AVnu Alliance™ Best Practices

Stream Reservation Protocol

Revision 1.0

2014 November 3

Levi Pearson

AVnu Alliance

1. Introduction and Scope

In 2007, the IEEE 802.1 “Audio/Video Bridging Task Group” started developing a series of specifications targeted to optimize time-synchronized low-latency media streaming services through IEEE 802 networks. One of these specifications, now a part of 802.1Q-2011, describes a protocol known as the Stream Reservation Protocol (SRP) that allows network endpoints to reserve bandwidth across a compliant network.

This document provides a high-level overview of the SRP protocol, focusing primarily on the application interfaces that it provides to end points and the way that the protocol propagates information about streams throughout the network. Detailed descriptions of the low-level operational details are mostly ignored, except in cases where the details place restrictions on how the higher-level interfaces can be used.

The descriptions in this document are meant as a guide to implementers of end point software and should not be taken as a normative specification. Other methods of interpreting the standards may also be valid, and readers of this document are encouraged to refer to 802.1Q-2011 for the precise specification details.

1.1. AVnu’s Relationship to IEEE AVB/TSN

AVnu is an industry alliance to foster, support and develop an ecosystem of inter-operable AVB/TSN devices. AVnu provides interoperability guidelines and compliance testing for device manufacturers, as well as technical guidance such as this document for implementing AVB/TSN systems. For more information, visit our website (www.AVnu.org).

Contents

2. Stream Reservation Protocol

__Note__: The following section is included to explain the system operation of one of the core protocols used with AVB - the Stream Reservation Protocol (SRP). Some material is included where necessary to describe typical, expected network bridge behavior. None of this material authoritatively defines behavior of an AVB bridge.

Stream Reservation Protocol (SRP) is an amendment to the IEEE 802.1Q standard (standardized separately as 802.1Qat) to provide end-to-end management of resource reservations for data streams requiring guaranteed Quality of Service (QoS) in Local Area Networks (LANs). The protocol allows stream endpoints to register their willingness to "Talk" or "Listen" to specific streams, and it propagates that information through the network. Network bridges between the endpoints maintain bandwidth reservation records when a Talker and one or more Listeners register their intentions for the same stream over a network path with sufficient bandwidth and other resources.

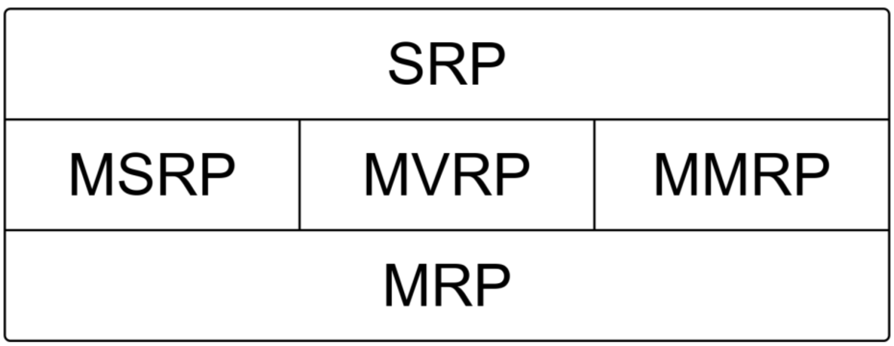

Figure 1 - SRP Layers

SRP has been implemented as a new protocol layer (see Figure 1) on top of an existing network management protocol called the Multiple Registration Protocol (MRP), which is also an amendment to the 802.1Q standard (separately standardized as 802.1Qak). MRP is a generic protocol that allows its participants to register attributes that may propagate in an application-specific way throughout a LAN. It is also the protocol underlying the Multiple VLAN Registration Protocol (MVRP) and the Multiple MAC Registration Protocol (MMRP), which both existed previously and may be used with SRP. The SRP standard provides a new MRP application called Multiple Stream Registration Protocol (MSRP) to manage attributes relating to stream bandwidth reservation. MSRP, MVRP, and MMRP provide all the network signaling for the SRP protocol.

As SRP operates in the network bridges of an AVB network, it maintains a set of tables in each bridge that record the information relevant to its operation. These include some general bridge-wide data, some per-port data, some per-stream data, and some per-reservation data. These will be discussed in more detail later, but their purpose is to keep track of which ports are participating in the SRP protocol; the communication latency between peers; the bandwidth requirements, rank, and priority of potential streams; and the status, description, and error reporting information for all registered reservation attempts. This information is used by the bridge to implement Forwarding and Queuing for Time Sensitive Streams (FQTSS, separately standardized as 802.1Qav) to provide guaranteed QoS for streams over reserved paths through the LAN.

SRP also provides a means for Talkers and Listeners to discover one another and to discover the VLAN and priority levels in use within an SRP Domain. This allows endpoints to activate their streaming data protocols only when the necessary network resources are guaranteed to be reserved and to configure those protocols for correct network operation.

SRP is not a request/response protocol in the typical sense. Instead, participants in the protocol propagate the intention of Talkers to transmit streams out unidirectionally through the network. The intention of a Listener to receive a stream is unidirectionally propagated back along the same network path by which the Talker's intention reached it. The bandwidth reservations are established as protocol participants on that path register the intention of both Talker and Listener for the same stream and verify that resources for the stream are available.

