



MRP Timers
Maximum attribute registrations

Craig Gunther (cgunther@harman.com)

03 February 2010

Problem statement

Obviously there are limits to the number of MSRP attributes a bridge or end-station can support. This is based on criteria you would expect: memory, CPU performance, bandwidth between CPU and switch fabric, etc.

There is also a worst-case limit based on the MRP timers and the LeaveAllEvent mechanism. Worst-case is defined as **non-contiguous attributes** that do not take advantage of the FirstValue + NumberOfValues condensed packing (i.e. one attribute per Vector Attribute).

The purpose of this presentation is to explore the LeaveAllEvent limit.

Concepts to understand

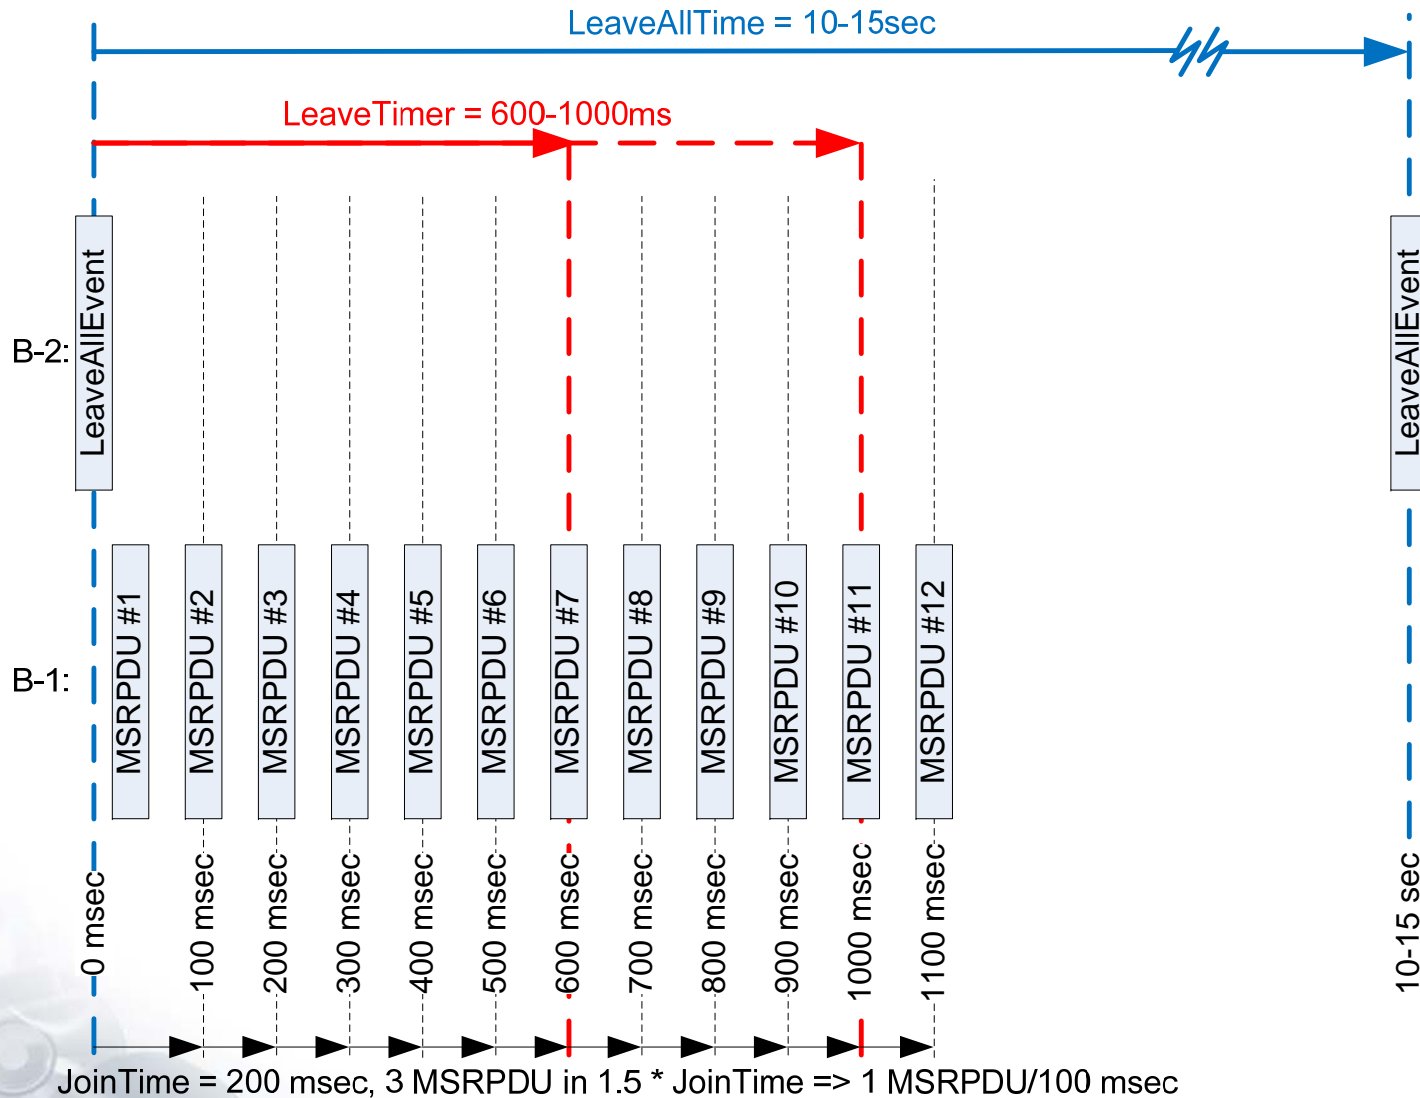
- Table 10-7 – MRP timer parameter values:
 - JoinTime = 20 centiseconds (200ms)
 - LeaveTime = 60-100 centiseconds (600-1000msec)
 - LeaveAllTime = 1000 centiseconds (10 seconds)
- Clause 10.7.11 Timer values: If operPointToPointMAC (6.4.3) is TRUE, a request for a transmit opportunity should result in such an opportunity as soon as is practicable, given other system constraints, and shall occur within the value specified for JoinTime (200 msec in Table 10-7) subject to **not more than three such transmission opportunities occurring in any period of 1.5*JoinTime.**
- Clause 10.8.2.4 Encoding of LeaveAllEvent: The LeaveAllEvent is interpreted on receipt of a MAD Leave All event to be **applied to the state machines for all Attributes of the type defined by the AttributeType field.**
- Table 35-1 – MSRP AttributeType Values: **Talker Advertise, Talker Failed, and Listener.**
- Clause 35.2.3.1: SRP uses MMRP for Talker Pruning.

