Cut-Through Forwarding (CTF) in Bridges and Bridged Network

– Forwarding, Learning and Active Topology (Enforcement) in Bridges –

Johannes Specht

(Self; Analog Devices, Inc.; Mitsubishi Electric Corporation; Phoenix Contact GmbH & Co. KG; PROFIBUS Nutzerorganisation e.V.; Siemens AG; Texas Instruments, Inc.)

DCN 1-22-0045-00-ICne

Introduction

- During the Nendica Meeting on October 20th 2022, I presented an updated version of the “Technical descriptions for Cut-Through Forwarding”
- The contained descriptions experienced intense discussions on “Active Topology Enforcement”, which is a processing stage present in the forwarding process of bridges.
- Upfront: Thank you for the discussions! This type of discussions help me to improve the document.
- The discussions led me to:
  1. Review the related descriptions in the aforesaid document (and several other descriptions).
  2. Review the specification of “Active Topology Enforcement” (8.6.1 of 802.1Q) and related specifications.
  3. Apply several updates to “Technical descriptions for Cut-Through Forwarding”.
  4. Discover various issues in “Active Topology Enforcement” in 8.6.1 of 802.1Q.
- This slide set is on item 4., issues in 8.6.1 of 802.1Q. While this matters for WG 802.1, it likewise matters for document “Technical descriptions for Cut-Through Forwarding”, because the therein described CTF bridges likewise perform active topology enforcement harmonized with 8.6.1 of 802.1Q.

What is “active topology”?

- In absence/before VLANs, the (Rapid) Spanning Tree Protocol [(R)STP] automatically disabled links to eliminate topological loops
- The resulting topology is called **“active topolology”**
  (can be represented a logical tree)
- Essentially implemented in Bridges by disabling Ports
- Essentially implemented in Bridges by disabling Ports for purposes of “**learning**” and “**forwarding**” separately.
  - **Learning**:
    Autodiscover’s new stations (arriving SAs at reception Ports), stores new SAs and associated Port in the filtering DB.
  - **Forwarding** (simplified):
    Relay arriving frames with DA=SA in the filtering DB to the associated Port.
- In presence of VLANs, there can be different (loop free) logical active topologies.

