S/UNI-2XGE

S/UNI™ - 2xGE

PM3386

TECHNICAL OVERVIEW

PRELIMINARY INFORMATION

ISSUE 1: OCTOBER 2000

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 PURPOSE AND SCOPE OF THIS DOCUMENT

PMC-Sierra has introduced the Industry's first Dual Gigabit Ethernet Controller with embedded serializer/deserializer (SERDES) to satisfy requirements of key network applications:

- Optical Switches and cross connects
- Multi-service switches
- Core and Edge Routers
- Ethernet over SONET Uplinks

This paper discusses how and why PMC-Sierra implemented a system level solution to satisfy the unique requirements of this marketplace. We first analyze the requirements driving the adoption of Gigabit Ethernet beyond the Enterprise market and describe how the features of the S/UNI-2xGE meet these requirements.

2 S/UNI-2XGE OVERVIEW

The S/UNI-2xGE is optimized for use in carrier grade, full-featured equipment that requires high-density Gigabit Ethernet ports. As shown in the following table, its feature set has been carefully designed to increase port density and reduce power for Gigabit Ethernet applications, while simplifying the implementation of Gigabit Ethernet ports on a wide variety of equipment such as Core Routers, Edge Routers, Multi-service switches, Optical Switches, Wireless base station router interconnect and access boxes.

Table 2-1 S/UNI-2xGE Features and Benefits

SUNI 2xGE Feature	System Benefit
Integrated Dual Gigabit Ethernet SERDES.	Increases port density while lowering power requirements. Allows direct connection to optical modules from the Dual Gigabit Ethernet Controller. To perform similar functions with discrete devices requires between 4-6 parts, which are now replaced with the S/UNI-2xGE. A single chip solution reduces manufacturing costs, simplifies design and testing, and increases reliability.
POS-PHY Level 3 system interface.	Typical Gigabit MAC devices have either a generic system interface or a proprietary bus. The POS-PHY Level 3 system interface has been standardized by the Optical Interworking Forum and the ATM Forum. POS-PHY Level 3 is a multi-service interface that is capable of handling Ethernet, Packet Over SONET, or ATM. This allows a wide set of physical layer devices to be developed with the same system interface. The benefit to system vendors is re-use of the back end of their line card design. Service cards implementing various physical layer protocols are easily integrated into a common architecture based on POS-PHY Level 3. PMC-Sierra, Inc. POS-PHY Level 3 devices include support for • OC-48c POS/ATM • Quad OC-12c POS/ATM

SUNI 2xGE Feature	System Benefit
	• Channelized OC-48 to STS-1 with sub rate capability. And now, • Gigabit Ethernet via the S/UNI-2xGE.
GMII Interface for Gigabit over copper support via an external Copper PHY.	Most high end Gigabit Ethernet applications use fibre optic cable rather than twisted pair cable. However, some applications may require copper twisted pair support. The addition of two GMII interfaces in the S/UNI-2xGE allows it to be connected to copper Gigabit Ethernet physical layer devices while leveraging the S/UNI-2xGE Gigabit MAC features and the POS-PHY Level 3 system interface.
Large system buffers to provide loss-less flow control operation.	The S/UNI-2xGE is targeted for high end Gigabit Ethernet applications. Many applications use external FIFO devices due to the lack of buffers in many Gigabit Ethernet MAC devices. The S/UNI-2xGE provides large system buffers to avoid extra external FIFO devices. The buffers have been sized to provide loss-less flow control for jumbo frames (9.6k bytes) up to a link distance of 5 Km. When one link must flow control the far end, there is a delay in halting the traffic due to propagation delay time, the reaction time of the far end GMAC (it cannot interrupt a packet transmission once it has started), and a return propagation time. If the GMAC that initiated flow control does not have sufficient buffer memory, packets will be dropped. The S/UNI-2xGE provides 16K byte transmit and 64K byte receive buffers for each channel.
Jumbo Frame support for frames up to 9.6K bytes.	Most Gigabit Ethernet controllers do not support jumbo frames. Between Servers and Routers, the use of Jumbo frames is desirable to increase data throughput by reducing the overhead. More data is transferred with the same fixed Ethernet overhead. While there is no standard for the

SUNI 2xGE Feature	System Benefit
	size of a jumbo frame, a survey of the industry showed that most applications plan to use a maximum of 9.6K byte frames.
POS-PHY Level 3 FPGA cores support.	POS-PHY Level 3 FPGA cores are being developed by a number of leading FPGA vendors. This simplifies the implementation of functions that are done in a device sitting behind the S/UNI-2xGE that would otherwise require an ASIC or the development of POS-PHY Level 3 interface code for an FPGA. Standard cores for POS-PHY Level 3 significantly improves time to market. PMC-Sierra, Inc. current partners for POS-PHY Level 3 cores are Xilinx and Altera. For contact names, please contact your local PMC Sales Representative.

3 ETHERNET – THE DOMINANT LAYER 2 PROTOCOL

Over 90% of enterprise links use some flavor of Ethernet. These links consist of 10 Mbit/s, 10/100 Mbit/s, Gigabit Ethernet, and soon 10 Gigabit Ethernet, which makes Ethernet the dominant Layer 2 protocol. As well, Ethernet is working its way into new applications throughout the entire network.

- Connections between POP Servers and Routers
- Gigabit Ethernet POP router interconnect
- Gigabit Ethernet ports for Optical cross-connects
- Ethernet Metro applications
- Gigabit Ethernet Access
- Ethernet over SONET
- Backplane or inter-shelf connections
- Wireless IP Router interconnects
- Ethernet over lambda's

This partial list of applications is driven by the economies of scale from leveraging Ethernet's existing market share within the Enterprise. In addition, protocol conversions are expensive within the network, requiring processing, management, and buffering (may be multiple devices). Reducing the number of protocol conversions within a network will lower the cost and simplify network management. Since the majority of traffic generated on the network is to and from Enterprise networks, Ethernet is the logical protocol choice. Ethernet is now being deployed in the POP, Metro, and Transport markets.

3.1 Gigabit Ethernet Market Drivers

- Router interconnect
- Data Centers / Web Hosting
- Gigabit Ethernet Metro
- Gigabit Ethernet Access

3.2 Gigabit Ethernet in Internet POPs

There are a number of applications that are driving the adoption of Gigabit Ethernet within Internet POPs.

The key requirements for equipment within an Internet Point of Presence (POP) are the availability of high-speed ports, high port density, and low power. Typically, Internet Service Providers (ISPs) work with limited rack or cage space with restrictions on size, power, and cooling. Thus, card port density is very important in determining the revenue a particular box will provide.

The number of Gigabit Ethernet ports being deployed in Internet Points of Presence POPs has increased significantly. This is due to the use of Gigabit Ethernet for router-to-router connections and router-to-server connections.

3.2.1 Router-to-Router Connections

For redundancy, dual homing techniques are used to ensure that every router has at least two connections to other routers within the POP. This requirement limits the number of routers within a POP due to the limited number of ports per router. To solve this problem switches are used to connect routers in a POP. Each router can connect to two switches to provide dual homing, while using a minimum number of ports for router interconnect. ATM switches were used initially, however, due to lower cost, Gigabit Ethernet switches are becoming the standard for router interconnect. This has lead to a significant increase in the number of Gigabit Ethernet ports deployed in POPs.

