BCM88830

FlexE Overhead, OAM, and IEEE 1588 Handling

Application Note

Table of Contents

Chapter 1: Introduction

1.1 About this Document

This document describes the mechanisms that process FlexE overhead (OH), FlexE IEEE 1588, and FlexE client operations, administration, and management (OAM) between the BCM88830 FlexE core and the external processing unit (FPGA or CPU), as well as the packet formats that enable this processing.

This document does *not* explain how to process FlexE OH, FlexE client OAM, or FlexE IEEE 1588.

1.2 Overview

To support FlexE client OAM, FlexE OH processing, and FlexE 1588, the adjacent FPGA (or CPU)should perform the processing.. For external processing, the external device must be connected to either one of the NIF Ethernet interfaces over a 10GbE or 25GbE interface or to both interfaces (the management ports). The extracted data that is received from FlexE ports is packetized and sent to the processing unit, while the inserted data is sent packetized from the processing unit to the FlexE framer.

The data path between the FlexE framer and the management ports can be L1 (directly connected to the FlexE framer) or L2 and L3 through the packet processor (PP) or traffic manager (TM).

Dedicate two Ethernet ports for the external unit. For L1, the Ethernet ports must be 25GbE for OH and OAM, and 10GbE for IEEE 1588 or Precision Time Protocol (PTP). For L2, the two ports should have the same bandwidth.

The high-level concept of operation is that the FlexE core in the BCM88830 extracts the relevant data from the incoming FlexE streams, packetizes the extracted messages, and sends them to the processing unit (for example, the external CPU or external FPGA). In the other direction, the FlexE core receives packets from the processing unit, unpacketizes those packets, and inserts the data in the relevant locations of the outgoing FlexE streams.

Chapter 2: IEEE 1588

2.1 IEEE 1588 Packet Format

IEEE 1588 messages are inserted into and extracted from the FlexE OH, packetized, and transmitted to or received from the processing unit.

Figure 1: Internal IEEE 1588 Packet Format

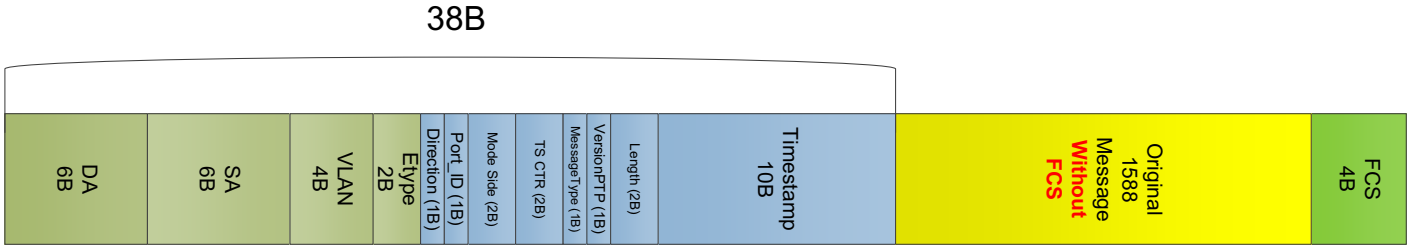

Table 1: IEEE 1588 Packet Fields

Field Name	Size	Description
DA	6B	Destination address (default value 0xFFFF_FFFF_FFFF).
SA	6B	Source address (default value 0x0).
VLAN	4B	The first 2 bytes are Vlan_Tag (default value 0x8100). The next 2 bytes are Vlan_ID (values 0x0 to 0x3b, OH management port).
Etype	2B	EtherType (default value 0x88F7).
Direction	1B	0x2 for processing unit to FlexE framer direction and 0x1 on the opposite direction.
Port_ID	1B	Values 0x0 to 0x3b (OH management port. Same values as vlan_id). **NOTE:** A maximum of 60 ports is allowed for IEEE 1588.
Mode_side	2B	The bits are as follows: ■ Bit [9:8] vlan_num: – 0x0: No VLAN tag. – 0x1: One VLAN tag (TPID = 0x8100). – 0x2: Two VLAN tags (outer-TPID = 0x8500, inner-TPID = 0x8100). ■ Bit[7:6] protocol_type: – 0x0:L2 Eth. – 0x1:IPv4/UDP. – 0x2:IPv6/UDP. ■ Bit[2]: For testing (also called pkt_side). When set, the FlexE core duplicates the PTP packet to transmit (standard path) and sends back to the processing unit. ■ Bit[1]: work_step:(value = 1: Latch TX two-step timestamp). ■ Bit[0]: udpchsum_update:(value = 1 update UDP checksum). ■ Other bits: Reserved 0.
TS_CTR	2B	The bits are as follows: ■ Bit[10] FPGA to chip Eth use txosts_en – One-step processing enable. ■ Bit[9] FPGA to chip Eth use txtsts_en – Two-step processing enable. ■ Bit[8] CF update ctrl. 1: update. ■ Bit[7:0] CF offset value. Others bits: Reserved 0.

