MRP-Based Stream Reservation

The Multiple Stream Registration Protocol (MSRP) for AVB stream reservation

Norman Finn

MPRP and MSRP

MRP, MVRP, and MMRP

- IEEE Std. 802.1ak-2007 Multiple Registration Protocol (MRP) is a robust, efficient protocol for declaring attributes to be registered in a database in each port of each bridge (optionally, station) in a bridged network.
- Two applications (so far) are based on MRP:

  MVRP: Attribute is a VLAN ID. Stations or configured Bridge Ports make (withdraw) declarations if they do (not) need to receive frames for a given VLAN ID. If a VLAN ID is registered on a Bridge Port by MVRP, the Bridge knows that that frames for that VLAN ID should be transmitted on that Bridge Port.

  MMRP: Attribute is a MAC address, often a multicast address. Stations or configured Bridge Ports make (withdraw) declarations if they do (not) need to receive frames for a given address. If an address is registered on a Bridge Port by MMRP, the Bridge knows that that frames for that address should be transmitted on that Bridge Port.

MPRP and MSRP

- This presentation offers two new protocols:

  A variant of MRP, the Multiple Payload Reservation Protocol (MPRP), with added capability, that can serve as the basis for specific applications.

  One application, Multiple Stream Reservation Protocol (MSRP) based on MPRP.

MPRP: Multiple Payload Registration Protocol

- MPRP separates the Attribute registered by MRP into a {key, payload} pair, and serves as the basis for a new class of applications.
- The payload is registered along with the key.
- On a given port, there is at most one {key, payload} registration for a given key K.
- Each application must have rules for:

  Registration: What {key K, payload R} is to be registered if a registration for {key K, payload P} exists (or not), and a declaration for {key K, payload Q} is made or withdrawn?

  Propagation: What {key K, payload X} declaration is to be made or withdrawn on a bridge port, given the port's registrations (if any) for {key K, payload Y}, and the other bridge ports' registrations for key K?

MPRP: Multiple Payload Registration Protocol

- The rules for propagating MPRP **keys** are the same as for MRP Attributes:

  Registrations with the same **key** are treated as if they were a single registration.

- The rules for propagating MPRP **payloads** can be rather arbitrary.

MSRP

- MSRP defines payloads and rules for propagating registrations for both Talkers and Listeners.
- The output queues are configured, and the Group Registrations are maintained, using these MSRP registrations.
- Full compatibility and interoperability with non-AVB devices' existing multicast distribution mechanisms is maintained, though without delivery guarantees, of course.
- MSRP is compatible with RSVP.
- The cases of multiple Talkers on the same stream, shared media LANs, unicast streams, and streams on duplicated destination addresses, are all handled.

MSRP vs. P802.1Qat/D0.8

- SAME: Both use MMRP (or GMRP or IGMP) to indicate the Listeners' needs for streams.
- DIFFERENT: MSRP uses an MRP/MPRP based protocol for sending Talkers' stream reservations, instead of defining a new time-based protocol. The queues are not configured by the stream reservations alone.
- NEW: MSRP also uses that MRP/MPRP based protocol for sending Listeners' responses. They:
  a) Cooperate with the Talkers' stream reservations to configure the output queues and enable transmission of the stream; and
  b) Notify the Talker(s) when the network is ready for the data stream.

MSRP primary control flow

MSRP control flow summary

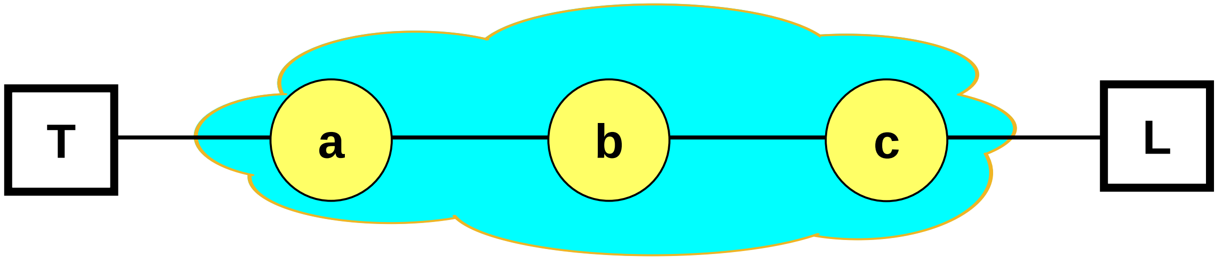

- Stream reservation requires that these two announcements be separate:

  I can supply stream S, and it has these characteristics.

  I have configured my queues to handle stream S, so transmission can begin.

- Without this interlock, the source can start transmission of the data stream before the network is configured to handle it, thus disrupting existing streams.

MSRP control flow summary

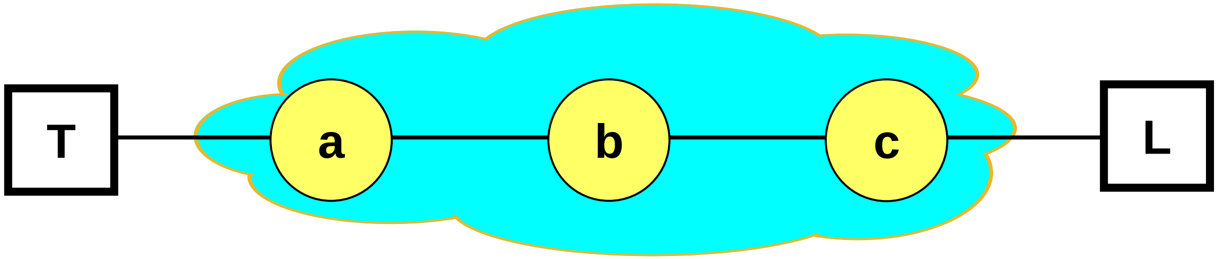

- MSRP allows any of MMRP, GMRP, IGMP, or MSRP to be used as a precursor to the “I can supply” step.
- MSRP adds a fourth step, failure notification, so that failed flows do not needlessly occupy resources.

MSRP control flow: Asking

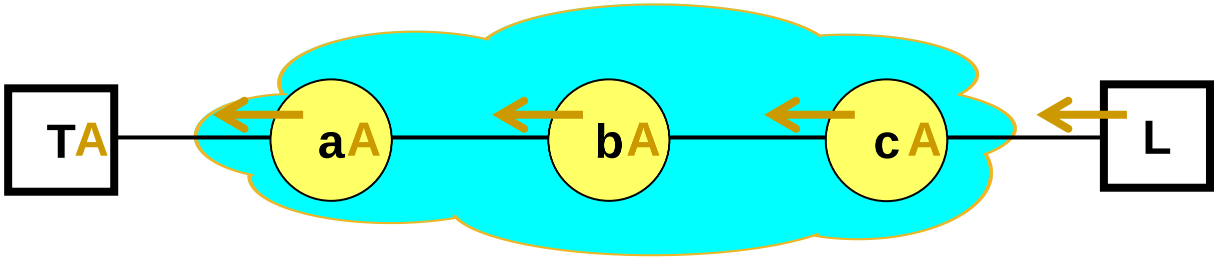

- One Talker **T**, one Listener **L**, three AV bridges **a**-**c**.
- L sends an asking declaration "A" that it wants a stream.
  - This could be an MSRP asking declaration.
  - This could be an MMRP declaration.
  - This could be a GMRP declaration.
  - This could be a "snooped" IGMP declaration.
- The AV Bridges and Talker **T** register this declaration.

MSRP control flow: Offering

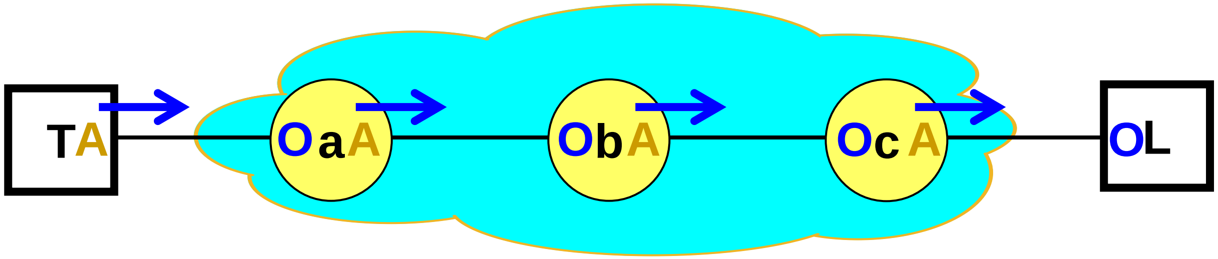

- Talker **T** issues an MSRP offering declaration.
- The offering declaration includes the ranking of the reservation's importance, its required bandwidth, its Ethernet priority level and destination MAC address, and optionally, a higher-layer stream ID.
- The AV Bridges and the Listener register the offer.
- The offering can be done before, during, or after the Listener's ready or asking registration.