The Media Access Control (MAC) functions are not used as a differentiator when adding GE to a router. Higher-level features such as classification, port density, and protocol support are used to differentiate. Thus, to achieve higher densities the integration of the GMAC and the other components, such as the FIFO and SERDES are required.

3.2.2 Router-to-Server Connections

Web hosting is becoming very popular in the new POP architectures. Companies who rely on the Internet for their business success do not want their web site speed determined by slow “last mile” access links to their offices. To avoid the slow data rates of access or avoid expensive high speed access, companies web sites are hosted in server farms located in Internet POPs or high speed Data Centers. The explosion of e-commerce applications is driving substantial development of Data and Web hosting Internet centers.

Typical connections between POP routers and servers are 100 Mbit/s Ethernet or Gigabit Ethernet that is fed to a 10/100 and Gigabit Ethernet switch. In both cases the majority of server connections are operating at 100 Mbit/s. As e-commerce grows these links are beginning to be upgraded to Gigabit Ethernet. This provides a significant opportunity for a large scale deployment of Gigabit Ethernet in Router to Server applications. Such deployment depends on port density; hence the drive to integrate the SERDES and Gigabit MAC in a single device.

4 S/UNI-2XGE TECHNICAL OVERVIEW

This section provides a detailed technical overview of the S/UNI-2xGE device. For a more complete description of the device please refer to the datasheet PMC-1001129 – Issue 4.

The S/UNI-2xGE is a dual Gigabit Ethernet controller with integrated SERDES, FIFO's, and a POS-PHY Level 3 system interface. Figure 4.1 shows a comparison of the size of the S/UNI-2xGE to a common discrete implementation. Shows a comparison of the power for each solution with the S/UNI-2xGE providing up to a 60% power savings.

Figure 4.1 Common Discrete Device Gigabit Ethernet Implementation

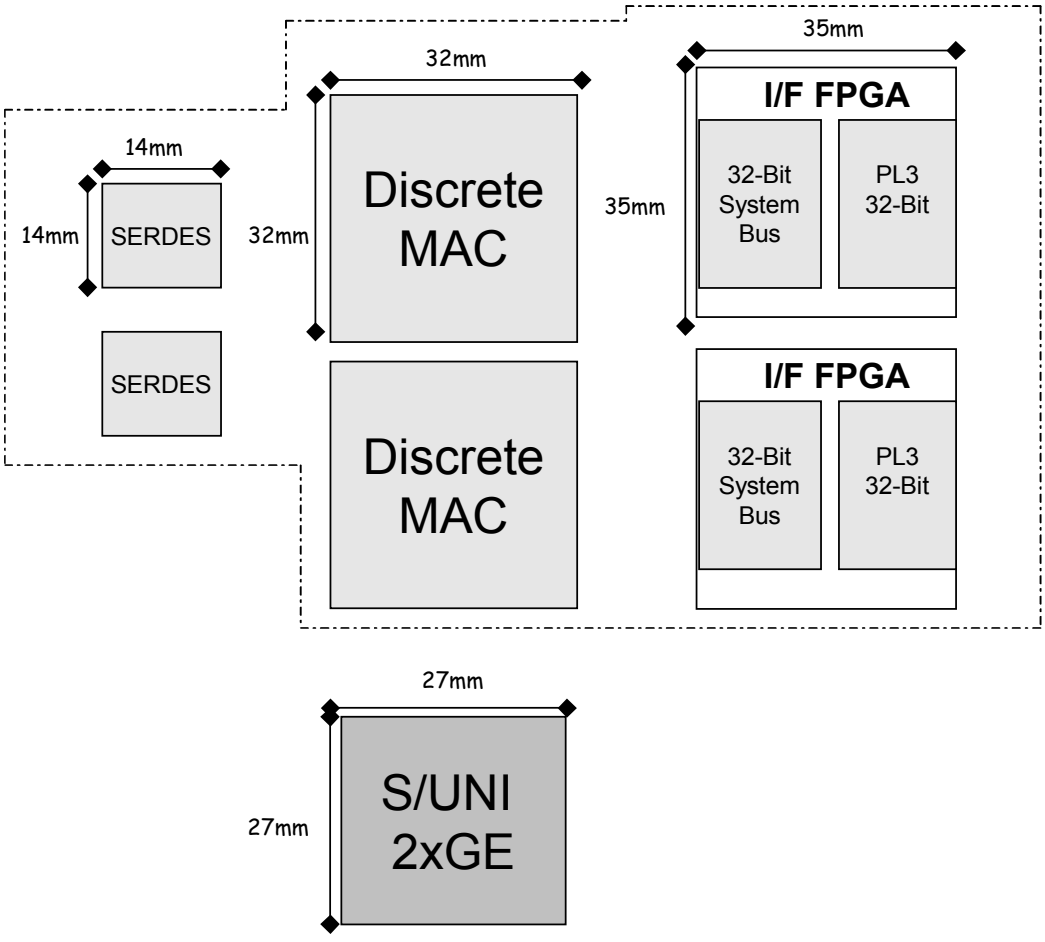

Table 4-1 S/UNI-2xGE versus Discrete Implementation Power Comparison

Description	*Components*	*Power*
Discrete Implementation	SERDES (x2)	~1.8 W
	MAC (x2)	~4.0 W
	FPGA (x2)	~3.0 W
	TOTAL	**~8.8 W**
S/UNI-2xGE	SERDES (x2)	
	MAC (x2)	
	PL3 interface + FIFO's	
	TOTAL	**<2 W/port**

4.1 Serializer/Deserializer

The S/UNI-2xGE incorporates a dual Gigabit Ethernet Serializer/deserializer (SERDES) to increase port density and lower the system power. The S/UNI-2xGE SERDES allows direct connection to optics modules and has high standards for jitter performance. Operation on each Gigabit Ethernet channel is programmable and independent, allowing for both fibre and copper support.

Figure 4.2 shows a diagram of the 2xGE with direct connection to optics on one channel and a connection to a copper PHY on the other channel. This diagram also shows the benefit of the PL3 system interface simplifying the development of multiple cards based on the same system interface. Dual Gigabit Ethernet, OC-48c Packet Over SONET (POS) and Quad OC-12c POS cards can be implemented with significant reuse of higher layer devices.

