Changes to 802.1Q necessary for 802.1Qbz (bridging 802.11 media)

Norman Finn

March, 2013 v01

Summary of changes

Significant changes

- (We will not list minor changes, such as adding “Access Point” to the definitions clause.)
- New conformance clause 5.4.1.8: Combined Access Point and VLAN Bridge.
  - Uses 802.1AC mapping of Bridge Port to Security Association, not Portal.
  - Reference Clause 44, it that is included in 802.1Q.
- 6.5.6 Frame Lifetime
  - End stations sleep, so frame lifetime can be exceeded.
  - Bridges are not allowed to sleep
- 6.7 Support of the ISS by specific MAC procedures
  - This clause has been moved to 802.1AC. Replace with placeholder.
- 6.9.1/6.9.2 Support of the EISS: See “Tagging format” slides

Significant changes

- 7.5 Locating end stations
  - 802.11 non-AP stations are discovered when they attach to the AP.
- 8.6.4 Egress
  - Just add a note mentioning that port ≠ Queue Set
- 8.6.5 Queuing frames: see “Queue Sets” slides
- 8.6.7 Queue management
  - An 802.11 station (AP or not) *can* retry the transmission of a frame.
- 9.4 Tag Protocol Identifier (TPID) formats
  - Mention that 6.9.1/6.9.2 can alter the MSDU when adding/removing tags.
- 34.4 Deriving actual bandwidth …
  - Mention that one frame can be replicated several times in a queue if a single Queue Set serves multiple Ports.

Significant changes

- Globally, but especially in 37 Enhanced Transmission Selection
  - Distinguish between “Queue Set” and “Port,” and use the right term.
- 37.3 ETS algorithm
- Add new Clause 44: see “Clause 44” slides
- Annex A: PICS
  - Update PICS with new requirements
- New clause G.3 Old and new LLC tag formats
  - Discuss reason for changing the way tags are done in LLC media, and discuss the compatibility issues

Tagging format

Tagging format

- At present, adding/removing a tag changes nothing about what follows the tag. That means that, on an LLC medium, a SNAP-encoded IP packet can be preceded by a SNAP-encoded VLAN tag and a SNAP-encoded CN-tag and a SNAP-encoded MACsec tag. **This is silly, especially for wireless media, where bandwidth is at a premium.**
- Suggested remedy: Change the way tags are inserted and removed on LLC media.
  - When inserting a tag into an LLC PDU, convert PDU to Type/Length and then add LLC tag.
  - When removing an LLC tag, convert enclosed (Type/Length) MSDU to LLC.
  - This requires an “LLC follows” Ethertype for LLC frames > 1536 bytes.
- Have a special managed parameter that says, “Do the old thing on this port”, but default is “do the new thing”.

Queue Sets

Queue Sets

- At present, 8.6.4 Egress says to use the vector of output ports that has been accompanying the frame to select one or more ports for output, and queue the frame in those ports' queues.
- When using Link Aggregation, one may have one set of queues per Bridge Port (as the standard suggests) or one may have one set of queues per physical port (more typical). There are weasel words in 802.1Q, but queues-per-physical-port is not what 8.6 says.
- In an 802.11 Access Point, you have one set of queues per BSS. Since one BSS can serve many stations, and there is one Bridge Port per station, an AP has many Bridge Ports per set of queues.
- So, we introduce the idea of Queue Sets to solve both problems.

Queue Sets

- A Queue Set is a group of:
  - One 8.6.6 Queuing frames entities
  - One to eight 8.6.7 Queue management entities (and their queues)
  - One 8.6.8 Transmission selection entity
- You may have multiple Queue Sets per Bridge Port, e.g. in the case of Link Aggregation. In this case:
  - There is something of an issue with Support of the EISS; you have to replicate this function several times, once for each Queue Set.
- You may have multiple Bridge Ports per Queue Set, e.g. in the case of an 802.11 Access Point. In this case:
  - The vector of output ports stays with the frame in the queue, and is used on output to select the port(s) on which the frame is transmitted.
  - There is something of an issue with Support of the EISS; see next slide.

Queue Sets

- In the case of multiple Bridge Ports per Queue Set, there is a problem with 6.9 Support of the EISS: This sublayer can generate a different VLAN tag on each port. Therefore, a single frame in an output queue can have to be transmitted multiple times, once for each VLAN tag value.
- A reasonable way to handle this is to say that knowledge of the differing requirements for tagging in 6.9 can result in multiple copies of a frame being placed in a queue, each with a subset of the port vector, so that all of the ports specified in a queue vector will be given the same tag.
- This could all be put in Clause 44, but I think that this is a common trait of Coordinated Shared Networks

Clause 44

44. Bridging 802.11 media

- The primary purpose of this clause is to reconcile the 802.11 architectural model with the 802.1 model. That is:
  - For the most part, they say the same thing in two different ways.
  - The differences are mapped to 802.11ak constructs.
- It is possible that this clause will not be needed, and that all we need is a reference to 802.11ak in Clause 5.
- It is also possible that Clause 44 shows how to make 802.11 look like 802.1 and 802.11ak shows how to map 802.1 into 802.11.
- The remainder of this section shows the 802.11 → 802.1 mapping.