MSRP control flow: Ready

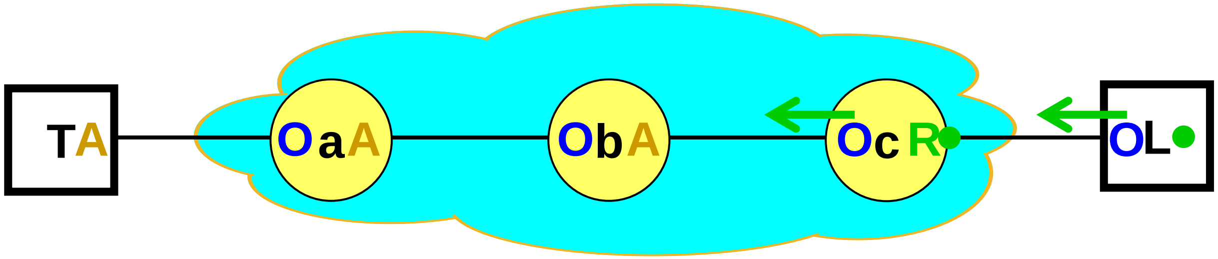

- When the Listener has configured itself to receive the stream, it issues an MSRP ready declaration.
- Only if a bridge has registered both the offering and the ready declarations, it configures its output queues on that port, and propagates the ready declaration.

  Only bridge **c** has, in this diagram, configured its output port.
- Regardless of the order in which the registrations are made, it is the presence of an offering registration on an input port, and a ready registration on an output port, that triggers the bridge to configure the output port.

MSRP control flow: Ready

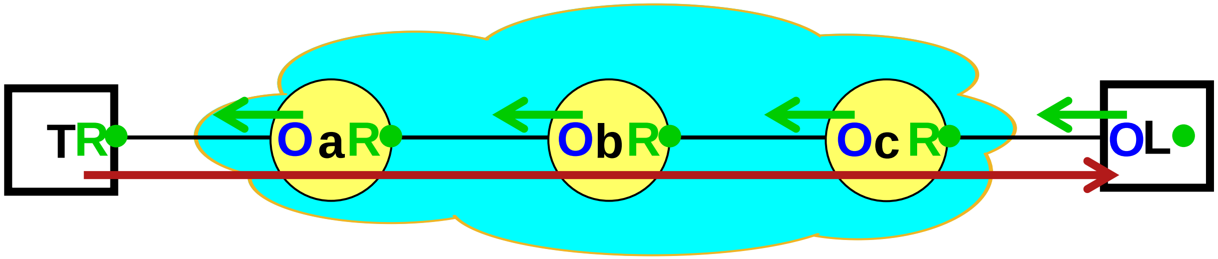

- When the ready declaration has been registered in Talker **T**, it knows that the network is ready to receive the stream.
- Talker **T** can start sending the data stream.

MSRP failure scenarios

MSRP control flow: Offering failure

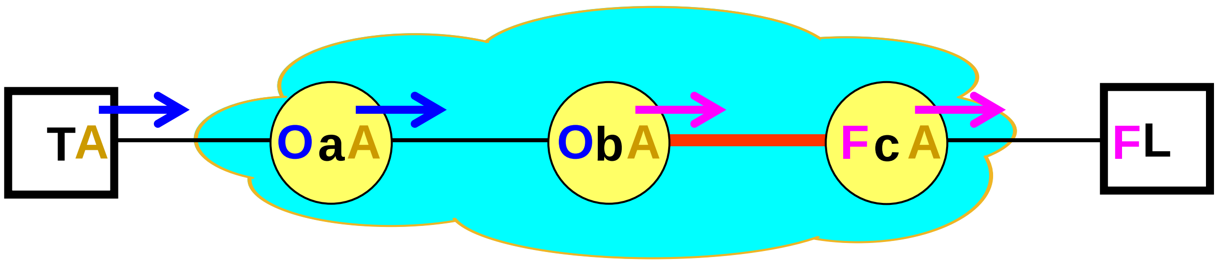

- Talker **T** issues an MSRP offering declaration.
- Link **b-c** has insufficient bandwidth to carry the flow.
- Therefore, AV Bridge **b** converts the offering registration to a failed declaration.
- Listener **L** knows that the stream reservation failed, so does not expect the data.

MSRP control flow: Offering failure

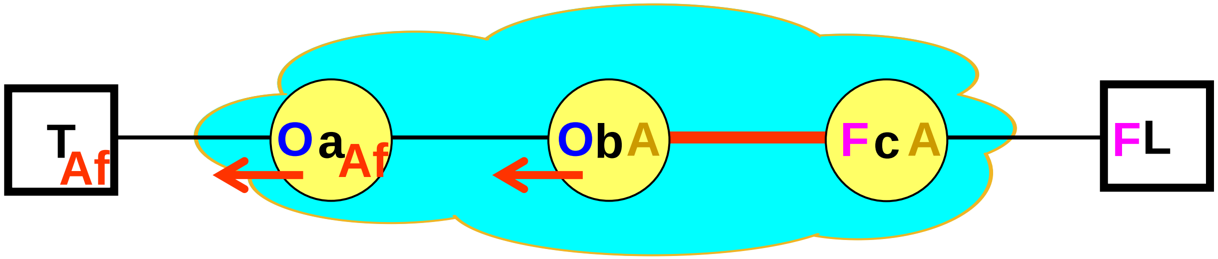

- In addition, AV Bridge **b** converts the asking registration to an asking failed declaration.
- All AV Bridges all know that the stream reservation failed, so none reserve any queue resources for it.
- Talker **T** knows that the flow reservation failed, and so does not send the data.

MSRP control flow: Ready failure

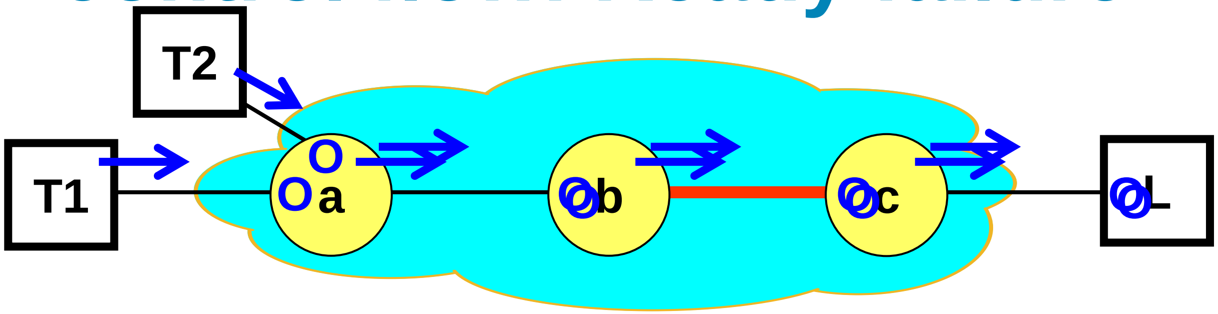

- No asking declarations have been made, yet.
- Talkers **T1** and **T2** both issue MSRP offering declarations. Link **b-c** can support only one of the two flows.
- AV Bridge **b** does not select which one to refuse, because it does not know whether it is actually on the path for either flow; it has not yet seen any asking or ready declarations.

MSRP control flow: Ready failure

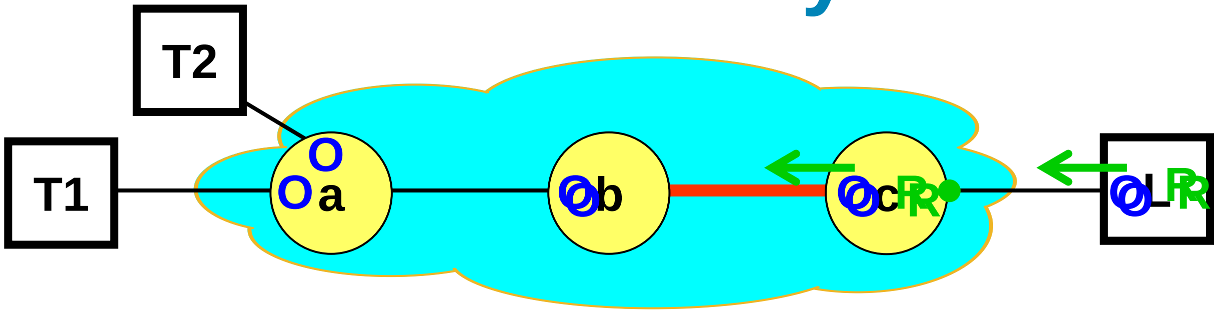

- When Listener **L** sends its ready declarations, AV Bridge **c** is able to set up the queues for both streams, but not so for AV Bridge **b**; the bandwidth of the link is insufficient.
- AV Bridge **b** uses the rankings of the two stream reservations to decide which one to grant. Let us say that **T1-L** succeeds, and **T2-L** fails.