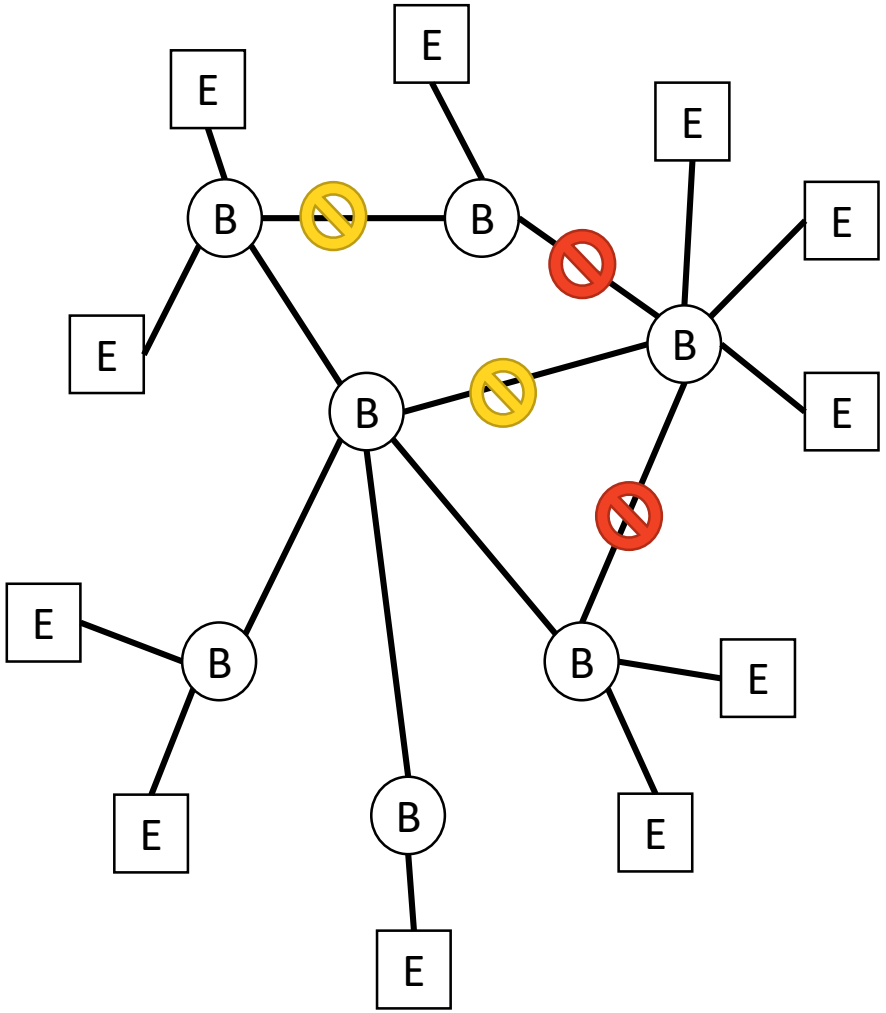

Disabling forwarding and learning in a more comprehensive picture

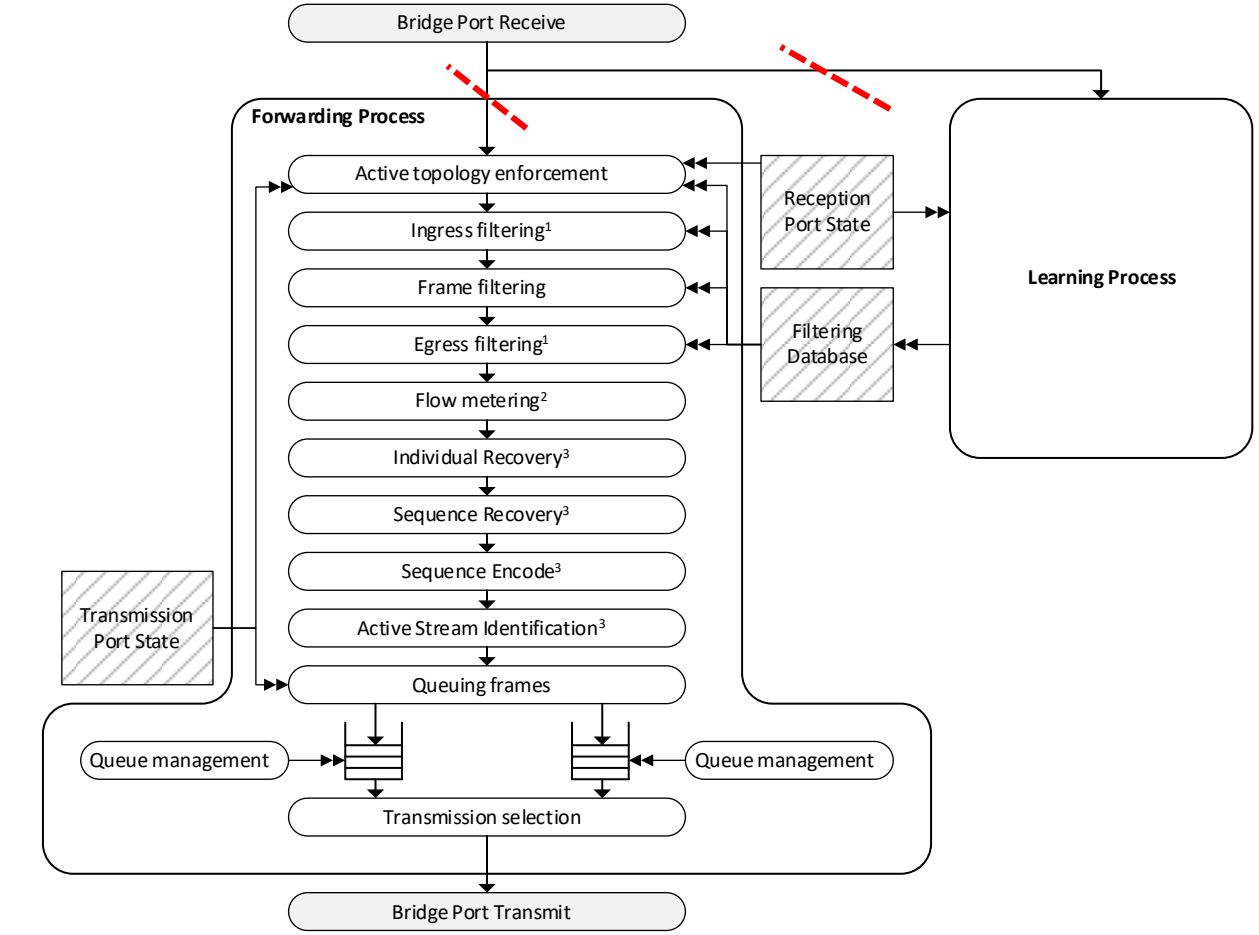

Notes

1: Optional - present in VLAN-aware CTF Bridges (absent in VLAN-unaware CTF Bridges).
2: Optional - present if PSFP is supported.
