Fragile bridges revisited

Mick Seaman

IEEE Std 802.1Q–2011 added a capability,[1] configurable per-Bridge Port, that avoids the creation of a data loop as a result of the failure of a neighbouring 'fragile' bridge's spanning tree protocol entity. I contributed a white paper [2] addressing the issue in 2008. Apparently some bridges still implement the white paper, while others follow the standard.[3] This note describes the background to, and differences between, these specifications.

Some existing IEEE Std 802.1Q protocol variable descriptions are incorrect. This note proposes maintenance and removal of an unnecessary and ineffective management control.

1. Fragile bridges

Failed fragile bridges were initially described as 'brain-dead'—a failure of spanning tree protocol operation while continuing to forward data frames.[4]

If a router becomes brain-dead, other routers and end stations will notice its lack of participation in routing protocol and stop sending it frames to forward.[5] The operation of a bridge is, by contrast, intended to be largely transparent. A bridge with a spanning tree Designated Port forwarding frames to and from a LAN to which another bridge's Root Port is attached is not necessarily aware of the latter's presence or failure to participate in the protocol,[6] and a data loop that is allowed to develop can persist indefinitely.

A brain-dead router thus has a temporary impact on local network connectivity, while a brain-dead bridge can bring down the whole network pending network management intervention. A router implementation whose complex control plane software can crash might be tolerated, but if bridging capabilities are added the impact can be less acceptable.[7]

2. Detecting live bridges

As part of rapid configuration, an RSTP-capable Designated Port [8] sends BPDUs with a Proposal flag to elicit BPDUs from a Root Port or Alternate Port of any other bridge attached to the same LAN.[9]

The persistent absence of responses to BPDU proposals can be used to identify 'Edge Ports', connecting LANs with no other attached bridges.[10,11] This is a vital part of rapid configuration: loop preventing 'cuts' in active topologies propagate to Edge Ports, which need not delay or interrupt connectivity to end station(s) to prevent data loops.

An absence of responses can also be used to detect the failure of a neighbouring bridge. Clearly a Bridge Port cannot be configured so that the apparent absence of a peer bridge will cause it to become Forwarding (on the supposition that it is an Edge Port) while at the same time allowing that apparent absence to cause it to stop Forwarding (on the supposition that its immediate neighbour bridge has experienced a partial failure). This note reviews management controls, which default to providing Edge Port detection.

---

[1] Fragile bridge detection uses procedures common to RSTP, MSTP, and SPB. See Clauses 13.23, 13.27.19, 13.27.27, 13.31, 13.33, Figure 13-25, and AutoIsolate/autoIsolatePort, isolate/isolatePort in IEEE Std 802.1Q-2022.

[2] https://www.ieee802.org/1/files/public/docs2008/aq-seaman-merged-fragile-bridges-0908.pdf

[3] Personal communication from Fred Gruman following interoperability testing by the OpenROADM group. Much appreciated.

[4] https://www.ieee802.org/1/files/public/docs2008/new-nfinn-mstp-issues-0908-v1.pdf

[5] Any bandwidth consuming looping of routed packets (mitigated more or less by packet time-to-live), can only persist until these other systems notice its lack of participation in routing protocol and develop a coherent view of the new topology.

[6] A bridge's spanning tree Root Port forwards frames but is not required to transmit periodic BPDUs. This was a deliberate design decision. The bandwidth consumed by the spanning tree protocol on any LAN is always a small fraction of the total, even if many bridges are connected to a shared-media LAN. Item (h) in 13.1 of IEEE Std 802.1Q-2022, item (f) in 8.1 of IEEE Std 802.1D-1998.

[7] The probability of an individual forwarding system crashing may seem minimal, but a single network might contain hundreds if not thousands of bridges, and a supplier might field maintenance reports for tens of millions.

[8] A Bridge Port whose role in the active topology is to forward frames to and from the network's spanning tree root.

[9] An STP Bridge Port (IEEE Std 802.1D prior to 2004) does not transmit BPDUs unless it is the Designated Port.

[10] The common assumption is that the port is connected, by a point-to-point LAN to an end station (which can be a router port). However the Edge Port detection logic also covers the case of connection to a shared media access LAN with zero or more attached end stations and bridges (the transition to PRT:DESIGNATED_PROPOSE is not qualified by operPointToPoint).

[11] A bridging function that cannot be part of a loop, e.g. one supporting a number of logical end stations within a server with only one port connected to the rest of the network or that does not bridge between its networks port's, need not operate a spanning tree protocol and the bridge port(s) within the network that provide(s) LAN connectivity to that bridging function should be identified as an 'Edge Port'.

3. Edge Port detection

The state machine variable operEdge is set (TRUE) if a port has been identified as an Edge Port. Its setting is managed by AdminEdge and AutoEdge, as follows.

If AdminEdge is TRUE, and the state machines are initialized or re-initialized (BEGIN) or the port is disabled (!portEnabled), operEdge is set. When operation begins, and the port is enabled, it will be a Designated Port and start Forwarding immediately.[12] The port will continue to set the Proposal flag in transmitted BPDUs, and operEdge will be cleared (FALSE) if it subsequently receives a BPDU (of any sort, and any priority). If the received BPDU is of worse priority, but does not claim to be Designated and Learning, the port will continue Forwarding.[13]

If AdminEdge is FALSE, a newly selected Designated Port will not become Forwarding immediately, but proposals will be transmitted [14] and, if AutoEdge is TRUE, operEdge will be set when the timer edgeDelayWhile expires (3 seconds) and the port will become Forwarding.

If the port receives an agreement,[15] the state machine variable proposing will be cleared and the Proposal flag not be set in subsequent BPDUs. If the port remains Designated, but becomes Discarding as consequence of propagating a loop preventing cut in the active topology with a worse spanning tree priority, proposing will be set again, and if AutoEdge is TRUE and no responses to proposals are received before an edgeDelayWhile timeout, operEdge will be set again.

Figure 1 shows the possible combined states of operEdge, forwarding, and proposing for a Designated Port with AdminEdge or AutoEdge set. Names reflect, but do not completely describe, reasons for each state. Figure 2 shows transition conditions, using the standard's state machine variables.[16] A sequence of state names provides a change history (trace) for operEdge, forwarding, and proposing.

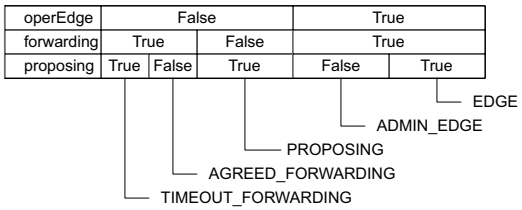

Figure 1—Designated Port states (not isolating)

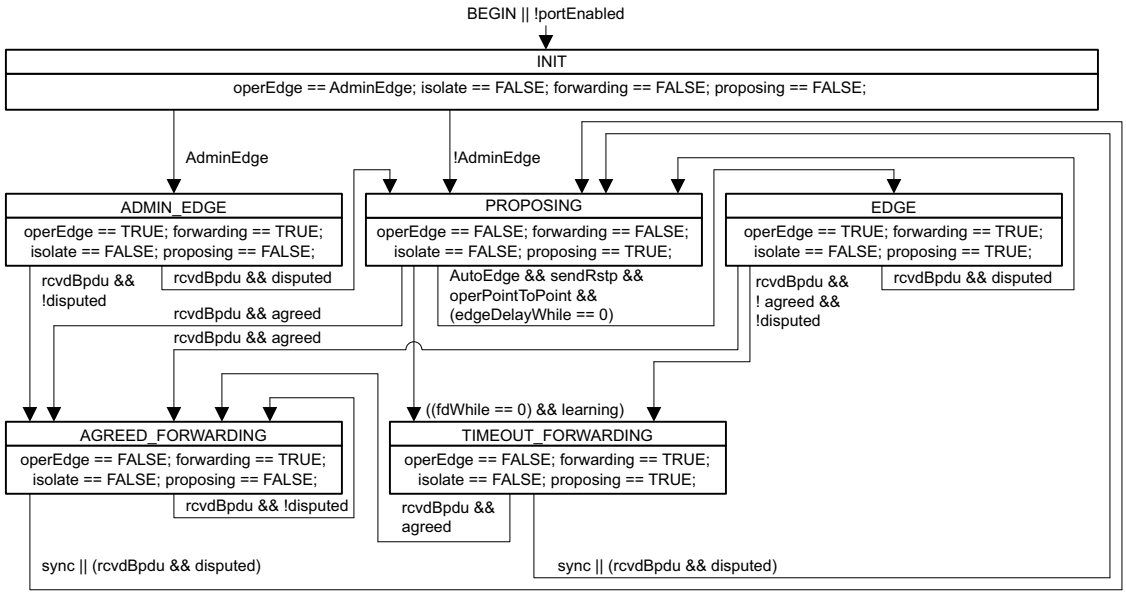

Figure 2—operEdge, forwarding, and proposing transitions for a Designated Port with AdminEdge or AutoEdge set

The scenarios described above have traces: INIT> ADMIN_EDGE>AGREED_FORWARDING, INIT>PROPOSING>EDGE, and INIT>PROPOSING>AGREED_FORWARDING>PROPOSING>EDGE. BEGIN corresponds to initialization of the state machines. The trace for a previously enabled Root, Alternate, or Backup Port will start with PROPOSING.

