

<http://aimsalliance.org>

AES67 / SMPTE ST 2110 COMMONALITIES AND CONSTRAINTS

Explanation of the relationship between the SMPTE ST 2110 standard
and the AES67 standard from the Audio Engineering Society

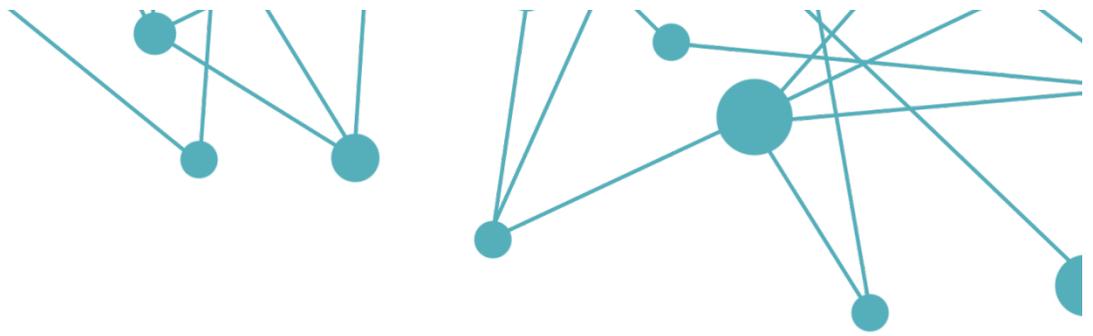
Updated – APRIL 2019

JOIN THE ALLIANCE

The Alliance for IP Media Solutions (AIMS) is a non-profit trade organization founded by leading companies to foster the adoption of one set of common, ubiquitous, standards-based protocols for interoperability over IP in the media and entertainment, and professional audio/video industries.



AIMS
Alliance for IP Media Solutions



CONTENTS

Contents.....	2
EXECUTIVE SUMMARY	3
AES67	4
SMPTE ST 2110	4
Time Synchronization	5
Synchronization and Media Clocks	6
Multicast	7
Other Protocols.....	7
Other Functions	8
Conformance Levels.....	8
Conformance to ST 2110	10
Conformance to AES67	10
Summary.....	12
Appendix A – Document Links	13
Appendix B – Background on reference clock, media clocks and RTP clocks	14

Revision 2 – change log:

- IGMP Requirements
- Multicast Address Range

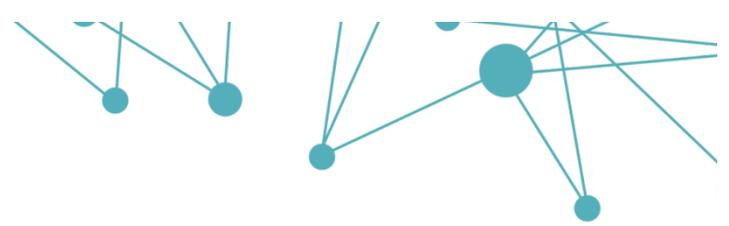
EXECUTIVE SUMMARY

The “AES67 Standard on High-performance streaming audio-over-IP interoperability” from the Audio Engineering Society is referenced by the “SMPTE ST 2110 Standard on Professional Media over Managed IP Networks” for transport of PCM digital audio (SMPTE ST 2110-30). This paper is intended to explain the key points of constraints in the SMPTE ST 2110 standard with respect to the AES67 standard. In addition, important commonalities are explained in order to provide a broader knowledge on how these constraints might affect interoperability between devices strictly adhering to SMPTE ST 2110 and devices following the requirements set forth in AES67.

