

ETSI TS 103 285 V1.1.1 (2015-05)



Digital Video Broadcasting (DVB); MPEG-DASH Profile for Transport of ISO BMFF Based DVB Services over IP Based Networks

EBU
OPERATING EUROVISION

DVB[®]
Digital Video
Broadcasting

Reference

DTS/JTC-DVB-342

Keywords

broadcasting, digital, DVB, IP, satellite, TV, video

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

The present document can be downloaded from:

<http://www.etsi.org/standards-search>

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the only prevailing document is the print of the Portable Document Format (PDF) version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status.

Information on the current status of this and other ETSI documents is available at

<http://portal.etsi.org/tb/status/status.asp>

If you find errors in the present document, please send your comment to one of the following services:

<https://portal.etsi.org/People/CommitteeSupportStaff.aspx>

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.

The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute yyyy.

All rights reserved.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.
3GPP™ and **LTE™** are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.
GSM® and the GSM logo are Trade Marks registered and owned by the GSM Association.

Contents

Intellectual Property Rights	7
Foreword.....	7
Modal verbs terminology.....	7
1 Scope	8
2 References	8
2.1 Normative references	8
2.2 Informative references.....	10
3 Definitions and abbreviations.....	10
3.1 Definitions.....	10
3.2 Abbreviations	12
4 DASH Constraints, HTTP Support, and Synchronization	13
4.1 DVB Profile of MPEG-DASH	13
4.2 Media Presentation Description Constraints	13
4.2.1 General.....	13
4.2.2 Constraints on Period elements (common Period constraints).....	13
4.2.3 Constraints on Period elements conforming to Live profile	13
4.2.4 Constraints on AdaptationSet elements (for Live Periods).....	13
4.2.5 Constraints on Representation elements (for Live Periods).....	14
4.2.6 Constraints on Period elements conforming to On Demand profile	14
4.2.7 Constraints on AdaptationSet element (On Demand).....	14
4.2.8 Constraints on Representation element (On Demand).....	15
4.3 Segment Format Constraints	15
4.4 Presence of Attributes and Elements	15
4.5 Dimension Constraints	16
4.6 Server Requirements	16
4.7 Availability Time Synchronization between Player and Server.....	16
4.7.1 Background.....	16
4.7.2 Service Provider Requirements.....	17
4.7.3 Player Requirements	17
5 DASH Specific Aspects for Video.....	18
5.1 DASH Specific Aspects for H.264/AVC Video.....	18
5.1.1 Profiles Supported	18
5.1.2 H.264/AVC Specifics	18
5.1.3 Signalling of the Codec Profile.....	18
5.1.4 Signalling within the AVCSampleEntry in an initialization segment.....	19
5.2 DASH Specific Aspects for HEVC Video	19
5.2.1 HEVC Specifics.....	19
5.2.2 Signalling of the Codec Profile.....	19
5.2.3 Profiles Supported	21
6 DASH Specific Aspects for Audio.....	21
6.1 Common DASH Specific Aspects for Audio	21
6.1.1 Adaptation Sets and Representations	21
6.1.2 Using the Role Scheme to Distinguish Between Different Adaptation Sets.....	22
6.2 DASH Specific Aspects for HE-AACv2 Audio.....	24
6.3 DASH Specific Aspects of Dolby Audio Technologies	25
6.4 DASH Specific Aspects of DTS Audio Technologies	25
6.5 DASH Specific Aspects of MPEG Surround	25
6.6 Service Continuity Considerations	26
6.6.1 Introduction.....	26
6.6.2 Default Operation for Seamless Service Continuity	26
6.6.3 Fallback Operation.....	26
7 DASH Specific Aspects for Subtitles.....	27
7.1 Carriage of Subtitles.....	27

7.1.1	General.....	27
7.1.2	Distinguishing Between Different Uses of Subtitles	27
7.2	Downloadable fonts.....	28
7.2.0	Introduction.....	28
7.2.1	Signalling within the MPD	28
7.2.1.1	DVB font download scheme	28
7.2.1.2	Scheme Identification in descriptors	29
7.2.1.3	Additional attributes on the EssentialProperty and SupplementalProperty descriptors	29
7.2.2	Font format	29
7.2.3	Errors	29
7.2.4	Player support	30
7.3	Example (informative)	30
8	Content Protection.....	30
8.1	Introduction	30
8.2	Background (informative)	30
8.3	Encryption of Different Representations	31
8.4	MPEG DASH MPD Content Protection descriptors	31
8.5	MPEG Common Encryption	31
8.5.0	Protection System Specific Header	31
8.5.1	Key Rotation	31
8.5.2	Random access requirement	32
8.6	Encryption Mode.....	32
8.6.0	Encryption Usage.....	32
8.6.1	Constraints on the SampleAuxiliaryInformationOffsetsBox	32
8.7	Mixing Non-Encrypted and Encrypted Content / Information.....	33
9	Carriage of Generic Streams and Events in DASH.....	33
9.1	Events	33
9.1.1	Background (Informative)	33
9.1.2	Content programme metadata	34
9.1.2.1	Definition	34
9.1.2.2	Event message syntax	35
9.1.2.3	Parental rating	36
9.1.2.4	Example	36
9.1.3	Application messages	36
9.1.4	Events for the Player.....	37
9.1.5	Long duration inband events.....	37
9.1.6	Monitoring of InbandEventStreams.....	38
10	Player Behaviour	38
10.0	Introduction	38
10.1	DVB Profile Support.....	38
10.2	ISO BMFF Player Constraints.....	38
10.3	Luminance Resolutions and Frame Rates	39
10.4	Audio/Video Switching Between Different Representations in the same Adaptation Set	40
10.5	Playback Across Periods	41
10.5.1	Background.....	41
10.5.2	Content Offering with Multiple Periods	41
10.5.2.1	General	41
10.5.2.2	Associated Adaptation Sets across Periods	41
10.5.2.3	Period Continuity	42
10.5.3	Player Requirements and Recommendations	43
10.6	Trick Mode Implementation (informative).....	43
10.6.0	Introduction to Trick Mode.....	43
10.6.1	Player Enabled Trick Modes.....	43
10.6.1.0	Introduction to Player Enabled Trick Modes	43
10.6.1.1	Fast Forward H.264/AVC Long GOP fast decoding and frame dropping	44
10.6.1.2	Fast Forward H.264/AVC picture discarding before decoding	45
10.6.1.3	Fast forward I picture access	45
10.6.1.4	Fast reverse I picture access	45
10.6.1.5	Segment random access	45
10.6.1.6	Partial Segment fast forward or reverse	45

10.6.2	Content enabled Trick Modes:	45
10.6.2.0	Introduction to Content enabled Trick Modes	45
10.6.2.1	Trick Play Representations	46
10.6.2.2	Sub Representations	46
10.7	Player Buffering	46
10.8	Player Resiliency	46
10.8.1	Introduction (Informative)	46
10.8.2	Handling of BaseURLs by Players	47
10.8.2.1	Use of BaseURLs and new attributes	47
10.8.2.2	Selection of BaseURL	48
10.8.2.3	Changing BaseURL	48
10.8.2.4	BaseURL Example (Informative)	49
10.8.3	DNS - HTTP Player	51
10.8.4	Types of error condition and recovery options (Informative)	51
10.8.4.0	Introduction	51
10.8.4.1	Heavy server load	51
10.8.4.2	Configuration errors	51
10.8.4.3	Authentication errors	51
10.8.4.4	Missing segments	51
10.8.4.5	Miscellaneous request errors	51
10.8.4.6	Transient connection errors or congestion	51
10.8.5	Specific errors and their categories	51
10.8.6	Requirements on Players	52
10.9	Live Services	53
10.9.1	Overview (informative)	53
10.9.2	Determining that an MPD represents a live service	53
10.9.3	Presentations using segment templates and no segment timeline	54
10.9.4	Considerations for 24/7 operation of a live stream	55
10.9.4.0	Introduction	55
10.9.4.1	AvailabilityStartTime	55
10.9.4.2	Large segment numbers	55
10.9.4.3	Large track fragment decode times	55
10.10	MPD Processing	55
10.11	HTTP Support	55
10.12	Reporting errors to content providers	56
10.12.1	Introduction (Informative)	56
10.12.2	Player support	56
10.12.3	DVB metrics reporting mechanism	56
10.12.3.1	Scheme Identification	56
10.12.3.2	Scheme URI	57
10.12.3.3	Additional attributes on the Reporting descriptor	57
10.12.3.4	Use of the @probability attribute	57
10.12.3.5	Reporting metrics	57
10.12.4	Definition of the DVBErrors metric	58
10.12.5	Example (informative)	59
10.13	Reporting media duration and current play position to higher level applications (informative)	60
11	Content Provider Guidelines (Informative)	61
11.1	Use of Profiles	61
11.2	Video Resolutions and Frame Rates in Adaptation Sets	61
11.2.0	Introduction	61
11.2.1	Video Resolutions	61
11.2.2	Frame Rates	62
11.3	Considerations for Audio in Adaptation Sets	62
11.3.0	Recommendation for Audio as a Percentage of Total Stream Bandwidth	62
11.3.1	Considerations for Advanced Audio Coding	62
11.3.2	Considerations for Dolby Enhanced AC-3 and AC-4	63
11.3.3	Considerations for DTS Audio Formats	64
11.4	Considerations for Interlaced and Progressive in Adaptation Sets	65
11.5	Mixing HD and SD in Adaptation Sets	65
11.6	Mixing H.264/AVC and HEVC in Adaptation Sets	65
11.7	Subtitles	65

11.8	Audio/Video Switching Across Period Boundaries	66
11.9	Signalling multiple serving locations using BaseURLs	66
11.9.1	Introduction.....	66
11.9.2	Prioritizing BaseURLs.....	66
11.9.3	Load balancing.....	66
11.9.4	Combining priority and load balancing	67
11.9.5	Relative and Absolute BaseURLs.....	69
11.10	Using Events	69
11.11	Asset Identifiers.....	70
11.12	Considerations for Running 24/7 services.....	70
11.12.1	A/V synchronization	70
11.12.2	Dual (or more) redundant encoding and hosting infrastructure	70
11.12.3	Encryption.....	70
11.12.4	Further information.....	71
11.13	Cache Control.....	71
	History	72

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<http://ipr.etsi.org>).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by Joint Technical Committee (JTC) Broadcast of the European Broadcasting Union (EBU), Comité Européen de Normalisation ELECTrotechnique (CENELEC) and the European Telecommunications Standards Institute (ETSI).

NOTE: The EBU/ETSI JTC Broadcast was established in 1990 to co-ordinate the drafting of standards in the specific field of broadcasting and related fields. Since 1995 the JTC Broadcast became a tripartite body by including in the Memorandum of Understanding also CENELEC, which is responsible for the standardization of radio and television receivers. The EBU is a professional association of broadcasting organizations whose work includes the co-ordination of its members' activities in the technical, legal, programme-making and programme-exchange domains. The EBU has active members in about 60 countries in the European broadcasting area; its headquarters is in Geneva.

European Broadcasting Union
CH-1218 GRAND SACONNEX (Geneva)
Switzerland
Tel: +41 22 717 21 11
Fax: +41 22 717 24 81

Founded in September 1993, the DVB Project is a market-led consortium of public and private sector organizations in the television industry. Its aim is to establish the framework for the introduction of MPEG-2 based digital television services. Now comprising over 200 organizations from more than 25 countries around the world, DVB fosters market-led systems, which meet the real needs, and economic circumstances, of the consumer electronics and the broadcast industry.

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document defines the delivery of TV content via HTTP adaptive streaming. This includes the following:

- A profile of the features defined in MPEG DASH specification [1] (referred to by MPEG as an "interoperability point") largely based on the "ISO/BMFF live" profile defined by MPEG.
- Constraints on the sizes or complexity of various parameters defined in the MPEG DASH specification [1].
- A selection of the video and audio codecs from the DVB toolbox that are technically appropriate with MPEG DASH constraints and/or requirements for the use of these, without mandating any particular codec.
- Using MPEG Common Encryption for content delivered according to the present document.
- Use of TTML subtitles with MPEG DASH.
- Requirements on Player behaviour needed to give inter-operable presentation of services.
- Guidelines for content providers on how to use MPEG DASH.

Amongst others, the following subjects are *not* covered in the present document:

- Mandatory codecs.
- Specific Content Protection schemes.

The normative XML schemas referenced by the present document are attached as separate files contained in archive ts_103285v010101p0.zip which accompanies the present document. The XML schemas included in the present document are informative.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at <http://docbox.etsi.org/Reference>.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

- | | |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [1] | ISO/IEC 23009-1:2014: "Information technology - Dynamic adaptive streaming over HTTP (DASH) - Part 1: Media presentation description and segment formats". |
| [2] | ISO/IEC TR 23009-3: "Information technology - Dynamic adaptive streaming over HTTP (DASH) - Part 3: Implementation guidelines". |
| [3] | ETSI TS 101 154: "Digital Video Broadcasting (DVB); Specification for the use of Video and Audio Coding in Broadcasting Applications based on the MPEG-2 Transport Stream". |
| [4] | ISO/IEC 14496-15:2014 (3rd Edition): "Information technology -- Coding of audio-visual objects -- Part 15: Carriage of network abstraction layer (NAL) unit structured video in ISO base media file format". |
| [5] | IETF RFC 6381: "The 'Codecs' and 'Profiles' Parameters for "Bucket" Media Types". |
| [6] | ISO/IEC 14496-12:2012 "Information technology -- Coding of audio-visual objects -- Part 12: ISO base media file format". |

- [7] ISO/IEC 23001-7, Second Edition: "Information technology -- MPEG systems technologies -- Part 7: Common encryption in ISO base media file format files".
 - [8] ISO/IEC 14496-3:2009/Amd 4:2013: "New levels for AAC profiles".
 - [9] ISO/IEC 14496-14: "Information technology - Coding of audio-visual objects - Part 14: MP4 file format".
 - [10] ETSI TS 102 366: "Digital Audio Compression (AC-3, Enhanced AC-3) Standard".
 - [11] ETSI TS 102 114: "DTS Coherent Acoustics; Core and Extensions with Additional Profiles".
 - [12] ISO/IEC 23003-1:2007: "Information technology -- MPEG audio technologies -- Part 1: MPEG Surround".
 - [13] EBU Tech 3380: "EBU-TT-D Subtitling Distribution Format".
 - [14] EBU Tech 3381: "Carriage of EBU-TT-D in ISO/BMFF".
 - [15] ETSI EN 300 468: "Digital Video Broadcasting (DVB); Specification for Service Information (SI) in DVB systems".
 - [16] ETSI ES 202 184: "MHEG-5 Broadcast Profile".
 - [17] IETF RFC 2616: "Hypertext Transfer Protocol -- HTTP/1.1".
 - [18] IETF RFC 6265: "HTTP State Management Mechanism".
 - [19] ISO/IEC 14496-30: "Information technology -- Coding of audio-visual objects -- Part 30: Timed text and other visual overlays in ISO base media file format".
 - [20] ISO/IEC 14496-22:2009: "Information technology -- Coding of audio-visual objects -- Part 22: Open Font Format".
 - [21] Recommendation ITU-R BT.709: "Parameter values for the HDTV standards for production and international programme exchange".
 - [22] W3C Recommendation (December 2012): "WOFF File Format 1.0".
- NOTE: Available at <http://www.w3.org/TR/WOFF>.
- [23] ETSI TS 103 190: "Digital Audio Compression (AC-4) Standard".
 - [24] IETF RFC 2782: "A DNS RR for specifying the location of services (DNS SRV)".
 - [25] IETF RFC 7164: "RTP and Lead Seconds".
 - [26] W3C Recommendation (November 2008) "Extensible Markup Language (XML) 1.0".
- NOTE: Available at: <http://www.w3.org/TR/REC-xml/>.
- [27] W3C Recommendation (September 2012): "Media Fragments URI 1.0".
- NOTE: Available at: <http://www.w3.org/TR/media-frags/>.
- [28] IETF RFC 5234: "Augmented BNF for Syntax Specifications: ABNF".
 - [29] ISO/IEC 14496-15:2014: "Information technology -- Coding of audio-visual objects -- Part 15: Carriage of network abstraction layer (NAL) unit structured video in ISO base media file format".
 - [30] ISO 8601:2004: "Data elements and interchange formats -- Information interchange -- Representation of dates and times".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] DASH-IF: "Guidelines for Implementation: DASH-IF Interoperability Points V3.0".

NOTE: Available at <http://dashif.org/guidelines/>

[i.2] DTS document 9302K62400 "Implementation of DTS Audio in Dynamic Adaptive Streaming over HTTP (DASH)".

NOTE: Available at <http://www.dts.com/resources>.

