BASELINE INTEROPERABILITY SPECIFICATION

JANUARY 30, 2019

REVISION 2.0A

AUTHOR

Fabian Braun (d&b audiotechnik)

CONTRIBUTORS

Andrew Elder (AudioScience)
Bob Noseworthy (UNH)
Christophe Calmejane (L-Acoustics)
Cole Peterson (Meyer Sound)
Davy Hermans (Luminex)
Frank An (Avid)
Genio Kronauer (L-Acoustics)
Henning Kaltheuner (d&b audiotechnik)
Kieran Tyrrell (Sienda)
Marc Illouz (L-Acoustics)
Morten Lave (Adamson)
Rekha Ramachandran (Cisco)
Richard Bugg (Meyer Sound)
All members of the Avnu ProAV Technical Workgroup

TABLE OF CONTENTS

1. INTRODUCTION

This document is Avnu's baseline specification for Milan endpoints.

Milan is based on AVB technology to guarantee deterministic transport of media. In addition a set of specifications is defined to allow seamless interoperability across ProA manufacturers.

The design goals of Milan for ProA devices include:

- A common wall clock for all devices
- Distribution of media clocks
- Guaranteed bandwidth for audio- and clock-streams
- Coexistence with non real-time traffic on the same network
- Low-latency and synchronized transmission of audio
- Common Discovery, Enumeration, Connection management and Control
- Operation on links with speeds at or above 100Mbps
- Co-existence with other AVB devices in a non-disruptive manner in a dynamic networking environment
- Plug-and-play
- Support for enhanced reliability through redundancy

To deliver these goals, this specification draws upon the various AVB standards and Milan specifications to specify a concise set of requirements for delivering the features necessary for the ProA use cases.

2. REFERENCES

Name	*Reference*
802.1BA	IEEE Std 802.1BA-2011, "Audio Video Bridging (AVB) Systems".
802.1Q	IEEE 802.1Q-2014, "Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks".
AVDECC	IEEE 1722.1-2013, "IEEE Standard for Device Discovery, Connection Management, and Control Protocol for IEEE 1722 Based Devices".
AVTP	IEEE 1722-2016, "IEEE Standard for Layer 2 Transport Protocol for Time-Sensitive Applications in Bridged Local Area Networks".
gPTP	IEEE Std. 802.1AS-2011 "IEEE Standard for Local and Metropolitan Area Networks – Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks".
gPTP-Cor1	IEEE Std. 802.1AS-2011/Cor 1-2013 "IEEE Standard for Local and metropolitan area networks— Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks— Corrigendum 1: Technical and Editorial Corrections".

gPTP-Cor2	IEEE Std. 802.1AS-2011/Cor 2-201X “IEEE Standard for Local and metropolitan area networks— Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks— Corrigendum 2: Technical and Editorial Corrections”.
FORMATS	Milan Formats interoperability specification - Revision 2.1.
CLOCKING	Milan Media clocking functional & interoperability specification - Revision 2.1
CONTROL	Milan Discovery, connection and control for talkers and listeners functional & interoperability specification - Revision 1.0.
REDUNDANCY	Milan Network redundancy interoperability specification - Revision 1.3.

3. GLOSSARY

Term	*Meaning*
AVB	Audio Video Bridging. The term to describe the suite of standards enabling precisely synchronized networking devices to communicate with guaranteed bandwidth and known worst-case latencies.
AVTPDU	Audio Video Transport Protocol Data Unit as defined in [AVTP].
BMCA	Best Master Clock Algorithm as described in [gPTP].
CBS	Credit Based Shaper Algorithm as defined in [802.1Q, Clause 34].
FQTSS	Forwarding and Queuing for Time-Sensitive Streams as defined in [802.1Q], Clause 34.
Listener	A PAAD which is able to receive and consume a Stream.
MAAP	MAC Address Acquisition Protocol as defined in [AVTP, Annex B].
Milan	Audio networking solution designed and promoted by the Avnu Alliance, based on the AVB and AVDECC standards.
MRP	Multiple Reservation Protocol as defined in [802.1Q, Clause 10].
MRPDU	Multiple Reservation Protocol Data Unit as defined in [802.1Q, Clause 10].
MSRP	Multiple Stream Reservation protocol as defined in [802.1Q, Clause 35].
MVRP	Multiple VLAN Reservation protocol as defined in [802.1Q, Clause 11.2].

