

Audio Video Bridging (AVB) Assumptions

IEEE 802.1 AVB Plenary

Sept 4, 2007 – Stockholm

Green Text = Agreed to on Various AVB Calls/Meetings

Black Text = Not Decided

Changes Marked with Red from last version

Don Pannell

Marvell

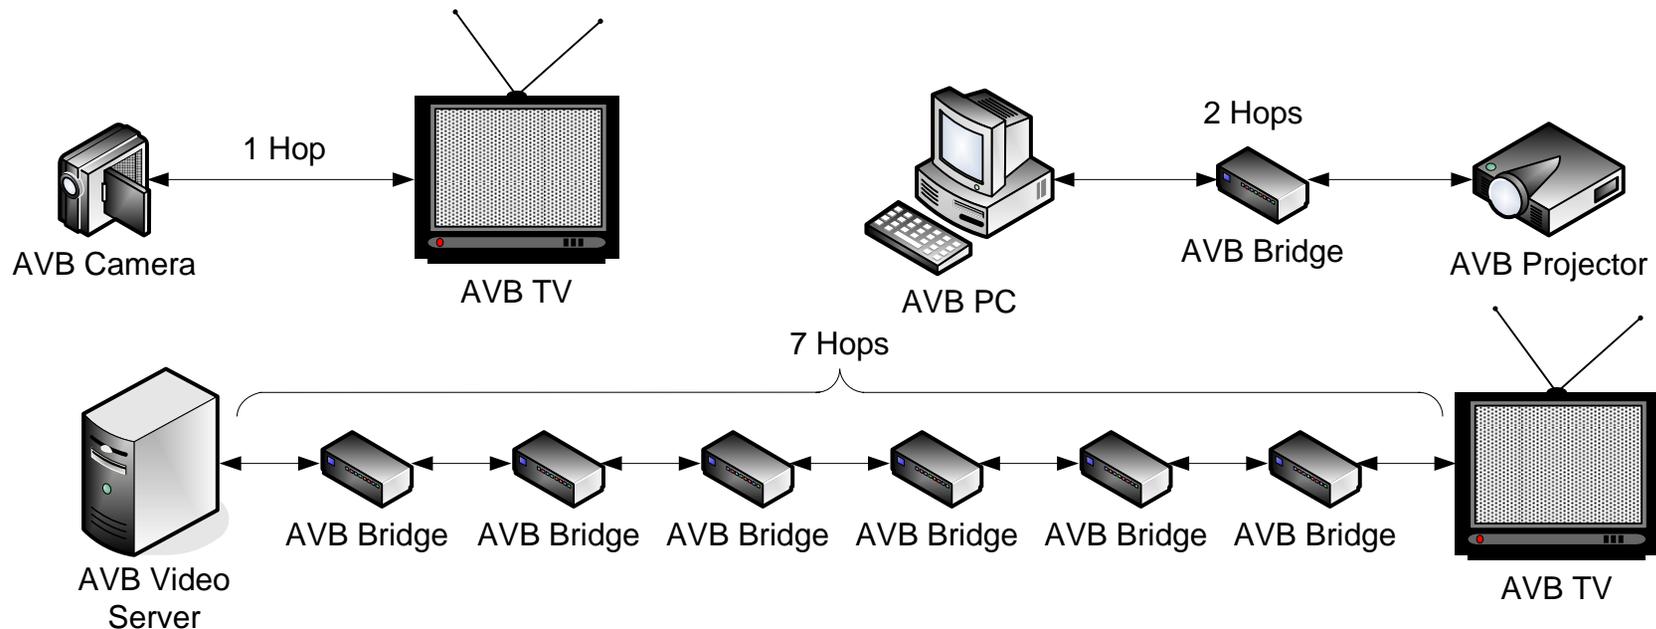
dpannell@marvell.com

Revision History

- Avb-pannell-assumptions-0907-v8: After many calls – 9/04/07 Stockholm
- Avb-pannell-assumptions-0707-v7: After many calls – 7/17/07 San Francisco
- Avb-pannell-assumptions-0607-v6: After Geneva Interim & after 5/23/07 call
- Avb-pannell-assumptions-0507-v5: After many calls
- Avb-pannell-assumptions-0407-v4: After 4/11/07 call
- Avb-pannell-assumptions-0407-v3: After 4/04/07 call
- Avb-pannell-assumptions-0307-v2: After 3/28/07 call
- Avb-pannell-assumptions-0307-v1: Before 3/28/07 call