Once the intention of a Listener to receive a Talker's stream reaches that Talker, the resources necessary to ensure the desired QoS parameters for the stream have been reserved along the entire path through the LAN from Talker to Listener, and the Talker may begin transmitting its stream.

Either party may later remove the reservation by removing its intention to participate in the stream, which causes the change in intention to propagate along the same path. When there are no longer matching Talker and Listener intentions registered in a participant for a stream, the reserved resources for that stream are freed by that participant.

2.1. MRP

Conceptually, the MRP protocol operates between a set of participants. A participant is unique per MRP application (e.g. MSRP, MVRP, or MMRP) and per port. This means that a network endpoint's MSRP application is a participant and a network bridge has one MSRP application participant per port, one MVRP application participant per port, etc.

2.1.1. Declarations, Registrations, and Attributes

The fundamental operation of a participant in an MRP application is to make or withdraw a **declaration** of an **attribute**. Making or withdrawing a declaration results in the addition or removal of an attribute **registration** in some other participants of that application. In other words, changing the declaration is the action of a participant, and a change of registration is the reaction in its link peer participant to that change of declaration.

In general, the state of a declaration or registration of an attribute is persistent, meaning that it does not change unless the declaring participant drops from the LAN or specifically changes the declaration.

Note: Changes to an attribute, such as changes to a stream reservation, require withdrawing the existing attribute, waiting two (2) LeaveAll timer periods, then declaring the attribute again. It is recommended that implementations handle this in the SRP protocol layer or below so that applications using SRP cannot accidentally violate this constraint.

2.1.2. Attribute Change Reactions and Propagation

This declaration-to-registration reaction always occurs in a single direction from the declaring participant to its network peer that is accessible via its port. If its peer is part of a network bridge with multiple participants in the application, the registration change that occurs may also **propagate** to the other participants in the bridge. This propagation is an application-defined process that transforms a change in attribute registration on one participant in the bridge to a related attribute declaration change in some other participants.

Propagation rules are specific to each attribute defined by an application. They are defined by a component of the application called the MRP Application Propagation (MAP) program. A MAP program might propagate the registration change of an attribute to an identical change of declaration by all other application participants in the bridge, or it might use the changed attribute registration to calculate a changed attribute declaration for a single other participant, or some other subset of participants.

These propagation rules, combined with the standard declaration-to-registration reaction, produce a chain reaction of declaration changes and registration changes throughout the LAN. This can result in the update of the application-defined attributes in all the participants reachable on the network by the participant making an attribute declaration change.

Figure 2 illustrates how a declaration at an endpoint participant results in a chain of registrations and propagated declarations that flow through a network. The nodes labeled “A” through “D” are endpoints, and “X” and “Y” are bridges. SRP participants are labeled with “D” if they are declaring an attribute and “R” if the attribute is registered in response to a link peer’s declaration. The arrows interior to the bridges represent the operation of the MAP program for SRP.

