Flooding Default Behavior

July18-21, 2005
San Francisco

Mike Takefman

Introduction

- Flooding / Cleave point rules have traditionally been implementation dependant
  - given 802.17-2004 bridging behavior there was no issue
- Implementation of 802.17b into existing bridge hardware changes this
  - bridge learning can cause permissive only behavior in split linecard systems
  - admittedly this is an implementation issue BUT
    - failure to address the issue may reduce market acceptance
    - 802.17b requires split linecard implementations to be possible

Bridge System Design

- 802.17 provides 2x the BW of a “normal” bridge port
  - use of existing bridge chips requires the use of 2 ports to achieve full BW
  - ringlet selection is handled by normal bridge learning but to achieve strict / relaxed behavior
    - bi-directional flooding is used by all ring stations
    - normal bridge behavior would be to flood out both ports anyway
  - bi-flood guarantees shortest path to other stations assuming cleave point is correctly selected

Flooding Behavior

- Typical access ring: stations 3 and 8 are head-ends
  - With uni-directional flooding stations 6-10 do not learn the shortest path to station 1
  - With bi-directional flooding stations 6-10 automatically learn the shortest path to station 1

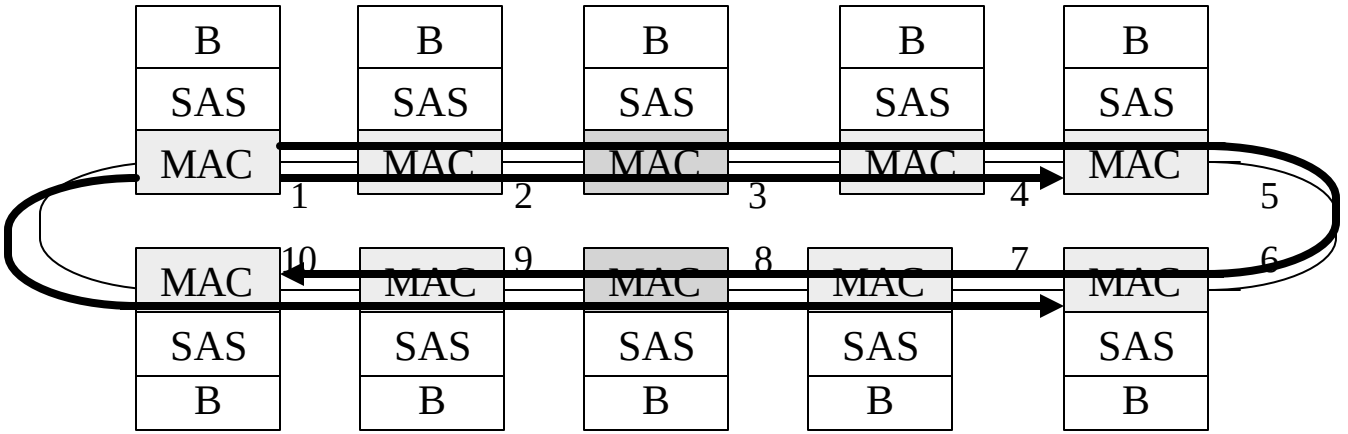

Issue with Uni-Flood

- Assume another 802.17b ring node performs uni-flood
  - bridge will learn the “wrong” path whenever the uni-flood arrives via the longest path
  - strict and relaxed traffic may change paths since a bi-directionally flooding station has sent the frame via the shortest path and then learns the longest path

Why Not Bi-Flood Always?

- If a station supports steering, then it must be able to bi-flood
  - all stations have to support steering
- Why not bi-flood always?
  - performance cost of bi-flood is alleviated for uni-cast traffic as SAS will limit flood duration
    - pay for the bi-directional flood once

Cleave Point Selection?

- The cleave point must be selected consistently in order for conversations to take the same path
  - For rings with an odd number of stations the cleave point is ½ way around the ring (equal number of stations on left and right)
  - For rings with an even number of stations, the cleave point is either to the left or to the right of the opposite station
    - Based on topology, the station with the lower MAC address will send on ringlet 0 and the station with the higher MAC address will send on ringlet 1

Advantages

- Enable lower complexity implementations of SAS onto existing bridge hardware
  - Improves acceptance of RPR
- Insures optimal inter-working of SAS enabled stations

Disadvantages

- The working group forces a particular implementation detail

Proposal

- As part of the TP message, SAS stations can indicate if they prefer bi-directional flooding
- If a station indicates the preference all stations will bi-directionally flood and use the cleave-point calculation as described
- If accepted this will cause change in clauses X and 11.

Straw Poll

- I would vote for a motion requiring this feature
  - Y: N: A
- I would vote against a motion requiring this feature
  - Y: N: A: