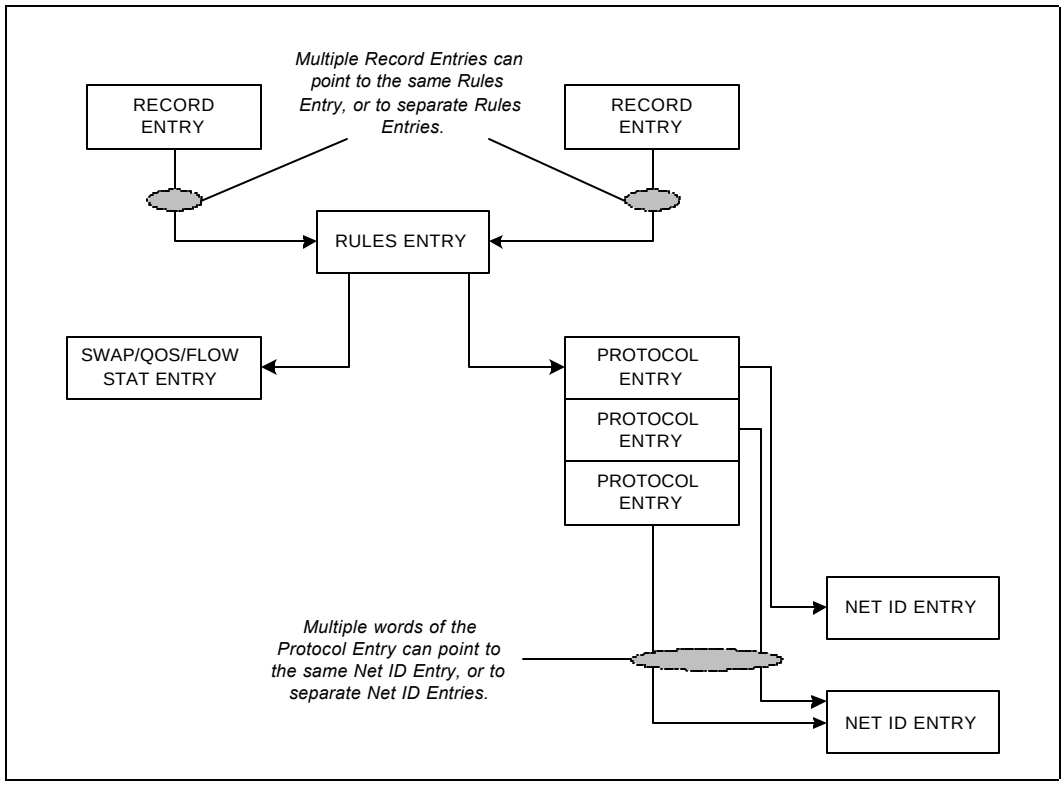

Table2 entries are maintained by the Switching Engine block for Layer 2, IP, and IPX addresses.

**Table 2. Switching Block Entries**

| Entry Type | Description |
|---|---|
| Record Entry | The basic entry associated with every address. Contains a pointer directing to another entry called the Rules Entry. |
| Rules Entry | Contains information about the port on which this address resides, as well as some of the filter, mirror, and priority rules. Contains further pointers to a QoS/Flow Stat Entry and a Protocol Entry. |
| QoS/Flow Stat Entry | Contains information that allows configuration of QoS flows and configuration of flow based statistics gathering (source address to destination address). For IP and IPX addresses, this QoS/Flow Stat Entry also contains the Destination Swap Address and therefore is called the Dst Swap/ QoS/ Flow Stat Entry. |
| Protocol Entry | Contains further filter, mirror, and priority configuration information and can be configured as a list of one or more words with different protocols selecting to different words in the same entry. Contains one or more pointers to a Net ID Entry. |
| Net ID Entry | Contains the Prepend Word feature which can be used to cascade other devices using the IXE5418 as a switching matrix. This entry also contains the source address and destination address based mirror port information for Global Source and Global Destination address based mirroring. |

## 5.2 Layer 2 Protocol Support

The Layer 2 protocols supported are: ARP, RARP, AppleTalk, DECNet, SNA, NetBios, DLC/LLC, IP, and IPX.

## 5.3 IP Switching/Routing features

The IXE5418 will forward the first packet of the flow to the CPU if it does not find a forwarding entry in its tables.

The CPU programs the outgoing port number, as well as the Ethernet address of the next hop or destination address, into the IXE5418 tables. The first packet must then be routed on the port that has the destination node connected through it.

Once the entries are created in the IXE5418 tables for the source and destination, all the packets belonging to the flow are routed in the hardware at wire speed. Packets belonging to protocols other than IP and IPX will be switched in the hardware at wire speeds using the Layer 2 switching algorithm.

### 5.3.1 Routing Domains

A routing domain is an IP network. An IP address and a subnet mask identify an IP routing domain.

Multiple routing domains can exist on the same port (in which case, they are identified by 802.1Q VLAN tags) or multiple ports can belong to a routing domain. All the ports that belong to a routing domain have the same IP address. Associated with each IP network is a routing domain number.

The IXE5418 device supports up to 256 routing domains. Each routing domain has an Ethernet address associated with it.

### 5.3.2 IP Protocol Support

The IP protocols supported are: TCP, UDP, ICMP, IGMP, EGP, OSPF, RSVP, and IGRP.