Performance Goals

- Max Latency (802.1Qav – MAC to MAC, not analog source to sync eg., microphone to speaker or hard drive to display)
 - 802.3: AVB Class **A**: Less than 2 mSec over 7 Hops
 - 802.3: AVB Class **B**: Less than 8 or 16 mSec over 7 Hops?
 - Prefer binary numbers
 - 802.11: ??
- Hops are defined in 802.1 terms, i.e., 1 LAN = 1 Hop



Performance Goals / Names

- Latency Variation (Jitter)
 - 802.3: Need to discuss objectives – affects shaper
- Class Observation Interval (may need a new name)
 - AVB Class A is 125 uSec
 - AVB Class B is 1-5 mSec?
- AVB Class A together with AVB Class B cannot use more than 75% of a link's bandwidth (i.e., not to exceed 75%)
 - The Remaining 25% (or more) is used for Legacy (non-AVB) streams
- Functional Device Type Names
 - AVB will use Talkers, Listeners & Bridges
 - Talker is the source of a stream, Listener is a receiver of a stream
 - A Bridge is an 802.1 Bridge
 - Any physical device could be any combination of these

802.3 Link Requirements

- Link Speed
 - 802.3: 100 Mbit/sec or faster (i.e., no 10 Mbit support)
- Link Duplex
 - 802.3: Full Duplex only (i.e., no half duplex support)
 - Works with clause 43 Link Aggregation (data and PTP)?
- Maximum Frame Size
 - Defined as Layer 2 bytes – or start of DA through end of CRC
 - 802.3: 1088 bytes for AVB Class A for 100 Mbit links?
1522 bytes? Or 2000 bytes for Legacy frames?
(i.e., no Jumbo frame support)
- Flow Control
 - 802.3x is not supported and cannot be used on AVB links

AVB Frame Format

- 802.1 Q Tagging
 - All AVB Streams will be Q Tagged
 - All PTP frames (for 802.1AS) will NOT be Q Tagged
 - All SRP frames (for 802.1Qat) will be Q Tagged
- VLANs
 - The VID is a VLAN and not a Stream Identifier
 - Stream Identifiers must be unique per VID
- Ether types
 - The Ether type of a frame is not a Stream Identifier
- Priorities
 - AVB Class **A** Streams will use a **default** Q Tag priority of 5 (PCP)
 - AVB Class **B** Streams will use a **default** Q Tag priority of 4 (PCP)
 - These Q Tag PRI defaults may change – voice your wish!