[12] PRT:DESIGNATED_LEARN>PRT:DESIGNATED_FORWARD.
[13] Otherwise the disputed flag will be set, and the port will become Discarding (possibly temporarily).
[14] Unless an STP BPDU is received, indicating the presence of a bridge that is not RSTP-capable.
[15] A BPDU of worse priority, received on a point-to-point link from a Root or Alternate Port with the Agreement flag set, setting agreed for the port.
[16] The notation used is similar to, but is not, that of the RSTP state machines. Each descriptive 'state box' contains variable values, not actions. The figure does not explain how transitions conditions occur, and omits intermediate state machine transitions. This figure is not a definitive description of RSTP, and should not be used to generate conformance tests.

4. Fragile bridge detection

Fragile bridges whose spanning tree operation has failed, are expected to continue to filter spanning tree messages (BPDUs), as that is part of normal data plane operation. The BPDUs that a potentially fragile bridge transmits on a LAN attest to its continued control plane spanning tree operation, and are not simply forwarded from another bridge.

If AdminEdge and AutoEdge are both FALSE, a persistent absence of responses to proposals on a point-to-point link is understood to indicate the failure of a connected fragile bridge's spanning tree protocol entity: isolate is set,[17] and the port stops Forwarding (becomes Discarding). The port will continue to transmit proposals and isolate will be reset (FALSE) if any BPDU is received.[18]

The principle difference between my white paper and the standard is when proposals are transmitted. This difference reflects different goals.

In the white paper, the goal is to detect the failure within a few seconds, and isolate that bridge,[19] even if the failure has no immediate impact on the network.

In the standard, the goal is to prevent data loops and consequent network meltdown. The port can only create a data loop if the port becomes a Forwarding Designated Port (having previously been an Alternate Port, or a Root Port) or remains Forwarding (having previously been a Root Port). It might become part of a data loop created by another port (elsewhere in the network) if it persisted in Forwarding as a Designated Port while propagating worse spanning tree information.[20]

The following description of the standard is based on its state machines, which take precedence over information in general text and variable descriptions.[21] Possible changes are discussed under Maintenance.

Figure 1 shows the possible combined states of operEdge, isolate, forwarding, and proposing for a Designated Port.

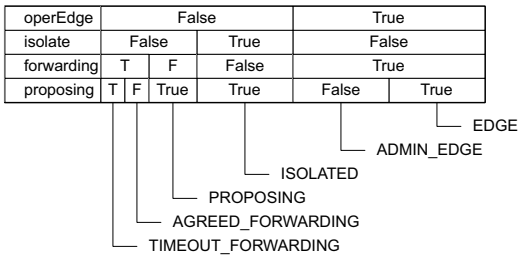

Figure 3—Designated Port states

Figure 4, showing transition conditions is based on proposed Clause 13 text I prepared in November 2008,[22] with the standard's goal of loop prevention. It pre-dates the addition of AutoIsolate and is a useful basis for explaining the white paper and the final standard. All three specifications provide failed fragile bridge isolation if, and only if, AdminEdge and AutoEdge are both FALSE. This leaves operEdge detection (enabled by setting AutoEdge) unchanged.

NOTE—If AdminEdge is set, and AutoEdge is clear, the port will initially be Forwarding without sending proposals or receiving an agreement. It will also remain Forwarding if a loop preventing cut in the active topology is propagated to the port. However once a BPDU has been received, further loop preventing cuts will force a transition to Discarding, with Forwarding only resuming after an agreement is received or the Forward Delay timer expires. The port will not be isolated.

[17] BDM:ISOLATED.
[18] PRX:RECEIVE.
[19] Strictly speaking, to transition the port that has detected the failure to Discarding, cutting its connectivity to a LAN connected to the failed bridge.
[20] Which is why the state machine variable sync is set and agreed is cleared, causing a port that is not an Edge Port to become Discarding until it receives an Agreement to the new information from its neighbour or (if fragile bridge isolation is not enabled) fdWhile expires (indicating that the information has been disseminated throughout the network).
[21] See the last paragraph of the introductory material of Clause 13 of IEEE Std 802.1Q-2022, and the initial paragraphs of 13.26 and 13.27.
[22] https://www.ieee802.org/1/files/public/docs2008/aq-seaman-merged-spanning-tree-protocols-1108.pdf, changes also in P802.1aq/D1.5 (December 2008).

