Flooding, Cleave Points, Learning

November 14, 2005
Vancouver, BC

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 line-card systems
  - admittedly this is an implementation issue BUT
    - failure to address the issue may reduce market acceptance
      - forklift upgrades of existing systems annoys customers
    - 802.17b requires split line-card implementations to be possible

Bridge System Design

- 802.17 provides and requires 2x the BW of a "normal" bridge port
  - use in existing bridges requires the use of 2 ports to achieve full BW

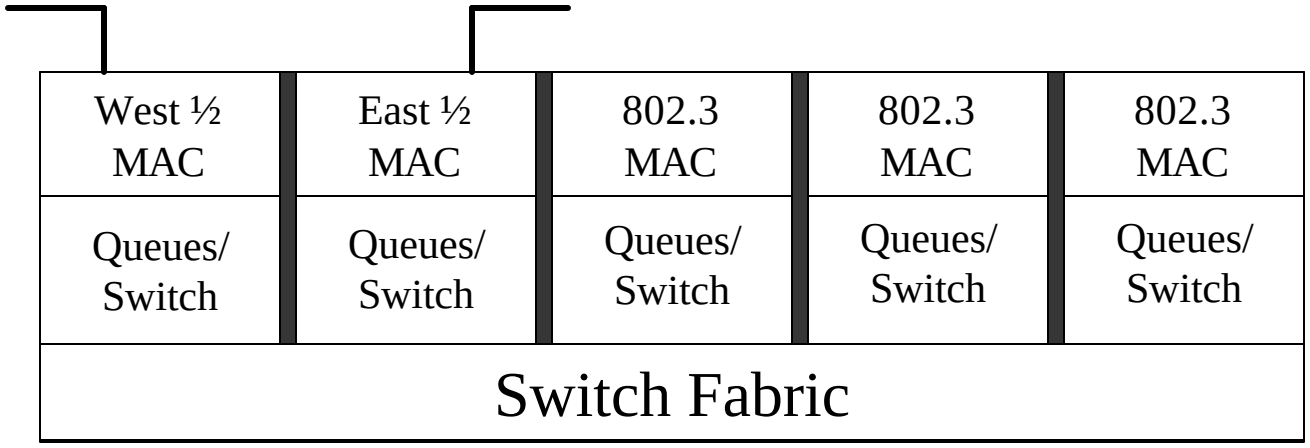

Bridge System Design

- When retrofitting into existing bridges that are completely 802.3 centric it is easiest to considering each ringlet as a separate port
  - bi-directional flooding comes for free
  - ringlet selection occurs based on normal learning

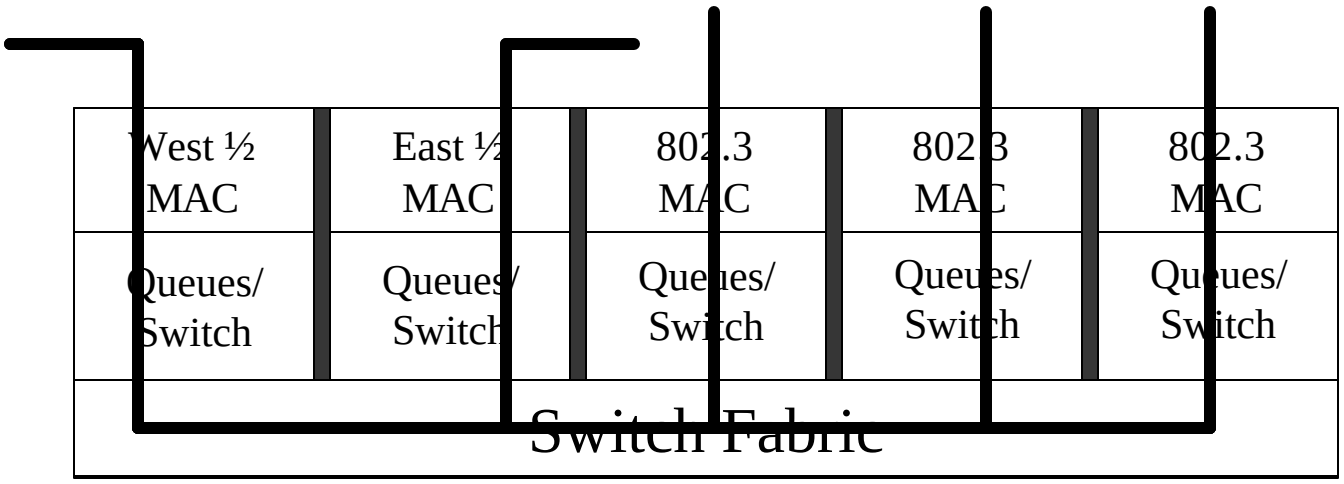

Bi-directional Advantage

- Bi-flood guarantees shortest path to other stations
  - assuming cleave point is correctly selected
  - optimal path is learned