Basically, SMPTE ST 2110-30 can be seen as a subset of AES67. The general operational principles and mandatory requirements for stream transport, packet setup, and signalization defined by SMPTE ST 2110 are identical to those of AES67. However, SMPTE ST 2110 defines further constraints to which AES67 implementations must adhere in order to ensure full compatibility:

- Support of the PTP profile defined in SMPTE ST 2059-2
- An offset value of zero between the media clock and the RTP stream clock
- Required option to force a device to operate in PTP slave-only mode
- Support of IGMPv3





AES67

AES67 is a standard defined by the Audio Engineering Society to enable high-performance audio-over-IP streaming interoperability between various IP-based audio networking products currently available. It is not a new technology; rather, it is an interoperability mode that enables bridging of various IP-based audio networking technologies.

The AES67 standard was first published in September 2013, and revised updates were published in 2015 and 2018. This standard provides comprehensive interoperability recommendations in the areas of synchronization, media clock identification, network transport, encoding and streaming, session description, and connection management.

SMPTE ST 2110

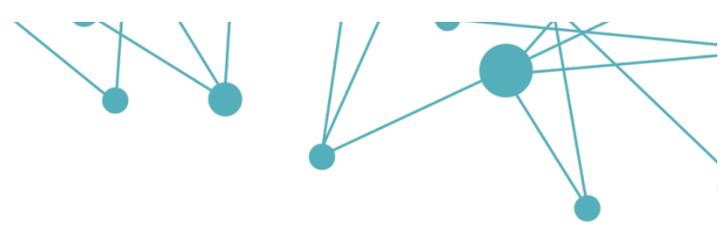
In September 2017, SMPTE announced the approval of the first set of standard documents within SMPTE ST 2110, “Professional Media Over Managed IP Networks”. This suite of standards specifies the carriage, synchronization, and description of separate elementary essence streams over professional internet protocol (IP) networks in real-time for the purposes of live production, playout, and other professional media applications.

SMPTE ST 2110 standards make it possible to separately route and break away the essence streams — audio, video, and ancillary data. Each essence flow may be routed separately and brought together again at the endpoint.

SMPTE ST 2110-10/-20/-30 — addressing system requirements and payload formats for uncompressed video and audio streams — were the first approved standards within this suite. With the revision of this document, ST 2110-21/-31/-40 have been published.

The section of the standard related specifically to PCM digital audio (ST 2110-30) refers to the well-established AES67 standard. However, the SMPTE standard does impose some constraints regarding AES67 that developers and end users need to consider.

The -31 document relates to the bit-transparent transport of AES3 signals, which goes beyond the definition of AES67 but can be used as an alternative transport format when PCM digital audio transport is not sufficient.



Time Synchronization

AES67 and ST 2110 both utilize the IEEE 1588-2008 Precision Time Protocol (PTPv2) for distribution of the reference clock. However, while AES67 mandates support of the PTP Default profile and recommends support of the AES67 PTP Media profile, defined in Annex A of the AES67 standard, ST 2110 requires support of the PTP profile defined in SMPTE ST 2059-2.

An examination jointly conducted by SMPTE and AES identified the commonalities of all three profiles. A report summarizes the differences and commonalities and provides recommendations for configuration when operating ST 2110 and AES67 devices in a common network. The report (AES-R16-2016) is available from the AES library (see link in Appendix A of this document).

Other relevant PTP requirements of ST 2110:

- Option to force a device to operate in slave-only mode:

ST 2110 requires the presence of a method or control to allow a user (administrator) to force the device into PTP slave-only operation by setting the PTP dataset member *defaultDS.slaveOnly* to a TRUE state, so that it will never even attempt to become an elected Grandmaster. AES67 does not have this requirement, and AES67 devices without this configuration capability rely solely on the IEEE 1588-2008 Best Master Clock Algorithm (BMCA) to decide whether to become Grandmaster for the network¹.

