



Intel® Media Switch IXE2424 10/100+Gigabit L2/3/4 Advanced Device

Data Sheet

Advanced Information

Product Features/Benefits

- Single-chip, 24-port 10/100 and 4-port Gigabit Ethernet Layer 2/3/4 switch/router
 - 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 or IP/MPLS routing modes
 - Delivers congestion-free performance through Enterprise switches during peak load periods
- Hardware assistance for several Layer 2 and Layer 2/3/4 protocols such as STP, Multiple Spanning Trees (802.1s), Rapid Reconfiguration (802.1w), Port-based Network Access Control (802.1x), GVRP, 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, Diffserv, WRED 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
- Advanced multicast, broadcast, and filtering capabilities
 - Enables video and voice multicasting on IP networks
 - Protects from broadcast storms
 - Enables high-performance intranet firewalls
- Support for multiple IP networks on a single port, as well as multiple ports on the same network
 - Accommodates adds, moves, and changes in network topology
- Supports MPLS label edge router and label switch router configurations
 - Enhances and simplifies packet forwarding through routers using MPLS labels or forwarding decisions
- Connections to other devices using standards-based interfaces such as SerDes/GMII, SMII, PCI, 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
- IXE2424EE version supports extended temperature range of -40° to +85° C
 - Operates at the very high and low temperatures required for telecommunications applications

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.



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 IXE2424 10/100+Gigabit L2/3/4 Advanced 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	IXE2424 General Description.....	1
2.0	Applications.....	1
3.0	Functional Description.....	2
3.1	Introduction.....	2
4.0	Interface Descriptions.....	4
5.0	Data Structures.....	5
5.1	Enabled Layer 2 Protocols.....	6
5.2	IP Switching/Routing Features.....	6
5.3	Port Based VLAN Features.....	7
5.4	Protocol Based VLAN Features.....	7
5.5	802.1Q VLAN and 802.1D, 1998 Edition Priority and Class of Service Features.....	7
5.6	Port Aggregation Features.....	7
5.7	Interrupt Generation and Handling.....	8
5.8	Packet Transfers To CPU.....	8
5.9	Unresolved Packet Transfers To CPU.....	8
5.10	Packet Transfers From CPU.....	8
6.0	Hardware Assisted Features.....	9
6.1	Address Learning.....	9
6.2	Address Aging.....	9
6.3	Learning of Socket Address.....	9
6.4	Aging of Socket Address.....	9
6.5	MDIO/MDC Scanning.....	9
6.6	Statistic Counters.....	9
6.7	Bandwidth Management.....	10
6.8	QoS.....	10
6.9	DiffServ.....	10
6.10	WRED.....	10
6.11	Broadcast and Multicast Storm Control.....	10
6.12	IP Multicast Routing.....	10
7.0	Electrical and Environmental Specifications.....	12
7.1	Absolute Maximum Rating.....	12
7.2	Thermal Resistance.....	12
7.3	Functional Operating Range.....	13
7.4	Clock Interface.....	13
8.0	Hardware Support.....	15
9.0	Software Support.....	16
9.1	APIs Supported.....	16
9.2	Protocols.....	18



10.0 Feature List 19

10.1 Interfaces 19

10.2 Layer 3/Layer 4 19

10.3 MPLS Switching 19

10.4 IEEE Standards..... 20

10.5 VLAN Configuration 20

10.6 Filtering 20

10.7 Class of Service 20

10.8 Port Mirroring 21

10.9 Bandwidth Management 21

10.10 Queues..... 21

10.11 Counters..... 21

10.12 Protocols 22

10.13 Low Cost Mode Of Operation..... 22

Figures

1 IXE2424 Block Diagram 1

2 IXE2424 System Block Diagram 3

3 Switching Flow in IXE2424..... 5

4 Clock Interface Timing Diagram 14

5 IXE2424 Reference Design System..... 15

6 Software Architecture Illustrating APIs Supported by the IXE2424..... 16

Tables

1 IXE2424 Interfaces..... 4

2 Switching Block Entries..... 6

3 Absolute Maximum Ratings 12

4 Thermal Resistance 12

5 Operating Conditions 13

6 Clock Interface Timing 13

7 APIs Supported by the IXE2424 17

8 Protocols Enabled by the IXE2424 18

Revision History

Date of Revision	Version	Description
Aug '01	1.0	First release.
Sept '01	1.1	Changes reflect B0 stepping.