## 5.4 IPX Switching/Routing Features

### 5.4.1 Network and Node Addresses

For IPX packet processing, the IXE5418 uses network or node addresses to perform the switching feature.

The IXE5418 maintains a separate address table for IPX network addresses and another address table for IPX node addresses, supporting one routing domain per port for IPX networks.

### 5.4.2 IPX Protocol Support

There are registers provided for each protocol that indicate the increment from the Protocol Offset where the protocol specific information for that protocol may be obtained.

The IPX protocols supported are: SPX, RIP, NCP, NLSP, SAP, and PROP.

## 5.5 Port-based VLAN Features

The IXE5418 provides Port-based VLAN features by allowing each port to specify the group of ports that belong in that VLAN. The Port based VLANs are the default means of creating VLANs when no other type of VLANs (for example, 802.1Q Tag) have been programmed.

A separate VLAN per port is provided for:

- Layer 2 packets
- IP packets
- IPX packets

### 5.5.1 Stack Port Net ID Entry

The stack port Net ID Entry is same as the port Net ID Entry. This table is used only in stacked configuration for the packets coming on the stack ports. The table is indexed using the Device Number and Port Number. The table is located in the external ram and uses the same window that is used for the Net ID entries.

The Stack Port Net ID Entry values are:

- Transmit Enable
- Transmit Enable CPU

### 5.5.2 Stack Port to Local Port Lookup

In Stack Mode, the Rules Entry has the device number and port number that corresponds to the destination device and port in the stack.

The stack port to local port lookup values are:

- Stack Port to Local Port
- Stack Port Don't Send Unresolved to CPU

## 5.6 Protocol-based VLAN Features

Protocol-based VLANs can be configured on the IXE5418 by having different protocols point to different words in the Protocol Entry. This allows the protocols to use different Net ID Entries, and thus create different VLANs.

The Protocol Entry used by each protocol is programmed in the protocol registers.

## 5.7 802.1Q VLAN and 802.1D, 1998 Edition Priority and Class of Service Features

The IXE5418 provides extensive support for VLANs and priorities based on the 802.1Q and 802.1D, 1998 Edition specifications.

It provides four levels of queues for all ports. The 3-bit Tag Priority field from the tagged packet can be mapped into the four priority levels by using the Map Priority Level register. The outgoing packet can be modified to regenerate the Tag Priority by using the per port Port Regenerate Priority Entry. The Map Priority Level register is a global mapping that applies to all ports, while there is a per port Regenerate Priority Entry. The Regenerate Priority Entry is indexed using the receive port number.

Also supported is regenerating the priority the packet should go out on. If there is not a need for regenerating the priority, then the user can program the outgoing priority to be the same as the incoming priority. The outgoing priority is also used to determine which internal queue the packet will be queued to.

## 5.8 Port Aggregation Features

The IXE5418 supports Port aggregation on all ports in groups of up to 4 ports. Port aggregation groups are restricted to start on port numbers 0, 2, 4, and 6, with ports being aggregated in increasing port numbers.

The two modes of port aggregation supported are:

- Ingress aggregation only (the default mode),
- Ingress and Egress aggregation.

## 5.9 Interrupt Generation and Handling

The IXE5418 provides flexible interrupt generation mechanisms. It provides two interrupt pins and two on-chip mask registers that enable the CPU to map any interrupt to either of the interrupt pins. This allows the CPU to treat certain interrupts as high priority and others as low priority.

The CPU can also mask off any interrupt from both registers and effectively convert the bit into a poll bit.

## 5.10 Packet Transfers to CPU

The IXE5418 provides extensive packet processing support in its hardware in order to reduce the burden of the CPU.

It provides four queues for packets that are to be sent to the CPU. These levels are described below. Each queue can be individually turned on or off to start and stop the CPU from receiving packets from that queue, and can be assigned a guaranteed bandwidth. This allows the CPU to prioritize the types of packets that it wants to process.

The DMA engines that are provided per queue allow the IXE5418 to directly transfer data into the CPU memory without CPU intervention.

### 5.10.1 CPU Queues

CPU Queue entries use the following values:

- Queue 3
  - — All packets in the reserved BPDU address range except 802.3x.
  - — Packets which are handled by the MAC and ASIC internally.
  - — GVRP and GMRP packets (if these protocols are enabled).
- Queue 2
  - — All packets (IP, IPX, L2) addressed to CPU.
- Queue 1
  - — All broadcast packets.
  - — All multicast packets (enabled to go to CPU).
  - — All packets that go to CPU through port mirroring.
- Queue 0
  - — All Other CPU packets including error packets.

### 5.10.2 Packet Headers

The IXE5418 prepends a header to every packet it writes into CPU memory. This header occupies one 32 byte section. The start of the packet is always aligned to a 32 byte boundary in the CPU memory. Thus, if a 65 byte packet is sent to the CPU memory, the header occupies the first 32 byte

section, and the packet occupies the next 3 sections, for a total of 128 bytes. The next packet header will then occupy the next 32 byte section in CPU memory. The CPU should decrement the Level Counter by the 4 for the above example.

There is a per queue bit that allows the CPU to turn off any queue from writing packets into the CPU memory.The IXE5418 prepends a 32-byte header to every packet sent to the CPU.

The IXE5418 expects the CPU to prepend a 32-byte header to every packet it receives from the CPU. This header allows the CPU to inform the IXE5418 on a packet by packet basis whether the IXE5418 should do an address resolution on the source and destination addresses, or whether it should use the information in the header to switch the packet.