- Signalization of reference clock:

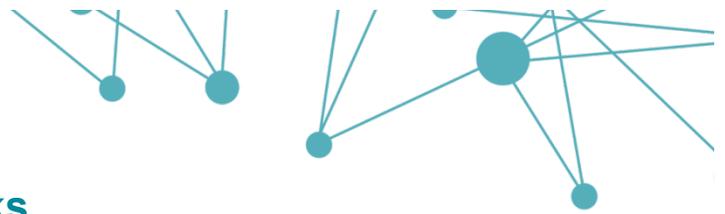
While it is expected that devices are usually referenced to a PTP Grandmaster clock or any other traceable time source, in some cases, a device may operate without a PTP reference on a local (internal) clock basis. This state is signaled in the related SDP of any outgoing stream by using the *localmac* token and the device's MAC address:

a=ts-refclk:localmac=<Ethernet MAC address of sender>

Since this is a “private” token not defined in the referring RFC 7273, an AES67 device might not be able to interpret this information. However, AES67 receivers are likely to react in an appropriate manner by not recognizing the clock reference and not trying to connect to such a stream at all.²

¹ Typically, an accordingly high value preconfigured in the PTP dataset members *defaultDS.priority1* and / or *defaultDS.clockQuality.clockClass* will translate into a very low BMCA ranking order, thus resulting in a PTP slave role in the presence of a dedicated or better suited GM device.

² AES67 standard states: “Receivers should not attempt to connect to senders if they are using a different PTP domain for their clock reference than the sender.”



Synchronization and Media Clocks

One important constraint in the ST 2110 standard compared to AES67 is the requirement for an **offset of zero** between the media clock and the RTP clock.

AES67 states that “*RTP clocks operate with a constant offset with respect to the media clock.*” The underlying RFC 3550 states that “*The initial value of the timestamp SHOULD be random, ...*”³

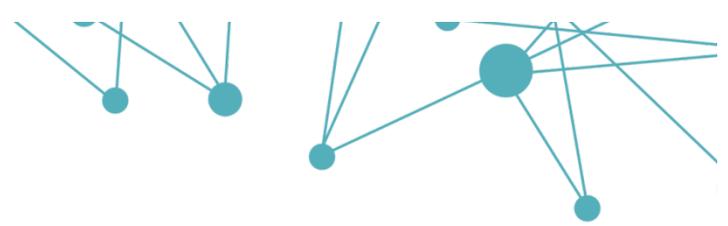
ST 2110-10 requires “*... RTP streams conforming to this standard shall utilize an RTP clock offset of zero from the Media Clock.*” The reason for this requirement is explained in a note stating, “*The requirement of a zero offset value in this standard allows fast restoration of signals after Sender restarts. Eliminating the random offset provision of IETF RFC 3550 allows the receiver to make use of the signal as soon as the packet stream is restored, without waiting for the systemic propagation of a revised SDP object.*”

This requirement only affects AES67 **senders**. To be ST 2110 compliant, they must generate streams with an RTP timestamp offset of 0 with respect to the media clock, effectively aligning the RTP clocks of all streams with the media clock (which is directly related to the reference clock). AES67 **receivers** are not affected because they are required to deal with any signaled offset.⁴

For further information on the relationship of reference clock, media clock, and RTP clock, the curious reader is directed to Appendix B.

³ The justification for a random offset is for security reasons in an encrypted RTP scenario.

⁴ ST 2110 states in a further note that “[ST 2110] Receivers designed to maintain compatibility with other RTP implementations will need to comply with the RTP Clock offset provisions in those RTP standards, specifically the possibility that the RTP Clock could be offset from the Media Clock.”



Multicast

In both ST 2110 and AES67, multicast support is required for stream data transport. However, the following differences need to be noted:

- IGMP: ST 2110 requires support of IGMPv3 (RFC 3376), while AES67 requires support of IGMPv2 (RFC 2236). However, IGMPv3 is designed to be interoperable with version 2 (and version 1); actually, RFC 3376 obsoletes RFC 2236. So, from a pure standards point of view, AES67 devices that only support IGMPv2 can operate in an ST 2110 environment.