Figure 4.2 S/UNI-2xGE Application Example with other PL3 Devices GMII Interface

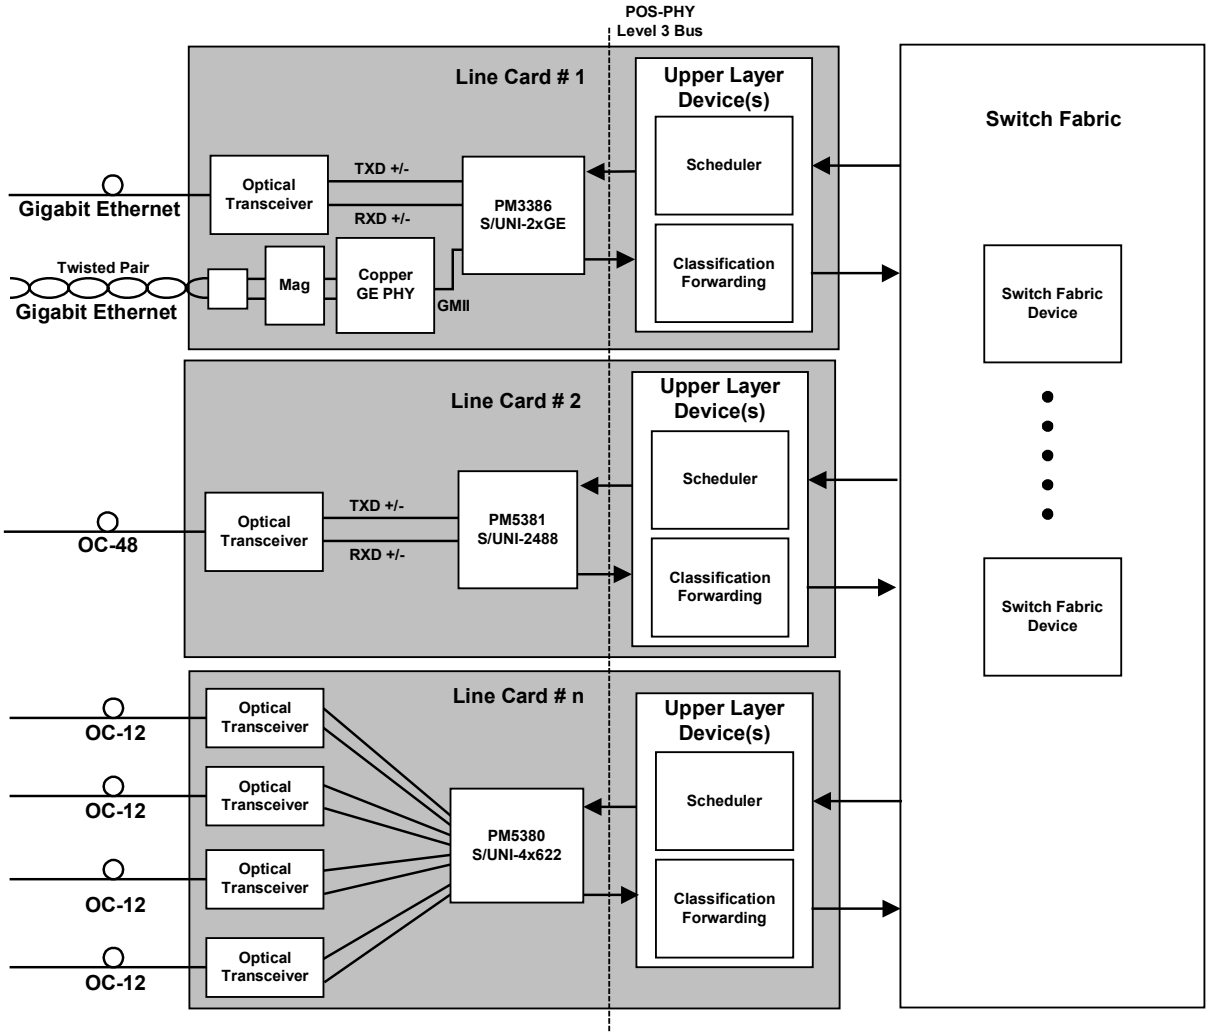

The S/UNI-2xGE has provided two GMII interfaces in addition to the SERDES interfaces. Although fibre implementations are expected to be the dominant application, Gigabit Ethernet over copper cannot be ignored. System vendors may be expected to provide cards that support Gigabit Ethernet over copper in addition to fibre based cards. Figure 4.2 shows an example of the S/UNI-2xGE connected to a copper Gigabit Ethernet PHY.

Another application for the GMII interface is to simplify implementation of Packet Over SONET interfaces on Gigabit Ethernet switches. The GMII interfaces of the S/UNI-2xGE are connected to existing GMII ports on a Gigabit Ethernet switch to provide access to the POS-PHY Level 3 interface. An FPGA is used to connect the POS-PHY Level 3 interface on the S/UNI-2xGE to the S/UNI 2488 device. The S/UNI 2488 provides the OC-48c POS functionality via a POS-PHY Level 3 system interface.

4.3 GMAC

The Gigabit Ethernet MAC performs basic frame checks, address filtering both multicast and unicast, flow control, jumbo frames up to 9.6K, statistics, and autonegotiation. The GMAC supports Ethernet 2.0, IEEE 802.3 LLC and IEEE 802.3 SNAP/LLC encoding formats and VLAN Tags. Loopbacks are also provided for diagnostic capability through the GMAC.

The GMAC also has both unicast and multicast filtering. Eight exact match filters are provided per channel for unicast packets. Filtering is performed on either the Destination Address (DA), the Source Address (SA), or the VLAN Tag ID (VID) + DA or VID+SA. Each filter can be programmed to either Accept or Discard the frame. These features allow significant design flexibility to either accept or reject specific address. Source Address filtering allows designers to implement security within a router port. For example, packets are only accepted from customers with a valid Service Level Agreement. In addition, certain source addresses can be filtered where no Service Level Agreement exists.

4.4 Statistics

The S/UNI-2xGE provides an extensive set of statistics counters to support RMON and SNMP applications. The counters are a minimum of 40 bits, providing a minimum roll-over of 58 minutes. This is done to provide a sufficient amount of time for the management software to retrieve and process these statistics.

A complete listing of the registers supported is provided in the S/UNI-2xGE datasheet (PMC-1991129).

4.5 WAN Sized FIFOs

The S/UNI-2xGE incorporates a significant amount of FIFO memory to accommodate increased distances to be supported in WAN applications. The S/UNI-2xGE provides 64K bytes per port on the ingress Gigabit Ethernet side and 16k bytes per port for the egress direction. This amount of buffering is sufficient on the ingress side to accommodate loss-less flow control for jumbo frames (9.6k bytes) on links up to 5Km. The egress direction can store a complete jumbo frame with room to spare. This buffer is sized to ensure that underflow does not occur on the line side.

4.6 Flexible Flow Control

The S/UNI-2xGE provides significant flexibility when implementing flow control. There are many different methods for flow control operation depending on the application requirements. The first method is to use the programmable watermarks within the S/UNI-2xGE to determine the upper and lower bounds for pause frame generation. This method allows the link layer device to de-assert the RENB, which will cause the S/UNI-2xGE FIFOs to fill and automatically send a pause frame when the high watermark is hit. The second method for flow control is for the higher level device to generate a pause frame and send it to the S/UNI-2xGE for transmission. The third method uses external side band flow control pins called PAUSE and PAUSED for each channel. A higher layer device can toggle

the PAUSE pin and cause the S/UNI-2xGE port to transmit a pause frame. When the pin is released a zero timed pause frame will be sent to remove the pause condition at the far end link.