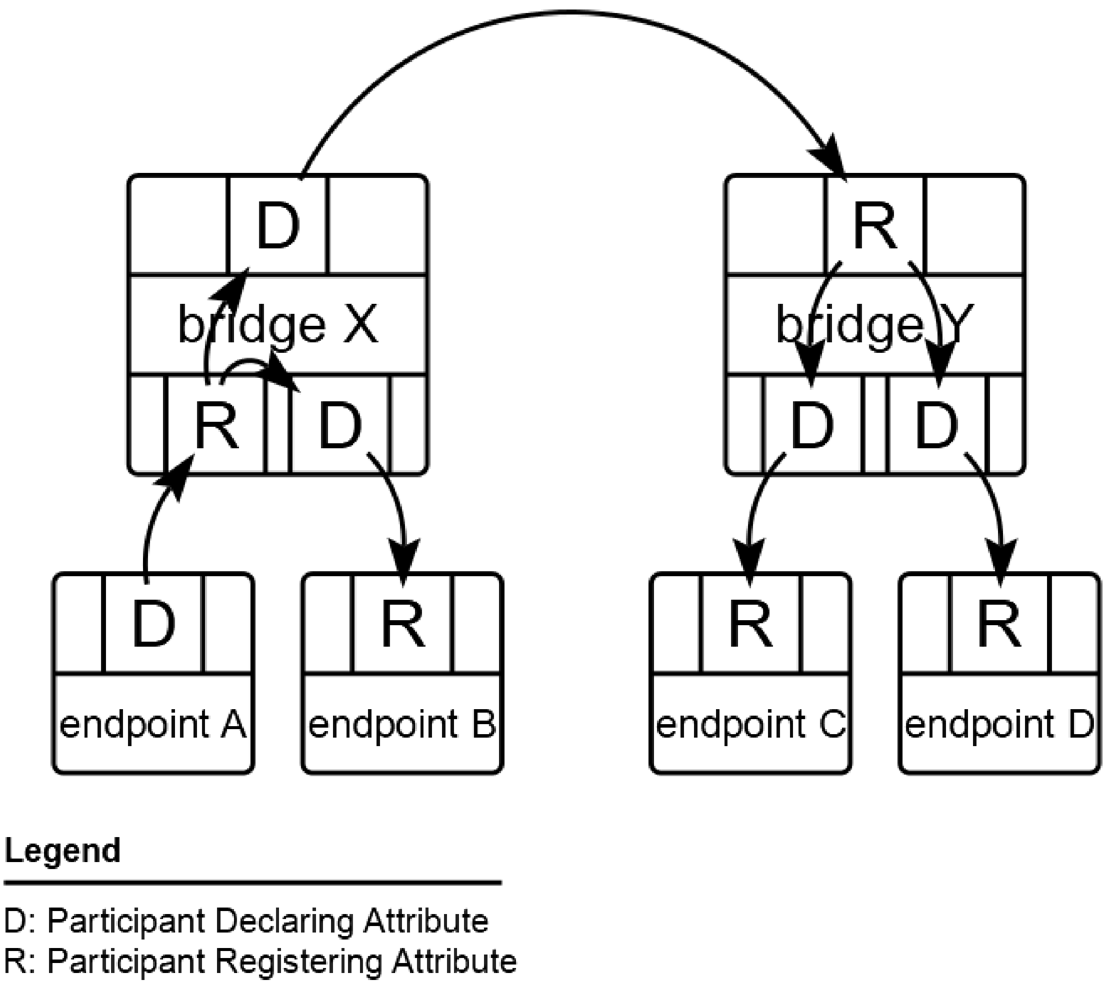

Figure 2 - Attribute Change Propagation

In addition to basic reaction and propagation, MRP provides a "new" flag that may be used to mark the initial declaration and resulting propagations as relating to a newly-declared attribute. This flag does not mean anything to the MRP protocol itself, but it is given special treatment in the MRP implementation so that it can provide a useful distinction to higher-level protocols built on MRP. It is up to those higher-level protocols to decide whether to use the "new" flag and how it should be interpreted.

Another operational detail that is a result of MRP's implementation is that the declaration-to-registration reaction across a network link does not necessarily happen immediately after the declaration is made to the MAD layer. Instead, declaration changes may be queued until a *transmit opportunity* is acquired, at which point all the queued declaration changes are passed to the link peer to create the corresponding registration changes there. The delay is not long, but this behavior is likely to be apparent in network traces and should be understood as normal. This behavior is only required for shared and simulated shared media, not for point-to-point protocols such as Ethernet. Point-to-point media receive a transmit opportunity as soon as feasible after one is requested, with the limitation that only 3 transmit opportunities may occur in any 300ms window (see IEEE 802.1Q 2011 section 10.7.11).

Although attributes are persistent at the MRP API level, it should be understood that their persistence is managed by MRP through continual refreshing of declarations based on timers. They can be removed both by explicit protocol actions as well as by failure of the declaring participant to maintain the declaration in a timely manner. This allows the attributes declared by participants that silently fail to be flushed from the system after a short delay, but there is also some delay in the removal of attributes throughout the LAN due to explicit withdrawal of declarations.

2.2. MRP Components and Interfaces

Each MRP application participant is composed of an Application component and a Multiple Attribute Declaration (MAD) component. Additionally, participants within a bridge may be connected to one another via the application's MAP program. Figure 3 shows how these components interact with one another and the Link Layer Controller (LLC), which is responsible for sending and receiving network traffic on a port.

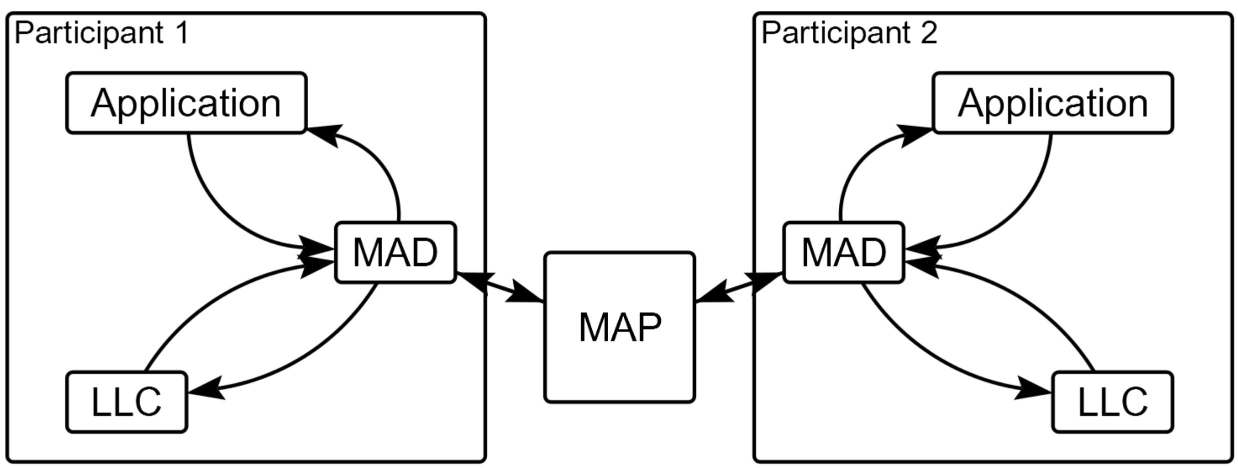

Figure 3 - SRP Components

2.2.1. Application

The Application component keeps track of the attributes understood by the MRP application and what they mean in terms of the functionality of the application protocol. It directs the MAD to make or remove attribute declarations, and it receives indications from MAD when registrations of attributes are made or removed. This control and communication of the MAD is guided by the application-specific higher-level interfaces.

- **Attributes**
  Each attribute defined by a MRP application is managed by its Application component. Each attribute has a type, which determines its representation, and a set of permissible values. The meaning of these values is defined by the application.

- **Interfaces**
  The Application has a set of software interfaces that are used to make and revoke attribute declarations.

2.2.2. MAD

The MAD component is responsible for the actual protocol execution. Other components control it through **request** interfaces and it notifies other components of attribute changes through its **indication** interfaces.

- **MAD_Join.request(attribute_type, attribute_value, new)**
  The Application and MAP components request the MAD to make a declaration of an attribute via this interface. The parameters indicate which of the application's attribute types is being declared and what value that attribute declaration should take. The "new" flag indicates that this is the initial declaration (or propagation) of the attribute and is used in an application-specific manner.

- **MAD_Leave.request(attribute_type, attribute_value)**
  The Application and MAP components request the MAD to remove a declaration of an attribute via this interface. The parameters indicate which of the application's attribute types is having its declaration removed and the value of the declaration to be removed.

- **MAD_Join.indication(attribute_type, attribute_value, new)**
  MAD uses this interface to notify the Application and MAP components of an attribute registration made as a result of its network peer performing a declaration.

- **MAD_Leave.indication(attribute_type, attribute_value)**
  MAD uses this interface to notify the Application and MAP components of the removal of an attribute registration made as a result of its network peer removing its attribute declaration.

2.2.3. MAP

The MAP component is shared between all of an application's participants in a bridge. It receives notification of registrations via the MAD's indication interfaces. It processes the indication from one participant and makes subsequent calls to the request interfaces of the other participants based on the MAP program of the application and the MAP Context of the indicating participant.

The MAP Context is a set of ports that constitute the active topology, such as defined by a VLAN or spanning tree. The contexts that an application may operate in are defined by the application.

2.2.4. Interactions

Figure 4illustrates the Application-level interactions that occur as a SRP attribute, such as a Talker attribute, propagates through a network. The network topology is identical to the one previously illustrated in Figure 2.