IEEE Std 802.1Q-2011 Figure 8-2

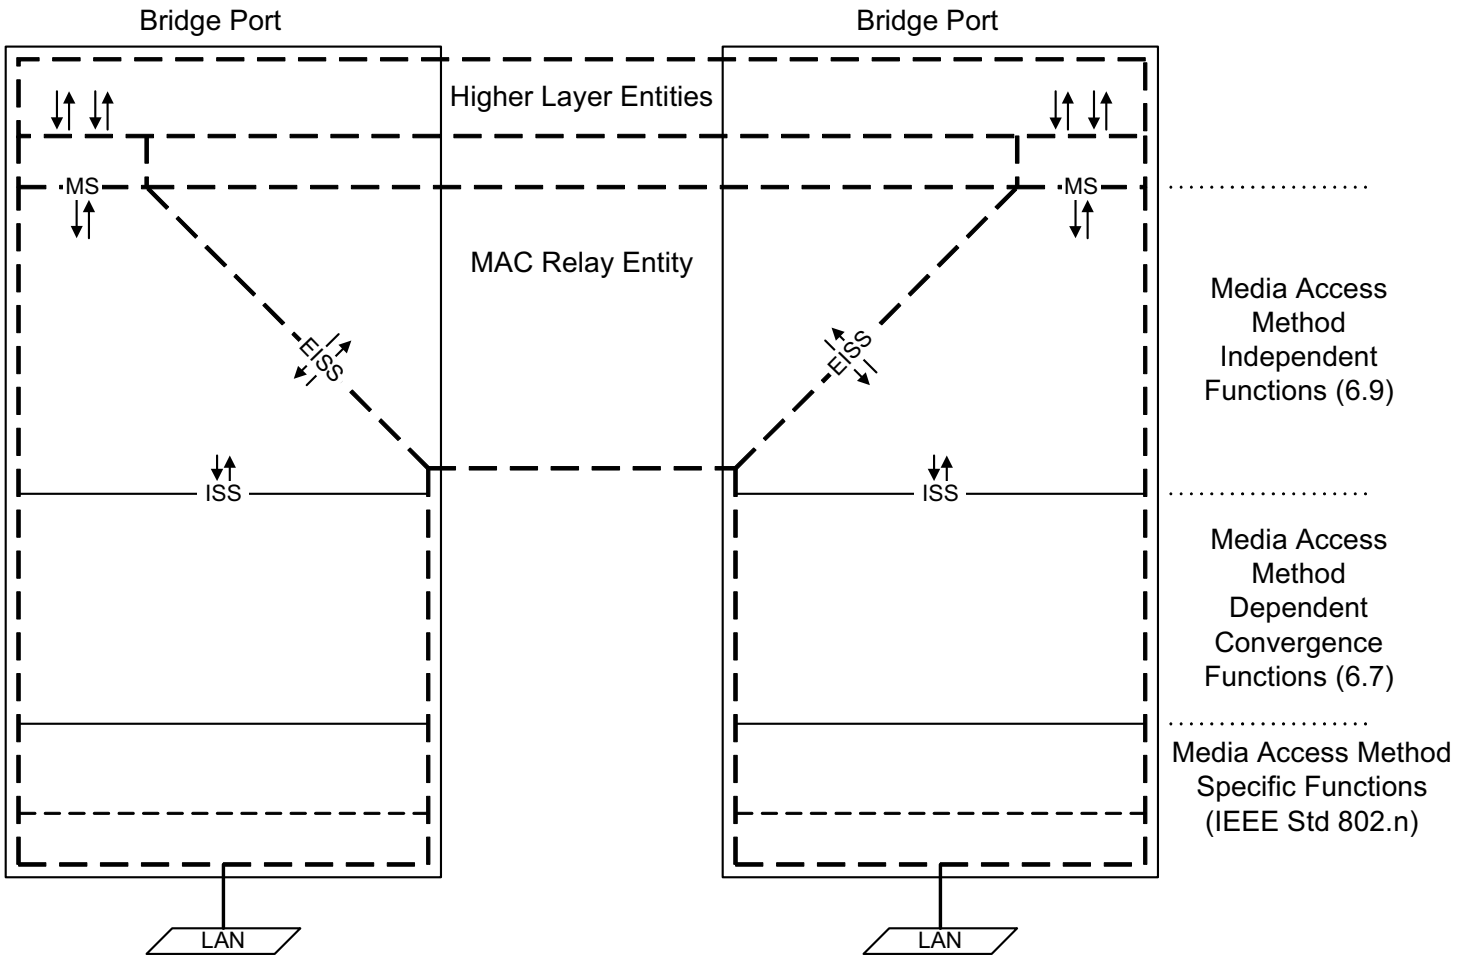

NOTE—The notation "IEEE Std 802.n" in this figure indicates that the specifications for these functions can be found in the relevant standard for the media access method concerned; for example, n would be 3 (IEEE Std 802.3) in the case of Ethernet.