Stream Identification - SRP

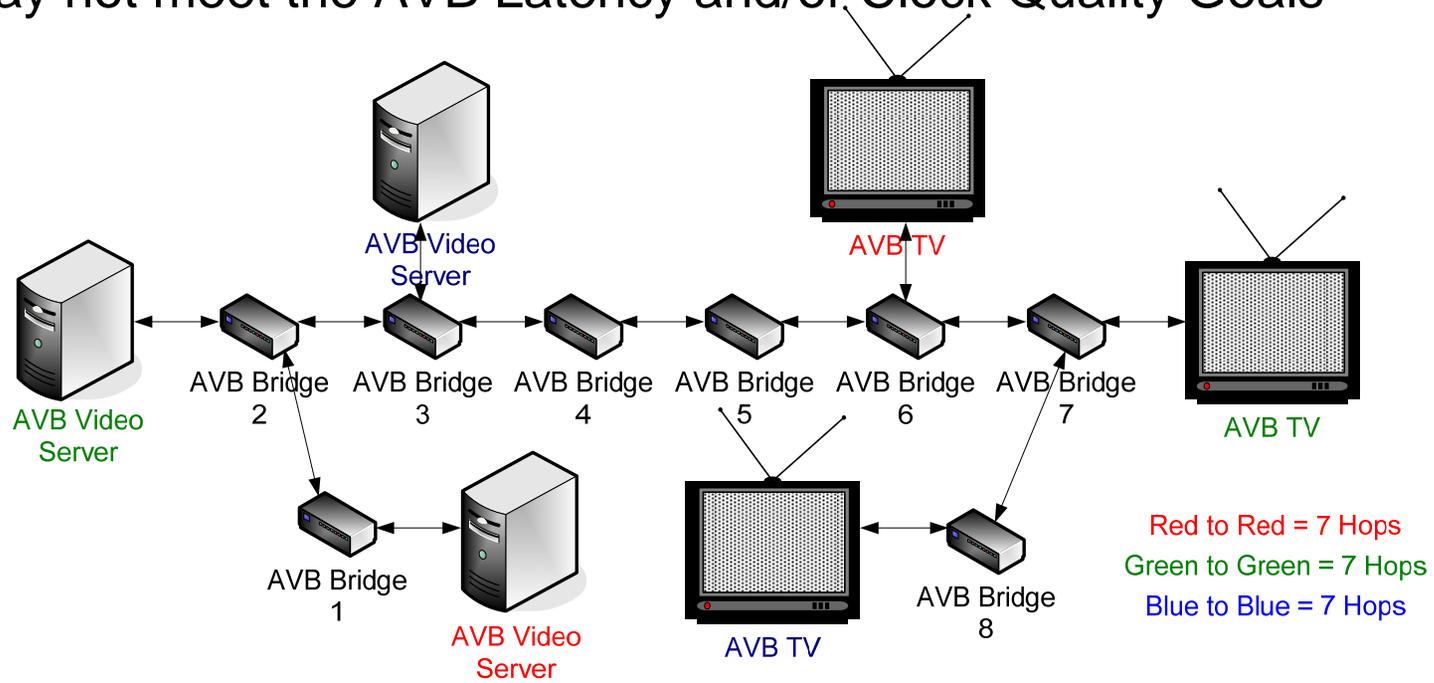
- Stream Identification (for Queue Usage in Bridges – Ref 1)
 - An AVB Stream Frame is Any Frame with a Q Tag priority 4 or 5 (assuming the defaults) entering an AVB port
 - An AVB Port is a port mode used to differentiate the port from being a Legacy port (AVB ports are part of the AVB Cloud, Legacy ports are at the edge of the AVB Cloud, not connected to an AVB aware device)
- Stream Policy (for Queue Usage in Bridges – Ref 1)
 - Only those frames that meet the Stream Identification (above) can be placed into the AVB Egress Queues
 - All other frames are placed into the Legacy Egress Queues
 - Frames that contain a Q Tag priority 4 or 5 entering a non-AVB port (i.e., a Legacy port) must have their Q Tag priority re-mapped to 2 or 3, respectively (i.e., PRI 4 goes to 2, PRI 5 goes to 3)
 - These Q Tag PRI defaults may change – voice your wish!

SRP Assumptions

- Stream Identification (for Mapping in Bridges)
 - Different MC Streams must use Unique MC Addresses (within the same VLAN)
 - Will the standard support Unicast DA for AVB streams?
- Each Stream can use only one PRI
- A Reservation (a DA MAC Address) is for one Stream only
- SRP will currently allocate a Static bandwidth for a stream (i.e., Dynamic, sometimes called Variable Bit Rate, support is dependent on contributions from those that need it and may need to become part of a separate PAR)
- Variable Bit Rate streams need to reserve the peak bandwidth that they need
- AVB Frame Priorities are changeable
 - These are changeable for 'engineered' networks
 - For Home networks the set of specified defaults simply must work!

