

Audio Video Bridging Gen 2 Assumptions

IEEE 802.1 AVB Plenary

July 2011 – San Francisco

Green Text = Agreed to at a Plenary (was Blue or Red)

Blue Text = Newly Agreed to (was Red at last Face 2 Face)

Black Text = Not Decided

Changes Marked with Red from last version

Don Pannell

Marvell Semiconductor

dpannell@marvell.com

Revision History

- **avb-pannell-gen2-assumptions-0511-v5: Work done in San Francisco**
- avb-pannell-gen2-assumptions-0311-v4: Work done in Singapore
- avb-pannell-gen2-assumptions-0111-v3: Work done in Kauai – not finished
- avb-pannell-gen2-assumptions-1110-v2: Work done in Dallas
- avb-pannell-gen2-assumptions-0910-v1: 1st grouping of all STDs – stolen from below
- at-cgunther-srp-rev2-assumptions: First draft presented July 2010, San Diego, CA

Overview

This document is a collection of concepts and ideas for *possible* inclusion in the next versions of SRP (802.1Qat) and/or the AVB Shaper (802.1Qav) or some new standard.

It should not be considered as a Work Item list yet. Each item needs contributions (i.e., presentations) before it can be agreed to and considered an item to be added to a draft standard. These presentations are needed immediately.

Automotive Needs (July 2011)

- Max Latency: 100 uSec w/4 FE hops for Control Frames
- Other Automotive Needs:
 - Max cable hop count: 8
 - Max number of nodes (bridges & end stations): 32
 - Max cable length: 24m
 - Max end to end cable length: 30m
- Control data attributes
 - Max data size (payload/Layer 2 Data size): 256 bytes
 - Max number of simultaneous transmissions: 8 to 32
 - Transmission period: 500 uSec
- Payload (Layer 2 Data) size for other traffic:
 - 256 to 1500 bytes

Industrial Needs (July 2011)

- Max Latency:
 - Interfering Frames (includes other same PCP frames) + Bridge Latency (not including Store Forward Latency) < 3 uSec / hop
- Other Needs:
 - Fixed Transmission Periods – 31.25u Sec to 1 mSec
 - Max 50% of Period for Low Latency Transmissions
 - At most 64 hops in a daisy chain
 - At most 512 devices off one controller
 - +/- 1uSec time sync between all nodes
 - At most 4096 streams
 - 10 to 300 byte control frame size

Consumer Needs

- Max Latency:
- Other Needs:
- The maximum time to make or break an SRP reservation in the absence of a topology change is:
 - This goal is defined per hop assuming a max of 7 hops
 - For consumer remote control applications this must not exceed 100 mSec?

Professional Needs

- Max Latency:
 -
- Other Needs:
- The maximum time to make or break an SRP reservation in the absence of a topology change is:
 - This goal is defined per hop assuming a max of 7 hops
 - For professional video applications this must not exceed 20 mSec?

Terms (July 2011)

- Pre-emption –
 - Pre-emptable –
 - Framelet – a portion of a Frame created by Pre-emption
 - Unassemble, Fission – the action of making Framelets
 - Assemble – the action of making Frames from Framelets
 - Suspend and Resume – what the queue does with the Pre-empted frame
-

Redundancy (July 2011)

- Reservation Ready – Discovery of all possible paths to a Listener and flows will propagate out all Bridge ports until a Blocked port is reached. Link Cost & Stream Reference Count can then be used to limit the discovered paths to two (MGMT can limit the available paths further).
- Seamless Redundancy – Where a Listener receives more than one copy of a stream on more than one port and it can select which one to use in real time.
- Reconfigure time =
 - Gen 1 = $T_{rec_routing}$ (RSTP time) + T_{rec_SRP} (SRP time)
 - With Reservation Ready = $T_{rec_routing}$ (RSTP time)
 - With Seamless Redundancy = 0

SRP Generation 2 Ideas

SRP – Possible New Work

- Dynamic bandwidth reservations (modify ‘on the fly’)
 - Done by requesting the same Stream ID with a new T-Spec?
 - Will not consider unless a contribution is supplied (Jan 2011)
- Variable bit rate reservations (statistical averaging)
 - Currently a video stream must reserve the max it will use
 - Still want to be able to Guarantee all streams are delivered (by sneaking into the 25%, which is the non-AVB bandwidth?)
 - Will not consider unless a contribution is supplied (Jan 2011)
- Dynamic changes to latency
 - Due to redundancy
 - Due to MGMT reconfiguration of a bridge
 - Change in Fan-in
 - Class % allocated
 - Due to Multiple Talkers – due to multi-Talker to one Listener
- Add the ability to get current worst case latency assuming no new reservations
 - Report Max size interfering frame that is smaller than 1522 if that is all a Talker node needs to Tx
 - Lock down the configuration by reporting SRP failed for any new requests?

