CCM and ECMP Path Latency Variation

Version 01

Stephen Haddock

September 3, 2013

ECMP CCM Timeline Diagrams

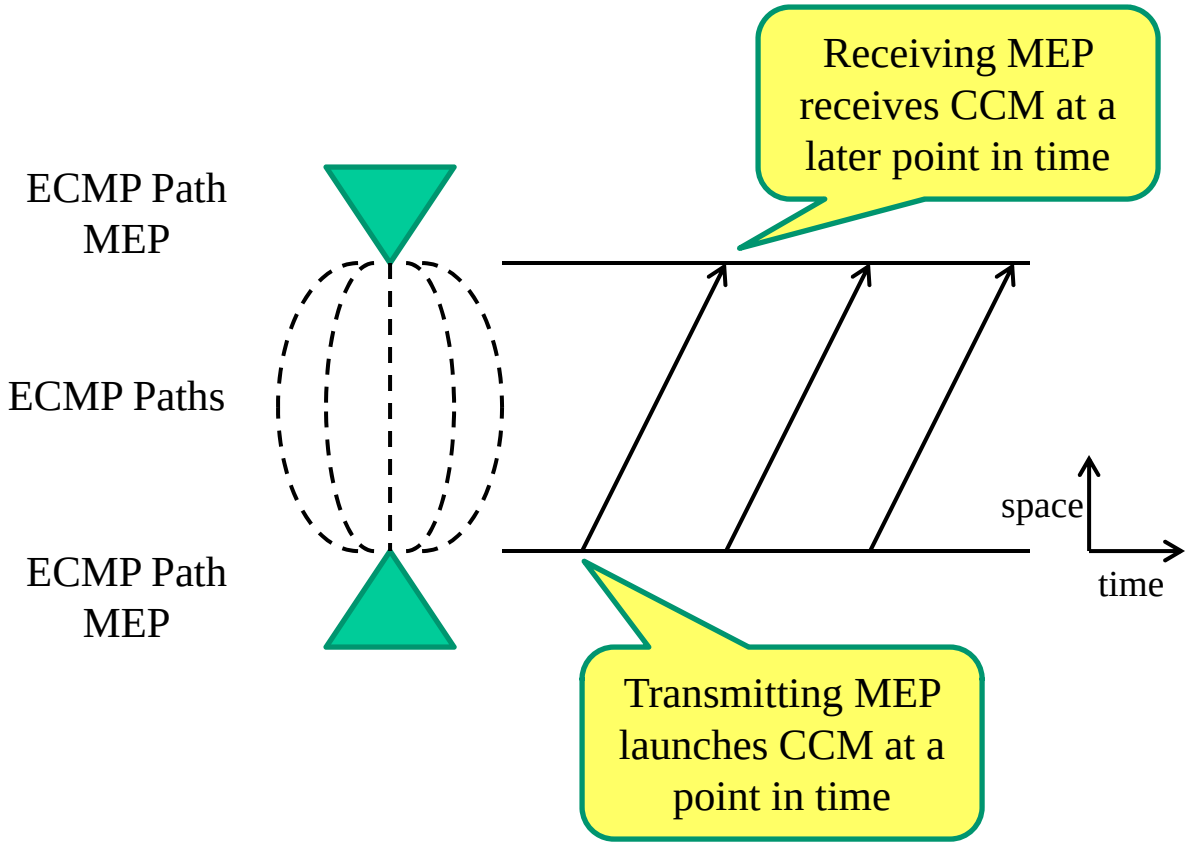

False Positive Fault Detection -- 1

- ECMP Path CCM transmission cycles through different paths (by using different Flow Hash values) every 4 frames.
- If ECMP Paths have significantly different latencies, it is possible to get a false positive detection of a connectivity fault with no loss of CCM PDUs.

