AVnu Testing Consortium

FQTSS Overview

What is FQTSS?

FQTSS is defined in the IEEE standard 801.Q Clause 34; Forwarding and Queuing Enhancements for Time-Sensitive Streams. The term FQTSS is used to describe a set of tools which are used to forward and queue time-sensitive streams. Since AVB frames cannot be dropped, there must be a mechanism in place to forward AVB frames quickly and efficiently. This is where FQTSS (aka Qav) comes into play.

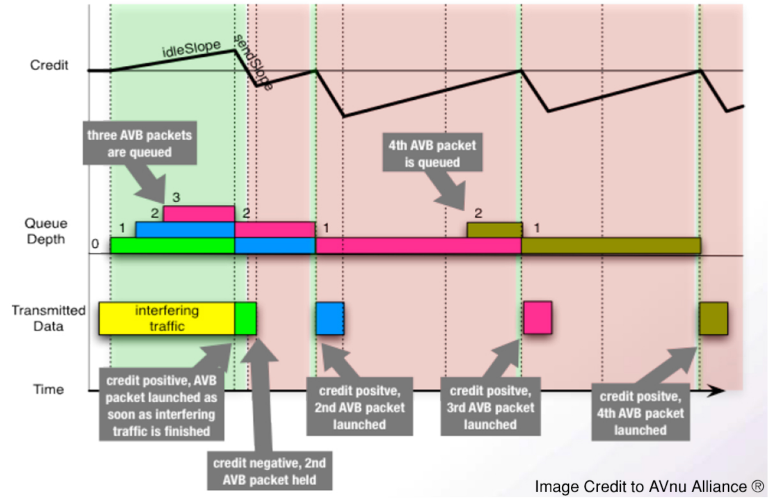

Image Credit to AVnu Alliance ®

The Components of FQTSS

FQTSS is made up of of six parts which each adress a specific need of AVB traffic in bridges:

1) A means of detecting the boundary between a set of bridges that support SRP (aka an SRP domain) and also with surrounding bridges that do not support SRP.
2) A set of bandwidth availability parameters for each port that are used to record the maximum bandwidth available to a given outbound queue, and the actual bandwidth reserved for that queue.
3) A credit-based shaper algorithm that is used to shape the transmission of stream-based traffic in accordance with the bandwidth that has been reserved on a given outbound queue.
4) Rules governing the relationship between the size of the layer 2 “payload” (the MSDU) carried in a frame and how that relates to the atual bandwidth that will be consumed when that MSDU is transmitted on a particular port.
5) An algorithm for determining the mapping of the priorities associated with received frames onto the traffic classes available on the transmission ports of a bridge.
6) A definition of the required behavior of an end station that acts as the source of a time-sensitive stream.