Pause frames that are received at the S/UNI-2xGE can be terminated at the GMAC or optionally passed to the higher layer device. When the pause frames are terminated by the GMAC, the PAUSED pin will be held high until the pause count down timer reaches zero. This provides indication to higher layer devices that a pause frame has been received.

The fourth method is to set register bits via the microprocessor interface to cause the S/UNI-2xGE to send a pause frame, or to monitor registers that indicate the reception of pause frames. Multiple different methods for implementing flow control provides the designer with significant flexibility to determine the methods that best suits the particular application.

4.7 POS-PHY Level 3 System Interface

The POS-PHY Level 3 system interface is a standard 32 bit interface running at 104 MHz. It can support a variety of data protocols including ATM, Packet over SONET, and now Gigabit Ethernet via the S/UNI-2xGE. PMC-Sierra provides a wide range of physical layer devices supporting different protocols and rates. This extensive set of devices that utilize POS-PHY Level 3 allow customers to develop a wide range of cards with a minimum of effort by allowing reuse of higher layer devices and board design.

Table 4-2 PMC-Sierra, Inc. POS-PHY Level 3 Physical layer devices

S/UNI-2xGE	Dual Gigabit Ethernet Controller with embedded SERDES
S/UNI-2488	Single OC-48c with embedded clock and data recovery
S/UNI-4x622	Quad OC-12c with embedded clock and data recover
S/UNI-MACH-48	Channelized OC-48

Note: The S/UNI-2xGE and other devices can be used in conjunction with the TSE 40 gigabit STS-1 cross connect element. The TSE can scale up to 160G non-blocking with four 40 gigabit TSE elements.

Table 4-3 POS-PHY Level 3 Related Standards

ATM Forum	Frame Based ATM Interface (Level 3)	STD# fb-phy-143.000
Optical Interworking Forum	System Parallel Interface Level 3	STD# oif2000.008

4.7.1 Single Gigabit Port Operation

If only one of the Gigabit Ethernet ports is required for specific applications, the POS-PHY level 3 interface does not need to be run at 104 MHz for wire-speed operation. Rather, the POS-PHY Level 3

interface can operate at 80 MHz. A power savings is realized when operating only one Gigabit Ethernet port.

4.8 Loopbacks

The S/UNI-2xGE provides three separate loopbacks to allow system developers to implement sophisticated diagnostic tests for their boards. These loopbacks are also useful for board production testing.

4.8.1 Line Side Loopback

This loopback allows a line side loopback to direct incoming traffic back out to the optics module. The 10-bit parallel data stream from the receive SERDES to be looped back into the Parallel In Serial Out (PISO) block within the transmit portion of the SERDES and back onto the line.

4.8.2 System Side Loopback

This loopback goes from the POS-PHY Level 3 interface through the Gigabit MAC and loops back towards the POS-PHY Level 3 interface. This loopback can be set per channel for diagnostic purposes.

A ring based Ethernet over SONET application could utilize the loopback function to allow a Gigabit Ethernet stream to bypass a particular node. In this case, there would be two S/UNI 2488 devices for East and West ring operation at OC-48. An FPGA would provide three POS-PHY Level 3 interfaces for the two S/UNI 2488's and the S/UNI-2xGE. If one, or both Gigabit Ethernet streams are required to bypass a particular node, the system side loopback could be enabled for that channel. Thus, one or both Gigabit Ethernet streams are looped back at the S/UNI-2xGE GMAC and sent to the East or West OC-48 device. The mapping of Ethernet frames into the OC-48c SONET payload is described in the Ethernet over SONET application note PMC-2001398.

4.8.3 System Side Loopback via GMII

This loopback goes from the POS-PHY Level 3 interface through the Gigabit MAC to the GMII interface logic and loops back towards the POS-PHY Level 3 interface.

4.9 Microprocessor Port

The S/UNI-2xGE incorporates a 16 bit generic Intel based microprocessor interface that is common to all POS-PHY Level 3 physical layer devices. A common interface simplifies the design when using multiple different PMC-Sierra physical layer devices.

4.10 Software Driver Support

In order to minimize the development effort and speed time to market, PMC-Sierra is providing a software driver for the S/UNI-2xGE. The driver provides an Application Programming Interface (API), which provides an extensive set of calls to simplify device initialization and configuration as well as statistics collection and interrupt management. Two sample calls are shown, a complete list of calls will be available in the software programmers guide .

4.10.1 Device Register Read API Example

The first sample call simplifies the process for reading the contents of a specific register in the S/UNI-2xGE device. The user provides the register number and the function derives the actual address location. An example call is shown below:

Table 4-4 Reading from Device Registers: Error! Unknown document property name.**Read**

This function can be used to read a register of a specific **Error! Unknown document property name.** device by providing the register number. This function derives the actual address location based on the device handle and register number inputs. It then reads the contents of this address location using the system specific macro, ***sysSuni2xgeRead***. Note that a failure to read returns a zero and any error indication is written to the associated DDB.

Prototype	`UINT2` **`Error! Unknown document property name.`**`Read(s`**`Error! Unknown document property name.`**`_HNDL deviceHandle, UINT2 regNum)`
Inputs	`deviceHandle` : device Handle (from *Error! Unknown document property name.**Add***) `regNum` : register number
Outputs	ERROR code written to the DDB
Returns	value read
Valid States	**`Error! Unknown document property name.`**`_PRESENT,` **`Error! Unknown document property name.`**`_ACTIVE,` **`Error! Unknown document property name.`**`_INACTIVE`
Side Effects	May affect registers that change after a read operation
Pseudocode	Begin End

4.10.2 Interrupt Service Routine Configuration API Example

Another example shown below provides a command for initializing Interrupt Service Routine processing. The S/UNI-2xGE software driver provides a flexible and configurable alarm update, status functions that can be tied to the interrupt service routine.

Table 4-5 Interrupt Service Routine: Error! Unknown document property name.**ISR**

Reads the state of the interrupt registers in the **Error! Unknown document property name.** and stores them in an ISV. Performs whatever functions are needed to clear the interrupt, from simply clearing bits to complex functions. This routine is called by the application code, from within ***sysSuni2xgeISRHandler***. If ISR mode is configured all interrupts that were detected are disabled and the ISV is returned to the Application. Note that the Application is then responsible for sending this buffer to the DPR task. If polling mode is selected, no ISV is returned to the Application and the DPR is called directly with the ISV.

Prototype	`void * Error! Unknown document property name.ISR(sError! Unknown document property name._HNDL deviceHandle)`
Inputs	`deviceHandle` : device Handle (from *Error! Unknown document property name.**Add***)
Outputs	None
Returns	(pointer to) ISV buffer (to send to the DPR) or NULL (pointer)
Valid States	`Error! Unknown document property name._ACTIVE`
Side Effects	None
Pseudocode	Begin

```
get an ISV buffer
update ISV with current interrupt status
if no valid interrupt condition
   return NULL
if in ISR mode
   disable all detected interrupts
   return ISV
else (Polling mode)
   call Error! Unknown document property name.DPR
   output NULL

End
```

4.11 S/UNI-2xGE Pinout

The pinout was designed to simplify board development using the S/UNI-2xGE. Care was taken to provide a direct connection to optical modules without having to cross high-speed data lines. In addition, the power escape patterns were analyzed to ease board design and improve signal integrity. Power and grounds have been grouped in addition to providing direct access for high-speed signal lines for connection to optics modules from the SERDES pins and to copper PHYs via the GMII interface. Please refer to **Figure 4.3** for a sample connection diagram for an optical module and the S/UNI-2xGE.