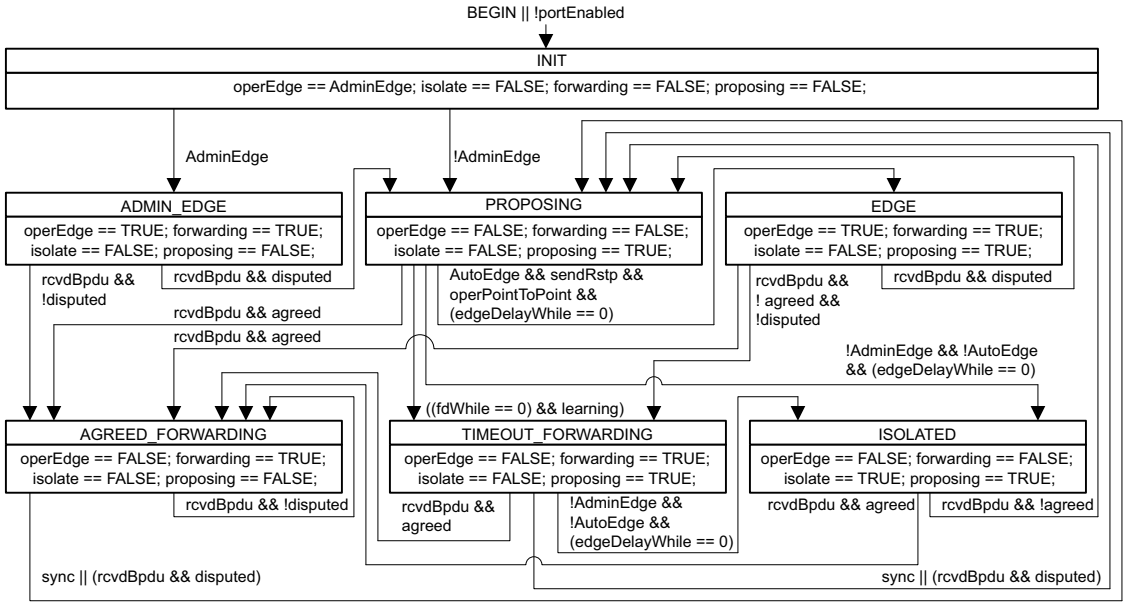

Figure 4—operEdge, isolate, forwarding, and proposing transitions for a Designated Port (P802.1aq/D1.5)

The white paper differs from Figure 4 (and the 2005 standard), by not clearing proposing in recordAgreement() it avoided AGREED_FORWARDING in favour of TIMEOUT_FORWARDING.[23]

The white paper also set proposing in PRT:DESIGNATED_PORT: unnecessarily as a newly Designated Port will be or will be made Discarding, and PRT:DESIGNATED_PROPOSE will set proposing.[24]

The behavior specified in the current standard only differs from Figure 4 if AutoIsolate is TRUE, with operPointToPoint and sendRstp set, and AdminEdge and AutoEdge both clear. Transitions setting forwarding then go to TIMEOUT_FORWARDING rather than to AGREED_FORWARDING.[25] This just delays clearing proposing (transitioning to AGREED_FORWARDING) until a second Agreement is received. AutoIsolate's impacts on bridge failure detection is thus negligible.

5. Fault coverage and remediation

The fragile bridge failure detection specifications in the standard and in the white paper both avoid transitioning a Designated Port to Forwarding if the neighbouring bridge's spanning tree protocol entity, or bidirectional control plane connectivity between the port and that entity, has failed.[26, 27]

The white paper also protects against the possibility of a port on a third bridge, connected to the fragile bridge but not configured for fragile bridge failure detection, becoming Designated and Forwarding (as a result of not receiving BPDUs from the failed bridge) and creating a data loop while the first Designated Port remains forwarding. However, the cost of this additional loop protection is not just the ongoing requirement for the monitored bridge to transmit responses to persistent proposals. Terminating preexisting connectivity means that the failed bridge will not be able to use the network as part of recovering from a failure,[28] until its spanning tree operation has been restored. In general, clearing (if at all possible) AdminEdge and AutoEdge thus enabling fragile bridge loop prevention on all bridge ports that are known to connect to other bridges is preferred.[29]

[23] Not a great name in this case, since agreements still transition to the state, the definition of the state is its combination of characteristic variables.
[24] If previously an Alternate or Backup Port it will be Discarding, if a Root Port, (reRoot && (rrWhile!=0) will transition to PRT:DESIGNATED_DISCARD.
[25] Because PRT:DESIGNATED_FORWARD transitions to PRT:DESIGNATED_PORT.
[26] The previously standardized dispute mechanism defends against one way control plane connectivity.
[27] There is no suggestion that coverage can be extended beyond non-operation of the protocol, and its transmission and reception of BPDUs, to faulty (i.e. unspecified) operation. The latter is a fool's errand for any protocol.
[28] Some systems implement(ed) 'dead man's handle' mechanisms, where a perceived operating system failure results in fall-back to a minimal trusted kernel which then recovers the system.

