

15 SEPTEMBER 2008
IEEE 802.1 - SEOUL



PBB-TE **Protection Switching**

Unidirectional switching scenarios

Oliver Thorp


Overview

- This contribution details scenarios where, following a unidirectional switch, the PBB-TE protection mechanism will forward a Backbone Service Instance on a different TESI in each direction
- 802.1Qay/D4.0 does not provide a mechanism that prevents or automatically corrects this state in all cases
- This issue was raised as ballot comment #196 against 802.1Qay/D3.5

Race condition on signal-fail for co-routed TESI

1) Initial state


- Assume there are two TESI configured in a protection group between BEBs A and F
 - Working TESI is ABCF
 - Protection TESI is ABDEF
- A single physical link (AB) is shared by both TESI
 - This link has failed
- The hold-off timer is disabled
- Both BEBs select the working TESI



Race condition on signal-fail for co-routed TESI

2) The link (AB) now recovers

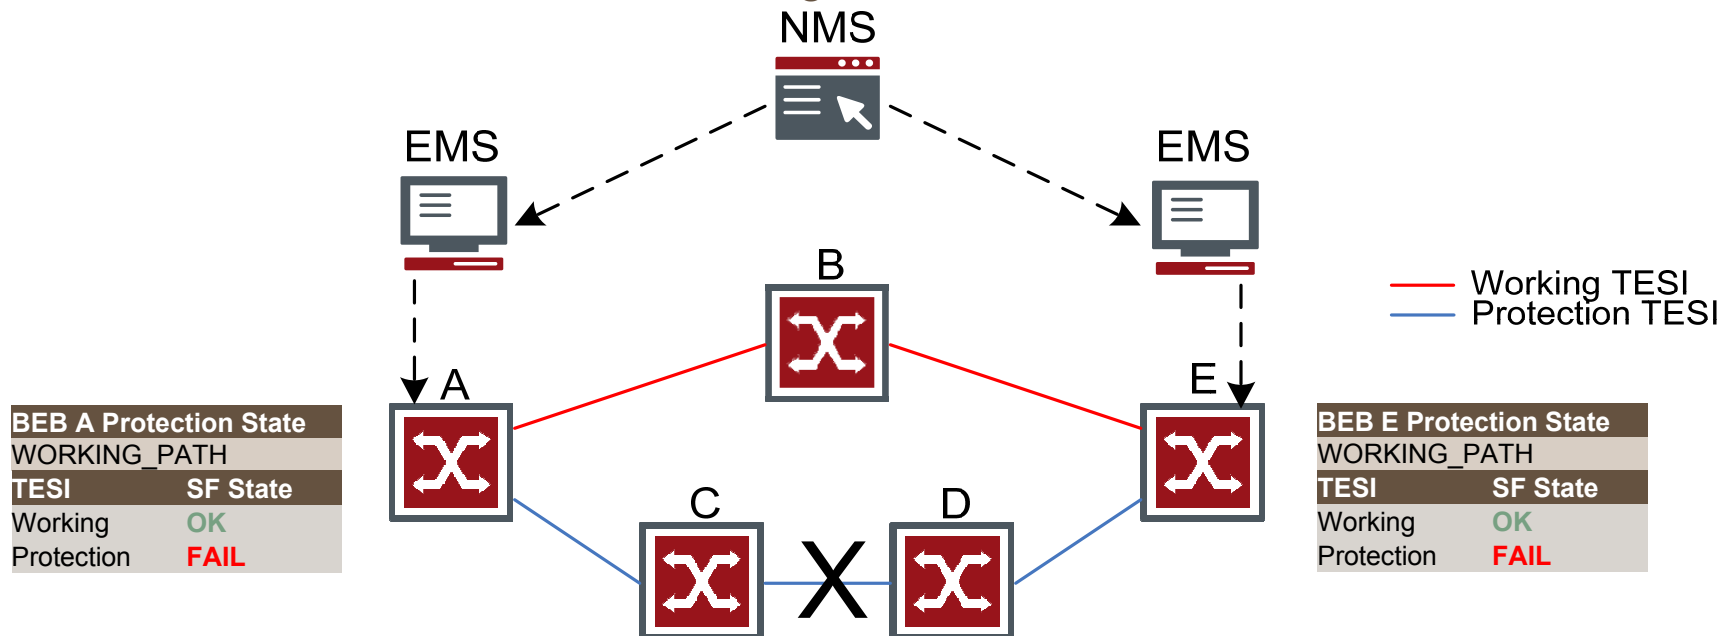
- The protection switching state machines on the BEBs at each end of the protection group will now independently switch to whichever TESI clears the SF condition first...
- In the case shown, there is 50% chance of the two BEBs picking different paths



