What LinkSec Should Know About Bridges

Norman Finn

What Do Bridges Do? (1)

- **A network of bridges emulates a single fat yellow coax, IEEE 802.3 Clauses 6-8, 10BASE5, good ol' Ethernet.**
  - — **Every frame transmitted is, to a *very* high degree of probability, received by every attached station.**
  - — **Each station filters out the traffic not addressed to that station.**
- **Repeaters (IEEE 802.3 Clause 9) can connect coaxes.**
  - — **There are physical limitations to how big this network can grow.**
  - — **A hubbed network must not have any closed loops, or every frame either 1) collides with itself (if no buffers in the hubs); or 2) loops forever (if buffers).**
  - — **At some point in network growth, it is no longer efficient for every station to receive and filter out every other stations' traffic.**

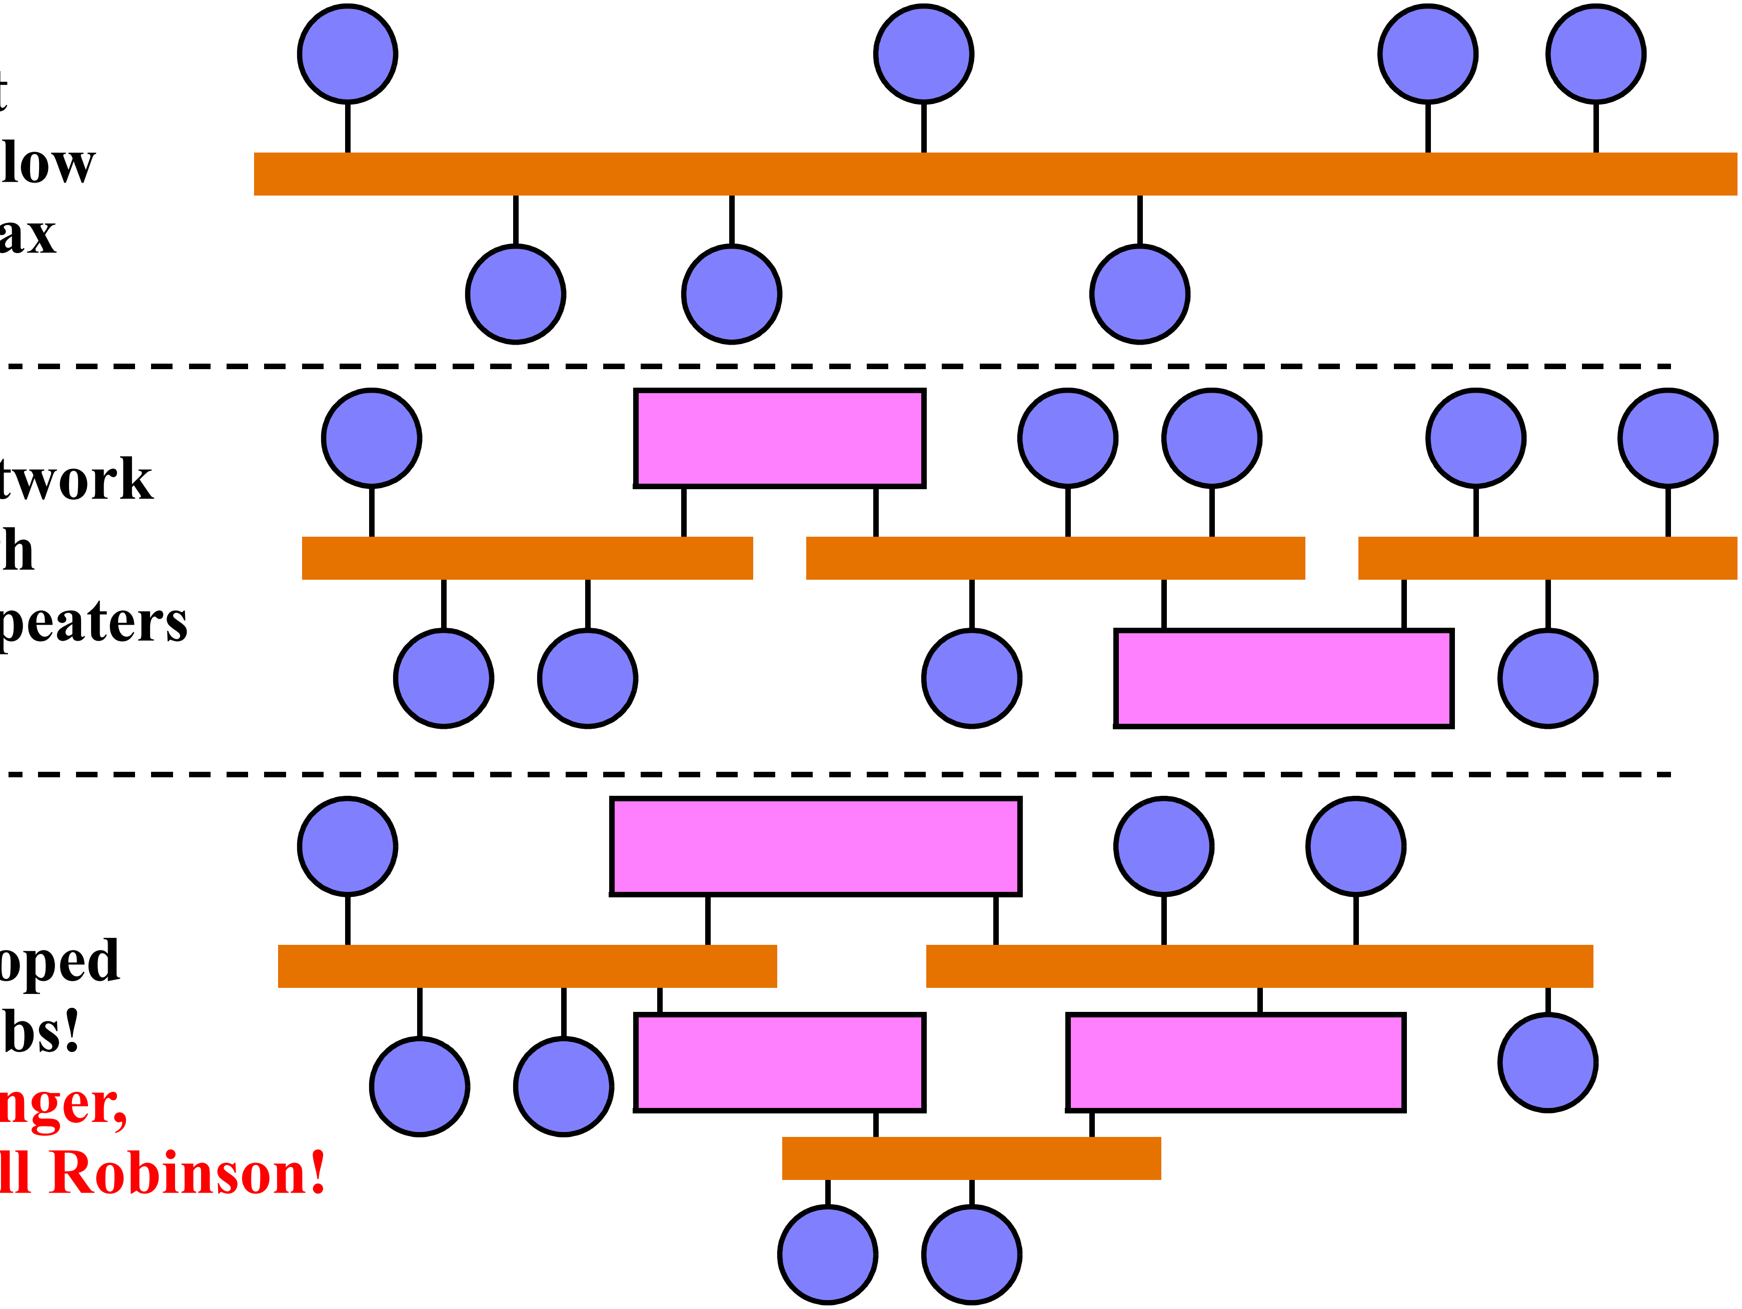

Fat Yellow Coax

Network with Repeaters

Looped Hubs! Danger, Will Robinson!

What Do Bridges Do? (2)

- Bridges connect “LAN Segments”. They are more capable than hubs, and scale to much larger networks.
  - — Bridges run spanning tree protocols so that physical loops do not necessarily cause forwarding loops, and so that the network can recover from failures of bridges or LAN segments.
  - — A bridge tries to avoid delivering any frame to a LAN segment on which it knows the frame is not needed.
  - — A bridge is buffered, and can support ports of different speeds in different parts of the network, i.e. high speed in the “core” and low speed at the “edge”.
  - — Thus, as a network expands, the traffic on any given LAN segment, especially at the network edge, grows more slowly than the total amount of traffic in the network.