MSRP control flow: Ready failure

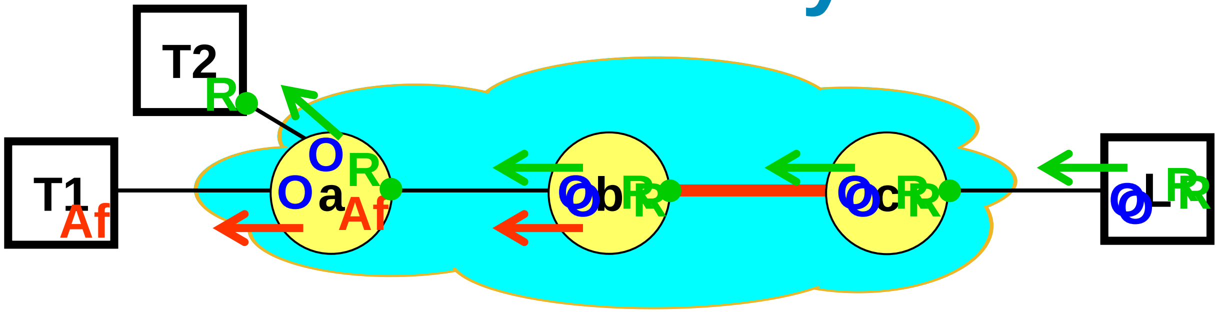

- AV Bridge **b** sends the requisite declarations to tell the two Talkers what happened.
- It sends a ready declaration for the **T2-L** stream, and an asking failed declaration for the **T1-L** stream.

MSRP control flow: Ready failure

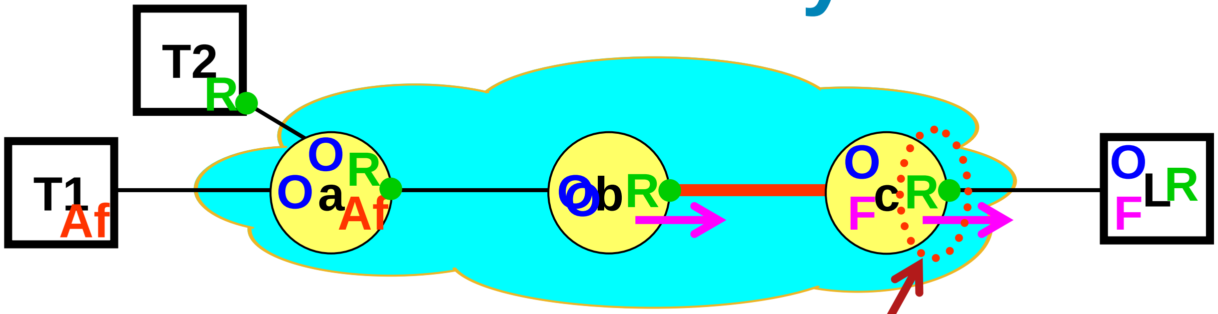

- AV Bridge **b** also tells Listener **L** about the failure.
- Note that AV Bridge **c** must revise its queue configuration; it thought that both registrations were going to succeed.
- The reconfiguration of the **c**-**L** link could cause a brief flap of some other flow(s), if the failed new flow is higher-ranking that some flow(s) already using the **c**-**L** link.

MSRP control flow: Multipath failure

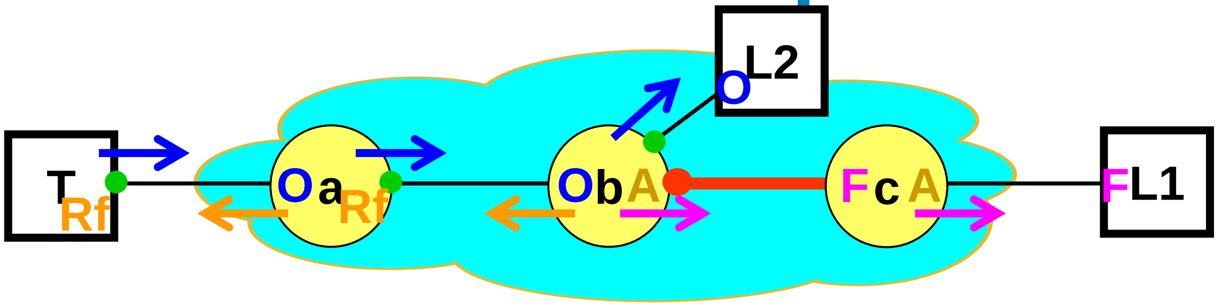

- If one path (**T-L1**) fails, and one path (**T-L2**) succeeds, then AV Bridge **b** returns a ready failed declaration towards Talker **T**.
- Talker **T** then knows that it can send the stream, but that not all Listeners will receive it.
- Listener **L1** knows not to expect the data, and Listener **L2** knows to expect it.

MSRP control flow: Listener to Talker

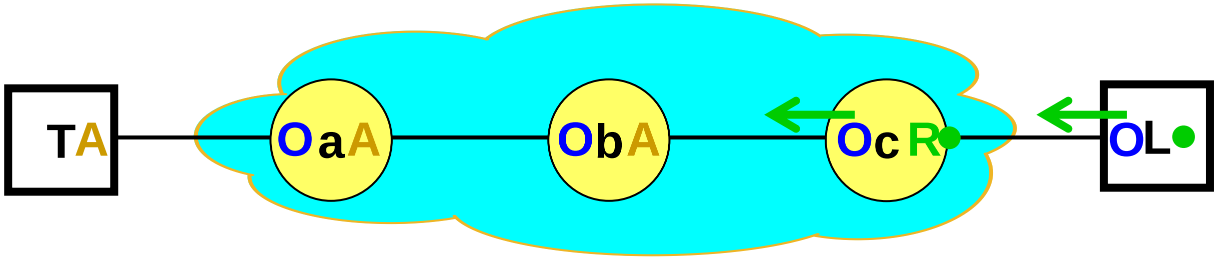

- MSRP builds the string of configured queues from the Listener towards the Talker.
- This takes, in the worst case, three passes across the network:

  “Asking”, “offering”, “ready”.

- Why not build the string of configured queues from the Talker towards the Listener?
- Wouldn’t that take only two passes?

  “Asking”, “granting”.

MSRP control flow: When does flow start?

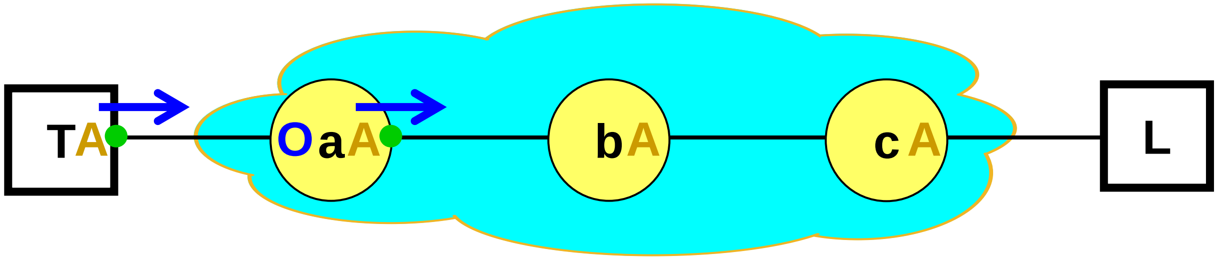

- There is one major problem with building the queues from the Talker to the Listener:

  When does the data flow start?

- If the Talker starts sending as soon as it receives the asking declaration, then it will transmit before AV Bridge **a** has configured its queues.

  (AV Bridge **a** cannot configure its queues before passing on the asking declaration, because it doesn't yet know where the Talker is; to do so would waste queue resources in AV Bridges that are not in the path of the stream.)

MSRP data flow: When does flow start?

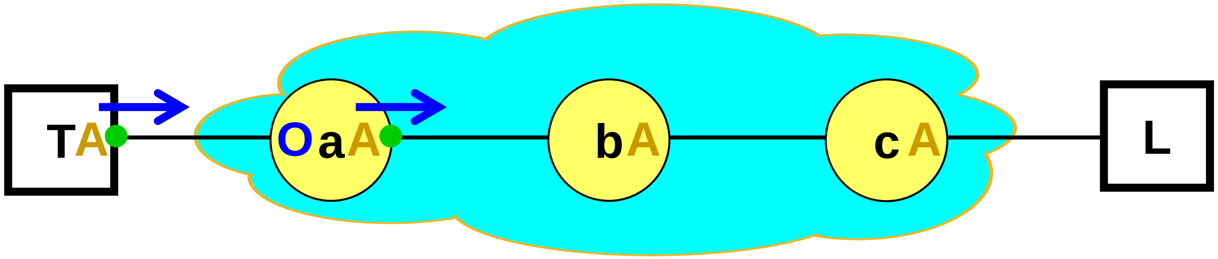

- This problem could be avoided, if AV Bridge **a** configures its group filtering database to discard the flow before it passes on the asking declaration, but then it would have to stop any existing flow using that same MAC address.
- This would make unicast stream reservation awkward, since all data flows would have to cease until the direction to the Talker was determined.

