



SRP AND RESERVATION TYPES

CRAIG GUNTHER

July 2016 IEEE 802 Plenary, San Diego, USA



This presentation discusses three types of reservations:

1. **DYNAMIC RESERVATIONS:** as defined by 802.1Q-2014, Clause 35
2. **PERSISTENT RESERVATIONS:** are just the SR part of SRP (i.e. no protocol is involved)
3. **PRELOADED RESERVATIONS:** can be stored/saved and immediately recreated upon power-up; these reservations can also be torn down

SRP (Stream Reservation Protocol) uses MSRP to create reservations. SRP will also be able to be configured via the UNI defined in P802.1Qcc-D1.0 Clause 99.

PRO-AUDIO (1ST GENERATION SRP)

Pro-Audio systems implement DYNAMIC RESERVATIONS today.

These are the reservations created by the SRP protocol defined in 802.1Q-2014, clause 35.

Reservations are created when an MSRP Talker Advertise and an MSRP Listener Ready are combined in a bridge.

MSRP attributes must be defended when the LeaveAll timer expires or the attribute(s) and the associated reservation(s) will be removed. The LeaveAll timer has a direct impact on the number of streams a network can support.

AUTOMOTIVE (1ST GENERATION)

Automotive systems implement PERSISTENT RESERVATIONS today. These reservations are not currently defined in an approved 802.1 standard.

Bridges and end stations load reservation¹ information from flash upon power-up and configure egress ports:

- MAC DA forwarding
- operIdleSlope
- VLAN membership

PERSISTENT RESERVATIONS could be configured via Management (Qcc); some type of central controller could push these reservations into the network without MSRP.

These reservations are persistent, they do not need to be “defended” (think LeaveAll timer); they are never torn down by MSRP.

Each bridge along the path from Talker to Listener must be similarly configured.

1. This information may be stored in a summarized format (e.g. a simple list of MAC DAs, total operIdleSlope of all streams in a class, and a list of VLANs)

AUTOMOTIVE (2ND GENERATION)

Automotive systems looking to add PRELOADED RESERVATIONS in the future.

Bridges and end stations load reservation information from flash upon power-up and configure egress ports (similar to PERSISTENT RESERVATIONS):

- MAC DA forwarding
- operIdleSlope
- VLAN membership

This allows bandwidth to be reallocated to mitigate the limitations of 100 Mbps (802.3bw) links.

PRELOADED RESERVATIONS behave similar to PERSISTENT RESERVATIONS until SRP is enabled; at that point they behave like DYNAMIC RESERVATIONS

- MSRP attributes are loaded from flash to recover the individual TSpec information; more detailed than summarized information used by PERSISTENT RESERVATIONS

When SRP is enabled:

- LeaveAll timer begins operation and PRELOADED RESERVATIONS must be defended
- MSRP Listener Ready and MSRP Talker Advertise attributes can be withdrawn to tear down a PRELOADED RESERVATION and free up associated bandwidth for another stream
- Other reservations can be added and withdrawn as desired

WHY THIS WON'T WORK TODAY (OR THE NEAR TERM)

100 Mbps: Because of work in various organizations and alliances the Automotive industry has embraced IEEE Std. 802.3bw-2015 (100BASE-T1). But 100 Mbps does not provide sufficient bandwidth for the capabilities that OEMs are desiring.

1000 Mbps: The capabilities of 802.3bp (1000BASE-T1) are very interesting and provide 10x the capability of the current Ethernet physical layer found in Automotive. However, that is still not enough bandwidth when you consider the possible new technologies that will use video feeds from HD cameras around the car. If that data is compressed the DSPs that do the signal processing will have to spend computation resources and time uncompressing the video stream as well as the possibility of lost resolution due to compression.

HDBase T (a.k.a IEEE 1911): A 10 Gbps CAT 5e or above is being considered as a solution, although the 4-pair cabling is a step backwards. This allows transmission of uncompressed 4K UHD @ 30Hz video which will allow the type of image processing desired by OEMs for the near future. It is my opinion that a 2-wire, short-haul solution would likely be requested from the 802.3 working group and if accepted that process could take **4 years** from Study Group to Standard publication (based on 802.3bp timeline of March 2012 to May 2016), **plus another 3 to 6 years** before it makes it into vehicles.

PRO-AUDIO (2ND GENERATION SRP)

Pro-Audio systems would also benefit from support for PRELOADED RESERVATIONS.

These reservations would be configured the same way as the previous slides have defined for Automotive reservations, but for a different set of use cases.

In the simple case of a concert venue, a garage band performance, or a video wall it is desirable to restore the A/V system as soon as possible after a power failure; without having to wait for the system controller (e.g. I722.1) to communicate with all the end stations and instruct them to reestablish the streams. Just allow the performance/presentation to resume from where it was interrupted.