In What Ways Does a Bridged Network Fail to Emulate a Fat Yellow Coax? (1)

- **Bridged network is less reliable, frame-by-frame.**
  - **Imagine three point-to-point 10 Mb/s links connecting three stations to one bridge. Two stations on two links pump out 10 Mb/s each, all destined for the third station. In the long run, half of the frames *must* be dropped by the bridge.**

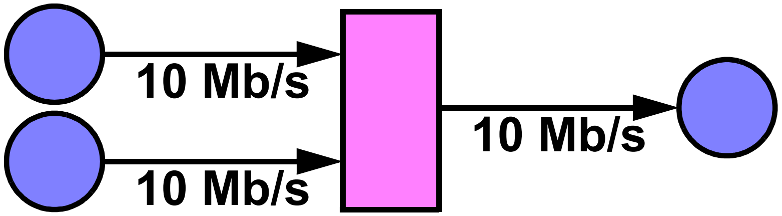

Something's gotta give!

- **Bridged network is more reliable second-by-second.**
  - **Redundant bridges and LAN segments make the network robust against the failure of any one bridge or LAN segment.**
  - **Intelligence in the bridges can make the network robust against the failure of a station, e.g. continuous transmission.**

In What Ways Does a Bridged Network Fail to Emulate a Fat Yellow Coax? (2)

- **A frame with a destination MAC address matching one of the 16 Bridge Reserved Addresses cannot be forwarded through the bridged network.**
- **Two stations on two different LAN segments cannot have the same MAC address in a bridged network.**

How Do Bridges Manage to Do That?

- **They run a number of Layer 2 protocols, including:**
  - — **One of the Spanning Tree Protocols to avoid forwarding loops**
  - — **GVRP, GMRP, and/or IGMP to restrict broadcasts, multicasts, and unicasts to only those LAN segments which need to see them**
  - — **802.1X (not limited to bridges!) to authorize the use of ports**
  - — **Link Aggregation Control Protocol to bundle physical links**
- **They learn source MAC addresses. Forward frames from one physical port to another based on destination MAC address. Make a best effort to guarantee that:**
  - — **No frame reaches a LAN segment that’s known to not need it;**
  - — **No frame is delivered to any LAN segment twice.**
- **They alter/add/strip VLAN tags, but alter nothing else.**

Spanning Tree Protocols

- **STP: Spanning Tree Protocol, IEEE Std. 802.1D-1998 Clause 8. To be omitted from IEEE Std. 802.1D-2003.**
  - **"Spanning Tree Classic", takes tens of seconds to converge.**
- **RSTP: Rapid Spanning Tree Protocol, IEEE Std. 802.1W-2001 Clause 17 or IEEE Std. 802.1D-2003 Clause 17.**
  - **Backward compatible with STP.**
  - **Converges in tens of milliseconds, supports larger networks.**
- **MSTP: Multiple Spanning Tree Protocol, IEEE Std. 802.1S-2003.**
  - **Interoperates with RSTP, backward compatible with STP.**
  - **Ties each VLAN to one of any number of spanning tree instances.**

What do the Spanning Tree Protocols do?

- **STPs create one or more virtual networks, or Spanning Trees, over a network of LAN segments and bridges, by blocking the passage of frames through some ports.**
- **They take up as little bandwidth as possible.**
- **Each resultant virtual network is:**
  - **"Spanning", because it reaches every bridge in the network; and**
  - **a "Tree", because there is only one unblocked path from any LAN segment to any other LAN segment.**
- **They reconfigure as quickly as possible after the failure or restoration of a bridge or LAN segment.**
  - **The loss of a frame is much preferred to its duplication or misordering during a reconfiguration.**

Limiting the Spread of Floods

- GVRP, IEEE Std. 802.1Q-1998 Clause 11, allows bridges to limit the spread of broadcast, multicast, and unicast frames on a particular VLAN to only those bridges that are required to handle them.
  - It limits tagged frames to the bridges with ports configured for that VLAN, and to the transit bridges required to reach them.
- GMRP, IEEE Std. 802.1D-1998 Clause 10, allows bridges to limit the spread of multicast frames.
  - Not all multicasts need be transmitted to all stations.
- IGMP, RFC 2236, is a router-endstation protocol which many bridges snoop and/or proxy to do the same job as GMRP.

