STANDARDS AND INFORMATION DOCUMENTS

Call for comment on DRAFT REVISED AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability

This document was developed by a writing group of the Audio Engineering Society Standards Committee (AESSC) and has been prepared for comment according to AES policies and procedures. It has been brought to the attention of International Electrotechnical Commission Technical Committee 100. Existing international standards relating to the subject of this document were used and referenced throughout its development.

Address comments by E-mail to standards@aes.org, or by mail to the AESSC Secretariat, Audio Engineering Society, PO Box 731, Lake Oswego OR 97034. **Only comments so addressed will be considered**. E-mail is preferred. **Comments that suggest changes must include proposed wording.** Comments shall be restricted to this document only. Send comments to other documents separately. Recipients of this document are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

This document will be approved by the AES after any adverse comment received within **six weeks** of the publication of this call on http://www.aes.org/standards/comments/, **2024-02-14** has been resolved. Any person receiving this call first through the *JAES* distribution may inform the Secretariat immediately of an intention to comment within a month of this distribution.

Because this document is a draft and is subject to change, no portion of it shall be quoted in any publication without the written permission of the AES, and all published references to it must include a prominent warning that the draft will be changed and must not be used as a standard.

[Notes]

AES67-XXXX
Revision of AES67-2018

DRAFT REVISED
AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability

Published by
Audio Engineering Society, Inc.

Abstract

High-performance media networks support professional quality audio (16 bit, 44,1 kHz and higher) with low latencies (less than 10 milliseconds) compatible with live sound reinforcement. The level of network performance needed to meet these requirements is typically available on wired local-area networks and is achievable on enterprise-scale networks. A number of networked audio systems have been developed to support high-performance media networking but until now there were no recommendations for operating these systems in an interoperable manner. This standard provides comprehensive interoperability recommendations in the areas of synchronization, media clock identification, network transport, encoding and streaming, session description and connection management.

Audio Engineering Society Inc. 697 Third Ave., Suite 405, New York, NY 10017, US.

www.aes.org/standards standards@aes.org

Contents

Foreword

This foreword is not part of the AES67-2013 *AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability*

This document was developed in project AES-X192, in the SC-02-12-H task group on high-performance streaming audio-over-IP interoperability, under the leadership of Kevin Gross.

Members of the writing group that contributed to this document in draft are: R. Abraham, J. Amate, M. Barbour, C. Becker-Foss, J. Berryman, M. Bishop, H. Blum, J. Boqvist, T. Borland, N. Borthwick, L. Bradshaw, P. Briscoe, L. Brito, J. Britton, C. Broad, N. Brunsgaard, D. Brutzman, E. Bukont Jr., A. Calvanese, R. Camprodon, M. Carter, A. Cedronius, A. Clark, H. Clarke, M. Coinchon, P. Cyrta, S. de Jaham, P. Demuytere, J. Dibley, A. Dickens, P. Dietrich, C. Dodds, S. Dove, A. Eales, E. Echols, R. Economaki, T. Edwards, A. Elder, L. Ellison, K. Fitzke, J. Fletcher, S. Flock, R. Foss, P. Foulkes, F. Gierlinger, R. Goforth, J. Grant, M. Graubner, D. Gravereaux, S. Gray, H. Hansen, B. Harshbarger, S. Heinzmann, A. Hildebrand, D. Hoch, M. Holtmann, Os. Igumbor, H. Jahne, T. Johansson, M. J. Teener, S. Johnson, L. Jonsson, A. Karamustafaoglu, K. Kearney, P. Keller, J. Koftinoff, P. Koftinoff, D. Koss, S. Langhans, M. Lave, D. Lazecko, S. Ledergerber, C. Lefebvre, J. Lindsay, G. Linis, S. Loehberg, A. Louko, A. Makivirta, J. A. Martinez, A. Mayo, W. McQuay, C. Merienne, A. Metz, J. Meunier, J. Meyer, R. Michl, S. Millan, D. O'Gwynn, N. O'Neill, H. Okai-Tettey, K. Parker, J. Passaniti, J. Peavey, J. Perez, W. Peters, S. Price, S. Pro, M. Quaix, F. Ragenard, R. Rayburn, D. Rice, T. Rohwedder, G. Rosenboom, M. Saito, M. S. Carreres, J. Sauter, R. Schoonbroodt, V. Schueppel, P. Schwizer, S. Scott, G. Shay, T. Shuttleworth, D. Silver, J. Simpson, M. Sims, A. Smimite, J. Snow, T. Staros, P. Stevens, J. M. Strawn, N. Sturmel, M. Szlapka, T. Thompson, P. Treleaven, A. Trevena, S. Turner, R. van der Zalm, B. van Kempen, I. Vysick, J. P. Waddell, Y. Wang, P. Warrington, H. Weibel, A. Williams, M. Willsher, A. Witham, J. Wood, J. A. Yeary, J. Yoshio, P. Yurt, U. Zanghieri, D. Zimmermann, R. Zwiebel.

This document was edited by Kevin Gross.

Richard Foss
Chair, working group SC-02-12, 2013-07-18

Foreword to second edition, 2015

This revision includes minor changes identified during 'plugfest' testing in October 2014 and was developed in task group SC-02-12-M. It includes updated references to RFC 7273, and clarifications in 6.3, 8.1, and 8.5.

Members of the task group that contributed to this revision in draft are: R. Abraham, J. Amate, L. Andrieu, R. Barbieri-Carrera, M. Barbour, C. Becker-Foss, F. Bergholtz, J.A. Berryman, J. Boqvist, J. Breitlow, D. Brulhart, B.J. Buchalter, A. Calvanese, C. Cellier, K. Cleary, M. Danielson, I. Dennis, C. Diehl, S. Dove, T. Duffy, J. Dunn, J. Evanson, J. Fletcher, S. Flock, R. Foss, J. Freyberger, G. Gebler, J. Grant, K. Gross, S. Heinzmann, E. Heurtel, A. Hildebrand, F. Hoyer, T. Johansson, L. Jonsson, B. Klinkradt, J. Koftinoff, S. Langhans, S. Ledergerber, C. Lefebvre, S. Leschka, Jo. Lindsay, G. Linis, A. Lucas, A. Makivirta, A. Metz, J. Meyer, R. Michl, G. Passador, J. Peavey, M. Quaix, C.R. Reed, T. Rohwedder, M. Schuchert, V. Schueppel, G. Shay, P. Stevens, N. Sturmel, M. Szlapka, P. Treleaven, A. van den Broek, B. van Kempen, J.P. Waddell, A. Williams, K. Wu, N. Yamaguchi, R. Zwiebel.

The task group was led by Kevin Gross.

Richard Foss
Chair, working group SC-02-12, 2015-07-27

Foreword to third edition, 2018

This revision contains clarifications and minor corrections and adds a Protocol Implementation Conformance Statement (PICS) as Annex G. A new sender keep-alive recommendation, has been added to clause 6.3. Minor clarifications and corrections include a specification of MTU requirements in the presence of allowed (but not recommended) additional information in the RTP header and correcting SDP examples in clause 8.5 to match an erratum issued by the IETF on RFC 7273. Corrected domainNumber range to conform to IEEE 1588-2008 requirements. These revisions were developed in task group SC-02-12-M.

Members of the task group that contributed to this revision in draft are: R. Abraham, J. Amate, N. Amaya, L. Andrieu, R. Barbieri-Carrera, M. Barbour, C. Becker-Foss, R. Bell, F. Bergholtz, J. A. Berryman, M. Blackburn, A. Boenninghoff, J. Boqvist, L. Bradshaw, D. Breithaupt, J. Breitlow, D. Brulhart, B. J. Buchalter, R. Bugg, R. Cabot, A. Caceres, A. Calvanese, C. Cellier, R. Charlesworth, K. Cleary, B. Cochran, A. Cooper, M. Danielson, T. deBrouwer, I. Dennis, C. Diehl, G. Diehl, S. Dove, T. Duffy, J. Dunn, J. Evanson, J. Fletcher, S. Flock, R. Foss, J. Freyberger, G. Gauthier, G. Gebler, J. Grant, E. Grossman, S. Heinzmann, M. Henke, M. Henry, E. Heurtel, A. Hildebrand, T. Holton, A. Holzinger, F. Hoyer, L. Huapaya, H. Jesuiter, T. Johansson, L. Jonsson, N. Keroe, A. Kitagawa, B. Klinkradt, J. Koftinoff, S. Langhans, J. Laundon, S. Ledergerber, C. Lefebvre, S. Leschka, E. Lestage, J. Lindsay, G. Linis, K. Lyver, A. Makivirta, C. Mannett, S. Mertens, A. Metz, J. Meyer, R. Michl, E. Miehs, T. Neuhaus, C. Nighman, N. Nzoyem, B. Olson, M. Overton, O. Palm, K. Parker, G. Passador, J. Peavey, J. Pruitt, M. Quaix, C. R. Reed, T. Rohwedder, A. Santos, G. Scherling, M. Schettke, M. Schuchert, V. Schueppel, S. Scott, A. Sharifan, G. Shay, M. Smaak, K. Soma, P. Stevens, N. Sturmel, M. Szlapka, J. Tikkanen, P. Treleaven, A. van den Broek, B. van Kempen, S. van Tienen, J. P. Waddell, P. Walker, D. Walters, C. Ware, P. Warrington, E. Wetzell, L. Whitcomb, A. Williams, K. Wu, N. Yamaguchi, M. Yonge, R. Zwiebel.

The task group was led by Kevin Gross.

The PICS (annex G) was edited by Gints Linis

Morten Lave
Chair, working group SC-02-12, 2017-12-11

Foreword to fourth edition, 2023

This revision contains multiple updates and corrections, categorized as indicated later in this foreword.

Authors of this revision have intentionally minimized the amount of text changes in the body of the standard. Those engaged in implementation of the standard are encouraged to also consult the PICS (Annex G) as this presents testable behavioral requirements that, in many cases, provide clarification of the normative statements found in the body of the standard.

The changes have resulted in deprecating or moving content of existing PICS statements, as well as inserting one new PICS section and a number of new PICS statements.

The existing PICS statements have retained their statement numbers within sections. New statements received new, formerly non-existing statement numbers. Where the contents of a PICS statement were moved between sections, the previous statement number is marked as deprecated, and a new statement number is created in the new position.

Some of the PICS sections have been renumbered.

Detailed list of changes by categories follows.

Changes directly affecting interoperability:

- 4.0, 8.2: Defined handling of PTP version IEEE 1588-2019. This is a new requirement added in this revision of the standard.
- 5: Added a discussion of using zero RTP offset to achieve interoperability with ST 2110-30:2017.
- 6.1.3 Multicast address range: Previous versions of this standard required use of the administratively scoped multicast address range, which puts all other multicast addresses outside the scope of this standard. It narrowed usability of AES67 for SSM applications and limited interoperability with ST 2110. To resolve this issue, the requirement to use the administratively scoped multicast addresses is replaced with a requirement to support them. Thus, using other addresses, although not necessarily supported by all conformant devices, becomes a valid AES67 application. Informative details are provided about usage of multicast addresses under ST 2110.
- 6.2: In previous versions, the requirement for senders to use the default DSCP settings in absence of a management interface was not required by the language of the standard. This has been clarified as a requirement.
- 6.3: Defined an exact range of transport port numbers to be supported by senders and receivers. This requirement was largely undefined before, only talking about “other or additional ports”, without any further details.
- 6.3, 7.2.0, G.3.4.3 (PICS): Clarified RFC 3551 requirements – explicit overrides are defined for silence suppression, channel numbering, ordering, content mapping, and packetization requirements. In particular, a new explicit requirement is added in this revision of the standard, which disallows use of silence suppression by senders.
- 7., Annex G: Reworked streaming interoperability criteria. The criteria specification method is changed from individual stream attributes to attribute vectors called *stream modes*. The concept of stream mode is introduced in 7.0, but most of the essential changes are applied to Annex G - PICS. Section 7 has been left mostly unchanged in this respect – it still provides a general description. Annex G specifies conformance criteria.

  As a result, a number of the previously existing PICS entries are deprecated, and new entries are created in G.3.5, to follow the new qualification criteria as defined in G.4. An instruction is provided in G.2 for documenting stream mode capability declarations in the PICS proforma.

These changes are aimed at resolving ambiguities and filling requirement gaps in previous versions of the standard. They are not aimed at changing existing interoperability requirements. The new specification method can, however, reveal possible misinterpretations of this standard's intentions in existing implementations.

- 7.5: Interoperability exception explained – when a sender of the lower timing accuracy class is connected with a receiver implementing only the minimally required buffering capacity, reliable reception of all streams cannot be guaranteed. An interoperability matrix summarizing the issue is provided.
- 8.2: Added guidance for handling of traceable time references. This option is defined in RFC 7273, but was not addressed in previous versions of AES67. ST 2110-30:2017 requires that traceable clock references are labeled so, therefore lack of this mechanism in AES67 devices can create a potential for interoperability problems. Changes applied:
    - Added an example of SDP signaling of the "traceable" property of a clock reference.
    - Stream connection setup rules are updated to reflect cases involving traceable references.
- 8.3: Added a discussion of an ST 2110-30:2017 requirement to specify mediaclk always at the media level of SDP. Not doing so could affect interoperability with ST 2110-30:2017.
- 10.1: Relaxed the strong requirement ("shall") to support SIP for unicast connections – changed to a recommendation ("should").
- 10.2 (was in 6.1): IGMPv3 support level elevated from "may" to "should"; behavior related to IGMPv3 is defined.
- 10.2 (was in 6.1): removed requirement for senders to use IGMP to request their own streams. See the new Annex F for a discussion of the motivation for this change.

Normative language improvements:

These editions include clarification of ambiguous or implied requirements, filling omissions, correcting unclear or misused terms or normative expressions. No changes of requirements are intended by this category of editions. In particular:

- Corrected improper or inconsistent use of media streaming terms, to achieve better alignment with the referenced RFC documents. Specifically:
    - The term "payload format" was sometimes used interchangeably with "encoding format" and "payload type". Use of these terms is corrected in this version.
    - The terms "sampling rate", "sampling frequency, and "sample frequency" are interchangeable, and they all have been used variously throughout the standard. Those cases are brought to uniformity now, using a single term "sampling rate".
- 0.1: Corrected improper use of normative language – verb "should" was used to express a generally valid but practically untestable goal. The "should" is replaced to clarify this sentence is informative.
- 4.0 The previous revisions of this standard used a term "common clock", which is inconsistent with PTP, where common time is shared between devices. "Common time" terminology is used in this revision.
- 4.1, 4.2: "IP network" and "1588 network" terms are replaced with "ordinary IP network" and "IP network with IEEE 1588-2008" respectively; "standard IP network" replaced with "ordinary IP network"; "use" (a user's act) replaced with "support" (device's property); "achieved" replaced with "achievable".
- 6.1.1: Added clause clarifying that senders are allowed to choose whether they support only multicast, only unicast, or both.

- 6.1.2: Corrected the terms describing ICMP signaling related to packet fragmentation and MTU size.
- 6.2 Table1: Clarified mapping of IEEE1588-2008 traffic types to QoS classes.
- 6.2 Corrected an unclear expression “shall make no assumptions ...”. It is replaced with “shall not depend on”, a more actionable statement.
- 7.0: Added clause to clarify devices may support receiving of audio streams, or sending, or both.
- 7.1 Clarified the scope of this standard with respect to combinations of sampling rates and payload formats.
- 7.2.2 Details added to the table listing the required and recommended packet times.
- 7.3 Clarified channel count requirements.
- 7.4 Clarified requirement for the link offset to be retrievable from the device. Previous versions used an expression referring to communication of the link offset within the device.
- 8.1 Term “milliseconds decimal” replaced with “milliseconds fractional part”.
- 8.1 Clarified the requirement to support alternative packet time signaling beyond the examples given in Table 4. Specifically, besides other values, valid alternative representations must be supported too, such as containing unnecessary precision digits or a decimal point not followed by the fractional part.
- 8.3: Added a discussion of specifying mediaclk at the media level of SDP as a precondition for interoperability with ST 2110-30:2017.
- 8.4 Reference redirected to section 7.1. In previous versions it pointed to table 2, which does not provide a complete specification.
- 10.1 Added explicit language allowing the use of other protocols or management interface for unicast connection management.
- Annex G (PICS): n/a option is now provided consistently, where relevant. In the previous version it was often missing.
- G.3.5 (PICS) – multiple statements: Streaming interoperability requirements reworked from individual attributes to attribute vectors called “stream modes”.
- G.3.6.1 (PICS) - 8.0-1: SDP support
    - Clarified applicability of SDP output generation and input interpretation requirements to senders and receivers, respectively.
    - Added a discussion of using off-the-shelf test suites as a tool for verification of SDP document handling.
    - Added a table row for information about the test tools used for verification of SDP document handling and the test configuration.

Informative and readability improvements:

These editions include partial restructuring of some sections, correcting inaccurate terminology and descriptions, general editorial corrections. In particular:

- 3. All definitions formerly divided between section 3 “Definitions” and Annex F “Glossary” are consolidated in section 3. New terms are added, and some definitions are reworked for better clarity.
- 5. Clauses refactored and language revised to improve general readability and provide better definitions of the related terms.

- 6.1, 7.6, 10.2 – content partially reordered and moved between sections to improve readability and consolidate sets of closely related requirements. Specifically:
    - Section 6.1 “Network layer” is now additionally structured by adding 3rd-level headings.
    - Discussion of multicast transport was formerly split between section 6.1 “Network layer” and section 7 “Encoding and streaming”. It is now all consolidated in section 6.1 “Network layer”.
    - All discussion of using IGMP for multicast connection management is moved from 6.1 “Network layer” and consolidated in 10.2 “Multicast connections”.
- E.6. Added a brief summary of RAVENNA discovery.
- New Annex F explains the context of using IGMP by senders to request their own streams. Such requirement was originally included in this standard, and it is removed in this revision.
- “Master” and “slave” terminology replaced by “timeTransmitter” and “timeReceiver” respectively, according to recommendations given in IEEE 1588g-2022.
- Multiple smaller editorial changes throughout the standard.

Maintenance of external document references:

- RFC 4566 has been revised and declared obsolete by IETF. Reference updated to RFC 8866.
- In previous versions, a number of documents have been misplaced between the lists of normative references and bibliography. Such misplacements have been corrected in this version as follows:
    - RFC 792 and RFC 4028 are moved from bibliography to normative references.
    - RFC 2974 is moved from normative references to bibliography.
- In previous versions, a number of documents appearing on the list of normative references or bibliography were not properly mentioned in the main text of this standard. Such omissions have been corrected in this version.
- A number of new document entries are added to the “Bibliography” section. Some were referenced in the text, and some are entirely new references. Newly added are a number of RFC documents by IETF, ST 2110-10 and ST 2110-30 by SMPTE, and an amendment to IEEE1588.
- Reference to IEEE 802.1Q-2011 is dated and is presented as such consistently in this version of AES67. Other references to IEEE 802.1 series documents are effectively undated, although release dates of these documents occasionally appear indicated in the text of this standard. They remain relevant and reflect the status at the moment of publishing the original version of AES67.

This revision was developed in task group SC-02-12-M.

Members of the task group that contributed to this revision in draft are: C. Becker-Foss, R. Cabot, O. Chambin, J. Fletcher, J. Grant, K. Gross, E. Heurtel, A. Hildebrand, T. Holton, I. Kostiukevych, A. Kuzub, G. Linis, S. Scott, P. Stevens, N. Sturmel, P. Treleaven, P. Tschiemer, R. Wadge, H. Wataru, L. Whitcomb, S. Zaho.

The task group was led by Kevin Gross.

This revision was edited by Gints Linis.

Morten Lave
Chair, working group SC-02-12, 2023-12-04

Note on normative language

In AES standards documents, sentences containing the verb “shall” are requirements for compliance with the document. Sentences containing the verb “should” are strong suggestions (recommendations). Sentences giving permission use the verb “may”. Sentences expressing a possibility use the verb “can”.

DRAFT REVISED

AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability

0 Introduction

0.1 General

High-performance media networks support professional quality audio (16 bit, 44,1 kHz and higher) with low latencies (less than 10 ms) compatible with live sound reinforcement. The level of network performance needed to meet these requirements is typically available on wired local-area networks and is achievable on enterprise-scale networks, but is generally not available on wide-area networks or the public internet.

The most recent generation of these media networks use a diversity of proprietary and standard protocols. Despite a common basis in Internet Protocol, the systems do not interoperate.

This standard provides specific recommendations for interoperability. The standard focuses on defining how existing protocols are used to create an interoperable system. No new protocols have been developed to achieve this.

The standard is expected to be useful for commercial audio applications including fixed and touring live sound reinforcement. It is also expected to be useful for distribution within broadcast, music production and post-production facilities.

This standard depends on established network protocols (see clause 2). These protocols can include additional options that are not required by this standard. Robust implementations of AES67 will tolerate these additional options.

Any behavior details not described in the main part of this standard are in some cases clarified in Annex G (PICS), by means of the respective evaluation criteria.

0.2 Patents

The Audio Engineering Society draws attention to the fact that it is claimed that compliance with this AES standard or information document can involve the use of a patent.

The AES holds no position concerning the evidence, validity and scope of this patent right.

The holder of this patent right has assured the AES that it is willing to negotiate licenses under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of this patent right is archived with the AES and listed on the public AES web site.