3: Optional - present if FRER is supported.

Just the forwarding process

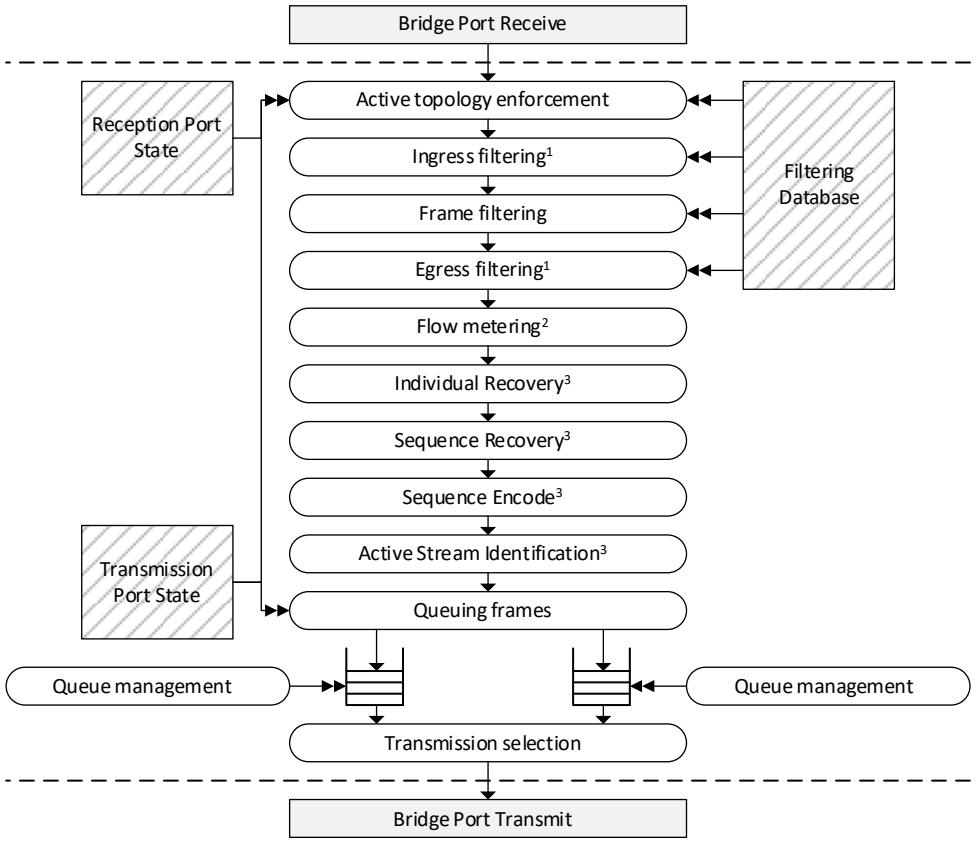

Notes
1: Optional - present in VLAN-aware CTF Bridges (absent in VLAN-unaware CTF Bridges).
2: Optional - present if PSFP is supported.
3: Optional - present if FRER is supported.