Table 1: IEEE 1588 Packet Fields (Continued)

Field Name	Size	Description
Message_Type	1B	The bits are as follows: ■ Bit[7:0] = 0x1f SSM. ■ Bit[7:4] = 0x0. ■ Bit[3:0] same as PTP message type.
versionPTP	1B	Default value 0x2.
Length	2B	Total bytes of the blue and yellow areas in Figure 1, Internal IEEE 1588 Packet Format. In other words, the total bytes starting with the Direction field and ending with the original IEEE 1588 message (without FCS).
Timestamp	10B	{28'b0, 52b timestamp}. 52b are in a format of {48b nanoseconds, 4b fraction of nanoseconds}.
1588 Message	Variable	Original PTP frame/SSM frame + update FCS.
FCS	4B	Updated FCS.

Chapter 3: FlexE OH and OAM

FlexE OH and client OAM are inserted into and extracted from the FlexE core, packetized, and then transmitted to and received from the processing unit. The processor of both OAM and OH must be the same unit because both share the same packet encapsulation.

The BCM88830 device supports the following FlexE OH and client OAM features:

- Up to 16 ports of FlexE overhead extraction and insertion.
- Up to 320 channels for OAM extraction and insertion in FlexE client layers in the line side.
- Up to 320 channels for OAM extraction and insertion in FlexE client layers in the system side.
- Alarm messaging for all ports and channels.
- Protection switching messaging.

3.1 Extraction Frame Format Definition

The encapsulation has two stages: the first stage is OH_SEG (overhead segment), and the second stage is OH_PKT (overhead packet). OH_SEG refers to cutting the overhead into many small slices, and OH_PKT refers to encapsulating the overhead slices of different ports and different channels into packets according to certain rules.

3.1.1 FlexE OH_SEG Frame Format

The following figure shows the FlexE overhead (OH_SEG) layer.

Figure 2: FlexE OH_SEG Format (Type OH)

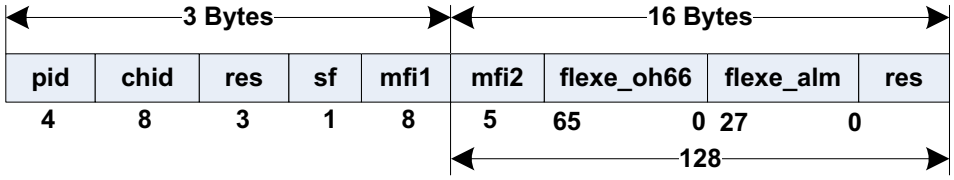

Table 2: FlexE OH_SEG Format (Type OH)

Field Name	Size	Description
pid	4b	Framer's OH/OAM processor port ID (0x6 for NIF FlexE port, 0x7 for backplane FlexE port).
chid	1B	FlexE interface ID range is 0 to 15.
res	3b	Reserved.
sf	1b	Status indication (SSF signal).
mfi1	1B	Overhead block indication, mfi[7:3] is the multi-frame number, and mfi[2:0] is the overhead block
mfi2	5b	(Not yet standard) 5-bit multi-frame indication (MFI) to support smaller particles. The range is 0 to 19.
flexe_oh66	66b	One OH block.
flexe_alm	28b	FlexE alarms (see Table 3, FlexE OH Alarms).
res	29b	Reserved.

The alarm (ALM) arrangement position and specific meaning are shown in the following table.