SRP - Possible New Work

- Configurable Max Latency parameter that can prevent a reservation
 - Add support for a new lowest latency Class (i.e, 2 or 3 Classes in one LAN)
 - On a per port and/or per bridge basis
- Add a Tear Down Rank bit?
 - So a newer stream can stay when bandwidth is needed elsewhere?
 - Need to consider comments received from previous Qat ballots (Mar 2011)
- Be able to create or pre-configure a reservation via MGMT/Flash
 - For quick boot up or setup via management objects
- Two-way reservations – decided no need to support
 - Must be handled at a higher layer
- Explicit path reservation – like Talker Advertise pruning to save network & CPU bandwidth by reducing flooding?
 - Needed for Redundancy?
 - Intent is to make things as simple as possible but built on SRP
 - Advertise Pruning on receipt Listener Ready

SRP - Possible New Work

- Link aggregation
 - With and without redundancy
- Redundancy
 - Spanning the range from no single points of failure to up to two completely independent paths with copied data
<http://www.ieee802.org/1/files/public/docs2010/at-kleineberg-goetz-AVB-redundancy-1110.pdf>
 - The redundant path may be statistically over subscribed
 - Protocol neutral interface to layer 2 redundancy mechanisms
 - Need to be able to determine stream recovery times & decision metrics
 - Need to be in sync with RSTP, MSRP or allow streams to egress Blocked ports?
 - **Need to look at Shortest Path Bridging & ECMP?**
- Energy Efficient Ethernet
 - Remove MMRP/MVRP periodic timers on EEE links (or all the links)
 - Rest may be solved in 802.1BA
- Unicast address Stream destination address
 - What is the real problem here? Makes Policing harder.
 - Streaming HTTP on top of TCP use an address passed to it by DNS
 - Can be detected and fixed? (i.e., make it a multicast on the AVB LAN)

SRP - Possible New Work

- Multiple Talkers per Stream (one streaming at a time)
 - Networked video switcher
 - Switch on a bit in a stream or switch at a specific time?
 - Or do the Talkers to all the turning on or off (i.e., the MUX'ing)?
 - Need the concept of a Group Reservation
- Multiple Talkers per Stream (time sliced approach)
 - Industrial control
 - <http://www.ieee802.org/1/files/public/docs2010/at-goetz-AVB-lowlatency-part1-0510.pdf>
- More SR Classes - Yes, new 'named' performances need to be defined
 - Some applications need better than 2mSec over 7 hops of FE
- Support More PCPs? Hopefully no more than 2 SR Classes at a time are needed.
 - Allow 'moving' specific applications to specific performance levels?
 - i.e., have many SR Classes but by default only two PCP's can be in use at one time?
 - Does any environment need more than 2 PCP's?
- Configurable SR class priorities and VIDs
 - Need service primitives (e.g. REGISTER_DOMAIN.request/indication), management (clause 12), or SNMP (clause 17) to do this currently

Other Q Enhancements

- Gateway between conflicting SR Class domains – already solved as Qat requires that both SR Class and PCP match
- Reduce Latency
 - Bursting concerns
 - Configure the characteristics of each Class's Qav Shaper
 - Preemption
 - Other options
- Automatic Talker pruning? – this is not easy
 - Should be addressed by 'Explicit path reservation'
- Need to fix SRP to support Ingress VLAN membership checking so a Talker needs to issue an MVRP join request to the VLAN it is using for flows.

Other Q Enhancements

- Enhance MRP to use difference-based updates rather than complete database updates (reduces bridge CPU overhead and control bus bandwidth usage)
 - Another goal is to support a larger attribute set
 - May have periodic updates of a portion of the database
 - Or don't do updates at all – only do Register / Deregister?

