Two & Three Step MSRP

Stream Registration overhead

Craig Gunther (cgunther@harman.com)

23 April 2008

Successful Stream Join

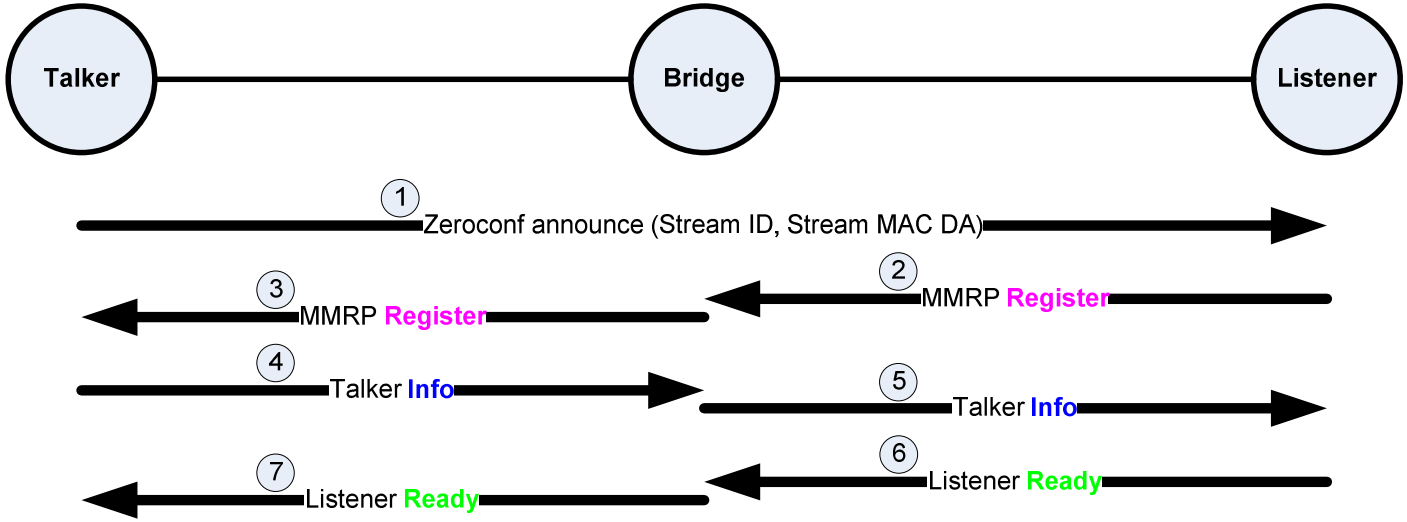

- Talker advertises stream via higher layer protocol (e.g. Zeroconf)
- Listener issues MMRP **Register**
- Talker responds with MSRP Talker **Info**
- Listener requests Stream with MSRP Listener **Ready**
- After receipt of the **Ready** the Talker can begin transmitting the audio/video stream at any time

2-step MSRP Approach (still supported)

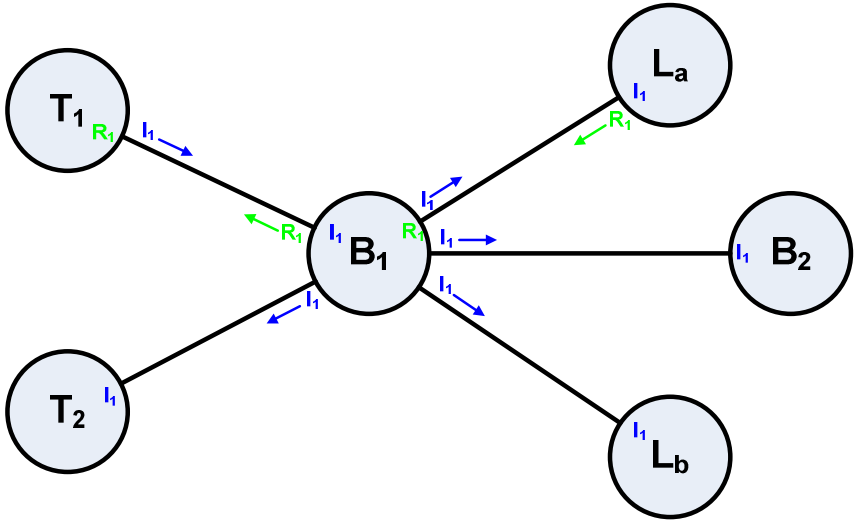

The 2-step approach causes MSRP Talker **Info**s to propagate throughout the network