Figure 4.3 PM3386 to ODL Interface

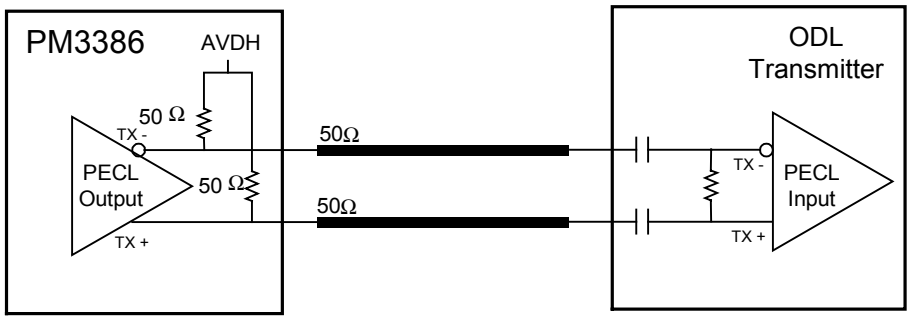

ODL to PM3386 Interface

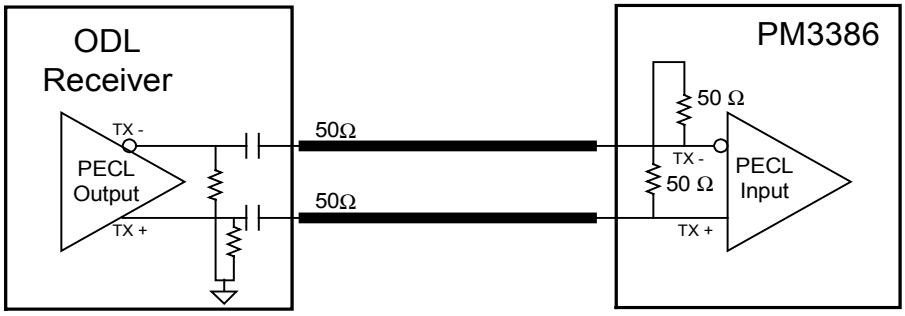

4.12 JTAG

A JTAG port is provided to test the digital I/O on the S/UNI-2xGE for board testing.

5 POS-PHY LEVEL 3 FPGA SUPPORT

Many system developments use FPGAs to complete designs and deliver equipment as fast as possible. FPGAs allow system designers to perform multiple spins without the delay of ASIC fabrication times. FPGAs also allow development of functions that are not completely specified, allowing future changes if required (i.e. tracking an emerging standard). After FPGA code has been tested and stable, ASICs can be developed to reduce the cost.

PMC-Sierra recognizes the importance of FPGAs in system design and time-to-market, which is why relationships have been developed with leading FPGA vendors to provide interoperable POS-PHY Level 3 cores in conjunction with PMC-Sierra's broad line of physical layer devices. PMC-Sierra is partnering with leading FPGA vendors. For more information, please contact your local PMC-Sierra representative.

6 EXAMPLE APPLICATIONS OF THE S/UNI-2XGE

6.1 Ethernet Ports in the WAN

The S/UNI-2xGE addresses a number of applications in the Wide Area Network and Internet POP. Figure 6-1 shows example applications for the S/UNI-2xGE in Internet POPs.

Figure 6-2 indicates where Ethernet ports will be found within the Enterprise WAN Edge.