However, depending on the topology of the network and the particular backward-compatibility implementation of the affected network switches and routers, parts of the network (or the network as a whole) might fall back to IGMPv2 operation when an IGMPv2-only device is connected to the network. This circumstance might negatively affect the possibility of running source-specific multicast (SSM) for some streams or in parts of the network (or in the network as a whole).^{5 6}

- Multicast address range: AES67 explicitly calls for support of the administratively scoped multicast address range (239.0.0.0 thru 239.255.255.255) for stream data transport. This requirement is in accordance with best practices as defined in “IANA Guidelines for IPv4 Multicast Address Assignments” (RFC 5771). While ST 2110 does not make any particular reference to the multicast address range to be supported, current practice has shown that use of other multicast address ranges in private or corporate networks appears to be common. Some AES67 devices strictly obey the multicast address requirements of AES67 and won’t allow operation outside that range, which would impose a constraint on some ST 2110 environments. (Or, to put it the other way around, those AES67 devices would not interoperate in ST 2110 environments where multicast address ranges outside the administratively scoped multicast range are intended to be used.)

Other Protocols

While ST 2110 references AES67 in most aspects, it does not require certain protocols to be supported:

- RTCP: While AES67 states that “*Both senders and receivers should transmit RTCP messages as specified in RFC 3550 clause 6*”, ST 2110 states that “*Senders and Receivers may implement RTCP, and Receivers shall tolerate the presence of RTCP.*” The underlying RFC 3550 states that “... [RTCP] SHOULD be used in all environments, but particularly in the IP multicast environment.” While RTCP appears to be useful for any kind of RTP streaming application, neither AES67 nor ST 2110 mandates its support. However, both standards mandate that devices tolerate the presence of RTCP messages.
- Session Initiation Protocol (SIP): While both ST 2110 and AES67 require support of unicast and multicast streaming, ST 2110 does not mandate support of SIP (or any other connection management protocol). SIP is widely used in VoIP applications, and the AES67

⁵ For details regarding SSM and IGMP compatibility mode operation see RFC 3376.

⁶ Note that in the most recent published JT-NM Test Plan (“JT-NM Tested Program Test Plan v.1.3”, 2019-03-24) presence of IGMPv3 communication is a test criteria for ST 2110 receivers.



standardization group found it particularly useful to mandate SIP support in a unicast connection scenario.

Other Functions

- Redundancy: Neither AES67 nor ST 2110 requires support for any redundancy functionality, but ST 2110 mentions that, in case redundancy is offered, it must be compliant to SMPTE ST 2022-7, which describes “Seamless Protection Switching of RTP Datagrams”. However, ST 2110 imposes one important constraint on ST 2022-7: *“Redundant streams shall not use both identical source addresses and identical destination addresses at the same time.”* The reason is: *“... while SMPTE ST 2022-7 allows for RTP streams with duplicated source and destination addresses (on separate physical networks); such a mechanism cannot be represented with SDP, and therefore the use of duplicate source and destination addresses is not supported by this Standard.”*
- Optional channel assignment map (SDP): ST 2110-30 introduces an optional channel-ordering signalization in the related SDP. While the number of channels in a given stream is signaled in the mandatory `a=rtpmap:` parameter (e.g. `a=rtpmap:96 L24/48000/8` for a stream containing eight channels of 24-bit PCM audio at 48 kHz), their labeling or meaning is undefined. The ST 2110 standardization group found it useful to define a channel-order mapping to provide basic information on the channel assignment. If channel ordering is signaled in the SDP, it must follow these guidelines:
 - `a=fmtp:<payload type> channel-order=<convention>.<order>`
 - Example: `a=fmtp:101 channel-order=SMPTE2110.(51,ST)` – this example refers to an 8-channel stream, where the first six channels form a 5.1 channel bundle and the last two channels are a (related) stereo bundle.

Conformance Levels

While ST 2110-30 follows AES67 with respect to mandatory support for payload formats, media clock, and packet setup, it specifies further optional choices, which are defined by *conformance levels*.