Information can be obtained from.

Audinate Pty Ltd.
PO Box 855
Ultimo, NSW 2007
Australia

Attention is drawn to the possibility that some of the elements of this AES standard or information document can be the subject of patent rights other than those identified above. AES shall not be held responsible for identifying any or all such patent rights.

1 Scope

This standard defines an interoperability mode for synchronization, encoding, transport, and connection management of high-performance audio over networks based on the Internet Protocol. For the purposes of this standard, high-performance audio refers to audio with full bandwidth and low noise. These requirements imply linear PCM coding with a sampling rate of 44,1 kHz and higher and resolution of 16 bits and higher. High performance also implies a low-latency capability compatible with live sound applications. This standard considers latency performance of 10 milliseconds or less.

2 Normative references

The following referenced documents are indispensable for the application of this standard. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

In case of a conflict between clauses of this standard and a referenced document, the clauses of this standard take precedence.

The document identification and versioning approach adopted by IETF for RFCs, where each update or addendum is identified as a new document, can produce multiple secondary references related to a single root document. This standard lists only the root documents in force at the moment of writing. For secondary references, readers are encouraged to visit the IETF document library and follow the links indicated in the respective RFC document headers. Secondary references to RFC documents, which are published after this revision of AES67, should be used judiciously. In case of conflicts, the status at the moment of publishing this revision of AES67 takes precedence.

AES11, *AES recommended practice for digital audio engineering - Synchronization of digital audio equipment in studio operations*; Audio Engineering Society, New York, NY., US

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, July 2008, Institute of Electrical and Electronics Engineers (IEEE), US

RFC 768, *User Datagram Protocol*", Internet Engineering Task Force

RFC 791, *Internet Protocol*, Internet Engineering Task Force

RFC 792, *Internet Control Message Protocol*, Internet Engineering Task Force

RFC 1112, *Host Extensions for IP Multicasting*, Internet Engineering Task Force

RFC 2236, *Internet Group Management Protocol, Version 2*, Internet Engineering Task Force

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, Internet Engineering Task Force

RFC 2616, *Hypertext Transfer Protocol - HTTP/1.1*, Internet Engineering Task Force

RFC 3190, *RTP Payload Format for 12-bit DAT Audio and 20- and 24-bit Linear Sampled Audio*, Internet Engineering Task Force

RFC 3261, *SIP: Session Initiation Protocol*, Internet Engineering Task Force

RFC 3264, *An Offer/Answer Model with the Session Description Protocol (SDP)*, Internet Engineering Task Force

RFC 3376, *Internet Group Management Protocol, Version 3*, Internet Engineering Task Force

RFC 3550, *RTP: A Transport Protocol for Real-Time Applications*, Internet Engineering Task Force

RFC 3551, *RTP Profile for Audio and Video Conferences with Minimal Control*, Internet Engineering Task Force

RFC 4028, *Session Timers in the Session Initiation Protocol (SIP)*, Internet Engineering Task Force

RFC 5939, *Session Description Protocol (SDP) Capability Negotiation*, Internet Engineering Task Force

RFC 7273, *RTP Clock Source Signalling*, Internet Engineering Task Force

RFC 8866, *SDP: Session Description Protocol*, Internet Engineering Task Force

3 Definitions and abbreviations

For the purposes of this standard, the following terms, definitions, and abbreviations apply.

3.1
Audio stream
See RTP stream.

3.2
Audio Video Bridging
AVB
Describes enhanced Ethernet networks specified in IEEE 802.1BA, IEEE 802.1Q-2011 and IEEE 802.1AS.

3.3
Bonjour
Bonjour is Apple's implementation of Zero Configuration Networking (Zeroconf), a group of technologies that includes service discovery, address assignment, and hostname resolution.

3.4
Best timeTransmitter clock algorithm
BTCA
(Formerly known as best master clock algorithm, BMCA)
An algorithm operated in PTP clocks to determine which clock is the best. The best clock becomes the timeTransmitter on the respective network segment. See IEEE 1588-2008.

3.5
Boundary Clock
A clock that has multiple Precision Time Protocol (PTP) ports in a domain and maintains the timescale used in the domain. It can serve as the source of time, that is, be a timeTransmitter clock; and can synchronize to another clock, that is, be a timeReceiver clock. See IEEE 1588-2008.

3.6
Byte
A unit comprising 8 bits of data. Over IP networks, data is transported in units of bytes.

3.7
CSRC
The contributing source (CSRC) is the source of a stream of RTP packets that has contributed to the combined stream produced by an RTP mixer. See RFC 3550.

3.8
DiffServ
Differentiated services (DiffServ) is a system for classifying traffic and providing quality of service (QoS) on an IP network. See RFC 2474.

3.9
Digital Audio Reference Signal
DARS
An audio clock signal defined in AES11.

3.10
Domain Name System
DNS
A hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network. See RFC 2474.

3.11
DNS Service Discovery
DNS-SD

A way of using standard DNS programming interfaces, servers, and packet formats to browse the network for services. It is one of the mechanisms used by Bonjour. See http://www.dns-sd.org/.

3.12
DSCP
The differentiated services code point (DSCP) is a 6-bit field in the IP packet header that is used for classification purposes. DSCP is part of the differentiated services architecture. See RFC 2474.

3.13
Encoding
The means in which audio is digitized and formatted into the sequence of packets that constitutes a stream.

3.14
Encoding format
See payload format.

3.15
End-to-end Transparent Clock
A transparent clock that supports the use of the end-to-end delay measurement mechanism between timeReceiver clocks and the timeTransmitter clock. See IEEE 1588-2008.

3.16
Ethernet
A physical and data link layer family of computer networking technologies for local area networks (LANs) standardized by IEEE.

3.17
EUI-64
A 64-bit globally unique identifier defined by IEEE, formed by combining a registered 24 or 36-bit company identifier and a company-unique device identifier. The EUI-64 is similar to the EUI-48 which is used to assign Ethernet media access control (MAC) addresses.

3.18
Expedited Forwarding
EF
One of the classifications used by DiffServ with the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services. EF traffic is often given strict priority queuing above all other traffic classes.

3.19
Grandmaster
The source of synchronization for clock distribution via PTP. The grandmaster is a network device and is identified by an EUI-64.

3.20
Grandmaster identifier
GMID
An EUI-64 used in IEEE 1588-2008 and IEEE 802.1AS synchronization standards to uniquely identify the grandmaster serving a synchronization domain.

3.21
IEEE
Institute of Electrical and Electronics Engineers is a professional association dedicated to advancing technological innovation and excellence. The IEEE publishes a wide range of communications standards.

3.22
IETF
Internet Engineering Task Force is the volunteer standards-developing organization responsible for the Internet Protocol suite.

3.23
IGMP
Internet Group Management Protocol (IGMP) is a communications protocol used by hosts to report their multicast group memberships to IPv4 routers. See RFC 2236, RFC 3376.

3.24
Internet Protocol
IP
Network layer protocol commonly used to transport data on Ethernet networks. See RFC 791, RFC 8200.

3.25
IPv4
Internet Protocol version 4 is the first and the most widely deployed version of the Internet Protocol. See RFC 791.

3.26
IPv6
Internet Protocol version 6 is the most recent revision of the Internet Protocol and is intended to replace IPv4 eventually. See RFC 8200.

3.27
Link offset
Link offset specifies the amount of time media spends on the network and in buffers at the sender and receiver as illustrated in figure 1. Link offset is also known as *network latency* or *playout delay*.

3.28
Local-area network
LAN
A computer network that interconnects computers in a limited area such as a home, school, computer laboratory, or office building.

3.29
Management interface
Any hardware and/or software means allowing the device to be configured, including but not limited to a hardware front panel, web-based graphical user interface, or remote control protocol.

3.30
Media clock
The clock that controls the rate at which senders and receivers pass audio samples to and from digital media streams, respectively. The media clock for audio streams reads in units of samples. The relationship between media clock and network clock is defined in section 5.

3.31
Media packet
One of the data packets carrying media data as part of a media stream. A media packet contains one or more samples for one or more audio channels.

3.32
Media stream
See RTP stream.

3.33
Maximum transmission unit
MTU
The size of the IP packet, measured in bytes, that can be transferred using a specific data link connection. The MTU for an Ethernet data link is 1500 bytes.

3.34
Multicast DNS
mDNS
A way of using the Domain Name System (DNS) programming interfaces and packet formats, without configuring a conventional DNS server. It is one of the mechanisms used by Bonjour.

3.35
Network clock
Source of the time delivered by the network synchronization mechanism defined in 4. The network clock reads in units of seconds.

3.36
Network layer
The network layer is layer 3 of the OSI model and is responsible for packet forwarding and routing of variable-length data sequences from a source to a destination.

3.37
Ordinary Clock
A clock that has a single Precision Time Protocol (PTP) port in a domain and maintains the timescale used in the domain. It can serve as a source of time, that is, be a timeTransmitter clock; or can synchronize to another clock, that is, be a timeReceiver clock. See IEEE 1588-2008.

3.38
OSI model
The Open Systems Interconnection Model published by International Telecommunication Union (ITU) characterizes and standardizes the functions of a communications system in terms of abstraction layers.

3.39
Packet time
The real-time duration of the media signal represented by data contained in a media packet. For example, a packet containing 12 samples of 48 kHz audio has a packet time of 12 ÷ 48 kHz = 250 microseconds.

3.40
Payload format
Audio sample encoding format used in RTP streams.

> NOTE In various documents, the terms “payload format” and “encoding format” are used inconsistently and sometimes interchangeably. Encountered meanings range from numeric representation of a single-channel audio sample value to a complete media stream format description including the packetization aspect. In this standard, the primary term is “payload format” as defined in 7.1, but “encoding format” remains being occasionally used in the text in the same meaning.

3.41
Payload type
A 7-bit RTP header field pointing to a payload format specification for this RTP stream. Payload format specifications can be statically mapped to payload type values by RFC 3551 (static payload types, values 0 to 95 decimal) or mapping can be defined in the SDP description of the current session (dynamic payload types, values 96 to 127 decimal).

3.42
Peer-to-peer Transparent Clock
A transparent clock that, in addition to providing Precision Time Protocol (PTP) event transit time information, also provides corrections for the propagation delay of the link connected to the port receiving the PTP event message. In the presence of peer-to-peer transparent clocks, delay measurements between timeReceiver clocks and the timeTransmitter clock are performed using the peer-to-peer delay measurement mechanism. See IEEE 1588-2008.

3.43
Precision time protocol
PTP
Time distribution protocol standardized in IEEE 1588-2002, IEEE 1588-2008 and IEEE 802.1AS-2011. In this standard, "IEEE 1588-2008" means the PTP version IEEE 1588-2008 or a functionally equivalent subset of IEEE 1588-2019.

3.44
Protocol Implementation Conformance Statement
PICS
A document indicating support of specific requirements by a given implementation.

3.45
Quality of service
QoS
Describes a system for classifying, marking and delivering traffic across a network in accordance with its performance requirements.

3.46
Receiver
A network device with ability to receive at least one media stream from the network.

3.47
Request for Comment
RFC
Request for Comments are memorandums published by the IETF relevant for the working of the Internet and Internet-connected systems. RFCs are referenced by number. RFC 791, for example, defines the Internet Protocol version 4 (IPv4).

3.48
RTCP
A companion protocol of the Real-time Transport Protocol (RTP), providing statistics and control information for RTP media packets. See RFC 3550.

3.49
Real-time Transport Protocol
RTP
RTP is defined in RFC 3550 and provides a means for applications to organize, mark and transport their media packets using UDP/IP networking.

3.50
RTP clock
RTP clock is a 32-bit representation of the least-significant part of the full media clock value. Individual RTP clocks each operate with a constant offset with respect to the media clock.

3.51
RTP timestamp
RTP timestamp is the value of the RTP clock at the sampling instant of the first audio sample included in the packet. RTP timestamps are carried in RTP packets containing the related stream data. See RFC 3550, RFC 3551.

3.52
RTP session
An RTP session is a media connection between sender and receiver. RTP sessions can be unicast or multicast. In teleconferencing RTP applications, multicast sessions can have multiple senders and receivers. However, under this standard, a session is allowed to have only one sender (see 7.6).

3.53
RTP stream
An RTP stream is a sequence of RTP packets with media data sent at regular interval. A stream can contain multiple channels. There can be multiple RTP streams per RTP session.

3.54
Session Announcement Protocol
SAP
An experimental protocol for announcing RTP sessions. SAP is defined in RFC 2974.

3.55
Session Description Protocol
SDP
A format for describing RTP sessions and their operating parameters including network addressing, encoding format and other metadata. SDP is defined in RFC 8866.

3.56
Sender
A network device with ability to source at least one media stream onto the network.

3.57
Session
See RTP session.

3.58
Session Initiation Protocol
SIP
A telecommunications connection management protocol defined in RFC 3261.

3.59
SIP URI
A SIP URI is a URI used by SIP to identify user agents. SIP URI take the form sip:<user>@<domain> or sips:<user>@<domain>. See 10.1.1 and RFC 3261.

3.60
Stream
See RTP stream.

3.61
Stream mode
Defined by a specific combination of the audio sampling rate, payload format, stream channel count, and packet time values of audio data carried by an RTP stream.

3.62
Streaming
The process of sending and receiving sequences of packets that constitute media streams, progressively in real time, at a rate matching the average rate of media data generation at the sender and playback at the receiver.

3.63
Stream reservation protocol
SRP
The AVB admission control protocol defined in IEEE 802.1Q-2011 clause 35.

3.64
TimeReceiver clock
A clock that is synchronized to a timeTransmitter clock (the provider of time) within an environment that uses the Precision Time Protocol (PTP). A timeReceiver can, in turn, be a timeTransmitter to another clock and can simultaneously be a boundary clock.

3.65
TimeTransmitter clock
A clock that serves as the provider of time within an environment that uses the Precision Time Protocol (PTP). TimeReceiver clocks are synchronized to the timeTransmitter clock operating on the respective network segment.

3.66
Traceable time (reference)
A clock is considered to provide traceable time if it can be proven to be synchronized to International Atomic Time (TAI). See RFC 7273.

3.67
Transport Layer Security
TLS
A cryptographic protocol for secure communication over IP networks.

3.68
Transparent clock
A switch or router device that measures the time taken for a Precision Time Protocol (PTP) event message to transit the device and provides this information to clocks receiving this PTP event message. See IEEE 1588-2008. See also definitions in this standard: End-to-end Transparent Clock; Peer-to-peer Transparent Clock.

3.69
Transport layer
The network layer is layer 4 of the OSI model and provides end-to-end communication services for network applications.

3.70
User datagram protocol
UDP
Constitutes a simple transport layer for the IP network layer. See RFC 768.

3.71
Uniform resource identifier
URI
An identifier for a network resource. An identification URI enables interaction with the resource over a network. See RFC 3986.

3.72
User agent
A SIP endpoint device such as a VoIP telephone.

3.73
UTF-8
A popular variable-length character encoding scheme that supports international character sets but is also backwards compatible with ASCII. UTF-8 is defined in RFC 3629.

3.74
Virtual LAN
VLAN
A single layer-2 network can be partitioned to create multiple distinct broadcast domains, which are mutually isolated so that packets can only pass between them via one or more routers; such a domain is referred to as a Virtual Local Area Network. See IEEE 802.1Q-2011.

3.75
Voice over IP
VoIP
The communication protocols, technologies, methodologies, and transmission techniques involved in the delivery of voice communications and multimedia sessions over Internet Protocol (IP) networks, such as the Internet.

4 Synchronization

4.0 General

The ability for network participants to share high-precision common time distinguishes high-performance media streaming from its lower-performance brethren such as Internet radio and IP telephony. Using a common time, receivers anywhere on the network can synchronize their playback with one another. A common time allows for a fixed and determinable latency between sender and receiver. A common time ensures that all streams are sampled and presented at exactly the same rate. Streams running at the same rate can be readily combined in receivers. This property is critical for efficient implementation of networked audio devices such as digital mixing consoles.

Synchronization of time shall be achieved using IEEE 1588-2008 Precision Time Protocol (PTP) or a functionally equivalent subset of IEEE 1588-2019. In this standard, both together are referred to as "IEEE 1588-2008".

IEEE 1588-2008 is profiled for use in different synchronization applications. A profile describes protocol attributes, available options and required device performance. IEEE 1588-2008 specifies default profiles for delay request-response (IEEE 1588-2008 annex J.3) and peer-to-peer (IEEE 1588-2008 annex J.4) mechanisms.

Devices, with the exception of certain AVB devices (see below), shall support the IEEE 1588-2008 default profiles. Devices supporting the default profiles shall use IPv4 encapsulation as described in IEEE 1588-2008 annex D.

As a single exception, devices that use the AVB synchronization mechanism described in 4.3, and that need to be connected to an AVB network in order to accomplish media streaming, are not required to implement the IEEE 1588-2008 default profiles.

4.1 Synchronization of ordinary IP networks

Devices intended to operate on ordinary (PTP-unaware) IP networks should support the media profile defined in annex A to ensure adequate performance for all applications. Devices may use the default profiles on ordinary IP networks at a cost of lock time and accuracy possibly being degraded.

4.2 Synchronization of IP networks with IEEE 1588-2008

On networks built using switches with IEEE 1588-2008 capabilities (boundary clocks or transparent clocks), adequate performance for audio transport is achievable using the IEEE 1588-2008 default profiles. Due to performance constraints, some IEEE 1588-2008 network equipment can be unable to support the media profile, therefore, on these networks, always choosing one of the default profiles is preferred.

4.3 Synchronization of AVB networks

Enhanced Ethernet networks, as specified in IEEE 802.1Q-2011, commonly known as Audio Video Bridging (AVB), deliver synchronization using IEEE 802.1AS. IEEE 802.1AS defines an IEEE 1588-2008 profile. AVB networks can use their native IEEE 802.1AS synchronization profile in preference to the default profiles or media profile. Methods for building heterogeneous synchronization networks using IEEE 1588-2008 and IEEE 802.1AS-2011 are described in annex D.

5 Media clock and RTP clock

The media clock controls the rate at which senders and receivers pass audio samples to and from digital media streams, respectively. The media clock for audio streams reads in units of samples.

The media clock has the following relationship to the network clock:

1) The media clock and the network clock shall share the IEEE 1588-2008 epoch of 1 January 1970 00:00:00 TAI, as defined in IEEE 1588-2008 clause 7.2.2. The value of the media clock shall be 0 at the IEEE 1588-2008 epoch.

NOTE Although TAI and UTC are mutually convertible by simply adding or subtracting leap seconds, they cannot be used interchangeably without such conversion. AES67 discusses and relies on TAI exclusively.

2) The rate at which the media clock increments shall be the same as the audio sampling frequency, referenced to the network clock. The value of the media clock shall change from 0 to 1 exactly one sampling period after the epoch, and further increase by 1 after every subsequent audio sampling period. The media clock for digital audio sampled at 48 kHz advances exactly 48 000 samples for each elapsed second on the network clock, for example.

NOTE If the sampling frequency of digital audio to be carried on the network does not match the media clock, it needs to be sampling-frequency converted to conform to the media clock.

RTP clock is a 32-bit representation of the least-significant part of the full media clock value. Individual RTP clocks each operate with a constant offset with respect to the media clock. The offset shall be conveyed through session description (see 8.3) on a per-stream basis.

NOTE The offset can be zero or a randomly selected 32-bit value. ST 2110-30:2017 defines a constraint to use a zero offset exclusively. To achieve interoperability with ST 2110-30:2017 devices, AES67 devices need to implement the ability to use a zero offset, either permanently or as a configuration option.

RTP timestamp is a 32-bit value of the RTP clock at the sampling instant of the first audio sample included in the packet. As per RFC 3550 and RFC 3551, RTP timestamps shall be carried in headers of RTP packets containing the related stream data.

The 32-bit RTP clock value will periodically overflow and roll over. The duration of the rollover period varies by the media clock rate. For example, at 48 kHz, the RTP clock will overflow approximately every 24,86 hours. To ensure proper phasing with respect to the network clock, media clock timestamps recovered at the receiver from the 32-bit RTP timestamps shall accurately take into account all such overflows (rollovers) between the epoch and the current time.

6 Transport

6.0 General

Transport aspects describe how media data, once encoded and packetized, is transported across the network. In terms of the OSI model, this clause defines operation on layer 3 (network layer) and layer 4 (transport layer). This standard does not specify how interoperability is achieved at lower layers in the model. It is assumed that best practices in transport of IP over the network technologies in question are employed.

NOTE Carriage of IP over Ethernet is described in RFC 894 - A Standard for the Transmission of IP Datagrams over Ethernet Networks.

6.1 Network layer

Editor's note on restructuring in AES67-2023:

- *In order to improve readability, this section introduces 3rd-level headings.*
- *Discussion of multicast transport was formerly split between this section and section 7 "Encoding and streaming". It is now consolidated in this section.*
- *Discussion of using IGMP for multicast connection management is moved from this section to section 10.2 "Multicast connections".*

6.1.1 General

Media packets shall be transported using IP version 4 as defined in RFC 791.

NOTE 1 Although care has been taken in design of this standard to facilitate future support for IPv6 (see RFC 8200), support for IPv6 is outside the scope of this revision of the standard.