Date of Revision	Version	Description
Nov '01	1.2	The following features have been removed from the IXE2424 specification: -IPX -IPX/RIP -RIP -VoIP -IPSec



1.0 IXE2424 General Description

The IXE2424 device supports 24 10/100 Mbps ports and 4 Gigabit ports, integrating the MACs, switching, routing, and queuing logic on-chip. The IXE2424 device is capable of wire speed Layer 2 switching and wire speed Layer 3 and 4 routing on all ports. It provides advanced filtering, mirroring, and prioritizing capabilities and Layer 4 bandwidth management features. The device also supports MPLS switching on all ports.

A system designed using a IXE2424 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 IXE2424 device interfaces to the CPU subsystem via a 32-bit PCI bus.

The IXE2424 is available as two parts:

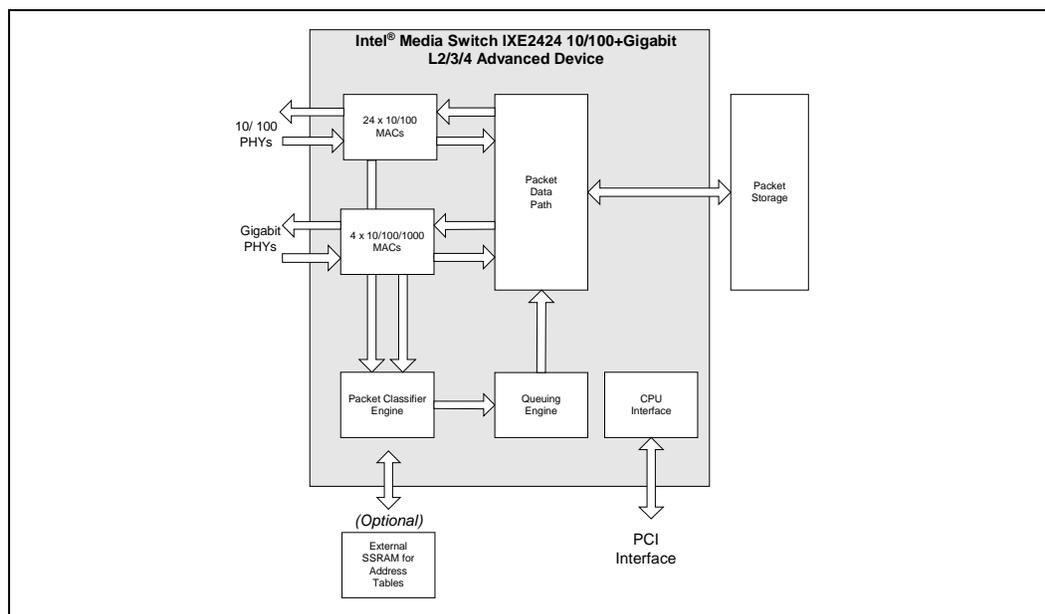
- IXE2424EC: Temperature range 0°C to 70°C
- IXE2424EE: Temperature range -40 °C to +85 °C

2.0 Applications

Key applications for the IXE2424 include:

- 24 + 4 Layer 2/3/4 workgroup and enterprise switches
- Layer 2/3/4 switch/router with Gigabit uplinks and advanced bandwidth management
- Cascadable high port count Layer 2/3/4 switch/router when using one or all Gigabit ports for cascading
- 24 + 4 MPLS label edge/switch router

Figure 1. IXE2424 Block Diagram



3.0 Functional Description

3.1 Introduction

The IXE2424 device is a highly integrated Layer 2 switch, MPLS label/edge switch/router and Layer 2/3/4 switch/router. It supports twenty-four 10/100 Mbps ports and four Gigabit ports with on-chip MACs. It also supports integrated switching and routing logic and on-chip packet queuing memory. The IXE2424 is capable of switching and routing packets at wire speed on all ports regardless of packet size.

