IEEE Standards Update

On-going work in Audio-Video-Bridging, Precise Time Synchronization & Time-Sensitive Networking

Technical Track

www.odva.org

IEEE Standards Update

Brief History of AVB/TSN

Motivation for TSN in Industrial

New Work in IEEE – 802 & 1588

Implications of TSN to EtherNet/IP

IEEE Standards Update

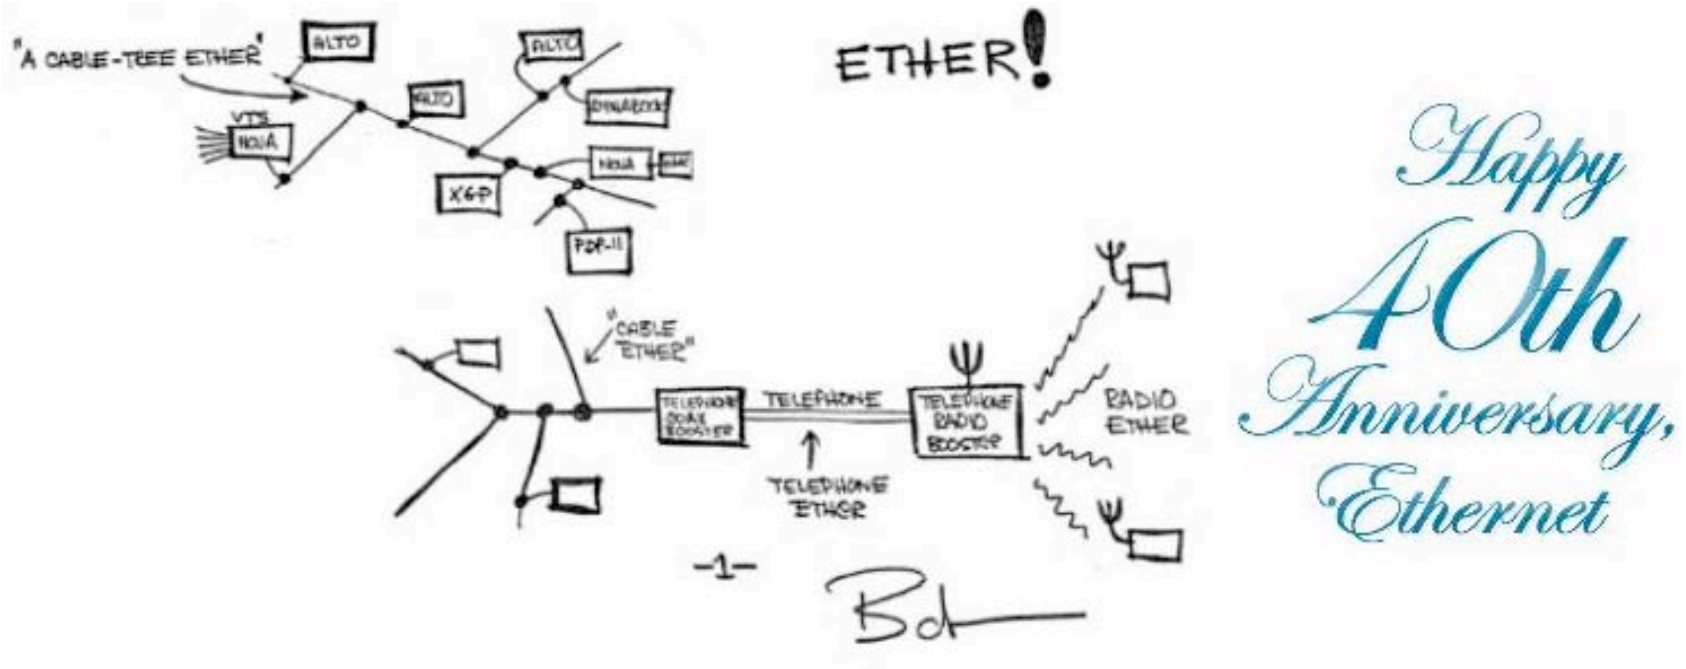

Evolution of Ethernet