Non Homogeneous Flooding

- 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

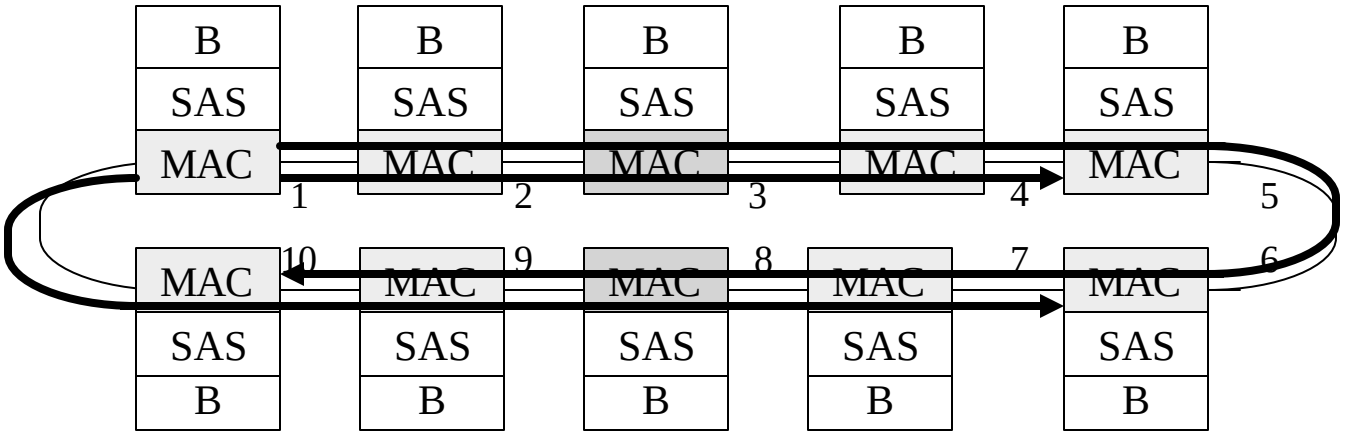

Issue with Uni-Flood

- 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 a few times until learning is complete
    - steady state has no flooding
  - Bi-flooding can be easily understood by customers are being optimal

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

.17b D1.1 Changes

- Clause 7.7.2 – change third paragraph to
  - For non SAS aware stations the determination of the cleave point is implementation specific when the ring is not currently experiencing a protection events.

.17b D1.1 Changes

- Clause 7.7.2.1 – add sub-clause
  - 7.7.2.1 Cleave point determination for SAS
  - In order to insure a interoperability of SAS implementations a station may request the use of a standardized cleave point calculation. This request is made as part of Station setting ATT message (11.4.3).
  - For rings with an odd number of stations, the cleave-point is selected such that an equal number of stations appear to the east and west of the station.
  - For rings with an even number of stations, the cleave point is selected such that for stations that are on opposing sides of the ring send traffic between each other over the same path. The station with the lower MAC address will send traffic on ringlet 0 and the opposing station will send traffic on ringlet 1.

.17b D1.1 Changes

- Clause 11.4.3 – change figure 11.18 to add bit 3 labeled cp
  - change 11.4.3.1 to 5-bit reserved field
  - change 11.4.3.2 ..4 to 11.4.3.3..5
  - add 11.4.3.2 cp: When a station running SAS receives a cp request it shall set its value of myFloodingForm to FLOOD_BIDIR and shall determine the cleave point as per 7.7.2.1

Advantages

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

Disadvantages

- Short term increase in bi-flooding which could have an performance impact for some implementations
- working group forces a particular implementation detail
  - steering is a clear example

Straw Poll

- A comment against P802.17b D1.1 has been entered requesting this feature.

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