## 5.11 Unresolved Packet Transfers to CPU

The IXE5418 also maintains four queues for unresolved packets.

Each queue can be individually turned on or off and assigned a guaranteed bandwidth. This allows the CPU to prioritize the types of packets that it wants to process.

The DMA engines provided per queue allow the IXE5418 to directly transfer data into the CPU memory without CPU intervention.

The only difference between “unresolved packet transfers to CPU” versus “resolved packet transfers to CPU”, is that unresolved packets have an option to not send the whole packet, but only part of the packet to the CPU. The value by which the read pointer should be incremented, and the value that should be written to the Level Counter register can be determined from the Unres Xfer Size field of the header.

For example, if the Unres Xfer Size field contains 0, then the Read Pointer should be incremented by 32 + 32 = 64 bytes, and the Level Counter register should be written with a value of 64.

### 5.11.1 Unresolved Queues

Unresolved Queue entries use the following values:

- Queue 3
  - — Layer 2 packets with source unresolved.
  - — Layer 2 packet with source port moved.
  - — Layer 2 packets include ARP and RARP packets. If IP and IPX switching and routing are disabled, all packets are treated as Layer 2 packets.
- Queue 2—If IP/IPX switching and routing is enabled:
  - — IPX and IP packets unresolved source.
  - — IPX and IP packets with source port move.
- Queue 1—packets include:
  - — IP and IPX packets that need to be routed, but for which the destination swap valid bit not set.
  - — Layer 4 packets with source or destination unresolved, given that the other is resolved.

- Queue 0
  - All Other Unresolved Packets
  - The remainder of the unresolved packets including all the destination unresolved cases and search counter expired packets.

The IXE5418 prepends a header to every unresolved packet to the CPU.

## 5.12 Packet Transfers From CPU

The IXE5418 provides two queues, which are supported by DMA, for packet transfers from the CPU.

Queue 1 is treated as the high priority queue, and is always processed before queue 0, unless a packet transfer had already started for queue 0. In this case, the complete packet from queue 0 is fetched from the CPU memory first, and then the packet from queue 1 is processed. The two queues allow a high priority process (for example, Spanning tree BPDUs) to not have its packets stuck behind other lower-priority packets.

# 6.0 Hardware Assisted Features

This sections describes the hardware assisted features provided in the IXE5418.

## 6.1 Address Learning

The IXE5418 facilitates address learning by using a hardware engine that provides a content addressable memory-like interface to the CPU.

## 6.2 Learning of Socket Address

The IXE5418 provides a separate address learning interface for Layer 4 Socket addresses in order to speed up the process of application based rules generation. The IXE5418 also provides an add, delete, and lookup state machine called the Layer 4 processor. This simplifies the adding and deleting of socket addresses. Finally, the IXE5418 has a separate aging state machine for aging Layer 4 Record Entries.

## 6.3 Address Aging

The IXE5418 contains hardware assistance for the entry aging process. It has three fundamental modes of operation: automatic, manual, and sparse manual. Each record entry has 3 bits which represent up to 8 age zones.In all modes, the aging controller steps through a programmable number of zones, broadcasting the current zone (also known as the “age zone”) throughout the ASIC.

## 6.4 Statistics Counters

The IXE5418 provides MAC-based and switch-based statistics gathering support on chip.

## 6.5 Flow-based Statistics

The Flow-based statistics features allow receive byte and packet count information to be kept for selected flows.

The statistic is kept in:

- “Layer 2 Qos/Flow” entry for layer 2 flows
- “IP Dest Swap/QoS/Flow” entry for IP based flows
- “IPX Dest Swap/QoS/Flow” entry for IPX based flows

## 6.6 Bandwidth Management

When performing bandwidth management, it is necessary to account for the size of the packets being arbitrated, not just the number of packets. Also, latency can become a problem. The IXE5418 supports two types of algorithms per port: the Credit Algorithm and the Strict Priority Algorithm.

## 6.7 QoS

The QoS logic in the IXE5418 allows certain flows to be mapped as QoS. These flows can be queued to a certain priority that has been assigned a guaranteed bandwidth. In addition, the rate at which the flow is sending is constantly monitored, and this is compared to the configured rate for this flow. If the rate is exceeded, then the packets that caused this exception are dropped.

## 6.8 Broadcast and Multicast Storm Control

The IXE5418 provides a per port configuration that allows broadcast and/or multicast storm control. The CPU can program a threshold value per port that indicates the number of broadcast and/or multicast packets that are allowed in a given time interval.

## 6.9 IP Multicast Routing

The IXE5418 supports up to 256 routing domains without placing any per port limitations. A port can have one or more routing domains (or IP networks) associated with it. An IP multicast packet that must be forwarded on such a port, may need to be sent out multiple times. Each time with a different 802.1Q tag associated with that routing domain. The number of times the packet must sent depends on the number of routing domains on that port that have members that belong to that multicast group. The IXE5418 allows up to 256 multicast group members on a port.

## 6.10 Prepend Word

The Prepend Word feature is fundamental to the stacking mechanism and is inserted at the head of every frame that is sent on the stack port. This word contains important information that allows all the stack devices to operate independently of each other.

# 7.0 Electrical and Environmental Specifications

This section summarizes the electrical and environmental specifications for the IXE5418. The IXE5418 supports both 5V and 3.3V signaling environments.