ST 2110-30 specifies six conformance levels as outlined below, with level A being the only mandatory level to be supported. Level A is a restatement of the mandatory requirements of AES67:

- A receiver must be able to receive and process streams with:
 - Linear 24-bit PCM encoding
 - 48 kHz sampling frequency (media clock)
 - 1 to 8 channels per stream
 - 1 ms packet time (48 audio samples per channel in each packet)



- Senders must be able to generate corresponding streams, whereas it remains the choice of the manufacturer how many channels per stream are supported, provided the number is within the receiver requirements.

Conformance levels B and C specify shorter packet times (mainly to reduce achievable end-to-end latency) and a higher channel count (to accommodate transport of full MADI connections). These options are described in AES67 but are not required. Conformance levels B and C include conformance level A (devices supporting conformance B or C must also support the requirements of level A):

Level	Supported by the Receiver
A (mandatory)	Reception of 48 kHz streams with 1 to 8 audio channels at packet times of 1 ms
B	Level A + 1 to 8 channels at packet times of 125 μs
C	Level A + 1 to 64 channels at packet times of 125 μs

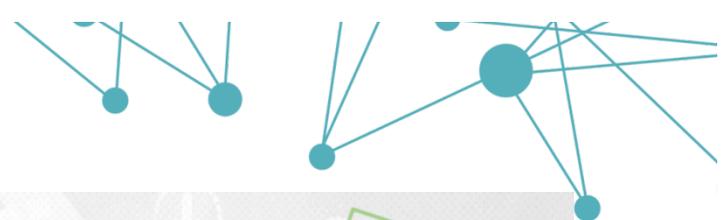
AES67 compliant

Table with requirements of Conformance Levels A-C (48 kHz)

In certain scenarios of the primary application space for ST 2110, audio sample rates of 96 kHz are also common, and further conformance levels have been specified:

- Conformance level AX is aligned with the AES67 recommendation for 96 kHz packet setup:
 - Linear 24-bit PCM encoding
 - 96 kHz sampling frequency (media clock)
 - 1 to 4 channels per stream⁷
 - 1 ms packet time (96 audio samples per channel in each packet)
- Level BX and CX specify shorter packet times and higher channel count, respectively.

⁷ The maximum channel count at 96 kHz / 1 ms packet time for a packet respecting the MTU of 1440 bytes RTP payload would be 5 (and 40 at 125 μs packet time). However, in order to harmonize conformance levels AX and CX with conformance levels A and C, SMPTE has chosen simply to cut the required channel number for AX and CX in half.



Level	Supported by the Receiver
AX	Level A (⇒ 48 kHz) + Reception of 96 kHz streams with 1 to 4 audio channels at packet times of 1 ms
BX	Level B + AX + 1 to 8 channels at packet times of 125 μs
CX	Level C + AX + 1 to 32 channels at packet times of 125 μs

Table with requirements of Conformance Levels AX-CX (96 kHz)

Conformance levels AX, BX, or CX always require compliance with conformance levels A, B or C, respectively.

Of course, manufacturers are free to offer more choices on packet setup beyond the requirements of the specified conformance levels, but there is no guarantee that other devices will be able to receive or process these streams.

Note: SDI allows the carriage of at least 16 embedded audio channels. Senders wishing to remain in compliance with level A can do so by sending groups of channels organized into multiple AES67 streams (i.e., two level A streams with eight channels each).

Conformance to ST 2110

Users of an AES67 implementation that wish to interoperate with an ST 2110 implementation must be prepared to synchronize using the AES67 Media Profile tuned to be compatible with ST 2059 as described in AES R16. No other special accommodations are required for receivers that will receive from ST 2110 transmitters operating in the ST 2110 mandatory conformance level A. AES67 transmitters must use zero offset between media clock and RTP clock to interoperate with ST 2110 receivers.

To interoperate with ST 2110 in other conformance levels, AES67 implementations should include support for 125 μs packet time and up to 64 channels per stream. Support for 96 kHz will open further ST 2110 interoperability opportunities.