[i.3] Recommendation ITU-T T.35: "Procedure for the allocation of ITU-T defined codes for non-standard facilities".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

access unit: unit of a media stream with an assigned Media Presentation time

accessibility: degree to which a media content or certain media content components are available to as many people as possible

adaptation set: set of interchangeable encoded versions of one or several media content components

adjusted segment availability start time: time instant in wall-clock time at which a Segment becomes an available Segment

asset: content including media and metadata together with the rights to use the content by the content provider

available segment: segment that is accessible at its assigned HTTP-URL with optionally an assigned byte range that when requested with an HTTP GET results in a reply with the Segment and a 2xx status code

bitstream switching segment: segment that if present contains essential data to switch to the Representation it is assigned to

complementary representation: representation which complements at least one dependent Representation

content protection: protection of content such that it can only be presented by authorized Devices

content provider: entity that owns or is licensed to sell content or content assets

continuous media: media with an inherent notion of time, for example, speech, audio, video, timed text or timed metadata

DASH metric: metric identified by a key and defined in Hi ISO/IEC 23009 [1]

dependent Representation: representation for which Segments from its complementary Representations are necessary for presentation and/or decoding of the contained media content components

earliest presentation time: smallest presentation time of any access unit of a Media Segment or Subsegment for a media stream

event: aperiodic sparse media-time related auxiliary information to the DASH client or to an application

event stream: sequence of related events

group: collection of Adaptation Sets that are not expected to be presented simultaneously

HTTP-URL: URL with a fixed scheme of "http" or "https"

index segment: Segment that primarily contains indexing information for Media Segments

initialization segment: segment containing metadata that is necessary to present the media streams encapsulated in Media Segments

media content: one media content period or a contiguous sequence of media content periods

media content component: one continuous component of the media content with an assigned media component type that can be encoded individually into a media stream

media content component type: a single type of media content such as audio, video, or text

media content period: set of media content components that have a common timeline as well as relationships on how they can be presented

media presentation: collection of data that establishes a bounded or unbounded presentation of media content

Media Presentation Description (MPD): formalized description for a Media Presentation for the purpose of providing a streaming service

media presentation timeline: concatenation of the timeline of all Periods which itself is common to all Representations in the Period

media segment: segment that complies with media format in use and enables playback when combined with zero or more preceding segments, and an Initialization Segment (if any)

media stream: encoded version of a media content component

media Subsegment: subsegment that only contains media data but no Segment Index

message: part of an event containing information that is exclusively handled by the event handler

MPD start time: approximate presentation start time of a Media Segment signalled in MPD

MPD duration: approximate presentation duration of a Media Segment signalled in MPD

period: interval of the Media Presentation, where a contiguous sequence of all Periods constitutes the Media Presentation

player: device responsible for receiving and rendering media services

presentation time: time associated to an access unit that maps it to the Media Presentation timeline

remote element entity: entity that contains one or more elements and is referenced in the MPD with an HTTP-URL contained in an @xlink:href attribute

representation: collection and encapsulation of one or more media streams in a delivery format and associated with descriptive metadata

segment: unit of data associated with an HTTP-URL and optionally a byte range that are specified by an MPD

segment availability end time: time instant in wall-clock time at which a Segment ceases to be an available Segment

segment index: compact index of the time range to byte range mapping within a Media Segment separately from the MPD

segment availability start time: latest time instant in wall-clock time at which a Segment becomes an available Segment

Service Provider (SP): entity providing a service to the end-user

Stream Access Point (SAP): position in a Representation enabling playback of a media stream to be started using only the information contained in Representation data starting from that position onwards (preceded by initializing data in the Initialization Segment, if any)

sub-representation: part of a Representation described in the MPD that is present in the entire Period

subsegment: unit within Media Segments that is indexed by a Segment Index

trick mode: playback in any mode other than forward playback at the recorded speed of the audio/video content ("1x").

NOTE: Examples include fast forward, slow motion, reverse, and random access.

valid Segment URL: HTTP-URL that is promised to reference a Segment during its Segment availability period

wall-clock time: time as stated by UTC

3.2 Abbreviations

For the purposes of the present document, the following abbreviation apply.

AAC-LC	Advanced Audio Encoding - Low Complexity
ABNF	Augmented Backus-Naur Form
AES	Advanced Encryption Standard
AOT	Audio Object Type
AVC	Advanced Video Coding
BMFF	Base Media File Format
CBR	Constant Bit Rate
CDN	Content Delivery Network
CENC	Common ENCryption standard
CFE	Common File Format
CVS	Coded Video Sequence
DASH	Dynamic Adaptive Streaming over HTTP
DNS	Domain Name System
DRM	Digital Rights Management
DVB	Digital Video Broadcasting
E-AC3	Enhanced AC3
EBU	European Broadcasting Union
EBU TT-D	European Broadcasting Union Time Text - Distribution
GeoIP	Geolocation by IP address
GOP	Group of Pictures
HbbTV	Hybrid Broadcast Broadband TV
HDTV	High Definition TeleVision
HE-AAC	High Efficiency Advanced Audio Coding
HEVC	High Efficiency Video Coding
HTTP	HyperText Transfer Protocol
IDR	Instantaneous Decoding Refresh
IETF	Internet Engineering Task Force
IRD	Integrated Receiver-Decoder
ISO BMFF	ISO Base Media File Format
KID	Key Identification
MHEG	Multimedia and Hypermedia information coding Expert Group
MHP	Multimedia Home Platform
MIME	Multipurpose Internet Mail Extensions
MPD	Media Presentation Description
MPEG	Moving Pictures Expert Group
NAL	Network Abstraction Layer
NTP	Network Time Protocol
PPS	Picture Parameter Set
RFC	Request For Comments
SD	Standard Definition
SEI	Supplemental Enhancement Information
SMPTE	Society of Motion Picture and Television Engineers
SNTP	Simple Network Time Protocol

SPS	Sequence Parameter Set
TCP	Transmission Control Protocol
TTML	Timed Text Markup Language
TV	TeleVision
TVA	TV Anytime
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
URN	Uniform Resource Name
VoD	Video on Demand
VUI	Video Usability Information
XML	eXtensible Markup Language

4 DASH Constraints, HTTP Support, and Synchronization

4.1 DVB Profile of MPEG-DASH

The DVB Profile of MPEG-DASH, known as an "interoperability point" by MPEG, is based on the merging of the ISO/IEC 23009-1 [1], ISO Base media file format live profile and ISO Base media file format On Demand profile. In addition it includes "xlink" but only in combination with the actuate property set to "onLoad".

The DVB profile excludes multiplexed representations, segment list addressing and "xlink" with the actuate property set to "onRequest".

The URN for the profile (MPEG Interoperability Point) shall be "urn:dvb:dash:profile:dvb-dash:2014".

4.2 Media Presentation Description Constraints

4.2.1 General

The MPD shall conform to the constraints in clause 4.2 in addition to the rules for the MPD as defined in ISO/IEC 23001-9 [1], clause 7.3.

NOTE: MPDs, or their constituent parts, which do not conform to the constraints in clause 4.2 may be ignored by the client.

4.2.2 Constraints on Period elements (common Period constraints)

- The Subset element may be ignored.
- The Period.SegmentList element shall not be present.
- If a Period element contains multiple Adaptation Sets with @contentType="video" then at least one Adaptation Set shall contain a Role element with @schemeIdUri="urn:mpeg:dash:role:2011" and @value="main".
- Remote elements with @xlink:actuate set to "onRequest" may be ignored.
- Each Period element shall conform to either clause 4.2.3 or 4.2.6.

NOTE: This means that each Period can be either based on ISO Base media file format live profile or ISO Base media file format On Demand profile, but not any combination of both.

4.2.3 Constraints on Period elements conforming to Live profile

- AdaptationSet elements that do not conform to clause 4.2.4 may be ignored.

4.2.4 Constraints on AdaptationSet elements (for Live Periods)

- The @group attribute may be ignored.

- Any **AdaptationSet** with a **ContentComponent** element may be ignored.

NOTE 1: This is so that Adaptation Sets with multiplexed media streams may be ignored.

- AdaptationSet element can be ignored unless AdaptationSet.SegmentTemplate is present and/or the Representation.SegmentTemplate element is present for each Representation within this Adaptation Set.

NOTE 2: SegmentTimeline is supported according to ISO/IEC 23009-1 [1].

- AdaptationSet element that contains more than one Representation can be ignored unless all of the following hold:
 - AdaptationSet@segmentAlignment is present and has value of 'true' or '1'.
 - AdaptationSet@startWithSAP is present and has value of 1 or 2.
 - The **MPD**@maxSegmentDuration attribute is present.
- Representation elements that do not conform to 4.2.5 may be ignored.

NOTE 3: AdaptationSets may legitimately include one or more pairs of Representations which, under the requirements of clause 10.4, conformant clients are not required to switch between. One example use case for such an AdaptationSet can be found in clause 11.4.

4.2.5 Constraints on Representation elements (for Live Periods)

- Representations with a value of the @mimeType attribute starting with a string other than "xxx/mp4" where xxx is one of "video", "audio", "application" or "text" may be ignored. Additional profile or codec specific parameters may be added to the end of the value of the @mimeType attribute.
- Representations not inferred to have an @profiles attribute equal to "urn:dvb:dash:profile:dvb-dash:isoff-ext-live:2014" may be ignored.

4.2.6 Constraints on Period elements conforming to On Demand profile

- The Period.SegmentTemplate element shall not be present.
- AdaptationSet elements that do not conform to clause 4.2.7 may be ignored.

4.2.7 Constraints on AdaptationSet element (On Demand)

- The @group attribute may be ignored.
- Any **AdaptationSet** with a **ContentComponent** element may be ignored.

NOTE 1: This is so that Adaptation Sets with multiplexed media streams may be ignored.

- AdaptationSet element can be ignored unless for each Representation within this Adaptation Set Representation.BaseURL is present.
- If either the AdaptationSet.SegmentList or the AdaptationSet.SegmentTemplate element is present in an AdaptationSet element then this AdaptationSet element may be ignored.
- If an AdaptationSet element contains more than one Representation element, and this AdaptationSet element can be ignored unless AdaptationSet@subsegmentAlignment is present and has value of 'true'; and AdaptationSet@subsegmentStartsWithSAP is present and has value of 1 or 2.
- Representation elements that do not conform to clause 4.2.8 may be ignored.

NOTE 2: AdaptationSets may legitimately include one or more pairs of Representations which, under the requirements of clause 10.4, conformant clients are not required to switch between. One example use case for such an AdaptationSet can be found in clause 11.4.

4.2.8 Constraints on Representation element (On Demand)

- Representations with a value of the @mimeType attribute starting with a string other than "xxx/mp4" where xxx is one of "video", "audio", "application" or "text" may be ignored. Additional profile or codec specific parameters may be added to the end of the value of the @mimeType attribute.
- Representations not inferred to have an @profiles attribute equal to "urn:dvb:dash:profile:dvb-dash:isoff-ext-on-demand:2014" may be ignored.
- If either the Representation.SegmentList or the Representation.SegmentTemplate element is present in a Representation element then this Representation element may be ignored.
- If the Representation element does not contain a BaseURL element then this Representation element may be ignored.
- If Representation consists of a single Segment that complies with Indexed Media Segment or Indexed Self-Initializing Media Segment, this Representation element can be ignored unless SegmentBase@indexRange is present.

4.3 Segment Format Constraints

For Representations and Segments referred to by the Representations in the profile-specific MPD for this profile, the following constraints shall be met:

- For Periods conforming to clause 4.2.3 or 4.2.6:
 - Representations and segments shall comply with the formats defined in ISO/IEC 23009-1 [1], clause 7.3.
 - In Media Segments, all Segment Index ('sidx') and Subsegment Index ('ssix') boxes, if present, shall be placed before any Movie Fragment ('moof') boxes.
 - Subtitle segments shall be available at or before the time at which other media segments with which they are presented become available.
 - All the initialization segments for Representations within an Adaptation Set shall have the same sample entry type (see ISO/IEC 14496-12 [6]). For example the inclusion of 'avc1' and 'avc3' based Representations or 'avc3' and 'hev1' based Representations within an Adaptation Set is not permitted.
 - The movie fragment box ('moof') shall contain only one track fragment box ('traf').
 - All Representations within an Adaptation Set shall have the same track_ID. The track_ID is located in the track header box and track fragment header box.
- For Periods conforming to clause 4.2.6:
 - Each Representation shall have only one Segment. This segment shall comply with the Self-Initializing Media Segment as defined in clause 6.3.5.2 of ISO/IEC 23009-1 [1].
 - The segment shall contain only one single Segment Index box ('sidx') for the entire segment.

4.4 Presence of Attributes and Elements

Elements and attributes are expected to be present for certain Adaptation Sets and Representations to enable suitable initial selection and switching.

Specifically the following applies:

- For any Adaptation Sets with @contentType="video" the following attributes should be present:
 - @maxWidth (or @width if all Representations have the same width)
 - @maxHeight (or @height if all Representations have the same height)
 - @maxFrameRate (or @frameRate if all Representations have the same frame rate)

- @par (Picture Aspect Ratio)

NOTE: The attributes @maxWidth and @maxHeight are expected to be used such that they describe the target display size. This means that they may exceed the actual largest size of any coded Representation in one Adaptation Set.

- For any Representation within an Adaptation Set with @contentType="video" the following attributes shall be present:
 - @width, if not present in the **AdaptationSet** element
 - @height, if not present in the **AdaptationSet** element
 - @frameRate, if not present in the **AdaptationSet** element
 - @scanType, if interlaced pictures are used within any Representation in the Adaptation Set
- For any Representation within an Adaptation Set with @contentType="video" the following attributes should be present or inherited from the Adaptation Set:
 - @sar (Sample Aspect Ratio)

4.5 Dimension Constraints

It is mandatory that:

- The MPD size before and after xlink resolution shall not exceed 256 Kbytes.
- The MPD has a maximum of 64 periods before and after xlink resolution.
- The MPD has a maximum of 16 adaptation sets per period.
- The MPD has a maximum of 16 representations per adaptation set.
- Segment duration shall be at least 1 second, except for the last segment of a Period which may be shorter.
- Subtitle segments shall have a maximum segment size of 512 kB.
- Where subsegments are not signalled, each video segment shall have a duration of not more than 15 seconds.
- Where subsegments are not signalled, each audio segment shall have a duration of not more than 15 seconds.
- Each video subsegment shall have a duration of not more than 15 seconds.
- Each audio subsegment shall have a duration of not more than 15 seconds.

4.6 Server Requirements

Servers shall support HTTP 1.1 as defined in RFC 2616 [17] and the use of Cookies as specified in RFC 6265 [18].

4.7 Availability Time Synchronization between Player and Server

4.7.1 Background

MPEG DASH includes provision for MPDs and segments that appear and disappear over time. In order to properly access these, DASH servers and Players should synchronize their clocks to a globally accurate time standard.

Specifically Segment Availability Times are expected to be accurately announced in the MPD and the Player needs to have access to the same time base as the MPD generation in order to enable a proper service. In order to ensure this, this clause provides server and Player requirements to ensure proper operation of a live service.

4.7.2 Service Provider Requirements

If the MPD is dynamic or if the **MPD@availabilityStartTime** is present then the service provider shall provide an MPD as follows:

- The MPD should contain at least one **UTCTiming** element with the @schemeIdURI attribute set to one of the following:
 - urn:mpeg:dash:utc:ntp:2014
 - urn:mpeg:dash:utc:http-head:2014
 - urn:mpeg:dash:utc:http-xdate:2014
 - urn:mpeg:dash:utc:http-iso:2014
 - urn:mpeg:dash:utc:http-ntp:2014
- If the MPD does not contain any **UTCTiming** element then the segments shall be available at the latest at the announced segment availability time using a globally accurate timing source with a tolerance of at most 200 ms.
- If the MPD contains an **UTCTiming** element then:
 - the announced timing information in the **UTCTiming** element shall be accessible to the DASH Player; and
 - the segments shall be available at the latest at the announced segment availability time in the MPD for any device that uses one of announced time synchronization methods at the same time.

If urn:mpeg:dash:utc:http-head:2012 is used, then the server specified in the **UTCTiming@value** attribute may be the server hosting the DASH segments such that with each request the Date general-header field in the HTTP header (see RFC 2616 [17], clause 14.18) can be used by the Player to maintain synchronization.

Leap seconds can be added or removed from UTC periodically. This may affect long running live DASH streams. A service provider may wish to take into account the considerations of RFC 7164 [25].

Servers may provide time using a different time zone to the one being used by Players or in the MPD.

4.7.3 Player Requirements

If the MPD is dynamic or if the **MPD@availabilityStartTime** is present, then the Player does the following:

- If the MPD does not contain any **UTCTiming** elements, the Player should acquire an accurate wall-clock time from its system. The anticipated inaccuracy of the timing source should be taken into account when requesting segments close to their segment availability time boundaries.
- If the MPD contains a **UTCTiming** element with the @schemeIdURI attribute set to "urn:mpeg:dash:utc:http-head:2012" or "urn:mpeg:dash:utc:http-xdate:2012" then the following requirements apply:
 - The Player shall use one of the timing information sources listed in the MPD to synchronize its clock.
 - The Player shall not request segments prior to the segment availability start time with reference to any of the chosen **UTCTiming** methods. The Player may take into account the accuracy of the timing source as well as any transmission delays when it makes segment requests.
- If the Player observes that segments are not available at their segment availability start time, the Player shall use the recovery methods defined in clause 10.8.
- Players should not access the server indicated in the **UTCTiming** element more frequently than is necessary. For example, it should not be requesting the time prior to retrieving every segment.

5 DASH Specific Aspects for Video

5.1 DASH Specific Aspects for H.264/AVC Video

5.1.1 Profiles Supported

Video encoded with H.264/AVC shall comply with the limitations set out in ETSI TS 101 154 [3], clause 5.7.1 "Specifications common to all H.264/AVC HDTV IRDs and Bitstreams" and either clause 5.7.2 "25Hz H.264/AVC HDTV IRD and Bitstream" or clause 5.7.3 "30 Hz H.264/AVC HDTV IRD and Bitstream".

Encoded video may use additional luminance resolutions for adaptive streaming listed in Table 17 and Table 18.

5.1.2 H.264/AVC Specifics

The encapsulation of H.264/AVC video data is based on the ISO BMFF as defined in ISO/IEC 14496-15 [4].

Segments shall start with SAP types 1 or 2 as described in ISO/IEC 14496-12 [6].

Players which support H.264 shall support both sample entries using 'avc1' and 'avc3' (both storage for SPS/PPS within the initialization segment or inband within the media segment).

Content should be offered using Inband Storage for SPS/PPS i.e. sample entries 'avc3' and 'avc4' based on ISO/IEC 14496-15 [4]. Content may be offered using either of the 'avc1' or 'avc2' sample entries. In this case, the Initialization Segment shall be common for all Representations within an Adaptation Set and the following shall hold:

- For video Representations, the width and height values in the track header box shall have the nominal display size in square pixels after decoding, H.264/AVC cropping, and rescaling.
- All information necessary to decode any Segment chosen from the Representations shall be provided in the Initialization Segment. For example, movie box for video Representation shall contain H.264/AVC decoder configuration records including all encoding parameters (i.e. Sequence Parameter Sets and Picture Parameter Sets) used for Representations in the Adaptation Sets for the entire period.

Initialization segments being common means that all Representations in an Adaption Set will have identically the same 'std' box. There will be one entry in the 'std' box for each Representation. Representations encoded with different "parameters" will use the sample description_index in the Track Fragment Header to identify which of the sample entries in the 'std' box is applicable to them.

5.1.3 Signalling of the Codec Profile

The video codec profile, level and constraints should be signalled within the MPD using the @codecs attribute. If present, the value of the @codecs attribute shall be set in accordance with RFC 6381 [5], clause 3.3.