## 7.1 Absolute Maximum Rating

Applying stresses beyond the absolute maximum may cause unrecoverable damage to the device. Operation of this product is not implied for any condition beyond the ranges specified in the functional operation range. Operating this product at the absolute maximum rating for a prolonged period can negatively impact device reliability. Table3 lists the absolute maximum ratings for the IXE5418.

**Table 3. Absolute Maximum Ratings**

| Parameter | Sym | Min | Max | Units |
|---|---|---|---|---|
| Supply Voltage Core | VDDCORE | -0.3 | 2.75 | V |
| Supply Voltage IO | VDDIO | -0.3 | 3.6 | V |
| Operating Temperature | $T_{op}$ | 0 | 50 | $^{o}C$ |
| Storage Temperature | $T_{st}$ | -65 | 125 | $^{o}C$ |

***Caution:*** Exceeding these values may cause permanent damage. Functional operation under these conditions is not implied. Exposure to maximum rating conditions for extended periods may affect device reliability.

## 7.2 Thermal Resistance

Table4 lists the thermal resistance ratings for the IXE5418

**Table 4. Thermal Resistance**

| Thermal Resistance | Parameter | Value |
|---|---|---|
| $\theta_{JC}$ | Junction to case | |
| $\theta_{CA}$ | Case to ambient | |
| $\theta_{JA}$ | Junction to ambient | |

Table5 lists the thermal resistance with air flow ratings for the IXE5418

**Table 5. Thermal Resistance with Air Flow**

| Thermal Resistance | Package | 0 LFPM | 100 LFPM | 200 LFPM | 300 LFPM | 500 LFPM |
|---|---|---|---|---|---|---|
| $\theta_{JA}$ ($^{o}$ C/W) | TBGA Type IA | | | | | |

Heat Sink may be required depending on the airflow in the system.

## 7.3 Functional Operating Range

Table 6 lists the functional operating range. Please refer to the “Absolute Maximum Rating” section for details about the absolute maximum ratings.

### Table 6. Operating Conditions

| Parameter | Sym | Min | Typ[1] | Max | Units | Test Conditions |
|---|---|---|---|---|---|---|
| Recommended Supply Voltage VDDCORE | VDDCORE | 2.375 | 2.50 | 2.625 | V[2] | |
| Recommended Supply Voltage VDDIO | VDDIO | 3.135 | 3.3 | 3.456 | V[2] | |
| Recommended Operating Temperature | $T_{op}$ | 0 | | 50 | °C | |
| Power Supply Current | $I_{DD}$ | | TBD | | A | |
| 3.3V/2.5V | $P_{tot}$ | | | 6.0 | W | |
| Maximum Junction Temperature | $T_{jmax}$ | | | 110 | °C | |

1. Typical values are at 25 °C and are for design aid only; not guaranteed and not subject to production testing.
2. Voltages with respect to ground unless otherwise specified.

## 7.4 3.3 Volt DC Specifications

Table 7 lists the 3.3 V signaling specifications for the IXE5418.

### Table 7. 3.3 V Signaling Specifications

| Parameter | Symbol | Condition | Minimum | Maximum |
|---|---|---|---|---|
| Input high voltage | Vih | - | 0.5 x Vdd | Vdd_clmp + 0.5 V[1] |
| Input low voltage | Vil | - | -0.5 V[2] | 0.3 x Vdd |
| Input leakage current | Iil | 0< Vin < Vdd | -15 µA | 15 µA |
| Output high voltage | Voh | Iout = -9.3 mA | 0.9 x Vdd | - |
| Output low voltage | Vol | Iout = 15 mA | - | 0.1 x Vdd |
| Pin capacitance | Cin | - | 3 pF | 7 pF |

1. Overvoltage protection maximum is +11 V for 11 ns, through a 25 Ω resistor.
2. Overvoltage protection minimum is -5.5 V for 11 ns, through a 25 Ω resistor.

Table 8 lists the current and loading specifications for the IXE5418.

**Table 8. Current and Loading Specification**

| Signal Group | Signal Name | IOL (mA) | Capacitive load (pF) |
|---|---|---|---|
| SRAM | SCS_N, SADSC_N, SADDR, SDATA, SWE_N, SOE_N | 8 | SRAM |
| PCI | | 4 | PCI |
| GMII/GPCS | TDS[9:0]/TXER, TXEN, TXD[7:0] $TXCLK_p$ | 6 | GMII/GPCS |
| MMI | MDC, MDIO | 4 | MMI |

## 7.5 GMII/GPCS Port DC Specifications

The GMII/GPCS port electrical specifications are compliant with the IEEE 802.3z standard.

## 7.6 LED Interface

The IXE5418 provides a serial LED interface which can be attached to four external HC595-type shift registers.

# 8.0 Mechanical Specifications

Figure 4 and Figure 5 illustrate the IXE5418 contained in a PBGA package.