SRP Assumptions

- The figure below is a valid AVB Cloud
- Therefore, SRP Needs to Count Hops and/or report the Max Latency or Port Speed reported from all the Hops?
 - So a Listener can determine if it wants to listen to the stream as it may not meet the AVB Latency and/or Clock Quality Goals



SRP Questions

- How do we handle the case where a Bridge or Listener is out of Address Resources?
 - Or it can't handle the current address but it can support others
- How are Stream Addresses going to be Allocated?
- How will SRP Respond to changes in the network topology (eg., Spanning Tree change) and/or usable bandwidth (eg., wireless and/or Energy Efficient Ethernet)?
- How will SRP Respond to changes in bandwidth request based on user requests (eg., when the user requests an increase in bandwidth)?
- How will SRP know the link's speed so that 100 Mbit links do not use frame sizes greater than specified?

SRP Questions

- What is the Stream identification for Talkers and Listeners?
- What is the Stream identification for the Reservation protocol?
- How are Changed AVB Priorities communicated through out the AVB Cloud?
- What is an SRP domain (i.e., an AVB Cloud) & its forwarding rules?
 - Assume it to be a subset of an 802.1AS domain
 - Is 802.1AS required if timing is not needed?
 - In the default/non-engineered case 802.1AS is needed to detect 'transparent bridges' however

PTP Assumptions

- PTP Clock Quality (802.1AS clock, Ref 2)
 - +/- 100ppm or better from a free running ≥ 25 MHz clock
 - Less than 4 ppm per Second drift
 - End point time synchronization accuracy for steady-state (up to 7 Hops) ≤ 1 uSec (i.e., any 2 PTP clocks separated by at most 7 Hops differ by no more than 1 uSec)
- Recovered Media Sample Clock Quality
 - Jitter/Wander: per MTIE in Ref 3
 - Endpoint media output synchronization accuracy ≤ 1 uSec
 - Media clock quality maintained during PTP clock discontinuity (e.g. change in Grandmaster) – CH 6/11/07
 - Startup settling time: 2 sec? – CH 6/11/07
- 802.3 PHY Latency Jitter/Wander (Ref 2)
 - No more than 86 ns per Hop or 43 ns per PHY (i.e., portion that is not known and corrected for)
- Spanning Tree (or equivalent) will eliminate routing loops in an AVB cloud – KS 7/17/07

PTP Assumptions

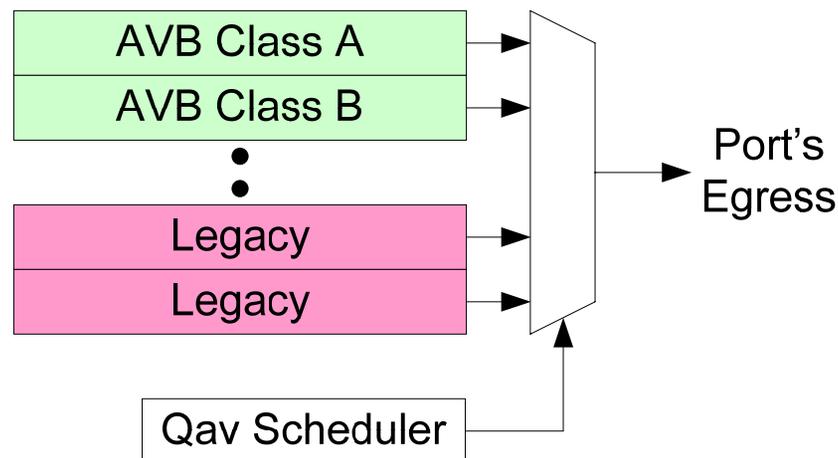
- PTP Clock Discontinuities – CH 6/11/07
 - Slave clocks will see discontinuities in Global Time (i.e. step changes in epoch) during media streaming operation.
 - Discontinuities may originate from selection of new Grandmaster, or may occur within a clock while it is acting as Grandmaster
 - A step change in epoch may be any value (usec to years)
 - The 802.1AS clock service will inform its client upon change in Grandmaster selection.
 - The clock service MAY inform its client of other clock discontinuities.
- An AVB domain (802.1Qat or SRP domain or AVB Cloud) must be contained inside an 802.1AS domain
 - Don't want an AVB Cloud using different time services
- Can't modify 802.1AS frames on-the-fly (i.e., in the egress path)
 - i.e., Implementations must use Follow-up messages