Figure 8-2—VLAN-aware Bridge architecture

IEEE Std 802.11-2011 Figure 5-1

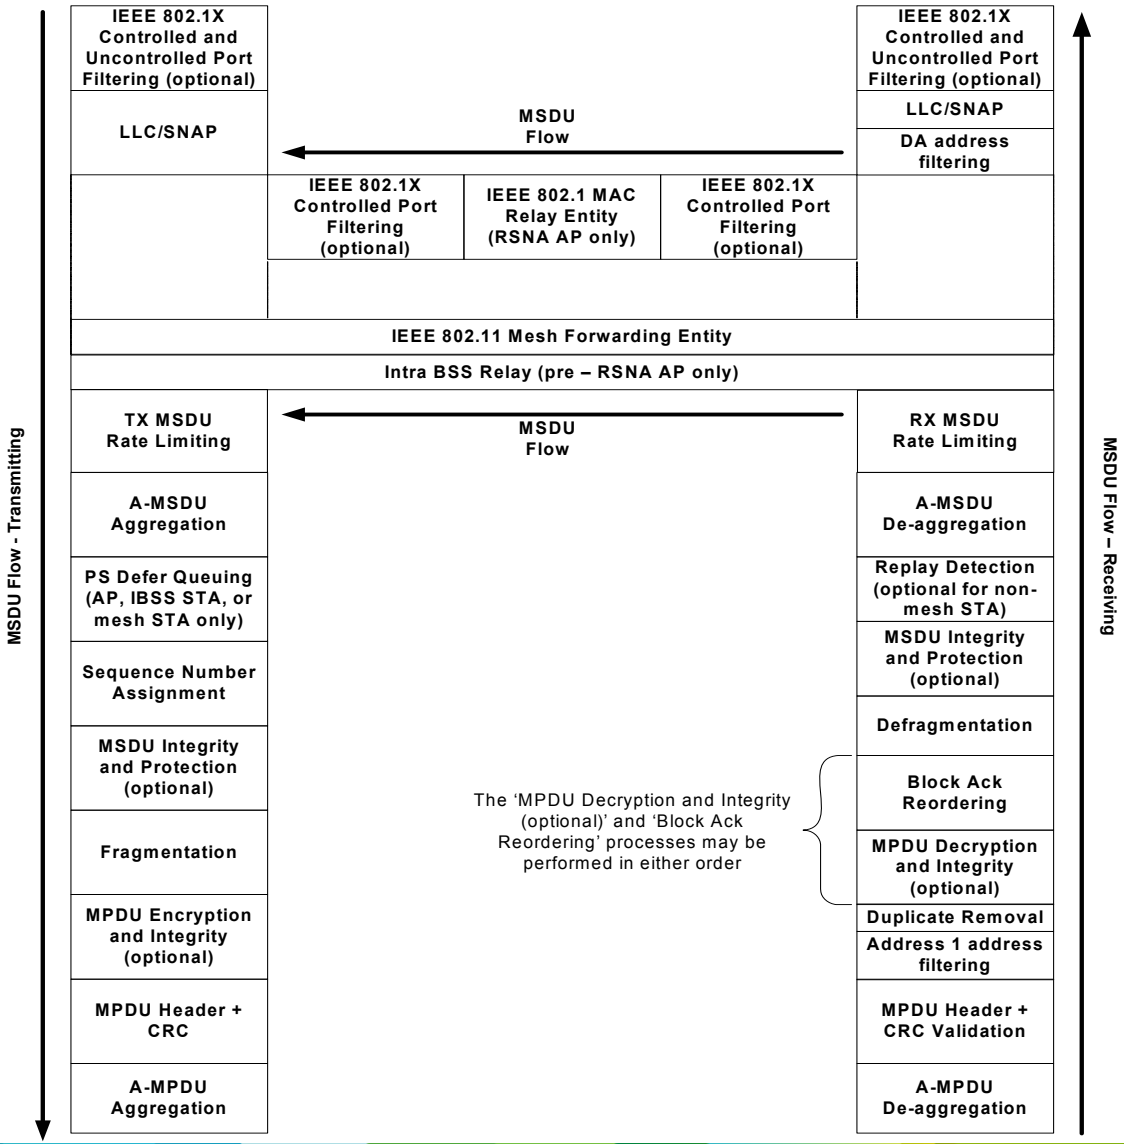

Figure 5-1—MAC data plane architecture

N ports or 1 port?