Figure 6.1 Ethernet Ports on POP WAN Equipment

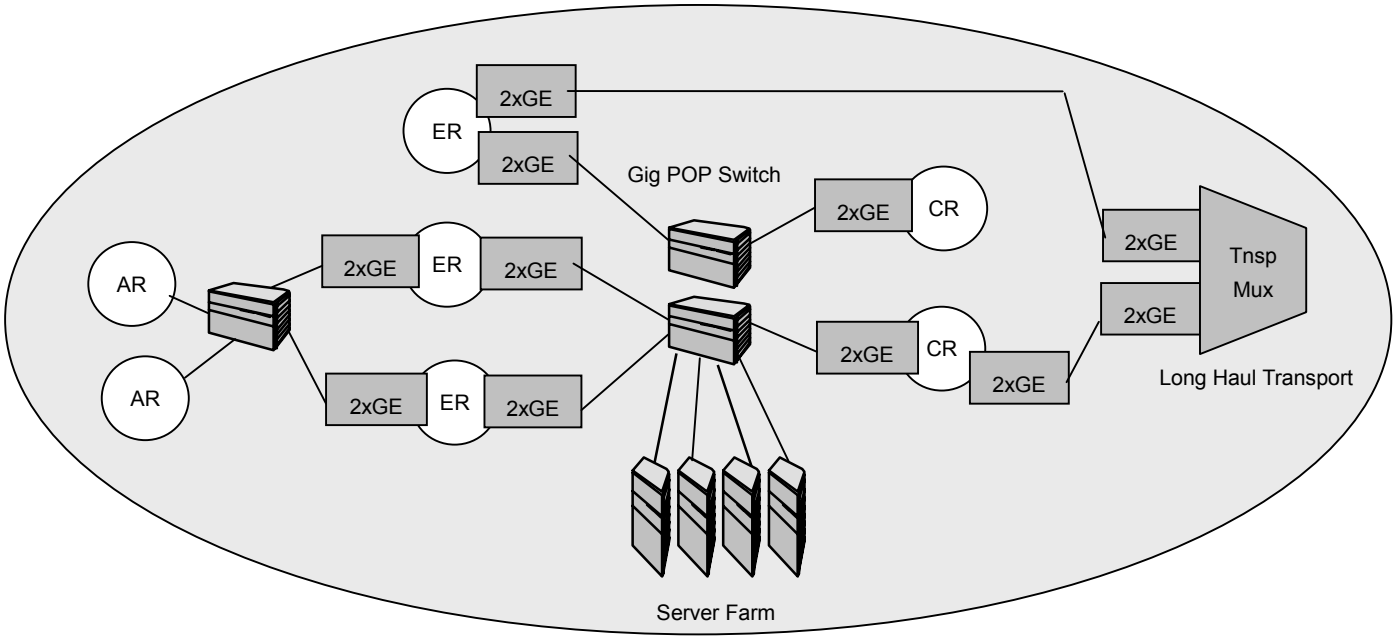

AR – Access Router ER – Edge Router EntR – Enterprise Switch/Router

CR – Core Router Tnsp Mux – Transport Mux

Figure 6.2 Ethernet Ports Enterprise WAN Edge

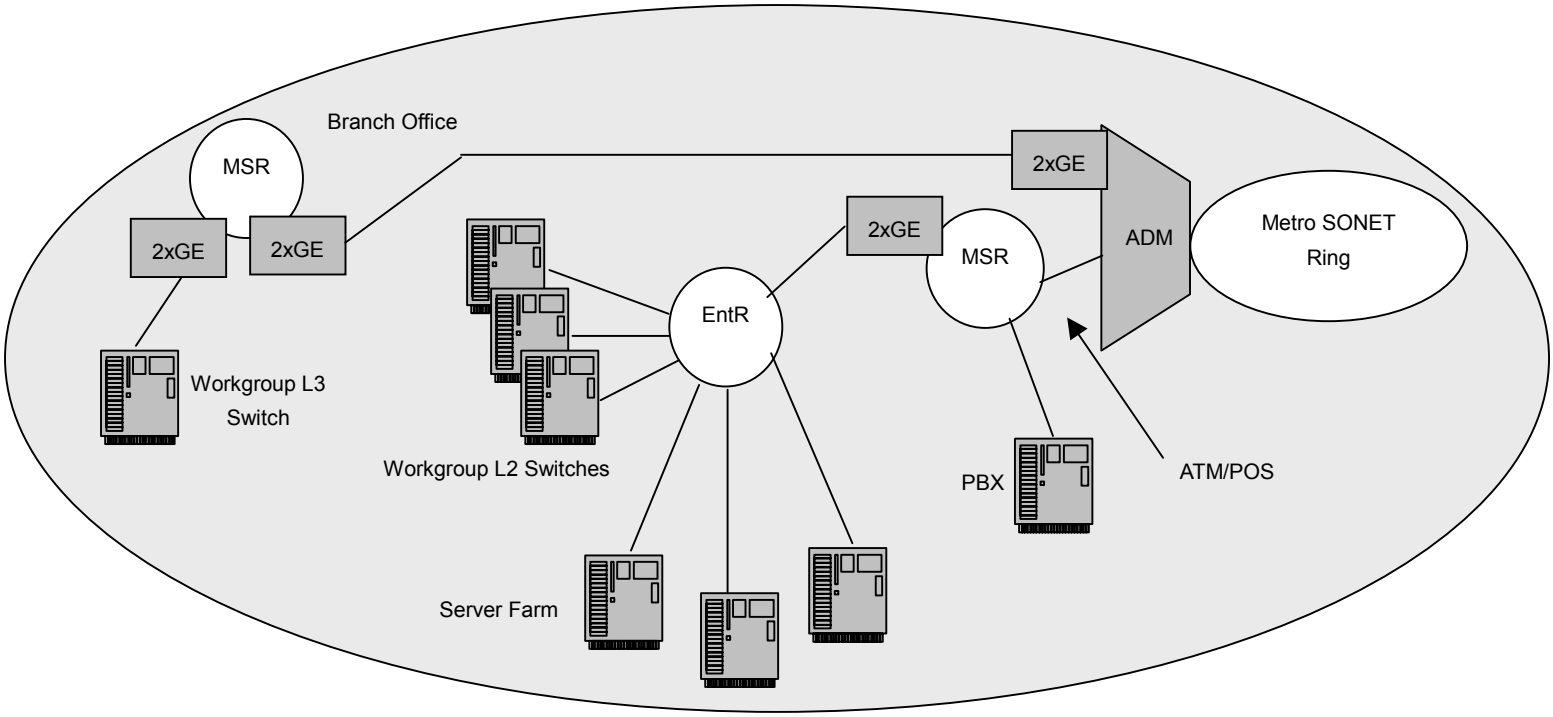

ADM – Add Drop Mux **EntR – Enterprise Switch** **MSR – Multi-Service Switch/Router**

6.2 The S/UNI-2xGE in Core and Edge Router Applications

The S/UNI-2xGE is taylored to meet the needs of Gigabit Ethernet designs in Core and Edge Router applications. By leveraging the POS-PHY Level 3 interface, the S/UNI-2xGE can seamlessly link to upper layer devices. Figure 6.3 provides an example of this application. Also shown in Figure 6.3 are two other PMC POS-PHY Level 3 devices that support OC-48c POS or Quad OC-12c POS allowing multiple line cards while maximizing reuse of the back-end card design.