10base2/5

Full Duplex, Switched

Convergence via TSN

Foresight of Layering

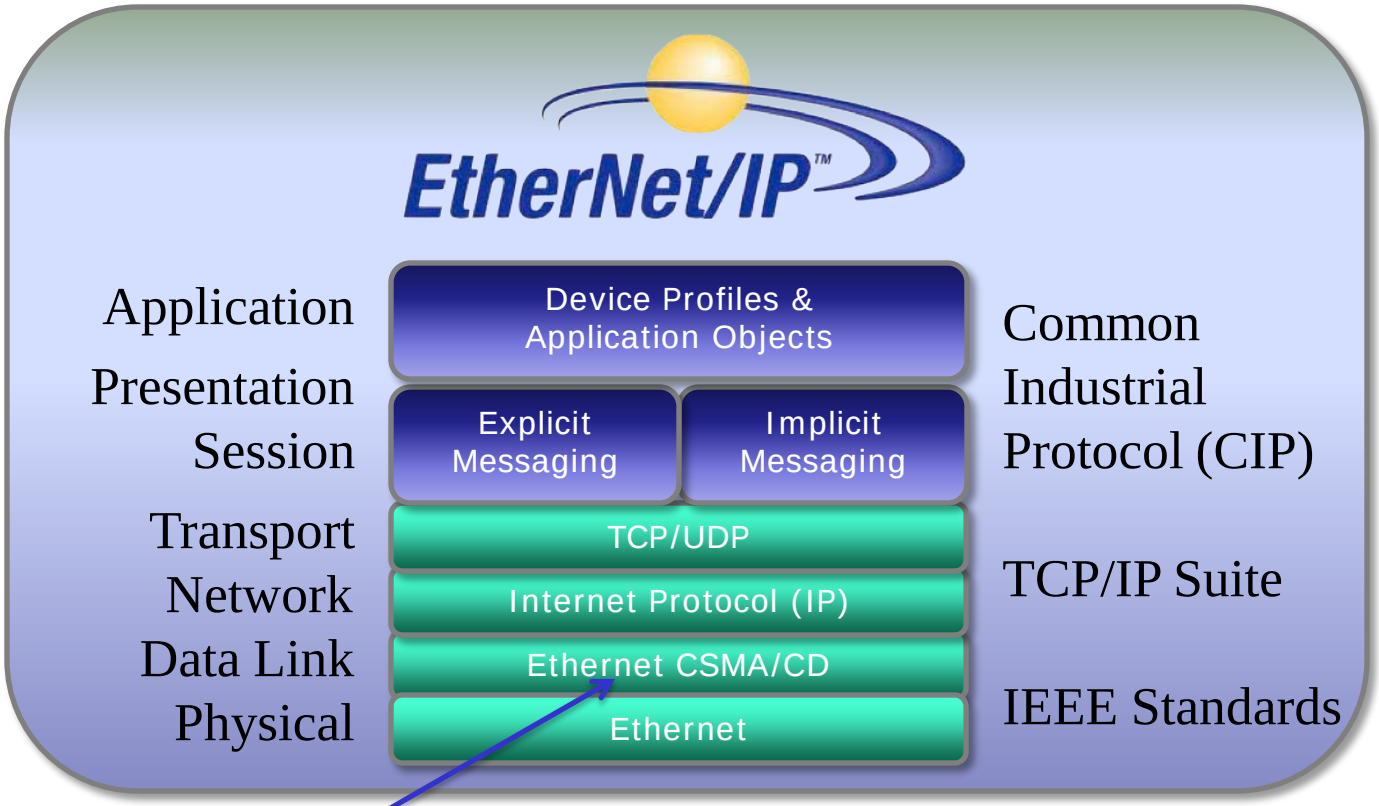

IEEE Audio Video Bridging (AVB) Standards

IEEE AVB/TSN History & Purpose

An IEEE 802.3 Residential Ethernet Study Group was formed in 2005, with major contributions by Broadcom, Nortel, Pioneer, Samsung, NEC, and Gibson (yes, the guitar people).