Senders shall support either unicast, or multicast, or both unicast and multicast media streams. Receivers shall be able to receive both unicast and multicast media streams.

6.1.2 MTU size and message fragmentation

Despite a requirement in RFC 791, receivers are not required under this standard to support reassembly of fragmented packets. A receiver that does not support reassembly shall ignore IP packet fragments.

Senders may set the Don't Fragment flag (DF) bit in the IP header of outgoing media packets. In the event that a packet marked DF needs to be fragmented by the network, it will instead be dropped and an ICMP message of type 3 ("Destination Unreachable"), code 4 ("fragmentation needed and DF set") will be sent back to the sender.

In all cases, in response to receipt of the aforementioned ICMP message, senders should terminate transmission of the offending stream.

6.1.3 Multicasting

Multicasting of stream data allows for efficient one-to-many distribution of audio. Multicasting is also an important component in simplified connection management described in 10.2 in which a sender multicasts a stream and receivers discover and then simply listen to the stream in progress.

All multicast messaging relevant to this standard, specifically synchronization and media streaming, shall be accomplished using IP multicasting as described in RFC 1112.

NOTE 2 Additional tutorial information on IP multicasting is available in RFC 3170.

Although RTP supports many-to-many connections, under this standard it is assumed that only a single device sends per multicast address.

Where used, multicasting shall support administratively scoped multicast addresses in the range 239.0.0.0 to 239.255.255.255. This range can be subdivided by network administrators, and a subset can be allocated for use by media networking.

NOTE 3 Readers are encouraged to consult RFC 5771 on IANA Guidelines for IPv4 Multicast Address Assignments.

NOTE 4 ST 2110-10:2022 allows any multicast addresses to be used, excepting the Local Network Control Block, address range 224.0.0.0 - 224.0.0.255, and the Internetwork Control Block, address range 224.0.1.0 - 224.0.1.255, as specified in RFC 5771. If an AES67 device is intended to be used in ST 2110 environment, it needs to be ready to handle multicast addresses from the extended range.

The destination address used for a particular stream shall be configurable through the management interface at the sender. It is assumed that each stream will be assigned a unique destination address within the scope. The nature of the management interface and the allocation scheme used are outside the scope of this standard.

NOTE Additional best-practice information on IP administratively scoped multicasting is available from the IETF in RFC 2365 - Administratively Scoped IP Multicast.

6.2 Quality of service

On a network shared with unregulated non-real-time traffic, time-critical media traffic generally requires prioritized handling known as QoS. In order to facilitate the implementation of suitable QoS in the network, devices shall implement the DiffServ method as described in RFC 2474. DiffServ uses the DSCP field in each IP packet header to mark packets according to their traffic class so that the network can easily recognize packets that need to be treated preferentially.

Minimally the three traffic classes described in table 1 shall be supported. Devices shall tag outgoing traffic with an appropriate DSCP value. Devices should use the default values for the DSCP field as given in table 1, but traffic may be marked with alternate DSCP values as provided by a network administrator or user through a

management interface. Senders may be configurable to use the same DSCP values for multiple classes - this has the effect of combining classes.

Although not required, devices may implement a management interface for DSCP configuration. If such management interface is not implemented, the device shall use the default values exclusively.

Media streams that require very low delivery latency can fail to traverse a network reliably when transported in the same QoS class with media streams using longer packet times. In order to differentiate media streams with different requirements, senders may be configurable to use classes in addition to those shown in table 1.

Table 1 - QoS classes and DiffServ associations

Class name	Traffic type	Default DiffServ class (DSCP decimal value)
Clock	The following IEEE 1588-2008 messages: *Announce, Sync, Follow_Up, Delay_Req, Delay_Resp, Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up*	EF (46)
Media	RTP and RTCP media stream data	AF41 (34)
Best effort	All other IEEE 1588-2008 messages, discovery and connection management messages, other traffic managed as a part of the AES67 ecosystem.	DF (0)

NOTE 1 The DSCP markings on packets do not define any particular behavior of network devices or imply particular policies the network must implement. As a security measure, a network can even ignore incoming DSCP markings in which case it could distinguish and prioritize the traffic through other means (for example, UDP port number, IP addressing). Such network issues are outside the scope of this standard, although annex B provides some informative guidelines for network administrators.

NOTE 2 This standard makes no specific recommendation as to DSCP marking of traffic outside the scope of this standard. Traffic generated by other applications is marked as DF (0) by most systems, a situation compatible with this standard.

Senders should mark outgoing RTCP packets with the same DSCP value as the respective RTP stream packets. Receivers should mark outgoing RTCP packets with the same DSCP value they would use on RTP packets if transmitting a similar stream.

Receivers shall not depend on DSCP markings on received packets.

NOTE 3 DSCP markings can be changed in route by the network, or assignments can be reconfigured at the sender without the receiver's knowledge.

6.3 Transport layer

The transport layer provides end-to-end communications between applications in devices on a network. The layer handles issues of packet loss and reordering and implements multiplexing so that a single network connection can serve multiple applications on the end station.

Devices shall use Real-time Transport Protocol as defined in RFC 3550. Devices shall operate in accordance with RTP Profile for Audio and Video Conferences with Minimal Control as defined in RFC 3551, subject to overrides defined in this standard.

Devices shall support and should use the default ports allocated for RTP and RTCP without requiring user configuration: 5004 for RTP and 5005 for RTCP (see RFC 3551, section 8).

Senders may use, and receivers should support the following ports if used by a particular sender:

- all ports in the range 1024-49151 (User Ports, also known as Registered Ports, as defined in RFC 6335, section 6)
- all ports in the range 49152-65535 (Dynamic Ports, also known as Private or Ephemeral Ports, as defined in RFC 6335, section 6)

NOTE: Implementers are encouraged to support ports besides the default ports to ensure greater flexibility for system integration. However, interoperability in the Registered Ports range (1024-49151) and Private or Ephemeral Ports range (49152-65535) is not guaranteed, as receivers are recommended but not required to support any of them. Also, particular ports can be occupied by other receiver services and therefore be unavailable to AES67 audio streaming. Senders seeking to use other than the default ports should preferably use ports in the Private or Ephemeral Port range (49152-65535). It is the responsibility of the system integrator to check that interoperability can be achieved if non-default ports are chosen.

Devices shall use UDP as defined in RFC 768 for transport of RTP.

Contrary to allowance in RFC 3551 section 4.1, AES67 senders shall not use silence suppression. Consequently, AES67 receivers are not required to support silence suppression.

Normative statements regarding channel numbering, ordering, and content mapping defined in RFC 3551 section 4.1 are not a requirement for this standard and are therefore out of scope.

Fragmentation is undesirable and, under this standard, receivers are not required to perform reassembly (6.1). The standard 1500-byte Ethernet MTU is assumed. To prevent fragmentation through a standard Ethernet infrastructure when using IPv4, and to assure future compatibility with IPv6, the maximum allowed RTP payload size shall be 1440 bytes.

NOTE 1 On connections offering lower MTU than Ethernet's 1500 bytes, senders can use a smaller maximum payload than specified here.

Encrypted streaming using TLS, while supported in the RTP protocol suite, is outside the scope of this interoperability standard.

Senders should not include contributing source (CSRC) identifiers in the RTP header. Senders should not add RTP header extensions (RFC 3550 clause 5.1). However, as per RFC 3551, receivers shall tolerate the presence of CSRC identifiers and header extensions.

If senders do include header extensions or CSRCs, the 1440-byte maximum allowed payload shall be adjusted downwards by the size of the added header material so as to avoid fragmentation through a standard Ethernet infrastructure.

Both senders and receivers should transmit RTCP messages as specified in RFC 3550 clause 6. Senders and receivers should allocate RTCP bandwidth as recommended in RFC 3551 clause 2 (RTCP report interval).

Unicast senders should monitor connectivity to their respective receivers in such a way as to detect failure of the receiver, and stop transmission, within 60 seconds.

NOTE 2 Senders that continue their unicast transmissions to a missing receiver unnecessarily consume network resources and can generate excessive ARP traffic and potentially have their transmissions flooded to all devices on the network segment.

If a receiver implements RTCP as recommended by RFC 3550, RTCP receiver reports are generally sufficient to achieve the recommended monitoring. For monitoring of receivers that do not implement RTCP, senders may use any other monitoring means available to them, including any of the following techniques:

- SIP session timers as described in RFC 4028
- SIP OPTION ping (see IETF draft-jones-sip-options-ping)
- ICMP Echo request (ping, see RFC 792)

7 Encoding and streaming

7.0 General

This section defines streaming interoperability requirements.

To be AES67 compliant, devices shall support receiving of audio streams, or sending, or both.

A stream mode is defined by a specific combination of the audio sampling rate, audio sample encoding, stream channel count, and packet time values. For detailed enumeration of stream mode requirements refer to Annex G (PICS).

7.1 Payload format and sampling rate

Payload format defines audio sample encodings. The following payload formats are supported:

L16 16-bit linear format defined in RFC 3551 clause 4.5.11

L24 24-bit linear format defined in RFC 3190 clause 4

All devices shall support 48 kHz sampling rate. Devices should support 96 and 44,1 kHz sampling rates. See also AES5.

When operating at 48 kHz sampling rate:

1. Receivers shall support both L16 and L24 encodings
2. Senders shall support either L16, or L24, or both encodings

When operating at 96 kHz sampling rate:

1. Both senders and receivers shall support L24 encoding

When operating at 44,1 kHz sampling rate:

2. Both senders and receivers shall support L16 encoding

While all devices shall be able to support 48 kHz sampling rate, they are not required to accept 48 kHz connections at all times. They can for instance, have a global sampling rate configuration and only accept 48 kHz connections when the user selects global 48 kHz mode. Although not required, devices may support multiple sampling rates simultaneously.

Devices may support other combinations of payload formats L16 and L24 with sampling rates 44,1 kHz, 48 kHz, and 96 kHz.

Payload format and sampling rate combinations beyond those defined in this section are outside the scope of this standard.

NOTE 1 The indication of the payload format and sampling rate in use for a given stream is given by a combination of the payload type field in the RTP header (RFC 3550 clause 5.1) and description information associated with the stream (see 8).

NOTE 2 While RFC 3190 mentions an “emphasis” signaling parameter, emphasis is generally considered to be a legacy issue dating back to the early 1980s. Emphasis is no longer generally used in digital audio signals and it is not expected to be used in these inter-operable networked streams.

7.2 Packet time

7.2.0 General

Packet time is the real-time duration of the media data contained in a media packet. Given the sampling rate and packet time, the number of samples per packet can be calculated.

Short packet times allow for lower latency but introduce overhead and high packet rates that can overtax some devices or networks. Long packet times imply higher latency and require additional buffering which can be unavailable on memory-constrained devices.

Packet time is determined by the sender, is specified in the session description (see 8.1) and may be negotiated through connection management (see 10). Senders shall not change packet time for the duration of a session. Although not required, receivers may adapt to packet time changes during a session. To enable interoperation with standard RTP implementations, receivers should not rely on the presence or accuracy of any packet time description. Receivers should be able to determine packet time based on the timestamps in received packets.

Interoperability is addressed by the requirement that devices implement the "1 millisecond" packet time defined in 7.2.1. Further interoperability is encouraged through additional packet time recommendations in 7.2.2.

Packetization interval recommendations as defined in RFC 3551 section 4.2 shall be overridden by the relevant requirements of this standard.

Product documentation for a device shall indicate which packet times are supported in send and receive directions.

> NOTE In the text of this standard, for better readability, instead of indicating precise values in physical time units, packet time options are referenced by names in quotation marks, for example "1 millisecond".

7.2.1 Required packet time

A packet time of "1 millisecond" offers the widest possible interoperability and compatibility with audio and network equipment.

Senders shall be able to load each audio packet with 48 samples of audio data, when operating at a sampling rate of 48 kHz or 44,1 kHz, and 96 samples, when operating at a sampling rate of 96 kHz. Receivers shall be able to receive and decode 48-sample packets, when operating at a sampling rate of 48 kHz or 44,1 kHz, and 96-sample packets, when operating at a sampling rate of 96 kHz.

While all devices shall support the packet time requirements specified above, they are not required to accept these connections at all times. They may, for instance, have a global packet time configuration and only accept these connections when so configured. Although not required, devices may support multiple packet times simultaneously.

7.2.2 Recommended packet times

For enhanced interoperability over a range of applications, senders and receivers should support one or more of the other packet times listed in table 2.

Senders and receivers may support additional packet times. Maximum packet time is limited by network MTU as described in 6.3.

Table 2 (informative) - Required and recommended packet times

Packet time	Sample count / sampling rate	Effective duration	Interoperability notes
"125 microseconds"	6 / 48 kHz 12 / 96 kHz 6 / 44,1 kHz	125 µs 125 µs ~ 136 µs	Compatible with class A AVB transport
"250 microseconds"	12 / 48 kHz 24 / 96 kHz 12 / 44,1 kHz	250 µs 250 µs ~ 272 µs	High-performance, low-latency operation. Interoperable with class A and compatible with class B AVB transport.

"333 microseconds"	16 / 48 kHz 32 / 96 kHz 16 / 44,1 kHz	333 $^1/_3$ µs 333 $^1/_3$ µs ~ 363 µs	Efficient low-latency operation
"1 millisecond"	48 / 48 kHz 96 / 96 kHz 48 / 44,1 kHz	1000 µs 1000 µs ~ 1088 µs	Required common packet time for all devices adhering to this standard
"4 milliseconds"	192 / 48 kHz 192 / 44,1 kHz (MTU exceeded at 96 kHz)	4000 µs ~ 4354 µs n.a.	For applications desiring interoperability with EBU Tech 3326 or transport over wider areas or on networks with limited QoS capability

NOTE 96 kHz is not discussed in EBU Tech 3326. MTU restrictions of clause 6.3 limit a 96 kHz audio stream using 4-ms packet time to a single channel.

7.3 Stream channel count

The maximum number of channels per stream is limited by the packet time, payload format, sampling rate, and network MTU as described in 6.3.

Receivers shall support reception of streams with any channel count in the range from 1 to 8. Receivers may support reception of streams with more than 8 channels. Senders shall be able to offer at least one stream with a channel count in the range from 1 to 8. Senders may support streams with more than 8 channels.

Table 3 - Examples: maximum channel capacities per stream, MTU 1500 bytes

Payload format, sampling rate	Packet time	Maximum channels per stream
L24, 48 kHz	"125 microseconds"	80
L16, 48 kHz	"250 microseconds"	60
L24, 48 kHz	"250 microseconds"	40
L24, 48 kHz	"333 microseconds"	30
L24, 96 kHz	"250 microseconds"	20
L24, 48 kHz	"1 millisecond"	10
L24, 48 kHz	"4 milliseconds"	2

NOTE Although bundling multiple channels in a stream can improve network and processing efficiency, it is recommended that bundling be used primarily in service of the application. Channels of related material (for example, stereo or surround sound) are good candidates for bundling. Bundling of unrelated channels destined for different receivers in an effort to reduce network overhead is discouraged as this complicates media routing configuration.

7.4 Link offset

Link offset describes the latency through a media network. It is defined as the difference in time between when audio enters the sender network system (*ingress time*) and when it leaves the receiver network system (*egress time*).

Ingress time is referenced at ingress to the *sender network system*. RTP packets are marked with origination timestamps in the *timestamp* field (RFC 3550 clause 5.1) based on this reference point. *Egress time* is referenced at egress from the *receiver network system*. *Link offset* is therefore the time difference between ingress at the sender and egress at the receiver. Link offset and ingress and egress reference points are illustrated in figure 1.

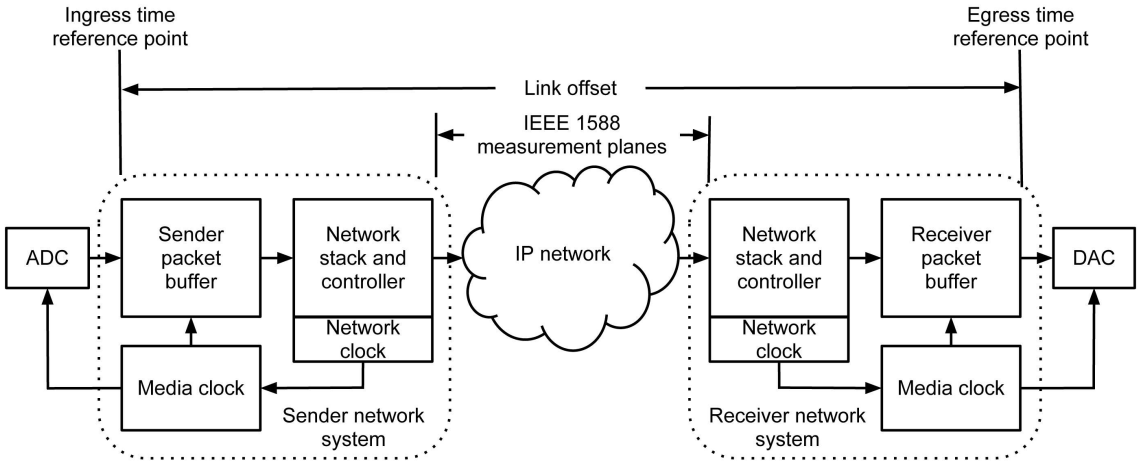

Figure 1 - Example network illustrating link offset and ingress and egress reference points

Link offset is determined at the receiver and is dependent on multiple factors, including packet time, propagation and queuing delays through the network, packet handling in the devices and buffering at the receiver. A receiver should attempt to maintain a constant link offset. At the same time it is recognized that unexpected changes to network conditions can require changing the buffering at the receiver resulting in a change of the link offset.

The link offset and any changes in link offset should be retrievable from the management entity of the receiver, if present.

> NOTE 1 Minimum possible link offset is packet time (see 7.2) plus network forwarding time. Forwarding time for a minimum-sized packet on a point-to-point gigabit Ethernet connection can be as low as 0,5 microseconds. Minimum link offsets in an actual implementation under realistic network conditions will approach twice the packet time and beyond.
>
> NOTE 2 Future work could specify a link offset management mechanism which is expected to require additional buffering and reporting of link offset to a central latency management server on the network and means for receivers to adjust link offset based on commands from the latency management server. See RFC 7272 for details.

7.5 Sender timing and receiver buffering

Buffering at the receiver is required to absorb jitter generated by packetization, network delivery and in the sender's and receiver's network stacks. The receiver's buffer must have capacity to also accommodate media for the duration of the link offset minus the minimum delivery time between sender and receiver. If buffering is too short, data can fail to arrive in time to be played, resulting in audio dropouts. Longer buffering improves robustness but introduces additional latency.

Receivers shall have a buffer capacity at least 3 times the packet time. Receivers should have a buffer capacity at least 20 times the packet time or 20 ms, whichever is smaller.

Senders nominally send packets associated with a stream at packet time intervals. Senders should transmit at the nominal transmission time with a variation of 1 packet time or less. Senders shall transmit data at the nominal transmission time with a variation of no more than 17 packet times or 17 ms, whichever is smaller.

The above requirements are designed to allow a range of implementations from hardware to applications running on desktop operating systems. Additional buffering capacity and more accurate transmission timing are encouraged and will produce improved robustness and interoperability.

Interoperability cannot be guaranteed between devices implementing only the minimal requirements at both ends, namely – between senders implementing the lowest transmission time accuracy and receivers implementing the smallest required buffer capacity. Such combinations can fail to deliver packets on time, causing audio dropouts. The interoperability cases are summarized in table n.

Table n: Sender timing and receiver buffering interoperability

Sender: Transmission time variation	**Receiver: Jitter buffer "Small"** 3 to 19 packet times, and less than 20ms	**Receiver: Jitter buffer "Large"** At least 20 packet times or 20 ms, whichever is smaller
"Strict" Not exceeding 1 packet time	YES	YES
"Loose" Not exceeding 17 packet times or 17 ms, whichever is smaller	NO	YES

7.6 Multicasting

Editor's note on restructuring in AES67-2023:

- *This section is moved to 6.1, to consolidate it with other requirements related to multicast transport.*

8 Session description

8.0 General

Session descriptions are used by discovery (see 9) and connection management (see 10) to specify critical information about each stream including network addressing (see 6), encoding format (see 7) and origination information.

SDP as specified in RFC 8866 shall be used to represent the sessions for connection management. Interoperability imposes additional SDP requirements and recommendations as set out in the following clauses.

NOTE RFC 8866 permits inclusion of additional attributes in session descriptions, besides those immediately required to fulfill the requirements of this standard. RFC 8866 requires that an SDP parser ignores any attributes it does not understand, instead of refusing the entire description.

8.1 Packet time

Packet time shall be represented in the session description by the following two SDP attributes defined in RFC 8866:

`a=ptime:`<milliseconds>[.<milliseconds fractional part>]

`a=maxptime:`<milliseconds>[.<milliseconds fractional part>]

Signaled packet time multiplied by sampling rate rounded to the nearest integer indicates the number of samples in each packet. The packet time descriptions shall be given with error less than half a sample period so that the calculated number of samples per packet rounds to the intended integer value. In many cases, this requires the inclusion of the <milliseconds fractional part> in the description. Where <milliseconds fractional part> is not required to accurately convey packet time, it shall be omitted from signaling.

Table 4 gives valid examples for the descriptions of packet times specified in 7.2. Packet times in descriptions shall be interpreted as a dotted decimal representation.