- If a station can be a bridge, 802.1 requires one Bridge Port per station attached to the Access Point. There is one leg of the baggy pants for each attached station, and an extra ISS for the SecY uncontrolled port.
- In the 802.11 model, there is exactly one “port” per BSS, with its receive and transmit sides shown separately in the diagram. The selection of which Security Association to use, and thus which “bridge port” (or the broadcast SA) is based on the destination/ receiver address in the frame. There is also a parameter specifying whether the frame is to be (or was) encrypted.
- **The resolution is not difficult – we define a new kind of port, the Security Association Port.**

Step 1: Up and Down vs Bi-directional

M-SAP to "physical port"	
TX MSDU Rate Limiting	RX MSDU Rate Limiting
A-MSDU Aggregation	A-MSDU De-aggregation
PS Defer Queuing	
Sequence Number Assignment	Replay Detection
MSDU Integrity and Protection	MSDU Integrity and Protection
Fragmentation	Defragmentation
	Block Ack Recording
MPDU Encryption and Integrity	MPDU Decryption and Integrity
	Duplicate Removal
	Address 1 Address Filtering
MPTU Header + CRC	MPDU Header + CRC Validation
A-MPDU Aggregation	A-MPDU De-aggregation

- Two legs simply become one "physical" port.
- "Physical port" has one M-SAP. (MAC Service Access Point).
- Note that this M-SAP is below the "Portal" that is offered the Bridge, today.

Step 2: Port per Security Association

- In 802.11, in the transmit direction, the BSS selection and destination address are mapped to a receiver address and security association. In the receive direction, the transmitter address selects the security association used for decryption.
- What is required in 802.1 is that each security association with an individual station is a separate Bridge Port.
- On input, this means that the transmitter address ultimately determines on which Bridge Port the frame was received.
- This means that either the transmitter address or some other kind of security association identifier must accompany the frame up the 802.11 stack to the MAC-Dependent Convergence Function, which can then use that parameter to split into multiple SAPs, one per Bridge Port

Step 2: Receive direction de-muxing

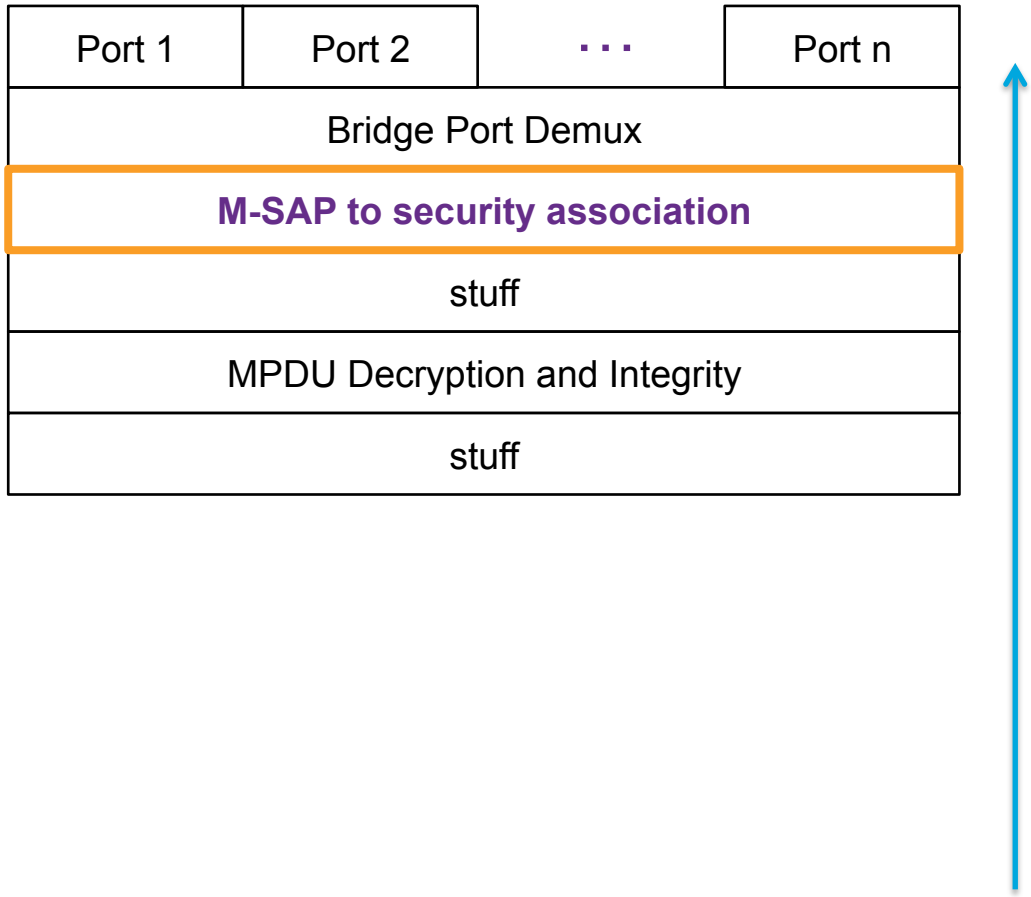

- Bridge port demux uses transmitter address (or equivalent security association ID) to select a Bridge Port on which to present the frame.
- The transmitter address or SAID must travel up the stack; this may be a new requirement for 802.11.