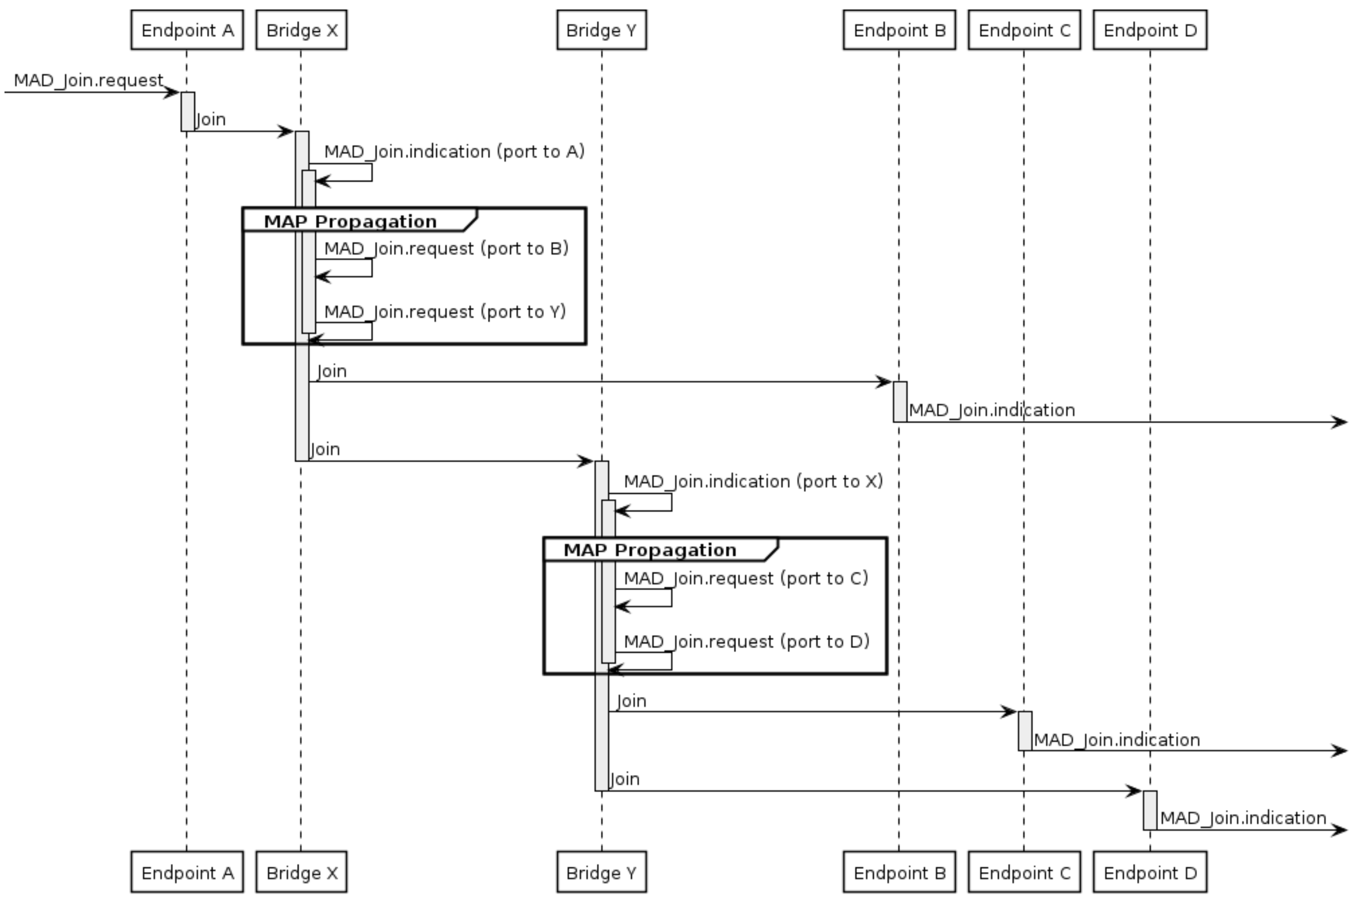

Figure 4 - Attribute Change Interactions

2.3. MMRP

SRP optionally uses MMRP to control the propagation of Talker registrations throughout the network. This option is known at talker pruning. When talker pruning is enabled, propagation of Talker attributes will only happen across ports that have been added to the MMRP group for the destination MAC address recorded in the Talker attribute. Talker pruning should be left disabled for standard AVB use as per 802.1BA-2001.

2.4. MVRP

In the context of SRP, MVRP is used by network bridges and end stations to declare membership in the VLAN in which a stream is being sourced. This allows the VLAN and Data Frame Priority that the Talker will use to tag its streaming data frames to be propagated from Talker to Listener. MSRP does not allow streams to be established across bridge ports that are members of the untagged set for the VLAN ID of that stream's source, so the Listener must use this information to join the correct VLAN if necessary. Of course, the Talker should be a member of the correct VLAN as well.

2.5. MSRP

The MSRP protocol is an MRP application to manage attributes that describe the intention of Talker and Listener end stations to transmit and receive streams. As Talker- and Listener-declared attributes are registered in the bridges between them, the bridges update their internal forwarding, queuing, and filtering state tables in order to support the QoS functionality specified by the AVB standards.

2.5.1. Attributes

The MSRP attributes are encoded in vector form, which means that a series of contiguous attributes can be encoded efficiently by providing a **first value** and the **number of attributes** being declared. Subsequent values are derived from their predecessors recursively by applying an **increment rule** determined by the attribute type.

A detailed description of how this encoding works is beyond the scope of this document[1], but a result of this is that devices that manage a large number of streams should make an effort to do so from a continuous range of resource values so that the stream reservation attributes can be represented and communicated efficiently by the MSRP protocol.

2.5.2. Talker Advertise attribute

This attribute represents the capability and intention of the declaring participant Talker to source a stream with the parameters described in the attribute. It will propagate throughout the SRP-enabled LAN except where insufficient resources to support the stream reservation are available.

The attribute carries the following parameter data:

- **StreamID**
  Eight octets uniquely identifying the stream. It may be subdivided into a 48-bit MAC address associated with the Talker and a 16-bit unique ID used to differentiate different streams sourced by the same Talker. Other encodings of these eight octets are possible, but are beyond the scope of this document.

- **DataFrameParameters**
  The addressing information for the stream that will be used to configure the bridge's filtering tables for reservation entries. This parameter is further subdivided into the following fields:
  - **Destination MAC Address**
    The destination MAC address of the streaming data packets. Only one Talker is allowed per

[1] See 802.1Q-2011 Section 10.8 for general MRP encoding rules and Section 35.2.2 for MSRP-specific information, including detailed descriptions of all MSRP attributes.

destination address. It must be a multicast or locally administered address. SRP does not support using a destination address for both best-effort and streaming traffic.

  - **VLAN Identifier**
    The VLAN that will be used for the streaming data packets.

- **TSpec**
  The traffic specification for a stream. These are used to configure the stream traffic shaping mechanism in the bridge on the ports associated with the stream. The TSpec is subdivided into the following fields:
  - **MaxFrameSize**
    The maximum frame size that a Talker will produce as part of the stream, excluding media-specific framing.
  - **MaxIntervalFrames**
    The number of frames that the Talker will produce per class measurement interval.

- **PriorityAndRank**
  Information about the priority class and the emergency status of a stream. It is subdivided into the following fields:
  - **Data Frame Priority**
    Used to generate the Priority Code Point (PCP) tag for the data stream.
  - **Rank**
    A single bit to identify emergency vs. non-emergency streams. Emergency streams use the value 0, non-emergency use 1.

- **AccumulatedLatency**
  The worst-case latency that a stream could encounter from Talker to Listener. This value should not change after it has been registered by a participant. If a participant is sent an attribute that has had the AccumulatedLatency change from its previously registered value, it will change the attribute propagation from Talker Advertise to Talker Fail with the FailureInformation code "Reported Latency Has Changed". The talker initializes this value with an estimate of maximum expected delay between the egress of a packet from the talker's network interface and when it reaches its network peer on its path towards the listener. Each bridge on the path will add the maximum expected delay between packet ingress on its own port and arrival at the next peer on the path. This means the value will be different at each bridge, but will remain constant at each point of registration. Because this is a worst-case estimate, it could be much worse than the actual average latency.

