

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18

**Proposed Text for  
A Reworked P802.1Qdd RAP Propagator**

**Individual Contribution by  
Feng Chen  
(Siemens AG)**

**Johannes Specht  
(Self, Siemens AG)**

**2022-02-21**

1	<b>Table of Contents</b>	
2	1 Introduction.....	4
3	1.1 Overview.....	4
4	1.2 Conventions.....	4
5	1.2.1 State Machine Diagrams .....	4
6	1.2.2 Pseud-Code Presentation .....	5
7	1.2.3 Variables and Scopes.....	5
8	1.2.4 create, delete, NULL .....	5
9	99.x Interfaces .....	5
10	99.x.1 RAP Participant service primitives (99.6.4.2).....	5
11	99.x.1.1 DECLARE_ATTRIBUTE.request(portID, attr).....	5
12	99.x.1.2 WITHDRAW_ATTRIBUTE.request(portID, attr).....	5
13	99.x.1.3 ATTRIBUTE_REGISTERED.indication(portID, attr).....	5
14	99.x.1.4 ATTRIBUTE_DEREGISTERED.indication(portID, attr) .....	5
15	99.x.1.5 DECLARATION_OPER_STATE.indication(portID, operState).....	5
16	99.x.2 Global service primitives.....	6
17	99.x.2.1 VLAN_CONTEXT_TOPO_CHANGE.indication(vid).....	6
18	99.x.2.2 STREAM_DA_REGISTERED.indication(portID, macAddr).....	6
19	99.x.2.3 STREAM_DA_DEREGISTERED.indication(portID, macAddr) .....	6
20	99.7 RAP Propagator .....	6
21	99.7.1 RAP Propagator overview.....	6
22	99.7.2 RAP Propagator state machine diagrams.....	6
23	99.7.3 RAP Propagator variables .....	14
24	99.7.3.1 taReg.....	14
25	99.7.3.2 taDec.....	14
26	99.7.3.3 laReg .....	14
27	99.7.3.4 laDec.....	15
28	99.7.3.5 localRaClass .....	15
29	99.7.3.6 neighborRaClass .....	15
30	99.7.3.7 port.....	15
31	99.7.3.8 portRaClass.....	15
32	99.7.4 RAP Propagator procedures .....	16
33	99.7.4.1 validateTaReg(ePortID, eTaAttr) .....	16
34	99.7.4.2 processTaIngress(eTaReg).....	16
35	99.7.4.3 processTaEgress(eTaDec).....	16
36	99.7.4.4 processLaIngress(eLaReg) .....	17

1	99.7.4.5 processLaEgress(eTaReg) .....	17
2	99.7.4.6 getTaRegDestPortIDs(eTaReg) .....	18
3	99.7.4.7 getTaStatus(eTaAttr) .....	18
4	99.7.4.8 failTa(eTaAttr, eFailureCode) .....	18
5	99.7.4.9 constructLaAttr(eStreamID, eVID, eLaStatus) .....	19
6	99.7.4.10 constructRaAttr().....	19
7	99.7.4.11 getLocalRaClass(ePriority) .....	19
8	99.7.4.12 setDomainBoundaryStatus(ePortID) .....	19
9	99.7.4.14 checkResources(eTaDec).....	20
10	99.7.4.15 adjustTa(eTaDec).....	20
11	99.7.4.16 allocateResources(eTaDec) .....	20
12	99.7.4.17 deallocateResources(eTaDec) .....	20
13	(Annex Z) Collected Issues during Development of this Document.....	21
14	Z.1 Camel-Case vs. Underscore-Case .....	21
15	Z.2 VLAN-aware LA attribute .....	21
16	99.4.4 Listener Attach attribute and TLV encoding.....	21
17	99.4.4.1 StreamID.....	21
18	99.4.4.2 VID .....	21
19	99.4.4.3 ListenerAttachStatus .....	21
20	Z.4 Sufficient validation of Talker Announce Registrations .....	22
21	Z.5 Issues related to RA class registrations .....	23
22		
23		

# 1 Introduction

## 2 1.1 Overview

3 This document contains a re-worked version of the RAP Propagator defined in subclause 99.7 of IEEE P802.1Qdd  
4 draft D0.5. Based on review feedback, several aspects were identified where clarity and readability of the technical  
5 contents defined so far in D0.5 can be enhanced.

6 The enhancements in the re-worked version include the following:

- 7 - New subclause structure, simplifying finding particular contents
- 8 - Contents ordered for top-down reading
- 9 - Removal of several synonyms/ambiguously used terms (“functions” vs. “procedures”, etc.)
- 10 - More formal description of protocol mechanisms
- 11 - Removal of unnecessary concurrency
- 12 - Various enhancements in non-formal descriptions (including shortening)

13 The re-worked contents are intended for incorporation into the next Qdd draft D0.6.

14 In this re-worked version, the functionality and operation of RAP Propagator is specified in terms of state machine  
15 diagrams, along with their variables and procedures. As an effective means of enhancing clarity and readability,  
16 pseudo-code has been used in describing the actions to be taken by a state or a procedure. The conventions for the  
17 descriptive methods used by this document are stated in 1.2.

18 In addition, several potential technical corrections and enhancements have been also identified during development  
19 of this document and collected in Annex Z at the end of this document. These technical issues need to be considered  
20 and discussed in the development of future Qdd drafts.

21 The remainder of this document is structured as follows:

- 22 - Clause 1.2 defines conventions for the proposed new representation.
- 23 - Clause 99.x summarizes the external interfaces between RAP Propagator and other entities (e.g., RAP  
24 Participant). A subset of these interfaces is found in P802.1/D0.5, in which case references are provided  
25 for further description.
- 26 - Clause 99.7 contains the re-worked version of the RAP Propagator.

## 27 1.2 Conventions

### 28 1.2.1 State Machine Diagrams

29 This document used state machine diagrams that contain pseudo-code when considered appropriate by the authors.  
30 These diagrams are aligned to the conventions in Annex E of IEEE Std 802.1Q-2018 to a large extent. The  
31 following extensions apply:

- 32 a) All transitions and operations are atomic and finish without any progress of time, as soon as the associated  
33 conditional expressions(s) of transition(s) is/are satisfied. The sole reason for steady states is that none of  
34 the conditional expressions of any outgoing transition of a particular state is satisfied.
- 35 b) Service primitives can occur in the conditional expressions of transitions. Invocation of such service  
36 primitives immediately leads to activation of a transition, provided that:
  - 37 1) All other sub-expressions are satisfied, if such expressions exist, and
  - 38 2) The state machine resides in the associated state prior to invocation.
- 39 c) Service primitives that cannot be processed immediately are queued in their order of invocation for  
40 subsequent processing (i.e., no service primitive is lost).

- 1 d) In case the conditional expressions of more than one transition are satisfied simultaneously, the taken  
2 transition is arbitrarily.

### 3 1.2.2 Pseud-Code Presentation

4 The pseudo-code found in this document is inspired by C++. The emphasis is on simplicity and clarity of  
5 specification and unambiguous description of the externally visible behavior. Efficiency (speed, memory usage,  
6 etc.) is left to the implementation (in software/firmware, hardware, combinations of the aforesaid, etc.).

### 7 1.2.3 Variables and Scopes

8 The RAP operation relies on several Bridge-local variables on different scopes, such as per-Bridge, per-Bridge per  
9 Port, or per-Bridge per Port per Stream. The pseudo-code uses an array notation (i.e., identifier followed by index  
10 expressions in brackets) to express such scopes, in which the array index expressions are used for addressing one  
11 or more entries in such arrays.

12 For example,

13 `xyz [<portID>, <StreamID>]` and `xyz [<portID>] [<StreamID>]`

14 are semantically identical and address an entry in an array `xyz` with index expressions for given Stream ID  
15 (<StreamID>) and a given Port ID (<portID>).

16 The asterisk character (`*`) is used as a wildcard for index expressions to address all items in the respective array  
17 dimension.

### 18 1.2.4 create, delete, NULL

19 Lifecycle management of entries in array variables in deeper Bridge-local scopes is realized by the two functions  
20 `create` and `delete`; such variables is initialized to special value `NULL`. The semantics of the two functions is  
21 as follows:

- 22 a) Invocation of `create <identifier> [...]` allocates an entry in an array variable and returns a reference,  
23 provided that the variable is `NULL` prior to this invocation. If it is not `NULL`, prior to the invocation of  
24 `create`, the existing reference is returned.
- 25 b) Invocation of `delete <identifier> [...]` deallocates an entry in an array variable and sets it back to  
26 special value `NULL`, provided that the variable is other than `NULL` prior to this invocation. The variable  
27 remains `NULL` if it had this value prior to the invocation of `delete`.