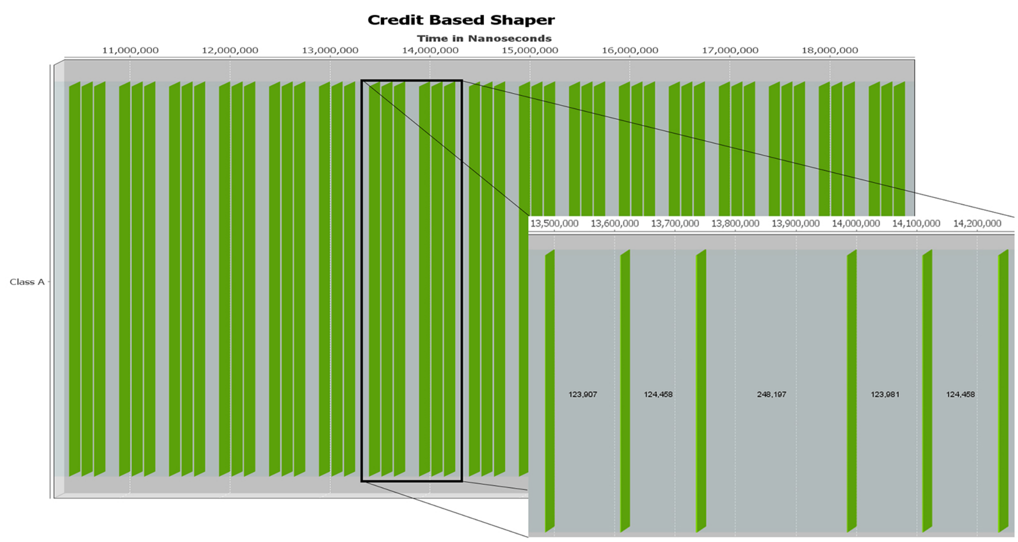

AVnu Testing Consortium
FQTSS Overview

1) Detection of SRP domains

Bridges detect the edge of an SRP domain by observing SRP behavior. If a bridge recieves SRP registrations using a particular priority, then it is reasonable to believe that they are being recieved from an SRP capable device; the SRP engine can therfore signal which ports of a bridge are at the boundary of an SRP domain.

2) Bandwidth availabity paramaters

There are four basic paramaters that exist for each port and for each traffic class which supports the credit-based shaper algorithm.

1) *portTransmitRate* - The transmission rate, in bits per second, that the underlying MAC service that supports transmission through the port provides.

2) *deltaBandwidth* - A percentage of *portTransmitRate* that can be reserved for use by the queue associated with a specific traffic class. For a given traffic class, the total bandwidth that can be reserved is the sum of the *deltaBandwidth* values for that traffic class and all higher traffic classes, minus any bandwidth reserved by higher traffic classes that support the credit-based shaper algorithm.

3) *adminIdleSlope* - The bandwidth, in bits per second, that has been requested by management to be reserved for use by the queue associated with a given traffic class. If SRP is in operation, this parameter has no effect; if SRP is not in operation, then the value of *operIdleSlope* is always equal to the value of *adminIdleSlope*.

4) *operIdleSlope* - The actual bandwidth, in bits per second, that is currently reserved for use by the queue asociated with a given traffic class. This value is used by the credit-based shaper algorithm as the *idleSlope* for the corresponding queue.

3) Relationships among availabity paramaters

The recommended default value of *deltaBandwidth* for the highest numbered traffic class supported is 75%, and for any lower numbered traffic classes, the recommended default value is 0%. The *deltaBandwidth* for a given traffic class, plus the *deltaBandwidth* values for any higher priority queues defines the total percentage of the port's bandwidth that can be reserved for that queue and all higher priority queues.

For the highest priority queue, this means that the maximum value of *operIdleSlope* is *deltaBandwidth%* of *portTransmitRate*. However, if *operIdleSlope* is actually less than this maximum value, any lower priority queue that supports the credit-based shaper algorithm can make use of the reservable bandwidth that is unused by the higher priority queue.

4) Bandwidth availability parameter management

If the stream reservation mechanisms defined in SRP are supported, then the values of *operIdleSlope* are determined solely by the operation of SRP. If SRP is not supported, then the values of *operIdleSlope* are equal to the values requested by management in the corresponding *adminIdleSlope*.

It is possible for the value of *portTransmitRate* for a port to change as a result of the normal operation of the underlying MAC service; it is also possible for managment actions to change the values of *deltaBandwidth* for a port. In either case, the consequence could be one of the following:

1) The sum of the *operIdleSlope* values for the port could now exceed the total reservable bandwidth allowed for the port, or the *operIdleSlope* value for a given queue could now exceed the reservable bandwidth allowed for the queue. Consequently, there could be streams currently active on the port that can no longer be supported.

4) Bandwidth availability parameter management cont.

2) The bandwidth now available to a given queue could mean that there are streams that are currently inactive that could be supported on the port.

3) Active streams that continue to be supported after the change could se their latency guarante change.

In either case corrective action either by management or by the stream reservation mechanisms defined in SRP are required to restore the parameters to a consistent set of values.

5) Deriving actual bandwidth requirements from the size of the MSDU

The forwarding and queuing mechanisms are defined in terms of the actual bandwidth used when frames are transmitted on the medium that supports the MAC service available through the port. In contrast, the SRP makes use of a traffic specification (TSpec) for each stream that defines the maximum number of bits per frame (*maxFrameSize*), and a maximum frame rate (*maxIntervalFrames*) for that stream. The TSpec takes no account of the per-frame overhead asociated with transmitting the MSDU over a given medium.