PAAD	A professional audio device with Ethernet AVB functionality.
ProA	Professional audio
Settled sink	A Sink which has successfully probed a remote Source via the Connection Management protocol, as described in [AVDECC, Clause 6.8.5].
Sink	A component within an AVDECC Entity that consumes a Stream, as described in [CONTROL].
Source	A component within an AVDECC Entity that produces a Stream, as described in [CONTROL].
SR class	Stream Reservation class.
Talker	A PAAD which is able to produce and transmit a Stream.
VID	The ID of a VLAN.
VLAN	Virtual Local Area Network.

4. SCOPE

This document's primary goal is to provide a baseline specification of Ethernet AVB functionality for Professional Audio AVB devices as a reference for manufacturers of endpoints.

All standards and specifications that are relevant for implementing an interoperable professional Audio Endpoint are referenced in this document. It also imposes additional requirements beyond those defined in the standards where necessary.

5. COMMON REQUIREMENTS

The following set of AVB standards and Milan specifications shall be considered the source of any additional requirements when such statements are made throughout this specification. This specification and Milan specifications supersede the standards referenced below.

5.1. AUDIO VIDEO BRIDGING

A PAAD shall implement all applicable requirements that are otherwise not discussed in this specification, as defined by [802.1BA].

5.2. AUDIO TRANSPORT

A PAAD shall implement all applicable requirements that are otherwise not discussed in this specification, as defined by [AVTP].

A PAAD shall implement all applicable requirements as defined by [FORMATS].

5.3. MEDIA CLOCKING

A PAAD shall implement all applicable requirements that are otherwise not discussed in this specification, as defined by [AVTP].

A PAAD shall implement all applicable requirements as defined by [CLOCKING].

5.4. CONNECTION MANAGEMENT AND CONTROL

A PAAD shall implement all applicable requirements as defined by [CONTROL].

A PAAD shall set the value of the protocol_version field in the GET_MILAN_INFO response to 1.

5.5. REDUNDANCY

A PAAD may implement seamless redundancy as defined by [REDUNDANCY].

If a PAAD implements redundancy compliant to [REDUNDANCY], it shall set the REDUNDANCY bit in the GET_MILAN_INFO response as defined in [CONTROL].

5.6. GPTP

A PAAD shall implement all applicable requirements that are otherwise not discussed in this specification, as defined by [gPTP], [gPTP-Cor1], and [gPTP-Cor2].

Differences between a standard gPTP implementation and a PAAD implementation as well as clarifications or emphasis on key standard-defined requirements are defined below.

5.6.1. gPTP Configuration

A PAAD shall support the BMCA as defined in [gPTP, Clause 10.3].

Note: the BMCA establishes a timing spanning tree for propagating Announce, Sync and Follow Up messages. This Timing Spanning Tree is independent of any RSTP or MSTP established by these protocols.

5.6.1.1. neighborPropDelayThresh [gPTP-Cor1, Clause 11.2.2]

As per [gPTP-Cor1], aPAAD utilizing a copper interface (specifically 100Base-TX or 1000Base-T) shall implement a neighborPropDelayThresh threshold with a default value of 800 ns on this port.

On a fiber port, the PAAD shall disable the neighborPropDelayThresh by default.

5.6.2. gPTP Operation

5.6.2.1. Priority1 [gPTP, Clause 8.6.2.1]

If a PAAD is Grandmaster Capable the default priority1 value shall be 248 (Other Time-Aware System).

