



Intel® Media Switch IXE5416 16-Port Gigabit Device

Data Sheet

Advance Information

Product Features/Benefits

- Single-chip, 16-port Gigabit switch/router with Layer 2/3/4 support
 - High integration, compact footprint, and low power dissipation enable the design of high port density systems at low per-port cost
- Integrated 10/100/1000 Ethernet MACs
 - Enables easy migration from 10/100 to Gigabit
- Wire-speed performance across all ports in switching or IP routing modes
 - Delivers congestion-free performance through enterprise switches during peak load periods
- Link aggregation
 - Enables meshed configurations with redundant paths for fail-safe networks
- Advanced traffic prioritization, QoS, and bandwidth management
 - Enables convergence of voice, video, and data traffic of Ethernet/IP networks
- Supports VLAN based on IEEE 802.1Q standards, ports, 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, and allows the building of high-performance intranet firewalls

Note: Key features and benefits for the IXE5416 device are given above. Please refer to Section 13.0 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 IXE5416 16-Port Gigabit 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	General Description.....	1
2.0	Applications	1
3.0	Interface Descriptions.....	2
4.0	MAC	3
5.0	Packet Decoding	3
6.0	Packet Forwarding	3
7.0	Packet Classification	4
8.0	Quality of Service	4
9.0	Packet Re-encapsulation	5
10.0	Electrical and Environmental Specifications.....	6
11.0	Hardware Support	7
12.0	Software Support.....	8
13.0	Feature List	11

Figures

1	IXE5416 Block Diagram	1
2	IXE5416 Reference Design.....	7
3	Software Architecture Illustrating APIs Supported by the IXE5416.....	8

Tables

1	Functional Operating Range	6
2	Supply Current and Power Dissipation.....	6
3	Thermal Resistance	6
4	APIs Supported by the IXE5416.....	9
5	Protocols Enabled by the IXE5416.....	10



1.0 General Description

The Intel® Media Switch IXE5416 16-Port Gigabit Device is a single-chip 16-port solution for layer 2 and layer 3 Gigabit Ethernet switching. The device uses a shared-memory architecture to achieve gigabit wire speed switching on all ports.

The highly integrated chip includes MACs, Address look-up CAM, Switch Engine, Primary Buffer Memory and programmable Quality of Service. The only external components required are the physical layer devices (Phys), and a management CPU. A secondary buffer memory can be used with the chip. The MACs interface to the physical layer through an IEEE compliant GMII/MII for 10/100/1000 Mbit/s operation or the TBI for 1000 Mbit/s operation.

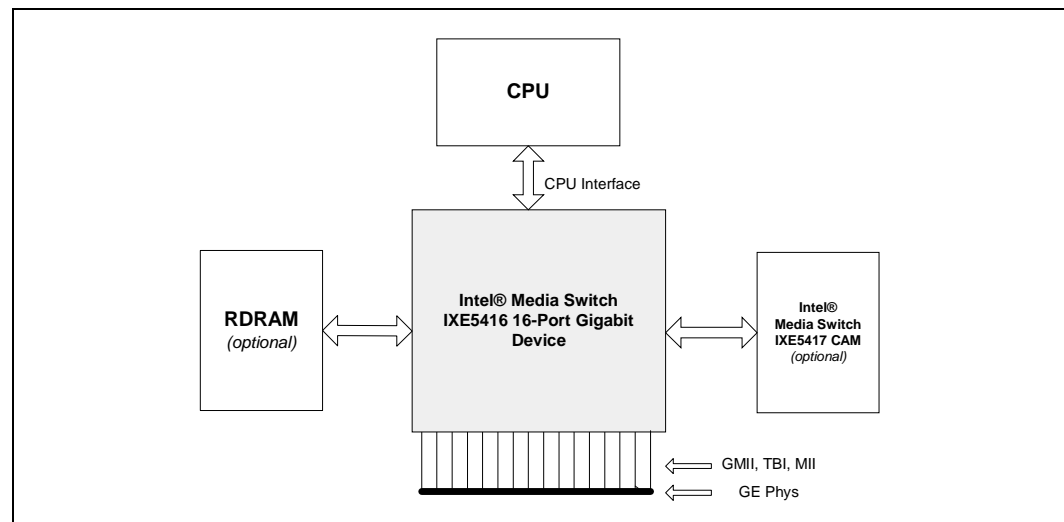
The IXE5416 supports layer 2 switching and layer 3 routing at wire-speed, along with support for VLANs, link aggregation, traffic prioritization and packet filtering based on layer 2 and layer 3/4 information.