MSRP data flow: When does flow start?

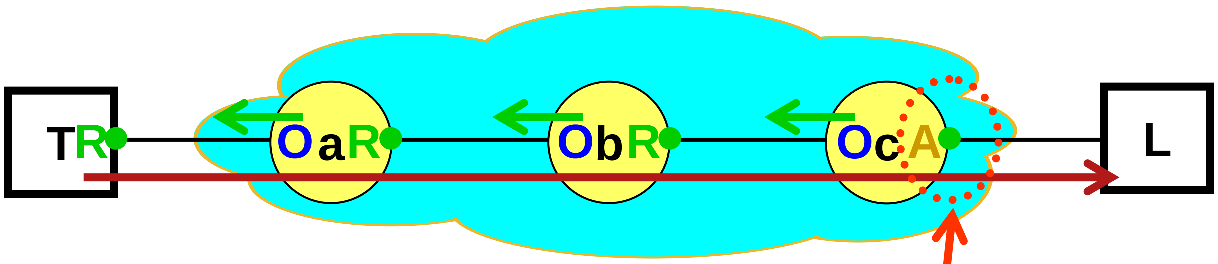

- Also, the scheme as presented allows an AV-capable Talker to supply a bandwidth-reserved stream to a non-AV-capable Listener.
- The Listener issues a normal IGMP, MMRP, etc.
- The Talker issues the offering Talker registration.
- The AV Bridge nearest the listener, knowing that the Listener is non-AV, can set up the queues.
- We see, then, that non-bridge Listener stations need not be AV-capable!

Flow description collisions

Collisions

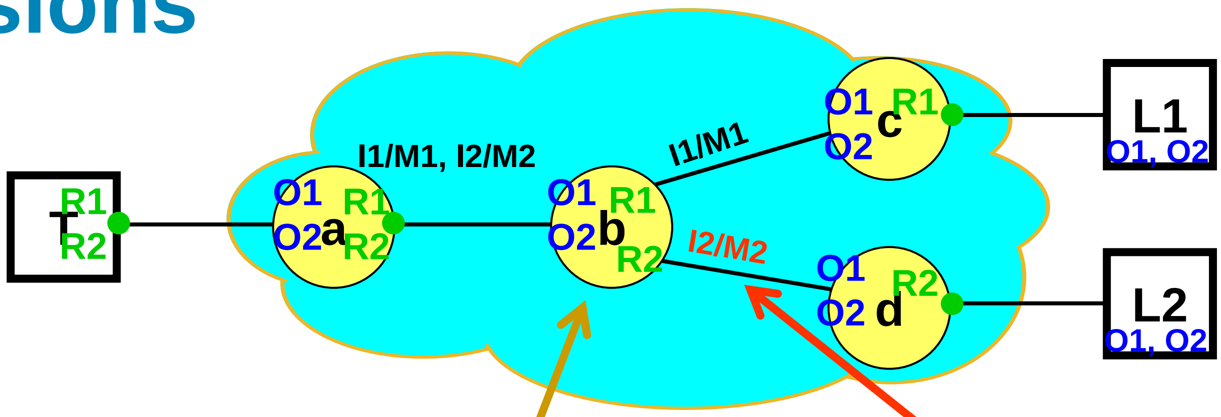

- Talker T and Listener L1 have a path up and running on IP multicast address I1, MAC address M1.
- Talker T and Listener L2 want to create a new path for IP multicast address I2, MAC address M2.
- Although I1 ≠ I2, M1 = M2. This can happen with IP.
- Because AV Bridge **b** can direct frames based only on destination MAC address, it must either send both streams on any given port, or it must send neither.

Collisions

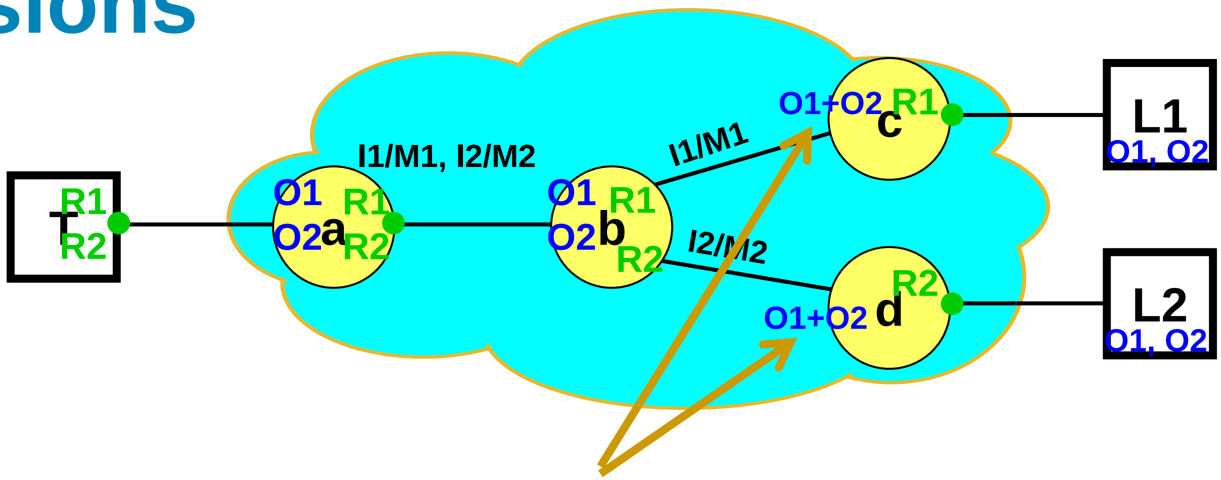

- Therefore, Bridge **b** must link the two MSRP offering declaration together, to say, "Sorry, but if I send one of these streams, I have to send both of them."
- This behavior is expected of a standard bridge, and since we need the destination MAC address as part of the declaration, anyway. We can reasonably accommodate this by simply adding a Boolean "MAC only" parameter indicating whether the bridge issuing the declaration can forward the data streams based on higher layer information, or must send all or nothing.

Collisions

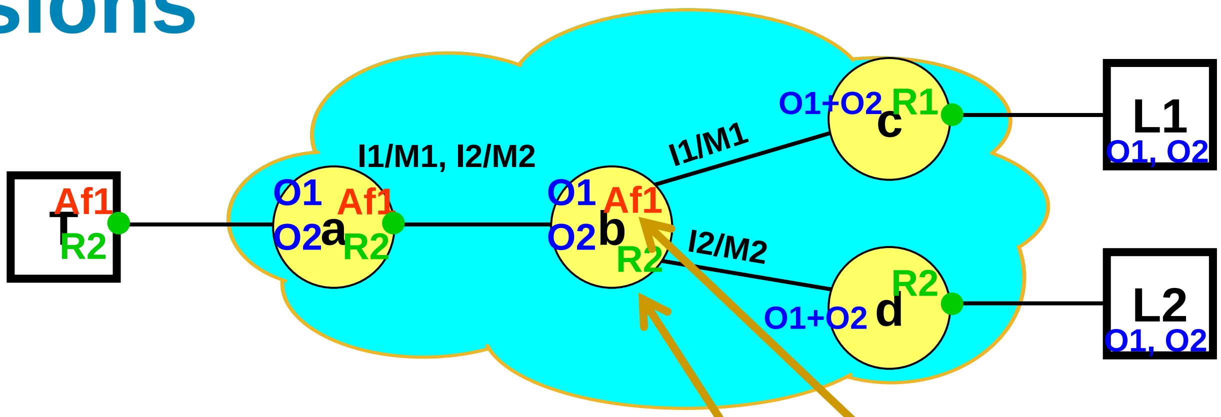

- In this example, AV Bridge **c** also cannot differentiate the streams, and the **c-L1** link has insufficient bandwidth to carry both of them, so **c** must refuse the **T-L1** stream.
- AV Bridge **d** can differentiate the streams, so it can suppress the unwanted stream, and accept the one it wants.

MSRP vs. RSVP

MSRP and RSVP

- The reader may have noticed that the MSRP Talker declaration is very similar to the RSVP (IETF RFCs 2205, 2750) “Path” message, and that the MSRP Listener declaration is very similar to the RSVP “Resv” message.
- This is intentional. RSVP is familiar to many network administrators and equipment developers, and the basic information flow that it supplies to routers is exactly what is needed for bridges, since the problems are very similar.