A video surveillance/security system would have the same type of requirements. In the event of a power failure or a partial network reset it would be important to immediately begin streaming the same video that was previously configured while the rest of the network continues to boot. Once the network is up and running SRP could begin administering the reservations as usual.

THE IDEAL SOLUTION

The preferred solution would be to support all three types of reservations (DYNAMIC, PERSISTENT and PRELOADED).

I would suggest that a combination of the UNI defined in Qcc plus some simple updates to SRP (via MRP) could provide all the desired capabilities.

The remainder of this presentation does not address the Qcc UNI PERSISTENT RESERVATIONS or SRP's DYNAMIC RESERVATIONS. The focus of the following 3 slides is on the complete set of MRP updates required to support PRELOADED RESERVATIONS within SRP.

PROPOSED CHANGES FOR MRP LEAVEALL TIMER

MRP does not prevent the LeaveAllTime from being set to zero, but what does that mean? I suggest the logical choice is to turn off the LeaveAllTimer. This needs to be clarified.

The LeaveAll state machine must continue to operate so Spanning Tree changes can be processed by MRP. Spanning Tree changes generate a leavealltimer! event (see 10.7.5.2).

Make the following modifications (see related 12.9.1.2.1):

10.7.4.3 leavealltimer

Change the text as shown:

The Leave All Period Timer, leavealltimer, controls the frequency with which the LeaveAll state machine generates LeaveAll PDUs. The timer is required on a per-Port, per-MRP Participant basis. If LeaveAllTime is zero the Leave All Period Timer is not started, otherwise the Leave All Period Timer is set to a random value, T, in the range $\text{LeaveAllTime} < T < 1.5 * \text{LeaveAllTime}$ when it is started. LeaveAllTime is defined in Table 10-7.

10.7.6.10 Start leavealltimer

Change the text as shown:

When LeaveAllTime is non-zero causes leavealltimer to be started, in accordance with the definition of the timer in 10.7.4.3.

10.7.9 LeaveAll state machine

Insert the following text after the first paragraph:

If the LeaveAllTime is changed from a value of zero to a non-zero value a Begin! event shall be signaled to the LeaveAll state machine to cause it to reinitialize.

ADDITIONAL MRP LEAVEALL TIMER CHANGES

More MRP modifications:

10.7.11 Timer values

*Rename Table 10-7 as shown:**

Table 10-7-MRP timer parameter default values

10.6 Protocol operation

Change NOTE-4 as shown:

NOTE 4—In the face of changing inputs, arbitrary loss, delay, reordering, and unsignaled interruption in participation, no protocol will meet its objectives. What can be said is that the objectives will be met after a known period of operation within specified limits, and what failures are most likely otherwise. MRP favors ensuring that registrations are made, at the expense of tolerating prolonged registrations. Thus, the LeaveAll mechanism, when it has an effect, most often implements “garbage collection.” At the same time it will also ensure that failed registrations are (re-)established.

These objectives cannot be ensured if the LeaveAllTime is set to zero; therefore it is recommended to set the LeaveAllTime to a non-zero value as soon as is practicable.

*Note that the change to the description of Table 10-7 is requested so that users of AVB systems will understand that SRP as defined in 802.1Q-2014 supports modifications to the MRP timers, which allows an increased number of reservations to be supported on existing AVB networks.

PROPOSED CHANGES FOR RESERVATION CONFIGURATION VIA MANAGEMENT

MSRP has several managed objects defined in Clause 12.22 Stream Reservation Protocol (SRP) entities. In particular *Table 12-16—SRP Stream Table row elements* defines the managed objects that correspond to the Talker Advertise, and *Table 12-17—SRP Reservations Table row elements* corresponds to the reservations built when a Talker Advertise is paired with a Listener Ready. In order to allow configuration of reservations through the management interface several of these objects need to be writable.

12.22.4 SRP Stream Table

Change all rows in Table 12-16—SRP Stream Table row elements to RW.

12.22.5 SRP Reservations Table

Change the following rows in Table 12-17—SRP Reservation Table row elements to RW: Direction, Declaration Type

HOW DO WE USE IT?

The previous three slides defined the entire set of (7) changes to SRP required to support PRELOADED RESERVATIONS. When those changes are standardized how would they be used?

A system designer would:

- Set the LeaveAllTime to zero during initial boot,
- If required, configure the MRP timers defined in Table 10-7 with appropriate values to support a large number of streams,
- Preload the SRP Reservations from FLASH memory via the changes to clauses 12.22.4 and 12.22.5. Similarly preload the MVRP and MMRP entries via the appropriate management variables,
- Upon command of the System Controller (or any other technique) configure all devices with the appropriate LeaveAllTime. There is no concern for “race conditions” related to which devices get the LeaveAllTime set first since all devices will respond appropriately to received LeaveAll packets even if their own LeaveAllTime is set to zero.
- Take snapshots of current stream configuration as required to prepare for next system boot/power-fail recovery.

Questions?

Thanks!