## 28 99.x Interfaces

### 29 99.x.1 RAP Participant service primitives (99.6.4.2)

30 << Note: This clause is provided in this document to provide references to P802.1Qdd/D0.5 and  
31 introduce a new per-RAP participant primitive `DECLARATION_OPER_STATE.indication(portID,  
32 operState)` in 99.x.1.5 not present in P802.1Qdd/D0.5.>>

#### 33 99.x.1.1 `DECLARE_ATTRIBUTE.request(portID, attr)`

34 See 99.6.4.2.1 of P802.1Qdd/D0.5.

#### 35 99.x.1.2 `WITHDRAW_ATTRIBUTE.request(portID, attr)`

36 See 99.6.4.2.2 of P802.1Qdd/D0.5.

#### 37 99.x.1.3 `ATTRIBUTE_REGISTERED.indication(portID, attr)`

38 See 99.6.4.2.3 of P802.1Qdd/D0.5.

#### 39 99.x.1.4 `ATTRIBUTE_DEREGISTERED.indication(portID, attr)`

40 See 99.6.4.2.4 of P802.1Qdd/D0.5.

#### 41 99.x.1.5 `DECLARATION_OPER_STATE.indication(portID, operState)`

42 This primitive is used by a RAP Participant on a Port (`portID`) to notify the service user of the operational state of  
43 its declaration function. The `operState` parameter contains a Boolean value that indicates whether the RAP  
44 Participant's declaration function is operational (`TRUE`) or not (`FALSE`).

1 << Note: This primitive is introduced to indicate a change to the value of portalCreated variable (99.6.2.2)  
2 in a RAP Participant.>>

### 3 **99.x.2 Global service primitives**

4 << Note: The following primitives are invoked by the local bridge in the events that could affect the  
5 current results of TA propagation, to signal the RAP Propagator that there is a need to reprocess  
6 propagation of the TA registrations. >>

#### 7 **99.x.2.1 VLAN\_CONTEXT\_TOPO\_CHANGE.indication(vid)**

8 This indication is invoked by the local Bridge to indicate a change in the set of Ports that forms the topology of a  
9 given VLAN context (99.2.4.1) identified by vid. Such a change can result from a spanning tree (7.3) or VLAN  
10 membership (7.4) reconfiguration.

#### 11 **99.x.2.2 STREAM\_DA\_REGISTERED.indication(portID, macAddr)**

12 This indication is invoked by the local Bridge to indicate that a destination MAC address (macAddr) is registered  
13 on a Port (portID).

#### 14 **99.x.2.3 STREAM\_DA\_DEREGISTERED.indication(portID, macAddr)**

15 This indication is invoked by the local Bridge to indicate that a destination MAC address (macAddr) is deregistered  
16 on a Port (portID).

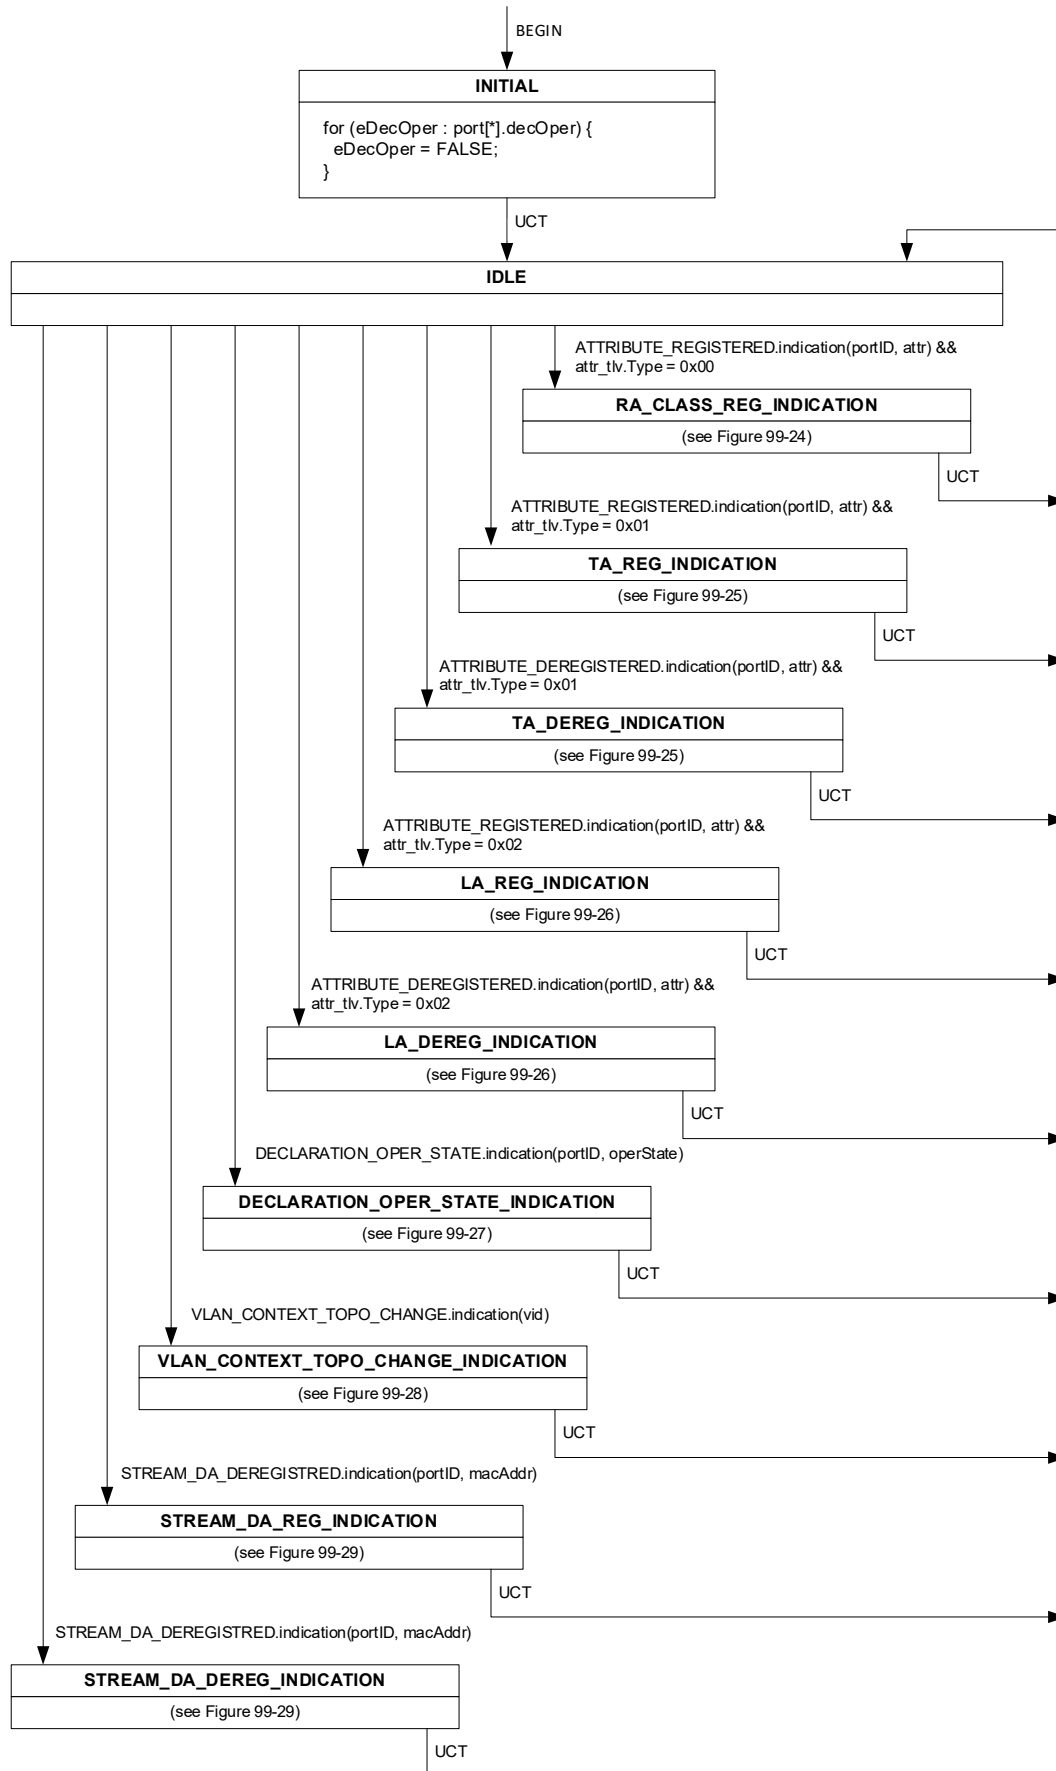
## 17 **99.7 RAP Propagator**

### 18 **99.7.1 RAP Propagator overview**

19 The operation of the RAP Propagator is defined by the state machine diagrams in 99.7.2 following the conventions  
20 in 1.2. Transitions in these diagrams are initiated by the events summarized in 99.7.3. The operations in the state  
21 machine diagrams operates on the variables specified in 99.7.4 and utilizes the procedures in 99.7.5.