In Layer 2 mode, the IXE2424 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
- 802.3u Fast Ethernet
- 802.3z Gigabit Ethernet standards
- 802.3ad Link Aggregation standards

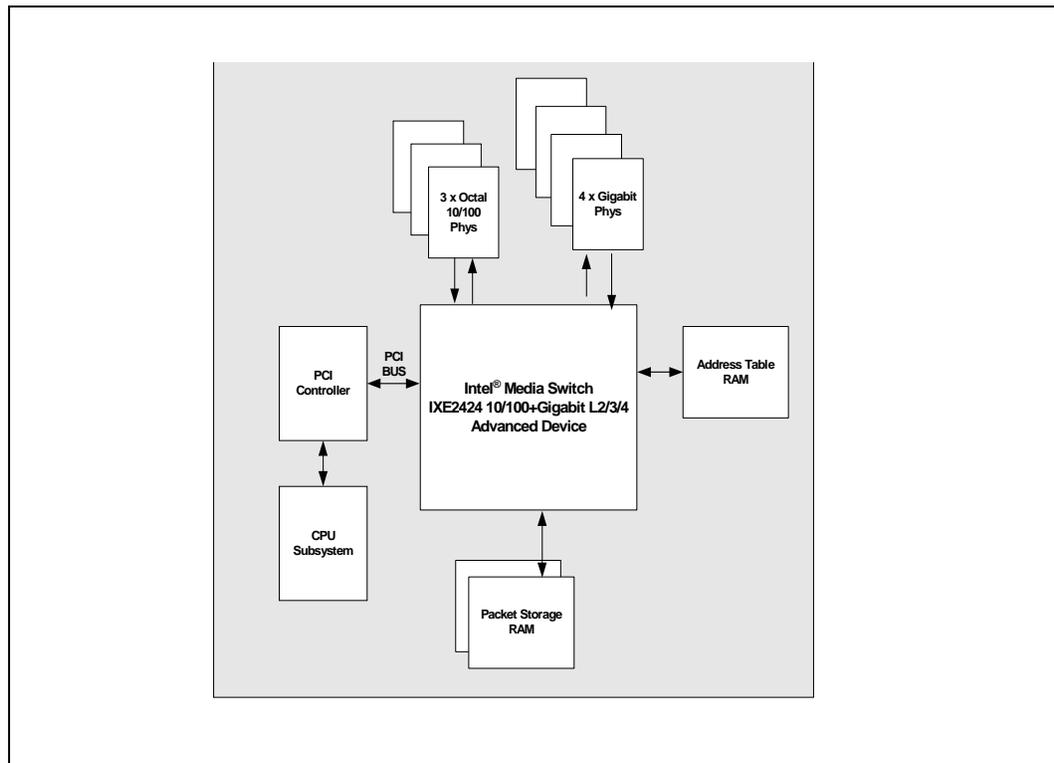
The IXE2424 provides on-chip 32-bit statistics counters to support the EtherStats group of RMON standards. In addition, it provides counters for monitoring flow-based statistics.

In addition to the IEEE standards support, the IXE2424 provides many advanced features such as filtering, mirroring, bandwidth management, and broadcast/multicast storm control, all in a single device.

The IXE2424 connects to other devices using standards based interfaces such as SerDes/GMII, SMII, PCI, and SSRAM. This simplifies the system design and provides flexibility when used with other devices.

Figure 2 illustrates a system block diagram of the IXE2424 device.

Figure 2. IXE2424 System Block Diagram



4.0 Interface Descriptions

Table 1 provides a description of the IXE2424 interfaces.