Active topology enforcement

- What it does:
  - Prevent from receive Ports being sent on transmission Ports that connect to a link that is not part of the active topology.
  - I.e., no transmission on disabled ports.
- What it **not does** (at least I could not find an explicit normative statement in 802.1Q):
  - Discard frames from disabled ports before being propagated to other stages in the forwarding process.

Some history – the Forwarding Process of 802.1D

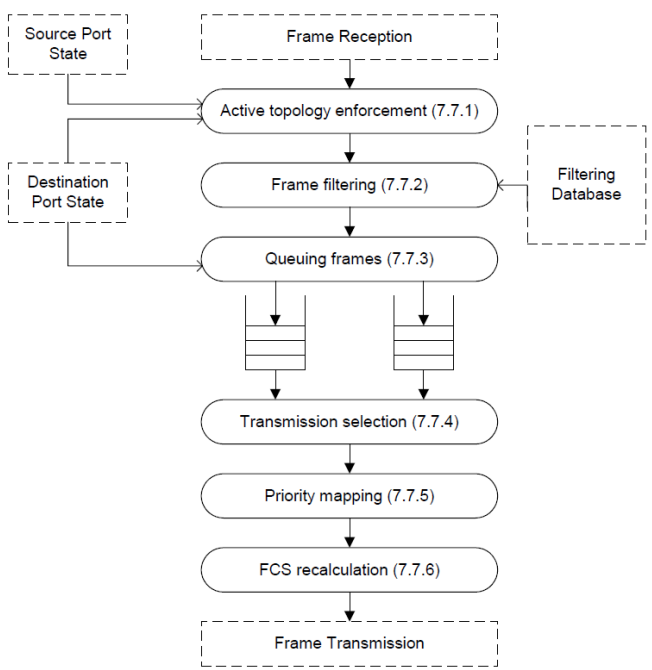

Figure 7-8—Operation of the Forwarding Process

7.7.1 Active topology enforcement

Each Port is selected as a potential transmission Port if, and only if

a) The Port on which the frame was received was in the Forwarding State (7.4), and

b) The Port considered for transmission is in the Forwarding state, and

c) The Port considered for transmission is not the Port on which the frame was received, and

d) The size of the mac_service_data_unit conveyed by the frame does not exceed the maximum size of mac_service_data_unit supported by the LAN attached to the Port considered for transmission.

For each Port not selected as a potential transmission Port, the frame shall be discarded.

Observations

1. Compared to 802.1Q, this is … short. But there are not VLANs in 802.1D, other spanning tree protocols or similar, PBB, etc.
2. The reception Port [item a)] state is part of the rule – if it is disabled, there is no (!!!) potential transmission Port.
3. There is an **explicit discarding rule**, and it is a normative requirement (shall)!
4. There is no explicit statement about **when discarding** is performed. Granted, it could be assumed to happen in 7.7.1, but there are some contradictions in the document elsewhere in detail… (I'll come back to this later)