Step 2: Transmit direction multiplexing

- On output, there is likely a two step process.
- First, the vector of output ports that accompany a frame through the bridge forwarding process select a BSS and receiver address.
  - If only one port is in the vector, that port's unicast receiver address is used.
  - If multiple ports are selected, the appropriate multicast or broadcast receiver address is used.
- The BSS and receiver address, in turn select the security association when the frame is delivered to the 802.11 MAC.
  - If the receiver address is a unicast, a unicast SA is used.
  - If the receiver address is a multicast, the broadcast SA for the BSS is used.
- NOTE: This deck leaves the question of how exactly the receiver address maps to the selection of receiving stations (e.g., a multicast receiver address of some sort) to P802.11ak.

Step 2: Transmit direction multiplexing

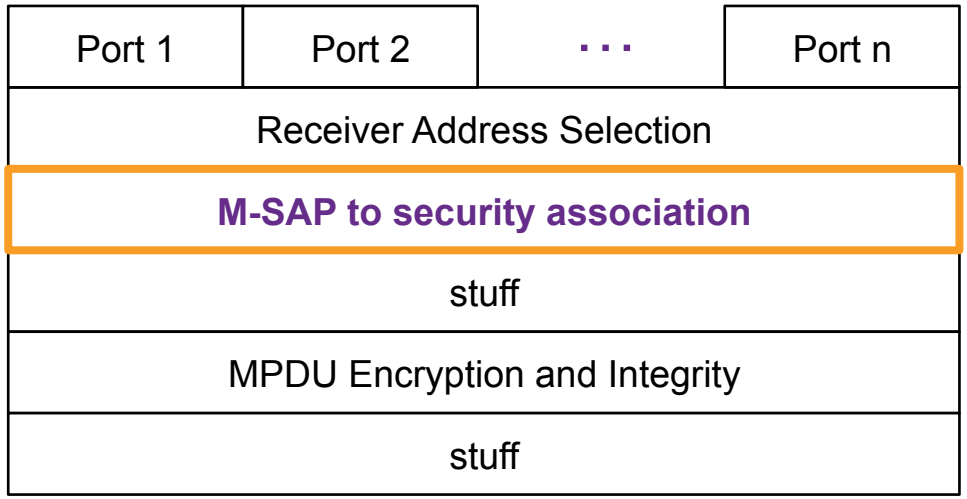

- A frame with a vector of ports is equivalent (to an observer outside the system) to an array of SAPs.
- The selection of ports determines the receiver address.
- MPDU Encryption uses the receiver address to determine the security association. (Is this a change, or how it is, now?)

Step 3: One Queue Set per BSS

- Each BSS has its own set of queues
- This merely means that we group together the Bridge Ports that are associated with a single BSS together with a single Queue Set.

Step 3: Where are the queues, really?

- It remains ambiguous, in 802.1Q, whether the queues are really above the physical port or down in the guts of the MAC. This question has not been resolved in 802.1Q, because different implementations are reasonable.
- The number of Queue Sets is more important, especially now that we have more complex dequeing algorithms. This deck proposes that we can and should settle that.
- If 802.11 wishes to specify a specific place for the queues, that is allowable within the 802.1 architecture, and 802.11ak will harm nothing by making that decision.

Step 4: MACsec and SecY

- In 802.1, encryption and decryption are above the 802.n MAC.
- In 802.11, encryption and decryption have to be below the MAC presented to the Bridge, because frames can be fragmented and reassembled, and each fragment is protected individually.

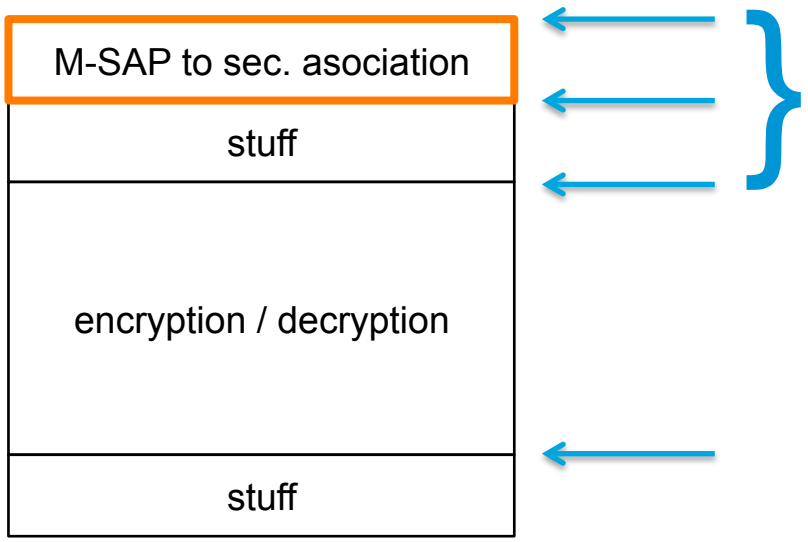

- Upper 802.11 Interface includes a parameter for “controlled/ uncontrolled port”
- Controlled/uncontrolled distinction is made in crypto block
- Lower 802.11 interface has no “controlled/uncontrolled” parameter