This morphed into the IEEE 802.1 Audio Video Bridging Task Group (AVB) in 2006. Key players: Intel, Broadcom, Marvell, Samsung,

Originally aimed at home electronics market, for when Ethernet and Wi-Fi replace all other signal wires in the home (HDMI, speaker wire, etc.)

Original goals:

- Make IEEE 1588 Precision Time Protocol Plug-and-Play.
- Offer the ability to reserve bandwidth for particular streams, and to give those streams a low, specified latency and near-zero congestion loss.
- Create an L2 transport protocol for time-stamped (display this packet on the speaker/screen at this moment) AV data. (This work taken to IEEE 1722.)
- Support these features no matter what best-effort traffic is also present.

Who Wants TSN?

Large interest in AVB from:

- Fox Network, ESPN, Formula I, etc., for the next generation television studio.
- Disney, Harman, for the next generation theme park and/or football stadium.
- Gibson, Meyer Laboratories, Riedel, for the next generation audio studio.

Large interest in TSN from:

- Audi, GM, Hyundai, BMW for automobiles
- General Electric, Rockwell, Siemens for industrial control processes.

Large interest in TSN from:

- Broadcom, Marvell, Intel, (chip vendors like standards).

Institute of Electrical and Electronics Engineers

IEEE

- (Communications Magazine …)
- IEEE Standards Association
  - (IEEE 1588 …)
  - IEEE 802 LAN/MAN Standards Committee
    - (IEEE 802.11 Wireless LANs Working Group …)
    - **IEEE 802.1 Higher Layer LAN Protocols Working Group**
      - (IEEE 802.1 Security TG, Data Center Bridging TG)
      - IEEE 802.1 Interworking (IWK) Task Group
      - **IEEE 802.1 Time-Sensitive Networking (TSN) Task Group (was AVB)**
      - (802.1 Task Groups are by area of expertise, not per-project.)

IEEE 802.1 AVB/TSN Finished Standards

IEEE Std 802.1AS-2011 Timing and Synchronization

- Plug-and-play profile of IEEE 1588 Precision Time Protocol
  - Provides automatic master clock selection and clock distribution tree construction.
  - No non-participating network nodes allowed in the distribution tree.
- Every participant is, effectively, a 1588 “boundary clock.”
- Practically speaking, 802.1AS supports only two-step, not one-step, clocking.
  - Inserting time stamps into data frames is hard to do and doesn’t play well with encryption.
  - Remembering when a “sync” frame entered or left the PHY using an internal mechanism, and then following the sync with a Follow-Up packet with that information translated by software, takes twice as many packets, but uses vastly less hardware to achieve the same accuracy as in-packet timestamps. (And, all of the variability for supporting new protocols is in the software, instead of the ASIC!)

IEEE Std 802.11v Wireless Network Management includes features to support 802.1AS.

IEEE 802.1 AVB/TSN Finished Standards

IEEE Std 802.1Qat-2010 Stream Reservation Protocol

- Adds TLVs to IEEE Std 802.1AB-2009 Link Layer Discovery Protocol (LLDP)
  - Allows one to discover which neighbors participate in AVB.
  - Bridges tell end stations what classes of latency are available on what priority levels, and what VLAN to use.
  - Used to ensure that 802.1AS doesn't use non-time-aware devices.
  - Bridges protect AVB priority levels by remapping any input from a non-AVB neighbor away from the AVB priorities.
- Runs a protocol, "Multiple Stream Reservation Protocol (MSRP)" to actually make reservations.
  - Talkers advertise streams, each with a network-unique multicast destination MAC address (likely obtained via IEEE 1722.1), via MSRP, a hop-by-hop flooding protocol.
  - MSRP advertisements can be flooded to all AVB-capable bridges and end stations, or can follow multicast listener reservations (made by either 802.1Q MVRP or by IGMP).
  - Hop-by-hop replies from Listener(s) actually enable hardware queues, from Listener to Talker..
  - When Talker gets reply, transmission can start with guarantees enforced.
- No longer an independent standard – rolled into IEEE 802.1Q-2012 edition.