Table 3: FlexE OH Alarms

Field Name	ALM	Description
ALM[27]	flexe_ccm19	FlexE Rx instance #19 calendar mismatch alarm
ALM[26]	flexe_ccm18	FlexE Rx instance #18 calendar mismatch alarm
ALM[25]	flexe_ccm17	FlexE Rx instance #17 calendar mismatch alarm
ALM[24]	flexe_ccm16	FlexE Rx instance #16 calendar mismatch alarm
ALM[23]	flexe_ccm15	FlexE Rx instance #15 calendar mismatch alarm
ALM[22]	flexe_ccm14	FlexE Rx instance #14 calendar mismatch alarm
ALM[21]	flexe_ccm13	FlexE Rx instance #13 calendar mismatch alarm
ALM[20]	flexe_ccm12	FlexE Rx instance #12 calendar mismatch alarm
ALM[19]	flexe_ccm11	FlexE Rx instance #11 calendar mismatch alarm
ALM[18]	flexe_ccm10	FlexE Rx instance #10 calendar mismatch alarm
ALM[17]	flexe_ccm9	FlexE Rx instance #9 calendar mismatch alarm
ALM[16]	flexe_ccm8	FlexE Rx instance #8 calendar mismatch alarm
ALM[15]	flexe_ccm7	FlexE Rx instance #7 calendar mismatch alarm
ALM[14]	flexe_ccm6	FlexE Rx instance #6 calendar mismatch alarm
ALM[13]	flexe_ccm5	FlexE Rx instance #5 calendar mismatch alarm
ALM[12]	flexe_ccm4	FlexE Rx instance #4 calendar mismatch alarm
ALM[11]	flexe_ccm3	FlexE Rx instance #3 calendar mismatch alarm
ALM[10]	flexe_ccm2	FlexE Rx instance #2 calendar mismatch alarm
ALM[9]	flexe_ccm1	FlexE Rx instance #1 calendar mismatch alarm
ALM[8]	flexe_ccm0	FlexE Rx instance #0 calendar mismatch alarm
ALM[7]	flexe_lf	FlexE Rx PHY Local fault alarm
ALM[6]	flexe_lof	FlexE Rx instance loss of frame alarm
ALM[5]	flexe_lom	FlexE Rx instance loss of multi-frame alarm
ALM[4]	flexe_gidm	FlexE Rx instance group id mismatch alarm
ALM[3]	flexe_pmm	FlexE Rx instance MAP mismatch alarm
ALM[2]	flexe_lop	FlexE Rx instance loss of pad alarm
ALM[1]	flexe_pidm	FlexE Rx instance id mismatch alarm
ALM[0]	flexe_rpf	FlexE Rx RPF alarm

3.1.2 OAM/ALM Frame Format

Two types of FlexE client OAM are supported: service principal name (SPN) OAM and metro transport network (MTN) path OAM.

The OAM/ALM on the FlexE client layer format (OH_SEG) is shown in the following figure.

Figure 3: FlexE OH_SEG Format (type OAM/ALM)

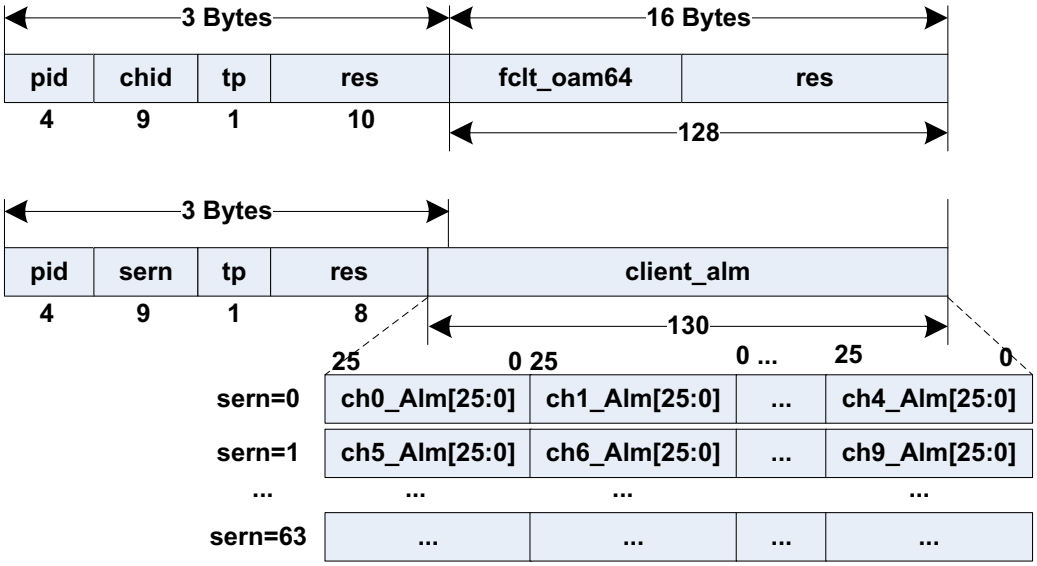

Table 4: FlexE OH_SEG Format (Type OAM/ALM)

Field Name	Size	Description
pid	4b	Framer's OH/OAM processor port ID (0x8 for NIF FlexE port OAM, 0x9 for backplane FlexE port OAM).
chid/sern	9b	Channel ID/serial number. For OAM, this field represents the channel ID; for alarm SEG, this field represents the sequence number.
tp	1b	Type indication field: ■ 0 represents that the SEG is an OAM SEG, and at this time, the 16-byte payload field has 64 bits as a padding field. ■ 1 means that the SEG is an alarm SEG. To ensure that the alarms of the same channel will not cross SEGs, each alarm SEG carries a total of 130 bits of alarm information; that is, each SEG carries five channel alarms. The 26 × 320-bit alarm requires 64 SEGs to be transmitted (see Table 5).
res	10b/8b	Reserved (10b for OAM, 8b for ALM).