### 22 **99.7.2 RAP Propagator state machine diagrams**

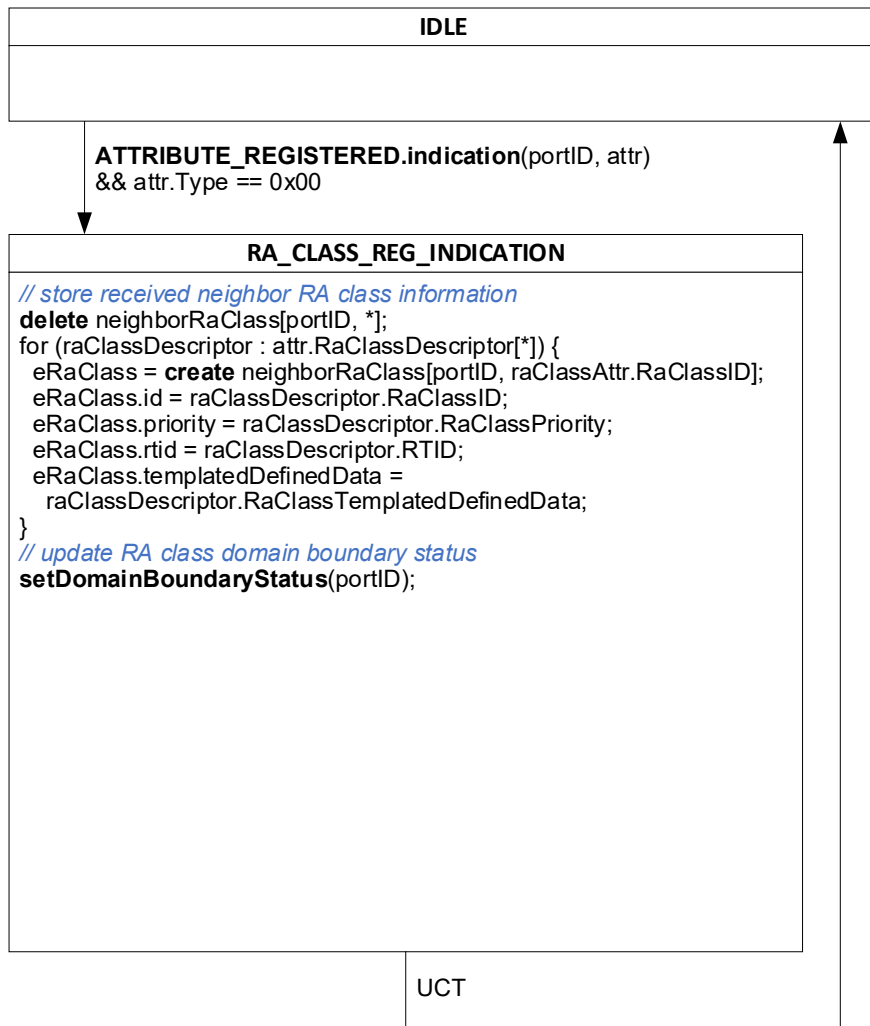
23 The operation of a RAP Propagator is defined by the state machine in Figure 99-23, where the procedures defined  
24 for the states (with the exception of the INITIAL and IDLE states) are illustrated in Figure 99-24 through Figure  
25 99-29.



1  
2

**Figure 99-23— RAP Propagator state machine**

1 The processing of RA class registration is defined by the state machine diagram in Figure 99-24.



2

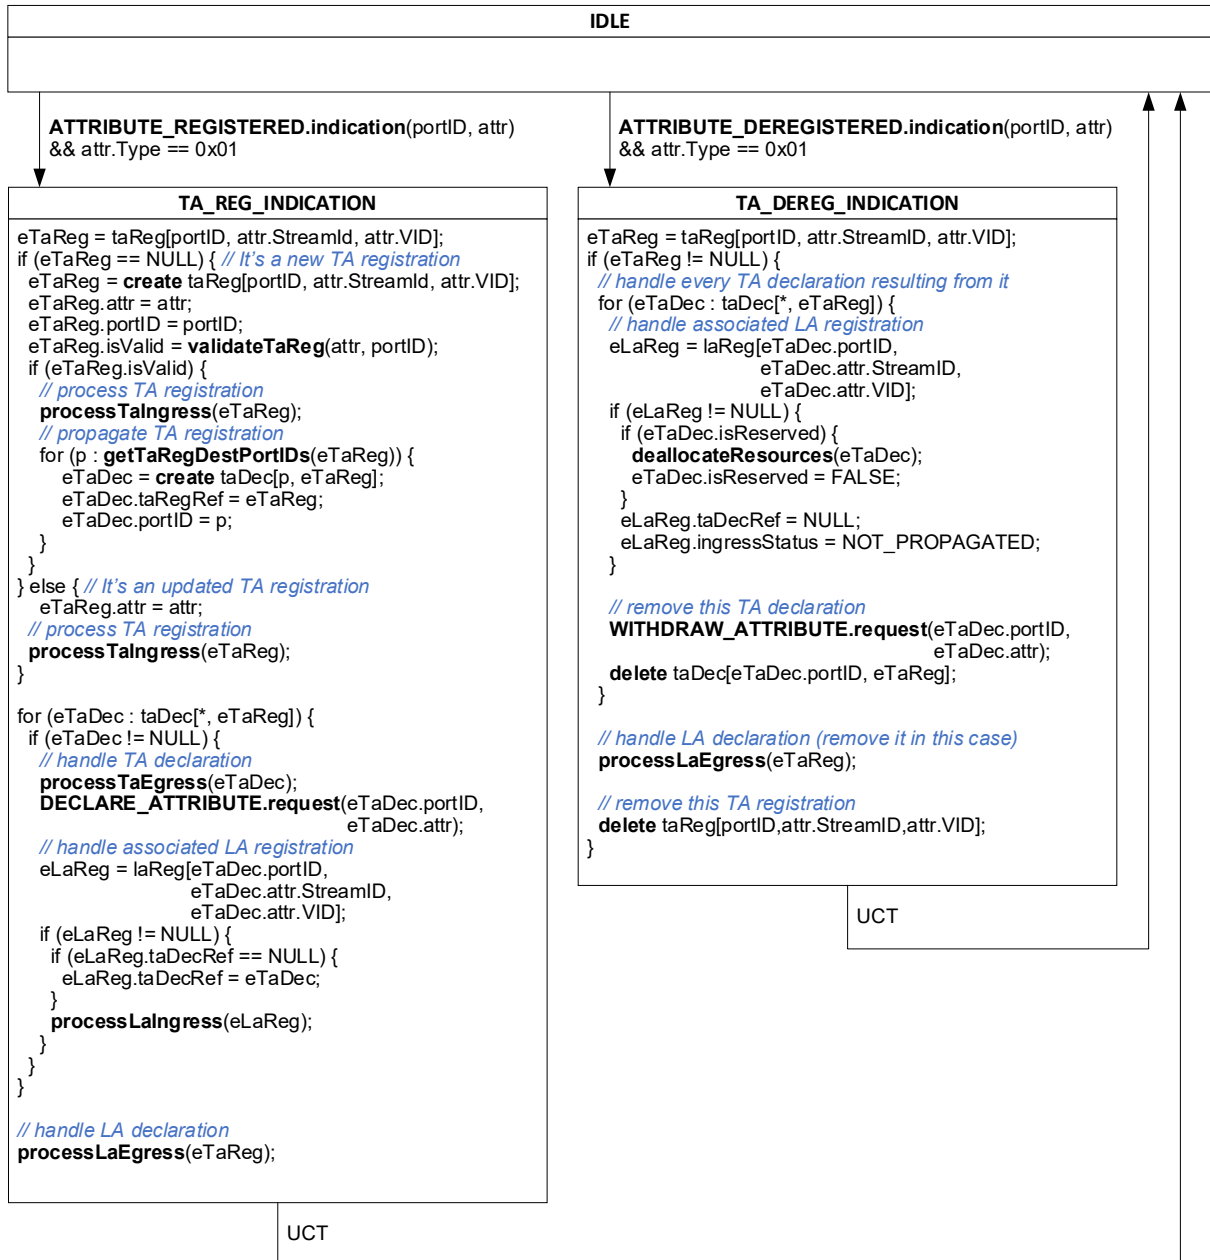
3

**Figure 99-24— Processing of RA class registration**

4

5 The processing of Talker Announce registration and deregistration is defined by the state machine diagram in  
 6 Figure 99-25.





1

2

**Figure 99-25— Processing of Talker Announce registration and deregistration**

3

The processing of Listener Announce registration and deregistration is defined by the state machine diagram in

4

Figure 99-26.