Values and representations beyond those enumerated in table 4 shall be correctly interpreted.

Table 4: Example, packet time signaling in SDP text

	Sampling rate		
Packet Time	**48 kHz**	**96 kHz**	**44,1 kHz**
“125 microseconds”	0.12	0.12	0.13
“250 microseconds”	0.25	0.25	0.27
“333 microseconds”	0.33	0.33	0.36
“1 millisecond”	1	1	1.09
“4 milliseconds”	4	4	4.35

NOTE Numbers used for packet time signaling are expressed in SDP text using the period character (decimal point).

Descriptions shall include a **`ptime`** attribute indicating the desired packet time. If more than one packet time is supported, a **`maxptime`** indicating the maximum packet time permitted shall be provided. The interoperable values for the <milliseconds> parameter for both **`ptime`** and **`maxptime`** are indicated in table 2.

The requirements of this description imply that the shorter packet time is always the preferred packet time. To override the implied assumption that a shorter packet time is always the preferred packet time, the capability negotiation attributes of RFC 5939 may be used to enumerate the supported packet times and order of preference.

If the range of packet times supported includes more than two of the standard packet times (table 2), the description should use the capability negotiation attributes of RFC 5939 to enumerate the supported packet times and order of preference.

NOTE The non-integral-millisecond descriptions can fail to be correctly understood by connection management partners not in compliance with this standard. The description could need to be confined to integer <millisecond> values when attempting connection to such partners.

8.2 Clock source

The **`ts-refclk`** attribute specifies the network clock reference used by the stream. **`ts-refclk`** supports specification of three variants of PTP in addition to other clock sources. The network clock source for each stream described shall be specified with one or more **`ts-refclk`** attributes as specified in RFC 7273.

RFC 7273 does not define a **`ptp-version`** option for IEEE 1588-2019. When IEEE 1588-2019 is in use, AES67 devices shall indicate “**`IEEE1588-2008`**”.

The following examples illustrate use of the attribute within the scope of synchronization options available in this standard.

EXAMPLE 1 Using an IEEE 1588-2008 network clock as discussed in 4.1 or 4.2. The GMID in this example is 39-A7-94-FF-FE-07-CB-D0 and the domain is 0:

```
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
```

EXAMPLE 2 Using an IEEE 1588-2008 network clock as discussed in 4.1 or 4.2. This example indicates that the network clock reference is traceable:

```
a=ts-refclk:ptp=IEEE1588-2008:traceable
```

EXAMPLE 3 Using an IEEE 802.1AS network clock as discussed in 4.3. The GMID in this example is 39-A7-94-FF-FE-07-CB-D0:

```
a=ts-refclk:ptp=IEEE802.1AS-2011:39-A7-94-FF-FE-07-CB-D0
```

Although, as discussed in RFC 7273, the PTP domain specification is optional, under this standard, when an RTP stream is referenced to IEEE 1588-2008, and a specific grandmaster clock is indicated, signaling shall include both GMID and PTP domain.

IEEE 802.1AS always uses domain 0 so no domain indication is supported or necessary. Two devices synchronized to IEEE 802.1AS shall be assumed to be using the same domain.

Receivers should attempt to connect to senders if they are using the same PTP domain and the same GMID clock reference as the sender, or when both are using traceable time. Receivers should not attempt to connect to senders if they are using a different PTP domain for their clock reference than the sender, and at least one of the used references is not known to be traceable.

The case of a PTP domain match and mismatched GMID can indicate either a transition state of the network or lack of a common clock reference between sender and receiver or different grandmasters referenced to the same traceable clock source (for example, GPS). Receivers may attempt to make a connection in this case even if one or both used references are not known to be traceable, but they should be prepared for a possible synchronization failure.

Under RFC 7273, senders may specify multiple equivalent clock sources. Receivers should evaluate all clock sources specified and should attempt to connect based on the recommendations in this clause.

Senders and receivers should monitor for changes in their synchronization status during transmission. Senders should update their clock source description when a change is detected.

When their synchronization status changes or an updated description is received from the sender, receivers should reevaluate their ability to continue receiving according to the recommendations in this clause. Receivers are not required to terminate reception on detection of synchronization signaling mismatch in an ongoing stream, but they should be prepared for possible synchronization failure.

8.3 RTP and media clocks

The relationship of media clock to RTP clock shall be described for each stream with an **a=mediaclk:direct=**<offset> attribute as specified in RFC 7273 clause 5.2. The offset specification shall be included in the description. A **mediaclk** attribute shall be provided for each stream described.

NOTE 1 RFC 7273 allows mediaclk to be specified at session, media or source levels. As described in section 5.4 of that document, a declaration at a higher layer satisfies the above requirement for all lower level streams.

NOTE 2 Contrary to RFC 7273 and AES67, ST 2110-10:2022, and consequently ST 2110-30:2017 too, require mediaclk to be always specified at the media level.

NOTE 3 The relationship of media clock to network clock is fixed and specified in 5.

The **mediaclk** attribute supports numerous media clock scenarios. The following example illustrates use of the attribute within the scope of this standard.

EXAMPLE media clock description - the RTP timestamp has a value of 1810024580 at the media clock epoch:

```
a=mediaclk:direct=1810024580
```

8.4 Payload types

Allocation of a dynamic payload type and associated **rtpmap** attribute is required to specify the interoperable payload formats defined in 7.1, as none of these formats are called out as static payload types in RFC 3551 (section 6, table 4 of that document). The receiver shall determine the payload format using the **rtpmap** attribute; it shall not assume any fixed relationship between payload type value and payload format. The relationship is defined on a stream-by-stream basis by senders using the **rtpmap** attribute.

8.5 Example descriptions

8.5.0 Errata

SDP errors in both this standard and in RFC 7273 have been corrected in standard revisions and published errata. It is possible to implement an SDP interpreter to tolerate the resulting errors in earlier implementations. Errors have included invalid **t=** specification, incorrect **a=sendonly** specification for multicast stream description, mis-ordered lines and specification of PTP domain with **domain-nmbr=** syntax.

8.5.1 Multicast session description example

Example simple SDP description for 8 channels of 24-bit, 48 kHz audio transmitted as a multicast stream with 1-millisecond packet time.

```
v=0
o=- 1311738121 1311738121 IN IP4 192.168.1.1
s=Stage left I/O
c=IN IP4 239.0.0.1/32
t=0 0
m=audio 5004 RTP/AVP 96
i=Channels 1-8
a=rtpmap:96 L24/48000/8
a=recvonly
a=ptime:1
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:direct=963214424
```

8.5.2 Unicast session description example

Example simple SDP description for 8 channels of 24-bit, 48 kHz audio transmitted as a unicast stream with 250-microsecond packet time.

```
v=0
o=audio 1311738121 1311738121 IN IP4 192.168.1.1
s=Stage left I/O
c=IN IP4 192.168.1.1
t=0 0
m=audio 5004 RTP/AVP 96
i=Channels 1-8
a=rtpmap:96 L24/48000/8
a=sendonly
a=ptime:0.250
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0
a=mediaclk:direct=2216659908
```

9 Discovery

Discovery is the network service which allows participants to build a list of the other participants or sessions available on the network. Such a list can be presented to users to assist with connection management. Discovery services can be used to deliver SDP description (see 8), a SIP URI (see 10.1), or equivalent session data in a different presentation format.

Devices are not required to implement discovery services. Devices may implement one or more discovery services. Examples of discovery methods are listed in Annex E.

10 Connection management

10.0 General

Connection management is the procedure and protocols used to establish one or more media streams between a sender and one or more receivers.

10.1 Unicast connections

Devices should support connection management for unicast streams using the Session Initiation Protocol (SIP) as defined in RFC 3261. SIP is widely used in IP telephony and is the connection management protocol used by EBU Tech 3326.

Devices may additionally support connection management for unicast streams using other protocols, or through the use of a management interface.

The following subsections apply to SIP connection setup, when such is used.

10.1.1 SIP URI

Under SIP, audio devices are SIP *user agents* with an associated SIP URI. SIP allows user agents to locate and make connections to other user agents by referencing the other's SIP URI. A SIP URI for a potential connection can be learned through discovery (see 9) or through other means (for example: static configuration, proprietary directory service).

The **`sip:`** URI form shall be used by devices using SIP under this standard. The **`sips:`** URI form indicates that connection management should be conducted securely using TLS. Secure connection management does not necessarily imply that stream transmission is done securely. Stream transmission security is negotiated as part of the connection management process. RFC 3261 recommends that user agents implement TLS. However, because this standard does not provide a recommendation for interoperability of secure media streaming (see 6.3), devices may support TLS and **`sips:`**, but they are not required for interoperability.

10.1.2 Server and serverless modes

SIP is conventionally used with the assistance and participation of SIP servers. Different types of servers perform different tasks for a SIP network. Servers can be located anywhere on the network where they are reachable by end stations. The use of servers creates a flexible and scalable connection management system.

Serverless mode is used to perform connection management between user agents in direct peer-to-peer fashion without the intervention of servers. The serverless mode is appropriate for modest installations where, due to limited scale, servers produce minimal benefit and the overhead of installing and configuring SIP servers introduces unnecessary complication.

In order to perform peer-to-peer connection management, the caller must have some means of determining network contact information (that is, host name or IP address) of the callee. This can be obtained through discovery (see 9), manual configuration, or higher-layer protocols. In peer-to-peer connection management, all SIP messages are directed to the target device instead of the server. A device using SIP under this standard shall respond to such requests.

Regardless of the support for serverless mode, devices using SIP shall be able to operate in a normal SIP environment featuring servers. Specifically, devices using SIP shall attempt to discover, and register with, SIP registration servers and respond to messages originating from servers.

10.1.3 User-Agent header field

The *User-Agent* header field in the SIP protocols is useful for conveying information about the end station that managers can use to expedite connection management and work around implementation-specific issues. The format of *User-Agent* data is defined in RFC 2616 clause 14.43. Devices should include a User-Agent header field in REGISTER and INVITE messages.

NOTE Security considerations for *User-Agent* in the context of SIP are discussed in RFC 3261 clause 20.41.

10.1.4 Format negotiation

The standard offer/answer model as described in RFC 3264 shall be used to negotiate the encoding format for a connection. Supported encoding formats are described in 7.1.

10.1.5 Packet time negotiation

The offer/answer model does not address negotiation of attributes such as packet time. As per 8.1, an offer supporting multiple packet times specifies a range with the **`ptime`** and **`maxptime`** attributes. When SIP is used for connection setup, an answer shall assume that **`ptime`** and **`maxptime`** packet times are supported. To increase flexibility and reliability, devices may implement the capability negotiation provisions of RFC 5939.

10.2 Multicast connections

Editor's note on restructuring in AES67-2023:

- *Discussion of using IGMP for multicast connection management is now consolidated in this section. Formerly, part of this was included in section 6.1 "Network layer".*

Multicast connection management may be accomplished without use of a connection management protocol. In this simple connection management scenario, the receiver is not required to make direct contact with the sender.

A receiver obtains a session description of the desired connection using discovery (see 9) or other means. When a receiver learns of a stream to which it would like to make a connection, it can use IGMP to inform the network of desire to receive and begin receiving the stream.

To ensure that desired multicasts are received and to allow the network to filter undesired multicasts, all devices shall support IGMPv2 as defined in RFC 2236 and should support IGMPv3 as defined in RFC 3376.

Devices supporting IGMPv3 (and consequently IGMPv2), as delivered from the manufacturer, when connecting to a network, shall start in the IGMPv3 mode and then follow the IGMPv3 rules according to the actual network conditions. Devices may implement a management interface allowing to set them to IGMPv2 mode.

NOTE 1 When operating on an IGMPv2 network, IGMPv3 devices can experience a two-minute startup delay looking for IGMPv3 services on the network. Networks can be configured to shorten this delay.

NOTE 2 The SMPTE ST 2110 suite includes a requirement to support IGMPv3. Addition of the recommendation to support IGMPv3 to AES67-2023 facilitates better alignment between the two standards.

NOTE 3 RFC 2236 and RFC 3376 include backwards-compatibility requirements. A device supporting IGMPv2 is able to correctly operate on a network supporting IGMPv1 or IGMPv2. A device supporting IGMPv3 is able to correctly operate on a network supporting IGMPv1, IGMPv2 or IGMPv3.

Devices shall use IGMP to request reception of any multicasts required. These include receipt of IEEE 1588-2008 synchronization messages (see IEEE 1588-2008 clause D.3), media streams using multicast addressing (see 7.6), and also messages of other application protocols that could be used on the device, such as discovery (clause 9), that use multicast messaging.

NOTE 4 IGMP registration data can be purged by the network in some routing reconfiguration scenarios. Such a purge can result in an interruption of streamed data. An effective method of expediting restoration of service is to retransmit IGMP membership reports. This can be achieved by closing and immediately reopening any affected multicast network sockets.

NOTE 5 The requirement for the senders to request receipt of their own multicast packets has been removed from AES67-2023. See Annex F for additional details.

Annex A (Normative) – Media profile

A.0 General

The IEEE 1588-2008 precision time protocol was designed to accommodate a range of synchronization applications. The differing requirements of different applications are accommodated through the definition and use of profiles. A profile is a set of operating parameter values and list of protocol components used in service of the application.

The media profile defined in this annex serves the requirements of media networking. Specifically, the profile enables short start-up time, high accuracy and compatibility with ordinary (PTP-unaware) IP networking equipment.

It is anticipated that additional profiles will be published as the IEEE 1588-2008 ecosystem grows and that some of these will be applicable to media networking. In response to these anticipated advances, devices may choose to implement profiles in addition to or instead of the media profile defined in this annex.

A.1 Media profile description

The Media profile specifies attributes and options required and allowed in use of IEEE 1588-2008 in synchronization for media transport under this standard. The basic structure and some of the text in this definition is taken from IEEE 1588-2008 annex J.

The Media profile differs from the Delay Request-Response Default PTP profile given in IEEE 1588-2008 clause J.3 in the following aspects:

- Identification information indicating AES origin of the profile.
- **`portDS.logSyncInterval`** and **`portDS.logMinDelayReqInterval`** are reduced to improve startup time and improve accuracy when using ordinary, as opposed to IEEE 1588-2008 enabled, network equipment.
- The only supported message encapsulation is UDP/IPv4
- Clock physical requirements compatible with AES11
- Additional **`clockClass`** values to signal AES11 DARS Grade
- A recommendation to implement the peer delay mechanism in addition to request-response

A.2 Media profile

A.2.1 Identification

The identification values for this PTP profile (see IEEE 1588-2008 clause 19.3.3) are as follows:

PTP Profile:
PTP profile for media applications.
Version 1.0
Profile identifier: 00-0B-5E-00-01-00

This profile is specified by the AES Standards Committee.

A copy can be obtained by ordering AES67-2023 from the AES Standards Store at:
www.aes.org/publications/standards/

A.2.2 PTP attribute values

Nodes shall implement all requirements in this PTP profile that specify default values or choices such that these default values or choices apply without requiring user configuration; that is, as delivered from the manufacturer.

Table A.1 – attribute values for use with Media profile

Attribute	Values
`defaultDS.domainNumber`	The default initialization value shall be 0. The configurable range shall be 0 to 127.
`portDS.logAnnounceInterval`	The default initialization value shall be 1. The configurable range shall be 0 to 4.
`portDS.logSyncInterval`	The default initialization value shall be -3. The configurable range shall be −4 to +1.
`portDS.logMinDelayReqInterval`	The default initialization value shall be 0. The configurable range shall be -3 to 5 or `portDS.logSyncInterval` to `portDS.logSyncInterval`+5, whichever is more restrictive.
`portDS.logMinPdelayReqInterval`	The default initialization value shall be 0. The configurable range shall be 0 to 5.
`portDS.announceReceiptTimeout`	The default initialization value shall be 3. The configurable range shall be in the range 2 to 10.
`defaultDS.priority1`	The default initialization value shall be 128.
`defaultDS.clockClass`	Table A.2 specifies additional values beyond those specified in IEEE 1588-2008 table 5 in support of AES11 physical clock specifications (see A.2.4).
`defaultDS.priority2`	The default initialization value shall be 128.
`defaultDS.timeReceiverOnly`	If this parameter is configurable, the default value shall be FALSE.
`transparentClockdefaultDS.primaryDomain`	The default initialization value shall be 0.
τ (see IEEE 1588-2008 clause 7.6.3.2)	The default initialization value shall be 1,0 s.

For each defined range, manufacturers may allow wider ranges.

Table A.2 - clockClass values for use with Media profile

clockClass (decimal)	Specification	Time scale	TimeReceiver capable	Specific to Media profile
6	A clock that is synchronized to a primary reference time source (for example, GPS)	PTP	No	No
7	A clock that has previously been designated as `clockClass` 6 but that has lost the ability to synchronize to a primary reference time source and is in holdover mode and within holdover specifications	PTP	No	No

13	A clock that is synchronized to an external media clock source	ARB	No	No
14	A clock that has previously been designated as **`clockClass`** 13 but that has lost the ability to synchronize to an external media clock source and is in holdover mode and within holdover specifications	ARB	No	No
52	Degradation alternative A for a clock of **`clockClass`** 7 that is not within holdover specification	PTP	No	No
58	Degradation alternative A for a clock of **`clockClass`** 14 that is not within holdover specification	ARB	No	No
150	A clock whose frequency is synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) and whose time has been previously synchronized to a primary reference time source	PTP	Yes	Yes
158	A clock whose frequency is synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) and whose time has been previously synchronized to a primary reference time source	PTP	Yes	Yes
166	A clock of unspecified tolerance that has been previously synchronized to a primary reference time source	PTP	Yes	Yes
187	Degradation alternative B for a clock of **`clockClass`** 7 that is not within holdover specification	PTP	Yes	No
193	Degradation alternative B for a clock of **`clockClass`** 14 that is not within holdover specification	ARB	Yes	No
220	A clock whose frequency is synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) and whose time has *not* been previously synchronized to a primary reference time source	ARB	Yes	Yes
228	A clock whose frequency is synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) and whose time has *not* been previously synchronized to a primary reference time source	ARB	Yes	Yes

248	Default. This **`clockClass`** shall be used if none of the other **`clockClass`** definitions apply; for example, a clock of unspecified tolerance that has not been previously synchronized to a primary reference time source.	ARB	Yes	No
255	A timeReceiver-only clock	n.a.	Yes	No

Holdover specification for all **`clockClass`** specifications are +/- 5 % of a 96 kHz word-clock period.

A.2.3 PTP options

Devices may implement any options of IEEE 1588-2008 clause 17. Devices may implement unicast negotiation as defined in clause 16.1 of IEEE 1588-2008. All of these options shall be inactive unless specifically activated by a management procedure.

Node management shall implement the management message mechanism of IEEE 1588-2008.

The best timeTransmitter clock algorithm shall be the algorithm specified by IEEE 1588-2008 clause 9.3.2.

The default path-delay measurement mechanism shall be the delay request-response mechanism specified by IEEE 1588-2008. The peer delay mechanism should also be implemented.

> NOTE Only a single mechanism is allowed per link. Boundary clocks should be used between links that use different path delay mechanisms.

A.2.4 Clock physical requirements

Clocks shall meet requirements of Grade 2 DARS set forth in AES11 clause 5.2. Clocks may conform to the Grade 1 requirements but in so doing shall indicate this to other network participants as described in A.2.2 **`defaultDS.clockClass`**.

> NOTE The fact that individual clocks meet AES11 requirements does not guarantee the whole clock distribution system meets AES11 specifications as network performance also contributes to the quality of the distributed clock.

Annex B (Informative) – Network QoS configuration recommendations

B.0 General

Networks supporting high-performance media streaming must provide the QoS required by these services. This standard recognizes the need for QoS and, in 6.2, specifies how traffic is to be marked for the network. Although network behavior is outside the scope of this standard, this annex gives general configuration guidelines for an IP network with QoS implemented in accordance with IETF DiffServ recommendations. See also RFC 4594 for general background on DiffServ configuration.

B.1 DiffServ network configuration

DiffServ is a framework for classification and differentiated treatment of network traffic. DiffServ is considered a course-grained QoS architecture as it operates on classes of traffic in aggregate rather than on individual traffic flows.

Per-hop behavior (PHB) is fundamental to DiffServ operation. At each router or switch, traffic is classified and retransmitted according to that classification. In the presence of congestion, higher priority traffic is retransmitted promptly while lower priority traffic waits in buffers inside network equipment and can be discarded altogether. The use of PHBs on classes of traffic, as opposed to individual traffic flows (for example, a media stream) produces a simple and scalable QoS solution.

Ethernet equipment typically supports PHB though use of multiple egress queues per port. Different classes of traffic are queued separately. When an egress port is available to transmit a packet, a selection algorithm selects a packet from one of the queues. There are several selection algorithms in use in network equipment. Arguably the simplest selection algorithm is *strict priority*. Under this scheme the oldest packet from the highest-priority non-empty queue is selected for transmission. This results in the lowest possible latency for high-priority traffic but can result in lower priority traffic being held off for long periods possibly resulting in lost data, a condition known as starvation. More sophisticated selection algorithms such as *weighted round robin* and *guaranteed minimum bandwidth*, address starvation by adding more balance to the process. Under these algorithms, no one is starved for bandwidth but latency for highest priority traffic is higher.