6. In-service upgrades

Once Figure 4's AGREED_FORWARDING state has been reached, the current standard (and the November 2008 draft specification) will keep forwarding TRUE, for a Designated Port even if the port's immediate neighbour's spanning tree protocol operation is suspended, provided there are no network topology changes (between the port and the network's spanning tree Root) that cause a loop preventing active topology cut to propagate to the port.

Contrary to what is said in the last sentence of 13.23 'Fragile Bridges', the setting of AutoIsolate and the management controls in Clause 12 that reference it have no bearing on this behavior and it is not mentioned in 13.22 'In-service upgrades'.

A port that is not receiving BPDUs from its neighbour, unless configured as a Layer 2 Gateway Port (13.40 of IEEE Std 802.1Q-2022),[30] will be (or will become when previously received information is timed out) a Designated Port. When its immediate neighbour is about to be the subject of an in-service upgrade, and can be expected to stop sending BPDUs, the network administrator can:

a) Set both AdminEdge and AutoEdge, so that the port will continue forwarding (with the briefest of interruptions) even if attached to a point-to-point link that is goes down (!portEnabled) temporarily as the neighbour completes its upgrade.

or

b) Clear both AdminEdge and AutoEdge, so that if the port stops forwarding it will not become forwarding again until it hears from its neighbour.

The first of these is an obvious choice if the port is a Designated Port before the upgrade begins. The second will stop a prior Alternate Port, from becoming a Forwarding Designated Port and creating a loop (provided it is not attached to a shared-media LAN, or the neighbour's upgrade completes within twice ForwardDelay).

Once the neighbour's upgrade is complete, AdminEdge can be reset and AutoEdge set to restore monitoring of the potentially fragile neighbour.

7. Specification differences

The principal differences between the current standard's specification for fragile bridge failure detection and that in the initial white paper are a consequence of their different goals: loop prevention vs immediate detection of spanning tree non-operation, as described above.

This section summarizes the reasons for all the Clause 13 differences relevant to fragile bridge detection. As compared to the white paper, the standard:

a) Did not change PIM:UPDATE to avoid clearing proposing when AdminEdge and AutoEdge are both FALSE.

   Reason: goal change as described above.

b) Did not change recordAgreement() to avoid clearing proposing when AdminEdge and AutoEdge are both FALSE.

   Reason: goal change as described above.

c) Did not use a different value for edgeDelayWhile() [5 seconds suggested] for isolation as opposed to edge detection [MigrateTime, 3 seconds].

   Reason: For most implementations 3 seconds should accommodate the potential loss of two messages, but more significantly the goal change means that proposing will be set far less frequently (once cleared by an agreement, it will remain clear until a further network topology change occurs) so there is much less chance of a loss interrupting connectivity. If isolate is incorrectly set, proposing will remain set, soliciting a further agreement that will clear isolate and restore connectivity.

   NOTE—Unrelated maintenance advised. See below.

d) Added the AutoIsolate management variable, and required it to be set as part of the additional setting of proposing in PRT:DESIGNATED_PORT.

   Minor error compounded: As pointed out in the discussion of Figure 4 above, the original setting of proposing in this state was and is not required.

   Unfortunately the standard's management specification in Clause 12 builds upon this error. See maintenance below.

[29] And, of course, the white paper's operation cannot defend against loop creation by a more distant bridge.

[30] I have not checked L2GP functionality as part of preparing this note. It was added by IEEE Std 802.1ah–2008, but revised in P802.aq and IEEE Std 802.1Q-REV (the original logic was incorrect).

8. MSTP considerations

In the white paper and the standard, a Designated Port sets the per port variable isolate [item i) in 13.27] when BPDUs, whose transmission has been solicited by proposing, are not received..

The standard improves on the white paper by not disrupting connectivity unnecessarily if a fragile bridge's control plane fails. A previously viable management path (with forwarding TRUE on a Designated Port that is configured to detect the failure) will remain viable as long as that is compatible with preventing loops. Prior to the failure occurring, that port could have been a Designated Port for the CIST and some MSTIs, but a Master Port, Root Port, or Alternate Port for other MSTIs. The failure will result in it timing out spanning tree information received on the port, and becoming a Designated Port for those other MSTIs with forwarding clear and proposing [item bh) in 13.27, with one instance per port for each tree (CIST, MSTI, and SPT)] set. If isolate is not set for the port, those MSTIs will become forwarding after twice FwdDelay.