MSRP and RSVP

- SAME: The RSVP “Path” message is very similar to the MSRP Talker declaration; the RSVP “Resv” message is very similar to the MSRP Listener declaration.
- DIFFERENT: RSVP operates at Layer 3 among routers, and uses individual PDUs with requests and responses, and individual timers for each stream. MSRP operates at Layer 2 among bridges, and uses the MRP mechanism to group registrations, and minimize timers and responses.
- NEW: MSRP adds facilities for notifying Talkers and Listeners of the success or failure of the reservations, and for grouping streams that cannot be differentiated by bridges’ data forwarding mechanisms.

MSRP and RSVP

- The scenario with RSVP is thus:

  Listener can send an IGMP message or an MMRP or MSRP asking declaration to prompt the Talker.

  Talker sends both an RSVP Path message and an MSRP Talker offering declaration.

  Listener sends the RSVP Resv message for a specific bandwidth, and also sends the MSRP Listener ready declaration to initiate queue configuration.

  When the Talker receives both of these messages, it can start transmitting the stream.

Granting/denying requests: Shared media disallowed

MSRP: Listener declarations

- Four types of **Listener** declarations can be made from an applicant port for stream S at bandwidth W, to be delivered to that port:

  asking: One or more Listeners want stream S through the applicant port, but no path to any Listener is known to be ready to receive it. Data cannot be sent.

  asking failed: One or more Listeners would like to receive stream S through the applicant port, at least one cannot, for lack of resources along the path, and no path to any Listener is known to be ready to receive it. Data cannot be sent.

  ready: All known paths to Listeners are configured to receive stream S through the applicant port. Data can be sent.

  ready failed: At least one path to a Listener behind this applicant is "ready", and at least one is "asking failed". Data can be sent.

MSRP: Talker declarations

- Two types of **Talker** declarations can be made from an applicant port, with regard to sending stream S, with rank R, from Talker T, requiring bandwidth W, on VLAN ID V, address M, priority P, to be sent from that port:

  offering: Applicant port can be made ready to send stream S, if and when a ready or ready failed Listener declaration has been registered on the port.

  failed: Applicant port cannot be configured to send stream S, for lack of resources somewhere along the path from the Talker through the port.

MSRP: Listener declarations

- The **Listener** declaration and registration **key** consists of the:
  - Direction: Listener.
  - Stream ID: Stream identifier **S** used to match Talker and Listener registrations. The stream ID must include the VLAN ID **V** and destination MAC address **M** of the stream's data frames (can be unicast or multicast) and may include higher-layer ID information that is opaque to bridges.
  - Declaring system ID: Identifies the system making the declaration.
- The **Listener** declaration and registration **payload** consists of the:
  - Type: Either asking, ready, or asking failed, or ready failed.
  - Bandwidth: Total bandwidth **Wr** that the declaring port is prepared to receive. Must not be 0 if type is ready or ready failed, must be 0 if type is asking or asking failed.

MSRP: Talker declarations

- The **Talker** declaration and registration **key** consists of the:

  Direction: Talker.

  Stream ID: Stream identifier **S** used to match Talker and Listener registrations. The stream ID must include the the VLAN ID **V** and destination MAC address **M** of the stream's data frames (can be unicast or multicast) and may include higher-layer ID information that is opaque to bridges.

  Talker ID: Uniquely identifies the Talker **T** offering the stream.

- The **Talker** declaration and registration **payload** consists of the:

  Type: Either offering or failed.

  Bandwidth: The maximum bandwidth **Wt** of the stream to be transmitted by Talker T.

  MAC only: A Boolean flag, indicating that the bridge cannot differentiate streams for forwarding purposes except by VLAN ID and destination MAC address.

  Other QoS information: Other information may be required to characterize the stream's requirements, e.g., a hop count or a delay.

  Rank: Value **R** used by bridges to resolve competition for bandwidth.

  Priority: Priority value **P** to be carried in data frames belonging to the stream.

MSRP: Stream IDs

- Stream IDs must contain the VLAN ID and destination MAC address to be used by the data stream.
- Stream IDs may also contain a higher-layer stream identification. This information:
  - Is a facility available to the higher layers for implementing multiple streams on one destination MAC address
  - Could be used by an AV Bridge to differentiate streams, but that capability is not required to meet the goals of P802.1Qat.
  - Should use an OUI encoding so that any organization can define a format for the IDs that can be compared for equality with any other ID as a binary number.

MSRP: Stream IDs

- When comparing stream IDs from different registrations, they either match in their entirety, do not; they are not matched piece-by-piece.

  A declaration with a higher-layer stream ID does not match a declaration with the same VID and MAC address, but without a higher-layer stream ID.

  Declarations with either different MAC addresses, different VIDs, or different higher-layer stream IDs do not match.

- Therefore, use of the higher-layer stream ID must be coordinated among the Talker(s) and Listener(s); inclusion of the higher-layer stream ID is not a choice local to a station.

MSRP: Bandwidth

- Each link in an AV Bridged network can have a different amount of overhead for an ethernet payload byte.
  - 802.11 wireless media have more addresses and more interframe gap, than do 802.3 links.
  - Some links may carry additional tags such as 802.1AE security headers.
- Therefore, the Talker and Listener express bandwidth in terms of bytes per second required by the stream, and in terms of average Ethernet payload size.
- Each system, when comparing different streams' bandwidths, must use its knowledge of the overhead of a given link to convert the bandwidth and payload size information into a bits-per-second bandwidth requirement.
- This bit-per-second bandwidth can be used for simple integer comparisons.

MSRP: MAC only

- If two streams are offered that use the same VLAN and destination MAC address, but have different higher-layer stream IDs, a bridge that implements no more than is specified in 802.1Q cannot select whether to forward one stream, but not the other, through a given port; 802.1Q provides only the destination MAC address for making this decision.
- If the MAC only parameter is true, then the bridge sending the declaration cannot differentiate streams by higher-layer stream IDs.
- If the MAC only parameter is false, then the bridge can differentiate all streams for which there is a MAC address collision.

MSRP: Rank

- The rank **R** carried in a Talker declaration or registration payload has two parts:
  - An “importance” field; and
  - A “time of issue” field.
- The importance field is derived from higher-layer information, and is provided by the Talker station.
  - It is an unsigned integer.
  - Default value is 2^(n-1), where n is the number of bits in the field.
- The time of issue field is in IEEE 1588 time stamp format.
  - If the Talker does not have an 802.1AS clock that is in synch with the AV Bridge(s) on the port to which the declaration is sent, it sends 0 for the time of issue, and the bridges fill in the field.

MSRP: Other QoS information

- Information in addition to the bandwidth may be required to meet the goals of IEEE 802.1Qat; this is for determination by the 802.1 Working Group.
- Such information could include:

  A maximum delay that data frames in the stream can endure. This unsigned value, presumably measured in microseconds, would be set to a worst-case value by the Talker station, and reduced by each AV Bridge propagating the Talker registration according to the worst-case delay that bridge is likely to impose.

  A hop count, set to a maximum value by the Talker station, and decremented by each AV Bridge propagating the registration.

  If either goes negative, the bridge cannot accommodate the stream, and the Talker registration's Type becomes failed.

  The maximum Ethernet payload size to be used by the stream could help an AV Bridge compute a more accurate estimate of the delay to be expected by a frame in another stream.

MSRP: stations

- A **Listener** station:
  - May wait for a Talker declaration to learn of the availability and characteristics (P, M, Wt) of stream S.
  - May send an asking type Listener declaration (bandwidth Wr = 0) to indicate its desire to receive a stream, and to solicit a Talker registration.
  - When ready to receive, must send a ready type Listener declaration for bandwidth Wr ≠ 0 specifying the bandwidth the Listener is ready to receive.
  - May send an asking failed type Listener declaration (Wr = 0) to indicate that it is unable to receive the stream, at least for the moment, or it may withdraw its Listener registration.
  - May register Talker declarations to know whether or not the stream reservation has succeeded.

MSRP: stations

- A **Talker** station:
  - May wait for a Listener declaration to learn that some Listener wants a stream.
  - When ready to send, must send an offering type Talker registration for bandwidth Wt ≠ 0 to request the allocation of bridge resources.
  - Must register all Listener declarations for at least the streams the Talker is declaring, and must not transmit any stream's data for which there is no ready or ready failed Listener registration, or send at a rate greater than min (Wt, Wr) for that registration's Wr.
  - May choose to send nothing for any stream such that Wt > Wr, or for any stream with an active failed or ready failed type Listener registration.