Within a network supporting DiffServ, traffic classes are identified by a 6-bit differentiated services code point (DSCP) value in every IP header. These values are assigned by the end stations generating the traffic as defined in 6.2.

Media streaming under this standard presents three traffic classes to the network. The recommended DSCP values specified in 6.2 for each class are suitable for use on networks configured in accordance with the recommendations in DiffServ RFCs.

> NOTE Not all networks are configured in accordance with the DiffServ recommendations. The RFCs allow considerable latitude and consequently, there is considerable variation in how DiffServ is configured and deployed on networks.

B.1.1 Clock

Clock traffic consists of relatively low-frequency transmissions (less than 100 packets/second) of small UDP packets (on the order of 100 bytes each) and low bandwidth (less than 100 kbits/second). Although clock traffic is not highly sensitive to packet loss, it is sensitive to latency and specifically to latency variation which is also known among network engineers as delay variation (DV). High levels of DV for clock traffic degrades the accuracy of clock delivery across the network.

It is recommended that IEEE 1588-2008 clock traffic be assigned highest priority for application traffic on the network recognizing that on some networks, network management and routing control traffic are assigned priority higher than any application traffic. On most DiffServ networks, this is achieved by assigning clock traffic to the Expedited Forwarding (EF) class, DSCP value 46. EF is typically implemented as a strict priority queue that is given transmission preference ahead of other queues.

NOTE EF is often also used for VoIP traffic on the assumption that, compared to previous network applications (file transfers, e-mail deliver, web browsing), VoIP is the most performance-critical network application. However, VoIP typically operates using a 20 ms packet time making it up to two orders of magnitude less time critical than media traffic associated with this standard. Fortunately, the small packet size (100 bytes typical) and low bandwidth (10 kbit/second per call) limit the potential interference from VoIP traffic.

Networks configured with an EF differentiated service and willing to trust the end station markings could accept the traffic as is. More security-conscious networks could wish to recognize and classify clock traffic based on its addressing and port assignments as shown in table B.1.

Table B.1 - Clock class traffic identification

Traffic	Destination address	Protocol	Destination port
Time-critical event messages	224.0.1.129	UDP	319

B.1.2 Media

Media traffic is characterized by high frequency (up to 8000 packets per second per stream) and high bandwidth (over 1 Mbit per audio channel). Packet size depends on number of audio channels carried in a stream and can range from less than 100 bytes for a single-channel stream to the full 1500-byte Ethernet MTU for a maximally loaded stream. Media packets must be delivered in a timely and reliable manner. Any packet loss in media data will manifest as audio dropouts. For the highest performance applications, packets delayed more than 250 microseconds by the network could arrive too late to be useful and this will also manifest as audio dropouts. The audibility and impact of audio dropouts is application dependent. Most professional applications have a low tolerance for dropouts.

This standard recommends in 6.2 that senders mark media traffic with the AF41 DSCP value of 34. The clause also allows implementations to use alternate DSCP markings for different streams based on QoS requirements.

It is recommended that media traffic be assigned to a high-priority assured-forwarding class. Assured-forwarding PHB is defined in RFC 2597. Some applications could benefit from multiple assured-forwarding classes for media. Implementation of assured forwarding could require network engineering to establish a subscribed rate that supports the traffic expected to be carried by the class. Best performance is achieved if assured forwarding is implemented as a second strict priority queue below EF. Assured forwarding is more typically implemented using a weighted round robin selection algorithm.

Networks configured with an AF41 differentiated service and willing to trust the end station markings could accept the traffic as is. More security-conscious networks could wish to recognize and classify media traffic based on its addressing and port assignments as shown in table B.2.

Table B.2 - Media class traffic identification

Traffic	Destination	Protocol	Destination port
Default RTP media	Any unicast or multicast	UDP	5004
Default RTCP management	Any unicast or multicast	UDP	5005
Additional or alternate media traffic addressing as allowed in 6.3	Any unicast or multicast	UDP	Any user port (1024 to 65535)

NOTE Details on additional or alternate port assignments for media traffic for a specific device can be available in product documentation or though inquiry of the manufacturer.

B.1.3 Best effort

Best-effort traffic constitutes any traffic related to this standard which is not classified as clock or media traffic. This includes messaging associated with discovery (see 9) and connection management (see 10).

End stations normally mark best-effort (BE) traffic with a DSCP value of 0. Because all networks implement a best effort traffic class and any traffic not otherwise classified is assumed to be best effort, no special configuration is typically required to accommodate BE traffic.

Annex C (Informative) – AVB network transport

C.0 General

This standard defines how to use an IP transport to carry media data. Ethernet is a common network used to support IP transport. AVB can be viewed as an improved Ethernet. The AVB network is described in a suite of IEEE standards. The standards relevant to network transport are:

- IEEE 802.1BA - Audio Video Bridging (AVB) Systems
- IEEE 802.1Q-2011 Clause 34 - Forwarding and queuing for time-sensitive streams
- IEEE 802.1Q-2011 Clause 35 - Stream Registration Protocol (SRP)

The AVB improvements can be used to improve performance of media transport. AVB improvements can also be used to improve clock distribution, a topic which is discussed separately in annex D.

C.1 AVB network transport

Two basic methods for transmission of interoperable media streams across an AVB network are available.

Transporting interoperable media as an AVB time-sensitive stream makes use of AVB improvements with the restriction that this method of transport must be confined to an AVB network domain and cannot exit to legacy portions of a network that do not support the AVB improvements and thus cannot connect to non-AVB devices.

Transporting interoperable media as “other traffic” allows the flexibility to connect to and through AVB and non-AVB portions of the network and to non-AVB devices. This mode of transport does not make use of AVB’s QoS and registration services and is confined to the proportion of network bandwidth not allocated to time-sensitive streams, typically 25%.

C.1.1 Interoperable media as AVB time-sensitive streams

Devices in compliance with this standard that also have AVB capabilities may choose to use AVB bandwidth reservation and QoS provisions for transport of media streams. The advantages of this approach include:

- Reserved bandwidth assures successful media transmission across the network
- Potential compatibility with other AVB devices using IP transport
- Premium bandwidth and latency performance

Devices using this method must be directly connected to an Ethernet switch supporting the AVB standards and protocols. Under AVB, there are two classes of media streams. Class A streams are highest priority and achieve a 2-ms latency guarantee on IEEE 802.1BA compliant networks. Class B streams are lower priority but still higher priority than any other non-stream data on the network. Class B streams achieve a 50-ms latency guarantee on IEEE 802.1BA compliant networks. Either class may be used for interoperable media transport.

Before a device can transmit interoperable media as AVB media traffic, bandwidth must be reserved for the traffic. This is accomplished using the Stream Reservation Protocol described in clause 35 of IEEE 802.1Q-2011. In the reservation, the sender specifies the class of traffic (A or B) and the maximum amount of traffic generated in the measurement interval for that class. The measurement interval is 125 microseconds for class A and 250 microseconds for class B.

Because the AVB network identifies streams by their Ethernet destination addresses, multicast destination addressing is generally used for AVB streams. Interoperable media transmitted should use multicast IP addressing as described in 7.6. Furthermore since, as per RFC 1112, multiple IP multicast addresses are mapped to a single Ethernet multicast address, additional care beyond what is described in 7.6 must be taken when selecting multicast destination addresses for use with AVB time-sensitive streams.

C.1.1.1 Sender behavior

A sender (talker in AVB terminology) uses SRP to provide the following information to the AVB network in a Talker Advertise message:

`StreamID`	A 64-bit globally-unique identifier comprised of the 48-bit Ethernet MAC address of the sender and 16-bits unique to the device generated by the sender.
Destination address	The Ethernet multicast destination address used for stream transmission.
VLAN identifier	All AVB media packets include an IEEE 802.1Q-2011 VLAN tag. The VLAN tag contains a 12-bit VLAN identifier. The default VLAN identifier for AVB streams is 2.
`MaxFrameSize`	The maximum size of the media stream packets created as defined in 7. **`MaxFrameSize`** includes the IP header but excludes any Ethernet overhead.
`MaxIntervalFrames`	Maximum number of frames a sender can transmit in one measurement interval. Since packet times allowed by 7.2 are greater than or equal to AVB measurement intervals, this is always 1.
Data Frame Priority	3 for class A, 2 for class B
Rank	1 for normal traffic, 0 for emergency traffic

SRP allocates a fixed amount of bandwidth equal to **`MaxFrameSize`** × **`MaxIntervalFrames`** per measurement interval. Note that because packet times for interoperable streams are equal to or longer than AVB measurement intervals, actual bandwidth consumed can be significantly less than bandwidth reserved. This should be taken into account in AVB network design.

Senders must receive positive acknowledgement in the form of *Listener Ready* or *Listener Ready Failed* from SRP before transmitting stream data packets. Senders transmit stream data packets formatted as IP packets as described in 7. AVB networks use multicast filtering to enforce bandwidth allocation and thus support only multicast destination addressing for stream data (see IEEE 802.1Q-2011 clause 35.2.2.8.3). Stream packets transported in this way across an AVB network must therefore use multicast IP destination addressing. The IP multicast destination address is mapped to an Ethernet multicast destination address in the range 01:00:5e:00:00:00 to 01:00:5e:7f:ff:ff according to the procedures of RFC 1112 clause 6.4.

The IP packets are framed into Ethernet packets. The Ethernet framing must include an IEEE 802.1Q-2011 VLAN tag with VLAN identifier and priority code point matching the corresponding SRP parameters: *VLAN Identifier* and *Data Frame Priority*. Transmission timing of stream data packets within a class by the sender must be in accordance with the credit-based shaper described in IEEE 802.1Q-2011 clause 8.6.82.

C.1.1.2 Receiver behavior

A receiver (listener in AVB terminology) uses SRP to provide a desired StreamID to the network in a *Listener Declaration*. If the stream is available, the receiver receives a *Talker Advertise* and must then send an MVRP membership request to join the VLAN specified in the *Talker Advertise*. At this point, stream data packets begin to arrive and the receiver should process them in the same manner they are processed on non-AVB networks.

C.1.2 Interoperable media as other traffic

Although, an AVB network reserves up to 75% of available bandwidth on each link for media traffic, this allocation is configurable. A larger percentage of bandwidth can be reserved for other traffic. The AVB QoS algorithm known as the credit-based shaper is considerate of other traffic giving it an opportunity to traverse the network each measurement interval. Measurement intervals are 125 microseconds for AVB networks carrying only class A streams and 250 microseconds for networks carrying class B streams or both class A and class B streams. Although other traffic is prioritized below AVB media traffic, the other traffic may be prioritized with

respect to itself. All of this creates an acceptable environment for interoperable media to be carried on or through an AVB network as other traffic. The advantages of this approach include:

- No requirement to implement SRP
- No requirement to implement credit-based shaper
- Ability to transmit media data across the AVB network boundary including bridging media between separate AVB network domains
- Works with or without AVB network equipment
- No wasted bandwidth due to mismatch between measurement interval and packet time
- No special sender or receiver behavior is required

C.1.2.1 Sender behavior

No special sender behavior is required. On some AVB networks, senders may wish to use an IEEE 802.1Q-2011 VLAN tag to convey QoS class. AVB networks are generally geared to Ethernet protocols and could not inspect the DSCP value supplied in the IP header.

The VLAN identifier value depends on the VLANs configured on the network. A value of 0 may be used here to indicate default VLAN membership though this use of a 0 VLAN identifier is not supported on all network equipment. Recommended values for the priority field in the VLAN tag are given in table B.3.

Table C.3 - Recommended PCP values for 802.1Q-2011 tagged media as "other traffic" through an AVB network

Traffic	PCP
Clock	6
Media	5
Best effort	0

NOTE PCP values 2 and 3 are used for AVB media traffic and will be remapped to 0 at the AVB domain boundary by the switch as per IEEE 802.1Q-2011 Clause 6.9.4.

C.1.2.2 Receiver behavior

No special receiver behavior is required. Receivers may wish to be prepared to receive media packets with IEEE 802.1Q-2011 tags.

Annex D (Informative) – Interfacing to IEEE 802.1AS clock domains

D.0 General

IEEE 802.1AS defines a profile for use of IEEE 1588-2008 on enhanced Ethernet networks, as specified in IEEE 802.1BA-2011. This type of networking is commonly known as Audio Video Bridging (AVB). As outlined in clause 4.3, AVB networks can use their native IEEE 802.1AS synchronization profile in place of the profiles otherwise required in clause 4 (the IEEE 1588-2008 default profiles and the Media profile specified in annex A). In many cases it can be desirable to interface the two profiles to build heterogeneous networks using both techniques.

Under both IEEE 1588-2008 and IEEE 802.1AS, a PTP clock is designated as an ordinary clock (OC), boundary clock (BC) or transparent clock (TC) though IEEE 802.1AS transparent clocks also have some boundary clock capabilities. A device may implement one or more of these capabilities. PTP clocks have one or more ports. OCs may have as few as one port. TCs and BCs must have two or more ports. Ports on BCs and OCs have an associated operating state, timeTransmitter or timeReceiver. An IEEE 1588-2008 profile is associated with each port of a clock. Due to its transparent nature, a TC is able to simultaneously associate itself with multiple clock domains and profiles.

D.1 Boundary clock interface

Typically, the same profile is used on all ports of a BC but it is also possible to associate different profiles with different ports. A boundary clock running different profiles on different ports creates an interface between portions of an IEEE 1588-2008 network operating with different profiles. The different profiles that may be used for media networking include the default profiles required by this standard and defined in IEEE 1588-2008 annex J, the Media profile recommended by this standard and defined in annex A, and the profile used by IEEE 802.1AS on AVB networks.

When a BC is used to communicate synchronization to different network segments in this way, all network segments participate in the best timeTransmitter clock algorithm (BTCA), and all segments must use the same version of IEEE 1588-2008 protocol, and all must use the same domain identifier. Use of BCs in this manner can be appropriate for connecting a segment running one of the default profiles with a segment running the Media profile. Because IEEE 802.1AS uses an Ethernet-specific BTCA this approach is not appropriate for connecting an IEEE 1588-2008 segment to an IEEE 802.1AS segment.

There are no clock source signaling issues when interconnecting clocks in this manner. The interface through the BC simply creates a larger PTP domain. Synchronization identification (see 8) is the same on both sides of the BC interface.

D.2 Ordinary clock interface

A multiport OC can run different profiles on different ports. With the following restrictions, a multiport OC can be used to synchronize multiple segments. An OC can have at most one port in timeReceiver mode. Clock distribution must be engineered to assure that an OC is not asked to be a timeReceiver on multiple segments and that the desired overall source for synchronization is selected for all networks.

A multiport OC can therefore be used simultaneously as a grandmaster for multiple PTP networks of different types, versions and profiles. The utility of this capability is limited by the requirement to collocate grandmasters and by the fact that there is no clock source signaling defined for this type of interconnection.

D.3 Traceable reference

Any of the clock sources supported by RFC 7273, excepting "local", can be designated "traceable". Doing so indicates that the clock is synchronized to the TAI global time reference. In theory, this allows any traceable clock to be used in place of any other traceable source. In practice, there is no means provided for signaling the accuracy of this synchronization. While accuracy of synchronization over a PTP network can be controllable, the accuracy between two traceable references is not specified and is generally unknown.

An important exception is when both references are GPS. GPS is typically accurate to about 100 nanoseconds, an accuracy comparable to that achieved with PTP. PTP networks with a GPS grandmaster can safely be considered to be operating as identical clocks. The clockClass information broadcast in *Announce* messages by the grandmaster or available through management messaging gives additional information about the synchronization quality beyond what is available in RFC 7273 signaling.

D.4 AVB network as a boundary clock

Because of the high-fidelity clock distribution within an IEEE 802.1AS clock domain, it is possible to use the entire IEEE 802.1AS domain as a distributed boundary clock for any other PTP domain. Implementation and the theory behind this capability is discussed in detail in a paper by Geoffrey M. Garner, Michel Ouellette and Michael Johas Teener. Although this technique doesn't necessarily synchronize the IEEE 802.1AS domain to any of the PTP domains passing through it, it does demonstrate another technique for integrating AVB technology and thus improving interoperability.

Annex E (Informative) – Discovery systems

E.0 General

Although no discovery service is mandated by this standard, clause 9 specifies that discovery services may be used to deliver SDP description (see 8), a SIP URI (see 10.1), or equivalent session data in a different presentation format. Developers of this standard are aware of the following discovery systems whose application to this standard is discussed below.

E.1 Bonjour

Bonjour is a collection of Zero Configuration Networking techniques developed by Apple Inc. and published as open documents. In the context of discovery, the relevant Bonjour techniques are multicast DNS (mDNS) described in an RFC 6762 and DNS service discovery (DNS-SD) described in an RFC 6763.

If these techniques are used for discovery, RTP sessions should be advertised by their SIP URI as described in IETF ***draft-lee-sip-dns-sd-uri.***

E.2 SAP

Use of SAP, or a similar mechanism to distribute SDP descriptions to potential receivers, enables a simplified connection management for multicast streaming (see 10.2).

When SAP is used, SAP version 2 as defined in RFC 2974 should be used. SDP descriptions as defined in clause 8 should be carried in the payload of SAP messages. Multicast sessions should be announced using destination address as specified in RFC 2974 section 3; globally scoped multicast sessions are announced using 224.2.127.254 destination address and administratively scoped multicast sessions are announced using the highest address in their scope.

E.3 Axia Discovery Protocol

The Axia discovery protocol operates over a system-wide dedicated multicast channel. Every device within the delivery scope for these multicasts and having subscribed to the multicast group, can hear the announcements of all other devices.

All devices periodically generate short presence announcements, and at longer interval, description advertisements. The description advertisements include attributes of the device, as well as a list and attributes of the streams they are able to transmit. When description data is updated, a new advertisement is transmitted without waiting for the next scheduled periodic message. Announcement data allows each participant to build a list of the other participants and streams available on the network, to assist with connection management (10).

E.4 Wheatstone WheatnetIP Discovery Protocol

The WheatnetIP discovery protocol operates over a system-wide dedicated IP multicast channel (“announce channel”). Every device within the delivery scope for these multicasts, having subscribed to the multicast group, can hear the announcements of all other devices.

The WheatnetIP protocol uses a *Route Master* elected from within the nodes in the system. Upon entering a WheatnetIP system, a new device issues an entry announcement (*Howdy*) on the announce channel. The Route Master acknowledges this message and initiates a roll call, whereby it asks each device in the system, in turn, to multicast its sources and destinations. The Route Master also assigns each device a range of multicast addresses to use for its output streams. Once the device has collected all of the source and destination information, it becomes active within the system. The Route Master periodically polls each device to ensure that all units listed are still in the system.

Since any device can be a Route Master, and every device contains all of the pertinent information to be a Route Master, loss of the Route Master is seamlessly transitioned to a new Route Master.

E.5 AMWA NMOS Discovery and Registration Specification (IS-04)

The Advanced Media Workflow Association (AMWA) has published a set of Networked Media Open Specifications (NMOS). These include a Discovery and Registration Specification known as AMWA Interface Specification 04 (IS-04).

In IS-04, logical hosts (known as Nodes) register their resources with a Registry. Clients can obtain this information from the Registry through querying or subscription. Clients then access the resources on the Node. For small numbers of Nodes, where a Registry is not available, IS-04 also provides a Peer-to-Peer Discovery mechanism in which each Node directly announces its resources.

The IS-04 specifications (Registration API, Query API, Node API) use protocols aimed at the web, including HTTP(S), WebSockets and JSON. DNS-SD is used for the discovery of the API endpoint. The specifications can be found at: https://github.com/AMWA-TV/nmos-discovery-registration.

E.6 RAVENNA Device and Stream Discovery

RAVENNA technology features a profile supporting AES67 interoperability. RAVENNA devices use DNS-SD over multicast DNS (mDNS) for advertisement and discovery of available devices and streams.

Each device advertises its presence on the network by publishing a corresponding HTTP configuration service allowing users to make configuration changes.

Senders advertise available streams with their individual session names by publishing corresponding RTSP services. Receivers can browse the service announcements and use the published information to interact with the sender's RTSP service, i.e., to receive the SDP data for a particular stream. This method works for unicast and multicast streams.

Detailed specification can be found in "RAVENNA Draft on Operational Principles" at: https://www.ravenna-network.com/resources/.

Annex F (Informative) – Senders using IGMP to request their own streams

Since the original publication of AES67-2013, this standard has required senders to use IGMP to request receipt of the multicast streams they produce. In the current revision this requirement has been removed.

A note in the previous versions of this standard acknowledged that this subscription would not result in senders receiving a copy of the stream data they send. The intent was to avoid issues associated with IGMP snooping implementations that flood multicast data if there is no registered listener to a group. Previous versions also noted that the stated requirement would normally be met by a sender subscription required to receive RTCP receiver reports.

AES67 allows implementations to use IGMPv3 and is silent about the use of the source-specific multicast (SSM) supported in IGMPv3. The silence on SSM was intended to avoid introducing unneeded limitations with respect to AES67 applications using SSM as it potentially came into widespread use.

Experience since 2013 has revealed some important considerations at the intersection of these issues.

Neither RFC 3550 nor AES67 require the implementation of RTCP receiver reports. Their use is far from universal. RTCP receiver reports introduce additional issues in an SSM context. RFC 5760 and RFC 6128 suggest that receivers use unicast messaging to deliver reports associated with SSM streams.