An H.264/AVC codec string is therefore of the form described by the following ABNF notation (see RFC 5234 [28]):

```
CODECSTRING = AVCVERSION "." PROFILE CONSTRAINTS LEVEL
AVCVERSION = "a" "v" "c" ("1" / "2" / "3" / "4")
PROFILE = HEXBYTE
CONSTRAINTS = HEXBYTE
LEVEL = HEXBYTE
HEXBYTE = 2(HEXDIG)
```

Where:

AVCVERSION corresponds to the name of the sample description entry within an ISO BMFF File. This is the name of the AVCSampleEntry stored in the 'std' box for the media track.

PROFILE, CONSTRAINTS and LEVEL come from an SPS NAL unit, and are also carried within the AVCDecoderConfigurationRecord:

- PROFILE is a hexadecimal representation of the AVCProfileIndication stored within the AVC Sample Entry.
- CONSTRAINTS is a hexadecimal representation of the profile_compatibility stored within the Visual Sample Entry.

- LEVEL is a hexadecimal representation of the AVCLevelIndication stored within the Visual Sample entry.

The values given for profile, constraints and level shall be such that a Player able to decode the identified profile and level can decode the media, but the values are not required to indicate the minimum capabilities required to decode the media.

The @codecs attribute may be signalled on the AdaptationSet, in which case it shall be possible to decode every representation within the AdaptationSet with a Player with capabilities matching the signalled requirements.

Examples of the signalling for various common profiles and levels are shown in Table 1.

Table 1: H.264/AVC Codec Parameter Examples

Profile	Level	@codec Parameter (avc1 sample entry)	@codec Parameter (avc3 sample entry)
Constrained Baseline	2.1	avc1.42c015	avc3.42c015
Constrained Baseline	3.0	avc1.42c01e	avc3.42c01e
Main	3.0	avc1.4d401e	avc3.4d401e
Main	3.1	avc1.4d401f	avc3.4d401f
High	3.0	avc1.64001e	avc3.64001e
High	3.1	avc1.64001f	avc3.64001f
High	3.2	avc1.640020	avc3.640020
High	4.0	avc1.640028	avc3.640028

5.1.4 Signalling within the AVCSampleEntry in an initialization segment

The AVCSampleEntry contained within the initialization segment of an H.264/AVC track contains a number of values indicating information about the encoded video. These values shall always be set such that the encoded video in all Representations using this initialization segment does not exceed the size, profile, level or compatibility indicated within the initialization segment, however the values are not required to indicate the exact size or nature of the specific stream.

5.2 DASH Specific Aspects for HEVC Video

5.2.1 HEVC Specifics

The encapsulation of HEVC video data in ISO BMFF is defined in ISO/IEC 14496-15 [4]. Players which support HEVC shall support both sample entries using 'hvc1' and 'hev1' (both storage for VPS/SPS/PPS within the initialization segment or inband within the media segment).

IDR pictures with nal_unit_type equal to IDR_N_LP and IDR_W_RADL are mapped to SAP types 1 and 2, respectively. BLA pictures with nal_unit_type equal to BLA_N_LP and BLA_W_RADL are mapped to SAP types 1 and 2, respectively.

NOTE: The mapping to SAP type 3 for ISO BMFF with HEVC deliberately remains undefined until MPEG reaches a conclusion. This includes the mapping of all other types of HEVC DVB_RAP pictures (including BLA pictures with nal_unit_type equal to BLA_W_LP, CRA pictures with nal_unit_type equal to CRA_NUT and pictures with nal_unit_type equal to TRAIL_R that contain only slices with slice_type equal to 2 (I slice), as specified in ETSI TS 101 154 [3], clause 5.14.1.8).

5.2.2 Signalling of the Codec Profile

The video codec profile, level and constraints should be signalled within the MPD using the @codecs attribute. If present, the value of the @codecs attribute shall be set in accordance with ISO/IEC14496-15 [29], clause A.3.

The video codec profile, level and constraints should be signalled within the MPD using the @codecs attribute. If present, the value of the @codecs attribute shall be set in accordance with ISO/IEC14496-15 [29], clause A.3.

NOTE 1: ISO/IEC 14496-15 [29] contains a significant bug with the codec parameters that will be corrected in Corrigendum 1. The text in the present document reflects the correction.

NOTE 2: An HEVC codec string is of the form described by the following ABNF notation (see RFC 5234 [28]):

```
CODECSTRING = CODEC "." PROFILE "." LEVEL "." CONSTRAINTS
CODEC = ( "h" "e" "v" "1" / "h" "v" "c" "1" )
PROFILE = PROFILE_SPACE PROFILE_IDC "." PROFILE_COMPATIBILITY
PROFILE_SPACE = " " / ALPHA
PROFILE_IDC = 1*3(DIGIT)
PROFILE_COMPATIBILITY = 1*8(HEXDIG)
LEVEL = TIER LEVEL_IDC
TIER = "L" / "H"
LEVEL_IDC = 1*3(DIGIT)
CONSTRAINTS = 2(HEXDIG) [ "." CONSTRAINTS ]
```

Where:

CODEC corresponds to the name of the sample description entry within an ISO BMFF File. This is the name of the HEVCSampleEntry stored in the 'std' box for the media track.

The other values are generated from values carried within the HEVCDecoderConfigurationRecord, which itself contains information from HEVC Parameter Sets:

- PROFILE_SPACE is the general_profile_space, encoded as no character (general_profile_space == 0), or 'A', 'B', 'C' for general_profile_space 1, 2, 3. General_profile_space is used to indicate the version of the HEVC specification in which the profile is defined.
- PROFILE_IDC is the decimal representation of general_profile_idc.
- PROFILE_COMPATIBILITY is general_profile_compatibility_flags, but in reverse bit order, in a hexadecimal representation (leading zeroes may be omitted).
- TIER represents the general_tier_flag, with 'L' indicating the flag is 0, and 'H' indicating the flag is 1.
- LEVEL_IDC is the decimal representation of general_level_idc.
- CONSTRAINTS is a hexadecimal representation of the general_constraint_indicator_flags. Each byte is separated by a '.', and trailing zero bytes may be omitted.

NOTE 3: The absence of a constraint being signalled in a @codecs attribute does not imply that the constraint is not met by the media.

The values used shall be such that a Player able to decode the identified profile and level can decode the media, but the values are not required to indicate the minimum capabilities required to decode the media.

The @codecs attribute may be signalled on the AdaptationSet, in which case it shall be possible to decode every representation within the AdaptationSet with a Player with capabilities matching the signalled requirements.

Examples of the signalling for various common profiles, level and constraint combinations are shown in Table 2.

Table 2: HEVC Codec Parameter Examples

Profile	Level	Tier	Constraints	@codec Parameter (hvc1 sample entry)	@codec Parameter (hev1 sample entry)
Main	3.1	Main	None	hvc1.1.6.L93.00	hev1.1.6.L93.00
			progressive_source, frame_only, non_packed	hvc1.1.6.L93.B0	hev1.1.6.L93.B0
			interlaced_source, non_packed	hvc1.1.6.L93.50	hev1.1.6.L93.50
Main	4.1	Main	None	hvc1.1.6.L123.00	hev1.1.6.L123.00
			progressive_source, frame_only, non_packed	hvc1.1.6.L123.B0	hev1.1.6.L123.B0
			interlaced_source, non_packed	hvc1.1.6.L123.50	hev1.1.6.L123.50
Main	5.1	Main	None	hvc1.1.6.L153.00	hev1.1.6.L153.00
			progressive_source, frame_only, non_packed	hvc1.1.6.L153.B0	hev1.1.6.L153.B0
			interlaced_source, non_packed	hvc1.1.6.L153.50	hev1.1.6.L153.50
Main 10	3.1	Main	None	hvc1.2.4.L93.00	hev1.2.4.L93.00
			progressive_source, frame_only, non_packed	hvc1.2.4.L93.B0	hev1.2.4.L93.B0
			interlaced_source, non_packed	hvc1.2.4.L93.50	hev1.2.4.L93.50
Main 10	4.1	Main	None	hvc1.2.4.L123.00	hev1.2.4.L123.00
			progressive_source, frame_only, non_packed	hvc1.2.4.L123.B0	hev1.2.4.L123.B0
			interlaced_source, non_packed	hvc1.2.4.L123.50	hev1.2.4.L123.50
Main 10	5.1	Main	None	hvc1.2.4.L153.00	hev1.2.4.L153.00
			progressive_source, frame_only, non_packed	hvc1.2.4.L153.B0	hev1.2.4.L153.B0
			interlaced_source, non_packed	hvc1.2.4.L153.50	hev1.2.4.L153.50

5.2.3 Profiles Supported

Video encoded with HEVC shall comply with the limitations set out in ETSI TS 101 154 [3], clause 5.14.1 "Specifications Common to all HEVC IRDs and Bitstreams" and either clause 5.14.2 "HEVC HDTV IRDs and Bitstreams" or clause 5.14.3 "HEVC UHD TV IRDs and Bitstreams".

Encoded video may use additional luminance resolutions for adaptive streaming listed in Table 17 and Table 19.

6 DASH Specific Aspects for Audio

6.1 Common DASH Specific Aspects for Audio

6.1.1 Adaptation Sets and Representations

All audio Representations shall either define or inherit the elements and attributes shown in Table 3.

Table 3: Mandatory Elements and Attributes For All Representations within an Audio Adaptation Set

Parameter	Description	Commonality Between Representations
Role	The main parameter to distinguish between the purposes of adaptation sets, e.g. "main".	Shall be common between all Representations in an Adaptation Set.
@mimeType	Specifies the MIME type of the concatenation of the Initialization Segment, if present, and all consecutive Media Segments in the Representation, e.g. "audio/mp4".	Shall be common between all Representations in an Adaptation Set.
@codecs	Specifies the codec for the Representation, e.g. "mp4a.40.5" for HE-AAC.	Should be common between all Representations in an Adaptation Set to ensure seamless transitions (see note).
@audioSamplingRate	Specifies the output sampling rate of the codec in units of Hz, e.g. "48 000" for 48 kHz.	Should be common between all Representations in an Adaptation Set to ensure seamless transitions (see note).
AudioChannelConfiguration	Specifies the channel configuration and layout for e.g. mono, stereo, or 5.1 surround.	Should be common between all Representations in an Adaptation Set to ensure seamless transitions (see note).
NOTE: Clause 10.2 sets out the requirements for seamless transitions between audio Representations.		

6.1.2 Using the Role Scheme to Distinguish Between Different Adaptation Sets

Every audio Adaptation Set shall include at least one Role element using the scheme "urn:mpeg:dash:role:2011" as defined in ISO/IEC 23009-1 [1]. The use of the @value attribute set to "main" for audio content indicates to the Player that the Adaptation Set is the preferred audio Adaptation Set by the Content Provider. If there is only one "main" then this Adaptation Set is then the default audio adaptation set.

If there is more than one audio Adaptation Set in a DASH presentation then at least one of them shall be tagged with an @value set to "main". It is possible to have multiple Adaptation Sets with @value set to "main", however, they shall be distinguished by other attributes such as @lang or @codec.

If multiple Adaptation Sets have an @value set to "main" then the Player will choose which one of these Adaptation Sets is the most appropriate to use and only if all of these are inappropriate, it may choose one with @value set to something other than "main".

If a programme has multiple audio Adaptation Sets with the same codec but with an original soundtrack translated into different languages, for example a film soundtrack originally in Spanish translated into English, then only the primary language shall have the @value set to "main" with all the other languages set to "dub". Players should then evaluate the @lang attribute of the Adaptation Set in order to confirm the audio language matches the dubbing language wanted by the user.

If a programme has multiple audio Adaptation Sets with the same codec but with different original soundtracks in different languages, for example a sports game commentated by multiple commentators in multiple languages, then all language Adaptation Sets shall have the @value set to "main". Players should then evaluate the @lang attribute of the Adaptation Set in order to confirm the audio language which matches the language wanted by the user.

If the programme has multiple audio Adaptation Sets with multiple codecs, the same original sound and the same language, but a Content Provider does not want to favour a codec, then they can set multiple Adaptation Sets with @value to "main" to let the Player choose the Adaptation Set.

If more than one role needs to be set then multiple role elements shall be used.

As shown in Table 4, the combined use of Role and Accessibility Descriptors shall identify Adaptation Sets containing audio description and clean audio streams. For receiver mixed Audio Description the associated audio stream shall use the @dependencyId attribute to indicate the dependency to the related Adaptation Set's Representations and hence also indicate that the associated audio stream shall not be provided as a Representation on its own. Players should ignore audio streams with other Role and Accessibility descriptor attributes that they do not understand.

Table 4: Role and Accessibility Descriptor Values for Audio Description and Clean Audio

Description	Role element element	Accessibility element
Broadcast mix AD	@schemeldUri = "urn:mpeg:dash:role:2011" @value = "alternate"	@schemeldUri = "urn:tva:metadata:cs:AudioPurposeCS:2007" @value = "1" for the visually impaired
Receiver mix AD	@schemeldUri = "urn:mpeg:dash:role:2011" @value = "commentary"	@schemeldUri = "urn:tva:metadata:cs:AudioPurposeCS:2007" @value = "1" for the visually impaired
Clean Audio	@schemeldUri = "urn:mpeg:dash:role:2011" @value = "alternate"	@schemeldUri = "urn:tva:metadata:cs:AudioPurposeCS:2007" @value = "2" for the hard of hearing

Table 5 shows some use cases based on the rules in this clause.

Table 5: Examples of the Use of Role, Accessibility and other attributes for audio Adaptation Sets

Use case	Description of the audio adaptation sets	@codecs	@audioSamplingRate	AudioChannelConfiguration	@lang	Role element @value	Accessibility element @value
Multiple languages with English original soundtrack.	English	mp4a.40.5	48 kHz	stereo	en	main	Element not used
	French	mp4a.40.5	48 kHz	stereo	fr	dub	Element not used
	German	mp4a.40.5	48 kHz	stereo	de	dub	Element not used
Multiple Channel configurations. Content Provider preferring 5.1ch.	5.1ch	ec-3	48 kHz	5.1ch	de	main	Element not used
	2.0ch	ec-3	48 kHz	stereo	de	alternate	Element not used
	7.1ch	ec-3	48 kHz	7.1ch	de	alternate	Element not used
Multiple codecs. Content Provider preferring Dolby and DTS.	Dolby	ec-3	48 kHz	5.1ch	en	main	Element not used
	DTS	dtsh	48 kHz	5.1ch	en	main	Element not used
	AAC	mp4a.40.5	48 kHz	5.1ch	en	alternate	Element not used
Multiple differences. Content Provider preferring AAC.	AAC	mp4a.40.2	44,1 kHz	stereo	en	main	Element not used
	Dolby multichannel	ec-3	48 kHz	7.1ch	en	alternate	Element not used
	DTS high resolution	dtsh	192 kHz	5.1ch	en	alternate	Element not used
Audio Description with Receiver mix AD.	main audio	ec-3	48 kHz	5.1ch	fr	main	Element not used
	receiver mixed AD	ec-3	48 kHz	mono	fr	commentary	1

Use case	Description of the audio adaptation sets	@codecs	@audioSamplingRate	AudioChannelConfiguration	@lang	Role element @value	Accessibility element @value
Audio Description with broadcast mix AD	main audio	dtsh	48 kHz	5.1ch	en	main	Element not used
	broadcast mixed AD	dtsh	48 kHz	5.1ch	en	alternate	1
Clean Audio	main audio	mp4a.40.2	48 kHz	stereo	de	main	Element not used
	clean channel	mp4a.40.2	48 kHz	stereo	de	alternate	2
Multiple languages, multiple codecs. Content provider prefers 5.1ch over stereo.	Dolby Track	ec-3	48 kHz	5.1ch	en	main	Element not used
	DTS Track	dtsh	48 kHz	5.1ch	en	main	Element not used
	Dolby Track	ec-3	48 kHz	5.1ch	fr	Main dub (see note)	Element not used
	AAC	mp4a.40.2	44,1 kHz	stereo	en	alternate	Element not used
	AAC	mp4a.40.2	44,1 kHz	stereo	fr	alternate	Element not used
Multiple languages with French original soundtrack shown in the UK where English is preferred language	English	mp4a.40.5	48 kHz	stereo	en	main dub	Element not used
	French	mp4a.40.5	48 kHz	stereo	fr	main	Element not used
	German	mp4a.40.5	48 kHz	stereo	de	dub	Element not used

NOTE: 2 separate role elements.

6.2 DASH Specific Aspects for HE-AACv2 Audio

The profile for basic stereo audio support shall be the MPEG-4 High Efficiency AAC v2 Profile, level 2 which also means that Players will be able to play AAC-LC, HE-AAC and HE-AACv2 encoded content. When using HE-AAC and HE-AACv2 bitstreams, explicit backwards compatible signalling shall be used to indicate the use of the SBR and PS coding tools. The content preparation shall ensure that each (Sub)Segment starts with a SAP type 1, according to ISO/IEC 23009-3 [2]. The signalling of the @codecs parameter is according to RFC 6381 [5] as shown in Table 6. Clause 6.1.1 of the present document requires the AudioChannelConfiguration element to be defined for all audio adaptation sets. Either of the schemes for identifying audio channel configuration mentioned in clause 5.8.4.7 of ISO/IEC 23009-1 [1] may be used.

The content should be prepared incorporating loudness and dynamic range information into the bitstream also considering DRC Presentation Mode in ISO/IEC 14496-3 [8].

Decoders shall support decoding of loudness and dynamic range related information, i.e. dynamic_range_info() and MPEG4_ancillary_data() in the bitstream.

Table 6: HE-AACv2 Codecs parameter according to RFC 6381 [5]

Codec	MIME type	@codecs attribute	ISO BMFF Encapsulation	SAP Type
MPEG-4 AAC Profile	audio/mp4	mp4a.40.2	ISO/IEC 14496-14 [9]	1
MPEG-4 HE-AAC Profile	audio/mp4	mp4a.40.5	ISO/IEC 14496-14 [9]	1
MPEG-4 HE-AAC v2 Profile	audio/mp4	mp4a.40.29	ISO/IEC 14496-14 [9]	1

NOTE: Since both, HE-AAC and HE-AACv2 are based on AAC-LC, for the above mentioned @codecs, the following is implied:

- mp4a.40.5 = mp4a.40.2 + mp4a.40.5

- $\text{mp4a.40.29} = \text{mp4a.40.2} + \text{mp4a.40.5} + \text{mp4a.40.29}$

6.3 DASH Specific Aspects of Dolby Audio Technologies

The signalling of the different audio codecs and the ISO BMFF encapsulation is in Table 7.