PTP Assumptions

- An 802.1AS domain is a contiguous set of bridges/end-stations that meet the 802.1as requirements. - DVJ 8/8/07
- The 802.1AS protocol packets are not propagated beyond the 802.1AS domain, except possibly for determining the edge of the 802.1AS domain. – DVJ 8/8/07
- An AVB domain is a contiguous set of bridges/end-stations that meet the AVB requirements, which include 802.1AS. – DVJ 8/8/07

PTP/Qav Priority

- Priority Models (do we need to spec this or just spec the latency)?
 - Concern is where to put PTP and/or Management (BPDU type) frames
 - Assume PTP is a MGMT (Management) frame & all MGMT are given the same scheduling priority?
 - At what point do MGMT frames get MUX'ed to the Port's Egress?



PTP Questions

- What is an 802.1AS domain?
 - Time may be larger distributed than an AVB (SRP) Cloud

Qav Assumptions

- Bridges do per class shaping
- Talkers do per stream shaping?
- Policing, if done, is mainly for DoS attacks, and is therefore much less accurate than Shaping since you don't want to discard any streams originating from conforming Talkers
 - This is because the Observation Interval (new name?) between devices is not synchronized
- In the absence of congestion, Talkers will not burst AVB stream frames (i.e., they will transmit 1 frame then an extended IFG to even out the frame spacing)

Qav Questions

- Observation intervals name – token bucket solutions... What are we really doing here?
- **Maybe we should call it a Shaping Interval?**

802.11 Requirements

- Link Speed
 - 802.11: 100 Mbits or faster?
- Link Duplex
 - 802.11: Half Duplex (no choice)
- Maximum Frame Size (Layer 2 bytes – DA through CRC)
 - 802.11: 1088 bytes for AVB Class **A** (depending on the speed of the link)?
1522 bytes? Or 2000 bytes for Legacy frames?
(i.e., no Jumbo frame support)
- Link Level Flow Control
 - 802.11: Won't be allowed

Other Issues

- How will MSTP select an AVB path over a CM (Congestion Management) path or a non-AVB/non-CM path using 'out of the box' defaults?
- How is the edge of the 'stream' cloud determined (i.e., the cloud that uses Qat and Qav)?
- Use 802.1AS to detect 'transparent bridges' and then indicate that the port it is connected to is not 'AVB Capable'?
 - Where does this fit into the building of STP?

Other Issues

- 802.1AE Environments
 - Any AVB Streams and PTP & SRP frames can be AE Tagged
 - Need to understand implications of MacSEC on PTP Timestamping
 - Will Need Help with 802.3 on this
- PONs are currently not specifically supported?
 - i.e., PON support is dependent on contributions from those that need it and may need to become part of a separate PAR
- Provider Networks are currently not specifically supported?
 - i.e., Provider Network support is dependent on contributions from those that need it and may need to become part of a separate PAR
- Other Assumptions ...
 - (this is a growing work in process)

References

- Ref 1: at-pannell-policies-0707-v04
 - Title: 802.1 Qat Policies Proposals
- Ref 2: as-garner-assumptions-for-error-sources-time-synch-0507-v03
 - Title: Assumptions for Sources of Time Synchronization Error in IEEE 802.1AS
- Ref 3: avb-garner-requirements-summary-r4-060217
 - Title: Summary of AVB Bridging Network Requirements