- Talker **T_1** advertises the Stream **Info** (**I_1**) towards Bridge **B_1**
- Bridge **B_1** forwards the Talker **Info** out <u>all</u> ports since the Stream MAC DA is not MMRP **Register**ed on any ports
- The Talker Stream **Info** declaration (64+ bytes) is registered on all devices attached to Bridge **B_1**
- Bridge **B_2** will also forward the Talker **Info** throughout the entire network
- Note that the **Ready** from Listener **L_a** only goes to Talker **T_1**

3-step MSRP Approach (preferred)

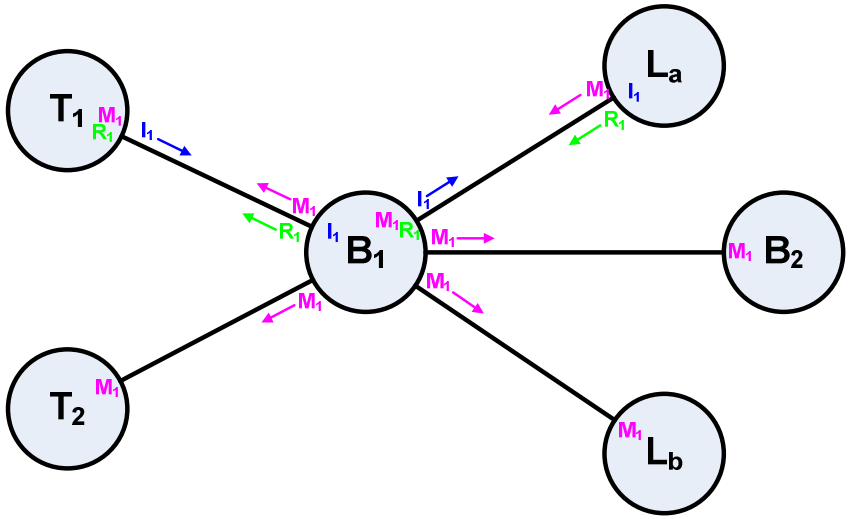

The 3-step approach greatly reduces the spread (and overhead) of MSRP Talker **Info**s

- Listener **L_a** sends an MMRP **Register** (6 bytes) which gets propagated on all ports
- Talker **T_1** responds with the Stream **Info** (**I_1**) towards Bridge **B_1**
- Bridge **B_1** recognizes the Streams MAC DA and only forwards the **Info** towards **L_a**
- Talker **Info** declarations only register along the Streams path from Talker to Listener
- Again, the **Ready** from Listener **L_a** only goes to Talker **T_1**