MSRP: Streams lists

- A **Talker** station's list of sendable streams on port **P** consists of:
  - All streams that the Talker's higher-layer protocols want to send on port **P**. (A Talker station has no unsendable streams.)
- A **Bridge's** list of sendable streams on port **P** consists of:
  - All offering Talker registrations on all ports other than **P**.
  - Multiple Talker registrations for the same stream ID are considered to be a single registration, using the maximum of the registrations' bandwidths values **Wr**., and the best rank **R**.
- The list of sendable streams on port P is sorted by the 2-tuple, {rank, talker ID}, from the most important (lowest integer value) to the least important (highest integer value).

MSRP: Collision groups

- The sendable streams (see previous slide) are examined modified, output port by output port, based on stream collisions.
  - If multiple offering registrations in the sendable streams list have the same VID and MAC address for their stream ID, and if the MAC only parameter is set in any of those registrations, and if the processing bridge is unable, on port **P**, to make an independent decision as to whether to output each of those streams on this port, then those streams form a "collision group". They are considered a single entry in the sendable streams list, ordered by the highest-ranking stream in the collision group.
  - A stream that does not collide with any other stream is considered a one-stream collision group.

MSRP: Propagating Talker registrations

- The collision groups are examined in order, from highest- to lowest-ranked, and the queue configuration for port **P** is built, starting from a basis of no reserved streams.

  If not all of the streams in the collision group can be added successfully to the queue configuration built so far, a failed declaration is sent for each of the streams on port **P**, and the streams are all disabled for transmission on this port. These are the streams that cannot be sent, if the quality of service guarantees are to be met.

  Otherwise, if there is an asking failed Listener registration on port **P** for every one of the streams in the collision group, a failed declaration is sent for each stream in the group, and all are disabled for transmission. This tells the receiver that the data for this collision group will not be sent from this port, so that the receiver can release any resources assigned to the collision group's streams.

  Otherwise, if there is a ready or ready failed Listener registration on port **P** for any stream in the collision group, then an offering declaration is sent for each of the streams in the collision group, all of the streams in the collision group are added to the queue configuration, and all are enabled for transmission. These are the streams that must be configured on port **P**.

  Otherwise, an offering declaration is sent for each of the streams in the collision group, but the streams are not added to the queue configuration, and they are all disabled for transmission. These are the streams that could be configured on port **P**, should a receiver on port **P** indicate its readiness to receive the stream.

  When computing the bandwidth of a stream, a bridge uses the maximum of **Wt**, from the Talker registration, and **Wr**, from the matching Listener registration (if present).

MSRP: Propagating Listener registrations

- To decide what Listener declarations to issue from port **P**, the bridge constructs a list of all Listener registrations all ports other than **P**.
- Registrations for the same stream ID S are combined into a single entry:

  The resultant **Wr** is the minimum of all non-0 **Wr** values, or 0 if all are 0.

  The propagated declaration type is either asking failed or ready failed, if any of the other ports' registrations are asking failed or ready failed; it is asking or ready, if all other ports' registrations are asking or ready.

  The propagated declaration type is either ready or ready failed, if any of the other ports' registrations are ready or ready failed; it is asking or asking failed, if all other ports' registrations are asking or asking failed.

  (The intersection of the last two points gives an unambiguous declaration type.)

Granting/denying requests:
Share Media allowed

MSRP: stations

- In the following slides, differences between ports attached shared and non-shared media are highlighted in magenta.
- Slides unchanged from the preceding point-to-point media section are not included in this section.

MSRP: stations

- A **Talker** station:
  - May wait for a Listener declaration to learn that some Listener wants a stream.
  - When ready to send, must send an offering type Talker registration for bandwidth $W_t \neq 0$ to request the allocation of bridge resources.
  - Must register all Talker and Listener declarations, and must not transmit any stream's data for which there is no ready or ready failed Listener registration, or send at a rate greater than min (W_t, W_r) for all Listener registrations' W_r.
  - May choose to send nothing for any stream such that $W_t > W_r$, or for any stream with an active failed or ready failed type Listener registration.

MSRP: Media stream list

- A **Bridge or Talker** must construct a media stream list on each port **P,** consisting of:

  All collision groups that a Talker's higher-layer protocols want to send on port **P**.

  All offering Talker registrations for collision groups on all ports.

- Every collision group on the list that has no asking, ready, or ready failed Listener registrations on port **P** that matches any of streams' stream ID **S,** is removed from the media stream list.

- The remaining collision groups in the media stream list on port P are sorted by the 2-tuple, {rank, talker ID}, from the most important (lowest integer value) to the least important (highest integer value). (Duplicates are eliminated, using the higher of their **Wt** values, and most important rank **R**.)

MSRP: Media load configuration

- For every shared medium port **P**, The collision groups in the media stream list are examined in order, from most-to least-important, and the media load configuration for the media to which port **P** is attached is constructed, starting from a basis of no configured streams.

  If the medium cannot accommodate the collision group, its streams are all deleted from the list. These are the streams that must be removed from the medium, if the quality of service guarantees are to be met.

  Otherwise, if there is a ready or ready failed Listener registration on port **P** for any stream in the collision group, all streams in the group are added to the media load configuration. These are the streams that must be carried on the medium.

  When computing the bandwidth of a stream, a bridge uses the maximum of **Wt**, from the Talker registration, and **Wr**, from the matching Listener registration (if present).

MSRP: Propagating Talker registrations

- The collision groups are examined in order, from highest- to lowest-ranked, and the queue configuration for port **P** is built, starting from a basis of no reserved streams.
  - If not all of the streams in the collision group can be added successfully to both the queue configuration built so far, and to the media load configuration, a failed declaration is sent for each of the streams on port **P**, and the streams are all disabled for transmission on this port.
  - Otherwise, if there is an asking failed Listener registration on port **P** for every one of the streams in the collision group, a failed declaration is sent for each stream in the group, and all are disabled for transmission. This tells the receiver that the data for this collision group will not be sent from this port, so that the receiver can release any resources assigned to the collision group's streams.
  - Otherwise, if there is a ready or ready failed Listener registration on port **P** for any stream in the collision group, then an offering declaration is sent for each of the streams in the collision group, all of the streams in the collision group are added to the queue configuration and to the media load configuration, and all are enabled for transmission. These are the streams that must be configured on port **P**.
  - Otherwise, an offering declaration is sent for each of the streams in the collision group, but the streams are not added to the queue configuration nor to the medial load configuration, and they are all disabled for transmission. These are the streams that could be configured on port **P**, should a receiver on port **P** indicate its readiness to receive the stream.
  - When computing the bandwidth of a stream, a bridge uses the maximum of **Wt**, from the Talker registration, and **Wr**, from all matching Listener registrations (if any).

Multiple Talkers

MSRP: Multiple Talkers

- The key of a Talker registration consists of both the stream ID and the talker ID.
- This is because MSRP assumes that the Talkers are not coordinated with each other, so that their respective bandwidth requirements are additive.
- Registrations for Talkers on the same stream are propagated individually.
- Multiple Talker registrations, from different Talkers, can match a single Listener registration, all with the same stream ID.
- When computing outgoing Talker declarations on port **P**, the different Talkers' offering registrations' bandwidths Wt are added together and compared to the Listener registration Wr.

  If the Talkers' Wt values add up to more than the Listener's Wr, then queue (or media load) configuration cannot support the least-important Talkers.

- In a group conference mode, where all are sending to all, it is very possible for not all Talkers to be enabled.

Changing stream bandwidth or additional QoS information.

MSRP: Increasing bandwidth

- If a stream’s bandwidth must increase, the Talker station changes the Talker registration’s payload.
- That information is propagated down toward the Listeners, who respond by increasing the bandwidth in their Listener registrations.
- The revised Listener registrations propagate back towards the Talker station, and the queues are reconfigured along the way back to the Talker.
- The Talker cannot send at the higher rate until the Listener registration is at least as large as the bandwidth to be sent.
- This increase may, of course, cause the stream to fail.

MSRP: Other QoS changes

- If a stream's bandwidth must decrease, the Talker station can simply transmit at the lower rate, and not change the registration.
- The Talker station should, however, reduce the bandwidth in its Talker registration. As this information is propagated towards the Listener stations, each AV Bridge along the way will adjust its transmission selection parameters to the new, lower value.
- This can cause other offered flows to be registered successfully.
- Similarly, changes to the delay or hop count are propagated by the AV Bridges, and may cause a stream to change it status, either to or from failed.

Issues

MSRP does not scale as well as MMRP

- Each MSRP Talker must register individually. In MMRP, Listeners register collectively.
- *N* Listeners on a shared medium result in *N* registrations on each port, instead of one registration on each port, because their payloads must be compared, minima computed, etc.
- The “listen before declaring” technique used by MRP for reducing the load of control traffic on a shared medium does not work.
  - One Talker’s registration will never appear twice.
  - Each system making a Listener declaration has a different Declaring system ID.
- These differences seem to be tied to the job to be done, rather than to any deficiency in MSRP, itself.
- Note, however, that the Listener ready and Listener ready failed declarations are sent only towards the source of a Talker offering registration, so are typically multipoint-to-point transmissions. This at least does not make the situation worse.

