



Intel® Media Switch IXE2426 10/100 + Gigabit L2 Device

Data Sheet

Advance Information

Product Features/Benefits

- Single-chip, 24-port 10/100 and 4-port Gigabit Ethernet Layer 2 switch
 - High integration, compact footprint, and low power dissipation enables the design of high-port density systems at the lowest per-port cost
- Wire-speed performance across all ports in switching mode
 - Delivers congestion-free performance through Enterprise switches during peak load periods
- Hardware assistance for several Layer 2 and protocols such as STP, Multiple Spanning Trees (802.1s), Rapid Reconfiguration (802.1w), Port-based Network Access Control (802.1x), GVRP, and GMRP packets
 - Reduces complexity and cost of CPU subsystem, which can be significant in stacks or chassis designs
- Link aggregation of all ports in groups up to 8 for 10/100 and 4 for Gigabit Ethernet ports
 - Enables meshed configurations with redundant paths for fail-safe networks
- Advanced traffic prioritization, QoS, and bandwidth management capabilities
 - Enables the convergence of voice, video, and data traffic of Ethernet/IP networks
- Fully compliant with VLAN implementation standards based on ports, tags, and addresses
- Enables flat plug-and-play networks that are easy to maintain
- Low-cost mode of operation
 - Enables a low-cost 24+4 Layer 2-only switch with no external address SSRAM
 - Provides support for up to 6K Layer 2 addresses and 64 802.1Q VLANs without external SSRAM
- Connections to other devices using standard interfaces such as SERDES/GMII, SMII, PCI, I²C, and SSRAM
 - Simplifies system design and provides flexibility when used with other devices
- Support for MAC-based and switch -based statistics gathering on chip
 - Enables effective network management using counters with SNMP

Note: Key features and benefits for the IXE2424 device are given above. Please refer to Section 10.0 for a complete feature list.

Notice: This document contains information on products in the sampling and initial production phases of development. The specifications are subject to change without notice. Verify with your local Intel sales office that you have the latest data sheet before finalizing a design.

August 2001

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel® products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel® products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Media Switch IXE2426 10/100+Gigabit L2 Device may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copyright © Intel Corporation, 2001

*Other brands and names are the property of their respective owners.

Contents

1.0	IXE2426 General Description.....	1
2.0	Applications.....	1
3.0	Functional Description.....	2
4.0	Interface Descriptions.....	4
5.0	Data Structures	5
6.0	Hardware Assisted Features	9
7.0	Electrical and Environmental Specifications.....	11
8.0	Hardware Support	14
9.0	Software Support.....	15
10.0	Feature List	18

Figures

1	IXE2426 Block Diagram	1
2	IXE2426 System Block Diagram	3
3	Switching Flow in IXE2426.....	5
4	Clock Interface Timing Diagram	13
5	IXE2426 Reference Design System.....	14
6	Software Architecture Illustrating APIs Supported by the IXE2426.....	15

Tables

1	IXE2426 Interfaces.....	4
2	Switching Block Entries	6
3	Absolute Maximum Ratings.....	11
4	Thermal Resistance	11
5	Thermal Resistance with Airflow	11
6	Operating Conditions.....	12
7	Clock Interface Timing.....	12
8	APIs Supported by the IXE2426.....	16
9	Protocols Enabled by the IXE2426.....	17



1.0 IXE2426 General Description

The IXE2426 device supports 24 10/100 Mbps ports and 4 Gigabit ports, integrating the MACs, switching, and queuing logic on-chip. The IXE2426 device is capable of wire-speed Layer 2 switching and wire-speed on all ports. It provides advanced filtering, mirroring, and prioritizing capabilities.