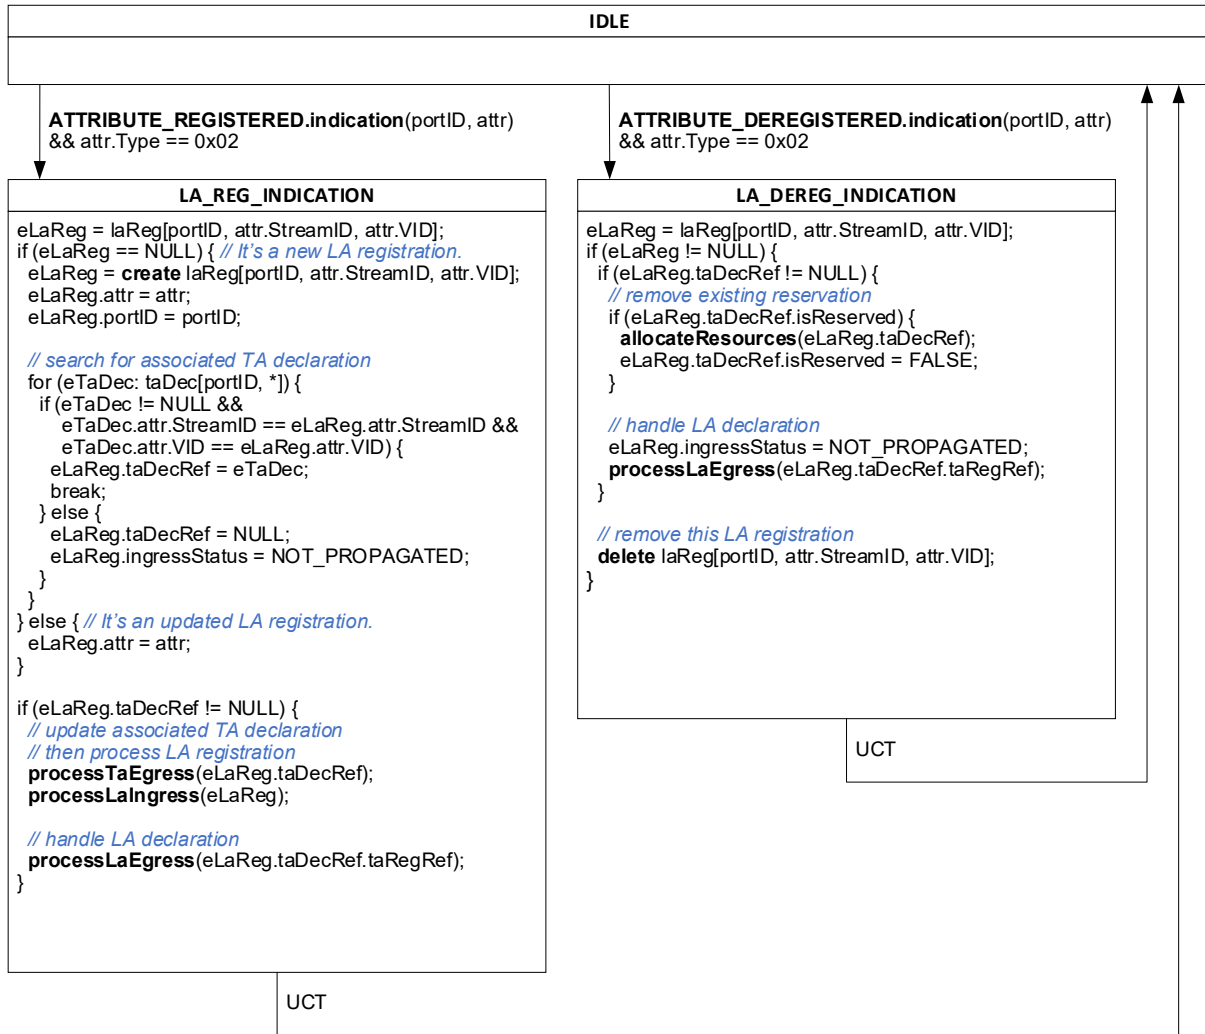
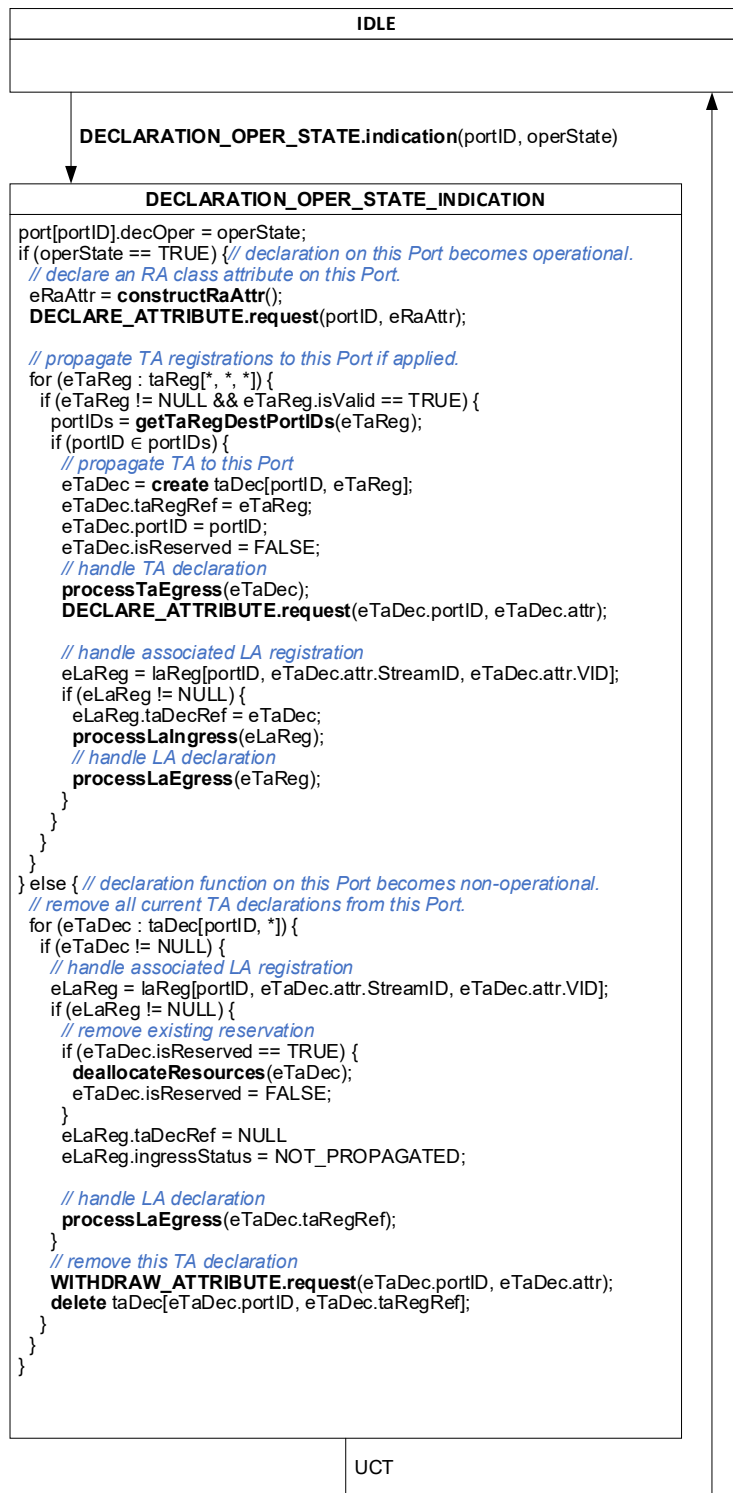


Figure 99-26 — Processing of Listener Attach registrations and deregistration

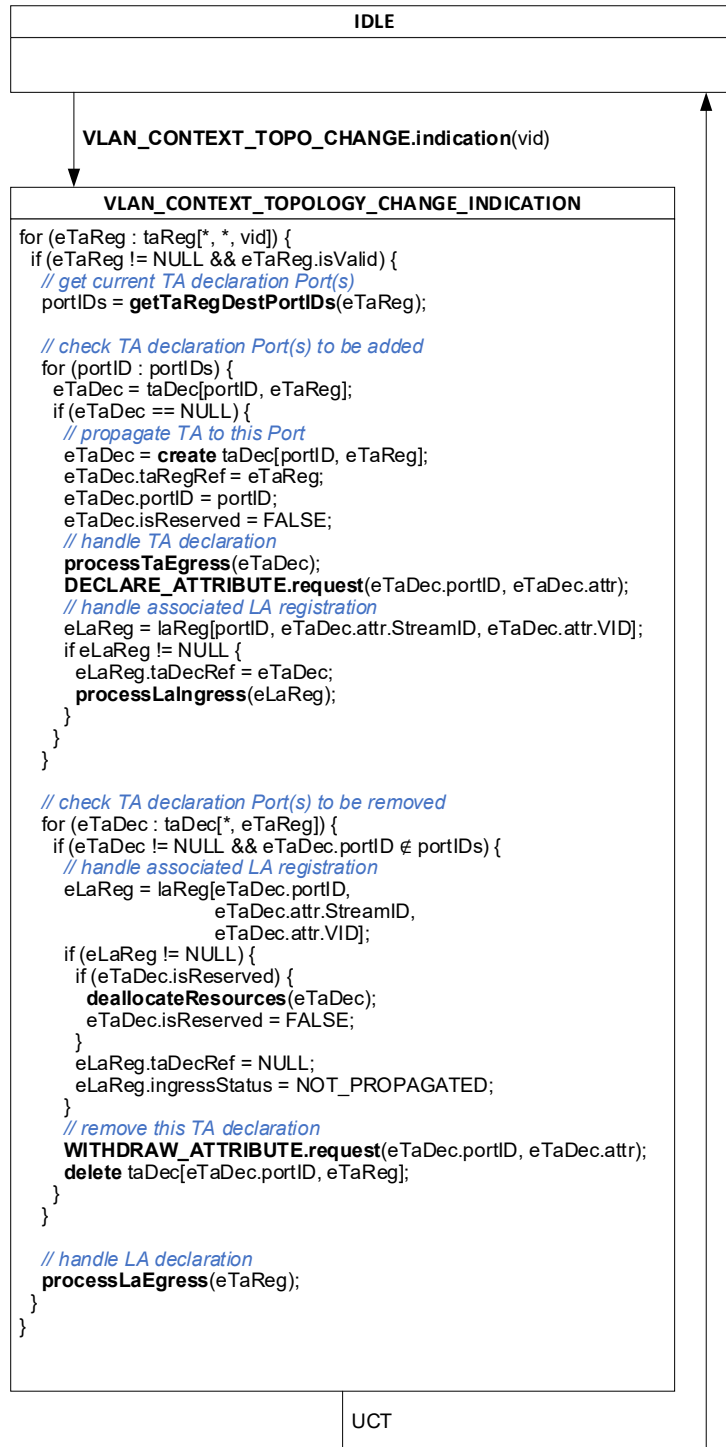
- 1
- 2
- 3
- 4 The processing of operational state change indications of the declaration function of RAP participants is defined
- 5 in Figure 99-27.