2.5.3. Talker Failed attribute

Talker Failed represents the intention of a Talker participant to transmit a stream along with a lack of some necessary resource for that stream to be reserved at some point between the Talker and Listener. A Talker Advertise may be propagated as a Talker Failed if there are insufficient resources in the bridge for the described stream to be reserved. In that case, at the point where the Talker Failed declaration is first made the FailureInformation code is set to the reason for the failure.

The parameters of a Talker Failed attribute are identical to that of a Talker Advertise, with the addition of one extra parameter:

- **FailureInformation**
  A failure code indicating why the Talker stream is not available despite the Talker's intent to offer it. There is a fixed mapping of failure codes to textual description of failure reasons defined by MSRP. The BridgeID, which indicates the network bridge the problem occurred in, is included in this parameter's data.

2.5.4. Listener Attribute

The Listener attributes all belong to the same basic type, but they are also divided into several subtypes. They represent the aggregate intention and ability of one or more Listener participants to consume a particular stream.

Listeners declare their intention individually, but the intermediate bridges between Listener and Talker **merge** the Listener registrations for a stream as they propagate back towards the Talker.

The Listener Attribute carries the following parameter data:

- **Subtype**
  The subtype parameter divides the Listener attribute into three different kinds. The encoding supports a fourth kind, Ignore, but it is only relevant to the vector encoding of the Listener attributes and isn't meaningful to the high-level protocol operation.
  - **Ready**
    This subtype indicates that there is at least one Listener that both intends to listen and has successfully reserved resources, and that there are no Listeners that intend to listen but were unable to reserve resources.
  - **Ready Failed**
    This subtype indicates that there is at least one Listener that both intends to listen and has successfully reserved resources, but at least one other Listener intends to listen but was unable to reserve resources.
  - **Asking Failed**
    This subtype indicates that there is at least one Listener that intends to listen but was unable to reserve resources, but there were no Listeners that both intended to listen and were successful at reserving resources.

- **StreamID**
  The only parameter (aside from subtype) associated with a Listener attribute is the Stream ID that the Listener(s) are declaring an intent to consume. It is the same as the StreamID parameter of the Talker attributes, so they can be directly compared to determine if the attributes refer to the same stream.

2.5.5. Domain Attribute

These parameters together identify which class ID, Priority, and VLAN ID are associated together for a SRP Domain. These determine the Domain boundary and allow endpoints to discover what VLAN-tag parameters to use when constructing their streaming packets. This attribute **does not propagate**, it is only used between immediate link peers. One SRP Domain exists per SR class, and they do not necessarily have the same boundaries.

MSRP-enabled bridges use this behavior to determine if the link partner of each port's participant is a member of the same SRP Domain. Any bridge port where there is no Domain attribute registration or where there is a Domain attribute that does not match the Domain being declared on that port is considered a **boundary port** for the Domain. No stream reservations are allowed to cross boundary ports, but MSRP attributes will propagate across them with modifications to indicate failure.

There is no standard mechanism in MSRP for auto-configuration of the Domain attribute; it must be programmed via bridge management functionality into each network bridge so that the attributes match for each port that will be used for stream reservations. However, end stations may use Domain registrations made in their MSRP participants to discover the Domain attribute parameters used throughout the network and then alter their Domain declarations to match.

SRP defines default Domain attributes that all compliant devices must support, and endpoints should declare the default values unless they have been configured otherwise. Regardless of how end station configuration is managed, the end station participant must declare the correct Domain attribute in order to participate in stream reservation.

Domain attributes are as follows:

- **SRclassID**
  This is a numeric representation of the SR class that the attribute is describing. The AVB standards define two SR classes, class A and class B. Class A is for streams with tighter latency requirements and it has a shorter FQTSS observation interval, meaning its streaming packets are smaller and transmitted more frequently.

SR Class	SR Class ID
A	6
B	5

- **SRclassPriority**
  This is the numeric value of the Data Frame Priority that will be encoded in the VLAN tag of stream packets for the associated SR class.

- **SRclassVID**
  This is the numeric value of the VLAN ID (SR_PVID) that is suggested for use in the VLAN tag of stream packets for the associated SR class. Unlike SRclassID and SRclassPriority, a Talker may choose to use a different VLAN ID for its streams and still receive a reservation in the domain’s SR class.

2.6. Propagation Rules

The propagation rules of MSRP attributes are unique among MRP applications and define a great deal of the overall protocol functionality. The propagation rules vary for each attribute type, so each type will be described separately, but there are a few common propagation rules.

MSRP attributes will only be propagated to ports running the MSRP application. Attributes for a particular SR class will only propagate to ports that support that SR class, e.g. ports that only support class B streams will not propagate attributes relating to class A streams.

While MMRP and MVRP attributes are propagated by bridges that are not capable of supporting these protocols (they are forwarded as standard network traffic), MSRP uses a special MAC address that is specified by 802.1Q to be non-forwarding, so a VLAN-aware bridge that is not MSRP-capable should not propagate the attributes. That being said, some bridges (typically small unmanaged ones that don't understand VLANs) ignore the non-forwarding address rules and forward them anyway.

2.6.1. Talker Attributes

A Talker registration change on one port will result in declaration changes being made on all other SRP-enabled ports. The value of the declared attribute will depend on the registered attribute and whether there are sufficient resources to make a reservation from the registering port to each declaring port.