IEEE 802.1 AVB/TSN Finished Standards

IEEE Std 802.1Qav-2009 Forwarding and Queuing for Time-Sensitive Streams

- Defines a credit-based shaper on which MSRP depends for actual data.
- The sum of the bandwidths of all streams exiting a given port at a given priority is set as the bandwidth of the shaper for that priority queue.
- The 802.1Qav shaper's properties make it easy to do the math when aggregating streams into classes of service.
- No longer an independent standard – rolled into IEEE 802.1Q-2012 edition.

IEEE Std 802.1BA-2011 Audio Video Bridging Systems

- Defines profiles of a number of standards for both end stations and bridges, including all of the above and 802.1Q, to obtain plug-and-play AVB.
- Defines specific values for "Class A" and "Class B" latency guarantees.
- Root standard for answering the question, "Are you AVB-capable?"

What do the AVB Standards Give You?

- Plug-and-play time synchronization with (typically) < 1µs error.
- A high probability that a reserved stream will not encounter congestion loss, which is, of course, the biggest cause of packet loss.
- A high probability that every frame in a reserved stream will be delivered within the promised latency.
- A converged networking infrastructure – no need to have two networks, one for production, and one for best-effort use.

New IEEE Work

IEEE 802.1 AVB/TSN Purpose

By the time the AVB standards were complete, the home electronics market had not climbed aboard the bandwagon.

But, the professional audio and video community had become so interested that products are being offered, today, that comply with AVB standards.

At the same time, the industrial control and vehicle control communities have started to realize that digital networks based on proprietary schemes (ProfiBus, CAN Bus, Flexray, and so on down to RS-232) were a losing proposition in the long term, and that they should shift to Ethernet, instead.

After creating a number of proprietary sort-of-Ethernet schemes, they have realized that truly standardized Ethernet would be even better. And AVB is very close to what they need!

IEEE 802.1 AVB/TSN Purpose

This new industrial constituency has caused the AVB TG to change its name to the IEEE 802.1 Time-Sensitive Networking Task Group (TSN). (Key players now Broadcom, Marvell, Intel, Cisco, Extreme, Ericsson, Gibson, Harman, Siemens, BMW, GM, Hirschman, Rockwell, GE.)

The goals are now to achieve the union of the audio/video and industrial markets' needs:

- Make IEEE 1588 Precision Time Protocol **simpler** to use, **more robust**, and **more accurate**, while keeping its plug-and-play capabilities.
- Offer the ability to reserve bandwidth for particular streams, and to **absolutely guarantee** those streams a low, specified latency and zero congestion loss.
- Allow a single data stream to take **multiple pinned-down paths** through the network to ensure that equipment failure will not cause packet loss.
- Allow **any data stream** to reserve this QoS, whether unicast or multicast, regardless of high-layer protocol, whether L2 or L3 end-to-end.
- Converge all the usual **existing Qualities of Service** into the TSN network (at reduced bandwidth) while maintaining the TSN guarantees.

What is Missing (to be Supplied)?

Robustness of the clock accuracy in the face of a master clock failure.

Fast bring-up after power-on.

The ability to handle a large number (1000s or 10,000s) of very slow streams.

100% guarantees of latency and congestion loss.

Frame loss rates in the 10^{-9} to 10^{-12} range over long periods of time (hours – years).

Wi-Fi for TSN.

These features in routers.

IEEE 802 Current TSN-Related Projects

P802.1ASbt (TSN) Enhancements for Timing and Synchronization

- Multiple clock sources
- Fast failover from source to source
- Robustness against master clock failure **including** complex failure cases.

P802.1Qbu (TSN) Frame Preemption (see P802.3br)

- Interrupt an 802.3 frame with lower latency requirements one or more times during its transmission, in order to send frames with more stringent latency requirements.
- This 802.1 effort is the (rather minor) bridge side of the (much larger) 802.3br effort.