Table 5: Client ID Alarm Position per Serial Number

sern	clientID				
0	ch0_alm	ch1_alm	ch2_alm	ch3_alm	ch4_alm
1	ch5_alm	ch6_alm	ch7_alm	ch8_alm	ch9_alm
2	ch10_alm	ch11_alm	ch12_alm	ch13_alm	ch14_alm
…	…	…	…	…	…
63	ch315_alm	ch316_alm	ch317_alm	ch318_alm	ch319_alm

The supported alarms per client are shown in the following table.

Table 6: FlexE Client Alarms

Field Position	Alram Name	Description
alm[0]	lpi_alm	The receiving direction detects an alarm that there is an LPI fault pattern in the data stream.
alm[1]	rf_alm	The receiving direction detects the alarm that there is a remote fault pattern in the data stream.
alm[2]	lf_alm	The receiving direction detects an alarm that there is a local fault pattern in the data stream.
alm[3]	rx_sdbei	Signal degradation (SD) alarm based on the backward error indication (BEI). The BEI is equivalent to the remote error indication (REI).
alm[4]	rx_sdbip	SD alarm based on bip8.
alm[5]	bas_rdi	Remote defect indication (RDI) alert for building automation system (BAS) OAM.
alm[6]	bas_csf_lpi	LPI alarm of BAS OAM.
alm[7]	bas_cs_rf	RF alarm of BAS OAM.
alm[8]	bas_cs_lf	LF alarm of BAS OAM.
alm[9]	bas_no_receive_alm	No BAS OAM alarm is received within the configured period. That is, a BAS loss alarm.
alm[10]	rx_sfbei	Signal failure (SF) alarm based on the BEI. The BEI is equivalent to the REI.
alm[11]	rx_sfbip	SF alarm based on bip8.
alm[12]	bas_period_alm	The alarm that the BAS period extracted from the receiving direction is inconsistent with the expected value.
alm[13]	cv_los_alm	CV loss alarm.
alm[14]	cs_los_alm	CS loss alarm.
alm[15]	aps_los_alm	APS loss alarm.
alm[16]	dm1_los_alm	DM1 OAM loss alarm.
alm[17]	dmm2_los_alm	DMM2 OAM loss alarm.
alm[18]	dmr2_los_alm	DMR2 OAM loss alarm.
alm[19]	mtn_bas_seq_alm	ALM of incorrect BAS sequence in MTN mode.
alm[20]	mtn_babl_lock_alm	MTN mode B-A-B-L sequence alarm.
alm[21]	mtn_low_seq_lock_alm	64 low-priority sequence error alarm.
alm[22]	mode_alm	Pattern matching error alarms.
alm[23]	oci_fault_alm	OCI pattern alarm.
alm[24]	client_fail_alm	Physical layer alarms summarized by the front-end module.
alm[25]	mtn_lowp_lock_alm	64 low-priority sequence error alarm (the difference from mtn_low_seq_lock_alm is that the monitoring range becomes smaller and the response time becomes shorter).

3.1.3 Extraction Packet Frame Format

The OH_PKT packet encapsulates several OH_SEG segments into an Ethernet-like packet. Nine OH_SEG segments are encapsulated. Valid OH_SEG segments have a pid value of 0x6 to 0x9. Invalid OH_SEG segments are marked with pid = 0xF.

The following figure shows the OH_PKT format.

Figure 4: OH_PKT Format

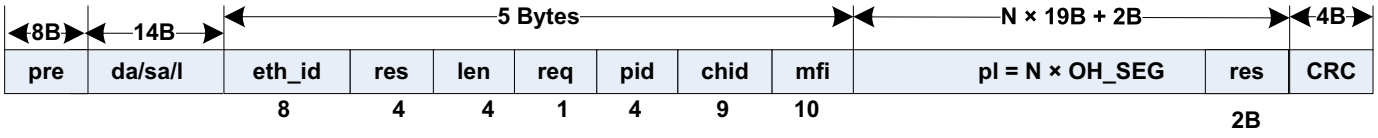

Table 7: OH_PKT Format

Field Name	Size	Description
da/sa	12B	Configurable. By default, all are 0. A single configuration exists for all packet types. Different Ethernet ports can be the same or different.
l	2B	16'd138.
eth_id	1B	Ethernet packet ID, 0–255 cyclically accumulated. A single *sequence number* is used for the interface (because all packet types traverse through this interface).
res	4b	Reserved.
len	4b	Length indication = 9.
req	1b	Indicates whether the overhead request signal is carried (only for OH request). ■ 1: OH request. ■ 0: No request
pid	4b	Valid only if req = 1. The port ID requested in the TX direction.
chid	9b	Valid only if req = 1. The channel ID requested in the TX direction.
mfi	10b	Valid only if req = 1. The channel MFI requested in the TX direction.
pl	171B	9 × OH_SEG (each of 19B).
res	2B	Reserved.
CRC	4B	Ethernet CRC.

3.2 Insertion Frame Format Definition

This section describes insertion frame formats for FlexE, OAM, ODU OH, and FlexO.

3.2.1 FlexE Insertion Frame Format

The FlexE framer sends a request once per overhead frame (~104 µs), and the processing unit replies with an Ethernet packet that provides the entire OH for the frame (eight blocks).

Figure 5: FlexE OH Insertion Frame Format

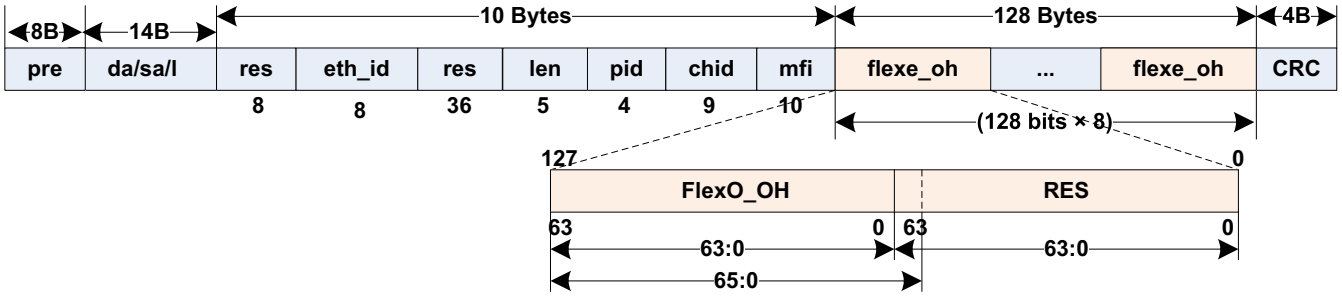

Table 8: FlexE OH Insertion Frame Format

Field Name	Size	Description
da/sa/l	14B	For da (destination address) and sa (source address), the value of the same Ethernet port must be identical and transmitted by the FPGA. The framer and external processor configuration should be consistent. For l (length), the value can vary according to the packet length. In fact, the framer does not care about this field because the relevant length is located in the len field.
res	1B	Reserved.
eth_id	1B	Ethernet packet ID, 0–255 cyclically accumulated. A single *sequence number* is used for the interface (because all packet types traverse through this interface).
res	36b	Reserved.
len	5b	The number of 128b elements in the payload (for FlexE OH 0x8).
pid	4b	The framer's OH/OAM processor port ID (0x6 for NIF FlexE port, 0x7 for backplane FlexE port).
chid	9b	Channel ID. The channel ID value is 0 to 15.
mfi	10b	Multi-frame indication.
pl	128B	The payload field carries eight OH blocks (corresponds to MFI, MFI + 1, MFI + 2...MFI + 7). There are a total of 8 × 128-bit slices. Each slice has 66b OH block + 62b reserved.
CRC	4B	Ethernet CRC.

3.2.2 OAM Insertion Frame Format

SPN OAMs (except of BAS) are generated by the processing unit with the processing unit's responsibility; however, MTN OAM is generated by the processing unit only as a response to a prior request from the processing unit.

NOTE: SPN BAS OAM is generated by the framer itself.

Figure 6: FlexE Client OAM/ALM Insertion Frame Format

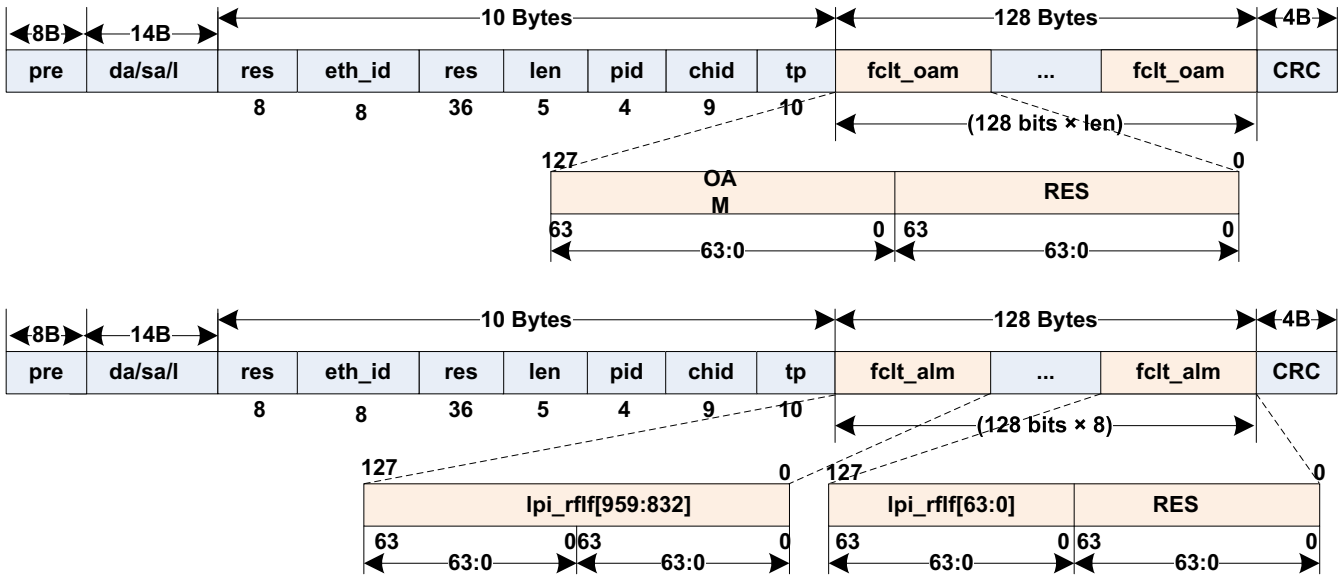

Table 9: FlexE Client OAM/ALM Insertion Frame Format

Field Name	Size	Description
da/sa/l	14B	For da (destination address) and sa (source address), the value of the same Ethernet port must be identical and transmitted by the FPGA. The framer and external processor configuration should be consistent. For l (length), the value can vary according to the packet length. In fact, the framer does not care about this field because the relevant length is located in the len field.
res	1B	Reserved.
eth_id	1B	Ethernet packet ID, 0–255 cyclically accumulated. A single *sequence number* is used for the interface (because all packet types traverse through this interface).
res	36b	Reserved.
len	5b	The number of 128b elements in the payload (valid values are 0x1 to 0x10 [decimal 1 to 16]).
pid	4b	The framer's OH/OAM processor port ID (0x8 for a NIF FlexE port, and 0x9 for a backplane FlexE port).
chid	9b	Channel ID. The channel ID value is 0 to 319.
tp	10b	The 10-bit type indicator. tp[9] = 0 indicates the packet carries OAM, and tp[7:0] indicates MFI. tp[9] = 1 indicates the packet carries a 3-bit alarm.

Table 9: FlexE Client OAM/ALM Insertion Frame Format (Continued)

Field Name	Size	Description
pl	128B	For payload, either the 3b alarm for all channels or 64b OAM blocks. 128b units applies here as well. For alarms, 3 × 320 bits + 64b reserved of the last 128b unit. For OAM, each 128b unit has 64b OAM block and 64b reserved. For SPN, all OAMs are sent from the processing unit except for BAS OAM, which is generated in the framer. For MTN, 64b OAM code block include BABLBABL according to the MTN standard as response to request from framer.
CRC	4B	Ethernet CRC.

The bit position of the various alarms per client are shown in the following table.

NOTE: For non-configured channel numbers, all alarms are zero.

Table 10: FlexE Client ALM Insertion Bit Position

Field Name	ALM	Bits	Corresponding Channels
lpi_rflf[959:640]	cs_rf	[959:896]	ch63 to ch0
		[895:832]	ch127 to ch64
		[831:768]	ch191 to ch128
		[767:704]	ch255 to ch192
		[703:640]	ch319 to ch256
lpi_rflf[639:320]	cs_lf	[639:576]	ch63 to ch0
		[575:512]	ch127 to ch64
		[511:448]	ch191 to ch128
		[447:384]	ch255 to ch192
		[383:320]	ch319 to ch256
lpi_rflf[319:0]	cs_lpi	[319:256]	ch63 to ch0
		[255:192]	ch127 to ch64
		[191:128]	ch191 to ch128
		[127:64]	ch255 to ch192
		[63:0]	ch319 to ch256

3.3 Protection Switching Message

3.3.1 66b Switch

Protection is done through the 66b switch. See Figure 7, 66b Switch Block Diagram.

In the diagram, FlexE#1 is for FlexE interface clients, FlexE#2 is for segment and reassembly (SAR) clients, MAC#1 is for L1 Ethernet for IEEE 1588 insertion and extraction, and BusB is for the terminated (to packets) clients.

Each input through software can be to any one of the streams (example A1, A2) or to both streams. MAC#2 can be to only one of the streams.

The switch itself is fully non-blocking, and each output port can use either round-robin (RR) or weighted round-robin (WRR) between the different FIFOs it holds.

Figure 7: 66b Switch Block Diagram

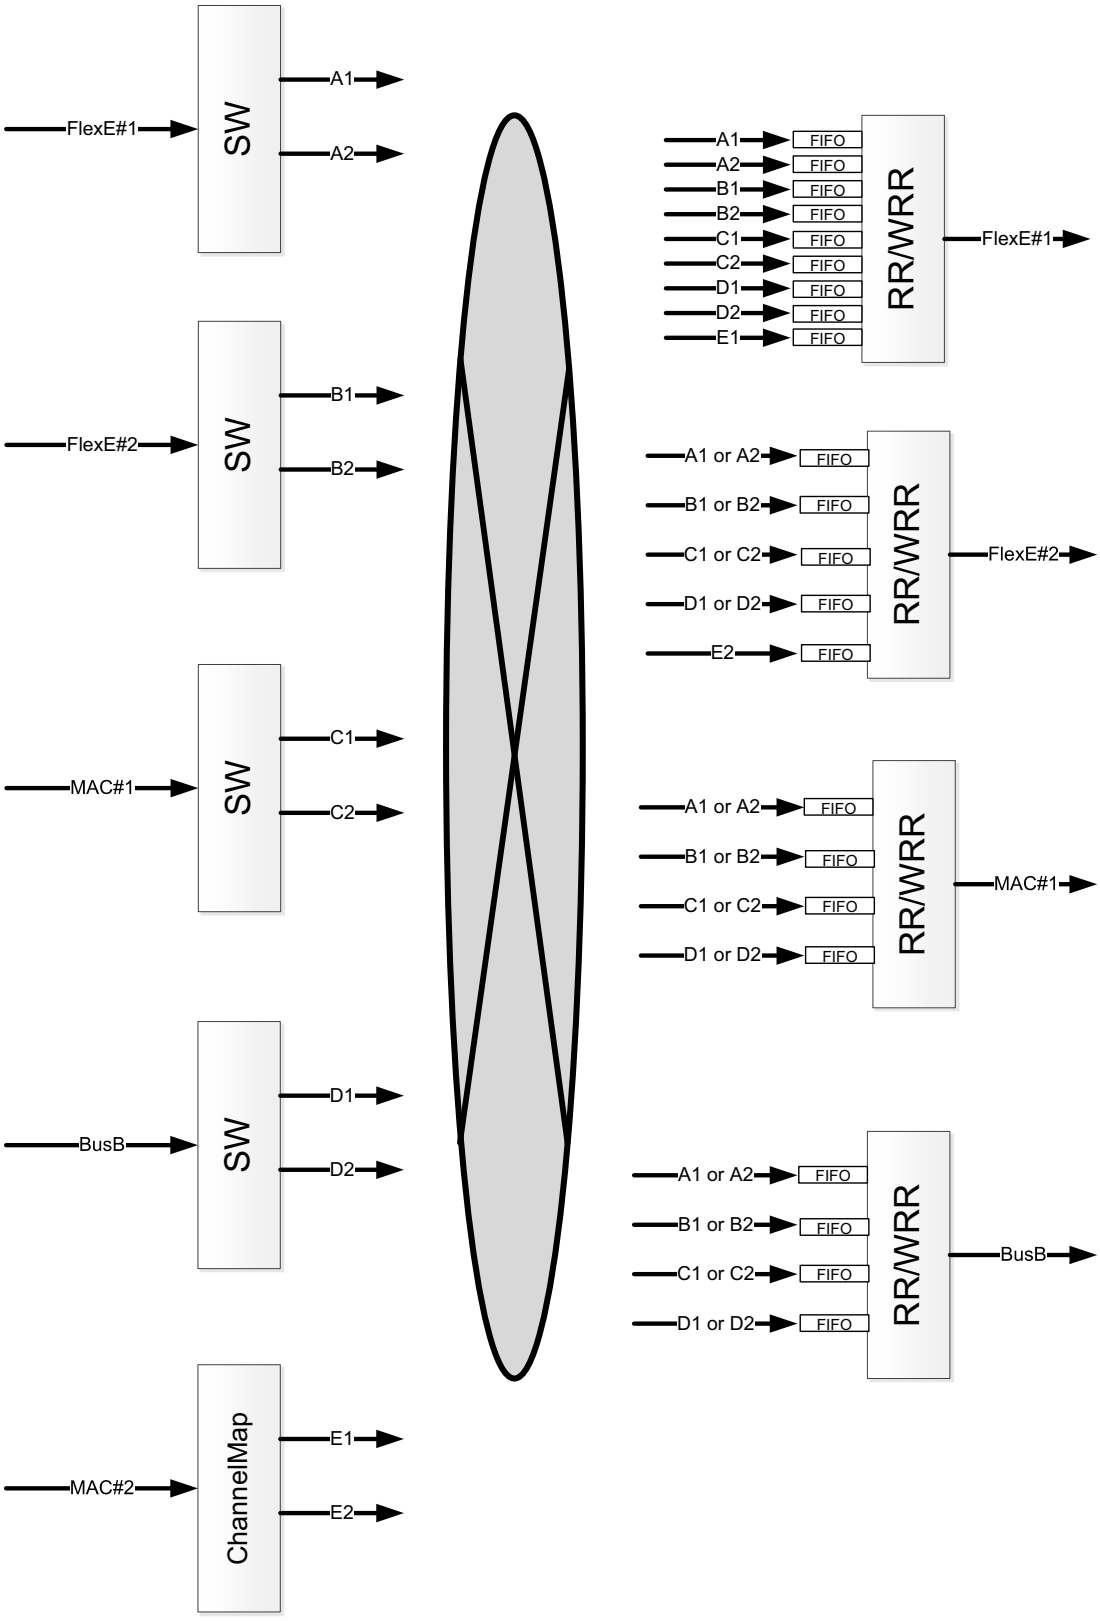

3.3.2 B66_sw Protection Switching

The 66B_SW module must accept the protection switching command sent from the processing unit. The structure of the protection-switching message is shown in the following figure.

Figure 8: b66_sw Frame Format

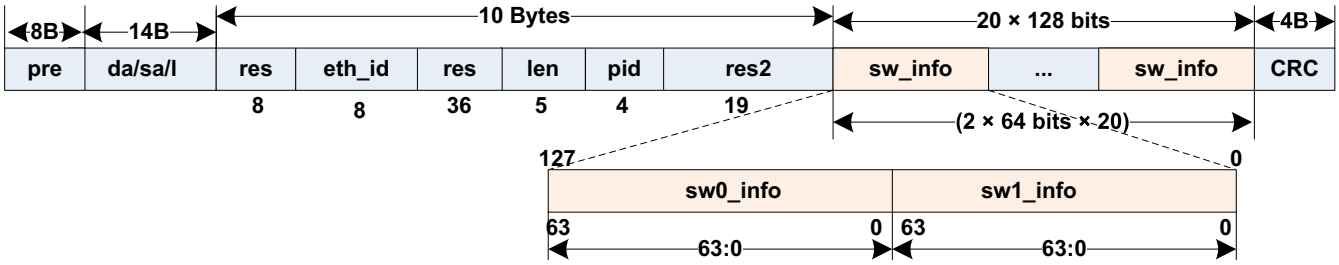

Table 11: b66_sw Frame Format

Field Name	Size	Description
da/sa/l	14B	For da (destination address) and sa (source address), the value of the same Ethernet port must be identical and transmitted by the FPGA. The framer and external processor configuration should be consistent. For l (length), the value can vary according to the packet length. In fact, the framer does not care about this field because the relevant length is located in the len field.
res	1B	Reserved.
eth_id	1B	Ethernet packet ID, 0–255 cyclically accumulated. A single *sequence number* is used for the interface (because all packet types traverse through this interface).
res	36b	Reserved.
len	5b	The number of 128b elements in the payload (For FlexE 66b SW protection switching frames, this field is a fixed value of 0x14).
pid	4b	The framer's OH/OAM processor port ID (the value is 0xa for software protection switching frames).
res2	19b	Reserved.
pl	320B	Protection switching command, which takes 128 bits as the base unit. There are 20 ×128 bits in total.
CRC	4B	Ethernet CRC.