The internal CAM holds all tables used in forwarding and VLAN tagging. Up to 8K entries are supported internally. For applications requiring larger address tables, there is an external interface for the Intel® Media Switch IXE5417 CAM (Content Addressable Memory). This memory provides storage for an additional 96K of layer 2 and layer 3 addresses.

2.0 Applications

- 16-24 port Gigabit Layer 2 workgroup switch
- 16-24 port Gigabit Layer 2/3/4 workgroup switch
- Daisy chain or Star stackable 10/100/1000 Layer 2/3/4 switch
- High port density Layer 2/3/4 switch/router solutions suitable for stackable and chassis switches. (Can be incorporated into stacking solutions with the Intel® Media Switch IXE2424 10/100+Gigabit L2/3/4 Advanced Device)

Figure 1. IXE5416 Block Diagram



3.0 Interface Descriptions

This section provides descriptions of the IXE5416 interfaces.

3.1 CAM Interface

Connects the IXE5417 CAM to the IXE5416 16-port Gigabit device. Cascaded IXE5417 CAMs (up to three) can be added in a tree structure. This gives up to 96K entries in external CAM.

3.2 LED Interface

Mainly targeted for TBI operation. For GMII operation, LED status is usually obtained from the Phy. The LED interface supports up to six LEDs per port.

3.3 CPU Interface

Used by an external processor to send control and configuration information to IXE5416 and to peripherals interfaced through IXE5416. Also used to read back information on the status and occurrence of certain events in IXE5416 and peripherals.

3.4 JTAG Interface

Supports the Test Access Port and Boundary Scan Architecture IEEE 1149.1 standards.

3.5 External Memory/Rambus Interface

The External Memory Controller (EMC) on the IXE5416 provides all functionality necessary for writing and reading packets to and from the off chip Direct RDRAM. It handles all address calculations and keeps track of every packet stored in external memory. It also hosts the Direct Rambus Memory Controller (RMC.d1) and the Direct Rambus ASIC Cell (RAC).

4.0 MAC

The IXE5416 provides built-in Ethernet MACs capable of 10/100/1000 Mbps operation on all ports. These MACs control the transmission and reception of data to and from the physical layer devices. The MACs support full and half duplex operation in 10/100M mode, and full duplex operation in 1000M mode. The MAC performs the following tasks:

- Interface to the physical layer device
- Frame synchronization
- Frame CRC check and generation
- Maximum and minimum packet length check
- Minimum Inter Frame Gap check
- Symmetric and asymmetric pause flow control
- RMON statistics
- SNMP statistics according to RFC 1213 (MIB-II)
- Allows status register access without interrupting packet transfer
- PCS functions (in TBI mode)
- Compliance with IEEE 802.3 standard.

5.0 Packet Decoding

The data flows through the receiving MAC to the Packet Decoder. The main task for this block is to decode the packet headers and to extract the fields (bytes) needed by the address look-up and classifier units such as L2 /L3 and L4 addresses, type of packet, error code information etc.

6.0 Packet Forwarding

The packet forwarding engine of the IXE5416 uses the packet header information extracted and decoded by the packet decoder. It processes this information and uses the result (list of destination port numbers, VLAN identifier etc.) to do:

- L2 Switching (bridging)
- L3 IP unicast and multicast routing

In addition, it supports:

- L2 address learning
- L2 address ageing
- 802.1Q VLANs, including port and protocol-based VLANs
- Spanning Tree
- Broadcast Storm Control

- 802.1p priority tags
- Port Mirroring
- Link Aggregation (802.3ad)

7.0 Packet Classification

The packet classifier looks into OSI layer 2, 3, and 4 in the packet data and classifies the packet according to user definable rules. The classification can be used to:

- Prioritize traffic by mapping packets to different traffic classes.
- Filter traffic by accepting, rejecting or denying packets on a packet by packet basis.
- Copy packets to the CPU.
- Copy packet headers to CPU for logging purposes.
- Count the number of packets that match each rule.

The classification block operates at wire-speed.