Tearing down queue chains

Tearing down queue chains: T to L

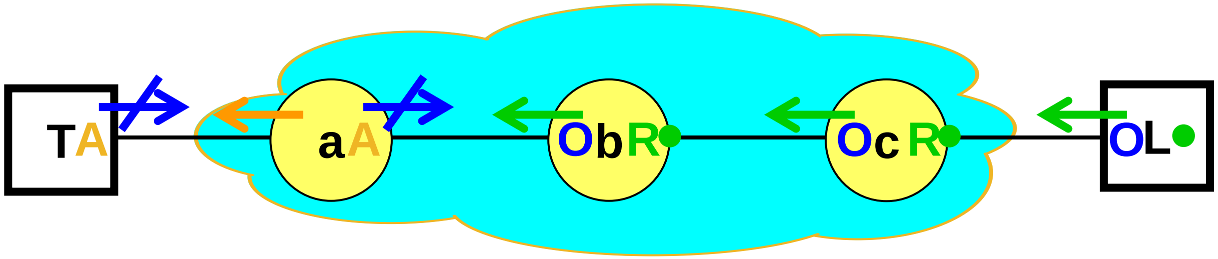

- If the Talker is finished, it stops sending data, and then it can withdraw the offering Talker declaration.
- As the offering registrations disappear, the ready registrations are propagated as asking, instead.
- When the offering registration is removed from the Listener, it withdraws its ready registration, and all the registrations disappear.

Tearing down queue chains: L to T

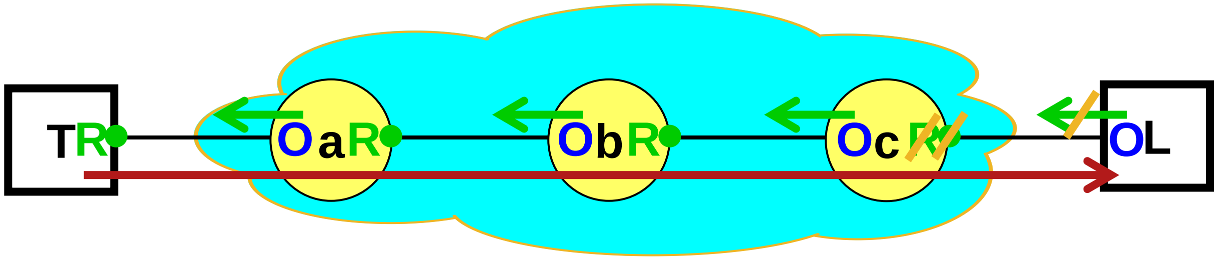

- If the Listener finishes first, perhaps before the stream ends, it withdraws its ready declaration.
- Each bridge, in turn:

  Disallows the flow from leaving the bridge port on which it used to flow;

  Removes the flow from the port configuration; and

  Adjusts all registrations according to the new situation.
- If and when the ready registration is removed from the Talker, it ceases to transmit.

MSRP and 802.11

MSRP in an 802.11 environment

- We will assume that the "Combined Mode" of http://www.ieee802.org/1/files/public/docs2007/avb-nfinn-point-to-multipoint-bridging-061307.pdf is used by the wireless medium, so that the wireless medium behaves like an 802.3 shared medium as far as the control protocols are concerned.
- In infrastructure mode (as assumed by the above reference), a frame sourced by one station that must reach another station must traverse the wireless medium twice – the access point must reflect the frame. A frame sourced by the access point, or that need be received only by the access point, traverses the wireless medium only once.
- Therefore, any stream that moves station-to-station requires twice the bandwidth of a station-to-access point or access point-to-station stream.

MSRP in an 802.11 environment

- Therefore, MSRP must be slightly modified for 802.11 wireless media.
- Both the **Talker** and the **Listener** declarations and registrations must include an additional **payload** term:

  Wireless Station: A Boolean value, true if the system transmitting the declaration is an 802.11 station.

- The access point reflects all MSRP PDUs received from any station back to all of the stations.
- When computing the media stream list (but not the sendable streams list), a system doubles the bandwidth for any stream such that both the Talker and at least one asking, ready, or ready failed Listener registration have Wireless Station set.

Registration propagation charts

Propagating registrations

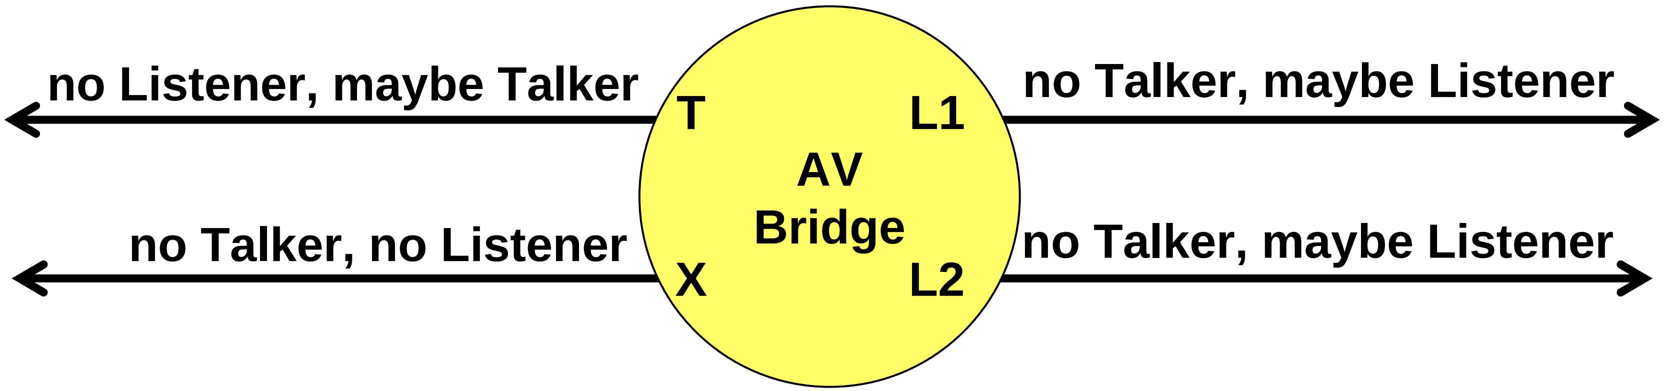

- In all tables below, we have a single AV Bridge with four ports:

  Port T may or may not have a Talker registration for stream S, but has no Listener registration.

  Ports L1 and L2 may or may not each have a Listener registration for stream S, but have no Talker registrations.

  Port X has neither a Talker nor a Listener registration for stream S.

Key

- In all tables, the following symbols are used:

Talker registrations	-	none
	O	offering
	F	failed
Listener registrations	-	none
	A	asking (MSRP, MMRP, GMRP, or IGMP)
	R	ready
	Af	asking failed
	Rf	ready failed

- The 0-bandwidth ready declaration is not included in this table.

Propagating Listener registrations

- Propagating Listener registrations in the absence of any Talker registration from port **T**.
- This table is merely the MMRP/GMRP propagation table, with all registrations treated the same.
- Normally, no ready, ready failed, or asking failed declarations should be received in this case, but this can happen briefly when registrations are changing.

Should this happen?			Y	Y	N	N	N	Y	Y	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Registration from	T	Talker	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	L1	Listener	-	-	-	-	-	A	A	A	A	A	R	R	R	R	R	Af	Af	Af	Af	Af	Rf	Rf	Rf	Rf	Rf
	L2	Listener	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf
Declaration to	T	Listener	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A
	L1	Listener	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A
	L2	Listener	-	-	-	-	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A
	X	Listener	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A

Propagating Listener registrations

- Propagating Listener registrations in the presence of an offering Talker registration from port **T** (Part 1 of 2).

  If ready for stream data, an R (R) is propagated; if not, an A (A).

  If there is a failed path, an f (f) is propagated, else not.

- The Yes/No question ("config.?") is whether or not port **L1** (or **L2**) has been configured for the stream.