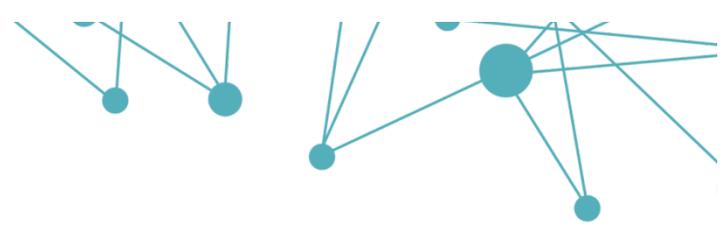
Conformance to AES67

AES67 has a single mandatory conformance level. Users of an ST 2110 implementation that wish to interoperate with any AES67 implementation must be prepared to operate using the IEEE 1588-2008 default profile. Receivers must be able to receive streams with an offset



between media and RTP clock timestamps. Transmitters need only operate in the ST 2110 conformance level A mode.

However, AES67 defines further mandatory requirements that would have to be added to SMPTE ST 2110 implementations in order to be fully AES67-compliant. Full details are available in the PICS section of the AES67-2018 standard.



Summary

The operational principles and mandatory requirements for synchronization, stream transport, packet setup, and signaling defined by SMPTE ST 2110 are aligned to those of AES67. However, SMPTE ST 2110 defines some constraints that would have to be supported by AES67 implementations to be SMPTE ST 2110-compliant. Further optional constraints defined in SMPTE ST 2110-30 are aligned with AES67 optional requirements or can be added to AES67 implementations without breaking AES67 compatibility.

The following diagram illustrates the relationship of SMPTE ST 2110 and AES67 and the most important compatibility constraints:

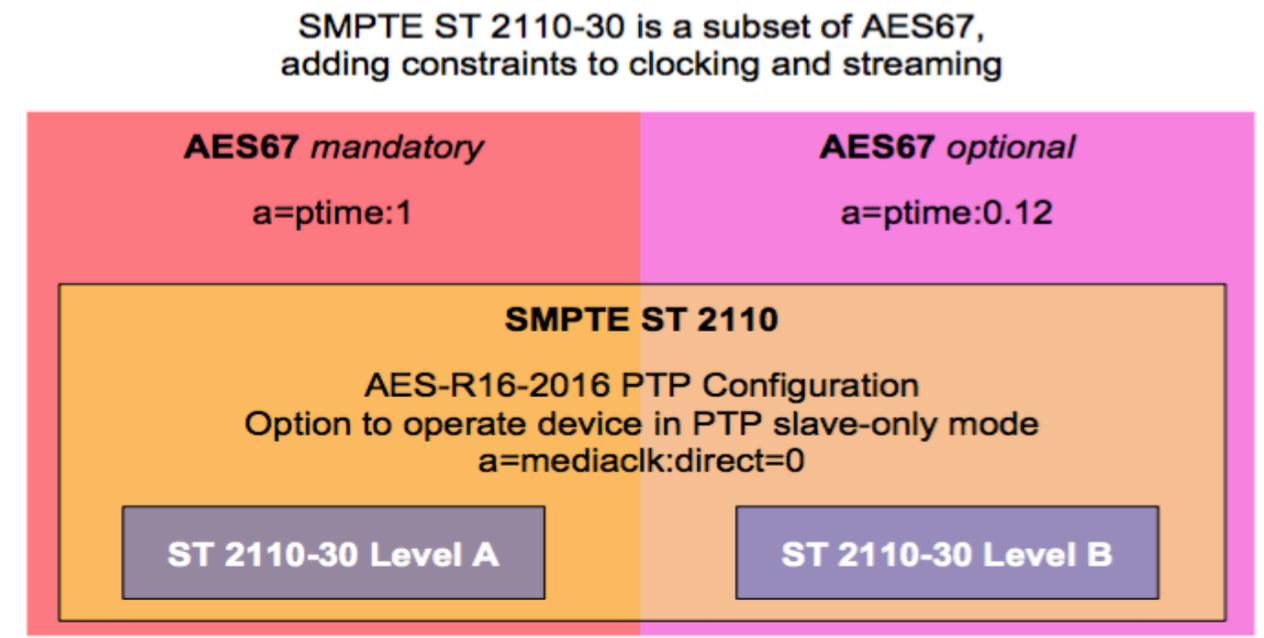
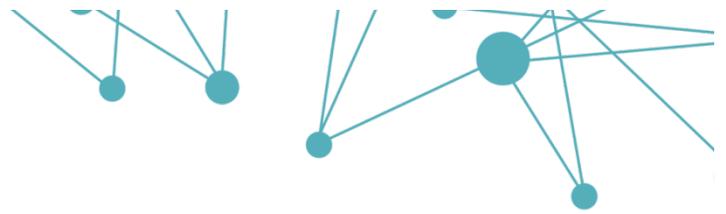