The classification unit classifies every packet passing through the device based on the L2 header, or the L3 and L4 headers. A set of rules is established to determine which packets to drop, which to switch/route, and which the processor should manage. Rules can also be set up to determine to which internal priority queue a packet should be assigned.

8.0 Quality of Service

The IXE5416 has a number of functions designed to provide Quality of Service support, making it possible to provide differentiated treatment of packets passing through the device. Put together, these functions can provide:

- Bandwidth guarantees
- Bandwidth limits
- Traffic priorities
- Scheduling
- Zero drop mode flow control

The QoS functions involve several IXE5416 blocks:

- Classifier - The results from the packet classifier can be used to prioritize a packet and to assign a traffic class to be used by the bandwidth distribution mechanism, the WFHBD.
- Prioritizer - The priority of a packet can be taken from different sources. The prioritizer determines which source to use for the internal queue priority and the outgoing 802.1p priority.
- WFHBD - The Weighted Fair Hashed Bandwidth Distributor schedules bandwidth on a port basis and a Traffic Class basis.
- Enqueuer - This block directs the flows between internal and external packet buffer memory during times of congestion.

- Scheduler - The IXE5416 has a strict priority scheduler operating on four queues per port.
- Flow control - The IXE5416 has an intelligent selective-pause with a guaranteed zero-drop algorithm. The IXE5416 can both send and receive pause frames.

9.0 Packet Re-encapsulation

The re-encapsulator restructures the packet for the output port.

Data is passed from the re-encapsulator to the MAC for transmission. The MAC adds the CRC and the L2 preamble, and then outputs the packet when the minimum IFG has expired since any previous activity on the media.

10.0 Electrical and Environmental Specifications

This section summarizes the electrical and environmental specifications for the IXE5416. The IXE5416 supports both 5V and 3.3V signaling environments.

10.1 Functional Operating Range

Table 1. Functional Operating Range

Parameter	Minimum	Maximum
Power supply, VDD1	3.0V	3.6V
Power supply, VDD2	2.3V	2.7V
Vref	1.4V (nom)	

10.2 Supply Current and Power Dissipation

Table 2. Supply Current and Power Dissipation

Power Supply	Maximum Current	Maximum Power
3.6V	1.25A	4.5W
2.7V	2.8A	7.5W

10.3 Thermal Considerations

IXE5416 is packaged in a thermally enhanced 42.5 mm EBGA package with an integrated copper heat spreader. IXE5416 is designed to operate with a junction temperature up to 100° C. The user must make sure that the temperature specification is not violated. The following thermal resistance parameters apply.

Table 3. Thermal Resistance


Symbol	Description	Value
θ_{jc}	Thermal resistance from junction to case	0.8° C/W
θ_{ja}	Thermal resistance from junction to ambient, zero air flow	9.8° C/W
θ_{ja-1}	Thermal resistance from junction to ambient, 1 m/s air flow	6.5° C/W
θ_{ja-2}	Thermal resistance from junction to ambient, 2 m/s air flow	5.8° C/W

For most cases some kind of cooling through a heat sink must be provided.

11.0 Hardware Support

Figure 2 illustrates the hardware supported in the IXC5416 reference design.