Step 4: MACsec and SecY

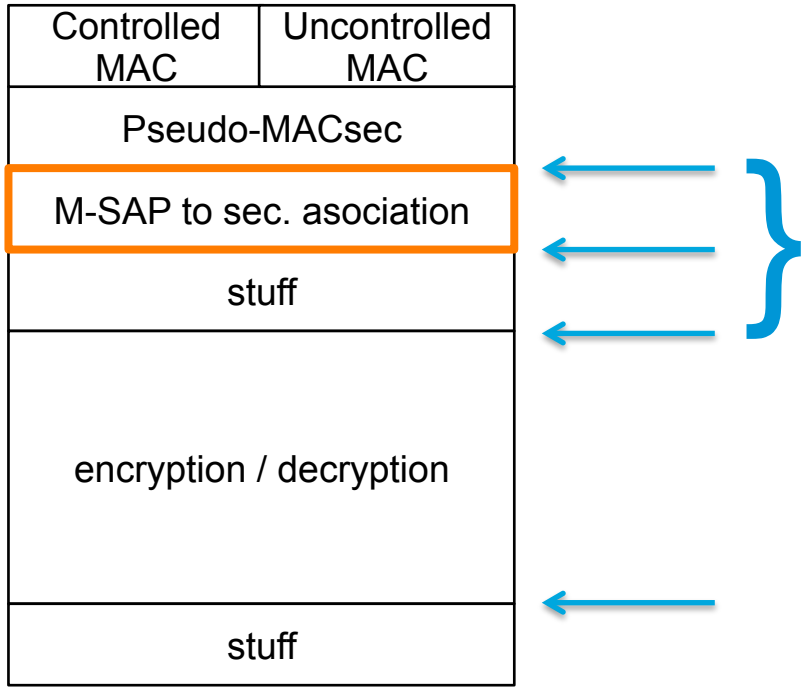

- Pseudo-MACsec layer splits controlled/uncontrolled ports using the parameter. To 802.1, everything looks normal.
- Rest of 802.11 stack remains the same.