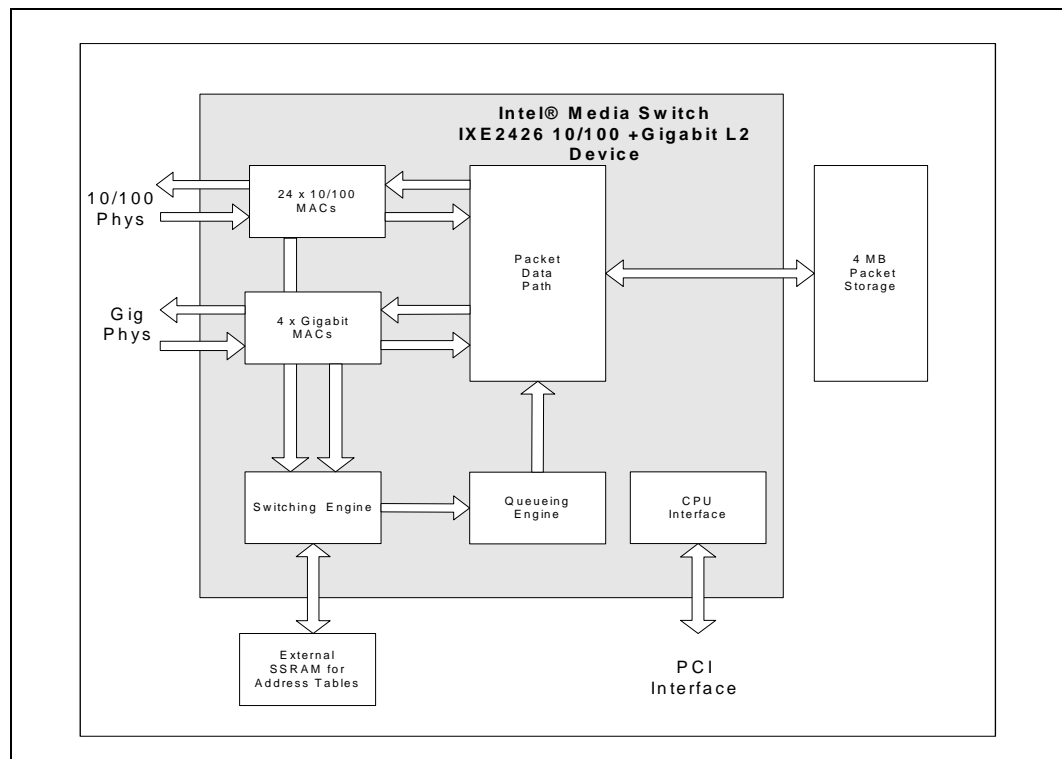
A system designed using a IXE2426 device requires transceivers on the 10/100 Mbps and Gigabit ports, memory for storing switching data structures, memory for storing packets, and a CPU subsystem. The IXE2426 device interfaces to the CPU subsystem via a 32-bit PCI bus.

2.0 Applications

Key applications for the IXE2426 include:

- 24 + 4 Layer 2 workgroup and enterprise switch with advanced bandwidth management
- Cascadable high port count Layer 2 switch when using one or all Gigabit ports for cascading

Figure 1. IXE2426 Block Diagram



3.0 Functional Description

3.1 Introduction

The IXE2426 device is a highly integrated Layer 2 switch. It supports 24 10/100 Mbps ports and 4 Gigabit ports with on-chip MACs. It also supports integrated switching logic and on-chip packet queueing memory. The IXE2426 is capable of switching packets at wire speed on all ports regardless of packet size.

IXE2426 supports:

- IEEE 802.1D, 1998 Edition standard
- IEEE 802.1Q VLAN standard

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

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

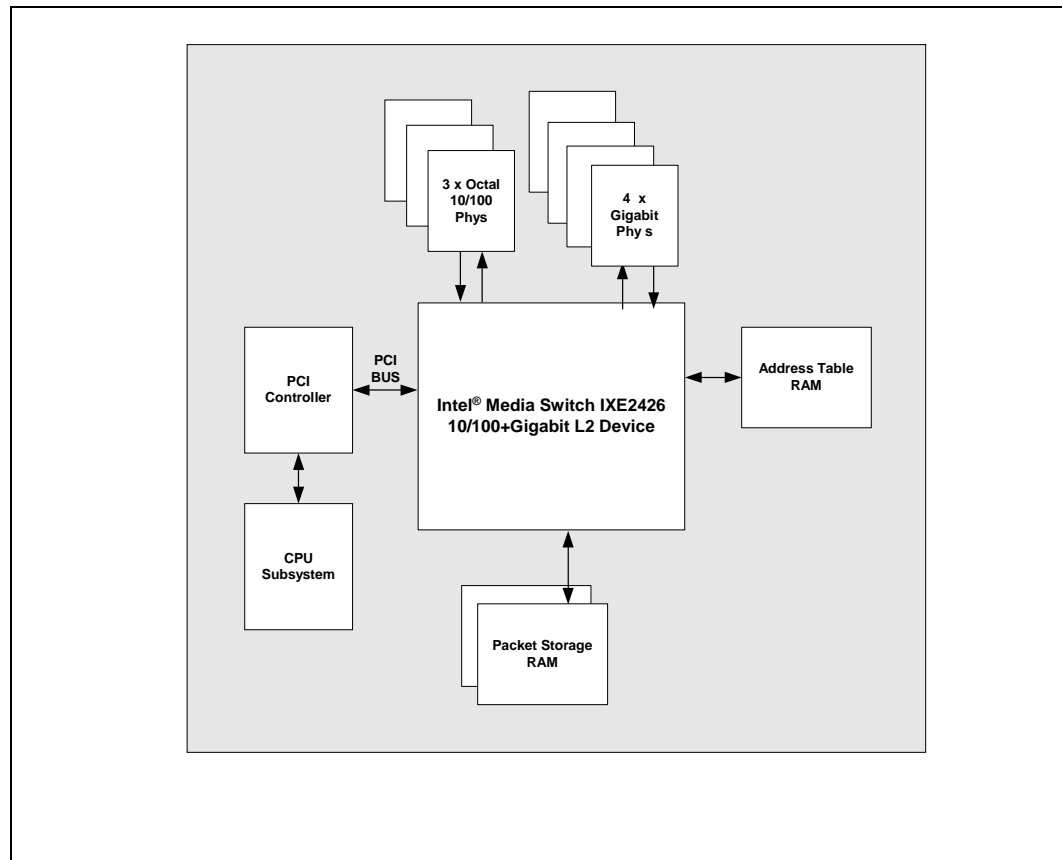
It provides on-chip 32-bit statistics counters to support the Etherstats group of RMON standards. In addition, it provides counters for monitoring flow-based statistics.

In addition to the IEEE standards support, the IXE2426 provides many advanced features such as Filtering, Mirroring, Bandwidth Management, Broadcast and Multicast Storm Control, all in a single device.

The IXE2426 interfaces to other devices by using only standards-based interfaces. This simplifies the system design and provides flexibility.

Figure 2 illustrates a system block diagram of the IXE2426 device.

Figure 2. IXE2426 System Block Diagram



4.0 Interface Descriptions

Table 1 provides a description of the IXE2426 interfaces.

Table 1. IXE2426 Interfaces

Interface	Description
PCI	PCI Rev 2.2 compliant interface with bus mastering capability for packet and unresolved entry transfers to CPU.
SMII	Two pin interface that allows the IXE2426 to communicate to the external 10/100 Mbps transceivers and provides three separate SYNC signals for point -to-point connections to three 8-port transceivers.
GMII/SERDES	Configurable interface used to link the IXE2426 to the four Gigabit ports.
Address Table SRAM	64-bit wide Pipelined Burst Synchronous SRAM interface for address storage.
Packet Storage	128-bit wide Pipelined Burst Synchronous Storage SRAM interface for packet storage.
MDIO	Serial MDIO interface for management of up to 32 Phy devices.
LED	Transfers the Gigabit ports' LED status. Can be attached to an external 74HC595-type shift register.

5.0 Data Structures

The Switching Engine block of the IXE2426 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).

Each type of “entry” contains different information that allows the IXE2426 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. Switching Flow in IXE2426