Figure 6.3 Core and Edge Router Applications

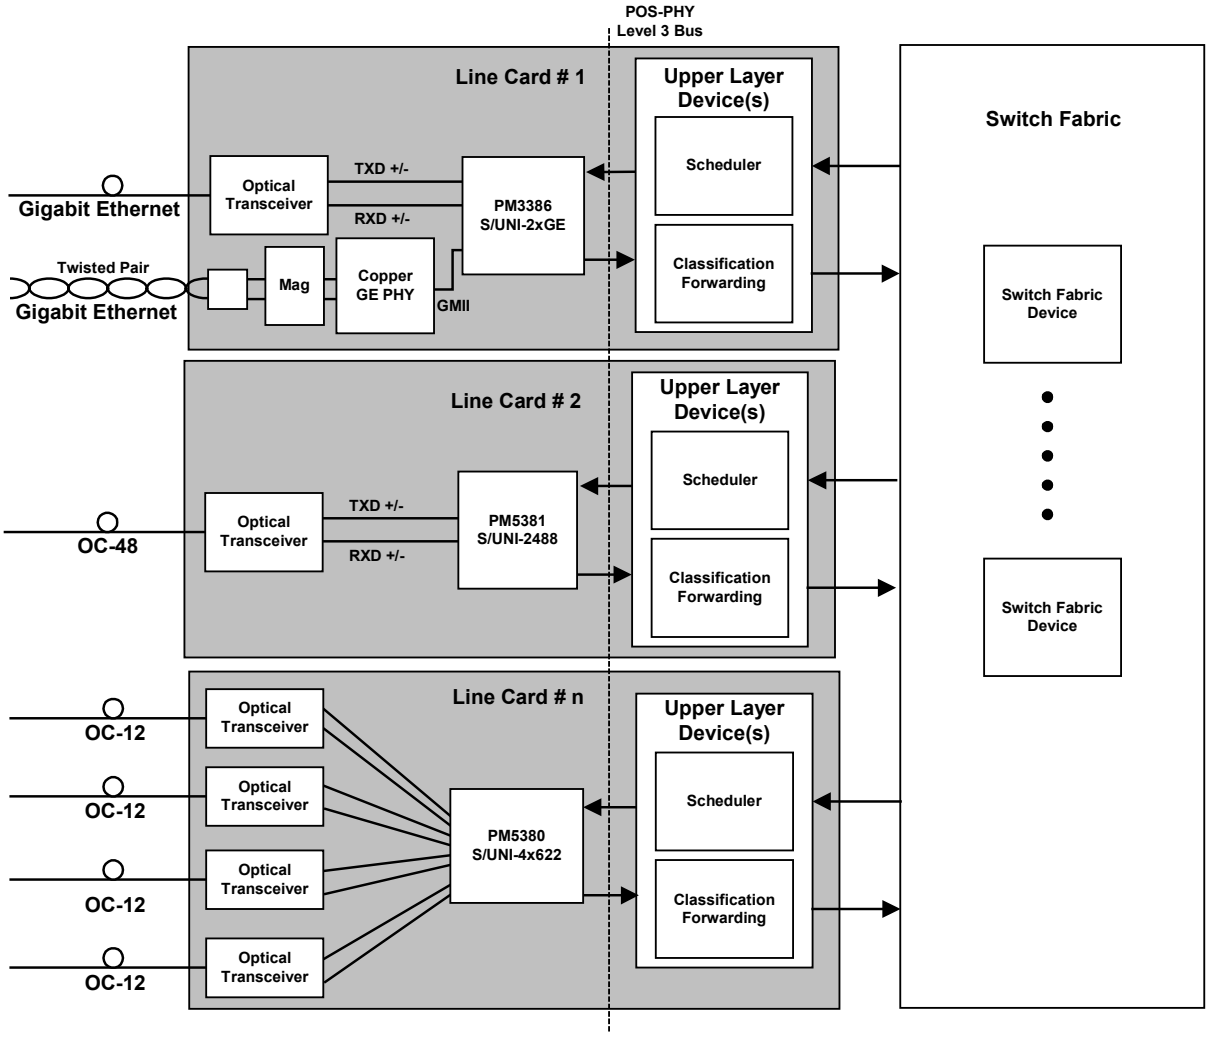

6.3 Ethernet over SONET Applications

As Ethernet begins to be deployed in Metro and Wide Area Networking (MAN and WAN) applications, Ethernet over SONET functionality becomes a requirement. This allows Ethernet to be transported across existing infrastructure. One example is to provide transparent bridging between geographically diverse sites and to allow enhanced VPN features to be implemented across the WAN.

The basic functionality required to implement EOS is to properly delineate the incoming or outgoing frames so that they can be processed by either the S/UNI 2488 (OC-48c POS device) or the S/UNI-2xGE (dual port SERDES/GMAC). As shown below, the FPGA uses two PL3 interfaces to connect the S/UNI-2xGE and the S/UNI 2488 to implement the EOS function.

For greater detail, please refer to the “Ethernet over SONET” application note PMC-2001398.

Figure 6.4 Ethernet over SONET Application

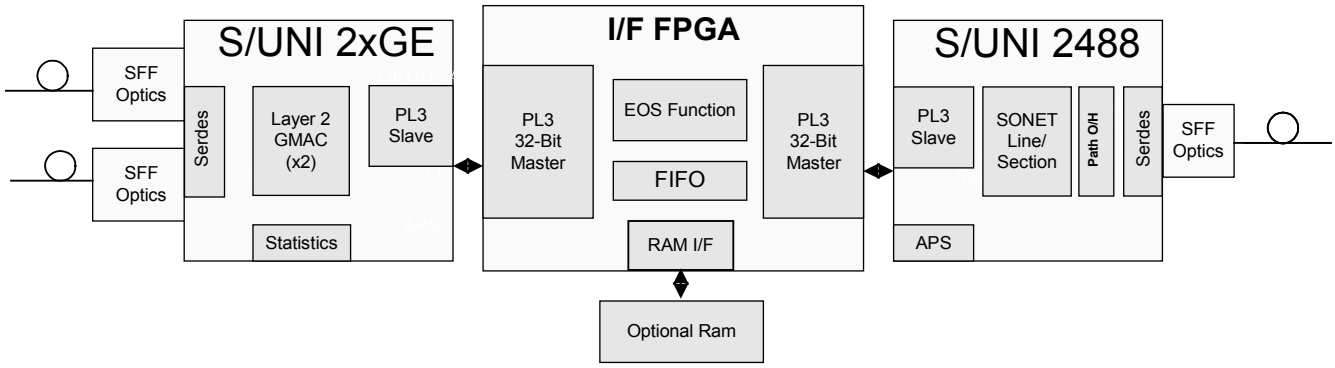

In the above application, the PPP over SONET capabilities of the S/UNI 2488 are used to reduce the packet processing requirements of the EOS functions implemented in the interface FPGA, this is one of a number of potential solutions. There are emerging alternative packet framing schemes that may be also be used. These schemes include Generic Frame Protocol (GFP) and Dynamic Packet Transport (DPT) which uses Spatial Reuse Protocol (SRP). These methods are being considered at the ITU (for GFP) and the IEEE Resilient Packet Ring Study Group (for both GFP and SRP).

Figure 6.5 Packet over SONET Applications

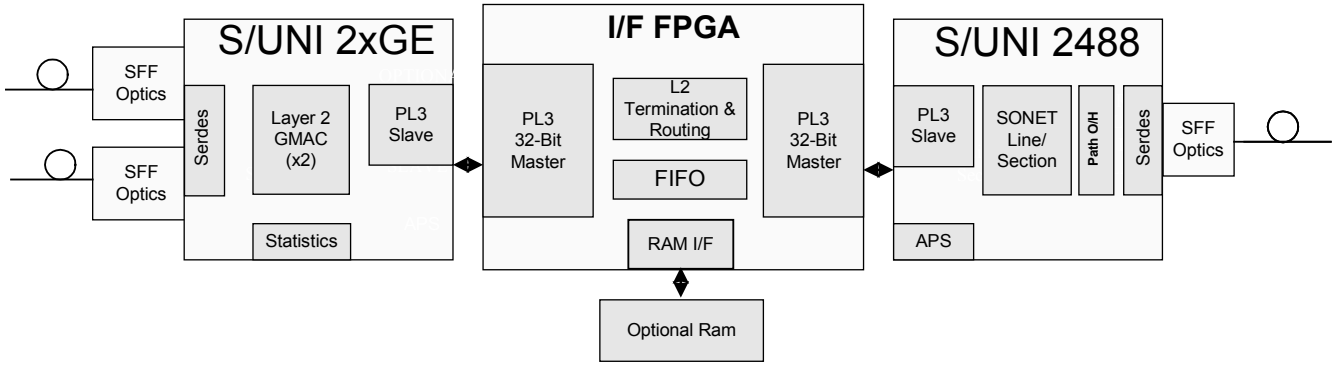

In Figure 6-5 we describe a Packet over SONET application. In this application, the FPGA must terminate and route the packets to either the S/UNI 2488 or the S/UNI-2xGE. The FPGA in this application also uses two PL3 interfaces to connect the S/UNI-2xGE and the S/UNI 2488 to implement the POS function.

6.4 Optical Cross-Connect Applications

Optical Cross-connects provide a new level of system hierarchy by providing Layer 1 optical stream cross-connecting. Many initial Optical Cross-connects address the SONET market space, however, Gigabit Ethernet is now becoming a requirement for these systems.