SRP - Other Ideas - 1

- Cloud diagnostics (devices along the path)
 - Perhaps 802.1ag? **Need a presentation here!**
- **Ingress policing/monitoring**
 - Someone's talking when they shouldn't be
 - Talking without a reservation
 - The stream's DA is not known in the filtering database
 - The frame's PCP is AVB to a unicast
 - Talking too much for the amount reserved
 - Exceeding the reservation
 - Is this perfect policing or best effort?
 - Must it stop a Denial of Service attack?
- 802.1AE (MACsec) environments?
 - Any AVB Streams and PTP & SRP frames can be AE Tagged?
 - Clean up the interface between the link and SRP?
 - What is this???
- PONs are currently not specifically supported
 - i.e., PON support is dependent on contributions from those that need it

SRP - Other Ideas - 2

- How will MSTP select an SRP path over a CM (Congestion Management) path or a non-SRP/non-CM path using 'out of the box' defaults?
 - For AVB with non-AVB devices: Use MSTP with at minimum one spanning tree instance and set AVB to AVB path costs low (match terms in capability vectors) and playing with root costs using MSTP's priority vector?
 - MJT will form an interest group to resolve.
- SRP for Layer 3? IETF issue?
 - Need an Internet Draft how RSVP can use SRP (Subnet Bandwidth Manager – SBM)

SRP - Other New Notes

- As per the discussion with 802.11aa (on March 17 in Singapore) they are targeted to be published in May 2012. We need to replace the duplicate Q Annex C diagrams and discussions relating to 802.11 with appropriate references and text.
- Preemption to be multi level (July 2011)? Two at most?
- Large (i.e, Jumbo) frames can be supported with AVB flows with Preemption by inference (July 2011).
- Max need of Two “Latency Controlled” Classes for automotive. Industrial? Others?

SRP - Other New Notes

- Introduce an SRP Class Z? that uses a non-AVB PCP but uses a high (6?) PCP for best effort flows. Needed for flows where the max bandwidth of the flow is not known (i.e., it doesn't have T-Spec). This allows these flows to get gPTP, SRP path selection & Gen 2 redundancy without the delivery guarantee. Target MKT is consumer.

AVB Shaper Generation 2 Ideas

Shaper Ideas

- Improve Latency
 - Configurable Shaper that defaults to the Non-Engineered LAN settings
 - Turn off the shaper in small Engineered LANs
 - Positive Based shaper
 - Time Aware shaper (TADS, TABS & TACS)
 - Ref <http://www.ieee802.org/1/files/public/docs2011/new-pannell-latency-options-0311-v1.pdf>

New Ideas Here

- Policing Goals or Issues (Jan 2011)
 - Someone's talking when they shouldn't be
 - Talking without a reservation
 - The stream's DA is not known in the filtering database
 - The frame's PCP is AVB to a unicast
 - Talking too much for the amount reserved
 - Exceeding the reservation
- SRP creates its own data path tree?

BACKUP SLIDES AND ADDITIONAL INFORMATION

gPTP Generation 2 Ideas

Green text is in the PAR
Black text is not in the PAR

gPTP Possible New Work

- Support for Link Agg (IEEE 802.1AX)
- Security? (need the requirements and level of needed security)
- Support for other media:
 - IEEE 1901 (if changes are needed)
 - WiFi Direct (if changes are needed)
 - Others?
- Alternate Timescales (e.g., transport time zone information)
- Time router (gPTP across a Layer 3 router) – Will not do (Jan 2011)
- Mapping between NTP and AS (applicable to 1588) – Will not do (Jan 2011)
- No One Step support on transmit?
- One Step Tolerant on receive
- Hardware Two Step (immediate follow up) – No spec change needed
- Look at improving performance for long daisy-chained time-aware systems (or long networks) that may be in a large ring
- Look at Faster Grand Master change over
 - Pre select a failover Grand Master so the selection when needed is faster

gPTP Possible New Work

- Redundancy
 - Short reconfiguration w/redundant paths when one path fails
 - Look for the holes/issues in a redundant/failover system?
- Automatic measurement of link delay asymmetry
- Detect buffered repeaters on other than 802.3 copper links
 - Add in a variable latency in the link delay as an enhanced mechanism? Maybe a MAC address discovery?
 - Need an alternate mechanism for long (fiber) links
- Reduce BMCA convergence time/work for large (>64 node) networks &/or when a loop exists
 - See Mick Seaman's work on loop detection – Will not consider unless a contribution is supplied (Jan 2011)
 - Large (64+) node networks force a lot of data examinations
- How to assess the synchronization performance of a node
 - For certification – Will not do (Jan 2011)
- Create an Annex to show Grand Master Re-election time