Note: Section 6.1.1 of this document defines that all Talker PAADs have to be Grandmaster Capable.

5.6.2.2. gPTP Message Intervals

Allowed gPTP Interval values shall be defined as follows:

gPTP Interval	*Tolerance*	*Default*	*Minimum*	*Maximum*
Pdelay transmission	+50%/-10%	1s	900ms	1500ms
Announce transmission	+50%/-10%	1s	900ms	1500ms
Sync transmission	+50%/-10%	125ms	112.5ms	187.5ms

Table 1 – gPTP Interval Tolerances

Note: These values are derived from adjusting the nominal interval time -10% and +50% to allow for variance in the observed intervals. The requirement is similar to the IEEE 1588 interval requirements of -30% and +30%, but is shifted to encourage systems to favor providing a slower interval, rather than a faster interval, to ease the processing burden on resource constrained systems.

5.6.2.3. gPTP Timeout Value Tolerances

Allowed gPTP timeouts values shall be defined as follows:

gPTP Timeout	*Tolerance*	*Default*	*Minimum*	*Maximum*
announceReceiptTimeout (3 times allowed Announce transmission interval)	+50%/-10%	3s	2.7s	4.5s
syncReceiptTimeout (3 times allowed Sync transmission interval)	+50%/-10%	375ms	337.5ms	562.5ms
followUpReceiptTimeout (same as Sync transmission interval)	+50%/-10%	125ms	112.5ms	187.5ms

Table 2 – gPTP Timeout Tolerances

Note: These values are derived from the -10% and +50% tolerances allowed on the corresponding intervals.

5.6.2.4. asCapable [gPTP, Clause 10.2.4.1]

A PAAD shall report asCapable as TRUE after no less than 2 and no more than 5 successfully received Pdelay Responses and Pdelay Response Follow Ups to the Pdelay Request messages sent by the device.

Note: This requirement ensures that all certified PAADs become asCapable within a bounded time.

5.6.2.5. Receipt of multiple Pdelay Responses for one Pdelay Request

A PAAD that receives more than one Pdelay Response messages from multiple clock identities in response to each of three successive Pdelay Request messages shall cease Pdelay Request transmission until one of the following conditions is met:

- The link state toggles
- portEnabled is toggled
- pttPortEnabled it toggled
- after 5 minutes

When one of the above conditions is met, the Pdelay Request transmission shall resume.

Note: this avoids storms of Pdelay Responses when a large number of PAADs is connected through non-AVB switches.

5.6.2.6. Pdelay Turnaround Time

Per [gPTP, Annex B.2.3], Pdelay turnaround time must be less than or equal to 10ms. This requirement shall be relaxed by 50% for Avnu purposes resulting in a maximum Pdelay turnaround time of 15ms.

Responses later than this time may or may not be processed by a gPTP device, but should not result in asCapable being set to FALSE if 3 or more consecutive responses are received later than 10ms but before pdelayReqInterval (typically 1 second).

5.6.2.7. Expected behavior of a gPTP clock when it computes a negative pdelay

A PAAD is expected to accept that the calculation of pdelay can result in negative values. Negative pdelay values between -80 ns and 0 ns shall not cause asCapable to be set to false.

5.7. MRP, MSRP, MVRP

5.7.1. MRP Operation

A PAAD shall support MRP.

5.7.1.1. MRP Timer Tolerances and Default Values

Allowed MRP Timer values shall be defined as follows:

MRP Timer	*Tolerance*	*Default*	*Minimum*	*Maximum*
periodictimer	+50%/-10%	1000ms	900ms	1500ms
joinTime	+20%/-10%	200ms	180ms	240ms
LeaveTime	+50%/-10%	5000ms	4500ms	7500ms
leavealltimer	±0.5s	10-15s	9.5s	15.5s

Table 3 - MRP Timer Tolerances

5.7.1.2. Handling of badly formed MRPDUs

A device that receives a badly formed MRPDU is expected by the standard to discard the entire PDU. This requirement is relaxed by this specification to allow for implementations that choose to process the portion of the MRPDU prior to the improper field.