A Talker Failed registration will propagate as a Talker Failed declaration with mostly the same parameters on the other ports.

Talker Advertise registrations require some calculation. For each declaring port, a determination is made whether there are sufficient resources to make a reservation for the stream on that port. If there are, a Talker Advertise is declared on that port. If not, a Talker Failed is declared with the reason for failure in the FailureInformation parameter. This resource calculation only takes into account active reservations, not other Talker registrations that don't yet have matching Listener registrations.

As the bridge adds and removes reservations, the available resources change. Each time one of these changes occurs, Talker Advertise propagation must be re-calculated. A new reservation may cause one or more Talker registrations for inactive streams to propagate as Talker Failed if there are no longer sufficient resources to support those streams. An ending reservation may likewise cause one or more Talker registrations for inactive streams to propagate as Talker Advertise again if sufficient resources have become available for them.

As Talker attributes are propagated, their AccumulatedLatency grows by the portTcMaxLatency value for the bridge, which represents the worst-case latency that the bridge could add to the total packet latency.

If a Talker Advertise attribute would normally propagate to a port as a Talker Advertise but that port has been marked as a boundary port for the SRP Domain, it will instead propagate as a Talker Failed attribute with the failure reason being "Egress Port is not AVB-capable". There is one exception to this: If the SRP Domain VLAN is specifically blocked on a port, the failure will not propagate but will instead be stored internally. An administrator may retrieve it via SNMP management tools. This is to prevent snooping of VLAN structure through SRP protocol probing.

Talker pruning can further limit propagation of Talker declarations. When it is supported and enabled, the attribute declarations will only be propagated to ports where the attribute's destination MAC address has previously been added via MMRP. Talker pruning should be left disabled for standard AVB use as per 802.1BA-2001.

2.6.2. Listener Attribute

Listener attributes have unique propagation rules, as they can be merged as they propagate. Their propagation is also highly dependent on the Talker attributes that are registered at the time of propagation calculation. Declaration of Listener attributes can also affect propagation of Talker Advertise attributes as resources available for reservations become exhausted.

When a Listener attribute is registered on a port, the other bridge port participants are checked for a Talker registration with a matching Stream ID. If no matching Talker registrations are found, the Listener attribute is not propagated.

If a participant is found with a Talker registration that has a matching Stream ID, then propagation to that participant will occur. If the Talker attribute is a Talker Advertise and the Listener attribute is a Listener Ready or Listener Ready Failed, the bridge will check for sufficient bandwidth and add the reservation to its tables.

If the check fails, then both the Talker and Listener attribute declarations will change to reflect the failure. This allows failure to propagate towards both the Talker and Listeners of a stream from the point at which it occurred, though only the Listener will receive detailed information as to why the failure occurred.

A reservation attempt for a stream with its rank bit set to the emergency level will succeed even if other reservations on the link reduced available resources below its reservation needs. In this case, one or more existing streams will have their reservations invalidated in order to free resources for the emergency reservation. Those invalidated resources will result in a propagation and merge of Listener failure towards the Talkers and Talker Failed with the reason encoded towards the Listeners.

If the participant already has a Listener registration due to propagation from a different Listener on a different port, the two Listener attributes will be merged according to the following rules:

- If the Talker attribute is a Talker Failed, the Listener attribute will propagate as Listener Asking Failed regardless of the attributes to be merged.
- If the current Listener attribute is Listener Ready:
    - Merging a Listener Asking Failed or Listener Ready Failed will result in propagating a Listener Ready Failed
    - Merging a Listener Ready will continue to propagate the Listener Ready attribute

- If the current Listener attribute is a Listener Ready Failed, the merge will result in a Listener Ready Failed as well
- If the current Listener attribute is a Listener Asking Failed
    - Merging a Listener Ready or Listener Ready failed will propagate a Listener Ready Failed
    - Merging a Listener Asking Failed will have no effect

The effect of these rules is to produce a Listener Asking Failed registration at the participant declaring a Talker Advertise if none of the participants declaring a Listener attribute can receive the stream, a Listener Ready Failed if at least one of the participants declaring a Listener attribute can receive the stream but at least one cannot, and a Listener Ready if all of the participants declaring a Listener attribute can receive it.

Since actual bandwidth reservation and population of the filtering entries in the bridge is a byproduct of the propagation of Listener attributes, a Talker that receives a Listener Ready or Listener Ready Failed attribute registration can be assured that when that registration is made, bandwidth has been reserved for that Stream ID to at least one Listener that is prepared to receive it.

According to the specification, a listener's intent to listen to a stream that it has not yet received a Talker registration for should be declared as a Listener Asking Failed rather than a Listener Ready. However, this declaration will have to be withdrawn and re-declared once the Talker Advertise arrives, so it is not currently recommended. This is an area under consideration for future protocol optimization, so in the future there may be a useful declaration for such an "anxious listener" to make.

2.6.3. Domain Attribute

There is no MAP program for the Domain attribute, so it **does not propagate** between bridge port participants. This means that declaration of this attribute only produces a registration on a participant's immediate link partner.

2.7. Usage Constraints

MSRP, due to its nature as an MRP protocol, is implemented by timer-driven state machines that ensure that the attributes it propagates represent the current state of the LAN's stream reservations. The timers ensure that changes in attribute registrations due to declaration changes, participants joining or leaving the LAN, or other dynamic events are propagated in a timely manner.