Race condition between MS & SF

1) Initial state

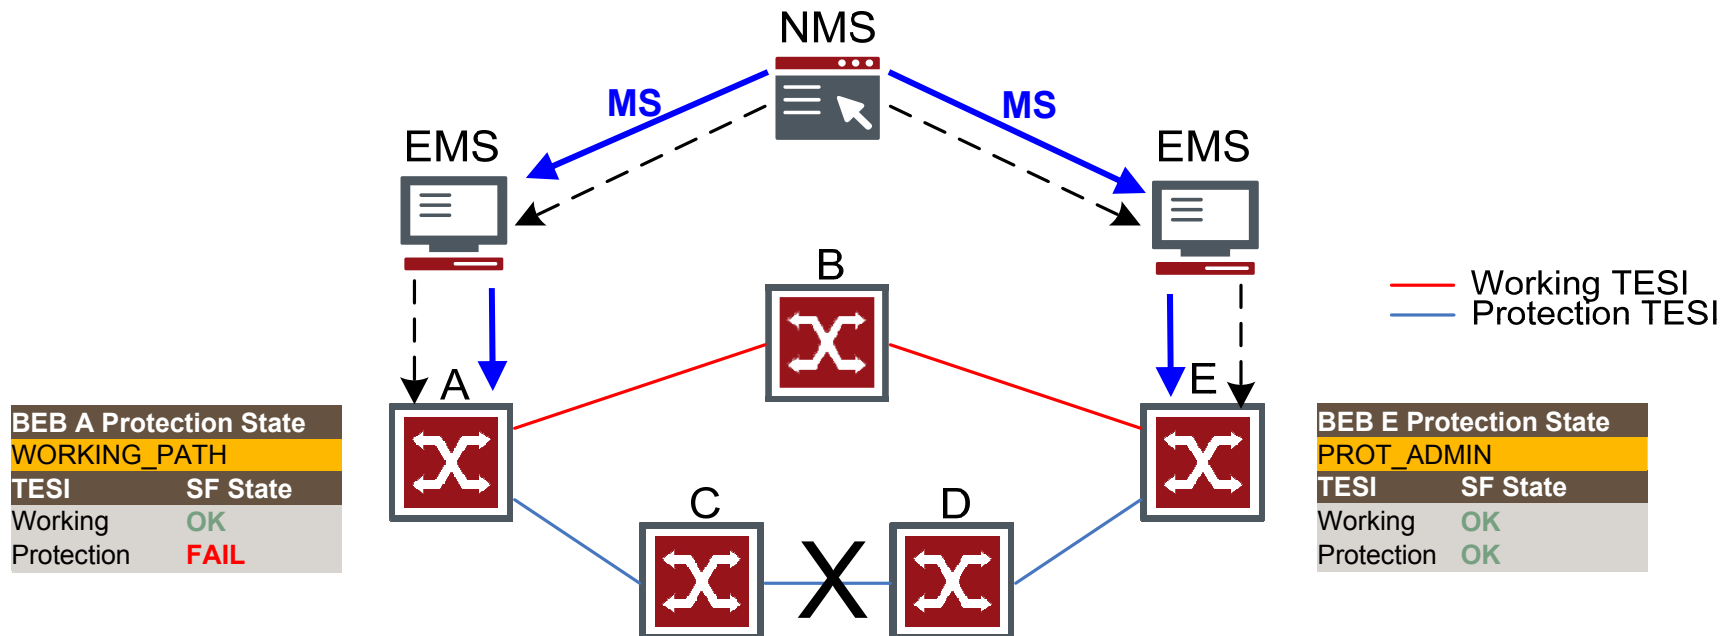
- Assume there are two TESI configurations in a protection group between BEBs A and E
 - Working TESI is ABE
 - Protection TESI is ACDE
- The protection TESI has failed
 - Both BEBs select the working TESI