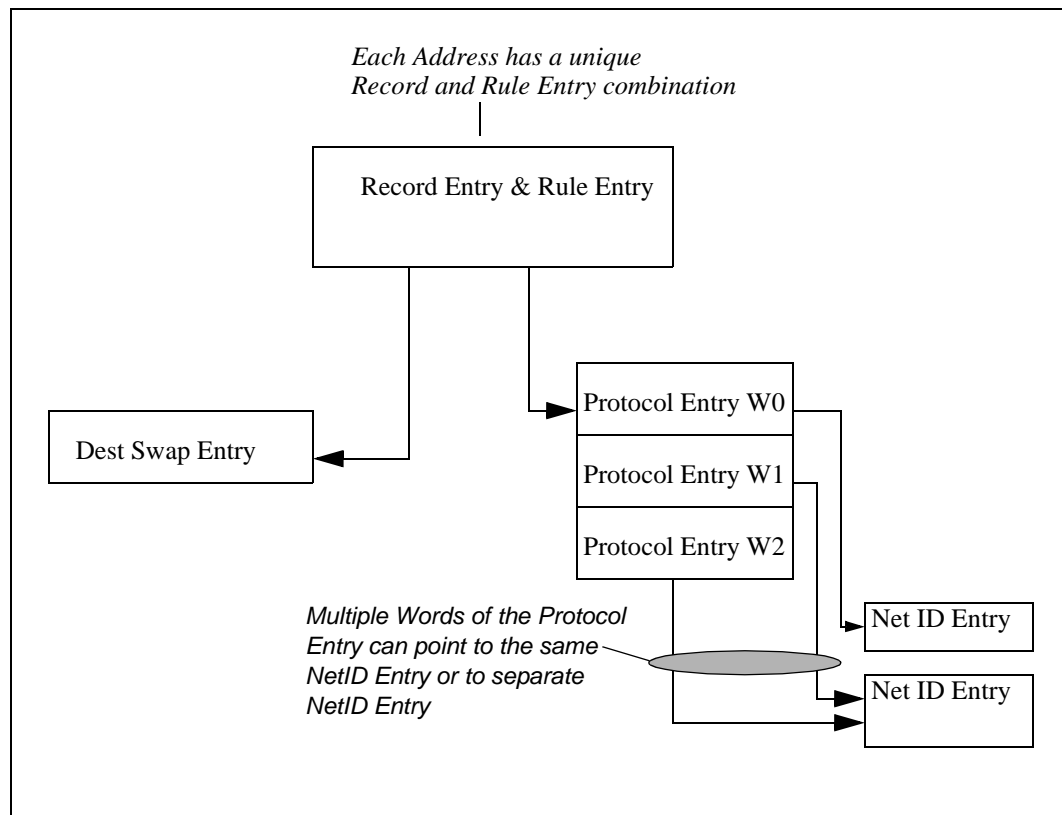


Table 2 entries are maintained by the Switching Engine block for Layer 2 addresses.

Table 2. Switching Block Entries

Entry Type	Description
Record Entry	The basic entry associated with every address. Is associated with another entry called the Rules Entry.
Rules Entry	Contains further information about the address including the port on which this address resides. Contains further pointers to a destination swap Entry and a Protocol Entry.
Destination Swap Entry	This entry contains the Destination Swap Address.
Protocol Entry	Contains filter, mirror, QoS, CoS, 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	This entry contains address-based VLAN information.

5.1 Enabled Layer 2 Protocols

The Layer 2 protocols enabled are: ARP, RARP, AppleTalk, DECNet, SNA, NetBios, DLC/LLC.

5.2 Port-based VLAN Features

The IXE2426 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.

5.3 Protocol-based VLAN Features

Protocol-based VLANs can be configured on the IXE2426 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.

5.4 802.1Q VLAN and 802.1D, 1998 Edition Priority and Class of Service Features

The IXE2426 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 Tag Priority by using the per port Port Regenerate Priority Entry. Map Priority Level register is a global mapping applying 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.

5.5 Port Aggregation Features

The IXE2426 supports port aggregation on all ports in groups of up to eight ports for the 10/100 Mbps ports, and in a group of four ports for the Gigabit ports.

The two modes of port aggregation supported are:

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

5.6 Interrupt Generation and Handling

The IXE2426 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.

5.7 Packet Transfers to CPU

The IXE2426 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. 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 IXE2426 to directly transfer data into the CPU memory without CPU intervention.

5.8 Unresolved Packet Transfers to CPU

The IXE2426 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 Scatter-Gather DMA engines provided per queue allow the IXE2426 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.

5.9 Packet Transfers From CPU

The IXE2426 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 section describes the hardware assisted features provided in the IXE2426.