Figure 6.6 Gigabit Ethernet Ports on an Optical Cross-Connect

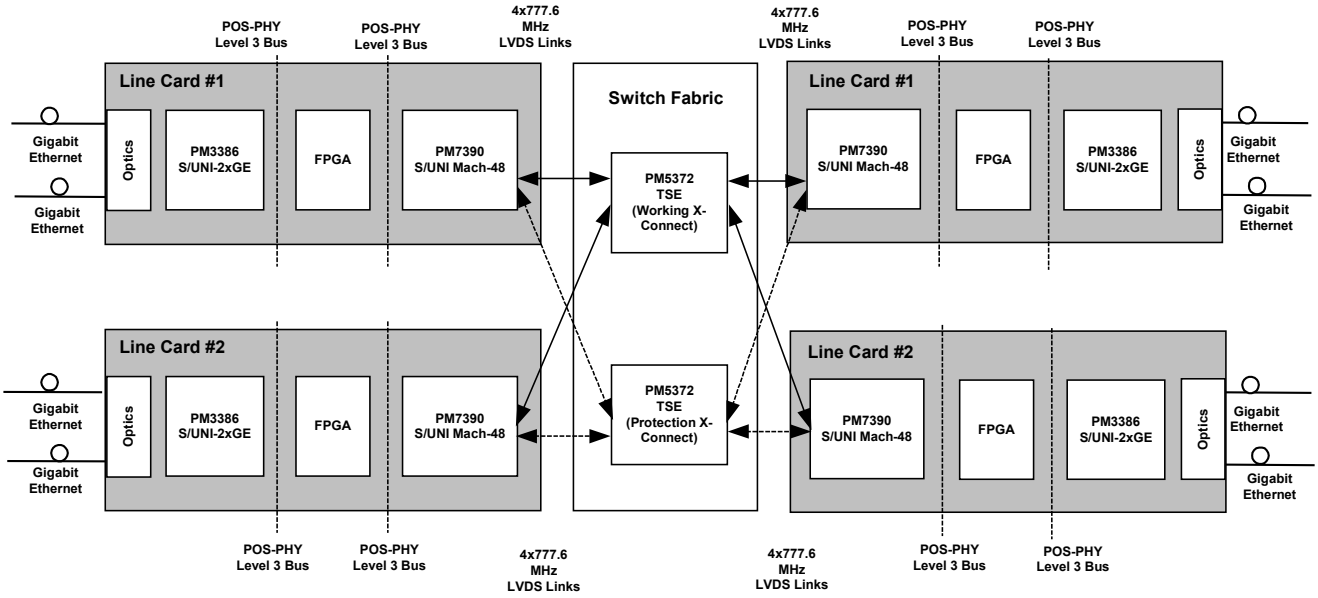

6.5 Optical Routers and Optical Multi-Service Switches

Gigabit Ethernet is becoming a new requirement for Multi-service switches. In Figure 6-7, we show how the S/UNI-2xGE can be deployed in this type of application. The S/UNI-2xGE connected via an FPGA to the S/UNI-Mach 48 now allows Gigabit Ethernet to be mapped across the cross connect fabric and allow grooming of traffic into various protocol cards. These protocol cards can provide multi-service switching or routing to a larger switching fabric.

Multi-service switches and routers require Optical solutions to manage the dramatic increase in bandwidth to be processed. These new types of equipment leverage the optical cross-connect to bypass traffic that is not destined for processing at their node. Thus, the equipment can support a much greater bandwidth than the switch or router can provide, by bypassing traffic destined for other nodes and only processing traffic destined for that particular node.

Figure 6.7 Gigabit Ethernet Ports on an Optical Multi-Service Switch or Optical Router

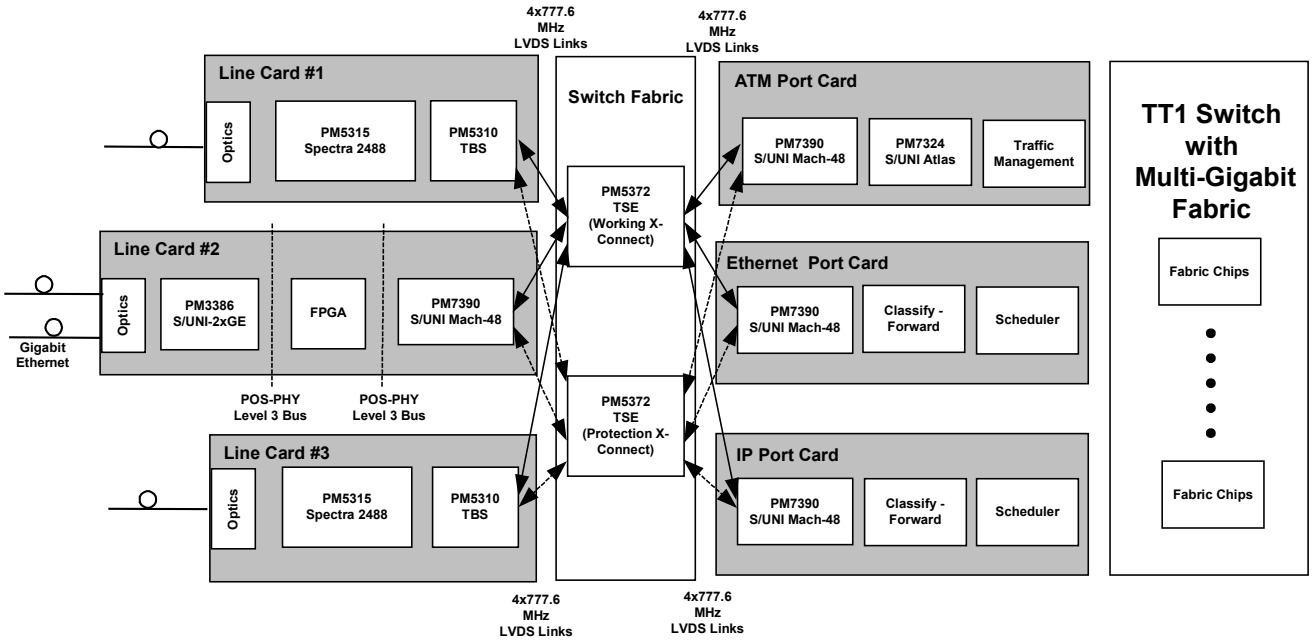

7 CONCLUSION

The S/UNI-2xGE fully addresses the front end requirements for Gigabit Ethernet in exploding markets of Router interconnect, Metro and Core Transport and Multi-Service switches. The key requirements of these markets are:

- High port density.
- Low power.
- Full MAC features including: statistics, frame filtering, VLAN and Jumbo frame support.
- Standard system interfaces to ensure compatibility.
- Extensive software driver support to simplify software development.

There is a single phrase that summarizes PMC-Sierra's goal in developing the S/UNI-2xGE device:

Ethernet is a persuasive protocol that will grow throughout the entire Global Network. Low power, high density and standard interfaces are the key to a successful deployment.

NOTES

CONTACTING PMC-SIERRA, INC.

PMC-Sierra, Inc.
105-8555 Baxter Place Burnaby, BC
Canada V5A 4V7

Tel: (604) 415-6000

Fax: (604) 415-6200

Document Information: document@pmc-sierra.com
Corporate Information: info@pmc-sierra.com
Application Information: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com

PMC-1991728 (P1) Issue date: October 2000