Race condition between MS & SF

2) Operator initiates a manual switch to protection

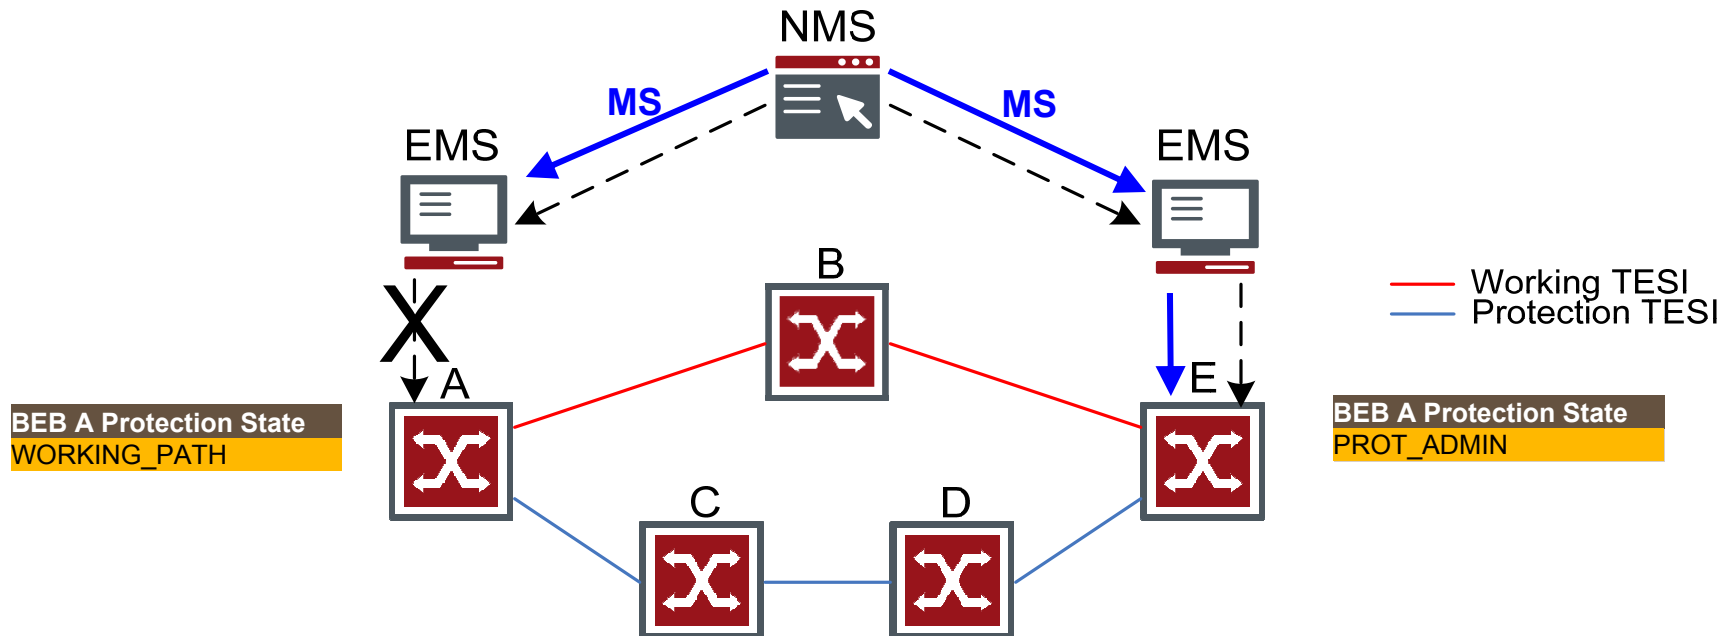
- At the same time the protection TESI recovers
- The order in which each BEB sees the two events will determine the resulting protection state



Loss of management path to BEB

• Scenario

- Assume there are two TESI s configured in a protection group between BEBs A and E
- The operator initiates a manual switch to protection
- The EMS is unable to send the command to BEB A
- BEB E performs a protection switch but BEB A does not



Mitigations

- Configure the hold-off timer to a period greater than the CCM detection time
 - Prevents signal fail race condition in the case shown
- Enable revertive protection switching mode
 - Any divergence will (eventually) clear itself
- Using mechanisms outside the scope of the 802.1Qay standard, verify that a operator request was performed at either both ends of a protection group or none
 - If necessary, apply a second command to one end of the protection group to fix any divergence in protection switching state.

Proposal


- Section 26.10.1 last paragraph, replace first sentence as follows:

"The description herein, defines and provides a scalable end-to-end resiliency mechanism that offers end-to-end 1:1 linear protection switching capable of load sharing for point-to-point TE service instances in a PBB-TE Region. This standard does not enforce bidirectional PBB-TE protection switching, however:

 - the someRDIdefect is used as a protection switching trigger which will normally cause a bidirectional switch; and
 - the Mismatch defect allows the network operator to check the remote mapping of traffic the working or protection TESI and if necessary request a manual switch before performing a maintenance action on one of the TESIs in a protection group."
- Annex Z.1: Add an additional bullet

"x) Bidirectional Switching
1) [McGuire]: item 4)"

to the list of requirements 'Not currently planning to address'



FUJITSU

Backup: How does G.8031 deal with this?

- ITU-T G.8031 does not have an issue with these race conditions since it co-ordinates the protection switching state using an APS protocol
 - G.8031 will select the working transport entity if the two endpoints disagree on the active transport entity and both transport entities are available
 - Eliminates divergence when endpoints of the protection group detect signal fail at different times
 - G.8031 implements operator requests by applying a command to one end of the protection group
 - Avoids race conditions in 802.1Qay where operator requests have only been applied to one end of the protection group