Table 1. IXE2424 Interfaces

Interface	Description
PCI	PCI Rev 2.2 compliant interface with bus mastering capability for packet and unresolved entry transfers to the CPU.
SMII	Two pin interface that allows the IXE2424 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 IXE2424 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 IXE2424 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 IXE2424 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 IXE2424

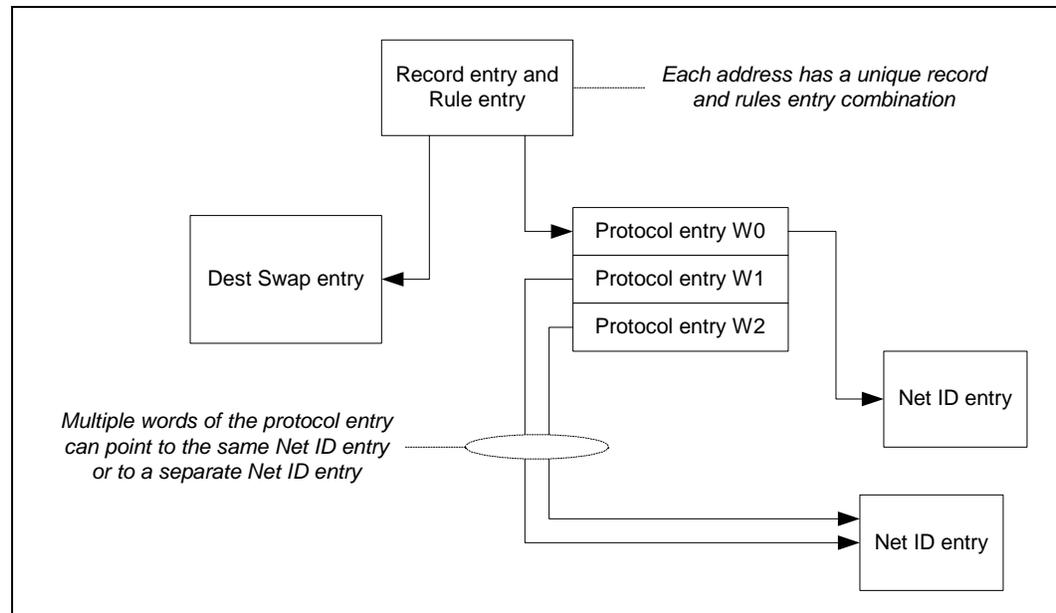


Table 2 entries are maintained by the switching engine block for Layer 2, IP and MPLS addresses.

Table 2. Switching Block Entries

Entry Type	Description
Record Entry	The basic entry associated with every address. 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. For IP and MPLS Multicast addresses, this entry contains pointers to multicast information entries.
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. For IP addresses, this entry contains information about any MPLS labels to be pushed into the packet.
Multicast Information Entry	Contains multicast packet replication information for IP and MPLS multicast addresses

5.1 Enabled Layer 2 Protocols

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

5.2 IP Switching/Routing Features

The IXE2424 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 and the Ethernet address of the next hop or destination address into the IXE2424 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 IXE2424 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 will be switched in the hardware at wire speeds using the Layer 2 switching algorithm.

5.2.1 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.

The IXE2424 device supports up to 256 routing domains.

5.2.2 Enabled IP Protocols

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

5.3 Port Based VLAN Features

The IXE2424 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

5.4 Protocol Based VLAN Features

Protocol based VLANs can be configured on the IXE2424 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.

5.5 802.1Q VLAN and 802.1D, 1998 Edition Priority and Class of Service Features

The IXE2424 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. The Map Priority Level register is a global mapping applied 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.

5.6 Port Aggregation Features

The IXE2424 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.

5.7 Interrupt Generation and Handling

The IXE2424 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.

5.8 Packet Transfers To CPU

The IXE2424 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 IXE2424 to directly transfer data into the CPU memory without CPU intervention.

5.9 Unresolved Packet Transfers To CPU

The IXE2424 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 IXE2424 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.

5.10 Packet Transfers From CPU

The IXE2424 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 IXE2424.

6.1 Address Learning

The IXE2424 facilitates address learning by using a hardware engine that provides a content addressable memory-like interface to the CPU. The IXE2424 needs CPU support to learn addresses.

6.2 Address Aging

The IXE2424 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 Learning of Socket Address

The IXE2424 provides a separate address learning interface for Layer 4 Socket addresses in order to speed up the process of application based rules generation. The IXE2424 also provides an add, delete, and lookup state machine called the Layer 4 processor. This simplifies the adding and deleting of socket addresses.

6.4 Aging of Socket Address

The IXE2424 has a separate aging state machine for aging Layer 4 Record entries.

6.5 MDIO/MDC Scanning

The IXE2424 provides a controller to manage up to 32 devices via the MDIO and MDC pins. Up to 32 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.

6.6 Statistic Counters

The IXE2424 provides MAC based and switch based statistics gathering support on-chip.

6.7 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 IXE2424 supports two types of algorithms per port: the Credit algorithm and the Strict Priority algorithm.

6.8 QoS

The QoS logic in the IXE2424 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. This is compared to the configured rate for this flow. If the rate is exceeded, then the packets that caused this exception are either dropped or remarked with new DiffServ code point and WRED precedence.

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.

6.9 DiffServ

The IXE2424 allows marking of packets belonging to certain flows. Packets belonging to such flows are assigned a particular value of the DiffServ code point. The DiffServ code point can also be changed based on QoS settings. Traffic prioritization and policing is done by mapping the DiffServ code point to internal priority levels, 802.1 p tag priority levels, and WRED parameter sets.

6.10 WRED

The IXE2424 supports two precedence classes on every priority queue. The WRED parameters are selectable in a per port per priority queue basis. WRED can be used along with QoS and DiffServ

6.11 Broadcast and Multicast Storm Control

The IXE2424 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.

6.12 IP Multicast Routing

The IXE2424 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 be sent depends on the number of routing domains on that port that have members that belong to that multicast group. The IXE2424 allows up to 256 multicast group members on a port.

7.0 Electrical and Environmental Specifications

This chapter contains the electrical and environmental specifications for the IXE2424. The IXE2424 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 13. Operating this product at the absolute maximum rating for a prolonged period can negatively impact device reliability. The following table, Table 3 lists the absolute maximum ratings for the IXE2424.

Table 3. Absolute Maximum Ratings

Parameter	Sym	Min	Max	Units
Supply Voltage Core	V_{DDCORE}	-0.3	1.98	V
Supply Voltage IO	V_{DDIO}	-0.3	3.6	V
Operating Temperature	T_{op}	0	70	°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

Package Type	Estimated Power W	Q_{JA} °C/W	Y_{JT} °C/W	$T_{case\ Max. (*)}$ °C
40mm-792TBGA Type II	3.5	9.6	1.0	121

(*) This is a not-to-exceed maximum allowable case temperature.

The thermal parameters defined above are based on empirical data of packages assembled on standard multi-layer 4s2p 1.0-oz. Cu. layer JEDEC test boards in a natural convection environment.

The maximum case temperature is based on the maximum junction temperature and defined by the relationship, $T_{case-max} = T_{jmax} - (Y_{JT} \times Power)$ where Y_{JT} is the junction-to-package top thermal characterization parameter. If the case temperature exceeds the specified $T_{case-max}$, thermal enhancements such as heat sinks or forced air will be required. Q_{JA} is the package junction-to-air thermal resistance.

7.3 Functional Operating Range

The following table, Table 5 lists the functional operating range.

