Harman International

MSRP Protocol Version 0 to Version 1 Interoperability

Craig Gunther (craig.gunther@harman.com)

November 2014, San Antonio

Interoperability Goals

When new features in MSRP Protocol Version 1 are published they must be architected such that there is a painless incremental upgrade path for existing Version 0 installations. Here are the goals:

1. Version 1 end-stations shall interoperate with Version 0 end-stations with quality of service as good or better than Version 0 provides
2. An intermix of Version 0 and Version 1 Bridges shall be supported while providing a quality of service that is as good or better than a pure Version 0 network
3. Version 0 end-stations shall work equally well with Version 0 and Version 1 Bridges
4. A stream running entirely though Version 1 devices shall support the extended capabilities described in MSRP Protocol Version 1*

* Note: Goal 4 is obvious and not discussed further

A New Version 1 Attribute: ListenerLatency

This presentation introduces a new MSRP AttributeType (Table 35-1, included below) that can be used instead of Version 0's Listener, we'll call it: ListenerLatency. This new attribute type will be used to implement the Listener specified AcceptableLatency introduced in the presentation: cc-cgunther-acceptable-latency-0314-v01.pdf.

Why is this needed? In SRP-2010 (MSRPDU Protocol Version 0) a stream will be stopped if a network reconfiguration occurs even if the resultant latency change is acceptable to the Listener device. Allowing Listeners to specify the maximum acceptable latency will solve this problem and provide the "network" will some guidelines when rerouting existing streams.

AttributeType	Value	Supported Protocol Version
Talker Advertise	1	All
Talker Failed	2	All
Listener	3	All
Domain	4	All
ListenerLatency	**5**	**Version 1 and beyond**

Meeting our Interoperability Goals (LISTENERLATENCY example)

Scenario 1: Upgrade Bridge to MSRP Version 1

- [10.8.3.5a] V1 Bridge will interpret Talker's V0 messages according to SRP-2010
- [10.8.3.5b] V0 Talker will interpret Bridge's V1 messages according to SRP-2010

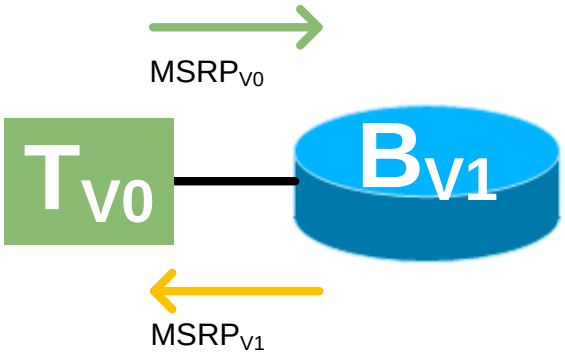

  - [10.8.3.5c1] New AttributeTypes (such as LISTENERLATENCY) are ignored, continuing with the next Message*
  - [10.8.3.5c2] New AttributeEvents (not New, JoinIn, In, JoinMt, Mt or Lv) are ignored, continuing with the next VectorAttribute

- Summary: V0 Talkers and V0 Listeners can communicate through a V1 Bridged network

*Note: Each Message contains only one AttributeType

Scenario 2: Upgrading Listener to Version 1

- Talker sends a TALKER ADVERTISE (*type*=1) in a Version 0 MSRPDU
- [10.8.3.5a] Listener receives the TALKER ADVERTISE and processes it according to SRP-2010
- [10.8.3.5c1] Listener only transmits Version 1 MSRPDU, however if Listener sends a LISTENERLATENCY READY (*type*=5) the Talker will discard it and the Listener would have no indication that this happened
- [10.8.3.5b] If Listener sends a LISTENER READY (*type*=3) the Talker will process it according to SRP-2010 and start streaming