Table 7: Dolby Technologies: Codec Parameters and ISO BMFF encapsulation

Codec	MIME type	@codecs attribute	ISO BMFF Encapsulation	SAP Type
Enhanced AC-3	audio/mp4	ec-3	ETSI TS 102 366 [10], Annex F	1
AC-4	audio/mp4	ac-4	ETSI TS 103 190 [23], Annex E	1

In the context of DASH, the following applies:

- The signalling of the Enhanced AC-3 codecs parameters is documented in ETSI TS 102 366 [10], which also provides information on ISO BMFF encapsulation.
- The signalling of the AC-4 codecs parameters is documented in ETSI TS 103 190 [23]. The document provides information on non-fragmented and fragmented ISO BMFF encapsulation.

AC-4 uses the concept of I-Frames (or sync frames) as known from video codecs. Each fragment of a representation that contains AC-4 shall start with a sync frame. For details please refer to ETSI TS 103 190 [23], clause E.5.

For E-AC-3 and AC-4 the Audio Channel Configuration element shall use the "tag:dolby.com,2014:dash:audio_channel_configuration:2011" scheme URI.

The @value attribute shall contain a four digit hexadecimal representation of the 16 bit field that describes the channel assignment as defined by Table E.5 in ETSI TS 102 366 [10] where left channel is MSB. For example, for a stream with L, C, R, Ls, Rs, LFE, the value shall be "F801" (hexadecimal equivalent of the binary value 1111 1000 0000 0001).

6.4 DASH Specific Aspects of DTS Audio Technologies

The signalling of DTS audio formats using ISO BMFF encapsulation is provided in Table 8.

Table 8: DTS Audio Technologies: Codec Parameters and ISO BMFF encapsulation

Codec	MIME type	@codecs attribute	ISO BMFF Encapsulation	SAP Type
DTS	audio/mp4	dtsc	ETSI TS 102 114 v1.4.1 [11] Annex E	1
DTS-HD (with legacy core)	audio/mp4	dtsh		1
DTS LBR	audio/mp4	dtse		1
DTS-HD (lossless, without legacy core)	audio/mp4	dtsl		1

For all DTS audio formats, the Audio Channel Configuration shall use

"tag:dts.com,2014:dash:audio_channel_configuration:2012" for the @schemeIdUri attribute. The @value attribute shall be set to the whole decimal number in the range of 1 to 32 indicated by channelcount in the DTSSampleEntry.

In addition to the constraints listed above in Table 8, the audio frame duration shall also remain constant for all streams within a given Adaptation Set. This will insure a seamless transition between streams on any audio access unit boundary.

More information on using DTS for DASH delivery can be found in "Implementation of DTS Audio in Dynamic Adaptive Streaming over HTTP (DASH)" [i.2].

6.5 DASH Specific Aspects of MPEG Surround

MPEG Surround, as defined in ISO/IEC 23003-1 [12], is a scheme for coding multichannel signals based on a down-mixed signal of the original multichannel signal, and associated spatial parameters. The down-mix shall be coded with MPEG-4 High Efficiency AAC v2 according to clause 6.2.

MPEG Surround shall comply with level 4 of the Baseline MPEG Surround profile.

The content preparation shall ensure that each (Sub)Segment starts with a SAP type 1, according to ISO/IEC 23009-3 [2].

Table 9: MPEG Surround Codecs parameters

Codec	MIME type	@codecs attribute	ISO BMFF Encapsulation	SAP Type
MPEG Surround	audio/mp4	mp4a.40.30	ISO/IEC 14496-14 [9]	1

NOTE: Since MPEG Surround is based on a down-mix coded with AAC-LC and HE-AAC, for the above-mentioned @codecs, the following is implied:

- mp4a.40.30 = AOT 2 + AOT 5 + AOT 30

6.6 Service Continuity Considerations

6.6.1 Introduction

Scaling the bit rate for maintaining service continuity under difficult network conditions is a key feature of DASH. In many situations it is sufficient to just scale the video bit rate. However, if the bit rate is severely constrained, it can be advantageous to also scale the audio bit rate as well in order to balance audiovisual quality.

The minimum bit rate under which a service should be continued depends on the service requirements but may also be as low as 500 kbit/s in practice. In such situations it is recommended to scale the audio bit rate as well, such that more of the remaining bit rate could be made available for video. In the given example with a total bit rate of 500 kbit/s, a low audio bit rate of 64 - 96 kbit/s is recommended, as modern audio codecs can still provide good audio quality at such bit rates.

6.6.2 Default Operation for Seamless Service Continuity

In order to allow service continuity with seamless bit rate adaptation it is recommended that the audio Adaptation Set used for main service operation includes Representations from the lowest required bit rate. If the service wishes to support bit rates below 96 kbps then it is recommended that the Player supports the HE-AACv2 profile.

6.6.3 Fallback Operation

If the audio codec used for main service operation does not provide sufficient capability to scale down bit rates low enough within an Adaptation Set, an additional audio Adaptation Set should be provided within the DASH presentation.

Adaptation sets with the purpose of providing media streams with low bit rates for fallback operation shall be indicated with an additional Supplemental Property Descriptor.

This descriptor shall have the @schemeIdUri attribute set to "urn:dvb:dash:fallback_adaptation_set:2014", and the @value attribute equal to the @id attribute of the Adaptation Set for which it supports the falling back operation.

This Adaptation Set should be selected by the Player as a fallback to maintain operation under low bandwidth conditions.

An additional low bit rate fallback Adaptation Set shall also be tagged with the same role as the Adaptation Set which it provides the fallback option for.

The two valid configuration options to consider for the fallback adaptation set are:

- An adaptation set using the same codec in a lower configuration scheme (e.g. just stereo/mono) in order to sufficiently scale down the bit rate used for audio.
- An adaptation set using another codec, for instance HE-AACv2, in an appropriate mono/stereo configuration in order to scale down the bit rate used for audio.

NOTE: It is not expected that Player implementations will seamlessly switch between Adaptation Sets with different configurations.

7 DASH Specific Aspects for Subtitles

7.1 Carriage of Subtitles

7.1.1 General

Where provided, subtitles shall be formatted using the EBU-TT-D Subtitling Distribution Format as specified in EBU Tech 3380 [13] and encapsulated in ISO BMFF in accordance with Carriage of EBU-TT-D in ISO BMFF specified in EBU Tech 3381 [14]. Where the subtitles require or may require the downloading of fonts, these shall be provided and signalled as described in clause 7.2.

NOTE 1: As the subtitles are carried as XML within the sample data, the use of HTTP compression is recommended. Clause 10.11 of the present document requires Player support for gzip compression to enable this.

NOTE 2: The present document does not preclude the carriage of subtitles by other means, for example using a standalone EBU-TT-D document provided to a Player by an application. However, such carriage, and the necessary signalling and media synchronization, are outside the scope of the present document.

The signalling of subtitling codecs for the codec parameter uses RFC 6381 [5] as shown in the Table 10:

Table 10: Subtitle Codecs parameter

Codec	MIME type	@codecs attribute	ISO BMFF Encapsulation
EBU-TT-D with ISO BMFF Encapsulation	application/mp4	stpp	EBU Tech 3381 [14] ISO/IEC 14496-12:2012 Amd2 [6] ISO/IEC 14496-30 [19]

The @contentType attribute indicated for subtitles shall be "text".

NOTE 3: At the time of publication of the present document no further specification of document type is possible within the RFC definition of the codecs parameter. Until such a specification is defined "stpp" only should be used. Once such a specification for identifying subtitle formats within the codecs parameter becomes available, other values may be used here providing the value starts with "stpp."

NOTE 4: Encapsulating TTML documents in ISO Media movie fragments, rather than using a stand alone document has three advantages:

- Subtitles for live content are transported using the same DASH mechanism through which audio and video components are delivered.
- The use of movie fragment encapsulation of TTML allows subtitle rendering to take advantage of the media timelines of the ISO Base Media File Format which is used to synchronize audio and video rendering.
- The ISO Base Media File Format timelines define the period (i.e. provide the 'external context' in TTML terminology) during which any given TTML document should be presented. If the TTML documents were not encapsulated in this way an alternative mechanism would be required to achieve the same functionality.

7.1.2 Distinguishing Between Different Uses of Subtitles

Subtitles are used to provide text which is synchronized to other media components and can be displayed on screen for the viewer. There are a number of scenarios in which subtitles might be used, for example to provide:

- access services - typically subtitles for the hard of hearing, which contain both a transcript of the dialogue and also mentions of background music and any significant non-spoken sounds;
- translations of the dialogue of a programme;
- textual commentary related to the programme.

In order to allow a Player to identify the primary purpose of a subtitle track, Role element and Accessibility element descriptors shall be used as necessary and the language attribute shall be set on the Adaptation Set. Table 11 shows the values to be set in these to indicate common subtitle types. There are also examples in Table 12.

Table 11: Signalling different subtitle types

Description	@lang	Role @schemeIdUri "urn:mpeg:dash:role:2011" @value	Accessibility @schemeIdUri "urn:tva:metadata:cs:AudioPurposeCS:2007" @value
Subtitles for the hard of hearing in the same language as the programme	Same as main audio for the programme	main	2 (for the hard of hearing)
Subtitles providing a translation of the main audio of the programme	Language of subtitles	alternate	No Accessibility element required
Subtitles providing additional information for the programme	Language of subtitles	commentary	No Accessibility element required
Programme uses subtitles to translate foreign soundtrack or parts of soundtrack into preferred language	Language of subtitles	main	No Accessibility element required

Table 12: Examples of Distinguishing Between Different Subtitles

Use case	@codecs	@lang	Role element @value	Accessibility element @value
English original soundtrack with subtitles for the hard of hearing and French and German translated subtitles.	stpp	eng	main	2
	stpp	fra	alternate	No Accessibility element required
	stpp	deu	alternate	No Accessibility element required
French and German original soundtrack for a French speaking audience with French translated subtitles and subtitles for the hard of hearing.	stpp	fra	main	No Accessibility element required
	stpp	fra	main	2

7.2 Downloadable fonts

7.2.0 Introduction

This clause defines a mechanism by which the availability of downloadable fonts may be signalled using descriptors within the MPD, the format of fonts made available and requirements on Players to support these.

7.2.1 Signalling within the MPD

7.2.1.1 DVB font download scheme

The availability of downloadable fonts is signalled using the EssentialProperty or SupplementalProperty descriptor. This descriptor shall use the values for @schemeIdUri and @value specified in clause 7.2.1.2. The descriptor shall carry all the mandatory additional attributes defined in clause 7.2.1.3. A descriptor with these properties shall only be placed within an AdaptationSet containing subtitle Representations. If the descriptor is placed elsewhere its meaning is undefined.

The choice of whether to use the EssentialProperty or SupplementalProperty descriptor is based upon whether the subtitles have to be presented with the signalled font only instead of the default font assumed to be embedded in the Player.

Where an EssentialProperty descriptor is used to indicate a font download then a Player shall ignore the containing AdaptationSet if it does not support downloading fonts.

Where a SupplementalProperty descriptor is used to indicate a font download then a Player shall ignore the SupplementalProperty descriptor if it does not support downloading fonts.

Note that a Player which does not conform to this specification will not recognize the @schemeIdUri and is required by ISO/IEC 23009-1 [1] to exhibit the behaviour in the previous two clauses.

7.2.1.2 Scheme Identification in descriptors

A descriptor is signalling a downloadable font in accordance with the present document when it has the @schemeIdUri attribute set to "urn:dvb:dash:fontdownload:2014" and the @value attribute set to 1.

7.2.1.3 Additional attributes on the EssentialProperty and SupplementalProperty descriptors

The following attributes are defined for use on the EssentialProperty and SupplementalProperty descriptors as follows:

```
<xs:attribute name="url" type="xs:anyURI"/>
<xs:attribute name="fontFamily" type="xs:string"/>
<xs:attribute name="mimeType" type="xs:string"/>
```

The semantics of the attributes are as follows:

url: URL of the font to download. This shall be interpreted in the same way as a media URL (i.e. relative URLs can make use of BaseURLs if desired).

fontFamily: the fontFamily name used in EBU-TT-D documents, [13] and [14].

mimeType: indicates the mime type of the resource available from the URL.

All of these attributes are mandatory when a descriptor indicates the scheme identified in clause 7.2.1.2.

7.2.2 Font format

Fonts for download shall be provided in one of the formats listed in Table 13. The mime type used in the descriptor shall be the appropriate one from Table 13.

Table 13: Supported font formats and their corresponding mime types

Font Format	Mime Type to specify in @mimeType attribute
Open Font Format specified in ISO/IEC 14496-22:2009 [20]	application/font-sfnt
WOFF (Web Open Font Format) File Format [22]	application/font-woff

Fonts that define glyphs for vertical scripts shall include the vhea table.

NOTE: The calculation of inter-line spacing in TTML depends on the selected font size and the value of the lineHeight attribute. Implementations are expected to use the appropriate ascender, descender and line gap values from the font's hhea, OS/2 or vhea tables when computing the inter-line spacing for readable text, defaulting to the OS/2 values sTypoAscender, sTypoDescender and sTypoLineGap, in the absence of other information. The font formats here all require the presence of ascender, descender and line gap values. The use of font formats which do not require them has not been permitted because if fonts are supplied that do not include these then the Player would be expected to make a non-deterministic estimate of the inter-line spacing. Such estimation will lead to inconsistent user experience between Players, and the possibility of unreadable text.

7.2.3 Errors

If a Player is unable to download a font for any reason or having downloaded a font is unable to process it, then:

- If the font download was signalled using the EssentialProperty descriptor, the Adaptation Set containing the descriptor shall not be presented.
- If the font download was signalled using the SupplementalProperty descriptor, the Adaptation Set containing the descriptor shall be presented as if the descriptor was not present.

7.2.4 Player support

Players should support the downloading of fonts which are contained in the formats specified in clause 7.2.2 and are signalled using the mechanism in clause 7.2.1. Players which do not support the downloading of fonts shall follow the rules in clause 7.2.1.1 regarding the action to take depending upon the descriptor used to reference the fonts. Players shall treat descriptors using the scheme identification specified in clause 7.2.1.2 as unrecognized if:

- they are missing mandatory attributes;
- they are placed other than inside an Adaptation Set containing subtitles.

In the event of an error using or accessing a downloadable font the rules in clause 7.2.3 shall be followed.

7.3 Example (informative)

The following extract from an MPD shows an Adaptation Set being used to deliver English subtitles for the hard of hearing where the main programme audio is in English and the content provider is making available a font to download.

```
<AdaptationSet id="sub1" startWithSAP="1" mimeType="application/mp4" lang="eng" codecs="stpp">
  <SupplementalProperty schemeIdUri="urn:dvb:dash:fontdownload:2014" value="1"
  dvb:url="http://fonts.example.com/easilyreadablefont.woff" dvb:mimeType="application/font-woff"
  dvb:fontFamily="easyread" />
  <Role schemeIdUri="urn:mpeg:dash:role:2011" value="main" />
  <Accessibility schemeIdUri="urn:tva:metadata:cs:AudioPurposeCS:2007" value="2" />
  <SegmentTemplate startNumber="1" timescale="1000" duration="10000"
  media="$RepresentationID$/Number$" initialization="$RepresentationID$/IS" />
  <Representation id="subs" bandwidth="20000" />
</AdaptationSet>
```

8 Content Protection

8.1 Introduction

The DVB MPEG DASH profile does not specify a full end-to-end content protection system (DRM), rather it specifies the usage of specific parameters that are defined within the MPEG DASH (ISO/IEC 23009-1 [1]) MPD and within the ISO Common Encryption "CENC" (ISO/IEC 23001-7 [7]). This enables multiple DRMs to protect the content, sharing information (such as KeyIDs and IVs), and then using an opaque container for each DRM to carry DRM specific data (such as licence or licence acquisition information).

The protection of content is optional in the DVB MPEG DASH profile but if used, it shall be in accordance with this clause in the present document.

8.2 Background (informative)

Media protected according to ISO/IEC 23001-7 [7] may require DRM specific information to decrypt. ISO/IEC 23001-7 [7] defines a Protection System Specific Header ('pssh') box to carry this information. The DRM is identified thanks to the SystemID parameter in the 'pssh'.

MPEG DASH also allows the carriage of DRM specific information within the MPD inside the ContentProtection element. Including such information in the MPD may allow a Player to start acquiring or processing a DRM licence in advance of downloading an Initialization Segment. One way to use the ContentProtection element is the UUID one (ISO/IEC 23009-1:2013 [1], clause 5.8.5.2 third bullet), which makes use of the same SystemID as CENC.

NOTE: Content Protection owners are encouraged to register their SystemID values in the list available at: <http://dashif.org/identifiers/protection/>.

Initialization Segments may contain one or more 'pssh' boxes within the Movie ('moov') box, one for each content protection system supported. Media segments may also include 'pssh' boxes in Movie Fragment ('moof') boxes to deliver media keys that are changed over time within a track (a technique often called 'key rotation'). The DRM may require information from both 'pssh' boxes from Initialization Segment and Media Segment in order to obtain media keys.

Key Rotation can be used to allow entitlement changes during live content. Key Rotation can be supported by the DVB DASH profile through the use of 'pssh' boxes carried in the Media Segments. Their role is to deliver information about new keys, licences or sub-licences that are required to continue playing the content. When Media Segment 'pssh' boxes are used in this way, care shall be taken to ensure that random access to each segment remains possible and that the DRM Player receives updated information sufficiently far in advance when playing normally to allow uninterrupted playback.

8.3 Encryption of Different Representations

All Representations in the same Adaptation Set shall be protected by the same license, and encrypted with the same key. That means all Representations shall have the same value of 'default_KID' in their 'tenc' boxes in their Initialization Segments. That also means that the ContentProtection descriptor shall be placed at the AdaptationSet level.

In the case of key rotation, key and licence use within an Adaptation Set shall be aligned. That is, no new licence acquisition or key change shall be required to play the same segment of a different representation.

In cases where HD and SD content are contained in one presentation and MPD, but different license rights are given for each resolution, then they shall be contained in different SD and HD Adaptation Sets, each with different **ContentProtection** descriptors in the Adaptation Set.

8.4 MPEG DASH MPD Content Protection descriptors

Players shall be able to process ContentProtection descriptors from the MPD. Players may ignore ContentProtection descriptors with @schemeIdUri attribute values that they do not understand.

Any Adaptation Set containing protected content shall contain one "mp4protection" ContentProtection descriptor as described in ISO/IEC 23009-1 [1], clause 5.8.5.2 first bullet with the following values:

- @schemeIdUri = "urn:mpeg:dash:mp4protection:2011"
- @value = "cenc"

Furthermore, this "mp4protection" Content Protection descriptor should include the extension defined in ISO/IEC 23001-7 [7] clause 11.2:

- @default_KID attribute under the "urn:mpeg:cenc:2013" namespace.
This reflects the default KID value found within the ISO BMFF 'tenc' box.

Any Adaptation Set containing protected content shall also contain one Content Protection descriptor for each supported DRM. Those Content Protection descriptors may contain DRM specific data. These ContentProtection descriptors shall use the UUID urn in the @schemeIdUri attribute as defined in IEC/ISO 23009-1 [1], clause 5.8.5.2, third bullet.

If an AdaptationSet includes one or more ContentProtection elements that the Player recognizes as identifying both a DRM system and the use of the "cenc" encryption scheme then the Player shall not reject or ignore the AdaptationSet solely because it does not also include an "mp4protection" content protection descriptor.

8.5 MPEG Common Encryption

8.5.0 Protection System Specific Header

Players shall be able to process 'pssh' boxes from both the Initialization Segment and from Media Segments. Players may ignore 'pssh' boxes for protection systems that they do not understand.

The initialization segment and/or media segments may carry 'pssh' boxes.

For each DRM SystemID, if DRM specific data is present both within a 'pssh' box in the initialization segment and within a Content Protection descriptor in the MPD, the data within each should carry equivalent information. However, the data carried within the MPD takes precedence over the data carried within the 'pssh'.

8.5.1 Key Rotation

For content using key rotation and conforming to this profile, the KID within an AdaptationSet shall not change more frequently than every 120 seconds.

8.5.2 Random access requirement

Content shall be encrypted such that it can still satisfy the requirement for a type 1 or type 2 SAP as defined by MPEG DASH. This means that a Player having a suitable licence shall be able to decrypt any Media Segment using only information carried in that Media Segment plus information from the Initialization Segment or MPD. Decryption of a Media Segment cannot depend on information only contained within a previous Media Segment.

8.6 Encryption Mode

8.6.0 Encryption Usage

Media data shall be encrypted using AES 128-bit in CTR mode (AES-CTR) using the 'cenc' scheme as defined in ISO/IEC 23001-7 [7], clause 9.

When encrypting NAL formatted video streams, ISO/IEC 23001-7 [7] requires that subsample encryption is used and NAL type and length fields shall remain unencrypted. In addition to that requirement the payload of NAL units shall remain unencrypted for the following NAL types:

- Access unit delimiter.
- Supplemental enhancement information, where the SEI message is listed as "Clear" in the encrypt column in Table 14. For SEI messages not listed in Table 14, or where the entry says "Application specific", consideration should be given to whether the SEI message needs to be accessed prior to decoding the media when deciding whether to encrypt the NAL unit payload.

Table 14: Encryption requirements for SEI Messages

SEI Message Type	Encrypted or Clear
Picture timing	Clear
Pan-scan rectangle	Clear
Filler payload	Clear recommended (known plaintext)
User data registered by ITU-T T.35 [i.3]	Application decision whether Encrypted or Clear
User data unregistered	Application decision whether Encrypted or Clear
Recovery point	Clear
Decoded reference picture marking repetition	Clear

8.6.1 Constraints on the SampleAuxiliaryInformationOffsetsBox

In order to ensure that the Player has access to the sample auxiliary information before it is needed to decrypt a sample, the offsets in any 'saio' box shall reference data that is stored after the start of, and before the end of, the 'moof' box that contains the 'traf' box that contains this 'saio' box.

As a means of meeting this requirement, a Sample Encryption Information box ('senc', defined below) may be included in the 'traf' box of tracks that contain encrypted media samples to provide the initialization vectors and subsample encryption information necessary to decrypt any encrypted media samples.

NOTE: Regardless of whether a 'senc' box is used, the common encryption sample auxiliary information shall be carried within the structure defined in clause 7 of ISO/IEC 23001-7 [7] and be correctly referenced using the 'saio' and 'saiz' boxes.

Sample Encryption Box ('senc')

Box Type	'senc'
Container	Track Fragment Box ('traf')
Mandatory	No (Option if track fragment is encrypted)
Quantity	Zero or one

The Sample Encryption Box contains the sample specific encryption data, including the initialization vectors needed for decryption and, optionally, alternative decryption parameters. It is used when the sample data in the fragment might be encrypted.

Syntax

```
aligned(8) class SampleEncryptionBox
    extends FullBox('senc', version=0, flags)
{
    unsigned int(32) sample_count;
    {
        unsigned int(IV_size*8) InitializationVector;
        if (flags & 0x000002)
        {
            unsigned int(16) subsample_count;
            {
                unsigned int(16) BytesOfClearData;
                unsigned int(32) BytesOfEncryptedData;
            } [ subsample_count ]
        }
    } [ sample_count ]
}
```

Semantics

- flags is inherited from the FullBox structure. The SampleEncryptionBox currently supports the following bit values:
 - 0x2 - UseSubSampleEncryption:
 - If the UseSubSampleEncryption flag is set, then the track fragment that contains this Sample Encryption Box shall use the sub-sample encryption as described in ISO/IEC 23001-7 [7], clause 9.6. When this flag is set, sub-sample mapping data follows each InitializationVector. The sub-sample mapping data consists of the number of sub-samples for each sample, followed by an array of values describing the number of bytes of clear data and the number of bytes of encrypted data for each sub-sample.
- sample_count is the number of encrypted samples in this track fragment. This value shall be either zero (0) or the total number of samples in the track fragment.
- InitializationVector shall conform to the definition specified in ISO/IEC 23001-7 [7], clause 9.2. Only one IV_size shall be used within a file. Selection of InitializationVector values should follow the recommendations of ISO/IEC 23001-7 [7], clause 9.3.
- subsample_count shall conform to the definition specified in ISO/IEC 23001-7 [7], clause 9.2.
- BytesOfClearData shall conform to the definition specified in ISO/IEC 23001-7 [7], clause 9.2.
- BytesOfEncryptedData shall conform to the definition specified in ISO/IEC 23001-7 [7], clause 9.2.

8.7 Mixing Non-Encrypted and Encrypted Content / Information

An MPD including encrypted video and/or audio may include information that is not encrypted such as subtitles (see clause 7 of the present document) and events (see clause 9 of the present document).

Where events are carried in-band in a video or audio stream using the 'emsg' box, common encryption does not support encryption of these.

Players need to respect the timing requirements for presentation of subtitles and events when video and/or audio content are encrypted taking account of the time taken to do the decryption.

NOTE: Encryption of subtitles is not defined by the present document.

9 Carriage of Generic Streams and Events in DASH

9.1 Events

9.1.1 Background (Informative)

MPEG DASH defines a generic event mechanism that supports two delivery methods. Events can be signalled in the MPD, or they can be carried in band within Media Segments. The semantics of the event are the same in each case.

Inband events might be used where the event is closely related to the A/V content and is known prior to the final packaging of the Media Segments. Inband events are particularly well suited for use with live streaming. Examples of their use could include signalling the scoring of points in a sports event, questions being asked on a quiz show, a programme or advert starting within a continuous TV channel, programme ratings changing or there being additional associated content available.

Events in the MPD might be used to indicate events relating to the way content has been assembled and combined for presentation. Examples could include boundaries between content items assembled into a playlist or notification points for reporting on advert viewing.

Since inband events are delivered in Media Segments, Players only have visibility of inband events for Media Segments they have parsed and have not yet discarded. Players joining a stream after an inband event has passed will not see it unless it is repeated. By contrast, MPD events are visible to Players for as long as they are present in the MPD.

This clause covers the use of the DASH event mechanism for the delivery of:

- programme metadata
- arbitrary messages typically intended to be handled by applications
- messages for the Player itself

9.1.2 Content programme metadata

9.1.2.1 Definition

A DASH presentation may consist of one or more consecutive programmes, each of which has its own metadata.

Content programme metadata may be delivered in the MPD by using an EventStream or in Media Segments using an InbandEventStream. It provides content identifiers and basic metadata relating to the current programme.

An EventStream or InbandEventStream element carrying content programme metadata is profiled as follows:

- @xlink attributes with value "on request" in an MPD event may be ignored by Players.
- @xlink attributes may be ignored with InbandEventStream by Players.
- The @schemeIdUri attribute shall be set to "urn:dvb:iptv:cpm:2014".
- The @value attribute for this scheme is defined in Table 15. Other values of the @value attribute are reserved for definition by a future revision of the present document. Players shall ignore event streams with an unrecognized value.

Table 15: Values defined for urn:dvb:iptv:cpm:2014 event scheme

@value attribute	Description
1	Indicates that the events carry Content Programme Metadata in the form of a TV-Anytime BroadcastEvent element, as specified in clause 9.1.2.2.

Events associated with the @schemeIdUri attribute "urn:dvb:iptv:cpm:2014" and with @value attribute of "1" are defined as follows:

- The presentation time (as indicated by the @presentationTime attribute of an MPD event or derived from the presentation_time_delta field of an inband event) shall be set to indicate a time at which the content programme metadata is applicable. It is recommended that where content programme metadata changes, the presentation time of the event carrying new content programme metadata be set as close as possible to the point from which the new metadata is applicable.
- A duration (as indicated by the @duration attribute of an MPD event or the duration field of an inband event) may be defined for the event, indicating the minimum duration for which the content programme metadata can be assumed to be valid. If the duration is undefined, the metadata can be assumed to be valid until the presentation time of the next content programme metadata event.

- An id (as indicated by the @id attribute of an MPD event of the id field of an inband event) may be defined for the event and is mandatory in the case of an inband event.

NOTE 1: If inband events and MPD events are used together, id values may be allocated independently for each delivery mechanism: there is no implied equivalence between events having the same id across different event streams.

- The value of the element (or the message_data field of an inband event) shall be as defined in Table 16.
- If multiple content programme metadata events overlap, the one with the latest presentation time can be considered to take precedence.

NOTE 2: For a live stream, in order to allow for Players joining the stream at any time, it is necessary to repeat content programme metadata periodically if using inband events alone. Alternatively, MPD events could be used to provide metadata for new Players joining the stream, with inband events providing a more accurately timed transition of metadata for existing Players.

9.1.2.2 Event message syntax

The format of the event payload carrying content programme metadata shall be one or more TV-Anytime BroadcastEvent elements that, when placed within the ProgramLocationTable element of the Default TVAMain fragment defined in Table 16, form a valid TVAnytime XML document.

Table 16: Default TVAMain fragment

```
<TVAMain xmlns="urn:tva:metadata:2012" xmlns:mpeg7="urn:tva:mpeg7:2008">
  <ClassificationSchemeTable />
  <ProgramDescription>
    <ProgramInformationTable />
    <GroupInformationTable />
    <ProgramLocationTable />
    <ServiceInformationTable />
    <CreditsInformationTable />
    <ProgramReviewTable />
    <SegmentInformationTable>
      <SegmentList />
      <SegmentGroupList />
    </SegmentInformationTable>
    <PurchaseInformationTable />
  </ProgramDescription>
</TVAMain>
```

An empty event payload indicates that no content programme metadata applies for the associated portion of the content timeline.

If more than one BroadcastEvent element is present, elements beyond the first contain metadata relating to a future time. The first element contains metadata that applies to the point on the content timeline indicated by the presentation time in the DASH event.

In order to carry XML structured data within the string value of an MPD Event element, the data shall be escaped or placed in a CDATA section in accordance with the XML specification 1.0 [26] Players following standard XML parsing rules need take no special action in order to extract valid BroadcastEvent elements from the Event element. No escaping is required when content programme metadata is carried in an inband event.

The BroadcastEvent may contain any metadata defined for it by the TVAnytime specification. This includes, for example, title and synopsis, associated identifiers, related material locators and parental guidance information.

9.1.2.3 Parental rating

Content programme metadata may include a parental rating in the BroadcastEvent/InstanceDescription/ParentalGuidance element. A parental rating delivered in this way shall be considered to take precedence over any parental rating of the same classification scheme that is defined in the MPD using the AdaptationSet/Rating or AdaptationSet/ContentComponent/Rating element for the duration that the content programme metadata event is active or until a different content programme metadata event becomes active, whichever is the earlier.

The following schemes are defined for including parental rating information in Content Programme Metadata:

- "urn:dvb:iptv:rating:2014" scheme where the id is a decimal number representing the minimum recommended age encoded as per ETSI EN 300 468 [15].
- "urn:dvb:iptv:guidance:2014" where the id is either:
 - 'G' indicating that guidance text relating to the content is provided in the ExplanatoryText element within the ParentalGuidance element.
 - 'W' indicating watershed (safe harbour) content that is only suitable for presentation after the watershed. Guidance text may also be provided [26].

Note that regulatory requirements and established practice mean that content can be expected to use other classification schemes and support for these may be required in specific territories.

NOTE: In the MPEG DASH Rating element, the rating scheme and the term identifier are separated into the @schemeIdUri and @id attributes whereas they are combined in the mpeg7:ParentalRating element used by TVAnytime as a single URI using a ':' separator in the case of a URN or a '#' separator in the case of a URL.

9.1.2.4 Example

```
<?xml version="1.0" encoding="UTF-8"?>
<MPD
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="urn:mpeg:dash:schema:mpd:2011"
  xsi:schemaLocation="urn:mpeg:dash:schema:mpd:2011 DASH-MPD.xsd"
  ... etc ...>

  <BaseURL>http://cdn1.example.com/</BaseURL>
  <BaseURL>http://cdn2.example.com/</BaseURL>

  <Period id="1">
    <EventStream schemeIdUri="urn:dvb:iptv:cpm:2014">
      <Event presentationTime="300" duration="1500" id="0"><![CDATA[<BroadcastEvent>
<Program crid="crid://broadcaster.example.com/ABCDEF"/>
<InstanceDescription>
  <Title xml:lang="en">The title</Title>
  <Synopsis xml:lang="en" length="medium">The description</Synopsis>
  <ParentalGuidance>
    <mpeg7:ParentalRating href="urn:dvb:iptv:rating:2014:15"/>
    <mpeg7:Region>GB</mpeg7:Region>
  </ParentalGuidance>
  </InstanceDescription>
</BroadcastEvent>]]></Event>
    </EventStream>
    <AdaptationSet>
      ... etc ...
    </AdaptationSet>
  </Period>
</MPD>
```

9.1.3 Application messages

An application message is some information that is intended to be consumed by a higher level application and which relates to a particular time (and optional duration) on the media timeline. These can be considered similar to Stream Events as defined in ETSI ES 202 184 [16] and used in MHEG, MHP and HbbTV application environments.

Application messages may be delivered in the MPD by using an EventStream or in Representations using an InbandEventStream.

The present document does not define a @schemeIdUri attribute value for application messages. There may be many event streams with many different @schemeIdUri attribute values. Typically, an API will be provided at the application level to allow an application to register to receive events from one or more @schemeIdUri and @value attribute pairs. It is not specified in the present document how this information is passed by the Player to the application.

9.1.4 Events for the Player

If the MPD@type attribute is set to 'dynamic' and the MPD contains a @minimumUpdatePeriod attribute, then this permits the server to update the MPD. Detailed MPD update procedures are provided in ISO/IEC 23009-1 [1], clause 5.4.

However, if the minimum update period is used alone, there is a conflict between wanting clients to discover changes to the MPD quickly, and trying to minimize the server load caused by frequent requests. DASH Inband Events can address this by informing the client of the need to refresh the MPD only when it actually changes.

If an InbandEventStream element with its @schemeIdUri attribute set to "urn:mpeg:dash:event:2012" and the @value attribute set to 1 or 2 is present on a Representation that is currently being decoded, then a DASH player shall only poll for a new MPD if it receives an MPD validity expiration event. If such an InbandEventStream is not present and the MPD is dynamic, the DASH Player shall determine when to reload the MPD according to the MPD@minimumUpdatePeriod attribute. Note that even if expiration events are intended to be used exclusively, the MPD@minimumUpdatePeriod attribute has to still be present in accordance with ISO/IEC 23009-1 [1]. Typically, it would be set to a low value.

Clients are not required to support patching of the MPD. Hence InbandEventStream elements with a @schemeIdUri attribute set to "urn:mpeg:dash:event:2012" with a @value attribute value of 2 may be treated as having a @value attribute value of 1, i.e. the client ignores the message_data and requests the MPD from the server.

If the DASH Player has an MPD which includes a value of the @publishTime attribute that is greater than or equal to the value of the publish time in the message_data of the event message box, then this event message box may be ignored. For example, this may happen when static Media Presentations are converted from a dynamic Media Presentation.

9.1.5 Long duration inband events

An inband event may have a duration that extends beyond the end of the media segment in which it is first delivered. In such cases, the content provider may choose to deliver the event only once, in which case a Player joining the stream at a later segment, or performing random access, will miss the event. Alternatively, the content provider may repeat the event information in a subsequent media segment. Where events are repeated in this way, the requirements defined in this clause apply.

The following requirements on content creation apply to inband events that are to be repeated:

- Each repetition of an event shall have the same @schemeIdUri, @value and @id attribute values.
- The presentation time of each instance of a particular event shall meet the requirements for presentation_time_delta defined in ISO/IEC 23009-1 [1], section 5.10.3.3.4.

NOTE: These constraints mean that each repetition will have a different indicated presentation time.

- The duration of each instance of a particular event shall be set such that when added to the event instance's presentation time, each instance of the event has the same end time. The event end time may be beyond the end of the media segment in which the event is carried.

Clients shall consider an inband event B to be a repetition of a previous inband event A if the @schemeIdUri, @value and @id attribute values of A and B are the same and the presentation time of event B falls before the end time of event A. Where events repeat multiple times, the comparison of presentation time for each received event shall be made against the end time of the repetition that immediately preceded it.

An event with unknown duration shall be considered to have indefinite length for the purposes of this requirement.

Players shall be able to track at least 10 concurrent or overlapping inband events from each inband event stream that it is monitoring. If a further overlapping event is encountered, Players may discard stored information about the oldest event. This may cause any further repetition of the discarded event to be registered as a new event.

9.1.6 Monitoring of InbandEventStreams

Players shall support monitoring of InbandEventStreams in all Representations that it is decoding at the time.

Players shall not download a Representation solely to gain access to an InbandEventStream contained within it.

InbandEventStreams with the same @schemeIdUri and @value attributes that are present in multiple AdaptationSets shall be considered equivalent and only one of them shall be processed at any particular time. Note that there is no requirement that the actual events carried on such InbandEventStreams be identical. For example, an InbandEventStream on a higher bitrate Representation may carry more data than one on a lower bitrate Representation and events accompanying an AdaptationSet in one language may be different to those for an AdaptationSet of a different language.

No equivalence is implied between an InbandEventStream and an EventStream that have the same @schemeIdUri and @value.

Content providers wishing to deliver InbandEventStreams that relate to the presentation as a whole shall ensure that they are duplicated in enough places that the player will receive them whatever selection of Representations it has made.

10 Player Behaviour

10.0 Introduction

This clause contains normative and informative text on client behaviour, however the normative aspects of client codec specific behaviour are dealt with in clause 5 (video), clause 6 (audio), clause 7 (subtitles). Clauses 4, 8 and 9 also contain some client requirements.

10.1 DVB Profile Support

Players shall support the DVB profile MPEG DASH as defined in the present document and all of the requirements and restrictions for that profile shall apply. Other profiles may also be supported.

Players shall be able play the content described by the profile-specific MPD (where the DVB profile is defined in clause 4.1) (but not necessarily other Adaptation Sets or Representations in the MPD discarded as part of the process of deriving the profile-specific MPD). The process of deriving a profile-specific MPD is defined in clause 8.1 of ISO/IEC 23009-1 [1].

10.2 ISO BMFF Player Constraints

Following are constraints on the Player with respect to ISO BMFF format:

- Players shall support the usage of the track fragment run box ('trun') with negative composition offsets in order to maintain audio visual presentation synchronization.

NOTE: Negative composition offsets were added to ISO/IEC 14496-12:2008 in Amendment 3, which is included in [6].

- Players shall support more than one sample entry in the 'std' box and shall support the use of the sample description index in the Track Fragment Header at the start of each segment to identify which of the sample entries is applicable to that segment.
- The Player shall support simultaneous presentation of multiple Adaptation Sets whose Representations have the same track_ID, for example video with track_ID 1 and audio also with track_ID 1.

10.3 Luminance Resolutions and Frame Rates

A Player that supports HD content shall support the decode and display of pictures with the resolutions in Table 17 and Table 18 at all supported frame rates.

NOTE 1: This does not preclude the use of other resolutions within an Adaptation Set, however, a limited number of resolutions are listed here to ease Player testability.

NOTE 2: The resolutions in the table are the resolutions in the Representations within an Adaptation Set. These may not be the same as the final display resolution, and are thus independent of region specific variations that are prevalent in Broadcast TV.

Table 17: Luminance Resolutions for progressive content

Horizontal @maxwidth	Vertical @maxheight
1 920	1 080
1 600	900
1 280	720
1 024	576
960	540
852	480
768	432
720	404
704	396
640	360
512	288
480	270
384	216
320	180
192	108

Table 18: Luminance Resolutions for interlaced content

Horizontal @maxwidth	Vertical @maxheight
1 920	1 080
704	576
544	576
352	288

A Player that supports UHDTV content shall support the decode and display of pictures with the resolutions shown in Table 19 in addition to the resolutions in Table 17 and Table 18.

NOTE 3: This does not preclude the use of other resolutions within an Adaptation Set, however, a limited number of resolutions are listed here to ease Player testability.

Table 19: Luminance Resolutions for UHDTV Progressive Content

Horizontal @maxwidth	Vertical @maxheight
3 840	2 160
3 200	1 800
2 560	1 440

For service continuity, reducing the frame rate may be beneficial at lower bitrates, so lower frame rates than are found elsewhere in the present document are needed. A Player shall support frame rates formed by a division by 2 and 4 of those of the frame rate families defined in clause 10.4 that it supports.

10.4 Audio/Video Switching Between Different Representations in the same Adaptation Set

Players are expected to transition between Representations in order to optimize the quality of the presentation for the network conditions.

It is important that transitions can be made without any disturbance to the video or audio such that if a sufficiently large number of Representations are available, the viewer does not notice the transition.

Where there is a Representation switch on one component, this shall not affect the playback of other components. For example, audio shall not pause, stutter or mute during a video Representation change.

Players shall support seamless switching between video Representations which differ only in any combination of the following properties:

- Frame rate, providing the frame rate is within one of the following families:
 - 25, 50 fps
 - 30/1,001, 60/1,001 fps
 - 30, 60 fps
 - 24, 48 fps
 - 24/1,001 fps
- Bit rate
- Profile and/or level
- Resolution, subject to maintaining the same picture aspect ratio

There is no requirement on Players to support seamless switching between Representations where other properties differ.

Players should only make switches between Representations that can be done seamlessly unless the switch is necessary to prevent interruption to the media presentation due to lack of data. Specifically a Player is not required to switch between interlaced and progressive video Representations in the same adaptation set unless it can do so seamlessly; however, a non-seamless Representation switch is preferred to terminating the presentation.

The video presentation shall be considered seamless if the following conditions are met:

- Temporal alignment with other media streams (e.g. audio) is maintained.
- There is no delay between the end of the last frame presented from a first Representation (as indicated by the presentation time and presentation duration of that frame) and the presentation of the first frame from a new Representation.
- There are no missing frames or black frames.
- There is no visible pixel shift during or as a result of a Representation switch.
- There are no visible changes in colour or luminance introduced by the Player.

Players shall support seamless switching between audio Representations which differ only in any combination of the following properties:

- Bit rate.

For audio presentation to be considered seamless the following conditions shall be met:

- There are no audible clicks, glitches or discontinuities.
- Audio presentation is continuous, with no pauses, missing or extra audio.

- Temporal alignment with other media streams (e.g. video, additional audio) is maintained.

The requirements of clause 4.5.1 of ISO/IEC 23009-1 [1] for seamless switching shall also be supported.

10.5 Playback Across Periods

10.5.1 Background

Content with multiple Periods may be created for different reasons, for example:

- to enable splicing of content, for example for advert insertion;
- to provide a synchronization point to avoid drift in segment numbering;
- to remove or add Representations in an Adaptation Set;
- to remove or add Adaptation Sets;
- to add or remove content offering on certain CDNs;
- to enable signalling of shorter segments, if produced by the encoder.

Typically, no continuity is necessary at a Period boundary in terms of content offering. The content may be offered with different codecs, language attributes, content protection and so on. It is expected that the Player plays the content continuously across Periods, but there may be implications in terms of implementation to make the playout fully continuous and seamless.

Generally, audio/video presentation across Period boundaries is not expected to be continuous. According to ISO/IEC 23009-1 [1], clause 7.2.1, at the start of a new Period, the playout procedure of the media content components may need to be adjusted at the end of the preceding Period to match the *PeriodStart* time of the next Period as there may be small overlaps or gaps in the content in the Representation at the end of the preceding Period. Overlaps (respectively gaps) may result from Media Segments with actual presentation duration of the media stream longer (respectively shorter) than indicated by the Period duration. Also at the start of a Period, if the earliest presentation time of any access unit of a Representation is not equal to the presentation time offset signalled in the *@presentationTimeOffset* attribute then the playout procedures need to be adjusted accordingly.

However, under certain conditions seamless continuation of the presentation across Periods may be achieved, especially if the content is offered properly. This clause provides an overview on content offering requirements and recommendations across Periods and the Player behaviour when downloading and presenting such content.

Note that this clause applies to static and dynamic Media Presentations and it is independent of the use of xlink. The rules apply for the case after remote Periods are dereferenced.

10.5.2 Content Offering with Multiple Periods

10.5.2.1 General

If content is offered with multiple Periods, then the content provider should offer the content such that the total media presentation time of a Representation is as close as possible to the actual Period duration. It is recommended that the Period duration is the maximum of the presentation durations of the Representations contained in the Period.

In certain circumstances the content provider may offer content in the next Period that is a continuation of the content in the previous Period, possibly in the immediately following Period or in a later Period. The latter case applies for example after an advertisement Period has been inserted.

10.5.2.2 Associated Adaptation Sets across Periods

The content provider may express that the media components contained in two Adaptation Sets in two different Periods are *associated* by assigning equivalent Asset Identifiers to both Periods and by identifying both Adaptation Sets with identical value for the attribute **AdaptationSet@id**.

If Adaptation Sets in two different Periods are *associated*, then the following parameters shall be identical for the two Adaptation Sets:

- the language as described by the @lang attribute,
- the media component type described by the @contentType attribute,
- the picture aspect ratio as described by the @par attribute,
- any role properties as described by the **Role** elements,
- any accessibility property as described by the **Accessibility** elements,
- any viewpoint property as described by the **Viewpoint** elements,
- for audio Adaptation Sets, all values and presence of all attributes and elements listed in Table 3.

NOTE: Additional requirements may apply to specific content protection schemes, the details of which are outside the scope of the present document.

10.5.2.3 Period Continuity

Content providers may explicitly signal that Adaptation Sets across Periods are *period-continuous*. It may do this by providing the following signalling:

- From the first Period:
 - PID means the Period@id attribute value.
 - AID means the value of the @id attribute of the Adaptation Set to be Period continuous.
- In subsequent Periods:
 - The Period contains an Adaptation Set with the following:
 - the @id attribute set to AID.
 - a SupplementalProperty descriptor with the @schemeIdUri attribute set to "urn:dvb:dash:period_continuity:2014" and the @value attribute set to PID.

If this is the case, then the following shall hold:

- All Representations in the Adaptation Set in the first Period shall share the same value EPT1 for the earliest presentation time.
- All Representations in the Adaptation Set in a subsequent Period shall share the same value EPT2 for the earliest presentation time.
- The Adaptation Sets with the value of their @id attribute set to AID in the first and subsequent Periods shall be *associated* as defined in clause 10.5.2.2.
- The presentation duration of each Representation in the Adaptation Set with the @id attribute set to AID in the first Period shall be EPT2 - EPT1, where the *presentation duration* of a Representation is identical to the difference between the end presentation time of the Representation and the earliest presentation time of any access unit.
- If a Representation exists in Adaptation Sets that have their @id attribute set to AID in the first and subsequent Periods where these Representations share the same value for their @id attributes, then the following shall hold:
 - the Representations shall have functionally equivalent Initialization Segments, i.e. the Initialization Segment signalled for the Representation in the first Period may be used to continue the play-out of the Representation in subsequent Periods, and

- the concatenation of the Initialization Segment for the Representation in the first Period and all Media Segments in the Representation in first Period and all Media Segments in the Representation in the subsequent Periods shall represent a conforming Segment sequence as defined in clause 4.5.4 of ISO/IEC 23009-1 [1].

Content providers should signal *period-continuous* Adaptation Sets.

10.5.3 Player Requirements and Recommendations

Typically, at a Period boundary no continuity in terms of content offering is ensured. The content may be offered with different codecs, language attributes, different protection and so on.

The Player should play the content continuously across Periods, but there may be implications in terms of implementation to provide fully continuous and seamless playout. It may be the case that at Period boundaries, the presentation engine needs to be reinitialised, for example due to changes in formats, codecs or other properties. This may result in a re-initialization delay. Such a re-initialization delay should be minimized.

If the Media Presentation has the @type attribute set to "static", then any delay caused by re-initialization should not lead to "missed" content, but instead the viewer should find the overall duration of playback has taken longer than expected. However the times reported by the Player to an application or the viewer shall not be affected by these delays (so media time stands still during the delay). If the Media Presentation has @type attribute set to "dynamic" then there is a risk that such re-initialization delays will cause the Player to slowly drift behind live. If this drift accumulates to a significant size, then the Player may need to jump forwards to rejoin the live edge or play it faster than real-time to rejoin the live edge.

If the Player presents media components of a certain Adaptation Set in one Period, and if the following Period has assigned an identical Asset Identifier, then the Player should identify an *associated* Adaptation Set (as defined in clause 10.5.2.2) and, in the absence of other information, continue playing the content in the associated Adaptation Set.

If furthermore the Adaptation Set is *period-continuous* (as defined in clause 10.5.2.2), i.e. the presentation times are continuous and this is signalled in the MPD, then the Player shall seamlessly play the content across the Period boundary under the constraints in clause 10.4. Most suitably the Player may continue playing the Representation in the Adaptation Set with the same @id, but there is no guarantee that this Representation is available. In this case the Player shall switch to any other Representation in the Adaptation Set.

10.6 Trick Mode Implementation (informative)

10.6.0 Introduction to Trick Mode

The terms "trick mode" refers to playback in any mode other than forward playback at the recorded speed of the audio/video content ("1x"). Examples include fast forward, slow motion, reverse, and random access. Random access may involve skipping to a presentation time, a scene, or a chapter to begin playback, and a menu or partial screen pictures may be presented to allow a user to select a position based on the content.

DASH streaming has rate limitations compared to file trick play due to constraints on Segment download speed, but DASH defines special I picture Representations, and can use other techniques that partially overcome download rate constraints.

10.6.1 Player Enabled Trick Modes

10.6.1.0 Introduction to Player Enabled Trick Modes

Trick mode playback is primarily implemented in players, and uses unspecified decoding and synchronization methods that sometimes require decoding partial bitstreams that violate H.264/AVC buffer and timing models. A player may pause or stop a Media Presentation. In this case the player simply stops requesting Media Segments or parts thereof. To resume, the player sends requests for Media Segments, starting with the next Segment after the last requested Segment.

Forward slow motion is simulated by decoding slower than normal and repeating decoded frames output for display. Audio fast forward can be pitch corrected at moderate speeds, but is usually not decoded during reverse or at fast forward video playback. Audio random access is easier than video because audio samples (sync frames) are randomly accessible, short (e.g. 10 - 32 ms vs. more than one second per H.264/AVC Coded Video Sequence), and can be decoded much faster than realtime to present a particular audio sample (e.g. a 48 kHz waveform sample).

When non-sequential or partial Media Segments are passed to a decoder, the `baseMediaDecodeTime` stored in the Track Fragment Decode Time Box ('`tfdt`') combined with the `Period@presentationTimeOffset` in the MPD can be sufficient to locate the movie fragment on the presentation timeline and provide synchronization between Adaptation Sets encoded with separate timelines, origins, and timescales in the same Period.

The '`tfdt`' indicates the decode time of the first decoded sample, and the presentation time of all samples in the Segment can be determined by adding composition offsets stored in the Track Run Box ('`trun`'). Negative composition offsets in a V1 '`trun`' box should be used to adjust the presentation time of the first presented sample to equal the '`tfdt`' time so that video sample presentation will not be delayed relative to audio samples. The composition offsets may differ between Representations due to different image sizes and the number of decoded frames held for reference in the H.264/AVC decoded picture buffer (more frames result in longer maximum removal delay if positive composition offsets are used). Delivery and decoding in ISO Media is not realtime, as it is in a transport stream, so the sole purpose of composition offsets is to re-sequence pictures from decode order to display order. Delaying presentation of the video track relative to audio is a side effect of reordering in only one direction (if positive offsets are used).

ISO Media edit lists or a '`sidx`' stored in each Initialization Segment can be used to offset the start of a video track in a file by the amount of its composition delay, but those options are problematic for adaptively switched Segments. Edit lists interfere with simple bitstream switching (`AdaptationSet@bitstreamSwitching="true"`), and require the DASH player to control presentation offsets of Segments and samples as they are scheduled for decoding in normal decoders. Loosely coupled players and decoders, such as an ECMAScript player adding Segments to an HTML5 Media Source Buffer, may not be able to compensate for variation in composition delay that are not signalled in the Media Segment header (i.e. '`tfdt`'). Special decoders would be required to override movie fragment header timing with Initialization Segment edit lists or '`sidx`' delivered to a decoder in sequence with each related Media Segment, I picture used for trick play, etc.

A player can create a "progress bar" by showing a timescale for the presentation duration in a static MPD, or the timescale from first to last available Segments in a dynamic presentation. Dynamic presentations should calculate Segment expiration when they become older than the PVR window, and Segment availability when the last sample in a Segment of `@maxSegmentDuration` is estimated to be older than current UTC time based on the `@availabilityStartTime` of the presentation according to the server clock, and a safe estimate of UTC at the player. Current play position is the composition time of the currently presented samples minus `@presentationTimeOffset`.

Where DASH is used in combination with an interactive application, it is also possible to implement some "trick modes" by modifying content within that app. For example, I frames can be extracted during the Segment download process and converted to JPEGs. The app can then show these JPEGs as required to construct scene menus, "film strips", etc.

No matter what trick play mode is used, once the desired location is found by the user, the player can use that presentation time to begin normal playback at the nearest access point to that presentation time.

10.6.1.1 Fast Forward H.264/AVC Long GOP fast decoding and frame dropping

Lower bitrate video Representations in an Adaptation Set are typically subsampled with fewer horizontal and vertical samples than the Adaptation Set's nominal display resolution, and can usually be decoded at a higher frame rate in a decoder that is initialized for a higher quality Representation. Output frames may be discarded if they exceed display system refresh rate. If a lower bitrate Representation can be downloaded at the faster decode rate, then fast forward playback can be maintained. For instance, a Representation that is 50 % subsampled horizontally and vertically could be downloadable and decodable at a 4x rate if a full resolution Representation is capable of 1x download and decode.

Representation or SubRepresentation elements may include the `@maxPlayoutRate` attribute to indicate the corresponding Representation or Sub-Representation may be used for the fast-forward decoding. This attribute indicates that the client may play the Representation or Sub-Representation with any speed up to the regular speed times the specified `@maxPlayoutRate` attribute with the same decoder profile and level required by the AdaptationSet. A player may decode lower resolution Representations at a faster rate whether the `@maxPlayoutRate` attribute is present or not.

10.6.1.2 Fast Forward H.264/AVC picture discarding before decoding

H.264/AVC pictures that are not referenced by other pictures may be discarded prior to decoding to play faster than 1x. H.264/AVC picture references are hierarchical, so discarding the highest reference tier removes dependencies on the next tier so that they may also be discarded. ISO Media files only identify independently decodable pictures and disposable pictures (top tier) by default, so discarding intermediate tiers, other than all but I pictures, may result in decoding errors. Sample Groups may be encoded in ISO Media files to provide additional dependency information to allow discarding a higher percentage of pictures for faster playback, if Segments can be downloaded at that rate.

