POPI Sideband MicroFrame Format Proposal

Contribution to IEEE P802.3dq Pin-Optimized PHY Interface Task Force

Jason Potterf, Cisco

2026-04-28

Special thanks to Brian Mowad, Brian Murray, Jacobo Riesco, and Peter Jones for early feedback on this concept

Baseline - MDIO

MDIO – Clause 22

- 64 bits per Management Frame
  - 32-bit Preamble (Optional)
  - 2-bit Start of Frame
  - 2-bit Operation Code
  - 5-bit PHY Address
  - 5-bit Register Address
  - 2-bit Turnaround Time
  - 16-bit Data

Table 22–12—Management frame format

	Management frame fields							
	PRE	ST	OP	PHYAD	REGAD	TA	DATA	IDLE
READ	1...1	01	10	AAAAA	RRRRR	Z0	DDDDDDDDDDDDDDDD	Z
WRITE	1...1	01	01	AAAAA	RRRRR	10	DDDDDDDDDDDDDDDD	Z

Key Observations

- Limited Address Space
  - 2^5=32 PHYs
  - 2^5=32 Registers
- Address in every transaction
- Multi-drop bus necessitates PHY Address

MDIO –Clause 45

- 64 bits per Management Frame
  - 32-bit Preamble (Optional)
  - 2-bit Start of Frame
  - 2-bit Operation Code
  - 5-bit Port Address
  - 5-bit Device Address
  - 2-bit Turnaround Time
  - 16-bit Address/Data

Table 45–426—Extensions to management frame format for indirect access

	Management frame fields							
Frame	PRE	ST	OP	PRTAD	DEVAD	TA	ADDRESS / DATA	IDLE
Address	1...1	00	00	PPPPP	EEEEE	10	AAAAAAAAAAAAAAAA	Z
Write	1...1	00	01	PPPPP	EEEEE	10	DDDDDDDDDDDDDDDD	Z
Read	1...1	00	11	PPPPP	EEEEE	Z0	DDDDDDDDDDDDDDDD	Z
Post-read-increment-address	1...1	00	10	PPPPP	EEEEE	Z0	DDDDDDDDDDDDDDDD	Z

Key Observations

- Same PHY Addressability
- Expanded Address Space for Each Port
  - 2^5=32 Devices * 2^{16} Addresses = 2^{21} (~ 2 Million) Registers
- Address in separate transaction

Achievable Clause 22 / Clause 45 Bandwidth

Worst-Case Bus Configuration

- 64 bits per Management Frame
- Management Data Clock (MDC)
  - Max 2.5 MHz per 802.3

Best-Case Bus Configuration

- 32 bits per Management Frame
  - Enable Preamble Suppression
- Management Data Clock (MDC)
  - > 10 MHz seen in field

Worst-Case Bandwidth @ 2.5 MHz

- 39,062.5 frames/s
- 625 kb/s

Best-Case Bus Transfer Rate @ 10 MHz

- 312,500 frames/s
- 5,000 kb/s

Common MDIO Data Flows

Payload Type	Typical Payload Length	Typical Transaction Frequency	Max Latency Tolerated
General MDIO Register Access	16 Bits	Bursty	Usually not latency sensitive
IEEE 1588 Timestamp Retrieval for Transmit	16 Bits – 16 Bytes	~10 Packets/s	On order of 10 ms for timestamp-related operations
IEEE 1588 Timestamp Retrieval for Receive	16 Bits – 16 Bytes	Every Packet Possible	On order of 10 ms for timestamp-related operations
Firmware Downloads	Kilobytes	Once at boot or upgrade	Affects power-loss recovery time
MACSEC (Re-)Key Exchange	256 Bits	On re-key, interval can range from seconds to minutes	Tight re-key timing budget must be met

Key Observations

- Timestamp retrieval drives MDIO latency
- MACSEC Re-Key Exchange drives bus speed

Problems to Solve

POPI Sideband Channel Challenges

- Challenging mix of data flows mirror Ethernet traffic
  - Long, speed sensitive flows want to be delivered as quickly as possible
  - Short, latency sensitive flows can’t wait for the long transactions to finish
  - IP networking constructs can solve this, but are overly complicated for this application
    - No desire to design a QoS scheme
    - No desire to reinvent TCP
- Per-port traffic patterns vary within a four-port group that shares a sideband channel
  - One port should not starve adjacent ports
  - Strict striping/TDM schemes waste bandwidth

Proposed Solution

POPI Sideband MacroFrames / MicroFrames

- Sideband transmits larger MacroFrames as a sequence of fixed length 8-byte MicroFrames
- Microframes have enough information to rebuild the MacroFrame
- Supports MDIO Management Frames in a single MicroFrame
- Allows preemption within a single MicroFrame for a maximum latency of one 8-byte MicroFrame