Original text from 802.1Q-2018

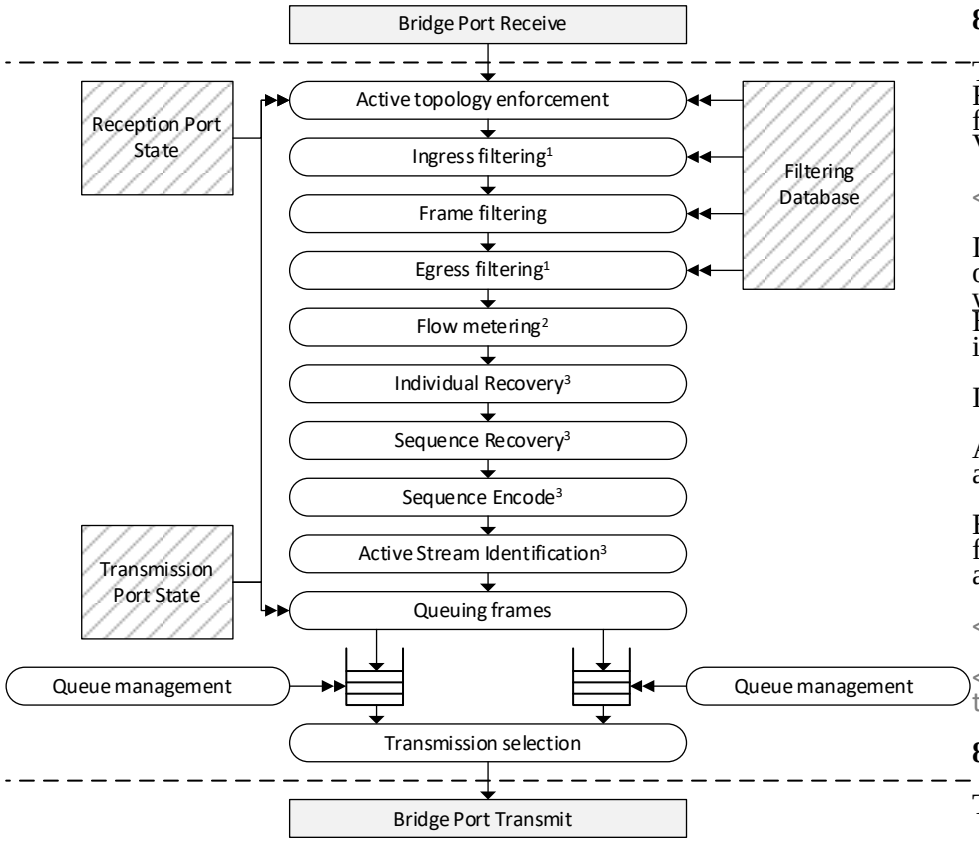

Notes
1: Optional - present in VLAN-aware CTF Bridges (absent in VLAN-unaware CTF Bridges).
2: Optional - present if PSFP is supported.
3: Optional - present if FRER is supported.

8.6.1 Active topology enforcement

To prevent data loops and unwanted learning of source MAC addresses, the Forwarding Process determines the values (TRUE or FALSE) of the learning and forwarding controls (8.4) appropriate to each received frame and Bridge Port. If learning is TRUE for the reception Port and ingress filtering (8.6.2) would not cause the received frame to be discarded, the source address and VID are submitted to the Learning Process.

<<NOTE 1 skipped - not normative by definition>>

If forwarding is TRUE for the reception Port, and either the EVBMode parameter value (40.4) for the Port is not "EVB Bridge" or the value of the operReflectiveRelayControl parameter for the Port is FALSE, each Bridge Port, other than the reception Port, with forwarding TRUE is identified as a potential transmission Port. If forwarding is TRUE for the receiving Port and the EVBMode parameter value (40.4) for the Port is "EVB Bridge" and the operReflectiveRelayControl parameter value for the Port is TRUE, each Bridge Port, including the reception Port, with forwarding TRUE is identified as a potential transmission Port.

In an edge relay (ER), the forwarding process may set learning FALSE for all frames.

An SST Bridge supports a single active topology, the CST. For each Bridge Port, RSTP determines a single value for forwarding and a single value for learning (8.4) for all frames.

Bridges with MST, PBB-TE, or SPB capabilities use the VID of the received frame to determine forwarding and learning for that frame for each Bridge Port that is enabled, i.e., has MAC_Operational True and an Administrative Bridge Port State of Enabled, as follows:

<<Fills the rest of clause 8.6.1 – skipped here (no [R]STP inside)>>

<<Skipping 8.6.2 through 8.6.5 – no "shall discard for potential transmission transmission Ports". Note that 8.6.2 only discards if the reception Port is part of the VLAN's topology, and not present in VLAN-unaware bridges>>

8.6.6 Queuing frames

The Forwarding Process shall queue each received frame to each of the potential transmission Ports (8.6.1, 8.6.3, 8.6.4).