In all cases, an MRP implementation shall discard the information following an invalid field within the same Vector Attribute list, and any subsequent Messages received in the same MRPDU.

Note: Refer to [802.1Q, Figure 10-5] to review MRPDU structure including Message and Vector Attribute List structure.

5.7.1.3. EndMarks for MRPDUs

When an MRPDU is to be transmitted, if sending the EndMark as "End of PDU" would result in PAD following the "End-Mark", then the EndMark shall be sent as 0x0000.

Note: This is a clarification of the requirements in 10.8.1.2 that define the EndMark to be either "End of PDU" or 0x0000.

An Ethernet frame that is smaller than the minimum frame size is required to be transmitted with a PAD field.

Note: The content of this field is undefined, but it is typically filled with bytes of 0x00. The requirement above ensures an EndMark is distinguishable from the PAD field.

5.7.2. MSRP Operation

A PAAD shall support MSRP.

5.7.2.1. Domain Declarations

After startup or a Link Up event, a PAAD shall use SR Class Priority 3 and Default VLAN ID 2 for Class A, and send MSRPDUs declaring an MSRP Domain for Class A using these parameters. If the PAAD receives an MSRP Domain attribute declaration for Class A specifying different parameters, it shall update its parameters using the received FirstValue and start declaring a matching MSRP Domain attribute.

Transmission of MSRP Domain Messages shall not be dependent on the state of gPTP on the port.

5.7.2.2. Instantaneous transition from IN to MT

For the MSRP application, the following transition of the MRP Registrar state machine:

IN / rLv! → (Start leavetimer) → LV

shall be changed by the following one:

IN / rLv! → (Lv) → MT

Note: this improvement allows the PAAD to use a 5-second value for the LeaveTime without requiring 5 seconds to detect that a Stream has been explicitly withdrawn from the network.

5.7.3. MVRP Operation

A PAAD shall support MVRP.

6. TALKER PAADS REQUIREMENTS

6.1. GPTP OPERATION

6.1.1. gPTP grandmaster capability

As per [802.1BA, Clause 6.7.2], a Talker PAAD shall be capable of operating as a gPTP grandmaster.

6.2. MVRP OPERATION

A Talker PAAD shall join the relevant VLAN via MVRP prior to sending any stream frames.

6.3. MSRP OPERATION

6.3.1. Talker attribute declaration

This section clarifies the statements of [CONTROL, Clause 6.7.4].

[CONTROL, Clause 6.7.4] states that "a PAAD shall declare a Talker attribute for each of its STREAM_OUTPUTs that has valid SRP parameters". It also states that the definition of "valid" is left open to the implementer.

This specification defines what "valid" means for the Stream Destination MAC Address parameter. This parameter is valid when both of the two following conditions are satisfied:

1) The Destination MAC Address has been allocated by the MAAP protocol and no conflict is being reported, and
2) A PROBE_TX_COMMAND message targeted to this STREAM_OUTPUT has been received during the last 15 seconds, or the PAAD is registering a Listener attribute that matches the Stream ID for this STREAM_OUTPUT.

Note: The goal is to avoid declaring talker attributes for streams that are not requested by any Listener, saving network resources.

6.3.2. Bandwidth reservation

A Talker PAAD shall use the following TSPEC parameters when declaring a Talker attribute for one of its output streams, and the following bandwidth when configuring the Credit-Based Shaper:

Stream format	***MaxFrameSize (bytes)***	***MaxIntervalFrames***	***Bandwidth (kbps)***
AAF PCM32, 48kHz, N channels	24*N + 24 + 1	1	1536*N + 4288
AAF PCM32, 96kHz, N channels	48*N + 24 + 1	1	3072*N + 4288
AAF PCM32, 192kHz, N channels	96*N + 24 + 1	1	6144*N + 4288
AAF PCM24, 48kHz, N channels	18*N + 24 + 1	1	1152*N + 4288
AAF PCM24, 96kHz, N channels	36*N + 24 + 1	1	2304*N + 4288
AAF PCM24, 192kHz, N channels	72*N + 24 + 1	1	4608*N + 4288
CRF, 1 ts/pdu	28 + 1	1	5632

Note: MaxFrameSize accounts for the size of the AVTPDU header and the size of the payload. One more byte is added to take the fact into account, that the sampling clock of the PAAD may be a bit faster than the nominal frequency.

Note: the rules used to compute the bandwidth occupied by a stream are as follows:

1) *Compute the size of the Ethernet frame by adding to MaxFrameSize the size of the Ethernet header (including the VLAN tag) and the FCS: F = MaxFrameSize + 22.*
2) *Ensure that F is not less than the minimum-sized Ethernet frame: if F<68 then set F=68.*
3) *Compute the size occupied by the frame on the wire, by adding to F the Ethernet preamble (8 bytes) and the Inter-Packet Gap (12 bytes): W = F + 20.*
4) *Multiply W by MaxIntervalFrames to obtain the size occupied per measurement interval, then by 8000 to obtain the size (in bytes) occupied per second, and again by 8 to obtain the size (in bits) occupied per second. This is the associated bandwidth in bits per second.*

6.4. FORWARDING AND QUEUING FOR TIME SENSITIVE STREAMS

A PAAD shall implement all applicable requirements that are otherwise not discussed in this specification, as defined by [802.1Q, Clause 34].

A Talker PAAD shall implement the CBS.

A Talker PAAD shall shape each individual stream, as well as the overall SR class, as described in [802.1Q, Clause 34.6.1].

6.5. AVTP OPERATION

6.5.1. MAAP

A PAAD shall implement Multicast Address Allocation Protocol (MAAP) as defined by [AVTP, Annex B].

A PAAD shall use destination MAC addresses allocated with MAAP for its output streams.

Note: A device using MAAP should not use only the low 3-bytes of the address when determining address usage. When detecting conflicts per IEEE Std 1722-2016 note b, table B.7, the device must compare the entire 48-bit address range with the block it is defending, versus the full 48-bit value of the address range of the MAAP_PROBE PDU, MAAP_DEFEND PDU, or MAAP_ANNOUNCE PDU received. A common mistake would be to assume that all MAAP address ranges used will have the same high three-bytes (0x91-E0-F0) which may not be the case if MAAP is in use for other purposes.

6.5.2. "tu" bit (timestamp uncertain field)

A Talker PAAD shall set the AVTP “tu” bit as described in [AVTP, Clause 4.4.4.7].

7. LISTENER PAADS REQUIREMENTS

7.1. MVRP OPERATION

A Listener PAAD shall declare an MVRP VID attribute for each VLAN used by its settled sinks.

7.2. AVTP OPERATION

7.2.1. Presentation Time supported by Listener

A Listener PAAD shall support buffering of received audio samples for up to 2,126 ms.

Note: This is derived from:
2ms (maximum presentation time offset) + 125us (timing uncertainty) + 1us (gPTP accuracy).

7.2.2. Filter the incoming AVTPDUs

A Listener PAAD shall discard the incoming AVTPDUs that don't respect the format which is configured in the associated STREAM_INPUT descriptor.

7.2.3. "tu" bit (timestamp uncertain field)

A Listener PAAD should use the AVTP "tu" bit as described in [AVTP, Clause 4.4.4.7], with the following addition:

- After detecting that the "tu" bit has been reset, the Listener PAAD should continue to let its media clock free-wheel for an appropriate amount of time. The free-wheel time should be long enough to cover the cases where the Listener PAAD is informed of the change of gPTP grandmaster/timing source after the Talker PAAD has already locked to this new gPTP grandmaster/timing source.