Port Control

- IEEE Std. 802.1X-2001 specifies how an Authenticator (a bridge, router or any other device) may employ an external (or resident) Authentication Server to control the access by a Supplicant device to a point-to-point connection to the Authenticator.
- Many bridges extend this standard to individually authenticate particular MAC addresses on a port.

Link Aggregation

- IEEE Std. 802.3-2002 Clause 43 specifies how two or more point-to-point 802.3 links may be combined so as to behave much as if they were a single, faster, link.
- LACP, the Link Aggregation Control Protocol, allows such aggregations to be constructed automatically.
- Bridges often employ Link Aggregation, but the technique is applicable to any pair of devices.

Learning, Forwarding, and Forgetting MAC addresses

- A bridge learns the source MAC address and arrival port of each frame received for forwarding.
- A bridge forwards each frame according to the destination MAC address of the frame, using the information learned from source MAC addresses.
  - If the address is unknown, bridge floods the frame to all ports.
  - Note that the source MAC address does not control forwarding!
- A bridge forgets any MAC address after it goes unused for a while. It forgets certain (usually, many) MAC addresses immediately when a significant change occurs in the Spanning Tree Topology.

802.1Q VLAN Tags

- **IEEE Stds. 802.1Q-1998 (and 802.3-2002 Clause 3.5) define VLAN tags.**
  - **An 802.1Q tag immediately follows the source MAC address of an 802.3 frame.**
  - **The 802.1Q tag contains an EtherType (0x8100), a 3-bit priority field, a 12-bit VLAN ID, and a 1-bit CFI bit.**
- **Before transmitting a received frame, a bridge may:**
  - **Add an 802.1Q tag to an untagged frame; supply a non-0 value for a 0 VLAN ID; or completely remove an 802.1Q tag; and/or**
  - **Alter the priority field in any way.**
- **A bridge may not:**
  - **Alter a non-0 VLAN ID; nor emit VLAN ID 0.**

Allowed transformations

Data	S	D		Data	S	D	
Data	S	D		Data	P=2, VID=8	S	D
Data	P=3, VID=8	S	D	Data	P=3, VID=8	S	D
Data	P=3, VID=8	S	D	Data	P=0, VID=8	S	D
Data	P=1, VID=0	S	D	Data	P=4, VID=9	S	D
Data	P, VID = any	S	D	Data	S	D	

Disallowed transformations

Any data frame				Data	P=any, VID=0	S	D
Data	P=1, VID=122	S	D	Data	P=1, VID=33	S	D

What is the Difference Between a "Bridge" and a "Switch"?

- **This author's answer:**
  - **Your company, *<insert some company name, not yours, here>*, makes those old-fashioned Bridges and Routers. [Boo! Hiss!]**
  - **My company, *<insert your company name here>*, makes those sexy new Switches. [Hooray!]**
- **A less judgemental answer:**
  - **A switch combines features of both bridges and routers in a manner that is indistinguishable, from a signaling and data forwarding standpoint, from some combination of standard bridges and routers.**
- **The answer required of a standards developer:**
  - **There is no such thing as a switch.**

Bridges and IEEE 802.10C-1998 (1)

- **IEEE Std. 802.10C-1998 Annex 3A defines a bridge-to-bridge model that may not be useful.**
  - — **Secure bridges are connected to secure endstations via secure LANs not using 802.10.**
  - — **Secure bridges protect insecure devices, carrying their traffic over insecure bridges. Secure and insecure bridges run “separate spanning trees”, i.e. secure STP runs over insecure network.**
  - — **The secure bridges form an overlay network over the insecure bridges. Secure bridges form 802.10 security associations with each other through the insecure network.**
  - — **Secure bridges use “Probe” (and 802.2 TEST) frames to figure out, for any given frame that must cross the insecure network, which secure bridge will handle that frame, and hence which security association to use.**
  - — **After a topology change, many addresses may need re-probing!**

Bridges and IEEE 802.10C-1998 (2)