The Bridge Detection Machine (BDM, 13.33) is responsible for setting isolate. It is a per port (not per tree) state machine but uses the per tree variable proposing without further clarification. For RSTP this is fine, as there is only one tree to consider, for MSTP isolate needs to be set if proposing is set for the CIST or any MSTI. This needs to be clarified in the standard.

A. Maintenance

The analysis above has identified items that need correcting in the standard's description of RSTP/MSTP and its management.

The standard itself is clear that the state machines and their associated procedures (13.28–13.40) take precedence over its description of timers and per bridge and per port variables (13.25–13.27),[31] that 13.24–13.40 take precedence over the general description (13.1–13.23) in Clause 13,[32] and that Clause 13 takes precedence over text in other parts of the standard.[13] This order of precedence is vital as it can be very hard to make a brief description of a variable or other piece part wholly accurate, as accuracy would require explaining its relationship to or combination with other parts of the specification. An attempt at a complete description of a single item can thus pull in most of the specification as each additional item is further qualified. A brief, but necessarily incomplete, description can easily become an unrealistic or unnecessary management control.[33]

Questions of precedence aside, rewording is always a sensitive issue, and can require significant discussion. We have a present forcing function in the form of the SA Ballot of P802.1Qdy 'YANG for the Multiple Spanning Tree Protocol'. We cannot, both as a matter of conscience and as a matter of not creating unnecessary future work for both ourselves and users of the proposed YANG model, standardize a model that is known to be broken. At the same time there is an urgent need for this model, quite independent of the fragile bridge issue.

I suggest that we split the necessary maintenance activity in two. The first part concerns P802.1Qdy, where our objective should be to avoid reference to current incorrect text and to avoid including anything that will need to be removed or changed when the rest of the standard has been corrected. The second part concerns the existing standard text, and should be considered as part of the present revision, P802.1Q-2022-Rev. In the following I have split the latter into several categories.

A.1 P802.1Qdy

The following suggested change will be part of an SA Ballot comment.

———

Comment

The leaf auto-isolate-port in the proposed RSTP YANG module does not in fact control automatic setting of isolate-port as stated in its description. The operation of the standard's Clause 13 state machines (which are clearly identified as taking precedence over other text in the standard) can set isolate-port if admin-edge-port and auto-edge-port (both in the module) are both true, independent of the setting of the corresponding AutoIsolate variable (referenced in the description of the leaf by 13.27.19). If admin-edge-port or auto-edge-port is false then isolate-port will not be set. The leaf auto-isolate-port serves no useful purpose. The accompanying file dy-seaman-fragile-bridges-0624.pdf provides supporting detail.

Suggested Remedy

Remove the leaf auto-isolate-port from the RSTP module, from Figure 48-24, and from the 48.5.26 schema.

———

NOTE 1—Removal of the auto-isolate-port leaf now will be less of a burden on implementers and other users of the standard than deprecating it in the future. An alternative would be for setting auto-isolate-port to have the side-effect of clearing AdminEdge and AutoEdge, while setting either of those would clear auto-isolate-port, but that would leave open what clearing auto-isolate-port should do (set followed by clear will not return AdminEdge and AutoEdge to their original values), and the discussion of wording, retention of history etc. is bound to be protracted. Anything other than simple removal of the leaf (which can be subsequently adjusted if necessary for a particular implementation by an augment, without revising the module) implies a constraint on, or risk of revision following, future discussion.

NOTE 2—There are really only two basic questions when configuring a port attached to a point-to-point link to take into account the possibility that it might not receive BPDUs from a neighbour when it transmits RST (or MST or SPT) BPDUs with the Proposal flag set.

a) Should the port be Discarding or Forwarding on link up or when (re)initialized, i.e. before it can be expected to receive any BPDU.

b) Should the port become Discarding or Forwarding when it has not received BPDUs for some time.

Both these are fully covered by AdminEdge and AutoEdge, there is no prospect of an additional control adding value.

[31] See the first paragraph of each of 13.25, 13.26, and 13.27.
[32] See the last paragraph of the introductory text of Clause 13 'Spanning tree protocols'.
[33] A 'don't bother me with details, these things have to be simple, just make it so'.

A.2 802.1Q maintenance—Clause 13

13.23 Fragile Bridges—Use of AutoIsolate

Comment

The last two sentences of 13.23 say (in reference to fragile bridge isolation):

> "The CIST Proposal flag is set in all BPDUs transmitted by a Designated Port, and used to solicit an Agreement from the neighbor (which might otherwise not transmit). This capability is controlled by the AutoIsolate (13.27.19) variable, is disabled by default to allow for in-service upgrades, and is only effective if the neighboring Bridge is capable of RSTP or MSTP operation."