The earlier AES67 requirement for senders to request receipt of their own multicast streams can create additional problems. Some network management systems assume that devices requesting receipt of a stream will actually receive it and could flag ports as overloaded due to these incorrect accounting assumptions.

Also, this requirement when implemented as an IGMPv2 any-source multicast (ASM) join request could have the effect of disabling SSM filtering of the stream on some networks, causing receivers to receive all streams in the group regardless of the IGMPv3-specified unicast source address.

IP multicasting was first defined in RFC 988 in 1986. This RFC includes a description of IGMPv1. The original purpose of IGMP was to communicate with multicast routers to allow multicast packets to be routed beyond the local subnet when needed. On an Ethernet network, IP multicast addresses are mapped to a range of Ethernet multicast MAC addresses. Since Ethernet standardization in 1982, the behavior for forwarding packets with a multicast destination address is to flood the packets to all devices on the network. Ethernet switch vendors realized IP multicast packet forwarding could be optimized by listening to (snooping) IGMP communications. Based on the snooped information, the switch can filter multicast packets and prevent flooding them to network devices that have not registered to receive specific multicast traffic. The implementation of this IGMP snooping feature was and remains manufacturer-specific, and exact behavior, especially in earlier implementations, in corner cases, such as whether to flood multicast traffic with no IGMP-registered receivers, can vary. In 2006, the IETF published RFC 4541, an informational document that helped bring consistency to how IGMP snooping is implemented. This situation has improved over time as network vendors have gained an appreciation of how IGMP snooping is used in real-world applications such as AES67. The flooding issue is not as prevalent as it was in 2013 when AES67 was first published.

Annex G (Normative) – Protocol implementation conformance criteria

G.1 Introduction

The supplier of an implementation of AES67 that is claimed to conform to this standard shall complete the following Protocol Implementation Conformance Statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. The PICS can be used for a variety of purposes by various parties, including the following:

a) As a checklist by the protocol implementer, to reduce the risk of failure to conform to this standard through oversight;

b) As a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the standard PICS proforma, by the supplier and acquirer, or potential acquirer, of the implementation;

c) As a basis for initially checking the possibility of interworking with another implementation by the user, or potential user, of the implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible PICS);

d) As the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation, by a protocol tester.

G.2 Instructions for completing the PICS proforma

G.2.1 General

The first part of the PICS proforma, Implementation Identification and Protocol Summary, is to be completed as indicated with the information necessary to identify fully both the supplier and the implementation.

The main part of the PICS proforma is a fixed-format questionnaire divided into subclauses, each containing one or more items. Answers to the questionnaire items are to be provided in the right-most column, either by simply marking an answer to indicate a restricted choice (usually Yes, No, or Not Applicable), or by entering a value or a set or range of values. (Note that there are some items where two or more choices from a set of possible answers can apply; all relevant choices are to be marked.). As noted in the PICS, some items do not require responses for certain compliant implementations. If no "n/a" response option is given in these cases, the requirement may be excluded from verification.

Each item is identified by a statement number reference in the first column that refers directly to a clause number in this standard; the second column describes the feature; the third column contains the level of the requirement (see table G.1); the fourth column indicates the question to be answered. The fifth column is used to record the answer to these questions.

Table G.1 - Requirement levels

Requirement level	Requirement language
1	Shall (requirement)
2	Should (strong suggestion)

3	May (permission)
0	Informative

Editor's note on statement numbering in AES67-2023:

- *This revision of the standard adds new PICS statements. Some of the new statement numbers, for example "5-X1", are constructed using an additional capital letter "X", to allow to preserve the existing sequential statement numbers, while avoiding plain new numbers appearing out of the natural order. There is no other meaning associated with the letter "X" appearing in statement numbers.*

G.2.2 Stream mode capabilities

Where required by the PICS proforma, detailed stream mode capabilities shall be indicated by means of one or more stream mode capability descriptors.

A stream mode capability descriptor is a text string in the following form:

{<sampling rates>}-{<audio sample encodings>}-{<channel counts>}-{<packet times>}

where

<sampling rates>	comma-separated list of sampling rate designators:	**48000, 96000, 44100**
<audio sample encodings>	comma-separated list of sample encoding designators:	**L16, L24**
<channel counts>	comma-separated list of channel counts, or channel count ranges in double-dot notation, or combination of both. Ranges are indicated with inclusive bounds.	
<packet times>	comma-separated list of packet time designators, or packet time designator ranges in double-dot notation, or combination of both. Ranges are indicated with inclusive bounds.	

A packet time designator is an integer number representing the packet time rounded to integer microseconds. For 48 kHz and 96 kHz sampling rates the designator indicates the effective packet time value. For 44.1 kHz sampling rate the designator indicates the packet time of a 48 kHz packet with an equal sample count.

A packet time designator range represents all possible packet times in single-sample increments.

Packet time designator examples:

Packet content	Packet time designator or range
48 samples at 48 kHz	**1000**
96 samples at 96 kHz	**1000**
48 samples at 44.1 kHz	**1000**
1 to 16 samples at 48 kHz	**21..333**

Stream mode capability descriptor examples:

```
{48000}-{L16,L24}-{1,2,3,4,5,6,7,8}-{1000}
{44100,48000}-{L16}-{1,2}-{1000,4000}
{48000}-{L24}-{1..8,16,32}-{125..250}
```

A specific stream mode capability descriptor indicates support of all stream modes resulting from all combinations of the stream mode attribute values included in it.

All the supplied stream mode capability descriptors combined shall fully and precisely reflect the range of stream modes actually supported by the device. As long as this requirement is met, the supplier of the PICS proforma is free to choose how many and what specific descriptors will be used to represent the entire range. Overlapping descriptors can sometimes be helpful and are generally allowed, but they should be used judiciously.

Stream mode interoperability requirements are defined in G.4.

G.3 PICS proforma

G.3.1 Identification

G.3.1.1 Implementation identification

Supplier (Note 1)	
Contact point for queries about the PICS (Note 1)	
Implementation Name(s) and Version(s) (Notes 1 and 3)	
Other information necessary for full identification - for example, name(s) and version(s) of machines and/or operating system names (Note 2)	

NOTE 1 Required for all implementations.

NOTE 2 May be completed as appropriate in meeting the requirements for the identification.

NOTE 3 The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (for example, Type, Series, Model).

G.3.1.2 Protocol summary

Identification of protocol specification	**AES67-2018, AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability**
Identification of amendments and corrigenda to the PICS proforma that have been completed as part of the PICS	

G.3.2 Synchronization

G.3.2.1 General

Statement Number	Feature	Requirement level	Notes	Supported
4.0-1	Synchronization of time shall be achieved using IEEE 1588-2008 Precision Time Protocol (PTP) or a functionally equivalent subset of IEEE 1588-2019.	1	Mark as supported if the device supports IEEE 1588-2008 Precision Time Protocol (PTP) and the functionally equivalent subset of IEEE 1588-2019. Indicate support per operation mode: • TimeReceiver: Yes/No • TimeTransmitter: Yes/No	 Yes [ ] No [ ] Yes [ ] No [ ]
4.0-2	Devices, except for devices using AVB synchronization (clause 4.3), shall support the IEEE 1588-2008 default profiles as described in annex J.	1	Applicable to non-AVB devices Mark as supported if the device supports IEEE 1588-2008 default profiles as described in annex J	 Yes [ ] No [ ] n/a [ ]
4.0-3	Devices using the default profiles shall use IPv4 encapsulation as described in IEEE 1588-2008 annex D.	1	Mark as supported if the device uses IPv4 encapsulation for the default profiles, as described in IEEE 1588-2008 annex D	Yes [ ] No [ ]

G.3.2.2 IP network synchronization

Statement Number	Feature	Requirement level	Notes	Supported
4.1-1	Devices should support the media profile defined in annex A, to assure adequate performance for all applications on ordinary (PTP-unaware) IP networks.	2	Mark as supported if the device supports the media profile defined in AES67 annex A. Note: Marking the media profile as supported neither requires nor implies support for the optional peer delay mechanism.	 Yes [ ] No [ ]
4.1-2	Devices may use the default profiles on IP networks	0	Device capability information: Indicate what profile is used on IP networks Indicate the selection method: permanent / manual / automatic If automatic, indicate the condition used	 [ ] [ ] [ ]

G.3.2.3 AVB network synchronization

Statement Number	Feature	Requirement level	Notes	Supported
4.3-1	AVB networks may use their native IEEE 802.1AS synchronization profile in preference to the default profiles or media profile.	3	Applicable to AVB devices Mark as supported if the device is capable of using the IEEE 802.1AS synchronization profile. Indicate the selection method: permanent / manual / automatic If automatic, indicate the condition used	Yes [ ] No [ ] n/a [ ] [ ] [ ]

G.3.3 Media clock and RTP clock

Statement Number	Feature	Requirement level	Notes	Supported
5-1, 5-5 combined	The media clock and the network clock shall share the epoch of 1 January 1970 00:00:00 TAI, as defined in IEEE 1588-2008 clause 7.2.2. The value of the media clock shall be 0 at the IEEE 1588-2008 epoch.	1	Mark as supported if the value of the media clock is 0 at the IEEE 1588-2008 epoch.	Yes [ ] No [ ]
5-2, 5-3, 5-5 combined	The rate at which the media clock increments shall be the same as the audio sampling frequency, referenced to the network clock. The value of the media clock shall change from 0 to 1 exactly one sampling period after the epoch, and further increase by 1 after every subsequent audio sampling period.	1	Mark as supported if the value of the media clock changes from 0 to 1 exactly one sampling period after the epoch and further increases by 1 after every subsequent audio sampling period.	Yes [ ] No [ ]
5-4	(skipped statement number)			
5-6	The RTP offset shall be conveyed through session description (see 8.3) on a per-stream basis.	1	Mark as supported if RTP offset is conveyed through session description on a per-stream basis.	Yes [ ] No [ ]
5-X1	RTP timestamps shall be carried in headers of RTP packets containing stream data	1	Mark as supported if proper RTP timestamps are carried in headers of RTP packets containing stream data	Yes [ ] No [ ]
5-7	Media clock timestamps recovered at the receiver from the 32-bit RTP timestamps shall accurately take into account all overflows (rollovers) between the epoch and the current time.	1	Mark as supported if the media clock timestamps recovered at the receiver from the 32-bit RTP timestamps accurately take into account all overflows (rollovers) between the epoch and the current time.	Yes [ ] No [ ]

G.3.4 Transport

G.3.4.1 Network layer

Statement Number	Feature	Requirement level	Notes	Supported
6.1-1	Media packets shall be transported using IP version 4 as defined in RFC 791.	**1**	Mark as supported if media packets are transported using IP version 4 as defined in RFC 791.	Yes [ ] No [ ]
6.1-X1	Senders shall support either unicast, or multicast, or both unicast and multicast media streams.	**1**	Mark as supported if the device can be configured to send, concurrently or non-concurrently, any one of, or both: • unicast streams in at least one mode within the scope of AES67 • multicast streams in at least one mode within the scope of AES67	Yes [ ] No [ ] Indicate what is supported
6.1-X2 (moved from 7.6-1)	Receivers shall be able to receive both unicast and multicast media streams.	**1**	Mark as supported if the device can be configured to receive, concurrently or non-concurrently, both: • multicast streams in at least one mode within the scope of AES67 • unicast streams in at least one mode within the scope of AES67	Yes [ ] No [ ]
6.1-2	A receiver that does not support reassembly shall ignore IP packet fragments.	**1**	Applicable if the receiver does not support IP packet reassembly Mark as supported if the receiver can safely ignore IP packet fragments	Yes [ ] No [ ] n/a [ ]
6.1-3	Senders may set the Don't Fragment flag (DF) bit in the IP header of outgoing media packets.	**3**	Mark as supported if the sender sets the Don't Fragment flag (DF) bit in the IP header of outgoing media packets.	Yes [ ] No [ ]
6.1-4	Senders should terminate transmission of the offending stream in response to receipt of an ICMP message of type 3 ("Destination Unreachable"), code 4 ("fragmentation needed and DF set").	**2**	Mark as supported if the sender terminates transmission of the offending stream in response to receipt of an ICMP message of type 3 ("Destination Unreachable"), code 4 ("fragmentation needed and DF set").	Yes [ ] No [ ]
6.1-5	Multicast messaging shall be accomplished using IP multicasting as described in RFC 1112.	**1**	Mark as supported if multicast messaging is accomplished using IP multicasting as described in RFC 1112.	Yes [ ] No [ ]
6.1-6, 6.1-7, 6.1-8	Moved to 10.2-2, 10.2-3, 10.2-6, respectively			
6.1-9	Deprecated			

6.1-10 (moved from 7.6-2)	Where used, multicasting shall support administratively scoped multicast addresses in the range `239.0.0.0` to `239.255.255.255`	**1**	Mark as supported if verified ability: • to receive streams marked with any multicast destination addresses in the entire range `239.0.0.0` to `239.255.255.255`, and • to mark output packets with any multicast destination addresses in the entire range `239.0.0.0` to `239.255.255.255`, when sending	Yes [ ] No [ ]
6.1-11 (moved from 7.6-3)	The destination address used for a particular multicast stream shall be configurable through the management interface at the sender.	**1**	Mark as supported if the destination address used for a particular multicast stream can be configured through the management interface at the sender.	Yes [ ] No [ ]

G.3.4.2 Quality of service

Statement Number	Feature	Requirement level	Notes	Supported
6.2-1	Devices shall implement the DiffServ method as described in RFC 2474.	**1**	Mark as supported if the device implements the DiffServ method as described in RFC 2474.	Yes [ ] No [ ]
6.2-2	The traffic classes given in table 1 shall be supported.	**1**	Mark as supported if all traffic classes given in table 1 are supported.	Yes [ ] No [ ]
6.2-3	Devices shall tag outgoing traffic with an appropriate DSCP value	**1**	Mark as supported if the device tags outgoing traffic with an appropriate DSCP value	Yes [ ] No [ ]
6.2-4	Devices should use the default values for the DSCP field as given in table 1.	**2**	Mark as supported if the device uses the default values without requiring user configuration.	Yes [ ] No [ ]
6.2-5	Although not required, devices may implement a management interface for DSCP configuration, allowing traffic to be marked with alternative DSCP values.	**3**	Mark as supported if DSCP field can be configured for one or more traffic classes defined in Table 1 through any means. Indicate what configuration means are provided	Yes [ ] No [ ] [ ]
6.2-6	Devices not implementing a management interface shall use the default values exclusively.	**1**	Applicable if the device does not implement a management interface for DSCP configuration Mark as supported if the device uses the default values exclusively.	Yes [ ] No [ ] n/a [ ]
6.2-7	Senders may be configurable to use the same DSCP values for multiple classes.	**3**	Test applicable if DSCP values are configurable Mark as supported if the device implements a method to configure the same DSCP values for multiple classes.	Yes [ ] No [ ] n/a [ ]
6.2-8	Senders may be configurable to use classes in addition to those shown in table 1.	**3**	Mark as supported if the device implements a method to configure additional traffic classes	Yes [ ] No [ ]

6.2-9	Senders should mark outgoing RTCP packets with the same DSCP value as the respective RTP stream packets.	**2**	Mark as supported if the sender marks outgoing RTCP packets with the same DSCP value as the respective RTP stream packets. Note: Senders are not required to implement RTCP	Yes [ ] No [ ] n/a [ ]
6.2-10	Receivers should mark outgoing RTCP packets with the same DSCP value they would use on RTP packets if transmitting a similar stream.	**2**	Mark as supported if the receiver marks outgoing RTCP packets with the same DSCP value it would use on RTP packets if transmitting a similar stream. Note: Receivers are not required to implement RTCP	Yes [ ] No [ ] n/a [ ]
6.2-11	Receivers shall not depend on DSCP markings on received packets.	**1**	Mark as supported if ability of the receiver to handle the stream does not depend on DSCP markings on received packets.	Yes [ ] No [ ]

G.3.4.3 Transport layer

Statement Number	Feature	Requirement level	Notes	Supported
6.3-1	Devices shall use Real-time Transport Protocol as defined in RFC 3550.	**1**	Mark as supported if the device uses Real-time Transport Protocol as defined in RFC 3550.	Yes [ ] No [ ]
6.3-2	Devices shall operate in accordance with RTP Profile for Audio and Video Conferences with Minimal Control as defined in RFC 3551, subject to overrides defined in this standard.	**1**	Mark as supported if the device operates in accordance with RTP Profile for Audio and Video Conferences with Minimal Control as defined in RFC 3551, subject to overrides defined in this standard.	Yes [ ] No [ ]
6.3-3a	Devices shall support the default port allocated for RTP: 5004	**1**	Mark as supported if the device supports the default port allocated for RTP: 5004	Yes [ ] No [ ]
6.3-3b	Devices should use the default port allocated for RTP without requiring user configuration: 5004	**2**	Mark as supported if the device uses the default port allocated for RTP without requiring user configuration: 5004	Yes [ ] No [ ]
6.3-4a	Devices shall support the default port allocated for RTCP: 5005	**1**	Mark as supported if the device supports the default port allocated for RTCP: 5005 Note: Devices are not required to implement RTCP	Yes [ ] No [ ] n/a [ ]
6.3-4b	Devices should use the default port allocated for RTCP without requiring user configuration: 5005	**2**	Mark as supported if the device uses the default port allocated for RTCP without requiring user configuration: 5005 Note: Devices are not required to implement RTCP	Yes [ ] No [ ] n/a [ ]

6.3-5a	Senders may use the following ports: • all ports in the range 1024-49151 (User Ports, also known as Registered Ports, as defined in RFC 6335) • all ports in the range 49152-65535 (Dynamic Ports, also known as Private or Ephemeral Ports, as defined in RFC 6335)	**3**	Mark as supported if the sender is capable to use any ports in the ranges 1024-49151 and 49152-65535, either fixed or configurable through the management interface or another method Indicate if one or both ranges are supported in entirety, or which parts of them are or are not supported.	Yes [ ] No [ ] n/a [ ] [ ]
6.3-5b	Receivers should support the following ports if used by a particular sender: • all ports in the range 1024-49151 (User Ports, also known as Registered Ports, as defined in RFC 6335) • all ports in the range 49152-65535 (Dynamic Ports, also known as Private or Ephemeral Ports, as defined in RFC 6335)	**2**	Mark as supported if the receiver supports any ports in the ranges 1024-49151 and 49152-65535, either fixed or configurable through the management interface or another method. Indicate if one or both ranges are supported in entirety, or which parts of them are or are not supported.	Yes [ ] No [ ] n/a [ ] [ ]
6.3-6	Devices shall use UDP as defined in RFC 768 for transport of RTP.	**1**	Mark as supported if the device uses UDP as defined in RFC 768 for transport of RTP.	Yes [ ] No [ ]
6.3-X1	Senders shall not use silence suppression.	**1**	Mark as supported if the sender does not use silence suppression.	Yes [ ] No [ ]
6.3-7	The maximum allowed RTP payload size shall be 1440 bytes, when no contributing source (CSRC) identifiers or header extensions are included.	**1**	Mark as supported if the RTP payload size does not exceed 1440 bytes, when no contributing source (CSRC) identifiers or header extensions are included.	Yes [ ] No [ ]
6.3-8	Senders should not include contributing source (CSRC) identifiers in the RTP header.	**2**	Mark as supported if the sender does not include contributing source (CSRC) identifier in the RTP header.	Yes [ ] No [ ]
6.3-9	Senders should not add RTP header extensions.	**2**	Mark as supported if the sender does not add RTP header extensions.	Yes [ ] No [ ]
6.3-10	Receivers shall tolerate the presence of CSRC identifiers and header extensions.	**1**	Mark as supported if in presence of CSRC and/or header extensions in an audio stream: • receiver continues normal operation, and no adverse effect on any of its functions is observed, and • no adverse effect on audio received with the same stream is observed	Yes [ ] No [ ]
6.3-10a	If senders include header extensions or CSRCs, the 1440 maximum allowed payload shall be adjusted downwards by the size of the added header material.	**1**	Mark as supported if the RTP payload size does not exceed 1440 bytes adjusted downwards by the size of the added header material, when contributing source (CSRC) identifiers or header extensions are included.	Yes [ ] No [ ] n/a [ ]

6.3-11	Both senders and receivers should transmit RTCP messages as specified in RFC 3550 clause 6.	**2**	Mark as supported if the device transmits RTCP messages as specified in RFC 3550 clause 6.	Yes [ ] No [ ]
6.3-12	Senders and receivers should allocate RTCP bandwidth as recommended in RFC 3551 clause 2 (RTCP report interval).	**2**	Mark as supported if the device allocates RTCP bandwidth as recommended in RFC 3551 clause 2 (RTCP report interval).	Yes [ ] No [ ] n/a [ ]
6.3-13	Unicast senders should monitor connectivity to their respective receivers in such a way as to detect failure of the receiver and stop transmission within 60 seconds.	**2**	Mark as supported if the unicast sender stops transmission to a missing receiver within 60 seconds.	Yes [ ] No [ ]
6.3-14	For monitoring of receivers that do not implement RTCP, senders may use any other monitoring means available to them including any of the following techniques. SIP session timers as described in RFC 4028 SIP OPTION ping (see IETF draft-jones-sip-options-ping) ICMP Echo request (ping)	**0**	Sender capability information: Indicate what monitoring means are used.	[ ]

