Issues/Concerns with MMRP
802.1 AVB Conference Call
March 28, 2007

Don Pannell

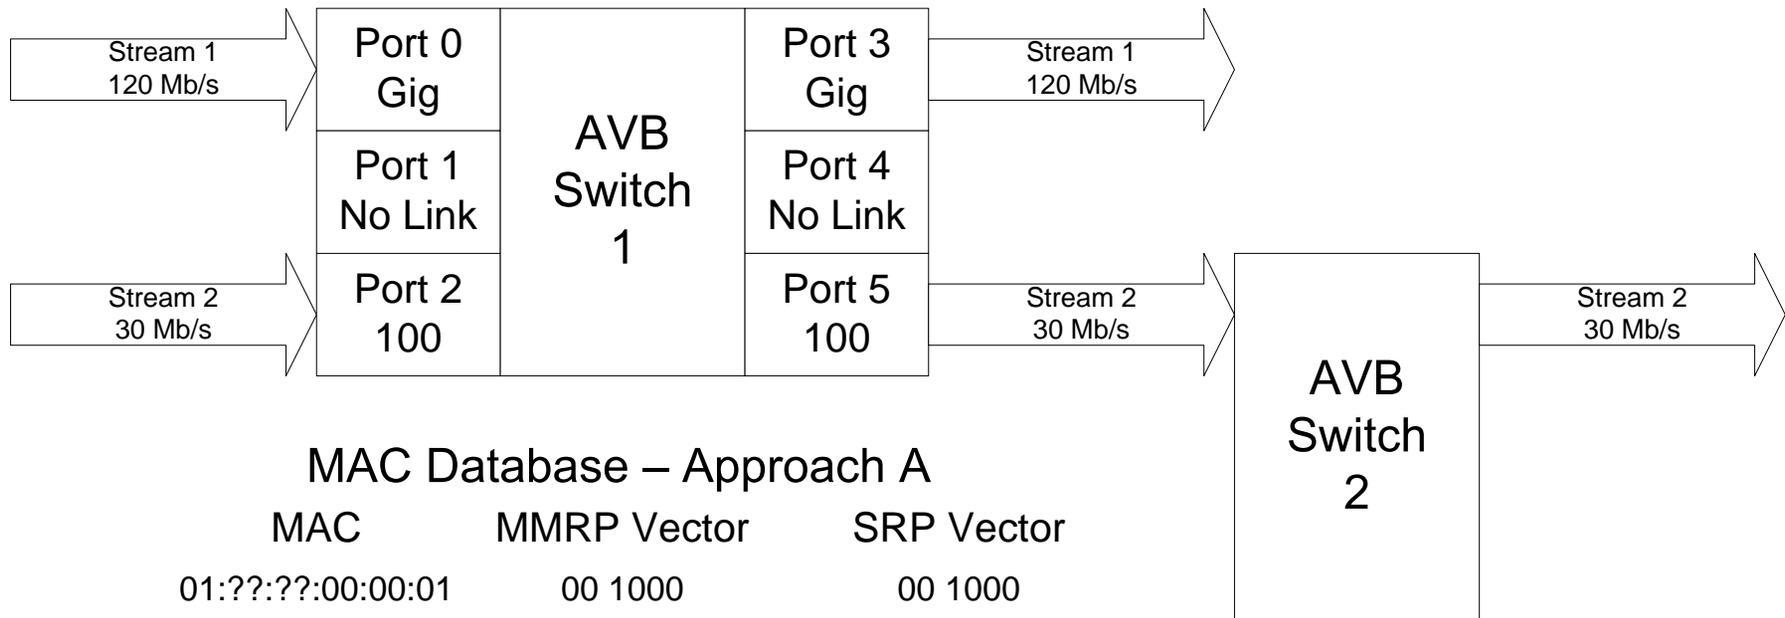Marvell

dpannell@marvell.com

Approach A

- Approach A's Assumptions (my understanding anyway):
  - MMRP is not 'touched by human hands'
  - MMRP runs 1st and then SRP gets its chance
  - MMRP maintains a per port bit (vector) for each multicast Stream address
  - MMRP will update a Stream's bit vector (the MMRP portion) in the hardware when requested to add or update a Stream's flow
  - SRP maintains a per port bit (vector) for each multicast Stream address
  - SRP will update a Stream's bit vector (the SRP portion) in the hardware when the requested Stream is 'approved'

Approach A

- Two bit vectors are needed
  - Do both need to be in hardware or can one be in software?
- Consider a 48 port switch if two bit vectors are needed in hardware
  - Any MAC Database will need 48-bits per entry just for MMRP
    - Architectures already support this to support IGMP multicast pruning
  - If two separate bit vectors are needed then an additional 48 bits per stream entry (MAC) is needed in the hardware
    - This adds too much cost when it can be easily avoided

Consider this Corner Case A

Stream 1 120 Mb/s → | Port 0 Gig | AVB Switch 1 | Port 3 Gig | → **Stream 1** 120 Mb/s

Port 1 No Link | Port 4 No Link

Stream 2 30 Mb/s → | Port 2 100 | Port 5 100 | → **Stream 2** 30 Mb/s → AVB Switch 2 → **Stream 2** 30 Mb/s

MAC Database – Approach A

MAC	MMRP Vector	SRP Vector
01:??:??:00:00:01	00 1000	00 1000
01:??:??:00:00:02	10 0000	10 0000

- Two streams 1 & 2 Setup by MMRP and Approved by SRP
- Everything is OK
- But then a device on AVB Switch 2 wants to join Stream 1