Table 5. 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 (IXE2424EC)	T _{op}	0		70	°C
Recommended Operating Temperature (IXE2424EE)	T _{op}	-40		85	°C
IDDAnalog			6	15	mA
IDDIO			200	300	mA
IDDCORE			1.1	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

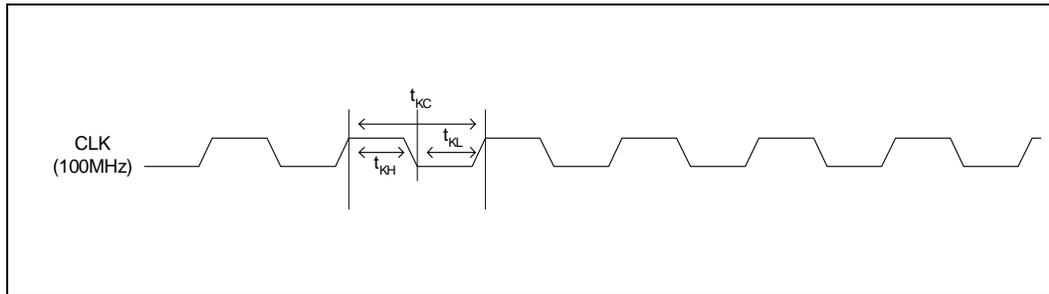
Table 6 describes the clock interface timing specifications.

Table 6. 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.

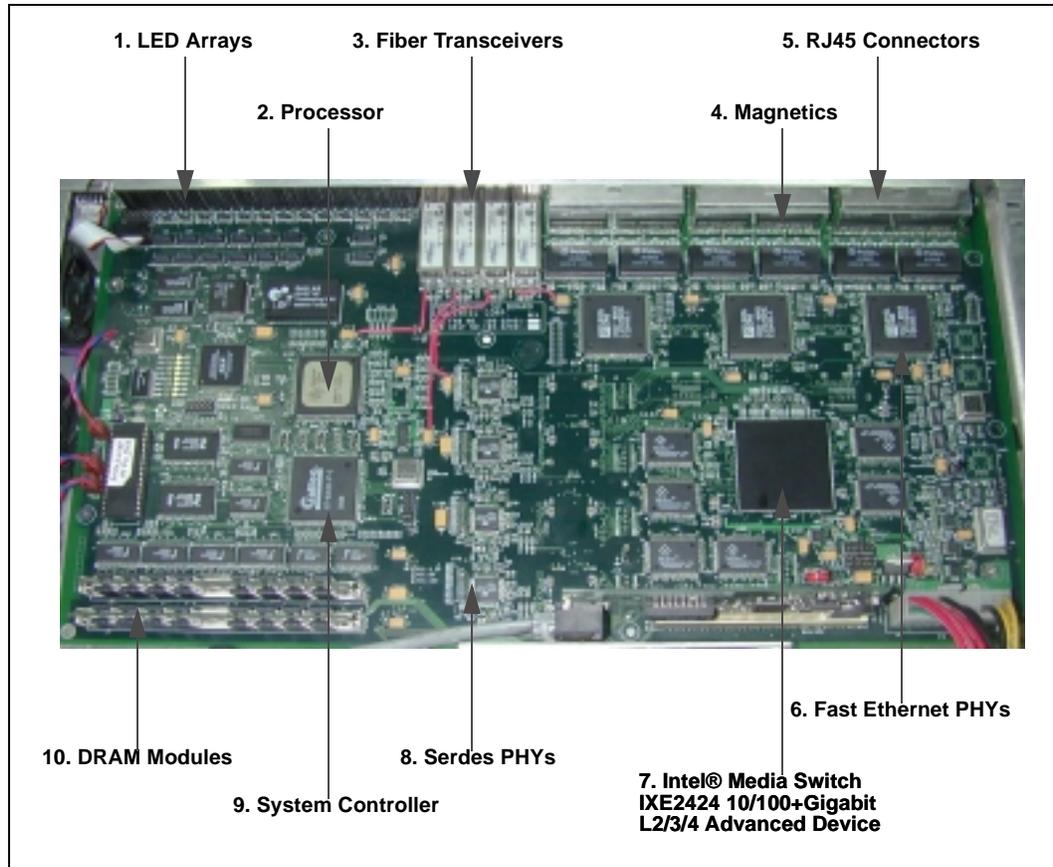
Figure 4. Clock Interface Timing Diagram



8.0 Hardware Support

Figure 5 illustrates the hardware supported in the IXE2424 reference design.