G.3.5 Encoding and streaming

Statement Number	Feature	Requirement level	Notes	Supported
7	Encoding and streaming		Indicate support per operation mode: • Device supports receiving: • Device supports sending: Note: Devices are allowed to implement only receiving or only sending. When testing such devices, exclude the unimplemented sending or receiving requirements from verification, and mark the respective features as not applicable throughout this section. See G.4 for explanation of qualification criteria.	Yes [ ] No [ ] Yes [ ] No [ ]

G.3.5.1 Payload format and sampling rate

Statement Number	Feature	Requirement level	Notes	Supported
7.1-1 .. 7.1-7	Deprecated. See new criteria in G.3.5.3’			
7.1-8	Although not required, devices may accept 48 kHz connections at all times	3	Mark as supported if the device can accept 48 kHz connections in all implemented operation modes and with any configuration settings.	Yes [ ] No [ ]
7.1-9	Although not required, devices may support multiple sampling rates simultaneously.	3	Mark as supported if more than a single sampling rate can be used simultaneously on different streams Indicate the combinations supported	Yes [ ] No [ ] [ ]
7.1-10	Deprecated. See new criteria in G.3.5.3’			

G.3.5.2 Packet time

G.3.5.2.1 General

Statement Number	Feature	Requirement level	Notes	Supported
7.2.0-1	Senders shall not change packet time for the duration of a session	1	Mark as supported if the sender does not change packet time for the duration of a session	Yes [ ] No [ ] n/a [ ]
7.2.0-2	Although not required, receivers may adapt to packet time changes during a session.	3	Mark as supported if the receiver is able to adapt to packet time changes during a session.	Yes [ ] No [ ] n/a [ ]
7.2.0-3	Receivers should not rely on the presence or accuracy of any packet time description. Receivers should be able to determine packet time based on the timestamps in received packets.	2	Mark as supported if streams are correctly received when no or incorrect packet time description is provided	Yes [ ] No [ ] n/a [ ]
7.2.0-5	Product documentation for a device shall indicate which packet times are supported in send and receive directions.	1	Mark as supported if product documentation indicates which packet times are supported in send and receive directions.	Yes [ ] No [ ]

G.3.5.2.2 Required packet time

Statement Number	Feature	Requirement level	Notes	Supported
7.2.1-1 .. 7.2.1-6	Deprecated. See new criteria in G.3.5.3’			

7.2.1-7	Although not required, devices may accept “1 millisecond” (48-sample at 48 kHz or 44,1 kHz, or 96-sample at 96 kHz) connections at all times.	**3**	Mark as supported if the device can accept “1 millisecond” connections in all implemented operation modes and with any configuration settings.	Yes [ ] No [ ]
7.2.1-8	Although not required, devices may support multiple packet times simultaneously.	**3**	Mark as supported if multiple packet times can be used simultaneously Indicate which packet times are simultaneously supported in which mode	Yes [ ] No [ ] [ ]

G.3.5.2.3 Recommended packet times

Statement Number	Feature	Requirement level	Notes	Supported
7.2.2-1 .. 7.2.2-2	Deprecated. See new criteria in G.3.5.3’			

G.3.5.3 Stream channel count

Statement Number	Feature	Requirement level	Notes	Supported
7.3-1 .. 7.3-4	Deprecated. See new criteria in G.3.5.3’			

G.3.5.4 Stream modes

NOTE Statement numbers in this section are forward references to the clause G.4 of this annex.

Statement Number	Feature	Requirement level	Notes	Supported
G.4-1	Stream mode interoperability for senders – required	**1**	Mark as supported if the device supports all the following capabilities as defined in Table G.2: • T1, “T-48k-1ms”	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what is supported
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: T-48k-1ms: **`{48000}-{L16,L24}-{1,2,3,4,5,6,7,8}-{1000}`**			

G.4-2	Stream mode interoperability for senders – recommended	**2**	Mark as supported if the device supports any one or more of the following capabilities as defined in Table G.2: • T2, "T-96k-1ms" • T3, "T-44k-1ms" • T4, "T-48k-non1ms" • T5, "T-96k-non1ms" • T6, "T-44k-non1ms"	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what is supported
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: T-96k-1ms: **{96000}-{L24}-{1,2,3,4,5,6,7,8}-{1000}** T-44k-1ms: **{44100}-{L16}-{1,2,3,4,5,6,7,8}-{1000}** T-48k-non1ms: **{48000}-{L16,L24}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}** T-96k-non1ms: **{96000}-{L24}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}** T-44k-non1ms: **{44100}-{L16}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}**			
G.4-3	Stream mode interoperability for senders – other	**3**	Mark as supported if the device supports any one or more of the following capabilities as defined in Table G.2: • T7, "T-other"	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what else is supported in addition to stream modes already covered by capabilities T1 to T6
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: T-other: **{44100,48000,96000}-{L16,L24}-{1,2,3,4,5,6,7,8,...}-{125,250,333,1000,4000,...}**			
G.4-4	Stream mode interoperability for receivers – required	**1**	Mark as supported if the device supports all the following capabilities as defined in Table G.3: • R1, "R-48k-1ms"	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what is supported
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: R-48k-1ms: **{48000}-{L16,L24}-{1,2,3,4,5,6,7,8}-{1000}**			

G.4-5	Stream mode interoperability for receivers – recommended	**2**	Mark as supported if the device supports any one or more of the following capabilities as defined in Table G.3: • R2, "R-96k-1ms" • R3, "R-44k-1ms" • R4, "R-48k-non1ms" • R5, "R-96k-non1ms" • R6, "R-44k-non1ms"	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what is supported
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: R-96k-1ms: **{96000}-{L24}-{1,2,3,4,5,6,7,8}-{1000}** R-44k-1ms: **{44100}-{L16}-{1,2,3,4,5,6,7,8}-{1000}** R-48k-non1ms: **{48000}-{L16,L24}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}** R-96k-non1ms: **{96000}-{L24}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}** R-44k-non1ms: **{44100}-{L16}-{1,2,3,4,5,6,7,8}-{125,250,333,4000}**			
G.4-6	Stream mode interoperability for receivers – other	**3**	Mark as supported if the device supports any one or more of the following capabilities as defined in Table G.3: • R7, "R-other"	Yes [ ] No [ ] n/a [ ] Use the next row to indicate what else is supported in addition to stream modes already covered by capabilities R1 to R6
	Provide the actual stream mode capability descriptors here. Edit or replace the templates below: R-other: **{44100,48000,96000}-{L16,L24}-{1,2,3,4,5,6,7,8,...}-{125,250,333,1000,4000,...}**			

G.3.5.5 Link offset

Statement Number	Feature	Requirement level	Notes	Supported
7.4-1	RTP packets are marked with origination timestamps in the timestamp field (RFC 3550 clause 5.1), referenced at ingress to the sender network system.	1	Mark as supported if the origination timestamps in RTP packets appear between: • time when audio producing the first sample in the packet enters the sender, and • time when the network packet carrying it leaves the sender, minus one packet time. Note: Accuracy of conformance verification is limited by the inability to observe events inside the device.	Yes [ ] No [ ]
7.4-2	A receiver should attempt to maintain a constant link offset.	2	Mark as supported if the receiver: • maintains a constant link offset during stable reception • recovers to the previous link offset when resuming reception after occasional stream interruptions Note: The receiver could change the link offset as a reaction to changes in network conditions, for example, changed propagation delay or jitter. Such reaction is not a reason to disqualify the device in this test.	Yes [ ] No [ ]
7.4-3	The link offset and any changes in link offset should be retrievable from the management entity of the receiver, if present.	2	Applicable if there is a management entity in the device. Mark as supported if the link offset and any changes in link offset are retrievable from the management entity of the receiver.	Yes [ ] No [ ] n/a [ ]

G.3.5.6 Sender timing and receiver buffering

Statement Number	Feature	Requirement level	Notes	Supported
7.5-1	Receivers shall have a buffer capacity at least 3 times the packet time.	1	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]
7.5-2	Receivers should have a buffer capacity at least 20 times the packet time or 20 ms whichever is smaller.	2	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]

7.5-3	Receiver capability information	**0**	Indicate the maximum supported size of the receive buffer in milliseconds. If varied by the sampling rate or other factors, list all cases or describe the dependency rules.	[ ]
7.5-4	Senders should transmit at the nominal transmission time with a variation of 1 packet time or less.	**2**	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]
7.5-5	Senders shall transmit data at the nominal transmission time with a variation of no more than 17 packet times or 17 ms whichever is smaller.	**1**	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]
7.5-6	Sender capability information	**0**	Indicate the transmission timing variation in milliseconds	[ ]

G.3.5.7 Multicasting

Statement Number	Feature	Requirement level	Notes	Supported
7.6-1	Moved to 6.1-1b			
7.6-2, 7.6-3	Moved to 6.1-10, 6.1-11			

G.3.6 Session description

G.3.6.1 General

Statement Number	Feature	Requirement level	Notes	Supported
8.0-1	SDP as specified in RFC 8866 shall be used to represent the sessions for connection management	1	Mark as supported when verified for SDP output generation on the sender device and SDP input interpretation on the receiver device with at least one connection management mechanism, which may be different for unicast connections and multicast connections. NOTE Full verification of an AES67 implementation against every normative clause of RFC 8866, as well as specification of such verification criteria in this standard, are considered needlessly complicated and time consuming for the goals of AES67. Based on common practice, off-the-shelf test suites are found to be a good alternative. As an example, the following freely available test suite for generated SDP has been designed with ST 2110-30:2017 and AES67 in mind: https://github.com/AMWA-TV/sdpoker Other similar test suites could exist and bring satisfactory results too.	Yes [ ] No [ ] Use the next row to indicate the test tools used and test configuration details
	Test tool name: <indicate the name or "none"> Download source: <indicate the URL or "none"> Test suite configuration description: <provide the description or indicate "none"> <Add the same information about other test tools, if such were used, information about peer devices used in the tests, and/or other information>			

G.3.6.2 Packet time

Statement Number	Feature	Requirement level	Notes	Supported
8.1-1	The packet time descriptions shall be given with error less than half a sample period.	1	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]
8.1-2	Where milliseconds fractional part is not required to accurately convey packet time, it shall be omitted from signaling	1	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]
8.1-3	Packet times in descriptions shall be interpreted as a dotted decimal representation.	1	Mark as supported if verified for all supported packet times	Yes [ ] No [ ]

8.1-4	Values and representations beyond those enumerated in table 4 shall be correctly interpreted.	**1**	Mark as supported if the device correctly interprets 1) values beyond those enumerated in table 4 2) the following alternative decimal representations of packet times, as long as they do not produce an incorrect calculated integer sample count: • containing unnecessary decimal point • containing unnecessary precision digits, whether they are correct or incorrect with respect to the exact fractional value	Yes [ ] No [ ]
8.1-5	Descriptions shall include a **`ptime`** attribute indicating the desired packet time.	**1**	Test applicable both to multicast signaling and to SIP negotiation Mark as supported if **`ptime`** is given in generated SDP in all relevant tests	Yes [ ] No [ ]
8.1-6	If more than one packet time is supported, a **`maxptime`** indicating the maximum packet time permitted shall be provided.	**1**	Test applicable to SIP negotiation Test applicable to devices that support more than one packet time Mark as supported if **`maxptime`** is given in generated SDP in all relevant tests	Yes [ ] No [ ] n/a [ ]
8.1-7	To override the implied assumption that a shorter packet time is always the preferred packet time, the capability negotiation attributes of RFC 5939 may be used to enumerate the supported packet times and order of preference.	**3**	Test applicable to SIP negotiation Test applicable to devices that support more than one packet time Mark as supported if the capability negotiation attributes of RFC 5939 are used to enumerate the supported packet times and order of preference in at least one test scenario.	Yes [ ] No [ ] n/a [ ]
8.1-8	If the range of packet times supported includes more than two of the standard packet times (table 2), the description should use the capability negotiation attributes of RFC 5939 to enumerate the supported packet times and order of preference.	**2**	Test applicable to SIP negotiation Test applicable to devices that support more than two of the standard packet times Mark as supported if the capability negotiation attributes of RFC 5939 are used to enumerate the supported packet times and order of preference in all relevant tests.	Yes [ ] No [ ] n/a [ ]

G.3.7 Clock source

Statement Number	Feature	Requirement level	Notes	Supported
8.2-1	The network clock source for each stream described shall be specified with one or more **`ts-refclk`** attributes as specified in RFC 7273.	1	Test applicable both to multicast signaling and to SIP negotiation Mark as supported if included in output SDP and correctly interpreted on input	Yes [ ] No [ ]
8.2-1a	When IEEE 1588-2019 is in use, AES67 devices shall indicate "**`IEEE1588-2008`**".	1	Mark as supported if IEEE 1588-2019 is indicated as "**`IEEE1588-2008`**"	Yes [ ] No [ ]
8.2-2	Signaling for RTP streams referenced to IEEE 1588-2008 shall indicate at least one of the following: a) both GMID and PTP domain, or b) property "traceable".	1	Test applicable both to multicast signaling and to SIP negotiation Mark as supported if included in output SDP and correctly interpreted on input	Yes [ ] No [ ]
8.2-3	Receivers should attempt to connect to senders if they are using the same PTP domain and the same GMID clock reference as the sender.	2	Mark as supported if receiver connects to a sender, which uses the same PTP domain and GMID clock reference.	Yes [ ] No [ ]
8.2-3a	Receivers should attempt to connect to senders if both are using traceable time.	2	Mark as supported if receiver connects to a sender, when both are using traceable time.	Yes [ ] No [ ]
8.2-4	Receivers should not attempt to connect to senders if they are using a different PTP domain for their clock reference than the sender, and at least one of the used references is not known to be traceable.	2	Mark as supported if receiver does not attempt connection to a sender which uses a different PTP domain, when the used reference is not known to be traceable for one or both of them.	Yes [ ] No [ ]
8.2-5	Receivers may attempt to make a connection in case of PTP domain match and mismatched GMID even if one or both used references are not known to be traceable.	3	Mark as supported if receiver connects to a sender which uses the same PTP domain and different GMID, when one or both used references are not known to be traceable.	Yes [ ] No [ ]
8.2-6	Receivers attempting to make a connection in case of PTP domain match and mismatched GMID, when one or both used references are not known to be traceable, should be prepared for possible synchronization failure.	2	Mark as supported if a receiver attempting to make a connection in case of PTP domain match and mismatched GMID, when one or both used references are not known to be traceable, is capable of detecting and reacting to the possible synchronization failure by invoking a signaling mechanism and/or disconnecting the faulty connection.	Yes [ ] No [ ]
8.2-7	(unused statement number)			
8.2-8	Senders should monitor for changes in their synchronization status during transmission and update their clock source description when a change is detected	2	Mark as supported if clock source description is updated in the SDP generated by Sender	Yes [ ] No [ ]

8.2-9	Receivers should monitor for changes in their synchronization status during reception. When their synchronization status changes, receivers should reevaluate their ability to continue receiving according to the recommendations in this clause	**0**	See test cases 8.2-9a and 8.2-9b.	
8.2-9a	Test applicable to receivers that refuse connections in case of PTP domain mismatch when one or both used references are not known to be traceable (case 8.2-4).	**2**	Mark as supported if the receiver disconnects from a previously connected sender when after its synchronization status change the PTP domain is different and one or both used references are not known to be traceable.	Yes [ ] No [ ] n/a [ ]
8.2-9b	Test applicable to receivers that refuse connections in case of PTP domain match and mismatched GMID when one or both used references are not known to be traceable (case 8.2-5).	**2**	Mark as supported if the receiver disconnects from a previously connected sender when after its synchronization status change the PTP domain is matched, GMID is different, and one or both used references are not known to be traceable.	Yes [ ] No [ ] n/a [ ]
8.2-10	Receivers should monitor for updated descriptions from the sender during reception. When an updated description is received from the sender, receivers should reevaluate their ability to continue receiving according to the recommendations in this clause	**0**	See test cases 8.2-10a and 8.2-10b.	
8.2-10a	Test applicable to receivers that refuse connections in case of PTP domain mismatch when one or both used references are not known to be traceable (case 8.2-4).	**2**	Mark as supported if the receiver disconnects from a previously connected sender when after receiving an updated description from the sender the PTP domain is different and one or both used references are not known to be traceable.	Yes [ ] No [ ] n/a [ ]
8.2-10b	Test applicable to receivers that refuse connections in case of PTP domain match and mismatched GMID when one or both used references are not known to be traceable (case 8.2-5).	**2**	Mark as supported if the receiver disconnects from a previously connected sender when after receiving an updated description from the sender the PTP domain is matched, GMID is different, and one or both used references are not known to be traceable.	Yes [ ] No [ ] n/a [ ]
8.2-11	Receivers are not required to terminate reception on detection of synchronization signaling mismatch in an ongoing stream.	**0**	Receiver capability information: Indicate what behavior is implemented	[ ]
8.2-12	Receivers continuing reception on detection of synchronization signaling_mismatch in an ongoing stream should be prepared for possible synchronization failure.	**2**	Test applicable to receivers that would attempt to continue reception under the given conditions. Mark as supported if the receiver, which is continuing receiving under these conditions, is capable of detecting and reacting to the possible synchronization failure by invoking a signaling mechanism and/or disconnecting the faulty connection.	Yes [ ] No [ ] n/a [ ]

G.3.7.1 RTP and media clock

Statement Number	Feature	Requirement level	Notes	Supported
8.3-1	The relationship of media clock to RTP clock shall be described for each stream with an `a=mediaclk:direct=<offset>` attribute as specified in RFC 7273 clause 5.2.	1	Mark as supported when verified for both SDP output generation and SDP input interpretation.	Yes [ ] No [ ]

G.3.7.2 Payload types

Statement Number	Feature	Requirement level	Notes	Supported
8.4-1	The receiver shall determine the payload type using the `rtpmap` attribute; it shall not assume any fixed relationship between payload type value and payload type.	1	Mark as supported when verified for SDP input interpretation.	Yes [ ] No [ ]

G.3.8 Discovery

Statement Number	Feature	Requirement level	Notes	Supported
9-1	Devices may implement one or more discovery services including Bonjour, SAP and others.	3	Mark as supported if at least one discovery service is implemented Indicate what discovery services are implemented	Yes [ ] No [ ] [ ]

G.3.9 Connection management

G.3.9.1 Unicast connections

Statement Number	Feature	Requirement level	Notes	Supported
10.1-1	Devices should support connection management for unicast streams using the Session Initiation Protocol (SIP) as defined in RFC 3261.	2	Mark as supported if SIP can be used to set up a connection. Verification of completeness of the SIP implementation is outside the scope of this PICS.	Yes [ ] No [ ]
10.1-2	Devices may additionally support connection management for unicast streams using other protocols, or through the use of a management interface.	3	Mark as supported if any additional connection management method for unicast streams is supported. Indicate the methods that are supported.	Yes [ ] No [ ] [ ]

G.3.9.1.1 SIP URI

Statement Number	Feature	Requirement level	Notes	Supported
10.1.1-1	The "**sip:** " URI form shall be used as defined in RFC 3261	1	Test applicable to connection setup using SIP. Mark as supported if the "**sip:** " URI form is used as defined in RFC 3261	Yes [ ] No [ ]
10.1.1-2	Devices may support TLS and "**sips:** ", but they are not required for interoperability.	0	Device capability information: Indicate support of TLS and "**sips:**"	Yes [ ] No [ ]

G.3.9.1.2 Server and serverless modes

Statement Number	Feature	Requirement level	Notes	Supported
10.1.2-1	In addition to serverless mode, devices using SIP shall be able to operate in a normal SIP environment featuring servers – attempt to discover, and register with, SIP registration servers, and respond to messages originating from servers.	1	Test applicable to connection setup using SIP. Mark as supported if devices can operate in a normal SIP environment featuring servers attempt to discover, and register with, SIP registration servers, and respond to messages originating from servers.	Yes [ ] No [ ]
10.1.2-2	A device using SIP under this standard shall respond to SIP messages sent directly from other user agents.	1	Test applicable to connection setup using SIP. Mark as supported if the device responds to SIP messages sent directly from other user agents. Note: Serverless mode	Yes [ ] No [ ]

G.3.9.1.3 User-Agent

Statement Number	Feature	Requirement level	Notes	Supported
10.1.3-1	Devices should include a User-Agent header field in REGISTER and INVITE messages.	2	Mark as supported if the device includes a User-Agent header field in REGISTER and INVITE messages.	Yes [ ] No [ ]

G.3.9.1.4 Format negotiation

Statement Number	Feature	Requirement level	Notes	Supported
10.1.4-1	The standard offer/answer model as described in RFC 3264 shall be used to negotiate the encoding format for a connection.	1	Mark as supported if the standard offer/answer model as described in RFC 3264 is used to negotiate the encoding format for a connection.	Yes [ ] No [ ]

G.3.9.1.5 Packet time negotiation

Note: No new requirements are specified in this clause. The relevant functionality is covered by tests in clauses 8.1-7 and 8.1-8