The protection switching information is 2560 bits in total and identifies the enable information (active high) of 1280 channels.

The bit to channel correspondence is as follows:

- Bits 1–0 correspond to FLEXE#1-CH0, bits 3–2 correspond to FLEXE#1-CH1,..., and bits 639–638 correspond to FLEXE#1-CH319.
- Bits 641–640 correspond to FLEXE#2-CH0, bits 643–642 correspond to FLEXE#2-CH1,..., and bits 1279–1278 correspond to FLEXE#2-CH319 (FlexE2 refers to SAR clients).
- Bits 1281–1280 correspond to MAC#1-CH0, bits 963–962 correspond to MAC#1-CH1,..., and bits 1919–1918 correspond to MAC#1-CH319 (MAC#1 refers to L1 Ethernet for IEEE 1588 insertion and extraction).
- Bits 1921–1920 correspond to BUSB-CH0, bits 1923–1922 corresponds to BUSB-CH1,..., and bits 2559–2558 correspond to BUSB-CH319.

Table 12: Arrangement of the Protection Bits

sw0_info	sw1_info	sw2_info	sw3_info	...	...	sw39_info
en[63:0]	en[127:64]	en[191:128]	en[255:192]	...	...	en[2559:2496]

Chapter 4: SPN OAM Supported Features and Format

The following OAM features are supported:

- Connectivity testing (CC)
- Connectivity verification (CV)
- Bit error detection (BIP)
- Remote error indication (REI)
- Remote defect indication (RDI)
- Delay measurement (1DM/2DM)
- Protection switching (APS)
- Customer signal type (CS)
- Customer signal failure indication (CSF)

For more information, refer to the BCM88480 (Q2a) *FlexE Overhead, OAM, and IEEE 1588 Handling* application note (88480-AN3xx).