- 1
- 2
- 3
- 4
- 5
- 6

**Figure 99-27 — Processing of operational state change indications of the declaration function of RAP participants**

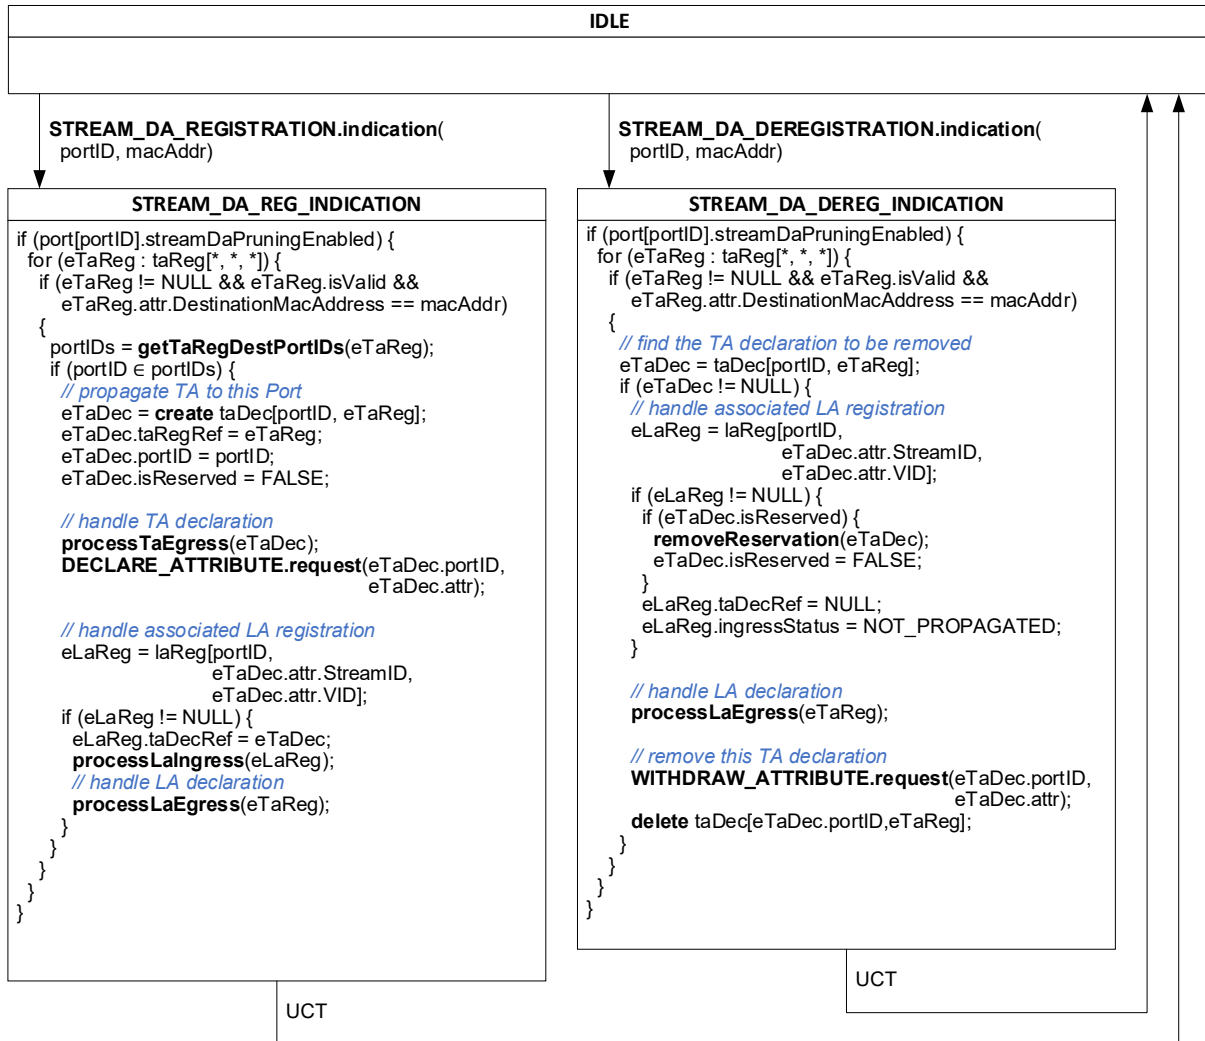
The processing of VLAN context topology change indications is define by the state machine diagram in Figure 99-28.



1  
2  
3  
4  
5  
6  
7

**Figure 99-28 — Processing of VLAN context topology change indications**

The processing of Stream Destination Address registration (99.x.2.2) and deregistration (99.x.2.3), respectively, is defined by the state machine diagram in Figure 99-29.



1

2

**Figure 99-29 — Processing of Stream Destination Address registration and deregistration**

## 1 **99.7.3 RAP Propagator variables**

### 2 **99.7.3.1 taReg**

3 The taReg variable is a 3-dimensional array, indexed by ingress Port, Stream ID (99.4.3.1) and VID (99.4.3.5.3),  
4 that contains entries representing Talker Announce registrations. Entries in taReg are records comprising the  
5 following elements:

- 6 a) **attr**: A registered Talker Announce attribute (99.4.3).
- 7 b) **portID**: The ingress Port on which the attribute in item a) is registered.
- 8 c) **isValid**: A Boolean value indicating whether the Talker Announce registration contained in this entry is  
9 valid (TRUE) or not (FALSE), as determined by the validateTaReg() procedure (99.7.4.1).
- 10 d) **ingressStatus**: The status of this Talker Announce registration on ingress determined by the  
11 processTaIngress() procedure (99.7.4.2), taking one of the following enumerated values:
  - 12 1) **TA\_RECV\_FAIL**: This Talker Announce registration contains a failure code generated by an  
13 upstream station.
  - 14 2) **TA\_INGRESS\_SUCCESS**: This Talker Announce registration contains no failure code and is not  
15 failed on ingress of this Bridge.
  - 16 3) **TA\_INGRESS\_FAIL**: This Talker Announce registration contains no failure code but failed on  
17 ingress of this Bridge with a failure code contained in ingressFailureCode [item e), below].
- 18 e) **ingressFailureCode**: A RAP failure code.

19 << Note: Item e) should be further detailed and discussed in D0.6 or later.>>

### 20 **99.7.3.2 taDec**

21 The taDec variable is a 2-dimensional array, indexed by egress Port and taReg entry, that contains entries  
22 representing Talker Announce declarations. Entries in taDec are records comprising the following elements:

- 23 a) **taRegRef**: A reference to a taReg entry from which this taDec entry results. The value of taRegRef is  
24 identical to taReg entry index of the taDec variable.
- 25 b) **attr**: A Talker Announce attribute determined by the processTaEgress() procedure (99.7.4.3) for  
26 declaration on the associated egress Port.
- 27 c) **portID**: The egress Port on which the Talker Announce declaration is made.
- 28 d) **isReserved**: A Boolean value indicating whether a reservation has been made for this TA declaration  
29 (TRUE) or not (FALSE).

### 30 **99.7.3.3 laReg**

31 The laReg variable is 3-dimensional array, indexed by ingress Port, Stream ID (99.4.4.1) and VID (99.4.4.2), that  
32 contains entries representing Listener Attach registrations. Entries in laReg are records comprising the following  
33 elements:

- 34 a) **attr**: A registered Listener Attach attribute (99.4.3).
- 35 b) **portID**: The Port on which the attribute in item a) of this entry is registered.
- 36 c) **taDecRef**: A NULL value indicating this Listener Attach registration has no associated Talker Announce  
37 declaration, or a reference to a taDec entry in the taReg that contains the Talker Announce declaration  
38 with which this Listener Attach registration is associated.
- 39 d) **ingressStatus**: The status of this Listener Attach registration on ingress determined by the  
40 processLaIngress() procedure (99.7.4.4), taking one of the following enumerated values:
  - 41 1) **NOT\_PROPAGATED**: This Listener Attach registration has no associated Talker Announce  
42 declaration and is not propagated.

