Freescale Semiconductor
Application Note

Document Number: AN3537
Rev. 1, 12/2007

Accommodating Layer 2 Padding (Shimming) with the Enhanced Three-Speed Ethernet Controller (eTSEC)

by Todd Kanning
Networking and Multimedia Group
Freescale Semiconductor, Inc.
Austin, TX

The enhanced three-speed Ethernet controller (eTSEC) offered on many PowerQUICC™ II Pro, PowerQUICC™ III, and other devices, allows for flexible manipulation of incoming and outgoing Ethernet data. One such feature is the ability to receive and propagate padded, or "shimmed," OSI layer 2 data to accommodate custom routing or direction of Ethernet data within a network. This application note describes what the shimming functionality does and the details of how to best utilize it.

This application note assumes that the reader already possesses fundamental knowledge of the OSI stack model and Ethernet protocols, as well as the primary hardware offload capabilities of the eTSEC and the proper configuration thereof. Please refer to the appropriate device reference manual for additional information regarding the offload capabilities of the eTSEC.

Currently, the following Freescale devices offer the ability to receive and propagate padded, or "shimmed," layer 2 frames: MPC8379E, MPC8315E, and MPC8572E.

Shimmed layer 2 data received on all other eTSEC-enabled devices prevents the parser from properly identifying the SFD-to-DA transition within the frame.

Contents

1 Shim Header Basics

In some cases, OSI layer 2 frames can be prepended with data intended to allow or assist with custom routing, parsing, or filing within a closed network. Such prepended "shim" data is defined by the owner of the network. The format, length, and content of the shim header is created within the network and used only within the network. Because this shim data is inserted between the SFD and the DA, any network element encountering such shimmed frames must be aware of at least its presence and length, if not its content and intended use, in order to properly process and parse the data.

2 Shim Headers and the eTSEC

The Receive Control Register (RCTRL), shown in Figure 1 and Table 1 (taken from the MPC8572E Reference Manual), contains a field called "L2OFF," which allows control of a layer 2 offset, also known as a "shim header," that is used to accommodate user-defined extension headers that exist between the SFD and DA. The value of this field represents the number of octet pairs from the start of the frame (SFD) that the parser should expect to see (and skip) before the first byte of the Ethernet DA.

Offset eTSEC: 0x2_4300 — Access: Read/Write

	0–6	7	8–10	11–15
R/W	L2OFF	TS	—	PAL
Reset	All zeros			

	16	17	18	19	20	21	22	23	24–25	26	27	28	29	30	31
R/W	CFA	LFC	VLEX	FILREN	FSQEN	GHTX	IPCSEN	TUCSEN	PRSDEP	—	BC_REJ	PROM	RSF	EMEN	—
Reset	All zeros														

Figure 1. RCTRL Register Definition

Table 1. RCTRL[L2OFF] Field Description

Bits	Name	Description
0–6	L2OFF	Layer 2 offset. The number of octet pairs from the start of the frame that the parser should expect to see before the first byte of the Ethernet DA. For frames received over Ethernet, the start of frame is regarded as the SFD symbol. For packets received through the FIFO packet interface the start of frame is regarded as the first octet of received data. The user may think of this value as representing the length—in multiples of 2 bytes—of a "shim" header that is inserted between the SFD and DA. By writing to RCTRL with a mask of 0xFE00_0000, the even byte length restriction is guaranteed. For normal frames, this field should be left as 0.

Valid values for RCTRL[L2OFF] are from 0x0 to 0x7F. Because this field represents the number of octet pairs expected between the SFD and the DA, the eTSEC can accommodate shim headers from 0 to 254 bytes.

Note that the receive control register (RCTRL) must be written either during initialization after a system reset or after a graceful receive stop has completed.

2.1 How to Properly Accommodate Reception of Layer 2 Shim Data

The primary configuration required for the eTSEC to properly receive and handle Ethernet packets containing a layer 2 shim header is the assignment of RCTRL[L2OFF] as noted above. Following are additional points of interest and requirements for optimal behavior in this situation.

- For proper shim header processing, the MAC must be programmed to operate in promiscuous mode by setting RCTRL[PROM]. Promiscuous mode allows the MAC to bypass all DA filtering provided directly by the MAC. In the presence of a properly anticipated shim header, all parser/filer capabilities based on DA are still supported except “miss” (RxBD[M]=1), “mulitcast” (RXBD[MC]=1), “broadcast” (RXBD[BC]=1 and PID1[EBC]=1).
- Because some shim headers may mimic control frame headers, the MAC cannot respond directly to MAC control frames, including 802.3x PAUSE frames. Flow Control must be disabled by clearing MACCFG1[Rx_Flow] (this is the default state).
- The management information base (MIB) counters evaluate frame size including the user defined shim header.
- The maximum frame length register (MAXFRM[Maximum Frame]) must be adjusted to include the size of the expected shim header (add RCTRL[L2OFF] x 2 bytes) to prevent correctly sized Ethernet headers from being incorrectly rejected or counted as oversized.
- The CRC calculation is done over the entire frame, including the user defined shim header, as shown in Figure 2.
- The arbitrary extraction of L2 data begins 8 bytes before the shim header. (See the description of the receive bit field extract control register (RBIFX) in the eTSEC chapter of the applicable device reference manual.) As shown in Figure 2, because the arbitrary extraction offset is limited to 64 bytes, and the shim header can be up to 254 bytes, some of the shim header and L2 header may become inaccessible through arbitrary extraction.