IEEE 802 Current TSN-Related Projects

P802.1Qbv (TSN) Enhancements for Scheduled Traffic

- Provide a rotating schedule of time slots that enable and disable Class of Service Queues.
- All 1 ≤ N ≤ 8 queues are gated independently, and the gating function is placed between “this queue is not empty” and “select a queue for transmission”.
- This makes possible just about any kind of scenario for overlap, send-this-unless-this-other-is-ready, use-unused-bandwidth-only-for-these-queues, etc.

P802.1Qca (IWK) Path Control and Reservation (IWK)

- Limited to modifications to IS-IS (as used by bridges – IEEE Std 802.1aq-2012)
- Can replace MSRP as a medium for distributing the stream requests.
- May be used to distribute the results from a Path Computation Element (implication: the IETF PCE) to establish pinned-down paths (we’ll see).
- May be used to distribute 802.1Qbv time schedules (we’ll see).

IEEE 802 Current TSN-Related Projects

P802.1CB (TSN) Frame Replication and Elimination for Reliability (Seamless Redundancy)

- Just getting started.
- Will probably define a new tag, following VLAN tag, that will carry a sequence number.
- Unicast or multicast stream data is sequence and replicated by host or by bridge, sent along all of more than one path, and the duplicates eliminated based on sequence number, either in the receiving host(s) or edge bridge(s).
- 802.1 may choose to provide, at intermediate points in the network, this replication / elimination service, in order to regenerate dual streams after a single failure, thus ensuring against multiple failures.

P802.1Qcc (TSN) Stream Reservation Protocol Enhancements

- Fix the user interface to MSRP and/or P802.1Qca so that we won't have to visit the problem again, even if we define new ways to make the data plane work better.

IEEE 802 Current TSN-Related Projects

IEEE P802.3br Interspersed Express Traffic (see P802.1Qbu)

- Finally on track, after much political (not in the bad sense) haggling.
- Interrupt an 802.3 frame with lower latency requirements one or more times during its transmission, in order to send frames with more stringent latency requirements.
- Supports lock-down, to prevent resumption of interrupted frame in a small pause between interrupting frames.
- The 802.1 effort is the (rather minor) bridge side of the (much larger) 802.3br effort.
- Only one level of interruption, meaning two buffers in the receiver.
- Software negotiation will be used to enable the feature.
- Changes go between the MAC and the PHY, to minimize ASIC impact.

IEEE 802 Current TSN-Related Projects

P802.1Qbz (TSN) Enhancements to Bridging of 802.11

- Small effort required to support P802.11ak.

IEEE P802.11ak Enhancements For Transit Links Within Bridged Networks (See also P802.1Qbz)

- Non-trivial effort to provide a bridge (or router or whatever) that is co-resident with an Access Point (a collection of logical AP functions) to be able to treat each association with a non-AP station like an independent point-to-point link for the purposes of multicast transmission.
- This means using a new frame format, similar to 802.11n A-MSDUs (aggregated frames), that carry a list of non-AP stations' association IDs to be includes or excluded from receiving the A-MSDU.
- This format also replaces the 802.2 LLC format with 802.3 Type/Length encoding.
- With this change, an AP station *or a non-AP station* can be part of an ordinary 802.1Q bridge, which means **that 802.11 can be used inside**, as well as at the edge of, **a Bridged LAN**.

Some non-obvious "Why"s

Transmission preemption is not simply to improve the latency of the preempting frames. At best, that only reduces the high-priority queue by a single frame transmission time.

- Preemption does, of course, allow the use of giant frames without killing the TSN frames' latency.
- But equally important, it allows the precise scheduling of TSN frames for 0 jitter, without worrying about the size of the gaps between scheduled frames.
- So, it both improves critical jitter, and reduces the impact of TSN on existing QoS algorithms.
- When combined with scheduling, supports end-to-end cut-through forwarding.

Industrial networks are built to extremely high standards for reliable data delivery.