MSRP attributes are also implemented as vectors, which allow very large numbers of attributes to be efficiently processed and propagated. This is important because, due to the timers involved in the implementation, there are strict deadlines for processing MSRP protocol data and some network bridges have very limited bandwidth between their ports and management processors even though they can pass vast amounts of data at near line speed between external ports.

These facts place additional constraints on the usage of the SRP interfaces, especially for systems with high port count bridges, many potential stream reservations, or many dynamic changes in reservations. In those situations, care must be taken in the choice of attribute parameters and the timing of dynamic changes in order to retain correct and efficient operation. The following section describes in more detail how that can be done.

2.7.1. Contiguous Attribute Parameters

MSRP vector attributes are represented by FirstValue and NumberOfValues fields along with an incrementing rule. This allows a potentially huge number of attributes to be represented by a single protocol record. By keeping in mind the incrementing rules for different attribute types, the attribute parameters for the system can be chosen to maximize the opportunities for this efficient representation.

Talker and Listener attributes share the same general increment rule. The FirstValue field contains all the parameters of the attribute, and the increment rule is as follows:

1. Add 1 to the Unique ID portion of the StreamID field
2. Add 1 to the Stream destination_address field

Of course, if an attribute does not have a Stream destination_address field, only the Unique ID portion of the StreamID field is incremented.

This increment rule implies that the following streams can be represented by a single record, as they follow a pattern where their Unique ID and destination_address fields each increment by 1 from one stream to the next:

```
StreamID:            10:20:30:40:50:60 00:00
destination_address: 01:02:03:04:05:00
other fields...

StreamID:            10:20:30:40:50:60 00:01
destination_address: 01:02:03:04:05:01
other fields...

StreamID:            10:20:30:40:50:60 00:02
destination_address: 01:02:03:04:05:02
other fields...

StreamID:            10:20:30:40:50:60 00:03
destination_address: 01:02:03:04:05:03
other fields...
```

The vector encoding for those four streams would resemble this:

```
FirstValue:
  StreamID:            10:20:30:40:50:60 00:00
  destination_address: 01:02:03:04:05:00
  other fields...
NumberOfValues: 4
```

The IEEE 1722 standard provides a protocol, the Multicast Address Allocation Protocol (MAAP), for reserving a sequential range of multicast MAC addresses within a larger range set aside for MAAP, or within a system-defined range. The MAAP protocol ensures that Talkers can use continuous sets of destination_address values without overlapping with other Talkers. MAAP or some other reservation scheme should be used to ensure that this vector attribute encoding can be taken advantage of when large number of streams are going to be provided by a Talker.

2.7.2. Reservation Change Timing

The state machine timers of the MSRP protocol impose some limitations on how quickly an existing attribute may be changed. If a Talker wishes to re-use a StreamID with different parameters from a currently advertised stream, it cannot just withdraw its current declaration and then immediately make a new one.

In order to successfully change an attribute, a participant must wait two full LeaveAllTime periods (see 802.1Q-2011 Table 10-7 for details) after withdrawing the declaration in order for the change to propagate throughout the LAN. Only after this time elapses should the participant re-use the StreamID in another attribute declaration.

Failure to follow this timing rule can result in corruption of streaming data.

2.7.3. VLAN Membership

When a Listener wishes to listen to an advertised stream, it must first ensure that its bridge port is a member of the VLAN that the stream will be transmitted on. This is communicated by the Talker Advertise attribute's VLAN Identifier field. It will often be, but is not required to be, the same as the default VLAN of the SRP Domain it is associated with.

When a Talker Advertise for the desired stream is indicated by the Listener's interface, the MVRP protocol must be employed by the Listener to join the VLAN before invoking the interface to declare a Listener Ready attribute. If this is not done, the stream's data frames could be dropped by the bridge instead of being forwarded to the Listener.

If a bridge does ingress VLAN filtering, a Talker on that bridge must be a member of the VLAN as well in order to get the streaming data onto the network. Future maintenance of the SRP protocol standard will therefore require Talkers to join the VLAN before invoking the interface to declare a Talker Advertise attribute.

2.7.4. Fan-In Limitations

This section describes a bridge configuration parameter that could cause otherwise difficult-to-interpret behavior if it is put into use on a large and heavily-provisioned AVB network. It does not place any implementation burden on endpoints, but the information may be useful for diagnostic purposes.

It may be that an AVB-capable bridge with a large port count will face increased maximum latency values for frames traversing the bridge. The mechanism for this hypothesized latency increase is the increased amount of possible interfering traffic when all ports may participate in a SR class. However, the fan-in factor has been removed from the 802.1BA latency formula, so it is not likely to cause real problems.

MSRP provides a parameter that may alleviate this potential problem in some cases if it does occur. By setting msrpMaxFanInPorts to a non-zero value less than the total number of ports, latency could be reduced. The price for the reduced latency is the reduction in the number of ports that can simultaneously have active listeners attached to them.

All ports still have the potential to forward streams to listeners, but once the number of fan-in ports allowed has been reached, Listener Ready registrations from other ports will be propagated as Listener Asking Failed and Talker Advertise registrations will propagate to them as Talker Failed. The failure code in the Talker Failed attributes will indicate that the reason was the msrpMaxFanInPorts parameter.