- 1           2) **ATTACH\_READY**: This Listener Attach registration has an associated Talker Announce  
2           declaration and is propagated with the Attach Ready status [item a) in 99.2.5.2].
- 3           3) **ATTACH\_FAIL**: This Listener Attach registration has an associated Talker Announce  
4           declaration and is propagated with the Attach Fail status [item b) in 99.2.5.2].
- 5           4) **ATTACH\_PARTIAL\_FAIL**: This Listener Attach registration has an associated Talker Announce  
6           declaration and is propagated with the Attach Partial Fail status [item c) in 99.2.5.2].

#### 7 **99.7.3.4 laDec**

8 The laDec variable is a one-dimensional array, indexed by taReg entry, that contains entries representing Listener  
9 Attach declarations. Entries in laDec are records comprising the following elements:

- 10           a) **taRegRef**: A reference to a taReg entry with which this laDec entry is associated.
- 11           b) **attr**: A Listener Attach attribute determined by the processLaEgress() procedure (99.7.4.5) for  
12           declaration on the associated egress Port (i.e., taRegRef.portID).

#### 13 **99.7.3.5 localRaClass**

14 The localRaClass variable is a one-dimensional array, indexed by RA class ID (99.2.2.1), that contains entries  
15 describing local RA classes. Entries in localRaClass are records comprising the following elements:

- 16           a) **id**: A 8-bit RA class ID (99.2.2.1).
- 17           b) **priority**: A 8-bit RA class priority (99.2.2.2)
- 18           c) **rtid**: A 32-bit RTID (99.2.2.3)
- 19           d) **templateDefinedData**: An octet string, possibly zero length, as the value to be carried in the RA Class  
20           Template Defined Data field (99.4.2.1.4) of the RA Class Descriptor sub- TLV for this RA class in an  
21           RA Class attribute declared by the RAP Propagator.

#### 22 **99.7.3.6 neighborRaClass**

23 The neighborRaClass variable is a two-dimensional array, indexed by Port and RA class ID (99.2.2.1), that  
24 contains entries storing RA class descriptions received from neighbor stations. Entries in neighborRaClass are  
25 records comprising the following elements:

- 26           a) **id**: A 8-bit RA class ID (99.2.2.1).
- 27           b) **priority**: A 8-bit RA class priority (99.2.2.2).
- 28           c) **rtid**: A 32-bit RTID (99.2.2.3)
- 29           d) **templateDefinedData**: An octet string, possibly zero length, containing the value carried in the RA Class  
30           Template Defined Data field (99.4.2.1.4) of the RA Class Descriptor sub-TLV for this RA class in the  
31           RA class attribute registered on the Port.

#### 32 **99.7.3.7 port**

33 The port variable is a one-dimensional array, indexed by Port, that contains entries for controlling the operation of  
34 the RAP Propagator on a per Port basis. Entries in port are records comprising the following elements:

- 35           a) **streamDaPruningEnabled**: A Boolean indicating whether Stream DA Pruning (99.2.4.2) on a given  
36           Port is administratively enabled (TRUE) or disabled (FALSE). The default value is FALSE.
- 37           b) **decOper**: A Boolean indicating whether the attribute declaration function on a Port is operational (TRUE)  
38           or not (FALSE). The default value is FALSE.

#### 39 **99.7.3.8 portRaClass**

40 The portRaClass variable is two-dimensional array, indexed by Port and RA Class ID (99.2.2.1), that contains  
41 entries for controlling the operation of the RAP Propagator on a per Port, per RA class basis. Entries in portRaClass  
42 are records comprising the following elements:

- 43           a) **domainBoundaryStatus**: A Boolean indicating whether a Port is a domain boundary port (99.2.2.4) for  
44           an RA class (TRUE) or not (FALSE).

- 1       b) **maxFrameSize**: A 16-bit unsigned integer, indicating the maximum frame size, in octets, of the streams  
2       that can be reserved in an RA class on a Port.
- 3       c) **minFrameSize**: A 16-bit unsigned integer, indicating the minimum frame size, in octets, of the streams  
4       that can be reserved in an RA class on a Port.
- 5       d) **maxBandwidth**: A 32-bit unsigned integer, indicating the maximum amount of bandwidth that can be  
6       reserved for use by an RA class on a Port. The bandwidth value is represented as a percentage of the  
7       Port's transmission rate determined by the operation of the underlying MAC and expressed as a fixed-  
8       point number scaled by a factor of 1,000,000; i.e., 100,000,000 (the maximum value) represents 100%.
- 9       e) **maxHopDelay**: A 32-bit unsigned integer, indicating the maximum delay, in nanoseconds, provided by  
10      an RA class for all the streams that are reserved with that RA class on a Port.

## 11 99.7.4 RAP Propagator procedures

### 12 99.7.4.1 validateTaReg(ePortID, eTaAttr)

13 This procedure determines whether a Talker Announce registration is valid or not.

14 << Note: The operation of this procedure appears non-trivial. Further discussion and some text for  
15 later technical consideration is found in Annex Z.4.>>

### 16 99.7.4.2 processTaIngress(eTaReg)

17 This procedure performs ingress processing for a Talker Announce registration in a taReg entry (eTaReg), as  
18 follows:

```

19 processTaIngress (eTaReg) {
20     if (getTaStatus (eTaReg.attr) == Announce Fail) {
21         // TA failed by an upstream station
22         eTaReg.ingressStatus = TA_RECV_FAIL;
23     } else {
24         // find the local RA class for this TA
25         eLocalRaClass = getLocalRaClass (eTaReg.attr.Priority);
26
27         // find the neighbor RA class for this TA
28         for (eNeighborRaClass : neighborRaClass[eTaReg.portID,*]) {
29             if (eNeighborRaClass != NULL &&
30                 eNeighborRaClass.priority == eTaReg.attr.Priority)
31             {
32                 break;
33             } else {
34                 eNeighborRaClass = NULL;
35             }
36         }
37
38         if (eLocalRaClass == NULL ||
39             // this TA is not associated with any local RA class
40             eNeighborRaClass == NULL ||
41             // this TA is not associated with any neighbor RA class
42             eLocalRaClass.id != eNeighborRaClass.id)
43             // unmatching RA class ID between local and neighbor RA class
44         {
45             // TA registration is received across an RA class domain boundary
46             eTaReg.ingressStatus = TA_INGRESS_FAIL;
47             eTaReg.ingressFailureCode = << Failure code TBD >>;
48         } else {
49             // TA registration is received within an RA class domain
50             eTaReg.ingressStatus = TA_INGRESS_SUCCESS;
51         }
52     }
53 }

```

### 54 99.7.4.3 processTaEgress(eTaDec)

55 This procedure performs egress processing for a Talker Announce declaration in a taDec entry (eTaDec), as  
56 follows:



```

1  processTaEgress (eTaDec) {
2      eTaReg = eTaDec.taRegRef;
3      if (eTaReg.ingressStatus == TA_INGRESS_SUCCESS) {
4          // TA is failed neither by an upstream station nor on ingress of this Bridge
5          if (!eTaDec.isReserved) {
6              // the stream not yet reserved -> check resources
7              resFailCode = checkResources(eTaDec);
8              if (resFailCode == 0) {
9                  // the stream is reservable (resource check ok).
10                 // update accumulated latency, Tspec.
11                 eTaDec.attr = adjustTa(eTaDec);
12             } else {
13                 // the stream is unreservable (resource check failed).
14                 // declare with the egress failure code.
15                 eTaDec.attr = failTa(eTaReg.attr, resFailCode);
16             }
17         } else {
18             // the stream already reserved -> resource checking once again not needed.
19             // eTaDec.attr already contains attribute previously adjusted for declaration.
20         }
21     } else if (eTaReg.ingressStatus == TA_INGRESS_FAIL) {
22         // TA is failed on ingress of this Bridge: declare with the ingress failure code
23         eTaDec.attr = failTa(eTaReg.attr, eTaReg.ingressResFailCode);
24     }
25     } else if (eTaReg.ingressStatus == TA_RECV_FAIL) {
26         // This TA registration is failed by an upstream Bridge -> declare as-is.
27         eTaDec.attr = eTaReg.attr;
28     }
29 }

```

#### 30 **99.7.4.4 processLaIngress(eLaReg)**

31 This procedure performs ingress processing for a Listener Attach registration in a laReg entry (eLaReg) that is  
32 associated with a Talker Announce declaration (eLaReg.taDecRef), as follows:

```

33 processLaIngress (eLaReg) {
34     if (eLaReg.attr.ListenerAttachStatus == Attach Fail ||
35         getTaStatus(eLaReg.taDecRef.attr) == Announce Fail)
36     {
37         // reservation not allowed
38         if (eLaReg.taDecRef.isReserved == TRUE) {
39             // remove the existing reservation
40             deallocateResources(eLaReg.taDecRef);
41             eLaReg.taDecRef.isReserved = FALSE;
42         }
43         // propagate LA as Attach Fail
44         eLaReg.ingressStatus = ATTACH_FAIL;
45     } else {
46         // current LA status is Attach Ready or Attach Partial Fail,
47         // and current TA declaration status is Announce Success, reservation allowed.
48         if (!eLaReg.taDecRef.isReserved) {
49             allocateResources(eLaReg.taDecRef);
50             eLaReg.taDecRef.isReserved = TRUE;
51         }
52         // propagate LA as-is, which is either Attach Ready or Attach Partial Fail
53         eLaReg.ingressStatus = eLaReg.attr.ListenerAttachStatus;
54     }
55 }