6.1 Address Learning

The IXE2426 facilitates address learning by using a hardware engine that provides a content addressable memory-like interface to the CPU. The IXE2426 needs CPU support to learn addresses.

6.2 Address Aging

The IXE2426 contains hardware assistance for the entry aging process. It has three modes of operation to support various levels of CPU intervention. 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 device.

6.3 MDIO/MDC Scanning

The IXE2426 provides a controller to manage up to 32 devices via the MDIO and MDC pins. Up to thirty-two 16-bit registers per device can be addressed. The controller allows the CPU to use the PCI bus to communicate with these devices, performing all serialization and deserialization. In addition, it allows the CPU to perform “gang” operations that work on several devices at a time. Finally, it can be set to continuously scan the devices looking for a change in their status, interrupting the CPU only when it detects such a change. The controller uses an internal RAM for all gang operations.

6.4 Statistic Counters

The IXE2426 provides MAC-based and switch-based statistics gathering support on chip.

6.5 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 IXE2426 supports two types of algorithms per port: the Credit Algorithm and the Strict Priority Algorithm.

6.6 QoS

The QoS logic in the IXE2426 allows certain flows to be mapped as QoS. These flows can be queued to a certain priority that has been assigned a guaranteed bandwidth. The rate at which the flow is sending is also constantly monitored, and this is compared to the configured rate for this flow.

Multiple flows can be mapped to the same transmit queue, and still be guaranteed the required bandwidth for that flow—the bandwidth management feature takes care of guaranteeing the bandwidth for the queue, and the QoS feature takes care of guaranteeing the bandwidth for all the flows mapped to the same queue.

6.7 Broadcast and Multicast Storm Control

The IXE2426 provides a per port configuration on the incoming and/or outgoing port basis 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/bytes that are allowed in a given time interval.

7.0 Electrical and Environmental Specifications

This chapter contains the electrical and environmental specifications for the IXE2426. The IXE2426 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 described in Section 7.3, “Functional Operating Range” on page 11. Operating this product at the absolute maximum rating for a prolonged period can negatively impact device reliability. Table 3 lists the absolute maximum ratings for the IXE2426.

Table 3. Absolute Maximum Ratings

Parameter	Sym	Min	Max	Units
Supply Voltage Core	V _{DDCORE}	TBD	1.98	V
Supply Voltage IO	V _{DDIO}	TBD	3.6	V
Operating Temperature	T _{op}	-40	85	°C
Storage Temperature	T _{st}	-65	+150	°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

Table 4. Thermal Resistance

Thermal Resistance	Parameter	Value
θ_{JC}	Junction to case	TBD
θ_{CA}	Case to ambient	TBD

Table 5. Thermal Resistance with Airflow

Thermal Resistance	Package	0 LFPM	100 LFPM	200 LFPM	300 LFPM	500 LFPM
θ_{JA} (°C/W)	TBGA Type IA	TBD	TBD	TBD	TBD	TBD
Note: Heat Sink may be required depending on the air flow in the system.						

7.3 Functional Operating Range

The following table, Table 6 lists the functional operating range.

Table 6. Operating Conditions

Parameter	Sym	Min	Typ ¹	Max	Units
Recommended Supply Voltage Core	V _{DDCORE}	1.71	1.8	1.89	V ²
Recommended Supply Voltage IO	V _{DD}	3.135	3.3	3.456	V ²
Recommended Operating Temperature (IXE2426EA)	T _{op}	0		55	°C
Recommended Operating Temperature (IXE2426EE)	T _{op}	-40		85	°C
IDDAnalog			TBD	TBD	mA
IDDIO			TBD	300	mA
IDDCORE			TBD	1.5	A
Total Power (3.3V/1.8V)	P _{tot}			3.7	W
Maximum Junction Temperature	T _{jmax}	-40		125	°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 Clock Interface