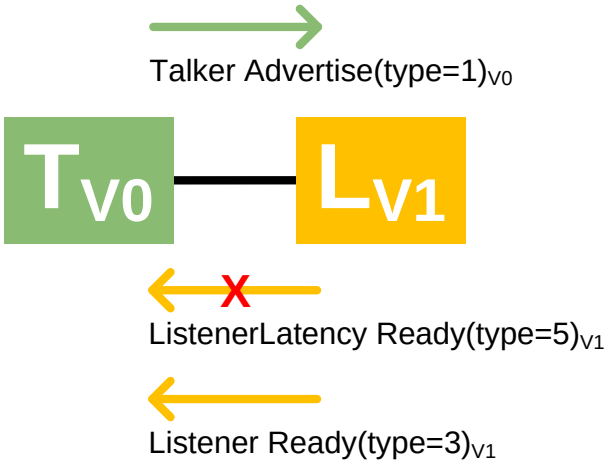

QUESTION: How can a device know whether to send LISTENER (*type*=3) or the new LISTENERLATENCY (*type*=5)? It can look at the protocol version in the peer's MSRP DOMAIN message and know if the peer will understand a *type* greater than 4.

Scenario 3: Upgrading Bridge and Listener to Version 1

- Talker sends a TALKER ADVERTISE (*type*=1) in a Version 0 MSRPDU
- [10.8.3.5a] Bridge receives the TALKER ADVERTISE, then propagates it in a Version 1 MSRPDU

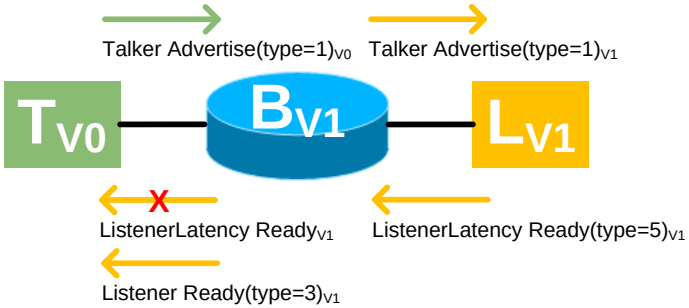

- Listener sends a LISTENERLATENCY READY (*type*=5) in a Version 1 MSRPDU
- Bridge receives it and propagates it in a Version 1 MSRPDU:
  - If Bridge propagates the LISTENERLATENCY READY (*type*=5) the Talker will discard it and the Listener would have no indication why this happened
  - If the Bridge converts the LISTENERLATENCY READY to a LISTENER READY (*type*=3) the Talker can process it according to SRP-2010 and start streaming

Scenario 4: Mixture of Version 0 and Version 1 Bridges

- Similar to Scenario 3: Bridge B_{v1} receives the TALKER ADVERTISE in a Version 0 MSRPDU from Bridge B_{v0} and propagates it in a Version 1 MSRPDU

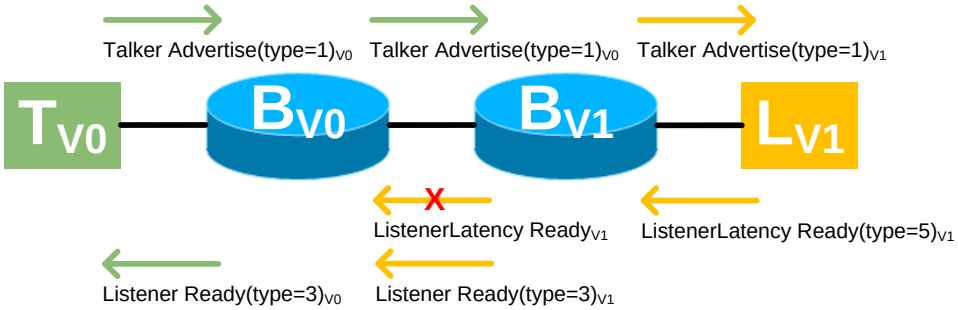

- Listener L_{v1} receives the TALKER ADVERTISE with its ACCUMULATEDLATENCY and responds with a LISTENERLATENCY READY with ACCEPTABLELATENCY set greater than or equal to ACCUMULATEDLATENCY
- Bridge B_{v1} receives the Version 1 MSRPDU LISTENERLATENCY READY from the Listener and converts it to a Version 1 MSRPDU LISTENER READY, which is then propagated by Bridge B_{v0} as a Version 0 MSRPDU LISTENER READY
  - Bridge B_{v1} will change TALKER ADVERTISE to TALKER FAILED if ACCUMULATEDLATENCY ever exceeds ACCEPTABLELATENCY
  - Bridge B_{v0} will change TALKER ADVERTISE to TALKER FAILED if ACCUMULATEDLATENCY changed (standard MSRP Version 0 behavior)