Registration from	T	Talker	O	O	O		O	O		O	O	O		O	O		O		O		O				O		O			
	L1	Listener	-	-	-		-	-		A	A	A		A	A		R		R		R				R		R			
		config.?															N	Y	N	Y	N		Y		N	Y	N		Y	
	L2	Listener	-	A	R		Af	Rf		-	A	R		Af	Rf		-		A		R				Af		Rf			
		config.?			N	Y		N	Y			N	Y		N	Y					N	Y	N	Y			N	Y	N	Y
Declaration to	T	Listener	-	A	Af	R	Af	Af	Rf	A	A	Af	R	Af	Af	Rf	Af	R	Af	R	Af	Rf	Rf	R	Af	Rf	Af	Rf	Rf	Rf
	L1	Listener	-	A	A	A	A	A	A	-	A	A	A	A	A	A	-	-	A	A	A	A	A	A	A	A	A	A	A	A
	L2	Listener	-	-	-	-	-	-	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A
	X	Listener	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A

Propagating Listener registrations

- Propagating Listener registrations in the presence of an offering Talker registration from port T (Part 2 of 2).

  If ready for stream data, an R is propagated; if not, an A.

  If there is a failed path, an f (f) is propagated, else not.

- The Yes/No question (“config.?”) is whether or not port **L1** (or **L2**) has been configured for the stream.

Registration from	T	Talker	O	O	O		O	O		O		O		O				O		O			
	L1	Listener	Af	Af	Af		Af	Af		Rf		Rf		Rf				Rf		Rf			
		config.?								N	Y	N	Y	N		Y		N	Y	N		Y	
	L2	Listener	-	A	R		Af	Rf		-		A		R				Af		Rf			
		config.?			N	Y		N	Y					N	Y	N	Y			N	Y	N	Y
Declaraition to	T	Listener	Af	Af	Af	Rf	Af	Af	Rf	Af	Rf	Af	Rf	Af	Rf	Rf	Rf	Af	Rf	Af	Rf	Rf	Rf
	L1	Listener	-	A	A	A	A	A	A	-	-	A	A	A	A	A	A	A	A	A	A	A	A
	L2	Listener	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A
	X	Listener	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A

Propagating Listener registrations

- Propagating Listener registrations in the presence of a failed Talker registration from port T.

  If there is a failed path, an f (f) is propagated, else not.

Registration from	T	Talker	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	L1	Listener	-	-	-	-	-	A	A	A	A	A	R	R	R	R	R	Af	Af	Af	Af	Af	Rf	Rf	Rf	Rf	Rf
	L2	Listener	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf	-	A	R	Af	Rf
Declaration to	T	Listener	-	A	Af	Af	Af	A	A	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af	Af
	L1	Listener	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A	-	A	A	A	A
	L2	Listener	-	-	-	-	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A
	X	Listener	-	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A

Propagating Talker registrations

- Propagating Talker registrations.
- Note that the second row of the table is an input, and is the result of the Listener registration propagated to port **T** from the previous tables. The next row is the registration received from port **T**.
- The Yes/No question ("Sufficient BW") is the result of the Talker propagation computation.

		Should this happen?	Y	Y	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Declaration to (from above)	T	Listener (A, R, Rf, Af)	any (or none)										none, A R or Rf (any except Af)										Af
Registration from	T	Talker (O,F)	-	-	-	-	-	F	F	F	F	F	O	O	O		O		O		O		O
	L	Listener (A, R, Rf, Af)	-	A	R	Af	Rf	-	A	R	Af	Rf	-		A		R		Af		Rf		any
		Sufficient BW available to pass this stream?											N	Y	N	Y	N	Y	N	Y	N	Y	
Declaration to	L	Talker (O,F)	-	-	-	-	-	F	F	F	F	F	F	O	F	O	F	O	F	O	F	O	F

Frame format

MPRP protocol definitions

- MPRP frames are based on the format of IEEE Std. 802.1ak-2007 MRP.
- MSRP is an MRP application and an MPRP application.
- The semantics of the term, “Attribute,” as used in 802.1ak-2007, must change for MPRP.
  - Whereas MRP registers Attributes, MPRP registers {Key, Payload} pairs.
  - The FirstValue field of 802.1ak-2007 clause 10.8.1.2 includes a {Key, Payload} pair, instead of just an Attribute.
  - Only the Key part of the pair is used when comparing values.
  - This will require more-or-less extensive editing to 802.1ak.

MSRP frame format

- A new 802 Ethertype, and a new 802.1ak MRP AttributeType and ProtocolVersion must be allocated for MSRP.
- MSRP has no need for the NumberOfValues field to have any value other than 1, since the key parts of an MSRP declaration are not simple integers. It is for decision by the 802.1 Working Group whether the NumberOfValues field could be eliminated, or made optional.
- The FirstValue defined in clause 10.8.1.2 of 802.1ak-2007 is defined by MSRP (as is the case for all MRP applications).

MPRP frame format: FirstValue

- The new syntax of an MPRP FirstValue:

  FirstValue ::= KeyLength, Key, PayloadLength, Payload

  KeyLength BYTE ::= *Length of Key field, in bytes*

  Key ::= *Defined by the specific MPRP application*

  PayloadLength BYTE ::= *Length of Payload field, in bytes*

  Payload ::= *Defined by the specific MPRP application*

MSRP frame format

- The syntax for MSRP:

  ProtocolVersion ::= *To be determined by P802.1Qat, e.g., 0*

  AttributeType ::= *To be determined by P802.1Qat, e.g., 12*

  Key ::= DirectionHsLength [, HigherStreamId], StreamId, SystemId

  DirectionHsLength BYTE ::= (RegType * 128) + *Length in bytes of HigherStreamId*

  RegType ::= Listener | Talker

  Listener ::= 0

  Talker ::= 1

  HigherStreamId ::= Oui, HsiSubtype, HsiValue

  Oui ::= *An Organizationally Unique Identifier (OUI) of the organization defining the HsiSubtype and HsiValue*

  HsiSubtype BYTE ::= *Defined by the organization owning the OUI, specifies the format and meaning of the HsiValue. Must be present, so that if the meaning of an HsiValue changes, a new OUI will not be required to distinguish it.*

  HsiValue ::= *Defined by the organization owning the OUI. The definition of the use of an HsiSubtype and HsiValue must be such that a system can perform a valid binary comparison of two HigherStreamId values for equality, without understanding the internal structure of the HsiValue.*

MSRP frame format

- The syntax for MSRP (continued):

  StreamId ::= LocalStreamId, StreamVlanId, StreamMacAddress

  LocalStreamId SHORT ::= *Integer value assigned by the Talker station to identify different streams sharing the same MAC address*

  StreamVlanId SHORT ::= *The VLAN ID to be used by data frames belonging to the reserved stream. 0, if the sender is not VLAN-aware, in which case the first VLAN-aware bridge to receive the registration supplies the VLAN ID. This field is ignored by a non-VLAN-aware receiver.*

  StreamMacAddress ::= *The destination MAC address to be used by data frames belonging to the reserved stream*

  SystemId ::= LocalSystemId, SystemMacAddress

  LocalSystemId SHORT ::= *Integer value assigned by the originator of the registration to identify different systems sharing the same MAC address*

  SystemMacAddress ::= *A global MAC address that, along with the LocalSystemId, uniquely identifies either the Talker (for a Talker registration) or the system transmitting the Listener registration*

MSRP frame format

- The syntax for MSRP (continued):

  Payload ::= ListenerPayload | TalkerPayload

  ListenerPayload ::= ListenerFlags, Bandwidth

  ListenerFlags BYTE ::= ((WirelessStation * 16) + ListenerType)

  WirelessStation ::= False | True

  ListenerType ::= Asking | AskingFailed | Ready | ReceivingFailed

  Asking ::= 0

  AskingFailed ::= 1

  Ready ::= 2

  ReceivingFailed ::= 3

  Bandwidth ::= BytesPerSecond1k, AvgPayloadSize

  BytesPerSecond1k LONG ::= (((*Number of bytes per second required by the stream*) + 500) / 1000)

  AvgPayloadSize SHORT ::= *The worst case for the average number of bytes in one frame of the stream measured over an interval to be determined by the 802.1 Working Group, e.g. 1 millisecond.*

MSRP frame format

- The syntax for MSRP (continued):

  TalkerPayload ::= TalkerFlags, Bandwidth, OtherQosInfo, Rank

  TalkerFlags BYTE ::= ((StreamPriority * 32) + (WirelessStation * 16) + (MacOnly * 8) + TalkerType)

  StreamPriority [0..7] ::= *Value of the priority to be used for the data stream frames*

  MacOnly ::= False | True

  TalkerType ::= Offering | Failed

  Offering ::= 0

  Failed ::= 1

  OtherQosInfo ::= *To be determined by the 802.1 Working Group, contents as described in an earlier slide, e.g. "OtherQosInfo BYTE ::= hop count"*

  Rank ::= GivenRank, TimeOfDeclaration

  GivenRank SHORT ::= *An unsigned value determined by higher layers, presumably other standards bodies than IEEE 802.1. Default value is 2^15.*

  TimeOfDeclaration LONGLONG ::= *An 8-byte time stamp in IEEE 1588 / IEEE P802.1AS format. 0 indicates undetermined.*