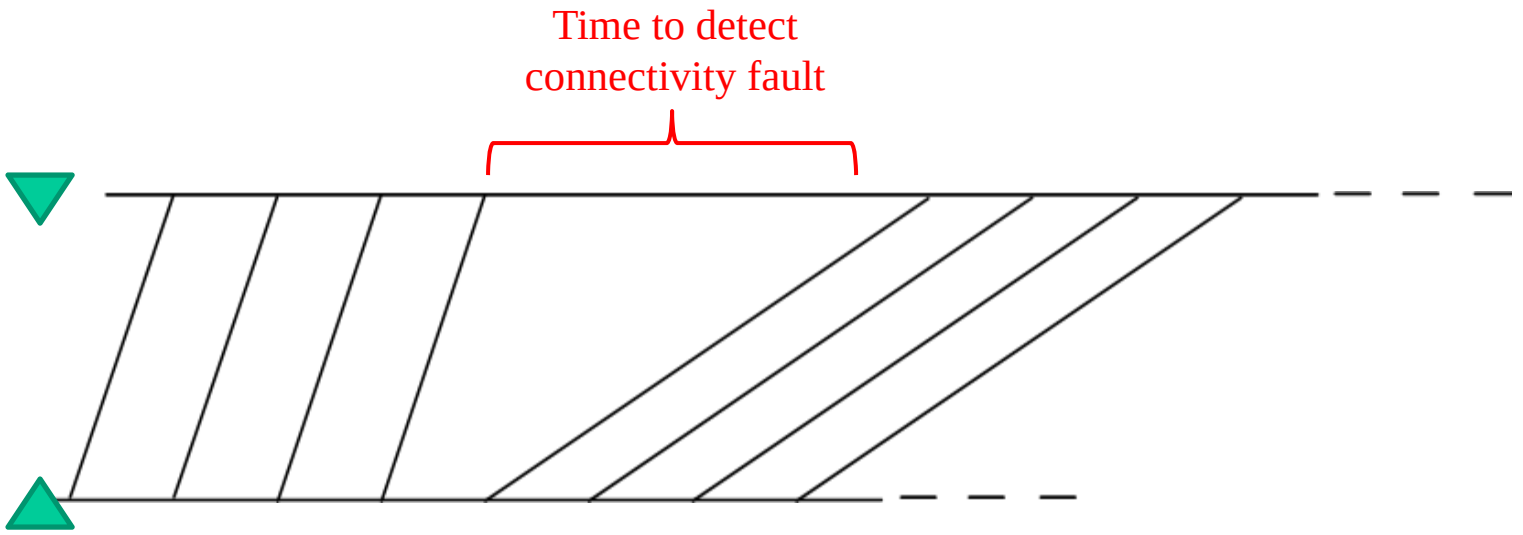

False Positive Fault Detection -- 2

- If ECMP Paths have slightly less significantly different latencies, it is possible to get a false positive detection of a connectivity fault with the loss of a single CCM PDU that is either the last transmission with one Flow Hash or the first transmission with the subsequent Flow Hash.