```

#### 56 **99.7.4.5 processLaEgress(eTaReg)**

57 This procedure performs egress processing for a Listener Attach declaration in a laDec entry that is associated  
58 with the given taReg entry (eTaReg), as follows:

```

59 processLaEgress (eTaReg) {
60     // collect statistics from all LA registrations to be merged into this LA declaration
61     numPropagated = numReady = numFailed = 0;
62     for (eLaReg : laReg[*, eTaReg.attr.StreamID, *]) {

```

```

1   if (eLaReg != NULL &&
2       eLaReg.taDecRef != NULL &&
3       eLaReg.taDecRef.taRegRef == eTaReg &&
4       eLaReg.ingressStatus != NOT_PROPAGATED)
5   {
6       numPropagated++;
7       if (eLaReg.ingressStatus == ATTACH_READY) numReady++;
8       else if (eLaReg.ingressStatus == ATTACH_FAIL) numFailed++;
9   }
10  }
11  // determine LA declaration status (LA merging)
12  eLaDec = laDec[eTaReg];
13  if (numPropagated > 0) {
14      // at least one LA registration is propagated
15      if (eLaDec == NULL) {
16          eLaDec = create laDec[eTaReg];
17          eLaDec.taRegRef = eTaReg;
18      }
19      eStreamID = eTaReg.attr.StreamID;
20      eVID = eTaReg.attr.VID;
21      if (numPropagated == numReady) {
22          // all propagated as Attach Ready -> declare Attach Ready
23          eLaDec.attr = constructLaAttr(eStreamID, eVID, Attach Ready);
24      } else if (numPropagated == numFailed) {
25          // all propagated as Attach Fail -> declare Attach Fail
26          eLaDec.attr = constructLaAttr(eStreamID, eVID, Attach Fail);
27      } else {
28          // Otherwise: declare Attach Partial Fail
29          eLaDec.attr = constructLaAttr(eStreamID, eVID, Attach Partial Fail);
30      }
31      DECLARE_ATTRIBUTE.request(eLaDec.taRegRef.portID, eLaDec.attr);
32  } else {
33      // no LA registration is propagated
34      // withdraw current LA declaration if existing
35      if eLaDec != NULL {
36          WITHDRAW_ATTRIBUTE.request(eLaDec.taRegRef.portID, eLaDec.attr);
37          delete laDec[eTaReg];
38      }
39  }

```

#### 99.7.4.6 getTaRegDestPortIDs(eTaReg)

This procedure determines the Port(s) to which the Talker Announce registration contained in a taReg entry (eTaReg) is to be propagated. It returns a set of portIDs, possibly empty, each of which meets all the following conditions:

- 44 a) portID != eTaReg.portID;
- 45 b) port[portID].decOper == TRUE;
- 46 c) portID is in the VLAN context identified by eTaReg.attr.VID, in accordance with 99.2.4.1;
- 47 d) If port[portID].streamDaPruningEnabled == TRUE, eTaReg.attr.DestinationMacAddress is registered on
- 48 the Port with portID, in accordance with 99.2.4.2.

#### 99.7.4.7 getTaStatus(eTaAttr)

This procedure returns either of the following two enumerated values to indicate the Talker Announce status (99.2.4.4) for the given Talker Announce attribute (eTaAttr):

- 52 1) **Announce Success**: if eTaAttr does not contain a Failure Information sub-TLV (99.4.3.9).
- 53 2) **Announce Failed**: if eTaAttr contains a Failure Information sub-TLV (99.4.3.9).

#### 99.7.4.8 failTa(eTaAttr, eFailureCode)

This procedure appends the given failure code (eFailureCode) to a Talker Announce attribute (eTaAttr), as follows:

- 1 a) Construct a Failure Information sub-TLV using the System ID of this Bridge and the eFailureCode value,  
2 in accordance with 99.4.3.9.
- 3 b) Append the Failure Information sub-TLV constructed in item a) above to eTaAttr, in accordance with  
4 99.4.3.
- 5 c) Return eTaAttr.

#### 6 **99.7.4.9 constructLaAttr(eStreamID, eVID, eLaStatus)**

7 This procedure constructs and returns a Listener Attach attribute (eLaAttr) whose Value field is filled as follows:

```
8     eLaAttr.StreamID = eStreamID;
9     eLaAttr.VID = eVID;
10    eLaAttr.ListenerAttachStatus = eLaStatus;
```

#### 11 **99.7.4.10 constructRaAttr()**

12 This procedure constructs and returns an RA class attribute (eRaAttr) as follows:

- 13 a) For each localRaClass entry (eLocalRaClass), construct an RA Class Descriptor sub-TLV (raDescTlv) in  
14 accordance with 99.4.2.1 and fill its Value field, as follows:

```
15     raDescTlv.RaClassID           = eLocalRaClass.id;
16     raDescTlv.RaClassPriority     = eLocalRaClass.priority;
17     raDescTlv.RTID               = eLocalRaClass.rtid;
18     raDescTlv.RaClassTemplatedDefinedData = eLocalRaClass.templatedDefinedData;
```

- 20 b) Construct an RA Class attribute (eRaAttr) and fills its Value field with raDescTlv(s) constructed in item  
21 a) above, in accordance with 99.4.2.

- 22 c) Return eRaAttr.

#### 23 **99.7.4.11 getLocalRaClass(ePriority)**

24 This procedure searches for a local RA class with the given RA class priority (ePriority), as follows:

```
25 getLocalRaClass (ePriority) {
26     for (eLocalRaClass: localRaClass[*]) {
27         if (eLocalRaClass.priority == ePriority) {
28             break;
29         } else {
30             eLocalRaClass = NULL;
31         }
32     }
33     return eLocalRaClass;
34 }
```

#### 35 **99.7.4.12 setDomainBoundaryStatus(ePortID)**

36 This procedure determines for the given Port (ePortID) the RA class domain boundary status (99.2.2.4) for each  
37 local RA class, as follows:

```
38 setDomainBoundaryStatus (ePortID) {
39     for (eLocalRaClass : localRaClass[*]) {
40         eNeighborRaClass = neighborRaClass[ePortID, eLocalRaClass.id];
41         if (eNeighborRaClass != NULL &&
42             eNeighborRaClass.priority == eLocalRaClass.priority)
43         {
44             portRaClass[ePortID, eLocalRaClass.id].domainBoundaryStatus = FALSE;
45         } else {
46             portRaClass[ePortID, eLocalRaClass.id].domainBoundaryStatus= TRUE;
47         }
48     }
49 }
```

50 NOTE—The RA class domain boundary status determined by this procedure is used by the Bridge to adjust the  
51 operation of priority regeneration, in accordance with 6.9.4.

52

1 **99.7.4.14 checkResources(eTaDec)**

2 << Note: It should be considered for D0.6 or subsequent draft to add explicit specification of the  
3 operation of the checkResources procedure.>>

4 This procedure performs resource checking for a Talker Announce declaration contained in the given taDec entry  
5 (eTaDec), to determine whether the announced stream is deemed “Reservable” or “Unreservable”, as follows:

6 a) eLocalRaClass = **getLocalRaClass**(eTaDec.attr.Priority);

7 b) Return a Zero value to indicate “Reservable”, if all the following conditions are met:

8 1) **Bandwidth**: The bandwidth required by this stream, plus that of all the streams currently reserved in  
9 the same RA class, does not exceed the bandwidth budget as indicated by the  
10 portRaClass[eTaDec.portID, eLocalRaClass.id].maxBandwidth value [item d) in 99.7.3.8].

11 2) **Latency**: The worst-case maximum latency, computed taking into account all the streams currently  
12 reserved in the target RA class plus this stream, does not exceed the delay budget as indicated by the  
13 portRaClass[eTaDec.portID, eRaClassID].maxHopDelay value [item e) in 99.7.3.8].

14 3) **Bridge resources**: There are sufficient resources in all the Bridge components (PSFP, queue buffer,  
15 shaper, FRER, etc) required by the announced stream.

16 c) Otherwise, return a RAP failure code (non-zero) associated with the reason why the stream is  
17 unreservable.

18 **99.7.4.15 adjustTa(eTaDec)**

19 << Note: The actual operation of this procedure appears shaper-specific and should be explicitly  
20 specified in D0.6 or later.>>

21 This procedure adjusts a Talker Announce attribute in preparation for a Talker Announce declaration referenced  
22 by eTaDec.

23 **99.7.4.16 allocateResources(eTaDec)**

24 << Note: This procedure implements 99.7.6.2 of D0.5. The operation of this procedure is intended to  
25 be specified explicitly in subsequent drafts (D0.6 and later) to ensure interoperability.>>

26 This procedure allocates resources of the underlying Bridge mechanisms (e.g., traffic shaper) in preparation for  
27 the stream referenced by eTaDec.

28 NOTE – In the normal operation of RAP, allocateResources finally allocates Bridge-internal resources that have been determined to be  
29 available by the checkResources procedure [item b)3) in 99.7.4.14] earlier.