A worst-case example

The following slide describes the MRP communications between two bridges and the relationship of the various MRP timers from Table 10-7. Bridge B-1 has declared several attributes to Bridge B-2. Bridge B-2 is issuing a LeaveAllEvent to Bridge B-1 and Bridge B-1 is re-declaring its attributes.

Any attributes not re-declared before the LeaveTime (600-1000msec) expiration will be transition to the LV state by Bridge B-2, which effectively removes those attributes from the bridge. MSRPDU #12 is such a failure. MSRPDU #11 is most likely going to be a failure as well.

MRP LeaveAllEvent and associated timers



How many attributes in 10 MSRPDU's?

- We should be able to reliably receive 10 MSRPDU's in 1 second (LeaveTimer)
- How many SRP attributes in 10 MSRPDU's:
 - 40 Talker Failed (1507 octets)/MSRPDU = 400 Talker Failed attrs
 - 53 Talker Advertise (1511 octets)/MSRPDU = 530 Talker Advertise attrs
 - 135 Listener Ready (1512 octets)/MSRPDU = 1350 Listener attrs
 - 165 Talker Pruning (1512 octets)/MSRPDU = 1650 Talker Pruning attrs

How to cause inconsistent behavior

- This scenario will cause multiple streams to be dropped when we only want to drop one:
 - Advertise 529 separate Talker Advertise attributes
 - Advertise the 530th Talker Advertise attribute which is really made up of 11 contiguous streams.
 - Tear-down the 2nd stream of the 11 contiguous streams, which now splits the 11 contiguous into a single stream plus 9 contiguous. This results in 531 Talker Advertise declarations and something must be dropped.
- Can we derive a formula so bridge can report an out of resources failure when the limit is hit?

How to get more attributes

- Utilize contiguous StreamID + StreamDA which removes the worst-case scenario.
- Don't combine LeaveAllEvent for TalkerAdvertise + TalkerFailed + Listener in a single packet. Send TalkerAdvertise LeaveAllEvent, then 5 seconds later send TalkerFailed LeaveAllEvent, then 5 seconds later send Listener LeaveAllEvent. That way the multiple attributes do not compete for PDU space in the same 10 frames.
- **Modifications to 802.1ak (distasteful!)**
 - Shorten JoinTimer (200msec) or modify the "1.5 * JoinTime" limit when responding to the LeaveAllEvent.
 - Lengthen LeaveTime (600-1000msec). **We would also have to address the problem of the LeaveTime bumping into the optional Periodic Timer** (re-declare your attributes every second).
 - Require two LeaveTime expirations which would allow us to re-declare half our attributes one time and half the next, effectively doubling the number of attributes.
 - Add some type of shortcut that allows for less octets to re-declare vs more octets for initial declaration.

Reducing other MSRP CPU overhead

- Do not use optional Periodic Timer for MSRP. 10x to 15x decrease in number of frames exchanged.
- Use 15 second LeaveAllEvent timer.
- Caching egress MSRP packets for less CPU overhead during declaration processing if MAD (MSRP Attribute Database) does not change.
- Caching ingress MSRP packets for quick binary comparison of previous registration vs current registration. This is more difficult since you may have received up to 10 packets to register all 530 Talker Advertise attributes.
- Use some type of 8-bit sequence number in the header (802.1ak change) that the transmitter increments if the packet contents changes. Listener could just keep a list of valid sequence numbers and know they don't need to reprocess identical packets. Must maintain some internal "sequence number" to "attribute" mapping.



BECKER



harman/kardon



lexicon




Harman International

Premium Solutions for Audio and Infotainment