<<Remainder of 8.6.6 skipped – it's deep in the pipeline anyway>>

8. 4 Active topologies, learning, and forwarding

<<…>> Any port that is not enabled <<…>> has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, has both forwarding and learning disabled for all spanning trees <<…>>

<<…>> Any port that has learning and forwarding disabled is assigned the Port State *Discarding*. <<…>>

Observations on Active Topology Enforcement

1. When determining the **potential transmission Ports**, the **reception Port is explicitly excluded** (in contrast to 802.1D).
2. The 802.1D normative **discard rule is gone.** Replaced by **selective (en)-queuing at the very end** (8.6.6) of the forwarding process.

8.6.1 Active topology enforcement

To prevent data loops and unwanted learning of source MAC addresses, the Forwarding Process determines the values (TRUE or FALSE) of the learning and forwarding controls (8.4) appropriate to each received frame and Bridge Port. If learning is TRUE for the reception Port and ingress filtering (8.6.2) would not cause the received frame to be discarded, the source address and VID are submitted to the Learning Process.

<<NOTE 1 skipped - not normative by definition>>

If forwarding is TRUE for the reception Port, and either the EVBMode parameter value (40.4) for the Port is not "EVB Bridge" or the value of the operReflectiveRelayControl parameter for the Port is FALSE, each Bridge Port, other than the reception Port, with forwarding TRUE is identified as a potential transmission Port. If forwarding is TRUE for the receiving Port and the EVBMode parameter value (40.4) for the Port is "EVB Bridge" and the operReflectiveRelayControl parameter value for the Port is TRUE, each Bridge Port, including the reception Port, with forwarding TRUE is identified as a potential transmission Port.

In an edge relay (ER), the forwarding process may set learning FALSE for all frames.

An SST Bridge supports a single active topology, the CST. For each Bridge Port, RSTP determines a single value for forwarding and a single value for learning (8.4) for all frames.

Bridges with MST, PBB-TE, or SPB capabilities use the VID of the received frame to determine forwarding and learning for that frame for each Bridge Port that is enabled, i.e., has MAC_Operational True and an Administrative Bridge Port State of Enabled, as follows:

<<Fills the rest of clause 8.6.1 – skipped here (no [R]STP inside)>>

<<Skipping 8.6.2 through 8.6.5 – no "shall discard for potential transmission transmission Ports". Note that 8.6.2 only discards if the reception Port is part of the VLAN's topology, and not present in VLAN-unaware bridges>>

8.6.6 Queuing frames

The Forwarding Process shall queue each received frame to each of the potential transmission Ports (8.6.1, 8.6.3, 8.6.4).

<<Remainder of 8.6.6 skipped – it's deep in the pipeline anyway>>

8. 4 Active topologies, learning, and forwarding

<<…>> Any port that is not enabled <<…>> has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, has both forwarding and learning disabled for all spanning trees <<…>>

<<…>> Any port that has learning and forwarding disabled is assigned the Port State *Discarding*. <<…>>

Isn’t selective queuing equivalent to discarding?

→ NO, I don’t think so!

802.1D

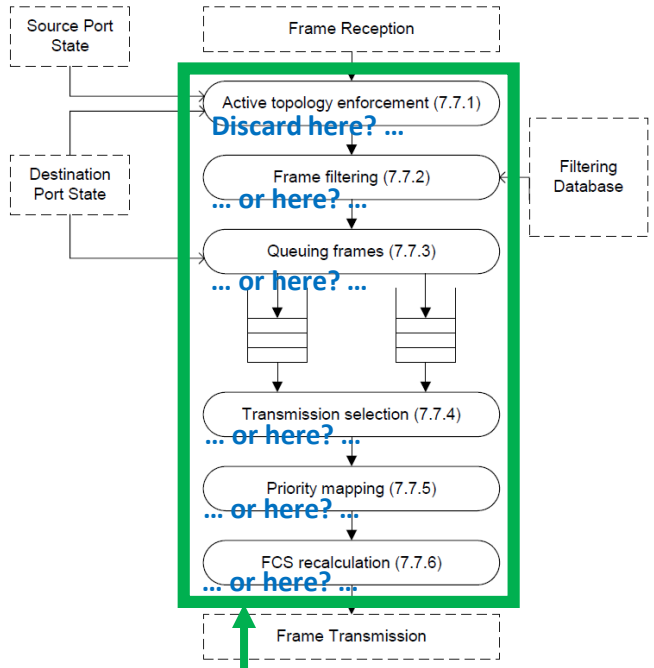

Figure 7-8—Operation of the Forwarding Process

It **does not matter when** the frame is **discarded** here:

1. Transmit Ports won’t send the frame.
2. None of the stages (7.7.1 … 7.7.6) updates bridge-internal state. They **just read bridge-internal state**.