30 **99.7.4.17 deallocateResources(eTaDec)**

31 << Note: This procedure implements 99.7.6.2 of D0.5. The operation of this procedure is intended to  
32 be specified explicitly in subsequent drafts (D0.6 and later) to ensure interoperability.>>

33 This procedure de-allocates resources of the underlying Bridge mechanisms (e.g., traffic shaper) at termination of  
34 the stream referenced by eTaDec

35 NOTE – In the normal operation of RAP, deallocateResources finally de-allocates Bridge-internal resources that have allocated by the  
36 allocateResources procedure (99.7.4.17) earlier.

37

38

39

40

41

1 **(Annex Z) Collected Issues during Development of this Document**

2

3 **Z.1 Camel-Case vs. Underscore-Case**

4 Several identifiers in P802.1Qdd/D0.5 use an underscore-case notation and should be changed to camel-case  
5 notation for consistency, compactness, and readability.

6 **Z.2 VLAN-aware LA attribute**

7 The Listener Attach attribute should be extended by a VID for FRER. The following contents are intended to  
8 replace clauses in 99.4.4 of P802.1Qdd/D0.5, followed by a figure to illustrate the issue for discussion.

9 **99.4.4 Listener Attach attribute and TLV encoding**

10 Listener Attach attributes are used in Listener Attach (99.2.5). A Listener Attach attribute encodes in the Value  
11 field a set of parameters, as illustrated in Figure 99-19.

	Octet	Length
StreamID	1	8
VID	9	1
ListenerAttachStatus	10	1

12

13 **Figure 99-19—Value of Listener Attach attribute TLV**

14 **99.4.4.1 StreamID**

15 An 8-octet field encoding the StreamID element as specified in 46.2.3.1.

16 **99.4.4.2 VID**

17 A 1-octet field encoding a VID.

18 **99.4.4.3 ListenerAttachStatus**

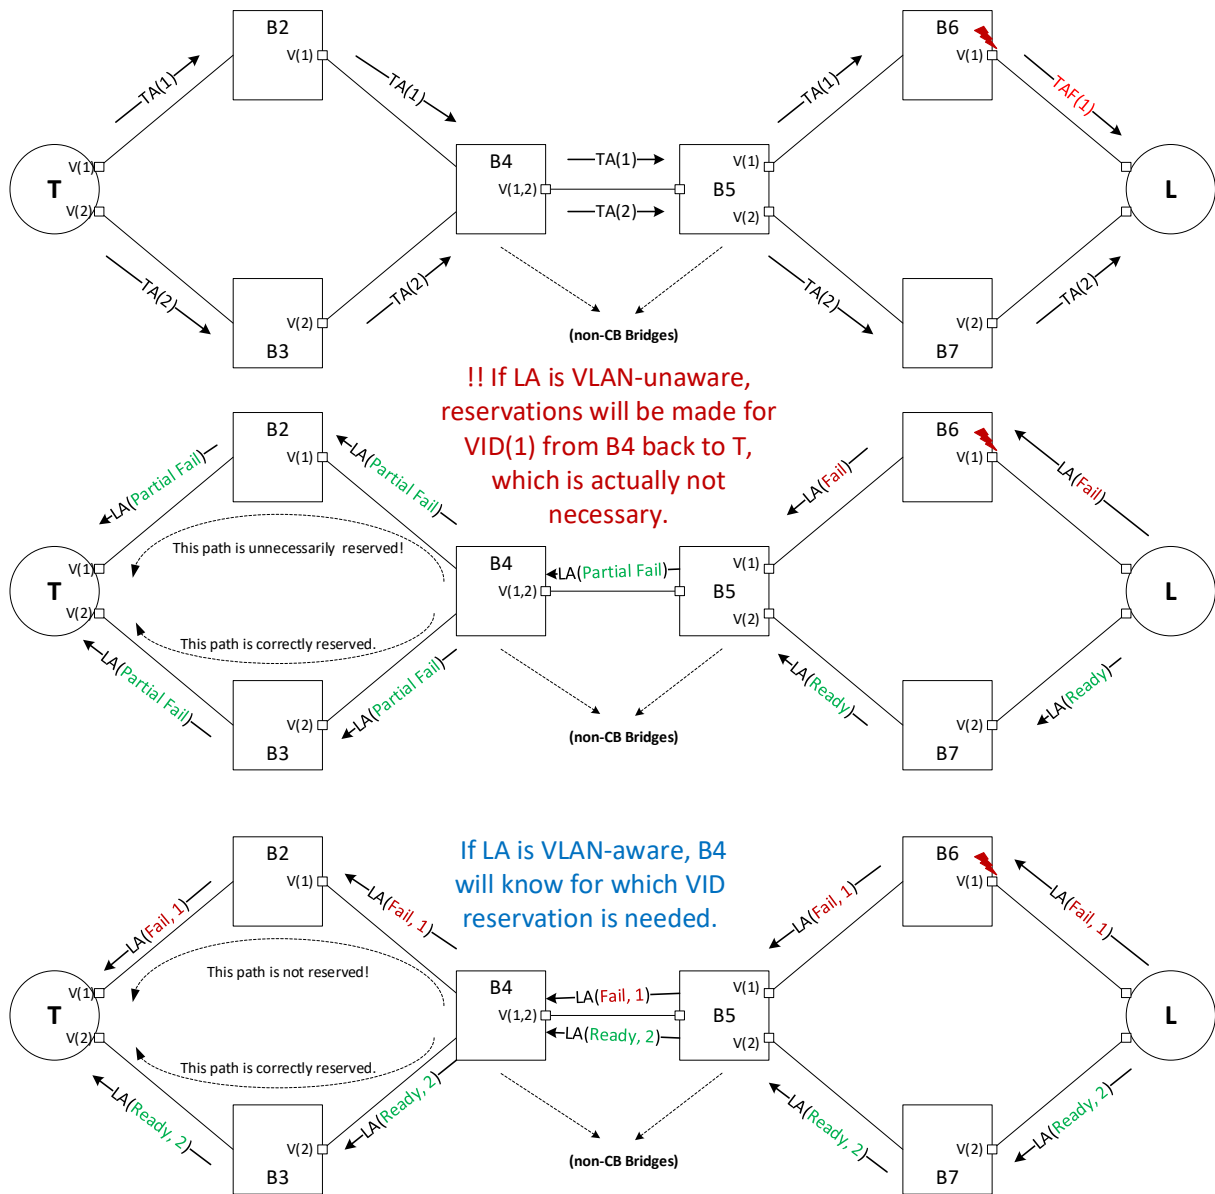
19 A 1-octet unsigned integer, taking one of the following three numerical values to indicate the status of the Listener  
20 Attach on the Port that declares the Listener Attach attribute:

21 **Table 99.x—Listener Attach status enumeration**

Name	Value	Reference
Attach Ready	1	item a) in 99.2.5.2
Attach Partial Fail	2	item b) in 99.2.5.2
Attach Fail	3	item c) in 99.2.5.2

22

**Scenario: E2E redundancy in a network with joint path  
between two non-CB Bridges**



1

2 The figure above illustrates what issues may be caused by VLAN-unaware Listener Attach, and how such issues  
3 can be resolved by use of VLAN-aware Listener Attach.

4

5 **Z.4 Sufficient validation of Talker Announce Registrations**

6 RAP Propagator needs to validate each Talker Announce registration for detection and appropriate handling of  
7 errors that may be caused by the following reasons:

- 8 a) Looping in the network topology (e.g., in the progress of spanning tree reconfiguration).  
9 b) Violation of uniqueness rules in use of Stream ID, Stream DA, etc.  
10 c) Misbehaving Bridge/end stations.

11 For single-context TA (non-redundant stream), the following error conditions may occur:

- 12 • `newTaReg.StreamId == oldTaReg.StreamId &&`  
13 `newTaReg.portID != oldTaReg.portID`

1       => keep old valid, set new invalid.

2       • newTaReg.StreamId == oldTaReg.StreamId &&

3             newTaReg.portID == oldTaReg.portID &&

4             (newTaReg.DA != oldTaReg.DA || newTaReg.VID != oldTaReg.VID ||

5             newTaReg.Priority != oldTaReg.Priority )

6       ⇒ set new as "hold" (NOT propagated) until the LA declaration on this

7       port has "LA failed" status or has been withdrawn.

8       ⇒ withdraw all existing TA declarations resulting from the old

9       Error cases for multiple-context TA (redundant streams) need to be further considered? Also, since LA is VLAN-

10       aware per Z.3, validation of LA registrations seems necessary as well.

11

## 12 **Z.5 Issues related to RA class registrations**

13 **Issue 1:** How a bridge deals with the existing TA registrations on a Port after receiving an updated RA class

14 registration from a neighbor station on that Port?

15 The current solution is to immediately reprocess all current TA registrations on that Port, based on the updated

16 neighbor RA class information.

17 To be discussed: should we consider enforcing the neighbor station, e.g. controlled by management, to perform a

18 reset to withdraw all TA declarations before it applies changes to its local RA class configurations.

19 **Issue 2:** How to deal with an RA class deregistration event?

20