However, when SRP determines the value to be used for the *operIdleSlope* associated with a given queue, it is necesary for this value to include the per-frame overhead that will be incurred when frames are transmitted on that port.

It is possible to determine the overhead that is added to the per-frame MSDU payload when a frame is transmitted. There are at least the following sources of per-frame overhead.

1) Any VLAN/secuirty tags that are added to the layer 2 payload as it passes through the protocol stack.

2) The MAC framing athat is added by the underlying MAC service.

3) Any physical layer overhead, such as preamble characters and inter-frame gaps.

The actual bandwidth needed to support a given stream is therefore defined as follows:

actualBandwidth = (perFrameOverhead + assumedPayloadSize) x maxFrameRate

6) Mapping priorities to traffic classes

The default mappings of priorities to traffic classes meet the following constraints:

1) Priority values that correspond to SR classes are mapped onto traffic classes that support the credit-based shaper algorithm as the transmission selection algorithm.

2) Traffic classes that support the credit-based shaper algorithm have a higher priority than traffic classes that support the strict priority (or any other) transmission selection algorithm.

3) At least one traffic class supports credit-based shaper algorithm, and at least one traffic class supports the strict priority transmission selection algorithm.

7) End station behavior

In order for an end station to successfully participate in the transmission and reception of time-sensitive streams, it is necessary for their behavior to be compatible with the operation of the fowarding and queueing mechanisms employed in bridges. The requirements for end stations that participate as “talkers” are different from the requirements that apply to “listeners”.

AVnu Testing Consortium
FQTSS Overview

7A) Talker Behavior

It is a requirement for a talker to use the priorities that the bridges in the network recognize as being associated with SR clases exclusively for transmiting stream data. It is also necessary for the talker, and the bridges in the path to the listener(s), to have a common view of the bandwidth required in order to transmit the talker's streams, and for that bandwidth to be reserved along the path to the listener(s). This latter requirement can be met by means of stream reservation mechanisms, such as defined in SRP, or by other management means.

Talkers shall exhibit transmission behavior for frames that are part of time-sensitive streams that is consistent with the operation of the credit-bsed shaper algorithm, both in terms of the way they transmit frames that are part of an individual data stream, and in terms of the way they transmit stream data frames from a port. The queueing model for a talker port, and for a given priority, can be considered to look like the figure below.

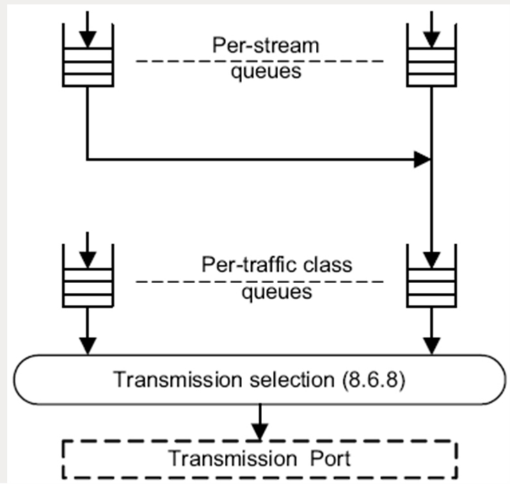

The talker places frames into the queue associated with an individual stream based on the TSpec for that stream. The queue associated with each individual stream uses the credit-based shaper algorithm as the means of determining the rate at which data frames for that stream are placed in the outbound queue for the priority that the stream is using.

The outbound queue for that priority makes use of the credit-based shaper algorithm as the means of determining the rate at which data frames for that stream are selected for transmission.

From this point of view, a talker can be thought of as a single-port bridge.

7B) Listener Behavior

The primary requirement for a listener station is that is is capable of buffering the amount of data that could be transmitted for a stream during a time period equivalent to the accumulated maximum jitter that could be experienced by stream data frames in transmission between talker and listener.

From the point of view of the specification of the forwarding and queueing requirements for time-sensitive streams, it is assumed that the listener will assess the buffering required for a stream as part of the stream bandwidth reservation mechanisms employed by the implementation.