Updates to Bridge-internal state – examples:

- Tokens in flow meters (state) can be “stolen” by frames coming from disabled reception Ports, less tokens left for frames from enabled reception Ports.
- Frames from disabled reception Ports can cause subsequent frames from enabled reception Ports being discarded by sequence recovery.

802.1Q+802.1CB

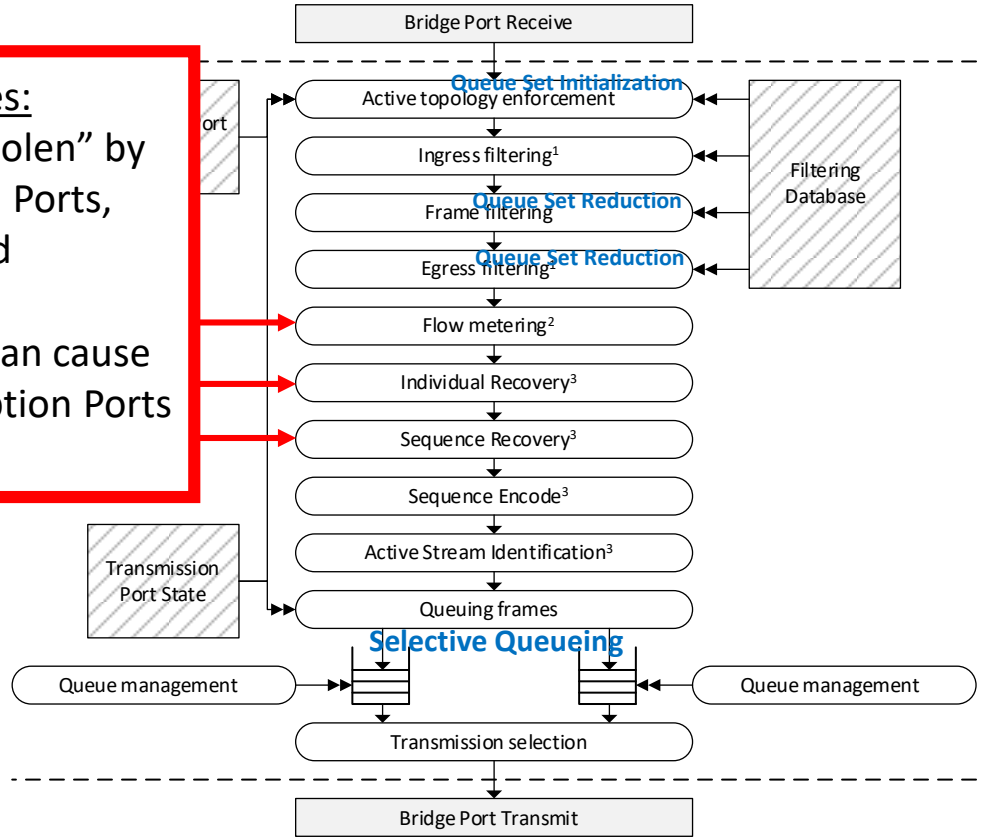

Notes

1: Optional - present in VLAN-aware CTF Bridges (absent in VLAN-unaware CTF Bridges).
2: Optional - present if PSFP is supported.
3: Optional - present if FRER is supported.