Figure 2. IXE5416 Reference Design

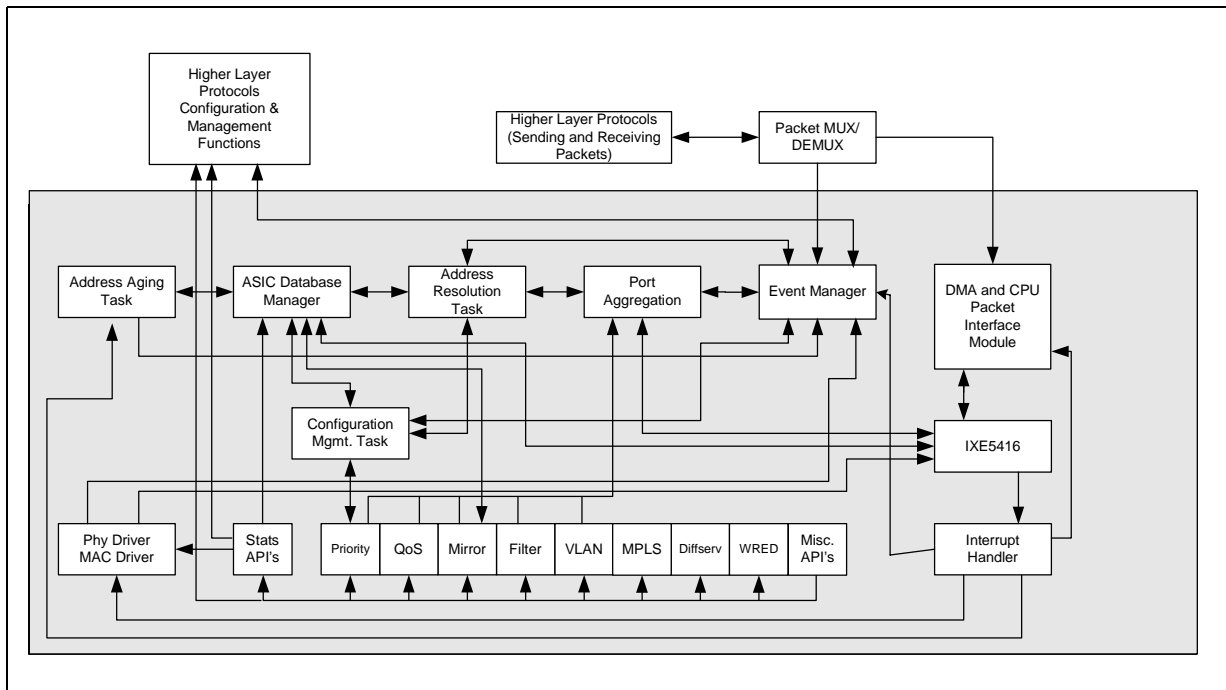


12.0 Software Support

12.1 APIs Supported

The IXE5416 supports the APIs illustrated in the software architecture block diagram shown in Figure 3.

Figure 3. Software Architecture Illustrating APIs Supported by the IXE5416



Summary descriptions and use of APIs supported by the IXE5416 are provided in Table 4.

Table 4. APIs Supported by the IXE5416 (Sheet 1 of 2)

API Module	Description/Use
Interrupt Handling	Handles the interrupt from the IXE5416.
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 IXE5416 registers for various configurations. MAC and Phy provide functions to initialize the Ethernet controllers built into the IXE5416 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 IP, IPX, and Layer 2 addresses.
IP 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 IXE5416 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 IXE5416 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 4. APIs Supported by the IXE5416 (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 and Ingress and Egress Aggregation modes of which Ingress Aggregation Only mode is the default.
MPLS	Used by Label Distribution Protocols to set up the MPLS forwarding plane in the IXE5416.
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.

12.2 Protocols

The IXE5416 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 5.

Table 5. Protocols Enabled by the IXE5416

Protocol	Description/Use
GxRP	Enabled GxRP family protocols 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) • Routing Information Protocol (RIPv1 and RIPv2) • 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.

13.0 Feature List

- 16-port 10/100/1000 Switch/Router, integrating MACs, CAM, packet buffer memory and Switching Engine
- Supports wire-speed L2 switching and L3 routing including L2 and IP multicast
- QoS provisioning on Layers 2/3/4 and 802.1p tag
- Flexible wire-speed packet classification and bandwidth management on L2 and L3/4
- Packet filtering support for 64 filters
- 128 Kbyte internal packet buffer
- External Direct RDRAM™ packet buffer, up to 512 Mbyte
- Integrated CAM with 8K addresses on chip
- Supports external CAM with up to 96K addresses
- Glueless CPU interface to PowerPC™ MPC8XX
- RTOS and CPU independent software device driver available (ANSI-C)
- Supports DMA slave capabilities for packet data
- Wire-speed MAC address learning on-chip
- VLAN support for 4K VLANs according to IEEE Std 802.1Q based on port and limited protocol
- SNMP, RMON and SMON statistics counters supported on-chip
- Link Aggregation for high-bandwidth interswitch links, supporting the IEEE 802.3ad standard
- Supports Spanning Tree, Broadcast Storm Control and Port Mirroring
- MII/GMII and TBI connections to external PHYs
- Full Duplex 1000 Mbit/s.
- Full and Half Duplex 10/100 Mbit/s
- IEEE Std 802.3x Pause Flow Control
- Thermally enhanced 836-EBGA package
- 3.3V I/Os, 5V tolerant