**Figure 4. IXE5418 Package Outline**

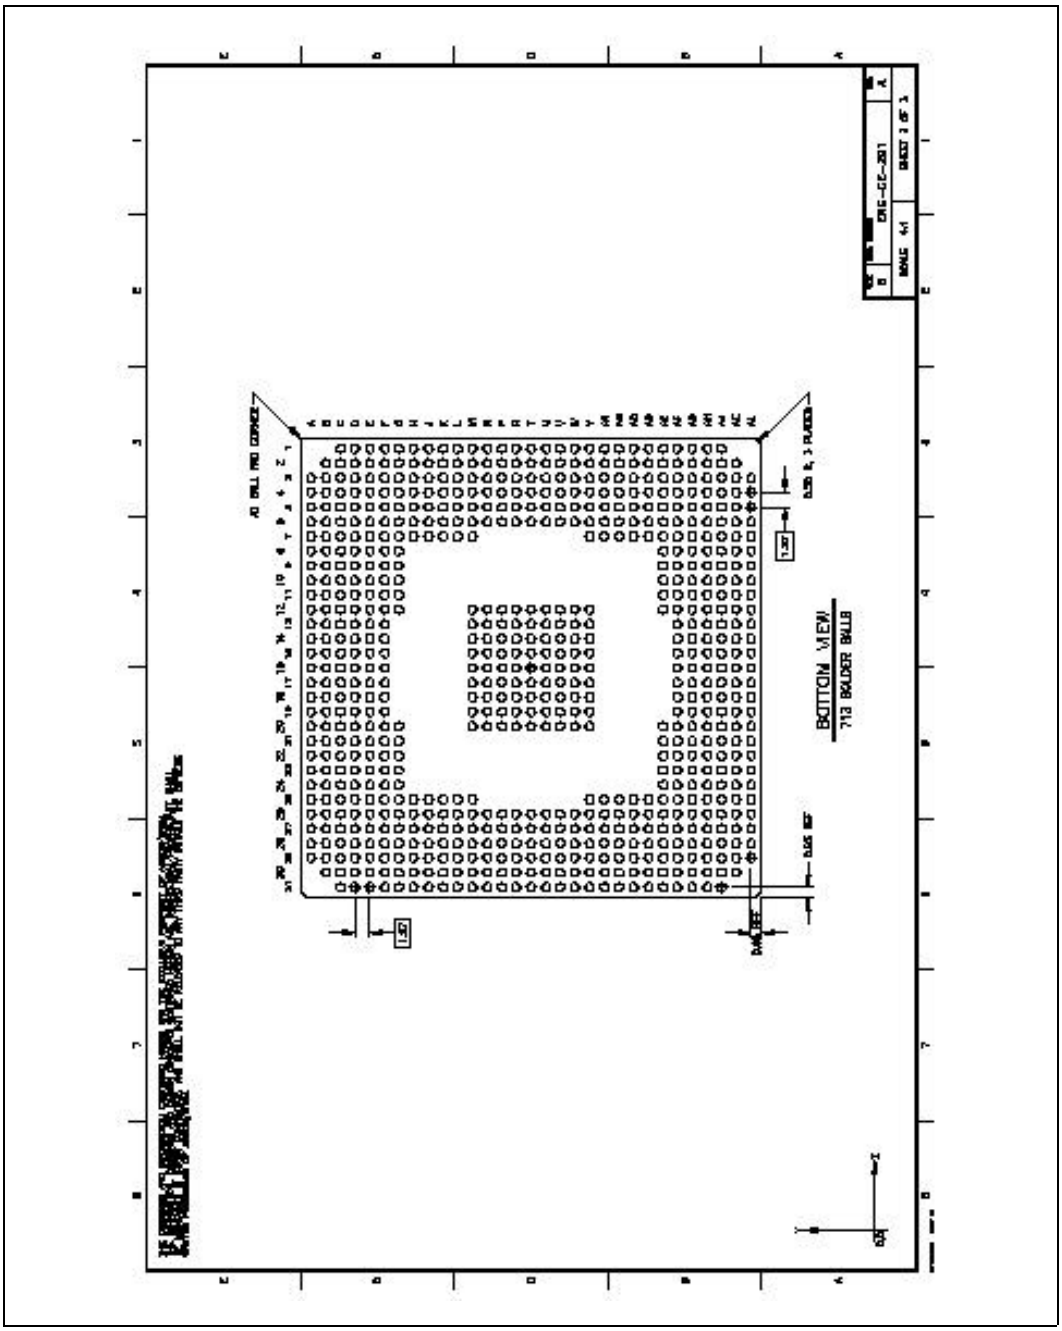

**Figure 5. IXE5418 Package Outline (Cont'd)**

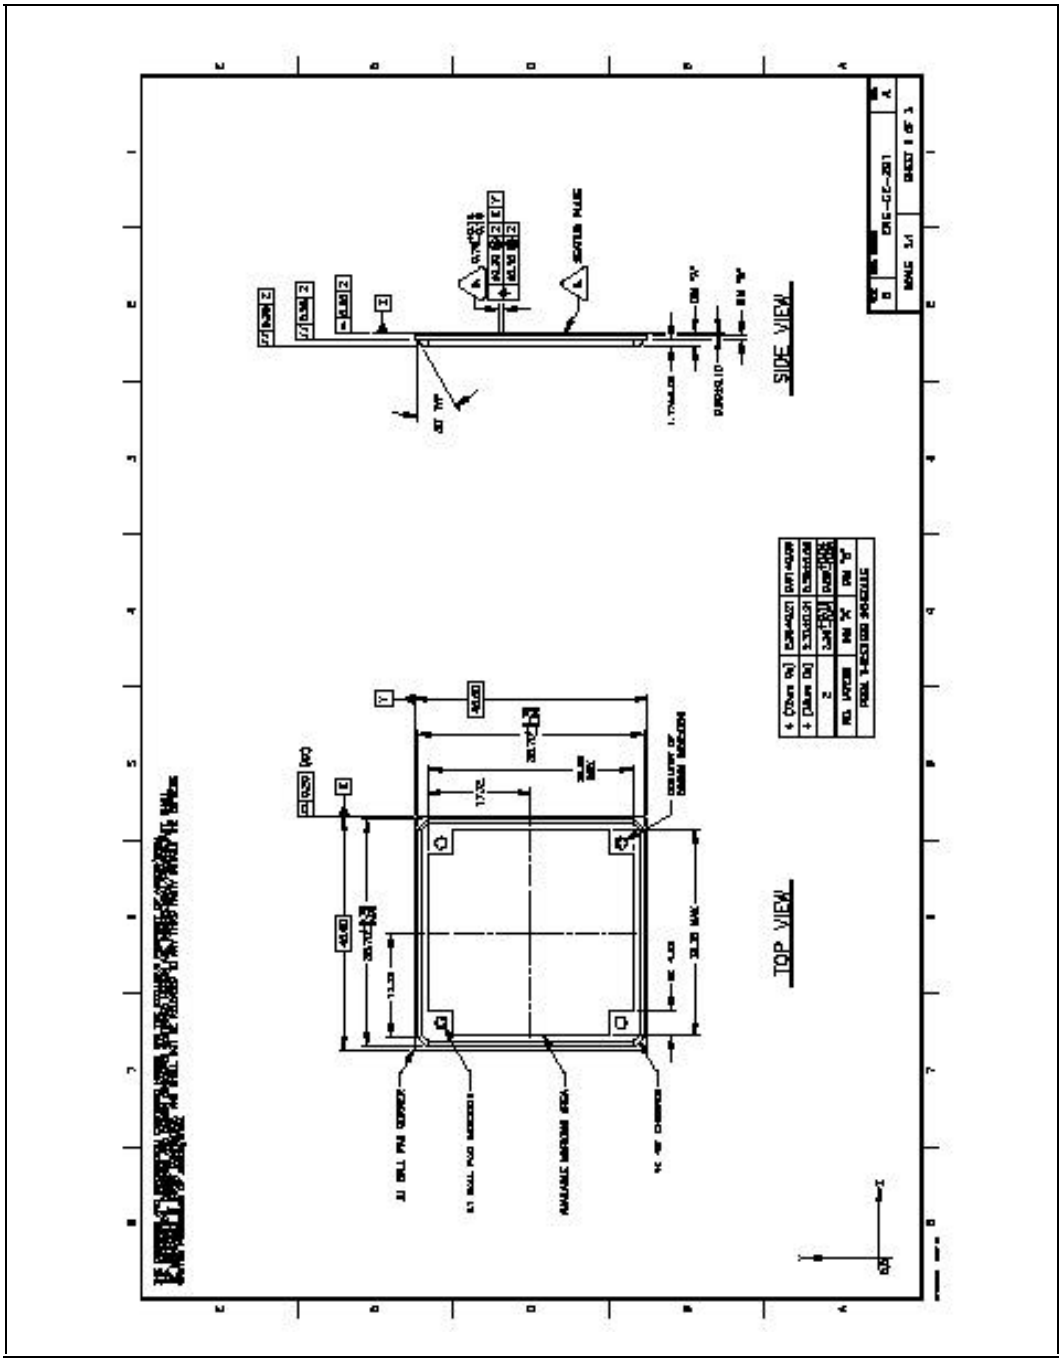

# 9.0 Software Support

## 9.1 APIs Supported

The IXE5418 supports the APIs illustrated in the software architecture block diagram shown in Figure 6.