These statements are not true, or misleading:

a) The Proposal flag is only set in all BPDUs until an agreement is received from the neighbour. If AutoIsolate is set (and AdminEdge and AutoEdge clear) it will be set once more after the port becomes Forwarding, and then will be cleared by receipt of the next agreement received. The Proposal flag will then only be set again if there is a change in the network topology that forces the port to Discarding.[34]

b) The Proposal flag is set whenever a Designated Port is Discarding (possibly as a result of propagating a loop preventing cut in the active topology), as part of its attempt to become Forwarding, and the port will remain Discarding (in the absence of received agreements) if it is attached to a point-to-point link and both AdminEdge and AutoEdge are clear. A port with operPointToPoint clear will transition to Forwarding after FwdDelay

c) If a port's neighbouring bridge is to be upgraded in-service, and the port is an Alternate Port then AdminEdge and AutoEdge need to be clear (selecting rather than disabling the isolation capability) so that the port does not become Forwarding (it will inevitably become Designated, after the prior received spanning tree information is time out).

Suggested Remedy (outline/draft)

Remove the reference to AutoIsolate, substitute a reference to AdminEdge and AutoEdge, remove in-service upgrade text (any new text on this subject should be part of 13.22 'In-service upgrades').

13.23 Fragile Bridges—Reference to 'this revision'

Comment

The reference in this clause to 'this revision' and to 'prior revisions' were applicable to IEEE Std 802.1Q–2011 and its predecessors respectively.

Suggested Remedy (outline/draft)

As with any other text that can be rendered historic by a future revision, this text should be updated to refer to specific revisions or their predecessors.

Figure 13-13 State machine overview

Comment

Fragile bridge detection and the accompanying isolation of non-responding bridges occurs independently of AutoIsolate, which should be removed from the specification.

operPointToPoint is used by several state machines but not shown in the figure.

Suggested Remedy (outline/draft)

Remove AutoIsolate from the figure, and add operPointToPoint showing where it is used.

13.27 Per port variables

Comment

AutoIsolate provides no useful functionality. It should be removed from the list, or (if it is thought desirable not to renumber list items, or to retain the item for information with an update to note its removal) marked as deprecated.

Suggested Remedy (outline/draft)

As per comment.

13.27.19 AutoIsolate

Comment

AutoIsolate provides no useful functionality. This clause (13.27.19) should be removed, or updated to show it as removed by the present revision. While an explanation of the reason for removal might seem desirable, it may be very hard to draft and agree suitable brief text. It would be better simply to state that the capability previously associated with the variable was (and continues to be supported by the appropriate setting of AdminEdge and AutoEdge.

Suggested Remedy (outline/draft)

As per comment.

[34] The port will not become Discarding if it was a Designated Port and Forwarding with operEdge set.

Figure 13-25 Port Role Transitions—Designate Port

Comment

AutoIsolate provides no useful ‘functionality’. The PRT:DESIGNATED_PORT conditional setting of proposing ‘if {cist} ..’ only serves to require two rather than one BPDU from its neighbouring bridge when this port transitions to Forwarding.

Suggested Remedy (outline/draft)

Remove the conditional setting of proposing from PRT:DESIGNATED_PORT.

Figure 13-25 Port Role Transitions—Designate Port

Comment

AutoIsolate provides no useful ‘functionality’. The PRT:DESIGNATED_PORT conditional setting of proposing ‘if {cist} ..’ only serves to require two BPDUs rather than one from its neighbouring bridge when this port transitions to Forwarding.

Suggested Remedy (outline/draft)

Remove the conditional setting of proposing from PRT:DESIGNATED_PORT.

A.3 802.1Q maintenance—Clause 12

12.8.2 Bridge Protocol Entity

Comment

AutoIsolate provides no useful ‘functionality’. The PRT:DESIGNATED_PORT conditional setting of proposing ‘if {cist} ..’ only serves to require two rather than one BPDU from its neighbouring bridge when this port transitions to Forwarding.

Suggested Remedy (outline/draft)

Remove autoIsolatePort [item n) in 12.8.2.1.3, item f) in 12.8.2.3.2].

NOTE—isolatePort appears in the list of CIST port parameters while it needs to apply to all trees, but there is no other useful place to put it, and grouping this information with CIST information is harmless.

A.4 802.1Q maintenance—Clause 17

IEEE8021-SPANNING-TREE MIB

The IEEE8021-SPANNING-TREE MIB and its current description in Clause 17 requires extensive maintenance, going well beyond the subject of this note. 17.2.3 ‘Structure of the ... MIB’ provides history, and states that the MIB is based on 802.1D-2004, although AdminEdgePort is included and an extension table (see Table 17-5 ...MIB structure) includes AutoEdgePort, AutoIsolatePort, and IsolatePort.