- In a factory with 1000s of Ethernet networks, where one machine's failure stops the line, requirements for 10^{-6} packet loss ratios are not uncommon, and 10^{-9} desired.
- **And, the critical network must be converged with the Enterprise network, in order to support virtual machine controllers in the Data Center.**

Relevance to the ODVA "style"

Scheduling and credit-based shapers are not the ODVA style.

- But, they offer a low-latency capability that, when you need it, is handy to have.
- These are tools in the IEEE 802.1 TSN toolbox, not mandatory usage models.
- TSN always leaves a guaranteed bandwidth and latency for the highest-priority best-effort traffic.

What More is Needed?

Layer 3, Layer 3, Layer 3

This AVB/TSN stuff only works if you have an end-to-end path through bridges.

Most users (though not the automobile people – YET) realize that this sort of capability has to be routable.

A number of people (mostly from Cisco, Ericsson and Broadcom) are working on casting these features in a form useable by routers.

- Preferably based on existing technologies such as MPLS and IP multicast.
- Preferably replacing the L2-only protocols with merged L2/L3 protocols that start from the base of L3 IETF work.

New Data Plane Mechanisms

The 802.1Qav shaper is not adequate as used on a Class of Service Queue. There are several proposals for new data plane mechanisms.

1. Peristaltic Scheduler: All ports on all bridges are in lock-step, and alternate between buffers at cycle boundaries. Any given stream is defined as using ≤ X bytes per cycle (including interframe gap), where cycle size ≥ X ≥ 64+20. Per-hop transmission and forwarding delay is negligible.

2. Various improvements on (1), being developed by Cisco (and presumably, others).

3. One 802.1Qav shaper per prioritized flow, instead of per CoS. Requires an NP-complete calculation to configure, but delivers the best possible latency numbers.

4. Every single input and every single output to/from a buffer is scheduled explicitly by a pre-run-time algorithm. (This is the method pioneered and deployed in AirBus planes and Audi automobiles by TTTech of Vienna, Austria.)

Other Standards of Interest

Other Standards of Interest

IEEE Std 1588-2002 Precision Time Protocol. (Version 2 in progress).

- Some people, especially Broadcom and Marvel, are trying to rationalize this standard so that "profiles" don't have to re-invent the wheel.

IEEE Std 1722-2011 Layer 2 Transport for Time Sensitive Streams

- This is not unlike RTP, but it uses the 802.1AS or 1588 time synch, instead of trying to do time sync and transport in the same protocol.

IEEE P1722a Amendments to 1722-2011.

- This may go as far as defining CAN Bus over Ethernet.

IEEE P1722.1 Discovery, Enumeration, Connection and Control

- Provides a means of obtaining multicast MAC addresses to use for AVB flows that are unique in a Bridged LAN. (At least, for a while.)

Other Standards of Interest

SAE AS6802: Time Triggered Ethernet

- Society of Automotive Engineers standard that enshrines, without actually specifying the details of, TTTech's clock algorithms.

IEC 62439-3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

- PRP: Allows a dual-homed host to connect to two completely independent bridged networks, and to talk redundantly with another, similar, dual-homed host.
- PRP: Defines a Redundancy Box (RedBox) that dual homes for some number of single-homed hosts.
- HSR: Defines an encapsulation for a dual-homed host to operate in a ring of similar hosts, sending data in both directions.
- HSR: Defines a "QuadBox" that interconnects two HSR rings.
- These protocols can be used together to build networks that are sort of Ethernet, but provide no redundancy at the routed (L3) layer.

AVnu Alliance

AVnu Alliance is an industry consortium.

Cisco is a founding member, along with Intel, Broadcom, and several other of the usual suspects (above).

This provides marketing and technical support for IEEE 802.1 AVB/TSN, rather like Wi-Fi Alliance does for IEEE 802.11.

They are in the process of expanding their membership to include major industrial networking suppliers and users, and automobile makers and equipment vendors.

They are expanding their focus to include L3 issues, as well as L2 issues.

One Last Note…

Participation in the development of these TSN standards is encouraged!

IEEE Standards Update

Thank You