G.3.9.2 Multicast connections

Statement Number	Feature	Requirement level	Notes	Supported
10.2-1	Multicast connection management may be accomplished without use of a connection management protocol.	3	Mark as supported if a connection can be set up without a connection management protocol Indicate what connection management method is implemented	Yes [ ] No [ ] [ ]
10.2-2 (moved from 6.1-6)	All devices shall support IGMPv2 as defined in RFC 2236.	1	Mark as supported if the device supports IGMPv2 as defined in RFC 2236.	Yes [ ] No [ ]
10.2-3 (moved from 6.1-7)	Devices should support IGMPv3 as defined in RFC 3376.	2	Mark as supported if the device supports IGMPv3 as defined in RFC 3376	Yes [ ] No [ ]
10.2-4	Devices supporting IGMPv3 (and consequently IGMPv2), as delivered from the manufacturer, when connecting to a network, shall start in the IGMPv3 mode and then follow the IGMPv3 rules according to the actual network conditions.	1	Applicable to devices supporting IGMPv3 Mark as supported if the device, when connecting to a network, starts in the IGMPv3 mode and then follows the IGMPv3 rules according to the actual network conditions	Yes [ ] No [ ] n/a [ ]
10.2-5	Devices may implement a management interface allowing to set them to IGMPv2 mode.	3	Applicable to devices supporting IGMPv3 Mark as supported if the device implements a management interface allowing to set it to IGMPv2 mode	Yes [ ] No [ ] n/a [ ]
10.2-6 (moved from 6.1-8)	Devices shall use IGMP to request reception of any multicasts required.	1	Mark as supported if the device uses IGMP to request reception of any multicasts required.	Yes [ ] No [ ]

G.3.10 Media profile (Normative)

G.3.10.1 General

Statement Number	Feature	Requirement level	Notes	Supported
A.0-1	Devices may choose to implement profiles in addition to or instead of the media profile defined in this annex.	3	Mark as supported if the device implements any other than the default and media profiles Indicate all profiles that are implemented in addition to the mandatory	Yes [ ] No [ ] [ ]

			default profiles	

G.3.11 Media profile

All tests in this section are applicable only if the media profile is supported (test item 4.1-1)

G.3.11.1 Identification

Statement Number	Feature	Requirement level	Notes	Supported
A.2.1-1	Profile identification: PTP profile for media applications. Version 1.0. Profile identifier: **00-0B-5E-00-01-00**.	1	Applicable to devices using the AES67 media profile Mark as supported if profile identifier is indicated in **profileIdentity** field of management messages (IEEE 1588-2008 15.5.3.1.2.10) as "**00-0B-5E-00-01-00**"	Yes [ ] No [ ] n/a [ ]

G.3.11.2 Table A.1 - PTP attribute values

Statement Number	Feature	Requirement level	Notes	Supported
A.2.2-1	Nodes shall implement all requirements in this PTP profile that specify default values or choices such that these default values or choices apply without requiring user configuration, as delivered from the manufacturer.	1	Mark as supported if the device implements all requirements in this PTP profile that specify default values or choices such that these default values or choices apply without requiring user configuration, as delivered from the manufacturer.	Yes [ ] No [ ]
A.2.2-2	For each defined range, manufacturers may allow wider ranges.	3	Mark as supported if a wider range is allowed for at least one of the PTP attributes	Yes [ ] No [ ]
A.2.2-3	**defaultDS.domainNumber** default 0	1	Mark as supported if the default initialization value of **defaultDS.domainNumber** is 0	Yes [ ] No [ ]
A.2.2-4	**defaultDS.domainNumber** range 0-127	1	Mark as supported if values 0 to 127 are allowed	Yes [ ] No [ ]
A.2.2-5	**portDS.logAnnounceInterval** default 1	1	Mark as supported if the default initialization value of **portDS.logAnnounceInterval** is 1	Yes [ ] No [ ]
A.2.2-6	**portDS.logAnnounceInterval** range 0 to 4	1	Mark as supported if the required range is supported Indicate the wider range, if allowed by the implementation	Yes [ ] No [ ] [ ]

A.2.2-7	**portDS.logSyncInterval** default –3	**1**	Mark as supported if the default initialization value of **portDS.logSyncInterval** is –3	Yes [ ] No [ ]
A.2.2-8	**portDS.logSyncInterval** range -4 to +1	**1**	Mark as supported if the required range is supported Indicate the wider range, if allowed by the implementation	Yes [ ] No [ ] [ ]
A.2.2-9	**portDS.logMinDelayReqInterval** default 0	**1**	Mark as supported if the default initialization value of **portDS.logMinDelayReqInterval** is 0	Yes [ ] No [ ]
A.2.2-10	**portDS.logMinDelayReqInterval** range -3 to 5 -- OR -- <**portDS.logSyncInterval**> to <**portDS.logSyncInterval**>+5 whichever is more restrictive	**1**	Mark as supported if the required range is supported Indicate the wider range, if allowed by the implementation	Yes [ ] No [ ] [ ]
A.2.2-11	**portDS.logMinPdelayReqInterval** default 0	**1**	Mark as supported if the default initialization value of **portDS.logMinPdelayReqInterval** is 0	
A.2.2-12	**portDS.logMinPdelayReqInterval** range 0 to 5	**1**	Mark as supported if the required range is supported Indicate the wider range, if allowed by the implementation	Yes [ ] No [ ] [ ]
A.2.2-11	**portDS.announceReceiptTimeout** default 3	**1**	Mark as supported if the default initialization value of **portDS.announceReceiptTimeout** is 3	Yes [ ] No [ ]
A.2.2-14	**portDS.announceReceiptTimeout** range 2 to 10	**1**	Mark as supported if the required range is supported Indicate the wider range, if allowed by the implementation	Yes [ ] No [ ] [ ]
A.2.2-15	**defaultDS.priority1** default 128	**1**	Mark as supported if the default initialization value of **defaultDS.priority1** is 128	Yes [ ] No [ ]
A.2.2-16	**defaultDS.priority2** default 128	**1**	Mark as supported if the default initialization value of **defaultDS.priority2** is 128	Yes [ ] No [ ]

A.2.2-17	**defaultDS.timeReceiverOnly** default FALSE	**1**	Test applicable to devices that support timeReceiver/timeTransmitter mode configuration Mark as supported if the default initialization value of **defaultDS.timeReceiverOnly** is FALSE	Yes [ ] No [ ] [ ]
A.2.2-18	**transparentClockdefaultDS.primaryDomain** default 0	**1**	Mark as supported if the default initialization value of **transparentClockdefaultDS.primaryDomain** is 0	Yes [ ] No [ ]
A.2.2-19	τ default 1,0 s	**1**	Mark as supported if the default initialization value of τ (see IEEE 1588-2008 clause 7.6.3.2) is 1,0 s	Yes [ ] No [ ]

G.3.11.3 Table A.2 values

Statement Number	Feature	Requirement level	Notes	Supported
A.2.2-20	Holdover specification for all **clockClass** specifications are +/- 5 % of a 96 kHz word-clock period.	**0**	Applies to all tests referring to the holdover mode	
A.2.2-21	**defaultDS.clockClass** 6 - synchronized to a primary reference time source (for example, GPS)	**0**	Applicable to devices supporting primary reference time source (for example, GPS) Mark as supported if the device supports **clockClass** 6 according to IEEE 1588-2008 clause 7.6.2.4 and table 5	Yes [ ] No [ ] n/a [ ]
A.2.2-22	**defaultDS.clockClass** 7 - **defaultDS.clockClass** 6 in holdover	**0**	Applicable to devices supporting primary reference time source (for example, GPS) Mark as supported if the device supports **clockClass** 7 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]
A.2.2-23	**defaultDS.clockClass** 13 - synchronized to an external media clock source	**0**	Applicable to devices supporting external media clock sources Mark as supported if the device supports **clockClass** 13 according to IEEE 1588-2008 clause 7.6.2.4 and table 5.	Yes [ ] No [ ] n/a [ ]
A.2.2-24	**defaultDS.clockClass** 14 - **defaultDS.clockClass** 13 in holdover	**0**	Applicable to devices supporting external media clock sources Mark as supported if the device supports **clockClass** 14 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]

A.2.2-25	**defaultDS.clockClass** 52 - degradation alternative A for a clock of **clockClass** 7 that is not within holdover specification	**0**	Applicable to devices supporting primary reference time source (for example, GPS) Mark as supported if the device supports **clockClass** 52 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]
A.2.2-26	**defaultDS.clockClass** 58 - degradation alternative A for a clock of **clockClass** 14 that is not within holdover specification	**0**	Applicable to devices supporting external media clock sources Mark as supported if the device supports **clockClass** 58 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]
A.2.2-27	**defaultDS.clockClass** 150 - A clock whose frequency is synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) and whose time has been previously synchronized to a primary reference time source	**0**	Applicable to devices which support a frequency reference and are capable of synchronizing to a primary reference time source either directly or indirectly as a PTP timeReceiver. Mark as supported if the device: • in absence of the intended primary reference time source, can synchronize (fall back) to a frequency reference source • indicates **defaultDS.clockClass** value 150, when, being previously synchronized to a primary reference time source, is now synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) • conforms to the relevant requirements of IEEE 1588-2008 clause 7.6.2.4 and table 5	Yes [ ] No [ ] n/a [ ]

A.2.2-28	**defaultDS.clockClass** 158 - A clock whose frequency is synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) and whose time has been previously synchronized to a primary reference time source	**0**	Applicable to devices which support a frequency reference and are capable of synchronizing to a primary reference time source either directly or indirectly as a PTP timeReceiver. Mark as supported if the device: • in absence of the intended primary reference time source, can synchronize to a frequency source (fall back to syntonized operation) • indicates **defaultDS.clockClass** value 158, when, being previously synchronized to a primary reference time source, is now synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) • conforms to the relevant requirements of IEEE 1588-2008 clause 7.6.2.4 and table 5	Yes [ ] No [ ] n/a [ ]
A.2.2-29	**defaultDS.clockClass** 166 - A clock of unspecified tolerance that has been previously synchronized to a primary reference time source	**0**	Applicable to devices which are capable of synchronizing to a primary reference time source either directly or indirectly as a PTP timeReceiver. Mark as supported if the device: • indicates **defaultDS.clockClass** value 166, when, being previously synchronized to a primary reference time source, is now of unspecified tolerance • conforms to the relevant requirements of IEEE 1588-2008 clause 7.6.2.4 and table 5	Yes [ ] No [ ] n/a [ ]
A.2.2-30	**defaultDS.clockClass** 187 - degradation alternative B for a clock of **clockClass** 7 that is not within holdover specification	**0**	Applicable to devices supporting primary reference time source (for example, GPS) Mark as supported if the device supports **clockClass** 187 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]
A.2.2-31	**defaultDS.clockClass** 193 - degradation alternative B for a clock of **clockClass** 14 that is not within holdover specification	**0**	Applicable to devices supporting external media clock sources Mark as supported if the device supports **clockClass** 193 according to IEEE 1588-2008 clause 7.6.2.4 and table 5, and the holdover specification given in AES67-2015 clause A.2.2 (PICS clause A.2.2-20).	Yes [ ] No [ ] n/a [ ]

A.2.2-32	**defaultDS.clockClass** 220 - A clock whose frequency is synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) and whose time has not been previously synchronized to a primary reference time source	**0**	Applicable to devices supporting frequency reference source Mark as supported if the device: • indicates **defaultDS.clockClass** value 220, when, not being previously synchronized to a primary reference time source, is synchronized to a reference with ± 1 ppm frequency accuracy (for example, a Grade-1 DARS according to AES11-2009) • conforms to the relevant requirements of IEEE 1588-2008 clause 7.6.2.4 and table 5 •	Yes [ ] No [ ] n/a [ ]
A.2.2-33	**defaultDS.clockClass** 228 - A clock whose frequency is synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) and whose time has not been previously synchronized to a primary reference time source	**0**	Applicable to devices supporting frequency reference source Mark as supported if the device: • indicates **defaultDS.clockClass** value 228, when, not being previously synchronized to a primary reference time source, is synchronized to a reference with ± 10 ppm frequency accuracy (for example, a Grade-2 DARS according to AES11-2009) • conforms to the relevant requirements of IEEE 1588-2008 clause 7.6.2.4 and table 5 •	Yes [ ] No [ ] n/a [ ]
A.2.2-34	**defaultDS.clockClass** 248 – default	**0**	Applicable to devices supporting timeTransmitter operation, if none of the other **clockClass** definitions apply Mark as supported if the device supports **clockClass** 248 according to IEEE 1588-2008 clause 7.6.2.4 and table 5.	Yes [ ] No [ ] n/a [ ]
A.2.2-35	**defaultDS.clockClass** 255 - a timeReceiver-only clock	**0**	Applicable to timeReceiver-only clocks (**defaultDS.timeReceiverOnly** is TRUE) Mark as supported if the device supports **clockClass** 255 according to IEEE 1588-2008 clause 7.6.2.4 and table 5.	Yes [ ] No [ ] n/a [ ]

G.3.11.4 PTP options

Statement Number	Feature	Requirement level	Notes	Supported
A.2.3-1	Devices may implement any options of IEEE 1588-2008 clause 17.	3	Mark as supported if at least one option is implemented Indicate the implemented options	Yes [ ] No [ ] [ ]
A.2.3-2	Devices may implement unicast negotiation as defined in clause 16.1 of IEEE 1588-2008.	3	Mark as supported if the device implements unicast negotiation as defined in clause 16.1 of IEEE 1588-2008.	Yes [ ] No [ ]
A.2.3-3	IEEE 1588-2008 clause 16.1 and 17 options shall be inactive unless specifically activated by a management procedure.	1	Mark as supported if IEEE 1588-2008 clause 16.1 and 17 options are inactive unless specifically activated by a management procedure.	Yes [ ] No [ ]
A.2.3-4	Node management shall implement the management message mechanism of IEEE 1588-2008.	1	Mark as supported if node management implements the management message mechanism of IEEE 1588-2008.	Yes [ ] No [ ]
A.2.3-5	The best timeTransmitter clock algorithm shall be the algorithm specified by IEEE 1588-2008 clause 9.3.2.	1	Mark as supported if the best timeTransmitter clock algorithm is the algorithm specified by IEEE 1588-2008 clause 9.3.2.	Yes [ ] No [ ]
A.2.3-6	The default path-delay measurement mechanism shall be the delay request-response mechanism specified by IEEE 1588-2008.	1	Mark as supported if the default path-delay measurement mechanism is the delay request-response mechanism specified by IEEE 1588-2008.	Yes [ ] No [ ]
A.2.3-7	The peer delay mechanism should also be implemented.	2	Mark as supported if the peer delay mechanism is also implemented.	Yes [ ] No [ ]

G.3.11.5 Clock physical requirements

Statement Number	Feature	Requirement level	Notes	Supported
A.2.4-1	Clocks shall meet requirements of Grade 2 DARS set forth in AES11 clause 5.2.	1	Mark as supported if the clock meets requirements of Grade 2 DARS set forth in AES11 clause 5.2.	Yes [ ] No [ ]
A.2.4-2	Clocks may conform to the Grade 1 requirements.	3	Mark as supported if the clock conforms to the Grade 1 requirements.	Yes [ ] No [ ]

G.4 Qualification criteria for encoding and streaming capabilities

A stream mode is defined by a specific combination of the audio sampling rate, audio sample encoding, stream channel count, and packet time values.

Devices implement specific stream modes, which determines their level of conformance. All requirements as used for conformance testing purposes are listed in Tables G.2 and G.3, for senders and receivers respectively, and enumerated as “Tx” for stream transmit and “Rx” for stream receive capabilities.

Each table row defines a single undividable requirement, which includes support of all stream modes resulting from all combinations of the stream mode attribute values indicated in the respective row.

Stream mode attribute values, which are outside the scope of this standard, are represented using a common option name “other”.

Stream modes that would exceed the maximal allowed RTP payload size 1440 bytes shall be excluded from verification.

Table G.2 - Stream mode interoperability requirements for senders

Requirement identifier / name	Sampling rate	Audio sample encoding	Stream channel count	Packet time	Requirement level
T1 “T-48k-1ms”	48 kHz	At least one of: L16, L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Shall”
T2 “T-96k-1ms”	96 kHz	L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Should”
T3 “T-48k-1ms”	44,1 kHz	L16	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Should”
T4 “T-48k-non1ms”	48 kHz	At least one of: L16, L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
T5 “T-96k-non1ms”	96 kHz	L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
T6 “T-44k-non1ms”	44,1 kHz	L16	At least one of: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
T7 “T-other”	At least one of: 48 kHz, 96 kHz, 44,1 kHz	At least one of: L16, L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8, “other”	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “1 millisecond” “4 milliseconds” “other”	“May”

Table G.3 - Stream mode interoperability requirements for receivers

Requirement identifier / name	Sampling rate	Audio sample encoding	Stream channel count	Packet time	Requirement level
R1 “R-48k-1ms”	48 kHz	All: L16, L24	All: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Shall”
R2 “R-96k-1ms”	96 kHz	L24	All: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Should”
R3 “R-44k-1ms”	44,1 kHz	L16	All: 1, 2, 3, 4, 5, 6, 7, 8	“1 millisecond”	“Should”
R4 “R-48k-non1ms”	48 kHz	All: L16, L24	All: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
R5 “R-96k-non1ms”	96 kHz	L24	All: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
R6 “R-44k-non1ms”	44,1 kHz	L16	All: 1, 2, 3, 4, 5, 6, 7, 8	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “4 milliseconds”	“Should”
R7 “R-other”	At least one of: 48 kHz, 96 kHz, 44,1 kHz	At least one of: L16, L24	At least one of: 1, 2, 3, 4, 5, 6, 7, 8, “other”	At least one of: “125 microseconds” “250 microseconds” “333 microseconds” “1 millisecond” “4 milliseconds” “other”	“May”

Annex H Bibliography

AES5, *AES recommended practice for professional digital audio - Preferred sampling frequencies for applications employing pulse-code modulation*. Audio Engineering Society, New York, NY., US

AES-R16-2021, *AES Standards Report - PTP parameters for AES67 and SMPTE ST 2059-2 interoperability*, Audio Engineering Society, New York, NY., US

EBU Tech 3326, *Audio contribution over IP - Requirements for Interoperability*, European Broadcasting Union, Geneva, Switzerland

IEEE 802.1BA, *Audio Video Bridging (AVB) Systems*, Institute of Electrical and Electronics Engineers (IEEE), US

IEEE 802.1Q-2011, *Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks*, Institute of Electrical and Electronics Engineers (IEEE), US

IEEE 802.1AS, *Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks*, Institute of Electrical and Electronics Engineers (IEEE), US

IEEE 1588g-2022, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. Amendment 2: Master-Slave Optional Alternative Terminology*, December 2022, Institute of Electrical and Electronics Engineers (IEEE), US

RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*, Internet Engineering Task Force

RFC 899, *Host Extensions for IP Multicasting*, Internet Engineering Task Force

RFC 2365, *Administratively Scoped IP Multicast*, Internet Engineering Task Force

RFC 2597, *Assured Forwarding PHB Group*, Internet Engineering Task Force

RFC 2974, *Session Announcement Protocol*, Internet Engineering Task Force

RFC 3170, *IP Multicast Applications: Challenges and Solutions*, Internet Engineering Task Force

RFC 3629, *UTF-8, a transformation format of ISO 10646*, Internet Engineering Task Force

RFC 3986, *Uniform Resource Identifier (URI): Generic Syntax*, Internet Engineering Task Force

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*, Internet Engineering Task Force

RFC 4594, *Configuration Guidelines for DiffServ Service Classes*, Internet Engineering Task Force

RFC 5760, *RTP Control Protocol (RTCP) Extensions for Single-Source Multicast Sessions with Unicast Feedback*, Internet Engineering Task Force

RFC 5771, *IANA Guidelines for IPv4 Multicast Address Assignments*, Internet Engineering Task Force

RFC 6128, *RTP Control Protocol (RTCP) Port for Source-Specific Multicast (SSM) Sessions*, Internet Engineering Task Force

RFC 6335, *Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry*, Internet Engineering Task Force

RFC 6762, *Multicast DNS*, Internet Engineering Task Force

RFC 6763, *DNS-Based Service Discovery*, Internet Engineering Task Force

RFC 7272, *Inter-Destination Media Synchronization (IDMS) Using the RTP Control Protocol (RTCP)*, Internet Engineering Task Force

RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*, Internet Engineering Task Force

draft-lee-sip-dns-sd-uri, *SIP URI Service Discovery using DNS-SD*, Internet Engineering Task Force. http://datatracker.ietf.org/doc/draft-lee-sip-dns-sd-uri/

draft-jones-sip-options-ping, *Using OPTIONS to Query for Operational Status in the Session Initiation Protocol (SIP)*, Internet Engineering Task Force. http://datatracker.ietf.org/doc/draft-jones-sip-options-ping/

*Geoffrey M. Garner, Michel Ouellette and Michael Johas Teener (2012-09-27). **"Using an IEEE 802.1AS Network as a Distributed IEEE 1588 Boundary, Ordinary, or Transparent Clock"**. 2010 International IEEE Symposium on Precision Clock Synchronization for Measurement Control and Communication (ISPCS) (IEEE).*

IS-04, NMOS Discovery & Registration, Advanced Media Workflow Association (AMWA), http://www.amwa.tv/projects/IS-04.shtml

ST 2110-10:2022, *Professional Media over Managed IP Networks: System Timing and Definitions*, Society of Motion Picture and Television Engineers

ST 2110-30:2017, *Professional Media over Managed IP Networks: PCM Digital Audio*, Society of Motion Picture and Television Engineers