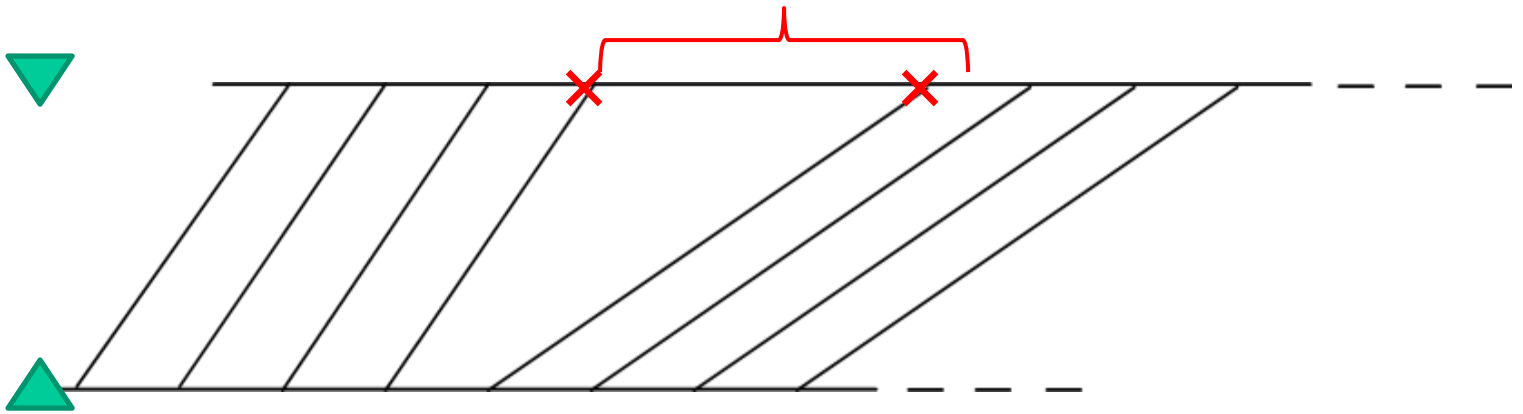

False Negative Fault Detection

- It is possible to loose all CCM PDUs with a given Flow Hash and still fail to detect a connectivity fault.

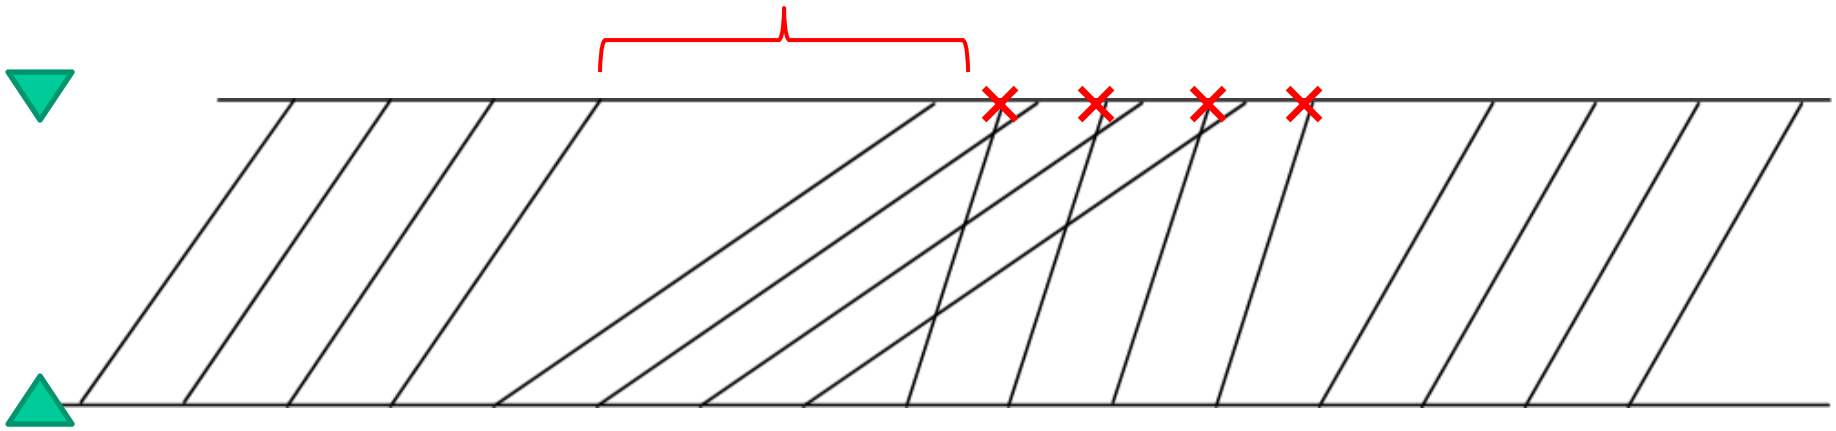

Proposed Resolution

Patient: "Doctor, it hurts when I do this."

Doctor: "Don't do that."

- The false positive and false negative connectivity fault detection can be avoided if the time between CCM transmissions is greater than the greatest difference in latency between any pair of ECMP paths.
  - This may be more constraining than strictly necessary. For example it may be that restricting the time between CCM transmissions to be greater than ½ the greatest difference in latency avoids the false positive and false negative cases. However this constraint assures there will never be a case where two consecutive CCM transmissions are received in the reverse order.
- A note to this effect in 802.1Qbp should be sufficient, no reason for it to be normative.

Thank You

.