

for Television — Declarative Data Essence — Internet Protocol Multicast Encapsulation



Page 1 of 10 pages

Table of contents

- 1 Scope and application
- 2 Normative references
- 3 Announcement protocol
- 4 Trigger protocol
- 5 Resource transfer: UHTTP
- Annex A Using enhanced television
- Annex B Example broadcast
- Annex C Glossary

1 Scope and application**1.1 Scope**

This standard defines the encapsulation of declarative data essence using internet protocol (IP) multicast. This is done in a transport-independent manner and relies solely on standard IP multicast techniques.

1.2 Application

This standard defines the transmission of declarative content across terrestrial (over the air), cable, and satellite systems as well as over the Internet. In addition, it will also bridge between networks — for example, data on an analog terrestrial broadcast must easily bridge to a digital cable system. This design goal was achieved through the definition of a transport-independent content format and the use of IP. Since IP bindings already exist for each of these video systems, the advantages of this work may be useful.

IP multicast is the mechanism for broadcast data delivery. Content creators should assume IP addresses may be changed downstream and, therefore, should not use them in their content. The transport operator is responsible only for making sure that an IP address is valid on the physical network where they broadcast it (not for any rebroadcasting). When possible, content creators should use valid IP multicast addresses to minimize the chance of collisions. Some systems may have two-way Internet connections. Capabilities in those systems are outside the scope of this standard and are described by the appropriate Internet standards.

Transport operators should use the standard IP transmission system for the appropriate medium (IETF, ATSC, DVB, etc.). It is assumed that when the user tunes to a television channel, the receiver automatically locates and delivers IP datagrams associated with the television broadcast. The mechanism for tuning video and connecting to the appropriate data stream is implementation and delivery standard specific and is not specified in this framework.

2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

SMPTE 363M-2002, Television — Declarative Data Essence — Content Level 1

SMPTE 364M-2001, Television — Declarative Data Essence — Unidirectional Hypertext Transport Protocol

IETF RFC 768, User Datagram Protocol (UDP)

IETF RFC 791, Internet Protocol — DARPA Internet Program Protocol Specification

IETF RFC 1112, Host Extensions for IP Multicasting

IETF RFC 2327, SDP: Session Description Protocol

IETF RFC 2974, Session Announcement Protocol (SAP)

ISO/IEC 11578:1996, Information Technology — Open Systems Interconnection — Remote Procedure Call (RPC), Annex A, Universal Unique Identifier (UUID)

3 Announcement protocol

Announcements are used to announce currently available programming to the receiver. The IP multicast addresses and ports for resource transfer and for triggers are announced using SDP announcements (IETF RFC 2327). The SDP header is preceded by an 8-byte SAP header (IETF RFC 2974). Announcements are sent on a well-known address (224.0.1.113) and port (2670). This address and port have been registered with the IANA.

v=0 SDP version, required to be 0.

o=username sid version IN IP4 ipaddress: Owner and session identifier, defined as usual in SDP specification. Username is "-", network type is IN, address type is IP4. *Sid* identifies an announcement for a particular broadcast (it can be a permanent announcement for all programming on a broadcast channel or for a particular show). *Version* indicates the version of the message. These values allow receivers to match a message to a previous message and know whether it has changed. Session ID and Version should be NTP values as recommended in SDP.

s=name Session name, required as in SDP specification.

i=, u= Optional, as in SDP specification.

e=, p= E-mail address or telephone number (at least one required in SDP specification).

b=CT:number Optional in SDP specification, but required here. Bandwidth in kilobytes per second as in the SDP specification. Bandwidth of the broadcast data can be used by receivers to choose among multiple versions of enhancement data according to the bandwidth the receiver can handle.

t=start stop As in SDP specification, gives start and stop time in NTP format. With programs stored on tape, at times it will not be possible to insert new announcements, so start times on tape could be incorrect. In this case, the start time should be set to the original broadcast time and the stop time set to 0. This is the standard for an

unbounded session. Assumptions are then made about the stop time (see IETF RFC 2327). A new announcement with the same sid and different version for the same broadcast station replaces the previous one. It is preferred that a tool read the tape and generate announcements with correct start and stop times, but it is not required. Content creators can choose to use only a station ID and not provide information about individual programs.

a=UUID:UUID Optional. The UUID should uniquely identify the enhancement (for example, a different UUID for each program), and can be accessed using the trigger receiver object. In analog television and many types of digital television broadcast data are tied tightly to A/V. Each virtual channel has its own private network associated with it. In other systems, enhancements for many virtual channels can be carried on the same network. These systems can use the UUID to link a television broadcast with a particular enhancement. How that association is indicated is beyond the scope of this standard. One technique would be to place the UUID in electronic program guide information. Use ASCII hex to encode UUIDs.

a=type:tve Required. Indicates to the receiver that the announcement refers to an enhancement related to this specification.

a=lang, a=sdplang Optional, as in SDP specification.

a=tve-type:<types> Optional. *tve-type*: specifies an extensible list of types that describe the nature of the enhancement. It is a session-level attribute and is not dependent on *a=tve-type:primary*. Optional. *tve-type:primary* specifies that this will be the primary enhancement stream associated with the currently playing video program whenever this enhancement's trigger stream is active. If *tve-type:primary* is not specified, the TVE stream is never the primary enhancement stream associated with video. This, like all *tve-type*: attributes, is a session level attribute. This attribute can be used by receivers to implement automatic loading of primary video enhancement streams. The actual display of and switching between enhancement streams are handled by the trigger streams.

a=tve-size:Kbytes Required. *tve-size*: provides an estimate of the high watermark of cache storage in kilobytes that will be required during the playing of the enhancement. This is necessary so that receivers can adequately judge whether or not they can successfully play an enhancement from beginning to end.

a=tve-level:x Content level identifier, where x is 1.0 for this version of the framework (optional, default is 1.0).

a=tve-ends:seconds Optional, specifies an end time relative to the reception time of the SDP announcement. *Seconds* is the number of seconds in the future that this announcement is valid. *Seconds* may change (count down) as an announced session progresses. This attribute, when present, overrides the default assumptions for end times in unbounded announcements.

m=data portvalue/2 tve-file/tve-trigger

c=IN IP4 ipaddress/ttl As in SDP specification. Compact form specifying two ports on the same address.

When there are multiple alternative enhancement streams for the same video program, they must all be announced at the media level of the same SDP announcement. All enhancement streams announced in the same SDP announcement are considered to be mutually exclusive variants of the primary enhancement stream. The receiver can choose between them based on media level attributes. For example, the *a=lang* field can be used at the media level to choose between language variants of the primary enhancement.

Each media section for the *tve-file* media type begins the next enhancement definition.

A longer form is available if the content creator or transport operator wants to use different IP addresses and ports for the data stream and trigger stream: