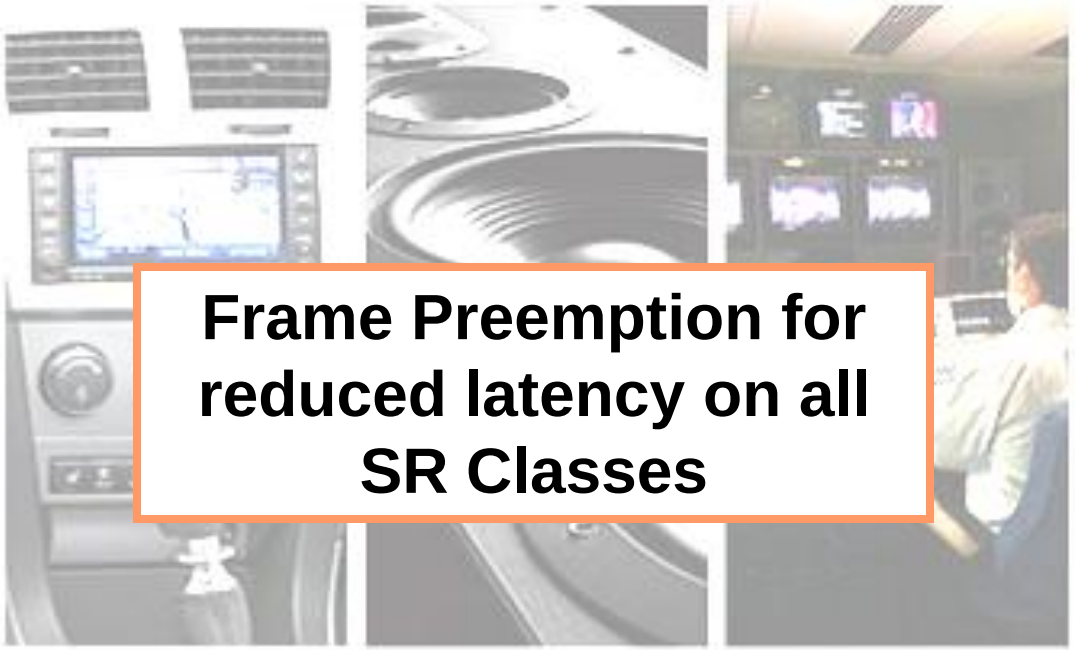

# Frame Preemption for reduced latency on all SR Classes

**Craig Gunther** (craig.gunther@harman.com)

May 2011 – Santa Fe interim (Updated 19May2011)

# Acknowledgements

Reference materials:

1. new-kim+goetz-Ultra-Low-Latency-Switching-v5.pdf
2. ba-kw-stream-latency-Improvements-0311.pdf
3. ba-pannell-latency-math-1110-v5.pdf
4. ba-boiger-per-hop-class-a-wc-latency-0311.pdf

# Introduction

- Ultra-low latency[1] & Preemption[1] are two separate topics
- The focus of this presentation is not ultra-low latency, but to explore the benefits of preemption on existing AVB Classes
- **Goal**: Make preemption available to all AVB shapers

# Definitions

- **Interfering Traffic (IT):** frames of a lower priority which cause delays to transmission of higher priority frames.
- **Preemption:** Suspending transmission of a lower priority frame so a higher priority frame can be transmitted, followed by resumption of the lower priority frame. This can occur more than once to a large low priority frame.

# Multi-Class Preemption (with separate UL frames)

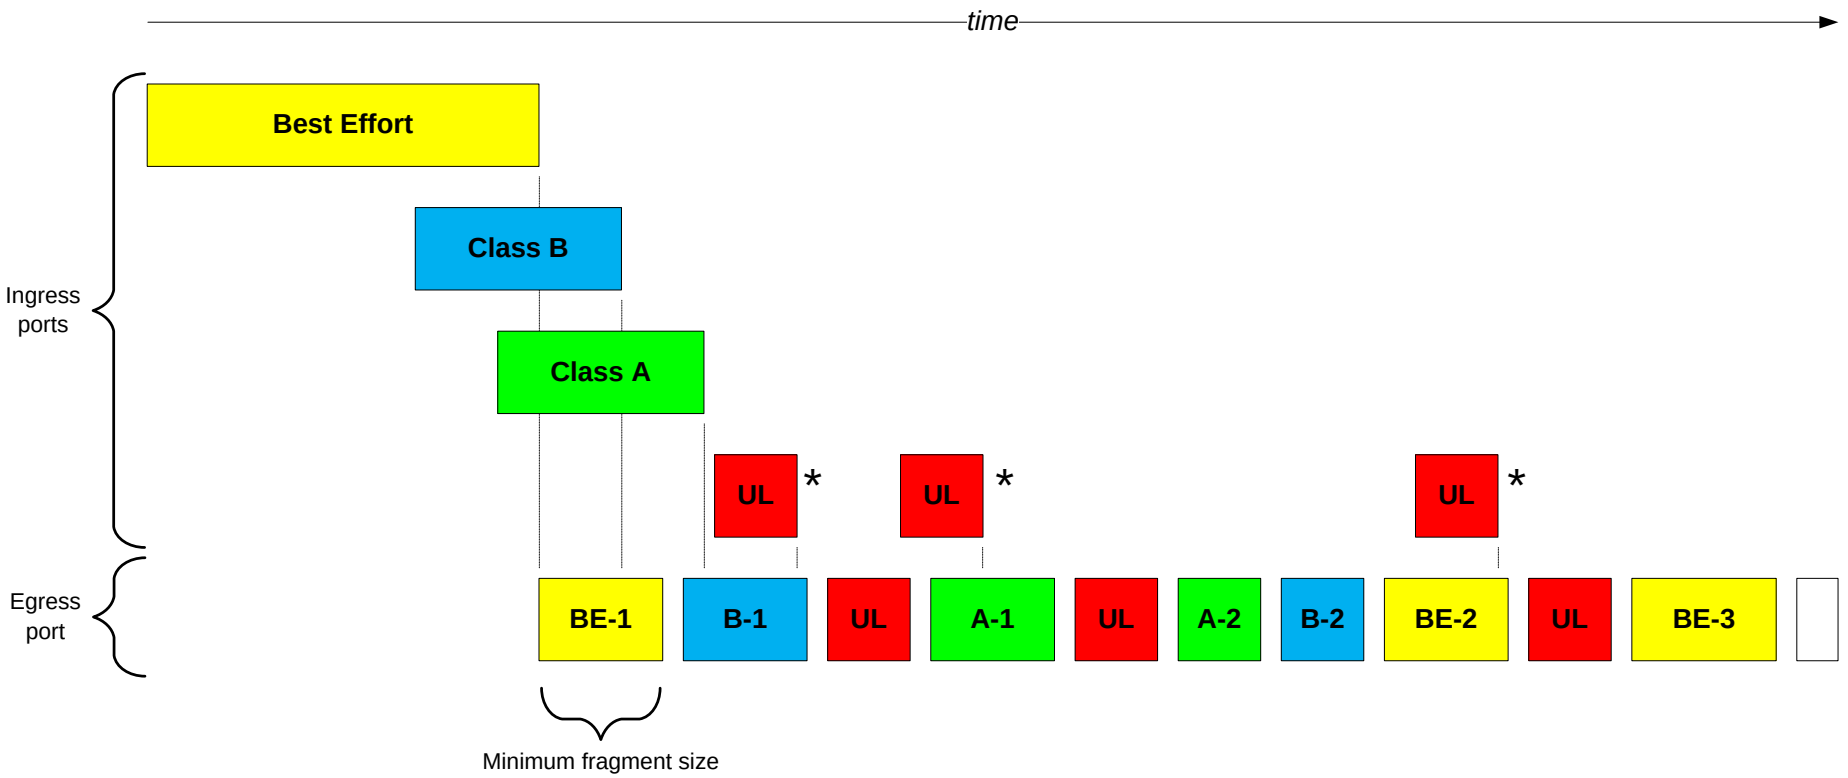


*Note: This slide assumes UL frames are separate from SR Class A & B frames

# Multi-Class Preemption (with UL in SR Class A)

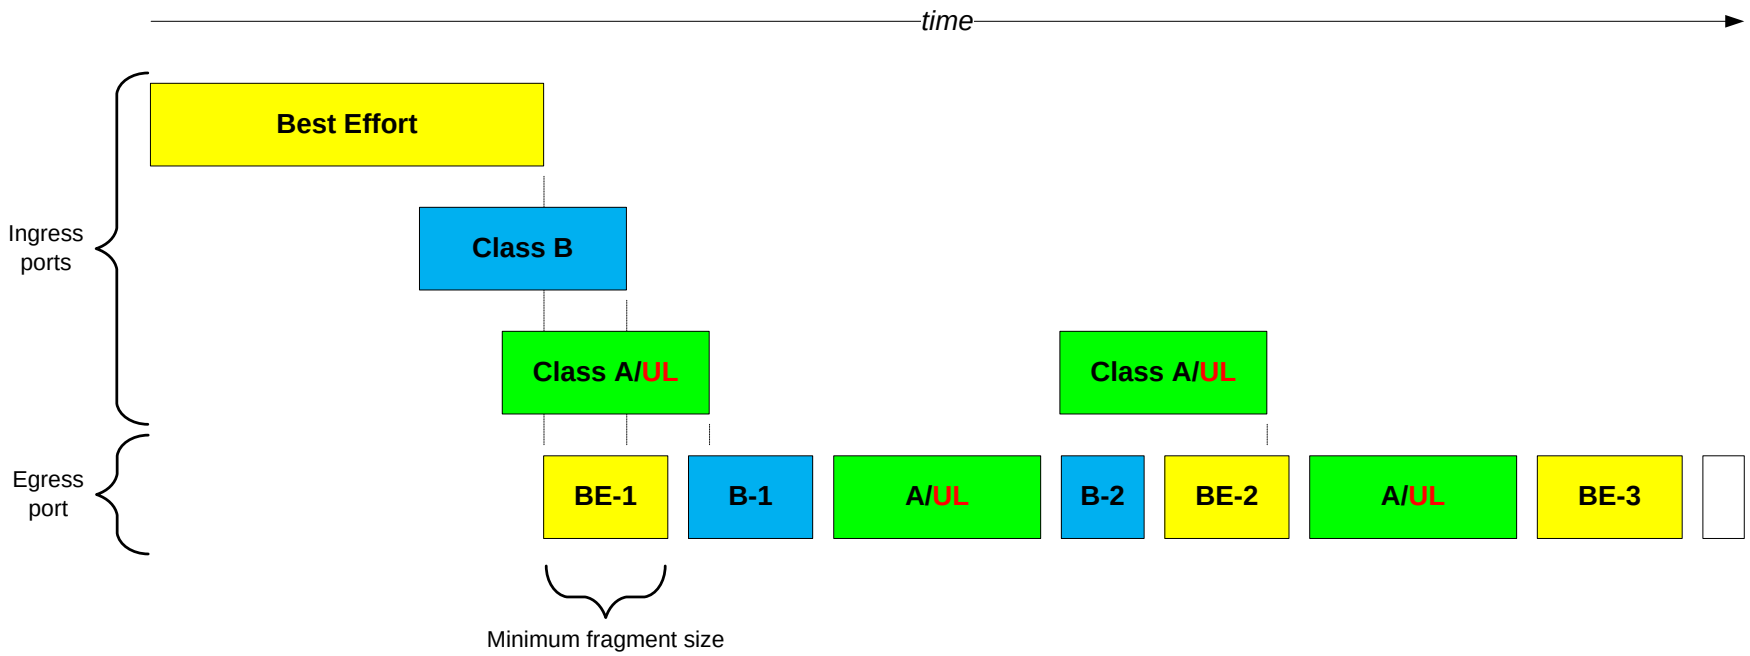

# Possible Preemption Marking[1]

- After peer Gen-2 devices agree they can do Preemption (via LLDP?) they know every packet sent between them has a new 8-bit header that defines characteristics of each preempted frame
- 8-bit header contains:
  - 2 flag bits: begin, previousEnd
  - Three preemption classes requires 2 bits to identify
    - Note: two preemption classes would require a 1-bit field
  - 4-bit sequence number per preemption class