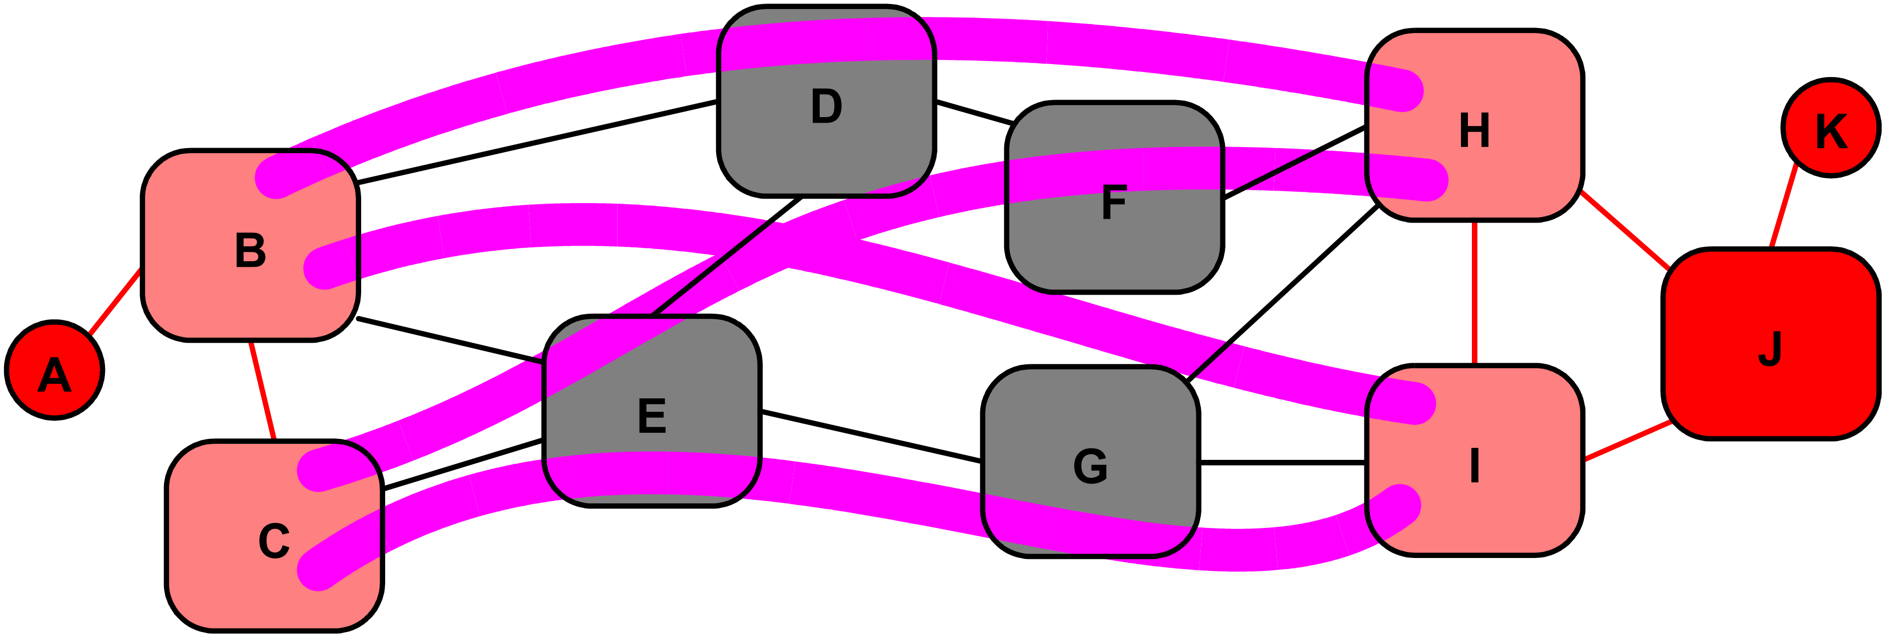

- Secure (pink) bridges B, C, H, I create a mesh of Security Associations across insecure (gray) bridges to securely carry traffic for insecure devices (red).

  — Topology change among H, I, J can require B to rethink which SA to use to reach K. Hence, Probe frames.

  — A group security association is also needed for multicast frames.

What is Authenticated (1)?