Putting it all together

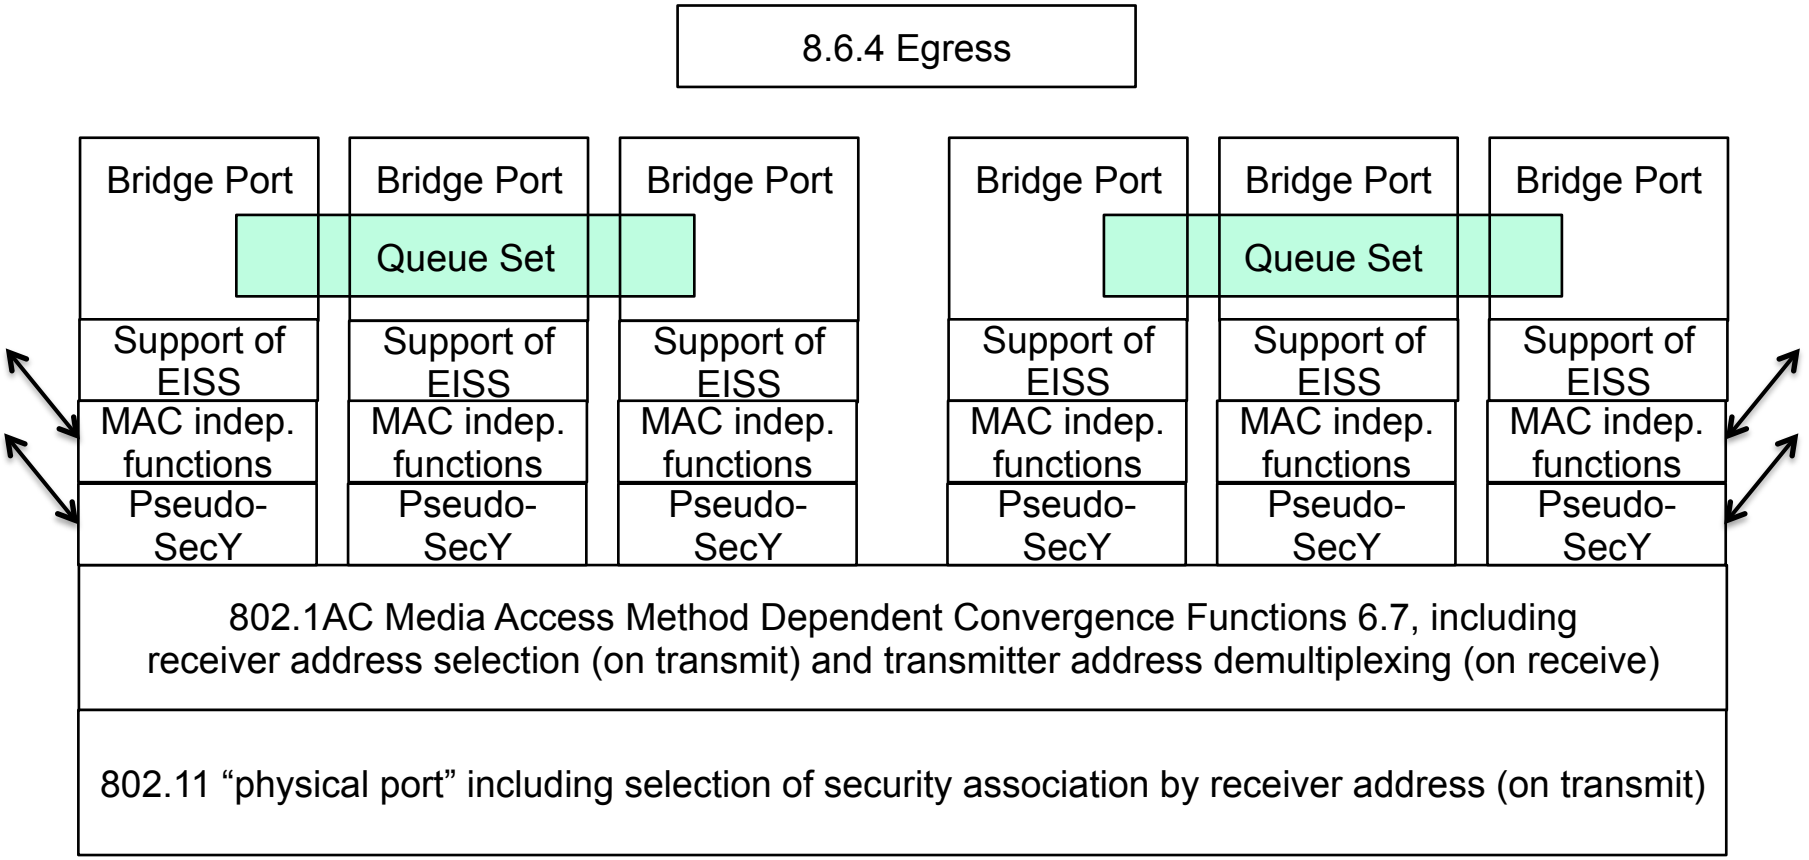

- This assumes one “physical port” covers all BSSs.

Is this more correct?