Figure 5. IXE2424 Reference Design System

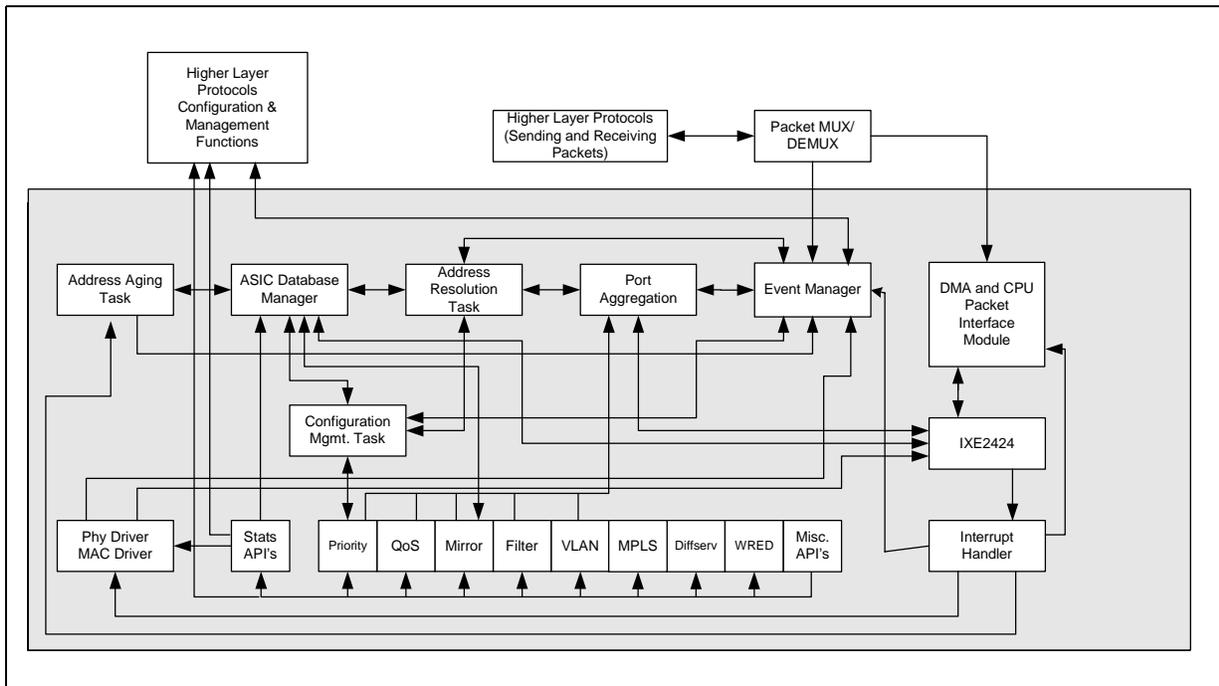


9.0 Software Support

9.1 APIs Supported

The IXE2424 supports the APIs illustrated in the software architecture block diagram shown in Figure 6.

Figure 6. Software Architecture Illustrating APIs Supported by the IXE2424





Summary descriptions and use of APIs supported by the IXE2424 are provided in Table 7.