Some conclusions here

- It seems Active Topology Enforcement in .1Q should do what the name suggests: Effectively discard frames coming from disabled reception Ports.
- What is unclear to me:
  - Why has the reception Port state been removed from the discarding/selective queuing logic?
  - Including it in Active Topology Enforcement would not fix the negative impact of state updates by frames from disabled reception Ports, …
  - … but it may enable shortcut fixes in flow metering, sequence recovery, etc. (e.g., “If the set of potential transmission Ports is empty, then the flow metering is not applied to the frame”).
- To me, the issue looks like a new maintenance item I may submit

Some other observations on Active Topology Enforcement

1. I don't think that the forwarding process set's the Port state. Spanning Tree Protocols (RSTP, etc.) do this[1].
   But OK, "determines ... appropriate to each received frame" may, with some interpretation, also be understood as "reads the port state and associates it with each received frame".

2. On the other hand, if the forwarding process itself would set the Port state, than this would be a **recursion**:
   1. 1st frame received on enabled Port state.
   2. Active Topology Enforcement, itself the first stage of the forwarding process, sets the Port state do disabled.
   3. 1st frame goes through the remainder of the forwarding process (not OK).
   4. 2nd frame received on disabled Port state. OK, if there would not be the issue of the missing discard requirement discussed in earlier slides.

3. **Look-ahead** into a later stage (non-causality).
   But OK, look-aheads are ok in 802 Stds[2].
   Here, implementations may execute FDB lookup first, then execute the pipeline stage operations (8.6.1, 8.6.2, ...).

4. Starting 8.6.1 with Edge Virtual Bridging (EVB) directly after an introductory paragraph seems like a special case that appears very early in the description.
   But OK, this does not change semantics.

1: The architecture in 8.2. of 802.1Q may not pass the VID parameter to higher layer entities (e.g., spanning tree protocols) via the MAC Service (MS) Interface, but this would not prevent higher layer entities from decoding it by themselves.

2: See https://mentor.ieee.org/802.1/dcn/22/1-22-0015-01-ICne-idealistic-model-for-p802-1du.pdf

8.6.1 Active topology enforcement

To prevent data loops and unwanted learning of source MAC addresses, the Forwarding Process determines the values (TRUE or FALSE) of the learning and forwarding controls (8.4) appropriate to each received frame and Bridge Port. If learning is TRUE for the reception Port and ingress filtering (8.6.2) would not cause the received frame to be discarded, the source address and VID are submitted to the Learning Process.

<<NOTE 1 skipped - not normative by definition>>

If forwarding is TRUE for the reception Port, and either the EVBMode parameter value (40.4) for the Port is not "EVB Bridge" or the value of the operReflectiveRelayControl parameter for the Port is FALSE, each Bridge Port, other than the reception Port, with forwarding TRUE is identified as a potential transmission Port. If forwarding is TRUE for the receiving Port and the EVBMode parameter value (40.4) for the Port is "EVB Bridge" and the operReflectiveRelayControl parameter value for the Port is TRUE, each Bridge Port, including the reception Port, with forwarding TRUE is identified as a potential transmission Port.

In an edge relay (ER), the forwarding process may set learning FALSE for all frames.

An SST Bridge supports a single active topology, the CST. For each Bridge Port, RSTP determines a single value for forwarding and a single value for learning (8.4) for all frames.

Bridges with MST, PBB-TE, or SPB capabilities use the VID of the received frame to determine forwarding and learning for that frame for each Bridge Port that is enabled, i.e., has MAC_Operational True and an Administrative Bridge Port State of Enabled, as follows:

<<Fills the rest of clause 8.6.1 – skipped here (no [R]STP inside)>>

<<Skipping 8.6.2 through 8.6.5 – no "shall discard for potential transmission transmission Ports". Note that 8.6.2 only discards if the reception Port is part of the VLAN's topology, and not present in VLAN-unaware bridges>>

8.6.6 Queuing frames

The Forwarding Process shall queue each received frame to each of the potential transmission Ports (8.6.1, 8.6.3, 8.6.4).

<<Remainder of 8.6.6 skipped – it's deep in the pipeline anyway>>

8. 4 Active topologies, learning, and forwarding

<<...>> Any port that is not enabled <<...>> has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, has both forwarding and learning disabled for all spanning trees <<...>>

<<...>> Any port that has learning and forwarding disabled is assigned the Port State *Discarding*. <<...>>

Thank You for Your Attention!

Questions,
Comments,
Opinions,
Ideas?