**Figure 6. Software Architecture Illustrating APIs Supported by the IXE5418**

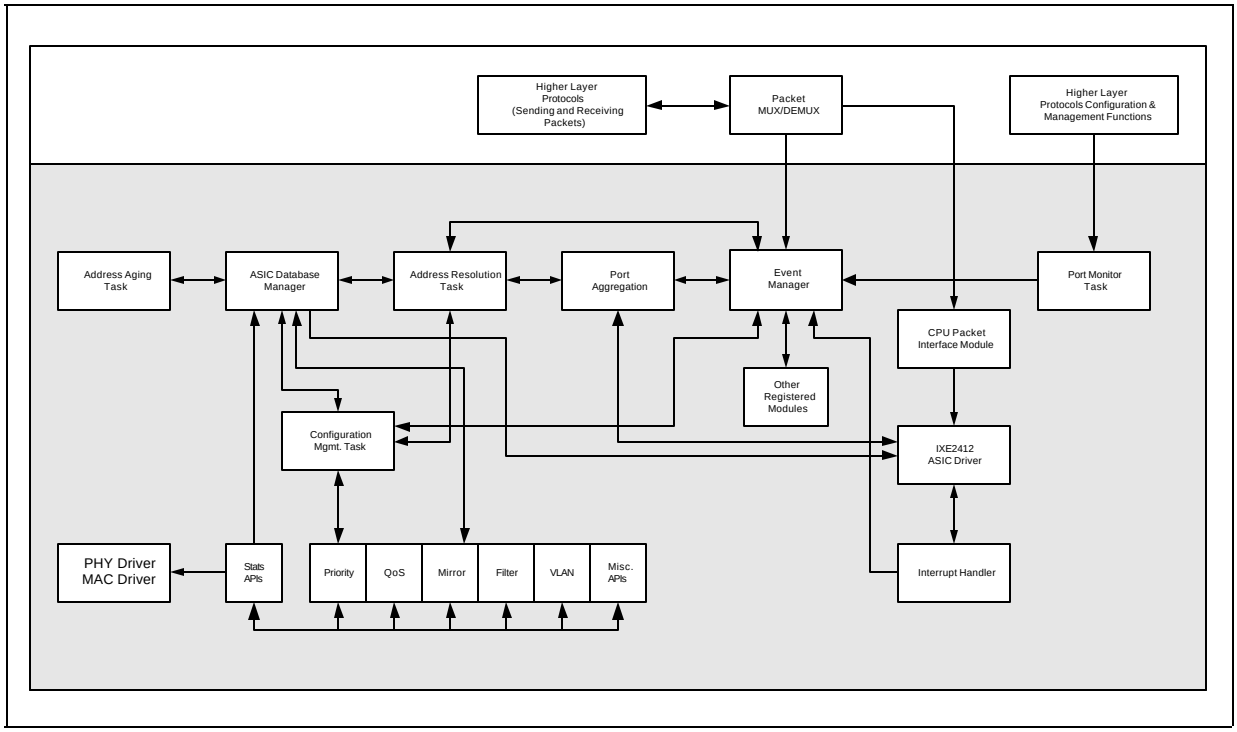

Summary descriptions and use of APIs supported by the IXE5418 are provided in Table9.

## Table 9. APIs Supported by the IXE5418

| API Module | Description/Use |
|---|---|
| Interrupt Handling | Handles the interrupt from the IXE5418. |
| OS Wrapper | Provides services needed from the underlying real time operating system. You can port these wrappers to your target operating system. |
| Notification Manager | Provides a generic method for distributing information within a system. Allows different software modules to register for events of interest and makes the information distribution transparent, modular, and flexible. |
| ASIC, Mac, and PHY | Enables you to program the performance of different functions. Provides functions for the ASIC, Mac, and PHY initialization, configuration, and management which include functions for ASIC initialization and routines for bit level manipulations of the IXE5418 ASIC registers for various configurations.<br>Mac and PHY provide functions to initialize the Ethernet controllers built into the IXE5418 and the Level One PHY (such as functions to change the speed, functions to change the duplex mode, etc.). |
| Address Resolution | Contains functions used for learning IP, IPX, and Layer 2 addresses. |
| IP and IPX Configuration and Management | Include the ability to perform a routing table lookup (for IP and IPX Address Resolution—call into IP and IPX routing modules), determine the Ethernet address of the destination station or the next-hop (for routed packets whose destination is unresolved to determine the address programmed into swap entry—call into ARP module for IP or SAP module for IPX), etc. |
| DMA Interface | Provides functions to send and receive packets between the ASIC and the CPU. The ASIC Driver provides a full set of functions supporting packet send and packet receive, and a separate set of functions that allow higher level software to manage the receive and send DMA buffer pools directly. |
| Address Aging | Provides functions for configuring an ageing interval. The IXE5418 ASIC tracks the address record entries that have been accessed over the ageing interval. Different aging time intervals can be specified for Layer 2, Layer 3, and Layer 4 entries. |
| Address Learning | Provides functions for address learning in the software. The hardware provides a CAM interface to facilitate fast learning of addresses. |
| Filters, Mirrors, Priorities, and Quality of Service | Supports configuring filters, mirrors, priorities, and quality of service for networks, nodes and ports. These APIs can be called from higher-layer software modules (such as SNMP agent) to configure these special rules for addresses. |
| VLAN | Provides APIs (based on ports, 802.1Q tags, and multicast addresses) to make VLAN configuration and management easier for higher-layer software modules such as GVRP, GMRP, or SNMP agent (for user-configured VLANs). |
| Statistics Gathering | Provides counters that count different events required for both standard and draft MIB implementations and functions for gathering Mac, PHY, and ASIC statistics, including RMON stats. The APIs provided for higher-layer software modules are for reading these counters. |
| Link Aggregation Configuration and Management | Support for port aggregation on all ports in groups of up to 4 ports. Supports Ingress Aggregation Only and Ingress and Egress Aggregation modes of which Ingress Aggregation Only mode is the default. |
| Miscellaneous | Provides functions for all other miscellaneous configurations, such as adding static entries to address tables, creating static routes, creating default routes, broadcast and multicast storm control, etc. The Miscellaneous API module interfaces to various modules including ASIC Database Manager, Configuration Management task, Address Resolution task, etc., to provide these functionalities to higher-layer protocol stacks. |

## 9.2 Protocols

The IXE5418 can perform extensive packet parsing and can obtain packet type (Type II, SNAP, etc.) and protocol (IP, IPX, GARP, GVRP, STP, etc.) information directly from the packet header.

Supported protocols are summarized in Table 10.