It remains to be seen whether anyone has the time and is prepared to update this MIB. It is, for example, beyond time to get rid of the ‘broken’ port state and use the ifStack correctly, for example.

A suitable starting point for updating this MIB, correcting references and anomalies, would be the RSTP YANG module once standardized.

IEEE 8021-MSTP MIB

The object ieee8021MstpCistPortAutoIsolatePort should be deprecated in this MIB. However if the MIB is to be revised it would be better to create a new MIB, starting with the MSTP YANG module, as the references in the latter are more accurate.

A.5 802.1Q maintenance—PICS

A.10 RSTP-23

This PICS item should be removed.

A.10 RSTP-23

This PICS item should be removed.

B. History

In the course of preparing this note I have searched the following:

—The publicly available documents on the IEEE 802.1 web server in the docs2008, docs2009, doc2010, and docs2011 directories.

—The publicly available IEEE 802.1 minutes for the period 2008 through 2011.

—The draft directories for P802.1aq and the revision of IEEE Std 802.1Q published in 2011. These directories include not only all drafts made available to the working group, but also all distributed proposed and final dispositions of comments.

—My own emails and possibly related directories prior to 2011.

NOTE—My personal records are not a complete history, and in fact added little to this history. To facilitate future searching and manage disk space, I routinely delete duplicates, prior emails in retained threads, emails with very large attachments that get stored in source files, and sources that I believe are completely superseded by past events. I replaced my personal PC with a MAC in 2012 and did some clean up at that time. I was not the editor for P802.1aq or for the IEEE Std 802.1Q-2011 revision, and cannot necessarily distinguish between source files that I originated and those I received later as the basis for future work.

The issue of 'fragile bridges', initially described as 'brain-dead' bridges, was raised by Norm Finn in September 2008.[35]

I prepared and distributed a note (referred to as the initial white paper above) addressing the issue in October 2008.[36]

I prepared a proposed update of Clause 13 in November 2008. This contribution included the shift in strategy from setting proposing continually to isolate a fragile bridge as soon as possible, to setting isolate when necessary to prevent data loops and retaining connectivity in the mean time. This update was a contribution, not an editor's draft, so was posted in the public directories. However, as it was based on copyrighted text and could be misinterpreted as a group work product, I retained the IEEE copyright and 'subject to change' footers on each page. The header on each page identifies it as 'P802.1aq/D1.0+suggested changes'.[37]

P802.1aq/D1.5 was distributed for a Task Group ballot 18th December 2008.[38] There were no significant changes to Clause 13 from the D1.0+ contribution. Ballot comments were discussed in the January 2009 interim meeting. No disposition of comments appears to have been posted for this ballot. I have checked the email record of ballot comments, and found none focused on fragile bridge isolation. There was a lot of other work on P802.1aq Shortest Path Bridging at this time, continuing through the following meetings, and I have checked the presentations in docs2009.

The proposed Clause 13 update identified as 'P802.1aq/D1.5++suggested changes' that I distributed in May 2009 included the conditional setting of proposing with AutoIsolate in the PRT:DESIGNATED_PORT state. The analysis in this note shows the latter additions to have been unnecessary (and AutoIsolate ineffective). The D1.5++ changes were included in the next draft, P802.1aq/D2.0, and balloted.

The AutoIsolate additions to P802.1aq were transferred to P802.1Q-REV/D1.1 (revision eventually published as 802.1Q-2011) as a result of a simple ballot comment on the D1.0 ballot.[39]

The text in 13.23 'Fragile bridges' referenced (with recommended removal) in A.2 above, was added following consideration of SA Ballot comments #52 and #53 on P802.1Q-REV/D1.3.[40] The discussion of those comments did not (so far as I am aware) dig into the detailed operation of AutoIsolate.

[35] https://www.ieee802.org/1/files/public/docs2008/new-nfinn-mstp-issues-0908-v1.pdf
[36] https://www.ieee802.org/1/files/public/docs2008/aq-seaman-merged-fragile-bridges-0908.pdf.
[37] https://www.ieee802.org/1/files/public/docs2008/aq-seaman-merged-spanning-tree-protocols-1108.pdf
[38] https://www.ieee802.org/1/private/email2/msg11295.html, https://www.ieee802.org/1/files/private/aq-drafts/d1/802-1aq-d1-5.pdf
[39] https://www.ieee802.org/1/files/private/q-2011-drafts/d1/802-1Q-REV-D1-1.pdf (July 12, 2010).
[40] https://www.ieee802.org/1/files/private/q-2011-drafts/d1/802-1Q-REV-d1-3-sponsor-dis.pdf (January 21, 2011).