10.6.1.3 Fast forward I picture access

For very fast forward, all but I pictures may be discarded before decoding. If I pictures are repeated once per second, the playback speed would be 50x for 50Hz frame display rate. That is probably not a sustainable Segment download rate. Decoded I frames may be repeated on output to produce lower speed increases, but motion will appear jerky. I pictures may be discarded before decoding for a faster play rate. Normal playback speed should normally be resumed on an IDR picture. H.264/AVC allows continuous decoding from other I pictures, but not all. ISO Media files do not distinguish between the three types of I pictures unless that information is added in sample groups. Resuming on a normal I picture could result in pictures with dependency on previous pictures failing to decode properly.

10.6.1.4 Fast reverse I picture access

Since I pictures are independently decodable, a player can play downloaded Segments in reverse order by parsing the movie fragment box ('moof') to locate I pictures and decode them in reverse order. The frequency of I pictures encoded and refresh duration of each decoded frame will determine the rate of reverse play. If Segments are already buffered as the result of forward play (e.g. instant replay scenario), the reverse rate may not be limited by Segment download rate.

10.6.1.5 Segment random access

Segments that are available according to the MPD@type="static", @availabilityStartTime, PVR buffer duration, current UTC time or Segment Timeline are normally requested and played from the start (first IDR picture). DASH Segments are optimized for simple concatenation and decoding, even when sequenced from different Representations. However, random access to a picture later in a Segment requires decoding from an I picture, and typically the first IDR picture, unless sample groups are present that additionally identify sample groups of I pictures that are followed by pictures that are all decodable from that entry point. Gradual decoding refresh is also possible in some streams when sample groups are present that provide that information. In general, random access to a picture within a Segment will not be seamless because earlier pictures will be decoded before the random access picture can be presented, resulting in a presentation time delay.

10.6.1.6 Partial Segment fast forward or reverse

ISO Media Segments require movie fragment header information, such as sample byte offsets, to locate and decode media samples (i.e. pictures, TTML documents, and audio sync frames). Video decoding can always start on the first sample of a DASH Segment because it is an IDR picture (it may not be presented first, but it is stored and decoded first). In order to increase download speed during IDR picture trick modes, just the first portion of a Segment containing the movie fragment box and a portion of the media data box ('mdat') containing the first (IDR) picture may be downloaded. The byte range of the Segment request may include an approximate estimate larger than the IDR size, or it may be an accurate size derived from 'sidx' and 'ssix' byte range indexes.

10.6.2 Content enabled Trick Modes:

10.6.2.0 Introduction to Content enabled Trick Modes

The trick play modes listed above do not require special content authoring or MPD description. This clause describes trick play modes that depend on special content and MPD authoring.

For instance, including the @maxPlayoutRate attribute in MPDs for lower resolution Representations does not require special content preparation. It is an optional MPD hint to players that they can decode that Representation at the faster frame rate (see clause 10.6.1.). This clause describes specially encoded Representations where @maxPlayoutRate is required to indicate special encoding.

10.6.2.1 Trick Play Representations

Trick Play Representations consist of long duration Segments containing only H.264/AVC IDR pictures that are typically low frame rate and low resolution to make Segments comparable to normal Segment size and download speed.

For example, a Segment duration could be one minute containing sixty 1 Hz images. In this case, if Segments can be downloaded at a rate of one Segment per second, a user can scan backward or forward at a maximum sustained frame rate of 60X.

Trick Play Representations will be stored in a separate Adaptation Set from the normal long GOP H.264/AVC Representations of the same source content to prevent automatic adaptive selection, and to maintain Segment alignment in the normal Adaptation Set. Trick Play Representations will indicate @maxPlayoutRate, and the Adaptation Set that contains them will include the attribute @codingDependency="false", and may contain an Essential Descriptor to prevent players that do not recognize Trick Play Representations from attempting to play the Trick Mode Adaptation Set for normal playback. A Role descriptor may be used to indicate this Adaptation Set is an "alternative" for the normal Adaptation Set. The URN "urn:mpeg:dash:role:2011" is defined to identify the role scheme defined in Table 22 of ISO/IEC 23009-1 [1], and **Role**@value is assigned to Adaptation Sets that contain a media component type to which this role is associated.

If a specific Representation or SubRepresentation element includes the @codingDependency attribute with value set to 'false', then the corresponding Representation or Sub-Representation may be used for both fast-forward and fast-rewind trick modes.

10.6.2.2 Sub Representations

If a specific Representation or SubRepresentation element includes the @maxPlayoutRate attribute, then the corresponding Representation or Sub-Representation may be used for the fast-forward trick mode followed by periodic frame dropping to prevent exceeding the display refresh frequency. The client may play the Representation or Sub-Representation with any speed up to the regular speed times the specified @maxPlayoutRate attribute with the same decoder profile and level requirements as the normal playout rate.

Scalable video coding methods can create SubRepresentations that can be decoded faster than 1X, or are temporal layers that normally result in a lower or higher frame rate at 1X playback, but a lower frame rate SubRepresentation decoded at the normal frame rate will result in fast forward playback. Disposable pictures are preselected by layer and are not decoded when the enhancement layer is not decoded. This has a similar result to discarding disposable pictures based on the coding dependency indicated for each sample in file metadata, except the work of setting appropriate coding dependencies with a steady frame rate is done in advance by the scalable video encoder.

In On Demand Profile Periods, Sub-Representations in combination with Index Segments and Subsegment Index boxes may be used for efficient trick mode implementation. Sub-Representations in combination with Index Segments containing Segment Index boxes ('sidx') and Subsegment Index boxes ('ssix') may be used for trick mode implementation. Given a Sub-Representation with the desired @maxPlayoutRate, ranges corresponding to SubRepresentation@level all level values from SubRepresentation@dependencyLevel may be extracted via byte ranges constructed from the information in Subsegment Index Box. These ranges can be used to construct more compact HTTP GET request.

10.7 Player Buffering

The Player should not buffer more data than equivalent to approximately 300 seconds of normal play in advance of the current play position.

10.8 Player Resiliency

10.8.1 Introduction (Informative)

Broadcasters generally go to great lengths to ensure that their output is reliable - in a broadcast environment this involves making use of multiple paths between operational sites, redundant equipment, and controls to automatically switch between these if a failure is detected. They can then provide a highly available signal from a transmitter. From a Player perspective this signal can either be received or not, it has no way to rectify a lack of signal.

In a world of DASH delivered content this is not the case. There are more opportunities for failures, either temporary or semi-permanent, between the Player and the broadcaster's point(s) of presence on the Internet. There is also much more scope for providing means of working around such failures. Since fault conditions may exist only in parts of the network and hence only be visible to some Players, it is up to the Player to detect problems and make use of the alternatives which broadcasters may offer.

Additionally Players attempting to recover from errors may unintentionally, but through poor fault handling, cause further problems (or in fact be the cause of a problem if the fault was actually a Player error). For example in response to an HTTP 404 (not found) status, a poorly configured Player might rapidly and incessantly retry the request. If a media segment is missing, it is unlikely to suddenly appear (even in the case of a live stream: if the segment is not available when it was supposed to be, that publication point, or something feeding it, has probably died). Such behaviour amongst a group of Players could cause a significantly higher number of requests needing to be serviced, even if the amount of data being transferred is small. In turn this high number of requests may lead to the server becoming unusable, which may affect other assets being distributed from the same system.

These issues can be addressed through the use of multiple server locations and appropriate Player responses to error conditions.

In addition to ensuring reliability, the use of multiple locations from which content can be accessed can be used to spread the load during times of heavy demand. This is another situation where Players will help, as the choice of which end point to use needs to be different for different Players. It can be addressed through the use of a random function in conjunction with weighting factors associated with each endpoint. Weighting factors are important as broadcasters with large audiences need to use CDNs which may have differing capacities.

The aim of this part of the document is to ensure that content providers can achieve the load balancing they require and where problems occur which can be resolved through the use of alternative servers this is done, but where it is not possible the failure is identified quickly. For more information about how these features are used see clause 11.9.

10.8.2 Handling of BaseURLs by Players

10.8.2.1 Use of BaseURLs and new attributes

According to IEC/ISO 23009-1 [1], MPDs may contain BaseURLs at many levels (within the MPD, Period, AdaptationSet and Representation elements). At each of these levels there can be more than one BaseURL. The presence of multiple BaseURLs at any given level indicates that the same content can be obtained from more than one location. This enables a content provider to increase resilience and provide load balancing by offering different distribution routes - for example different CDNs.

Where Players find, having evaluated the rules as specified in ISO/IEC 23009-1 [1] clause 5.6 for resolving relative URLs (if necessary), that multiple BaseURLs are present for an item then they shall make use of the alternatives offered in the event of an error with the first one used, as described in this clause.

A new attributeGroup, containing two attributes, is defined by the present document. It is part of the schema with the namespace `urn:dvb:dash:dash-extensions:2014-1` which is included in the ZIP file attached to the present document. This attributeGroup, called `baseUrlAttributes`, shall only be used to extend the BaseURL element as indicated in clause 10.8.2.1.

The `baseUrlAttributes` group is defined as follows:

```
<xs:attributeGroup name="baseUrlAttributes">
  <xs:attribute name="priority" type="xs:positiveInteger" use="optional" default="1"/>
  <xs:attribute name="weight" type="xs:positiveInteger" use="optional" default="1"/>
</xs:attributeGroup>
```

The two attributes from the `baseUrlAttributes` group are used together with the `serviceLocation` attribute on BaseURL elements to provide a mechanism by which content providers can signal to Players the priority, weight and independence of different BaseURLs. The attributes are defined as:

- `@priority` is a positive integer. It has a default value of 1. Lower values indicate higher priority levels.

- @weight is a positive integer. It has a default value of 1. The weight describes the relative weighting of BaseURLs with the same priority, with the semantics that a BaseURL with a weight of 2 is twice as likely to be chosen by a Player as one with a weight of 1. This is used to indicate how the content provider needs traffic to be distributed between BaseURLs of a given priority and is usually related to the relative capacity of the infrastructure serving the different BaseURLs.
- @serviceLocation (defined in ISO/IEC 23009-1 [1]) is an optional string described further below.

Players shall carry out BaseURL reference resolution as specified in clause 5.6.4 of ISO/IEC 23009-1 [1], with the following caveats relating to the additional attributes:

- The document base URI shall be treated as if it were a BaseURL element with the default values for the @priority and @weight attributes, and no @serviceLocation attribute.
- Players shall associate values of @priority, @weight and @serviceLocation attributes with resolved BaseURL elements. These shall be taken from the BaseURL element that provided an absolute URL for the resolution process, or the document base URI if that was used as the absolute URL. Hence, Players shall ignore any @priority, @weight or @serviceLocation attributes that are present on BaseURL elements containing relative URLs.

The @serviceLocation attribute is used to implement a blacklisting of failed BaseURL locations. To do this the Player shall maintain a list of @serviceLocation values which have failed (see clause 10.8.2.3). When an MPD is first loaded in a session the blacklist shall be empty. The blacklist is retained when the MPD is reloaded by the Player, but discarded when a different MPD is loaded or at the end of the current session.

Whenever the Player needs to construct a list of BaseURLs, using the rules above, all URLs in the list which have a @serviceLocation attribute matching an entry in the blacklist shall be removed from the available list of BaseURLs. Additionally, whenever a BaseURL is removed from the available list of BaseURLs, any other BaseURL with the same @priority value as the BaseURL being removed shall also be removed.

This reduced list of BaseURLs shall be used when selecting a BaseURL as described in the following clauses.

When checking a BaseURL against the entries in the blacklist, an empty or missing @serviceLocation attribute shall not match any entry in the blacklist.

NOTE: The removal of BaseURLs sharing a @priority attribute value with a BaseURL having a blacklisted @serviceLocation is to prevent unintended changes to the relative weightings in clause 10.8.2.2. Also see clause 11.9.4 for more explanation of this.

10.8.2.2 Selection of BaseURL

When a Player needs to use a BaseURL to resolve a reference to external content, such as may be found inside a Segment Template, it shall pick the BaseURL as follows:

- It shall begin by taking the set of resolved BaseURLs present or inherited at the current position in the MPD, resolved and filtered as described in clause 10.8.2.1, that have the lowest @priority attribute value.
- If there is more than one BaseURL with this lowest @priority attribute value then the Player shall select one of them at random such that the probability of each BaseURL being chosen is proportional to the value of its @weight attribute. The method described in RFC 2782 [24] or picking from a number of weighted entries is suitable for this, but there may be other algorithms which achieve the same effect.
- If there are no BaseURLs after applying blacklisting, the Player shall stop playback and report an error.

Once a random selection has been carried out amongst a group of BaseURLs with the same @priority attribute value, then that choice should be re-used if the selection needs to be made again unless the blacklist has been modified or the available BaseURLs have changed.

10.8.2.3 Changing BaseURL

At any point where a Player needs to change BaseURL as directed in clause 10.8.6, the Player shall:

- If the previously used BaseURL had a non-empty @serviceLocation attribute, add that @serviceLocation attribute value to the blacklist. This BaseURL is removed from the list of available BaseURLs, as are any other BaseURLs in the list with the same @priority attribute value.

- Select the set of BaseURLs from the list available BaseURLs that have the lowest @priority attribute value.
- If there is more than one BaseURL in this set, use the rules in clause 10.8.2.2, using the @weight attribute, to select between them.
- If a BaseURL with the same @serviceLocation attribute is in use elsewhere by the Player (for example in accessing content for a different Adaptation Set) then that BaseURL should be replaced following the rules given in this clause when the next request for media is to be made.

10.8.2.4 BaseURL Example (Informative)

Take the following MPD excerpt as an example:

```
<MPD>
  <BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/</BaseURL>
  <BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/</BaseURL>
  <BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/</BaseURL>
  <BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/</BaseURL>
  <BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="B">http://cdn2.example.com/</BaseURL>
  <BaseURL dvb:priority="5" serviceLocation="D">http://cdn4.example.com/</BaseURL>
  <BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/</BaseURL>
  <Period id="p1" start="PT0S">
    <BaseURL>period/</BaseURL>
    <AdaptationSet>
      <SegmentTemplate media="$RepresentationId$/$Number$" initialization="1/IS"/>
      <Representation id="rep1" ... />
      ...
    </AdaptationSet>
  </Period>
</MPD>
```

Initially the Player has an empty serviceLocation blacklist.

When the Player needs to access the media within period "p1" it will resolve the BaseURL within the period element and end up with the following set of BaseURLs:

```
<BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/period/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/period/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/period/</BaseURL>
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/period/</BaseURL>
<BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="B">http://cdn2.example.com/period/</BaseURL>
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="D">http://cdn4.example.com/period/</BaseURL>
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/period/</BaseURL>
```

Note that the BaseURL with serviceLocation "D" has been assigned the default weight of "1". No BaseURLs are excluded since the serviceLocation blacklist is empty.

From this it takes the BaseURL(s) with the lowest priority:

```
<BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/period/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/period/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/period/</BaseURL>
```

It then needs to pick one of them to use. It adds up the weight values of all these BaseURLs which gives a value of 100. It then picks a random integer in the range from 0 to 99 (inclusive). It then selects the BaseURL using that number (called rn here) as follows:

- $0 \leq rn < 10 \rightarrow \text{http://cdn1.example.com/period/}$
- $10 \leq rn < 40 \rightarrow \text{http://cdn2.example.com/period/}$
- $40 \leq rn < 100 \rightarrow \text{http://cdn3.example.com/period/}$

In this case consider the random number it picked is 30, which gives the value of rn. The BaseURL picked is "http://cdn2.example.com/period/" with @serviceLocation "B".

It would then access the initialization segment for the Adaptation Set shown in the example as "http://cdn2.example.com/period/rep1/IS" and the media segments as "http://cdn2.example.com/period/rep1/\$Number\$"

Now consider a fault occurs which triggers the Change URL behaviour described in clause 10.8.2.3. The blacklist of serviceLocations is now updated to include "B".

Upon examining the list of BaseURLs the Player now discards any BaseURL with a serviceLocation of "B" or with a priority of "1".

```
<BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/period/</BaseURL>
<BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="D">http://cdn4.example.com/period/</BaseURL>
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/period/</BaseURL>
```

Now the lowest priority which it can use is 3, so the BaseURL(s) with that priority are selected:

```
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/period/</BaseURL>
```

There is only one BaseURL, so the random number selection can be skipped. The BaseURL value is now "http://cdn3x.example.com/period/".

Now, consider again a condition requiring the Change BaseURL behaviour. This leads to "C" being added to the serviceLocation blacklist. That now contains "B" and "C".

```
<BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="D">http://cdn4.example.com/period/</BaseURL>
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/period/</BaseURL>
```

The Player now finds the lowest priority value is 5 and selects the BaseURLs with that value:

```
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="D">http://cdn4.example.com/period/</BaseURL>
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/period/</BaseURL>
```

As there are two, the Player will pick at random again. It adds up the weight values of all these BaseURLs which gives a value of 2. It then picks a random integer between 0 and 1 (inclusive). It then selects the BaseURL using that number (called rn here) as follows:

- $0 \leq rn < 1 \rightarrow \text{http://cdn4.example.com/period/}$
- $1 \leq rn < 2 \rightarrow \text{http://cdn5.example.com/example/period/}$

In this case consider the random number it picked is 1. The BaseURL picked is "http://cdn5.example.com/example/period/" with serviceLocation "E".

Finally consider a further fault requiring the Change BaseURL behaviour. The value "E" is now added to the serviceLocation blacklist and all BaseURLs with serviceLocation of "E" or priority of "5" are removed from the list of available BaseURLs. The list of the available BaseURLs now contains:

```
<BaseURL dvb:priority="1" dvb:weight="10" serviceLocation="A">http://cdn1.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="1" dvb:weight="30" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="1" dvb:weight="60" serviceLocation="C">http://cdn3.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="C">http://cdn3x.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="B">http://cdn2.example.com/period/</BaseURL> This serviceLocation has been blacklisted
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="D">http://cdn4.example.com/period/</BaseURL> This priority has a serviceLocation which has been blacklisted.
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="E">http://cdn5.example.com/example/period/</BaseURL> This serviceLocation has been blacklisted
```

There are no more available so the Player reports an error to the application controlling it and ends the session.

10.8.3 DNS - HTTP Player

DNS records for a given hostname may provide more than one address. Players should be able to make use of at least one additional address in the event that the first address received in the response is unreachable or refuses connections.