7.4.1 Signal Timing

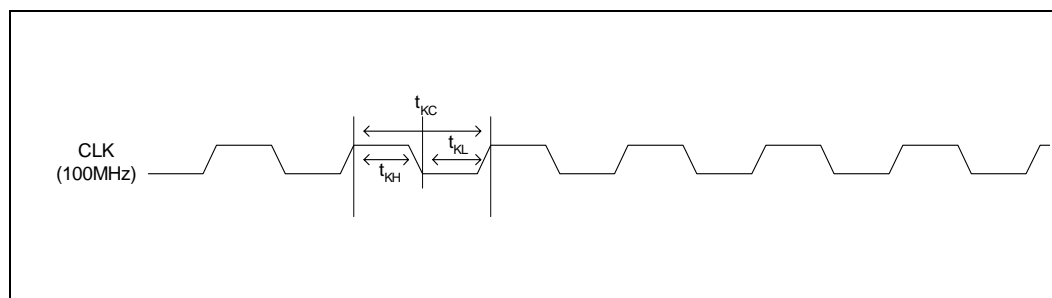
Table 7 describes the clock interface timing specifications.

Table 7. Clock Interface Timing

Symbol	Parameter	Min	Max
f _{FREQ}	Frequency		100MHz
	Duty Cycle	40%	80%
	Cycle-to-Cycle Jitter		200ps
t _{KH}	Time High	4ns	
t _{KL}	Time Low	4ns	
t _R	Rise Time		1.5ns
t _F	Fall Time		1.5ns

Figure 4 illustrates the Clock interface timing diagram.


Figure 4. Clock Interface Timing Diagram



8.0 Hardware Support

Figure 5 illustrates the hardware supported in the IXE2426 reference design.

Figure 5. IXE2426 Reference Design System

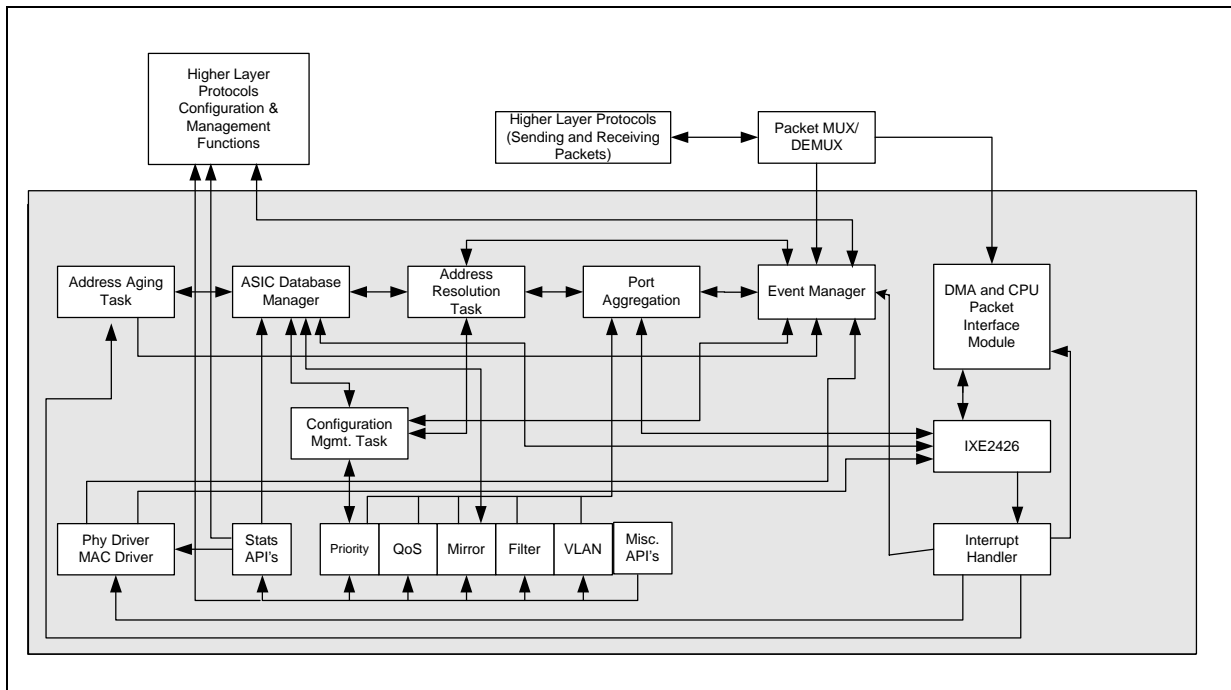


9.0 Software Support

9.1 APIs Supported

The IXE2426 supports the APIs illustrated in the software architecture block diagram shown in Figure 6.