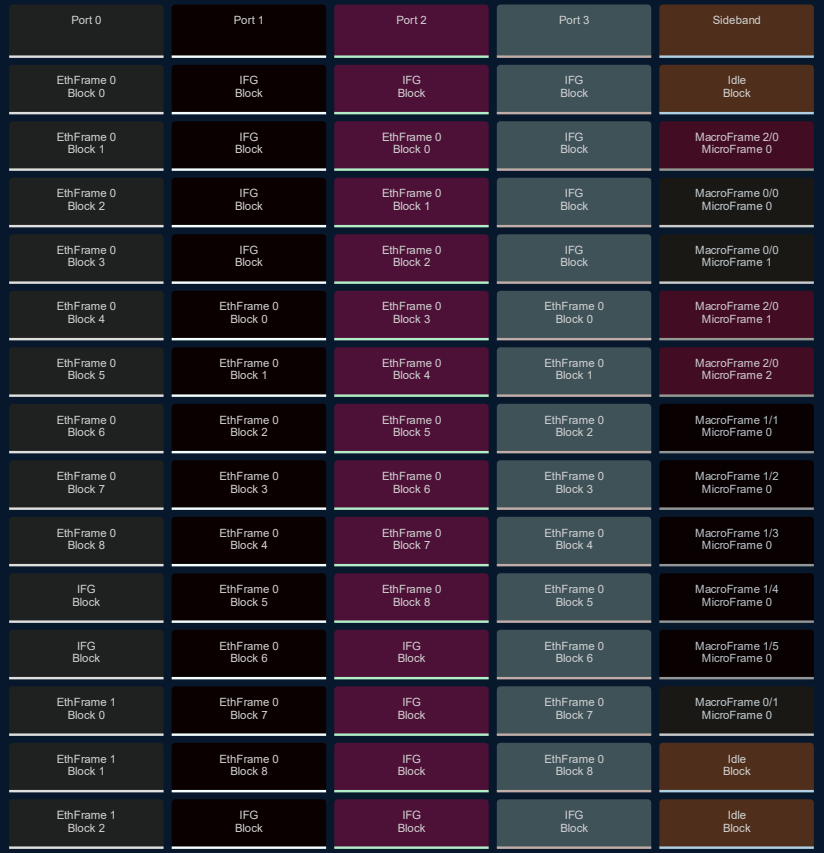

POPI Sideband MacroFrames / MicroFrames

- 64 bits per MicroFrame, up to 256 MicroFrames per MacroFrame
- 8-bit MacroFrame Type
  - Basis for carrying unlimited diverse payload formats
- 2-bit Port Number
  - Assumes fixed 4:1 Port:Sideband ratio
- 6-bit MacroFrame ID
  - When combined with Port Number, provides primary key to reassemble a chain of MicroFrames to form a MacroFrame
- 8-bit MacroFrame Sequence Number
  - Ensures ordered delivery of MicroFrames
- 32-bit Payload
  - Payload defined by the MacroFrame Type
- 8-bit Checksum
  - Protect MicroFrame integrity

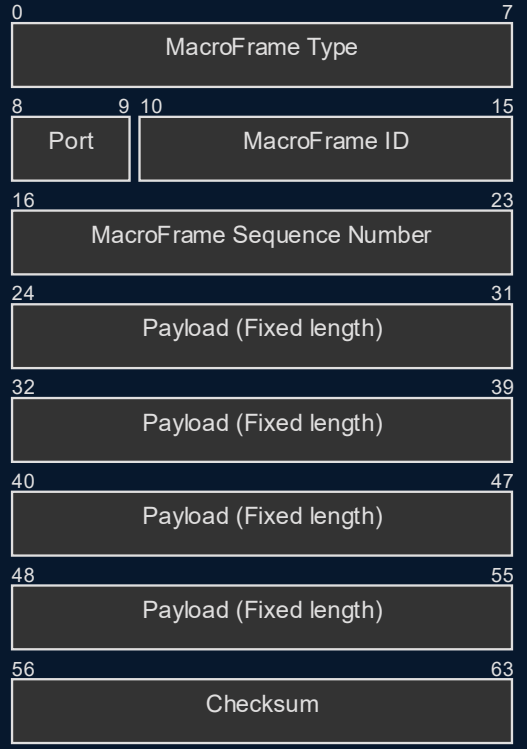

POPI Side Channel MicroFrame

Additional Benefits

- MacroFrames can be used to
  - Fragment a large transaction while ensuring context is maintained for reassembly
  - Create a “stream” of data to or from a particular endpoint in a PHY
- MicroFrames occupy the same worst-case footprint as an MDIO Management Frame but provide twice the payload, resulting in a net increase in available bandwidth
- In the worst-case 2.5 Mbit/s sideband configuration, which occurs at 25% of the bandwidth of a single 10 Mbit/s port, the proposed MicroFrame scheme provides performance that meets or exceeds the best-case 802.3 compliant MDIO bus
- The inclusion of a 2-bit port designator allows idle ports to cede sideband bandwith to active ports, as opposed to a fixed per-port sideband data schedule
- MacroFrame Type ID can be extended within the Payload to create a virtually unlimited number of sub-types

Questions and complaints about the chosen names