Proposal for meeting Goal 1 and 2

Scenarios 2, 3 and 4 demonstrate that a Version 1 device can receive a stream from a Version 0 peer device as long as the Version 1 device converts the ListenerLatency attribute to a Listener attribute. This applies to end-stations as well as bridges along the path.

The associated text for the amendment could say something like this:

> *"A ListenerLatency attribute shall be propagated as a Listener attribute if the peer device implements MSRP Protocol Version 0 (as determined by the Protocol Version used in the peer's Domain message)."*

The end result is that the path from the Version 0 device all the way back to the Talker will behave like an SRP-2010 network, which means the stream will be stopped if the AccumulatedLatency changes. However, the path from that device to the Listener will support the AcceptableLatency originally requested.

Therefore, the quality of service is as good or better than Version 0 provides.

Goal 1 and 2

Interoperability Goals Are Met!
Other Rules For LISTENERLATENCY **Propagation.**

Propagating Listener Attributes

- It is unnecessary to forward a LISTENER READY (LR) and a LISTENERLATENCY READY (LLR) on the same egress port
  - Convert the LISTENER READY to a LISTENERLATENCY READY with ACCEPTABLELATENCY set equal to ACCUMULATEDLATENCY

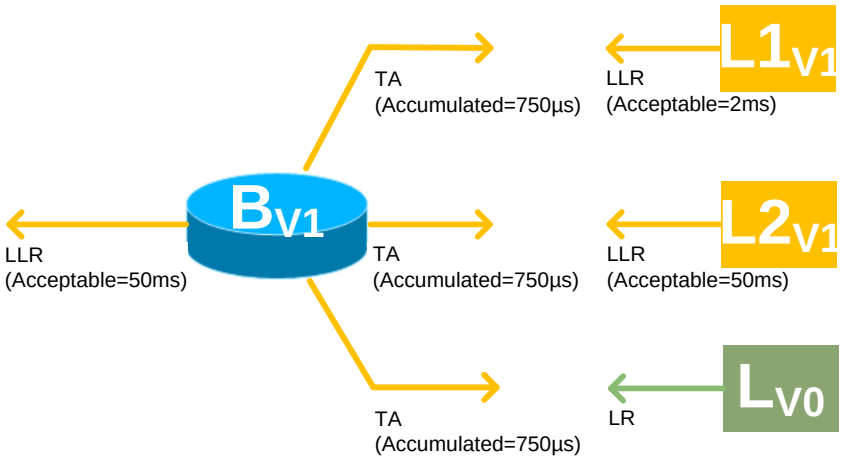

- Propagated LISTENERLATENCY READY shall have the ACCEPTABLELATENCY set to the largest ACCEPTABLELATENCY of all received values from all other ports
  - This assures the Bridge's streaming ingress peer will support the most flexible Listener requirement
  - Each of the Bridge's egress stream ports can fail the associated TALKER ADVERTISE if the individual ACCEPTABLELATENCY requirements cannot be met

Other Thoughts…

- Entire stream path from farthest away Version 0 Bridge back to the Talker will not be able to take advantage of Version 1 AcceptableLatency capabilities.
- Version 1 MSRPDU ListenerLatency attributes will likely not pack as well as Listener attributes.
- If we introduce new FailureCodes (Table 35-6) for ListenerLatency then we will not be able to combine Listener and ListenerLatency failure responses as shown in slide 10
- Is there any interest in introducing TLVs into MSRP Version 1 to make future version enhancements easier to implement? If so, we should do it now.
  - If we had done this in Version 0 we would not need to do the ListenerLatency to Listener conversion presented here!


Harman International

Premium Solutions for Audio and Infotainment