MSRP and Legacy Listeners

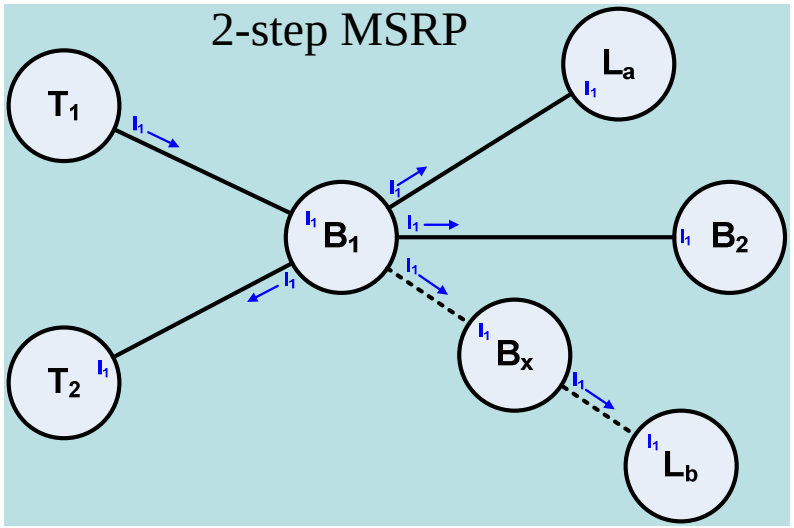

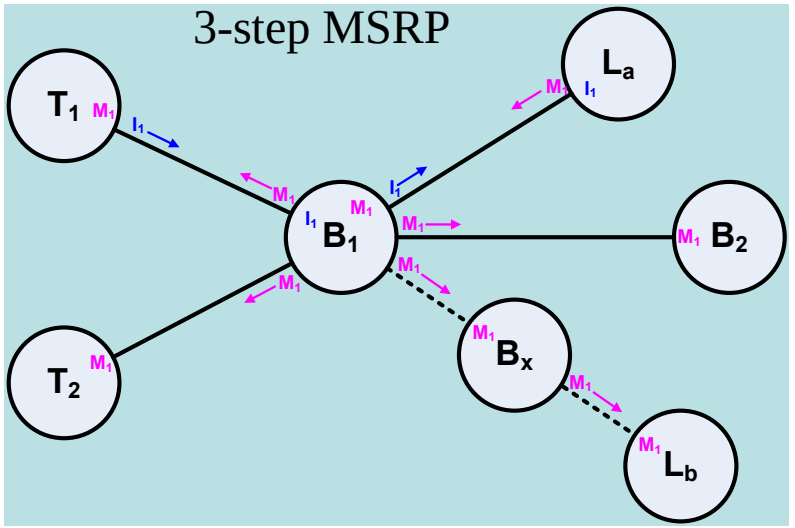

- Assume Listener **L_b** is a legacy Listener that does not support MMRP **Register** declarations
- 2-step approach shows that Listener **L_b** and all Bridges (**B_x**) along the path to Listener **L_b** will learn about Talker **T_1**'s Stream
- 3-step approach excludes Listener **L_b** and all Bridges (**B_x**) along the path to Listener **L_b** from learning about Talker **T_1**'s Stream

Legacy Listeners and 3-step MSRP Solutions

- Two possible solutions:
  - A clever Bridge manufacturer may add a switch that disables the 3-step related pruning of Talker **Info**s
    - Defeats all benefits of 3-step pruning
    - Don't forget costs of unnecessary Talker **Info** propagation
  - Bridges could also provide an MMRP proxy capability for Legacy Listeners
    - 3-step pruning benefit is still intact
    - If Bridge is doing MMRP proxy (e.g. on behalf of RSVP) then it can assume it must do MSRP proxy as well.
- These solutions are out of scope but may be discussed in an informative annex

Automatic Priority Downgrades

- Bridges will not automatically downgrade a Stream's priority to Best Effort (within an AVB cloud) when there is insufficient bandwidth left on an associated outbound port
  - Since we are using MMRP in the 3-step approach this "priority downgrade" would cause all links downstream to be sent the Best Effort Stream whether they wanted it or not
  - If a station desires a Best Effort Stream it must ask for it explicitly by a means other than MSRP (could be as easy as using MMRP by itself)
  - Defining Stream behavior as it leaves the "cloud" is out of scope but may be discussed in an informative annex

MSRP Attribute Size Considerations

- MRP LeaveAllTimer causes a Declaration/Registration refresh every 10-15 seconds (802.1ak Table 10-7). There are roughly 64+ bytes per **Info**, 24+ bytes per **Ready**.
  - We have good reason to be prudent about the amount of information we pass around in these declarations
    - A single ride in DisneyWorld contains 450+ audio channels (http://livedesignonline.com/mag/show_business_blast_off_epcot)
    - Large mixing consoles control 1700 channels with multiple consoles per facility

Thanks