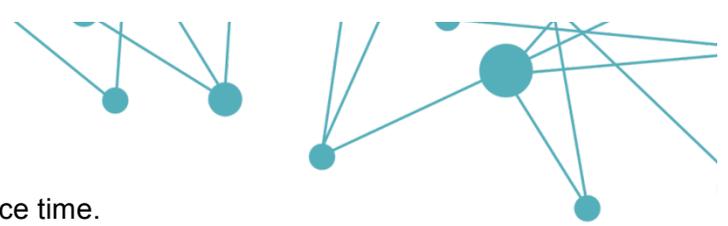
Diagram illustrating SMPTE ST 2110 and AES67 relationship

Consequently, SMPTE ST 2110-30 can be seen as a subset of AES67, provided that the constraints outlined in this document are matched.



Appendix A – Document Links

- **AES67-2018:** <http://www.aes.org/publications/standards/search.cfm?docID=96>
- **SMPTE ST 2110:** <https://www.smpte.org/st-2110>
- **AES-R16-2016:** <http://www.aes.org/publications/standards/search.cfm?docID=105>
- **SMPTE ST 2059-2:** <https://www.smpte.org/news-events/news-releases/smpte-publishes-first-two-parts-standard-enabling-deployment-ptp-timed>
- **IEEE 1588-2008:** <https://standards.ieee.org/findstds/standard/1588-2008.html>
- **IETF RFC 3376:** <https://www.ietf.org/rfc/rfc3376.txt>
- **IETF RFC 3550:** <https://www.ietf.org/rfc/rfc3550.txt>
- **IETF RFC 5771:** <https://www.ietf.org/rfc/rfc5771.txt>
- **IETF RFC 7273:** <https://www.ietf.org/rfc/rfc7273.txt>



- The reference clock provides a common reference time.
- The reference time is periodically transmitted to the devices that maintain a local copy of the reference clock.
- Upon node startup, a media clock is generated as required (i.e., 48 kHz).
- The media clock is aligned to the reference clock.
- At the time a stream is instantiated, it receives a random (but known) offset **R** with respect to the media clock.
- This offset will be constant throughout the stream's lifetime.
- The value of the offset is expressed in media clock cycles and will be conveyed to the receiver via SDP (`a=mediack:direct=<offset>`), so that the receiver can exactly align the incoming RTP stream to the reference clock as required.
- In AES67, **R** may be a random value (following the recommendation of RFC 3550).
- **SMPTE ST 2110 requires this offset value to be "0"**, effectively aligning any RTP stream clock with the related media clock and the common reference clock origin.

Note: The SMPTE ST 2110 zero-offset requirement **does not affect** the ability to align streams with each other, nor does it affect **any alignment precision**. It is merely a provision to allow receivers to easily restart processing of temporarily interrupted streams, avoiding the need for a (new) random offset value to be propagated from sender to receiver.

CONTACT US

Tina Lipscomb - ADMINISTRATOR
23117 39th AVE SE
Bothell, WA 98021 US
+1 425 870 6574
tina@aimsalliance.org