Note that the appearance in memory of each of the portions shown in Figure 2 may be dependant upon that state of individual parsing enable bits. (For example, MACCFG2[PreAM RxEN] must be set for the preamble and SFD to appear in memory.)

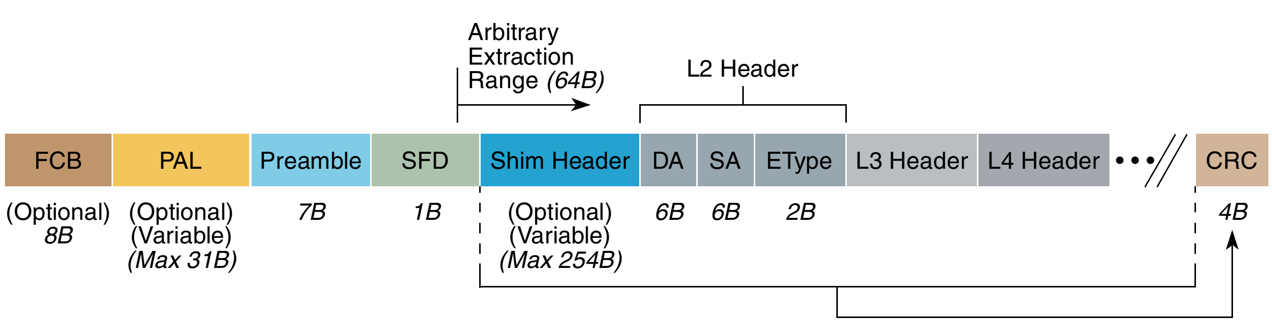

Figure 2. Ethernet Frame with Shim Header Present

- Packet alignment padding (configured with RCTRL[PAL]) is inserted before the user defined shim header, rather than the L2 header, if preamble extraction is disabled, as shown in Figure 2.
- VLAN extraction does not take the shim header into account. If shim headers are used, VLAN extraction must be disabled by clearing RCTRL[VLEX]. Note that normal parsing recalculates the VLAN offset, so the following VLAN filer properties all function properly with shim headers: VLAN tag seen (PID1[VLN]), canonical format indicator (PID1[CFI]), VLAN network identifier (PID8[VID]), and VLAN user priority (PID9[PRI]).
- Shim headers are not supported in FIFO mode unless L2 parsing is enabled (RCTRL[PRSFM] = 1).

2.2 How to Properly Propagate (Transmit) Layer 2 Shim Data

Following are points of interest and requirements for proper propagation of received Ethernet packets containing a layer 2 shim header.

- Transmitted pause frames do not include the shim header and would be incorrectly processed by a receiver expecting to see a shim header. As such, flow control must be disabled by clearing MACCFG1[Tx_Flow].
- VLAN insertion does not take the shim header into account. As such, VLAN insertion must be disabled by clearing TCTRL[VLINS] if shim headers are present.

3 Terminology

DA = Destination address

eTSEC = Enhanced three-speed Ethernet controller

FCB = Frame control block

FCS = Frame check sequence

FEC = Forward equivalence class

MIB = Management information base

MPLS = Multi-protocol label switching

OSI = Open systems interconnect

OSI 7-layer model = the standard model for data abstraction within Ethernet

PAL = Packet alignment padding

PDU = Protocol data unit

RMON = Remote network monitoring

SFD = Start frame delimiter

TOE = TCP/IP offload engine

VLAN = Virtual LAN (IEEE Std 802.11Q™)

4 Revision History

Table 2 provides a revision history for this application note.

Table 2. Document Revision History

Rev. Number	Date	Substantive Change(s)
0	11/2007	Initial release.
1	12/2007	Corrected maximum shim header length (from 128B to 254B) noted in Figure 2.

THIS PAGE INTENTIONALLY LEFT BLANK

THIS PAGE INTENTIONALLY LEFT BLANK

How to Reach Us:

Home Page:
www.freescale.com

Web Support:
http://www.freescale.com/support

USA/Europe or Locations Not Listed:
Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or
+1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120 191014 or
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:
Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
+1-800 441-2447 or
+1-303-675-2140
Fax: +1-303-675-2150
LDCForFreescaleSemiconductor
@hibbertgroup.com

Document Number: AN3537
Rev. 1
12/2007