Figure 6. Software Architecture Illustrating APIs Supported by the IXE2426



Summary descriptions and use of APIs supported by the IXE2426 are provided in Table 8.

Table 8. APIs Supported by the IXE2426

API Module	Description/Use
Interrupt Handling	Handles the interrupt from the IXE2426.
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 IXE2426 registers for various configurations. MAC and Phy provide functions to initialize the Ethernet controllers built into the IXE2426 and the LXT9782 Phy (such as functions to change the speed, functions to change the duplex mode, etc.).
Address Resolution	Contains functions used for learning Layer 2 addresses.
DMA Interface	Provides functions to send and receive packets between the IXE2426 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 IXE2426 tracks the address record entries that have been accessed over the ageing interval. Different aging time intervals can be specified for Layer 2 entries.
Address Learning	Provides functions for address learning in the software. The hardware provides a CAM interface to facilitate fast learning of addresses.
Filter, Mirror, Priority, 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 eight ports for the 10/100 Mbps ports, and in a group of two ports for the Gigabit 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 IXE2426 can perform extensive packet parsing and can obtain packet type (Type II, SNAP, etc.) and protocol (GARP, GVRP, STP, etc.) information directly from the packet header.

Enabled protocols are summarized in Table 9.

Table 9. Protocols Enabled by the IXE2426

Protocol	Description/Use
GxRP	<p>GxRP family protocols enabled 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.</p> <p>GxRP 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.</p>
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 Control	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	Allows one or more links to be aggregated to form a link aggregation group so that a MAC client can treat the link aggregation group as a single link.

10.0 Feature List

- 24 10/100 Mbps ports and 4 Gigabit port switching in a single 792 pin TBGA chip
- Integrated Gigabit and 10/100 Ethernet MACs
- On-chip storage for port transmit queues
- 10/100 Mbps ports configurable to half or full duplex
- Wire speed switching on every port
- Hardware assisted address learning and aging
- Link aggregation of all ports in groups up to 8 for 10/100 Mbps ports and 4 for Gigabit ports
- Scatter-Gather based DMA support
- Broadcast and multicast storm control with configurable per port settings
- 66 MHz PCI Interface

10.1 Interfaces

- SERDES & GMII interface for Gigabit ports and SMII interface for 10/100 Mbps ports
- 10/100 Mbps MACs for Gigabit Ports
- PCI Rev 2.2 compliant 66MHz CPU interface with bus mastering capability for packet and unresolved entry transfers to CPU
- SSRAM interface for address table sizes can be as large as 40K entries

10.2 IEEE Standards

- 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
- Proposed 802.3ad standard

10.3 VLAN Configuration

The VLAN configuration is based on:

- 802.1Q tags
- Ports
- Multicast Ethernet Addresses
- Protocols

10.4 Filtering

The filtering is based on:

- Ports
- Source Ethernet address
- Destination Ethernet address
- Source-destination Ethernet address pairs
- Protocols

10.5 Class of Service

The class of service is based on:

- Ports
- Source Ethernet address
- Destination Ethernet address
- Source-destination Ethernet address pairs
- Protocols

10.6 Port Mirroring

The port mirroring is based on:

- Ports
- Source Ethernet address
- Destination Ethernet address
- Source-destination Ethernet address pairs
- Protocols

10.7 Bandwidth Management

The bandwidth management is based on:

- Source Ethernet address
- Destination Ethernet address
- Source-destination Ethernet address pairs

10.8 Queues

- Two queues for packets from CPU; to allow prioritizing of packet types

- 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

10.9 Counters

The counters are based on:

- EtherStats group of RMON
- Proposed SMON standard
- RFC 1573 and Ethernet interfaces MIB (RFC 1643)

10.10 Protocols

- Hardware support for BPDU, GVRP, GMRP, IGMP, ICMP packets

10.11 Low Cost Mode Of Operation

- Supports building a low cost 24 + 4 Layer 2 only switch with no external address table SSRAM
- Up to 6K Layer 2 addresses and 64 802.1Q VLANs can be supported without external address table SSRAM