Table 7. APIs Supported by the IXE2424 (Sheet 1 of 2)

API Module	Description/Use
Interrupt Handling	Handles the interrupt from the IXE2424.
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 IXE2424 registers for various configurations. MAC and PHY provide functions to initialize the Ethernet controllers built into the IXE2424 and theLXT9782 PHY (such as functions to change the speed, functions to change the duplex mode, etc.).
Address Resolution	Contains functions used for learning IP and Layer 2 addresses.
IP Configuration and Management	Includes the ability to perform a routing table lookup (for IP Address Resolution—calls into IP routing modules), determines the Ethernet address of the destination station or the next-hop (for routed packets whose destination is unresolved to determine the address programmed into the swap entry—calls into ARP module for IP), etc.
DMA Interface	Provides functions to send and receive packets between the IXE2424 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 IXE2424 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.
Filter, Mirror, Priority, DiffServ and Quality of Service	Supports configuring filters, mirrors, priorities, differentiated services, 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.

Table 7. APIs Supported by the IXE2424 (Sheet 2 of 2)

API Module	Description/Use
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 mode and Ingress and Egress Aggregation mode where Ingress Aggregation Only mode is the default.
MPLS	Used by label distribution protocols to set up the MPLS forwarding plane in the IXE2424.
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 IXE2424 can perform extensive packet parsing and can obtain packet type (Type II, SNAP, etc.) and protocol (IP, GARP, GVRP, STP, etc.) information directly from the packet header.

Enabled protocols are summarized in Table 8.

Table 8. Protocols Enabled by the IXE2424

Protocol	Description/Use
GxRP	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. 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.
IP	Enabled IP protocol modules are the: <ul style="list-style-type: none"> • Internet Protocol (IP) • Address Resolution Protocol (ARP) • Internet Control Message Protocol (ICMP) • Internet Group Management Protocol (IGMP) • Internet Control Message Protocol (ICMP) • User Datagram Protocol (UDP) IP modules can be organized as independent tasks or be grouped together into a single task depending upon the target environment and user requirement.
Spanning Tree	Spanning Tree Algorithm and Protocol (STAP) functionality, which computes single spanning tree of all nodes in the arbitrary bridged network.
Stacking Control	Enables 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 switch/router 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 and routing 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 and GMII interface for Gigabit ports and SMII interface for 10/100 Mbps ports
- 10/100/1000 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 up to 40K entries (16K each for Layer 2, IP, and 8K MPLS entries) with no per port limits. In Layer 2 Only mode, address table sizes can be as large as 40K entries.

10.2 Layer 3/Layer 4

- Packet by packet IP routing in hardware
- Layer 4 application level intelligence for IP switching and routing
- Layer 2 switching for protocols other than IP
- Automatic recognition of Ethernet Type II, 802.3, and SNAP packets
- 256 IP networks 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 MPLS Switching

- Supports MPLS switching on all ports
- Device can work as an Ingress/Egress MPLS/IP router in addition to being used as a pure MPLS forwarding node

- QoS, WRED, DiffServ are also supported for MPLS packets

10.4 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 802.1ad standard

10.5 VLAN Configuration

based on:

- 802.1Q tags
- Ports
- Addresses
- Protocols

10.6 Filtering

based on:

- Ports
- Source Ethernet, IP address
- Destination Ethernet, IP address
- Source-destination Ethernet, IP address pairs
- MPLS label
- End-to-end IP applications
- Protocols

10.7 Class of Service

based on:

- 802.1p tags
- Ports
- Source Ethernet, IP address
- Destination Ethernet, IP address
- Source-destination Ethernet, IP address pairs

- MPLS label
- End-to-end IP applications
- Protocols
- Support for Differential Services based on IETF RFC 2474

10.8 Port Mirroring

based on:

- Ports
- Source Ethernet, IP address
- Destination Ethernet, IP address
- Source-destination Ethernet, IP address pairs
- MPLS label
- End-to-end IP applications
- Protocols

10.9 Bandwidth Management

based on:

- Source Ethernet, IP address
- Destination Ethernet, IP address
- Source-destination Ethernet, IP address pairs
- MPLS label
- End-to-end IP applications

10.10 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
- WRED implemented for all queues

10.11 Counters

based on:

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

10.12 Protocols

- Hardware enabling for BPDU, GVRP, GMRP, IGMP, ICMP packets
- Hardware enabling for routing protocols such as OSPF

10.13 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