- The 802.10 Integrity Check Value (ICV) protects most of the frame:

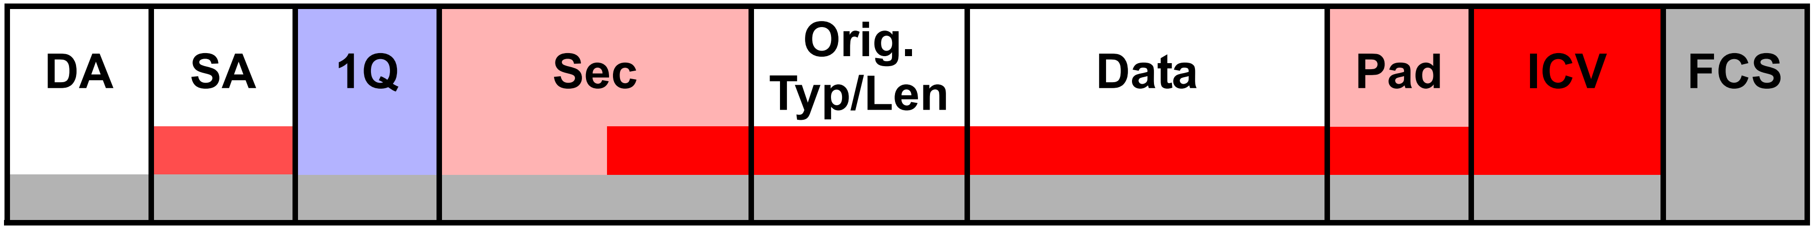

— (The ICV covers a copy of the SA placed in the Security header.) VLAN tag is not mentioned in 802.10, but it would follow the SA.

- Bridge-to-bridge LinkSec needs the ICV to protect all of the frame, including the 802.1Q VLAN tag:

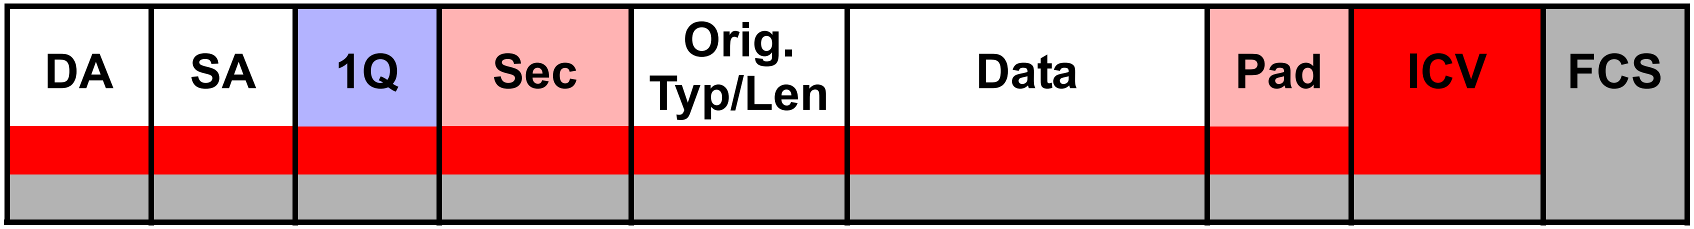

— The copy of the source MAC address would then not be needed.

What is Authenticated (2)?

- But, because 802.1Q bridges may add or delete an 802.1Q tag, end-to-end security associations transported through bridges must not authenticate the 802.1Q tag:

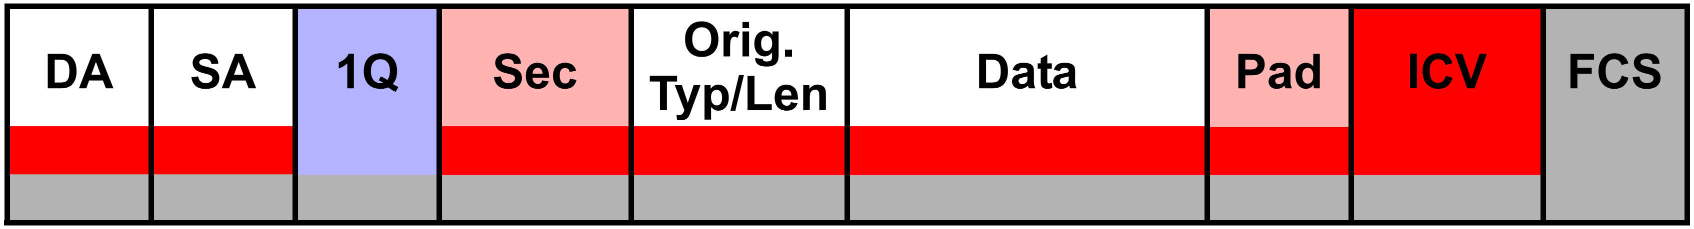

DA | SA | 1Q | Sec | Orig. Typ/Len | Data | Pad | ICV | FCS