# Reassembling the pieces

- If “previousEnd” bit is set then previous frame has been completely reassembled; pass it on
  - Sequence numbers can be used to detect missing pieces. Note that there are only 16 sequence numbers so this can fail if there are 16 missing pieces in a row.
- If “begin” bit is set then reset the class reassembly buffer pointer to the beginning
- Append piece to per-class reassembly buffer

# Multi-Class Preemption Concerns

- One Ingress buffer for each preemption class
  - Class A and Class B buffers are limited size
  - Best Effort buffer must support Jumbo frames
- MACsec, etc, concerns?
- Effects on PHY/MAC/CAM?
- Will 8-bit preemption header work?
  - Note: an alternative was suggested that would use a 32-bit header with 16 bits used for a new EtherType

# Multi-Class Preemption Benefits

- Jumbo frames are back!
- Talker burst limit of two back-to-back frames[4]
  - Can we now define a latency formula?
- Gen-1 and Gen-2 switches can co-exist between Talkers and Listeners
  - Obviously preemption (and reduced latency) can only occur between Gen-2 devices
- Reduced latency for higher priority frames

# Bridge Port Latency Math with Preemption[3]

Max Latency = $t_{Device} + t_{Interval} + t_{MaxFrameSize} + t_{Stream} - t_{(Stream+Gap)}$*1.333

$t_{Device}$ = 5.12µs

$t_{Interval}$ = 125µs

$t_{MaxFrameSize}$ = 6.72µs (for 64 bytes + IFG + preamble), 7.68µs (for 96 bytes + IFG + preamble)

$t_{Stream}$ = 5.12µs (assuming 64-byte frames)

$t_{(Stream+Gap)}$ = 5.12µs + 1.6µs

Max $Latency_{100\ MB/s}$ = 5.12µs + 125µs + 6.72µs + 5.12µs - ((5.12µs + 1.6µs) * 1.333) = **133.00µs**

Max $Latency_{1000\ MB/s}$ = 0.512µs + 125µs + 0.672µs + 0.512µs - ((0.512µs + 0.16µs) * 1.333) = **125.80µs**

| | 100 MB/s [x7] | 1000 MB/s [x7] |
|---|---|---|
| Without SaR | **249.64µs [1747µs]** | **137.46µs [962µs]** |
| With 64-byte SaR | **133.00µs [931µs]** | **125.80µs [881µs]** |
| With 96-byte SaR | **135.56µs [949µs]** | **126.06µs [882µs]** |

# Thanks

# Change history

- v01 original presentation at Santa Fe
- v02 presented on 18May2011 AVB weekly call
  - Slide 2: Update presentation references
  - Slide 3: Fragmentation should only be used by UL & SR Classes
  - Slide 5: Fix packet diagram on ingress ports
  - Slide 6: Sequence #s are not optional
  - Slide 8: Address reassembly buffer requirements
  - Slide 9: Jumbo packets are more important than originally thought
- V03 changes suggested on 18May2011 AVB weekly call
  - Don't say "fragmentation" or "SaR", use "preemption" to reduce confusion
  - Slide 2: Number Acknowledgements as footnotes
  - Slide 3: Make goal of presentation more obvious
  - Slide 6: Insert new diagram for UL contained in SR Class A
  - Slide 7: Add note about a 1-bit field for preemption class ID
  - Slide 9: Add note about possible 32-bit header w/EtherType