Corner Case A, part 2

Stream 1
120 Mb/s

| Port 0 Gig |
| Port 1 No Link |
| Port 2 100 |

AVB Switch 1

Stream 2
30 Mb/s

| Port 3 Gig |
| Port 4 No Link |
| Port 5 100 |

Stream 1
120 Mb/s

Stream 2
30 Mb/s

AVB Switch 2

Stream 2
30 Mb/s

Stream 1 - Request
0 Mb/s

MAC Database – Approach A

MAC	MMRP Vector	SRP Vector
01:??:??:00:00:01	10 1000	00 1000
01:??:??:00:00:02	10 0000	10 0000
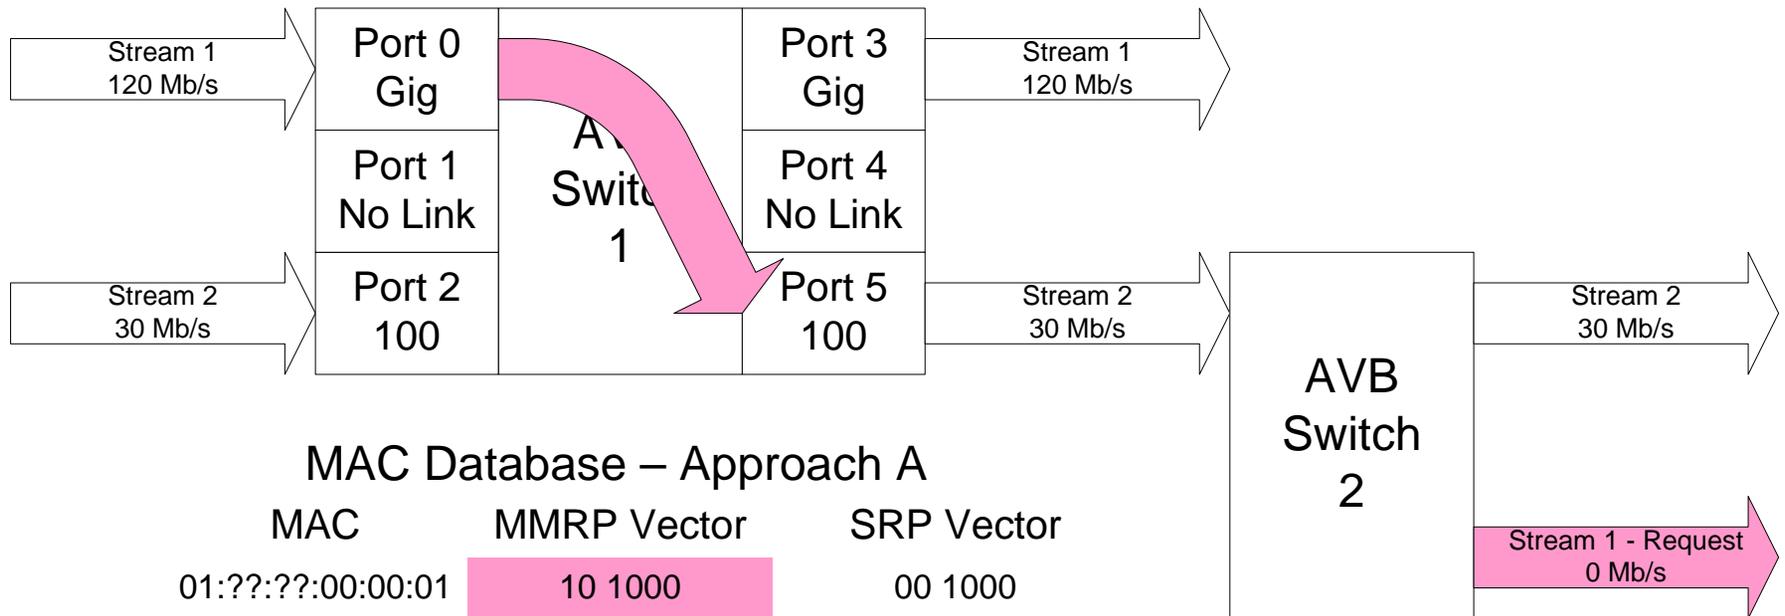
- MMRP, not knowing if SRP will 'approve' the flow, allows Stream 1 to be mapped to port 5 (in its vector).
- Then SRP 'sees' that there is not enough bandwidth so it does not update Stream 1's SRP vector.
- If the MMRP vector is the only one used then the new Join request will cause Stream 2 to have drops until SRP 'fixes' the problem

Corner Case A's Issues

- Requires **_Two_** bit vectors in hardware!
  - Remember the 48 port switch – This will cost a lot
- Requires the hardware knows which bit vector to use
  - Remember other 'legacy' ports on the switch could be using MMRP for 'standard' MMRP flows
    - MMRP flows need to look at the MMRP bit vector
    - SRP flows need to look at the SRP bit vector
  - What does the hardware look at to know which vector to use?
    - Can use PRI 4 or 5 on an AVB port – but this has nothing to do with the address database
    - It is better to use the DA as its part of the address database

Approach B

- Approach B's Assumptions:
  - MMRP is only slightly 'touched by human hands'
    - MMRP does not update the hardware Address Database bit vector for a MAC if the Stream is an SRP stream. That is ONLY change!
    - MMRP can know if the Stream is an SRP stream by a range of MAC addresses
      - Something that is good for MMRP anyway!
  - MMRP runs 1st and then SRP gets its chance
  - SRP will update a Stream's bit vector in the hardware Address Database when the requested Stream is 'approved'

- Only *One* hardware bit vector is needed
  - Architectures already support this to support IGMP multicast pruning
  - Saves the need and cost of an additional bit per port per address database entry!