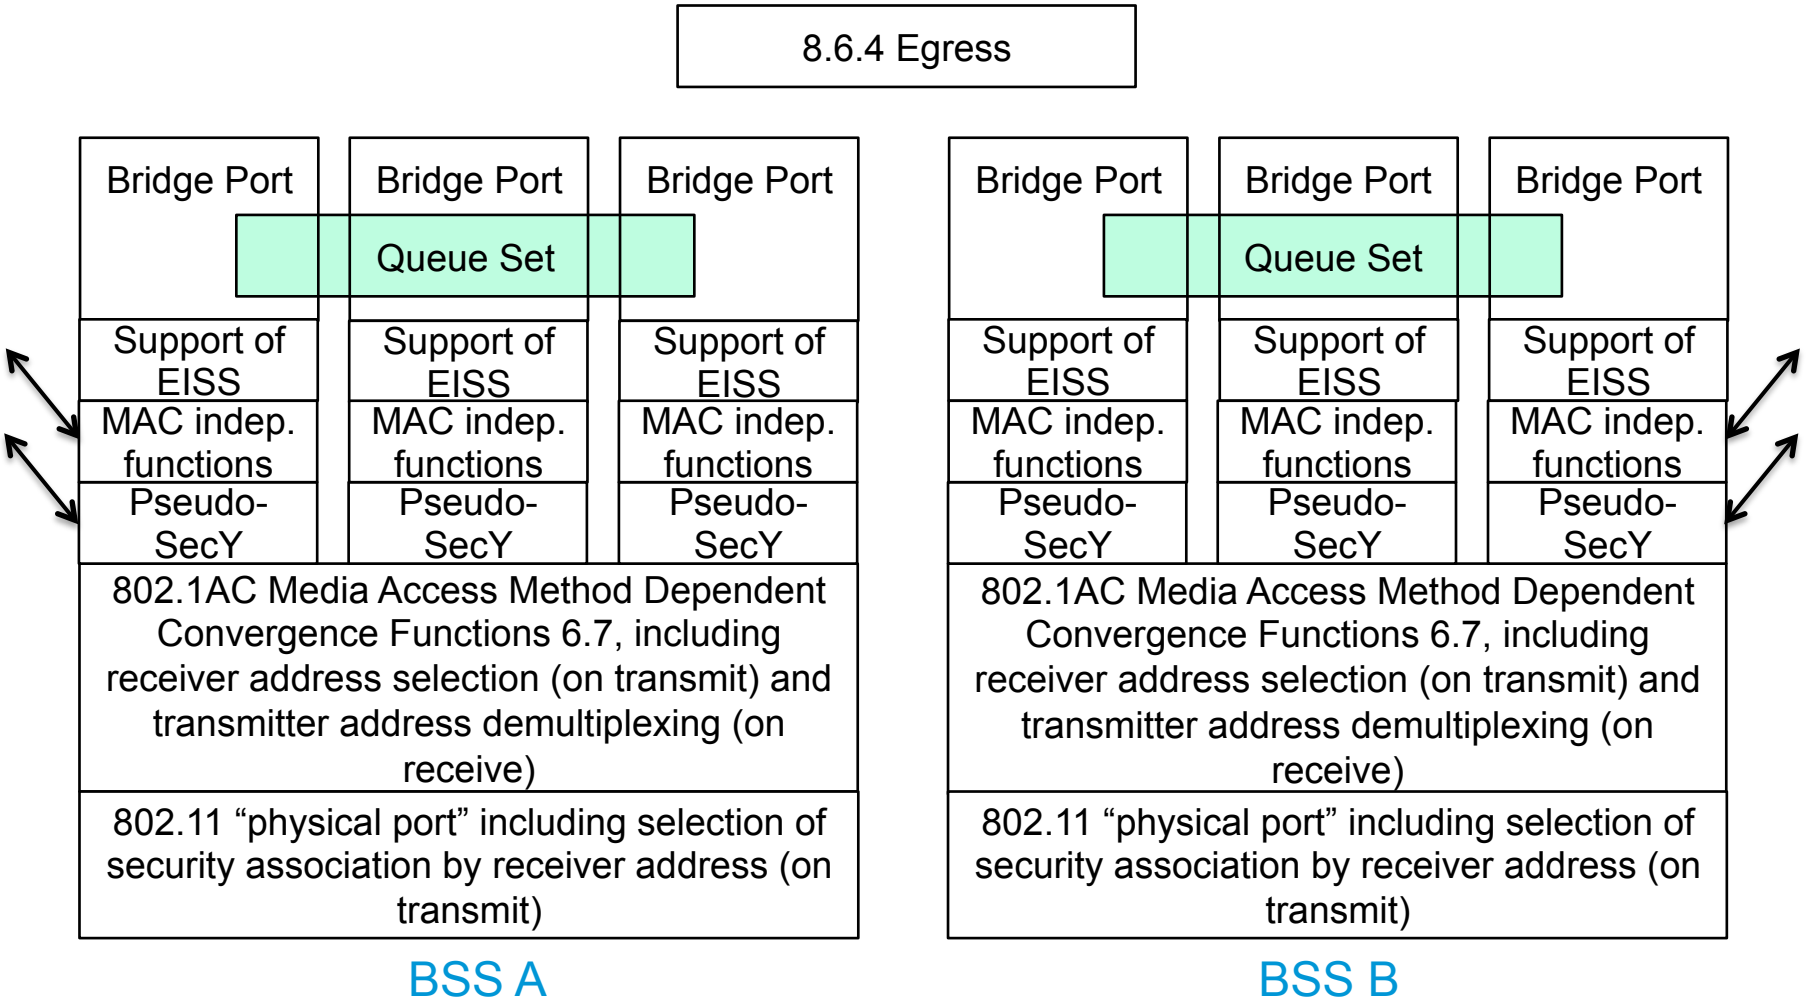

- This separates the “physical port” by BSS

Non-Access Point stations

Non-Access Point stations

- A non-AP station must present one instance of the MAC service to the Bridge (for use as a Bridge Port) for each security association that the station has with another station, whether that station is an AP or not.
- This Service Access Point must not reflect frames back to the transmitter.
- This could be called a “Portal,” a “Security Association Port,” or some third kind of object. If a third object, 802.1AC must include that, also.

Open questions

Questions

- **Do we include Clause 44?**
- How to handle basic LLC/Length-Type conversion when bridging between wired and wireless Bridge Ports?
  - This editor's opinion: Add LLC/Length-type parameter to EISS, and have "Support of EISS" translate to local format on output, if necessary.
- Is Congestion Notification allowed for 802.11 media?
  - This editor's opinion: Say nothing. It's an option, anyway.
- Is Priority-based Flow Control allowed for 802.11 media?
  - This editor's opinion: Say nothing. It's an option, anyway.
- How to handle re-transmission of frames in 802.11?
  - Since queue position is ambiguous, we can leave this to 802.11.
  - **Is there a problem with out-of-order delivery that needs to be mentioned?**

Thank you.