### Table 10. Protocols Supported by the IXE5418

| Protocol | Description/Use |
|---|---|
| GARP | GARP family protocols supported include GARP, GARP Multicast Registration Protocol (GMRP), and GARP VLAN Registration Protocol (GVRP), which comprise a task within a MAC bridge system that interacts with other tasks and the driver through messages and events.<br>GARP services are the basis for implementing GMRP and GVRP. GMRP provides the ability to register (and de-register) Group Address membership in a MAC bridge. GVRP provides the functionality to register (and de-register) VLANs dynamically. |
| IP | Supported IP protocol modules are the:<br>• Internet Protocol (IP)<br>• Address Resolution Protocol (ARP)<br>• Internet Control Message Protocol (ICMP)<br>• Internet Group Management Protocol (IGMP)<br>• User Datagram Protocol (UDP)<br>• Routing Information Protocol (RIPv1 and RIPv2)<br>• Internet Control Message Protocol (ICMP)<br>• User Datagram Protocol (UDP)<br>• Trace Route utility<br>IP modules can be organized as independent tasks or be grouped together into a single task depending upon the target environment and user requirement. |
| IPX | Supported IPX LAN implementation protocol modules are the:<br>• Forwarding module<br>• Routing Information Protocol (RIP) module<br>• Service Advertisement Protocol (SAP) module<br>The IPX protocol<br>• validates the received packet by verifying checksum, length of packet, socket number, network address etc.<br>• determines whether the packet is destined to the router or is to be forwarded over LAN interfaces.<br>• maintains extensive statistics with respect to both LAN interfaces.<br>• generates IPX Ping Request packets and processes the incoming IPX Ping Request packets, sends back Ping Responses, and maintains useful statistics related to IPX Ping protocol operation. |
| Spanning Tree | Support for Spanning Tree Algorithm and Protocol (STAP) functionality, which computes single spanning tree of all nodes in the arbitrary bridged network. |
| Stacking | Support for multiple Ethernet switches to be interconnected and managed as if they were a single larger switching device. All features supported as it would for a stand-alone system. |
| LACP | Link Aggregation Control Protocol (LACP) allows one or more links to be aggregated together to form a link aggregation group, such that a MAC client can treat the link aggregation group as if it were a single high bandwidth link. |

# 10.0 Feature List

- Eight10/100/Gigabit port switching/routing in a single 713 pin PBGA chip
- Integrated 10/100/Gigabit Ethernet MACs
- On-chip storage for port transmit queues
- On-chip packet storage
- 10/100/Gigabit ports in full duplex mode
- Broadcast and multicast storm control with configurable per port settings
- Link aggregation on all ports in groups up to 4 ports
- Wire speed switching and routing on every port
- Hardware assisted address learning and aging

## 10.1 Interfaces

- GPCS (SERDES) and GMII interface for Gigabit ports
- PCI compliant CPU interface with bus mastering capability for packet and unresolved entry transfers to CPU. Also contains hooks to connect to a low cost I2C interface.
- SSRAM interface for address table sizes up to 40,000 entries (16,000 each for Layer 2, IP and 8,000 IPX entries) with no per port limits. Address table sizes can be as large as 40,000 entries in Layer 2 mode only.

## 10.2 Layer 3/Layer 4

- Packet by packet IP and IPX routing in hardware
- Layer 4 application level intelligence for IP switching and routing
- Layer 2 switching for protocols other than IP and IPX
- Automatic recognition of Ethernet Type II, 802.3, and SNAP packets
- 256 IP networks per device with no per port limit. The same IP network can span multiple ports
- IP nodes can belong to up to 256 multicast groups on any port

## 10.3 IEEE Standards

supported:

- 802.1D, 1998 edition which includes 802.1p Priority and Class of Service standard
- 802.1Q VLAN standard
- 802.3, 1998 edition which includes 802.3u Fast Ethernet standard, 802.3z Gigabit standard, and 802.3x Flow control standard
- 802.3ad standard

## 10.4 VLAN Configuration

based on:

- 802.1Q tags
- ports
- Multicast Ethernet address
- protocols

## 10.5 Filtering

based on:

- ports
- destination port
- source Ethernet, IP, or IPX address
- destination Ethernet, IP, or IPX address
- source-destination Ethernet, IP, or IPX address pairs
- end to end IP applications
- protocols

## 10.6 Class of Service

based on:

- 802.1Q tags
- receive port
- destination port
- source Ethernet, IP, or IPX address
- destination Ethernet, IP, or IPX address
- source-destination Ethernet, IP, or IPX address pairs

- end to end IP applications
- protocols

## 10.7 Port Mirroring

based on:

- receive port
- destination port
- source Ethernet, IP, or IPX address
- destination Ethernet, IP, or IPX address
- source-destination Ethernet, IP, or IPX address pairs
- end to end IP applications
- protocols

## 10.8 Bandwidth Management

based on:

- source Ethernet, IP, or IPX address
- destination Ethernet, IP, or IPX address
- source-destination Ethernet, IP, or IPX address pairs
- end to end IP applications

## 10.9 Queues

- Four queues for CPU packets; to allow prioritizing of packet types
- Four queues for unresolved packets to CPU; to allow prioritizing of unresolved types
- Four queues per port with user configurable priorities and weighted fair queuing

## 10.10 Counters

based on:

- EtherStats group of RMON
- proposed SMON standard
- MIB-II (RFC 1213), RFC 1573 and Ethernet interfaces MIB (RFC 1643)

## 10.11 Protocols

- Hardware support for BPDU, GVRP, GMRP, IGMP packets
- Hardware support for routing protocols such as RIP, IPX/RIP, IPX/SAP, OSPF, DVMRP

## 10.12 Standalone Operation

- No CPU required for Layer 2 remotely managed switch
- Supports Layer 2 address learning/aging in hardware
- Supports statistics collection/VLAN configuration from remote switch/workstation

***Note:*** In an unbalanced port configuration 1000 to 10, etc., a significant packet drop would be experienced due to packet buffer overflow.