ALANTEC™

Intelligent Switching Hubs

An Efficient Frame-Tagging Format for VLANs

John Wakerly, CTO
2115 O'Nel Drive
San Jose, CA 95131
1-800-ALANTEC
wakerly@alantec.com

Pros and Cons of VLAN Tagging

- That's not what I'm here to talk about!
- BUT, if you want to do tagging for any purpose, here's a good way to do it

Some Existing Tagging Styles

802.10

- Split open the frame, insert 802.10 tag
- Send the tagged frame, recalculate a new CRC "on the fly"
  - Breaks the integrity of the untagged frame's CRC

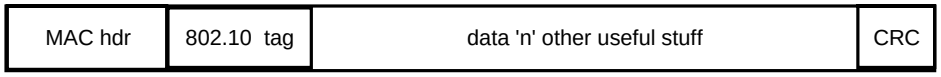

U.S. Patent 5,394,402 (Floyd Ross)

- Tag prepended, doesn't say what to do with the CRC
- If not recalculated, then tagged frame cannot be checked
- If recalculated, same drawback as 802.10
- Otherwise, see next slide

Some Existing Tagging Styles

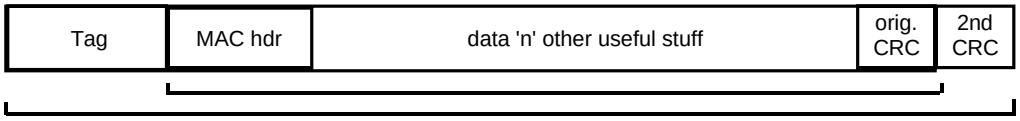

Double CRC

- Prepend tag to untagged frame and append new CRC
- Preserves integrity of untagged frame
- But still has a drawback:
  - Any device (e.g., VLAN-aware server or router) that creates or modifies a tagged frame must calculate two CRCs
  - Typical server or router would have to calculate the first CRC using software
  - First CRC is a function of the entire untagged frame

Proposed Generic Tagging Format

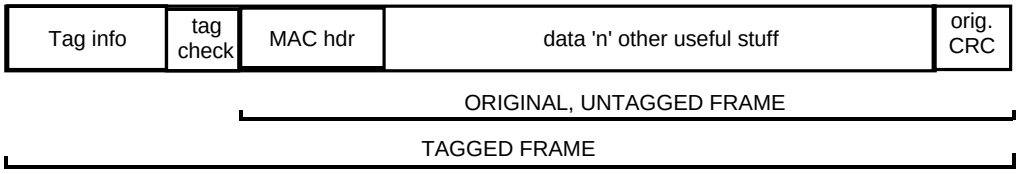

"Tag info" has arbitrary length and content

"Tag check" is same size as an Ethernet CRC (32 bits) and is a function of "Tag Info" only

Is This Possible?

Classical CRC codes are linear codes

- Prepend 0s to a coded message ==> still a coded message
- Append 0s to a coded message ==> still a coded message
- Add two coded message (bit-wise XOR) ==> still a coded message

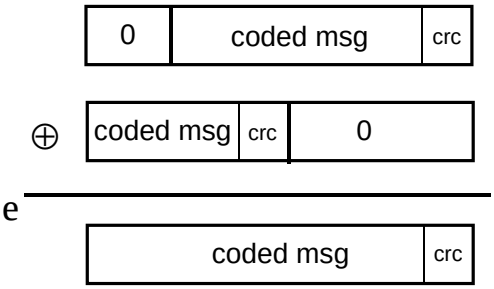

Problem: The Ethernet CRC is a modified classical CRC and is NOT linear

Dealing with the Ethernet CRC

Ethernet CRC modifications

- Complement the first 32 bits of the frame
  - Equivalent to initializing the shift register to FFFFFFFF
- Complement the 32-bit CRC before sending

. . . several pages of math omitted . . .

- You can still make it work!

Tag check equals standard Ethernet CRC, except ...

- Add "magic pattern" to the tag's standard Ethernet CRC
  - Magic pattern is a constant, independent of tag contents or length

Calculating the Tag Check

Using hardware

- Standard Ethernet CRC generator
- Add (bit-wise XOR) magic pattern to calculated CRC
  - Equivalently, initialize the shift register to a backed-up magic pattern instead of FFFFFFFF

Using software

- Bit serial (slow)
- Byte-serial: 256x32 table, one table look-up and 32-bit XOR per tag info byte
- Either way, XOR the result with magic pattern

Benefits of the Proposed Format

The original CRC is not modified in any way

The original CRC is valid for the tagged frame as well as for the untagged frame

Both the original frame and the tagged frame have complete CRC coverage

- Including "software errors" in that inadvertently modify the wrong buffer

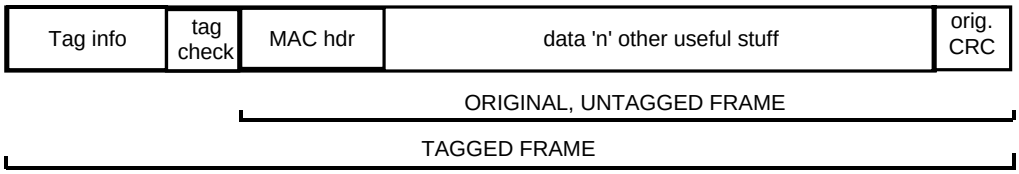

Further Benefits

What about devices that create or modify a tagged frame?

- Tag check is calculated by hardware or software
- New CRC at the end of the frame is calculated "on the fly" by the standard Ethernet CRC generator
  - CRC generator is unaware that it's working on a tagged frame

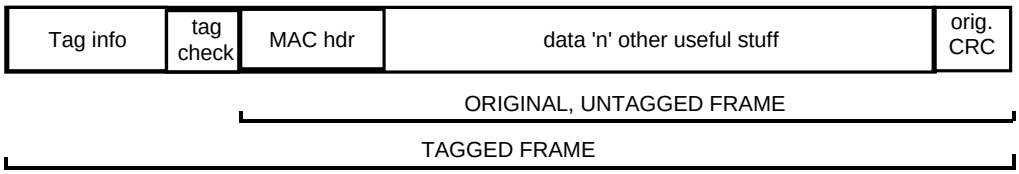

A Few Implementation Notes

Remember, tag check is a function of the tag only

- Tag checks can be pre-computed by software for each different tag value (i.e., per VLAN connection)
- Scatter/gather DMA controllers can re-use the same (fixed) tag for multiple frames following the same path
- Bridge table can contain a pointer (index) to a pre-computed tag

Frames are not split

- No header copying
- Easy to send both tagged and untagged versions of the same frame on hybrid links

Other Issues

Does this really work?

- I did the math
  - Well, that's no guarantee, even though I did publish a book on error-detecting codes in 1978...
- Also, wrote a C program to calculate CRCs and verify consistency of tagged and untagged frames/CRCs
- Math and source code will be included in full write-up

Full write-up

- 85% drafted
- Will post to exploder within a week

Other Issues

Patent application will be filed (U.S. only)

- Strictly to protect against future Soderbloms
- If granted, will license on non-discriminatory, inexpensive terms (i.e., less than it costs to send an engineer to 802.1 for a year)
- Prior art welcome (saves me the $50K-$100K to file and prosecute!)

Any questions?

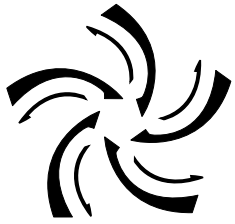

ALANTEC™

Intelligent Switching Hubs