10.8.4 Types of error condition and recovery options (Informative)

10.8.4.0 Introduction

There are a number of types of error condition which may lead to problems. It may not necessarily be obvious which has occurred, but by grouping them into categories a strategy for dealing with each category can be used.

10.8.4.1 Heavy server load

Network congestion in the home or the Player's ISP is not an error condition and also is something which adaptive streaming is designed to work around. However congestion at the server end, or servers suffering from high loading, may be avoidable if alternative servers with more available capacity can be used.

10.8.4.2 Configuration errors

These types of error may appear when MPDs incorrectly contain links to CDNs which have not been configured to serve the content within the MPD or which have lost connectivity with the content provider. Although this is a situation which should not happen, the possibility needs to be taken into account.

10.8.4.3 Authentication errors

Authentication errors could occur for these reasons:

- by mistake - for example as a configuration error,
- intentionally - for example a GeoIP check failing,
- a time limited token having expired.

10.8.4.4 Missing segments

Three problems are collated here as missing segments, however the causes are significantly different and so the response to such errors needs care to avoid causing additional problems. Problems which may lead to segments being irretrievable are:

- Equipment failure at the content provider, leading to media segments from one content production path being unavailable. A Player may be able to recover from this by trying alternate BaseURLs as other paths may still be operating (see clause 10.8.6).
- Player/server time misalignment when playing a live stream. This may cause the Player to attempt to retrieve segments before they become available on the server, or after they have left the timeshift buffer.
- Intermittent faults in the content provider or distribution network causing occasional segments to be unavailable.

10.8.4.5 Miscellaneous request errors

There are some HTTP statuses which indicate an error, but are not expected to be encountered. If they are encountered and the Player is unaware of why it has happened then they should be treated in the same way as a configuration error.

10.8.4.6 Transient connection errors or congestion

This category covers errors which lead to dropped or stalled connections, without an identified cause and which do not recur if the connection is reopened. It also covers erratic data transfer caused by congestion.

10.8.5 Specific errors and their categories

This clause lists errors as they are likely to be seen at a Player and the category they belong in. For descriptions of the categories see clause 10.8.4. Player behaviour is defined in clause 10.8.6.

Table 20: Specific errors and their categories

Connection	HTTP Status Code	Error category	Notes
DNS resolution failed	N/A	Heavy server load	Although this may be caused by something other than heavy server load, the actions to take in response to this are the same.
Host unreachable	N/A	Heavy server load	
Connection refused	N/A	Heavy server load	
Connection or packet transfer ('socket') timeout	N/A	Transient connection problems or congestion	These are problems which are likely to be caused by network errors, possibly close to the Player, which may disappear if the connection retried.
Congestion related problems	N/A	Transient connection problems or congestion	For example unexpectedly slow or bursty connections.
OK	401 (Unauthorized)	Authentication errors	
OK	402, 403	Authentication errors	Might indicate a token is invalid
OK	404 (Not found)	Missing segments	
OK	405 (Method not allowed) 406 (Not acceptable) 407 (Proxy authentication required) 409 (Conflict) 411 (Length required) 412 (Precondition failed) 413 (Request entity too large) 414 (Request-URI too long) 415 (Unsupported media type) 417 (Expectation failed)	Miscellaneous request errors	These should not really occur in this use.
OK	408 (Request timeout)	Miscellaneous request errors	Although this is a retry-able error, it may be the Player is not sending the whole request.
OK	410 (Gone)	Missing segments	Might be used on a live stream to indicate something has been deleted, but this is not guaranteed.
OK	416 (Requested range not satisfiable)	Missing segments	
OK	500 (Internal server error)	Heavy server load	
OK	501 (Not implemented)	Miscellaneous request errors	These should not really occur in this use.
OK	502 (Bad gateway)	Configuration errors	
OK	503 (Service unavailable)	Heavy server load	
OK	504 (Gateway timeout)	Heavy server load	
OK	505 (HTTP version not supported)	Miscellaneous request errors	These should not really occur in this use.

10.8.6 Requirements on Players

Where Players encounter an error identified in clause 10.8.4 as being in one of the categories described in clause 10.8.5 then they need to take action to try to enable them to continue.

This clause gives some mandatory requirements to ensure that failover between BaseURLs happens in the event of a failure, and to avoid flooding content servers with requests in the event of a configuration error or complete stream failure.

Players shall follow the action specified in the "Action to take" column of when they encounter an error, from the row with the appropriate values in the "Error category" and "MPD@type" columns of that table. Where an action specifies "retry" the Player should close and reopen the connection to the server before making the HTTP request. Players should use a suitable delay between retries, balancing normal backoff etiquette with the need to avoid disrupting playback. Where an action specifies "change BaseURL", there is no requirement for the Player to remember previous retries after changing BaseURL.

Table 21: Action to take in reaction to errors in the different categories

Error Category	MPD@type	Action to take	Maximum Number of retries
Heavy server load	static or dynamic	The Player may retry the request up to the max number of retries specified. If the problem persists it shall change BaseURL as specified in clause 10.8.2.3.	1
Missing segments	static	The Player shall change BaseURL as specified in clause 10.8.2.3.	
Missing segments	dynamic	The Player shall reload the MPD. If the MPD indicates a source of time as specified in clause 4.7.3 the Player shall resynchronize to one of the time sources as described in clause 4.7.3. If, as a result of reloading the MPD and performing any required time synchronization, the Player determines the request is no longer appropriate, it shall adjust its position in the media to reflect the new MPD and any new time value. If the request is still valid the Player may retry the request up to the max number of retries specified. If trying the above does not result in success the Player shall change BaseURL as specified in clause 10.8.2.3.	2
Configuration errors	static or dynamic	The Player may retry the request up to the max number of retries specified. If the problem persists it shall change BaseURL as specified in clause 10.8.2.3.	1
Miscellaneous request errors	static or dynamic	The Player may retry the request up to the max number of retries specified. If the problem persists it shall change BaseURL as specified in clause 10.8.2.3..	1
Authentication errors	static or dynamic	The Player may retry the request up to the max number of retries specified. If the problem persists it shall change BaseURL as specified in clause 10.8.2.3.	1
Transient connection problems or congestion	static or dynamic	Players should attempt to deal with these problems through the adaptive bitrate system, retrying requests where connections time out and changing bitrate if congestion causes poor throughput.	

10.9 Live Services

10.9.1 Overview (informative)

Live services in DASH are provided through the use of mechanisms in the MPD to allow the signalling of availability times and the media segments being short files which can be encoded and then made available over HTTP. The Player retrieves these when they become available according to the MPD. See also [30] for more information on live services using MPEG-DASH.

10.9.2 Determining that an MPD represents a live service

MPEG DASH defines two types of Media Presentations: static and dynamic. A static Media Presentation has all the segments available at the value of the **MPD@availabilityStartTime** attribute whilst a dynamic media presentation has the segments available over time.

By default, an MPD with the @type attribute set to "dynamic" suggests that the Player would want to join the stream at the live edge, therefore to download the latest available segment (or close to, depending on the buffering model), and then start playing from that segment onwards.

However there are circumstances where a dynamic MPD might be used with content intended for playback from the start, or from another position. For example, when a content provider offers 'start again' functionality for a live programme, the intention is to make the content available as an on-demand programme, but not all the segments will be available immediately. This can be signalled to the Player by including an MPD Anchor in the MPD URL provided to the Player. The format and behaviour of MPD Anchors is defined in clause C.4 of ISO/IEC 23009-1 [1].

The position from which a Player begins playback depends upon the MPD attributes @type and @suggestedPresentationDelay, together with any signalling in an MPD Anchor. Specifically:

- If an MPD Anchor is present, the Player shall begin playback at the position indicated by that Anchor.
- If there is no MPD Anchor and the MPD@type attribute is set to "static" the Player shall begin playback at the beginning of the MPD.
- If there is no MPD Anchor and the MPD@type attribute is set to "dynamic":
 - If the MPD@suggestedPresentationDelay attribute is not present, then the Player shall begin playback from a point such that the media is being presented no more than 45 seconds behind the time at which it becomes available.
 - If the MPD@suggestedPresentationDelay attribute is present then the Player shall begin playback from a point such that the media is being presented no more than the greater of the value of the @suggestedPresentationDelay attribute and 45 seconds behind the time at which it becomes available.

NOTE: The value of 45 seconds used in the above bullet points has been deliberately chosen to be three times the maximum segment duration for audio or video segments specified in clause 4.5.

For example, to start from the beginning of the MPD using an MPD Anchor, the following would be added to the end of the MPD URL provided to the Player:

#t=0

Or to start from somewhere other than the start, in this case 50 minutes from the beginning of the period with period ID "programme_part_2":

#period=programme_part_2&t=50:00

NOTE 2: As per clause C.4 of ISO/IEC 23009-1 [1], if the period parameter is not used, the time indicated using the t parameter is as per the field definition of the W3C Media Fragments Recommendation v1.0 clause 4.2.1 [27].

NOTE 3: As per clause C.4 of ISO/IEC 23009-1 [1], the period ID has to be URL encoded/decoded as necessary and needs to match one of the **Period@id** fields in the MPD.

Where an MPD Anchor is used it should refer to a time which is currently available according to the times in the MPD. If the Player receives an MPD Anchor for a time which is not available the error handling behaviour shall be as defined in the W3C Media Fragments Recommendation. Under no circumstances shall the Player request media segments outside of their availability window.

10.9.3 Presentations using segment templates and no segment timeline

Where segments are identified using a segment template using the \$Number\$ replacement with a constant segment duration, the Player may generate segment names up to the time at which the MPD is to be next be updated (according to the MPD@minimumUpdatePeriod attribute). However not all these segments will be available on the server. The Player shall calculate whether a segment is available before requesting it. Players shall not attempt to read segments which, according to the MPD, are not available.

Requests for unavailable segments, particularly those which will become available in the future, cause significant problems for content providers due to difficulties in caching the HTTP 404 status without later preventing access when the resource becomes available.

10.9.4 Considerations for 24/7 operation of a live stream

10.9.4.0 Introduction

Where live television channels are carried over DASH, the stream may run continuously without Periods or any total duration or end time ever being signalled. This does not pose any conceptual problems for a DASH system, but there may be areas to be considered.

10.9.4.1 AvailabilityStartTime

The availability start time could be a considerable time in the past. Players need to be aware of this possibility. Players shall support availability start times at least one hour in the past. This field in the MPD uses the XML Period type, which indicates a time using the UTC format.

The MPD time does not track leap seconds. If these occur during a live service they may advance or retard the media against the real time.

10.9.4.2 Large segment numbers

For a continuously running stream it is highly likely that segment numbers may become very large. When using segment templates with leading zeroes, these may be quite long. Players shall support segment numbers of at least 10 decimal digits.

10.9.4.3 Large track fragment decode times

Track fragment decode times, particularly for large track timescale values, are likely to require the use of 64 bit values. Therefore Players shall support track fragment decode time using 64 bit numbers.

10.10 MPD Processing

Players shall support the DVB DASH profile URN "urn:dvb:dash:profile:dvb-dash:2014". Players shall be able to play the content described by the profile-specific MPD (as defined in clause 8.1 of ISO/IEC 23009-1 [1]) (but not necessarily other Adaptation Sets or Representations in the MPD discarded as part of the process of deriving the profile-specific MPD).

Where there are multiple Adaptation Sets of the same component type (e.g. 2 x video Adaptation Sets), Players shall by default select the Adaptation Set that is signalled with a Role element with a value of "main" from the "urn:mpeg:dash:role:2011" scheme. There is no requirement for a Player to render the "main" Adaptation Set if it understands the logic and signalling of other potentially more appropriate Adaptation Sets or is required by an application to select a different Adaptation Set.

As suggested in the first note in ISO/IEC 23009-1 [1] 5.8.4.6 concerning Frame-Packing, Players shall ignore Representations if the scheme or the value for all FramePacking elements are not recognized by the Player.

Players shall ignore a parent element if the scheme or the value for the EssentialProperty child element is not recognized by the Player (as suggested by the note in ISO/IEC 23009-1 [1] clause 5.8.4.8).

10.11 HTTP Support

Players shall support HTTP 1.1 as defined in RFC 2616 [17] and the use of Cookies as specified in RFC 6265 [18]. Specifically:

- Players shall support byte range requests.
- Players should follow the rules regarding HTTP status and error codes as defined in clause A.7 of ISO/IEC 23009-1 [1].
- Players shall support gzip content coding according to RFC 2616 [17].

- Where Players receive an HTTP redirect status code of 301, 302 or 307 together with a Location header, they shall follow the redirect for that URL as described in clause 10.3 of RFC 2616 [17]. In the event of successive redirects (that is the initial request returns a redirect, and the request for the new location also returns a redirect), Players shall follow at least three redirects. If a Player caches a redirect, in accordance with RFC 2616 [17] that shall not affect requests for different URLs.
- Players shall use the MPD.Location element URL for all MPD updates and not the URL used to initially retrieve the MPD.
- Players shall support keep-alive on TCP connections where further media segments are to be downloaded from the same server.
- Where SSL is supported, SSL session caching shall also be supported.
- Players should use suitable timeouts when opening a TCP connection. It is recommended this should be between 1 and 3 seconds.

NOTE: In addition to connection timeouts, inter-byte timeouts of a similar value may be appropriate, but care is needed to avoid connections which are waiting due to the Player limiting the throughput or connections idle between HTTP requests being 'timed out'.

10.12 Reporting errors to content providers

10.12.1 Introduction (Informative)

In order to manage services in an efficient manner, content providers need to know if Players are encountering difficulties in accessing their services. Automated monitoring is expected to be used to discover complete failures of systems, but, as noted in clause 10.8, problems may be unique to a Player's connection. Additionally transient errors may be difficult to track down, but having these reported when they occur will increase the likelihood of them being resolved.

The present document makes use of the DASH metrics mechanism (see Annex D of ISO/IEC 23009-1 [1]) for the purpose of reporting errors. A list, DVBErrors, is added to the metrics attribute; this list contains information about error events. Additionally a metrics reporting mechanism (see clause 5.9 of ISO/IEC 23009-1 [1]) in which objects within a list are converted into URL query strings and sent using HTTP GET requests is defined. To reduce the number of reports received, an optional attribute @probability is added to the Reporting element to ensure that only a fraction of the Player population submits reports.

10.12.2 Player support

Players shall support the metrics reporting mechanism as defined in clause 10.12.3 and the DVBErrors metric described in clause 10.12.3. Players may ignore other reporting mechanisms or requests to report other metrics systems.

10.12.3 DVB metrics reporting mechanism

10.12.3.1 Scheme Identification

The DVB Metric reporting mechanism defined by the present document is indicated as described in clause 10.12.3.1.

Where the DVB Metric reporting mechanism is indicated in a Reporting descriptor, it shall have the @reportingUrl attribute and may have the @probability attribute defined in clause 10.12.3.2. If a required attribute is missing, the Reporting descriptor may be ignored by the Player, and if it was the only Reporting descriptor then the Metrics element in which it was present may be ignored.

In addition to any time restrictions specified by a Range element within the Metrics element (see clause 5.9.2 of ISO/IEC 23009-1 [1]), Players shall also restrict their reporting based on the probability attribute of the Reporting descriptor, as described in clause 10.12.3.4.

Reports shall be made as described in clause 10.12.3.5.

10.12.3.2 Scheme URI

The DVB Metric reporting scheme defined by the present document is indicated by the @schemeIdUri attribute of "urn:dvb:dash:reporting:2014" and @value attribute of 1.

10.12.3.3 Additional attributes on the Reporting descriptor

The following are in the DVB schema to provide additional attributes required by this metric reporting scheme:

```
<xs:attribute name="probability" type="ReportingProbabilityType" use="optional" default="0"/>
<xs:attribute name="reportingUrl" type="ReportingURLType"/>

<xs:simpleType name="ReportingProbabilityType">
  <xs:restriction base="xs:positiveInteger">
    <xs:maxInclusive value="1000"/>
  </xs:restriction>
</xs:simpleType>
<xs:simpleType name="ReportingURLType">
  <xs:restriction base="xs:anyURI">
    <xs:pattern value="https?:.*"/>
  </xs:restriction>
</xs:simpleType>
```

The attributes are defined as follows:

- @reportingUrl - An absolute HTTP or HTTPS URL. This shall be present when the scheme type is indicated as in clause 10.12.3.2.
- @probability - A positive integer between 0 and 1000, indicating the probability, in thousandths of a whole, of this Player submitting error reports for this session. This enables "sampling" of the Player base for error reports to reduce the number of reports received. If absent it takes the value 0.

10.12.3.4 Use of the @probability attribute

The DVB reporting mechanism allows sampling of the Player population. This enables content providers to reduce the amount of infrastructure required for handling reports, whilst still obtaining enough data to identify and resolve problems. The sampling is enabled through the use of the probability attribute on the Reporting descriptor. The value of this is between 1 and 1000, indicating how many Players per 1000 are to report errors.

When the Player receives an MPD which indicates that the DVB metrics reporting mechanism is to be used, it shall determine whether it is a reporting Player as follows:

- If the @priority attribute is set to 1 000, it shall be a reporting Player.
- If the @priority attribute is missing, the Player shall not be a reporting Player.
- For any other value of the @probability attribute, it shall decide at random whether to be a reporting Player, such that the probability of being one is @probability/1000. For example the Player could pick a random number from 1 to 1000 inclusive and if the number is greater than or equal to @priority, the Player is a reporting Player.

A Player's status, as a reporting Player or not, shall remain static for the duration of the MPD, regardless of MPD updates.

If the Player is not a reporting Player, then the Player shall not report any errors.

NOTE: A Player becoming a reporting Player is itself something which is recorded by the DVBErrors metric.

10.12.3.5 Reporting metrics

This reporting mechanism operates by creating one HTTP GET request for every entry in the top level list of the metric. For example if the metric being reported is "DVB Errors" (see clause 10.8.4) then one GET request is issued for each Entry in the list (with each entry being an individual error) .

Players shall report metrics as soon as practical after the event which gives rise to the metric has occurred. If many entries occur at the same time or in quick succession separate reports shall be made for each.

For each entry in the top level list within the metric (in the case of the DVBErrors metric each entry corresponds to an "error event" described in clause 10.8.4) the Player shall:

- Take each (key, value) pair from the metric entry and create a string consisting of the name of the key, followed by an equals (=) character, followed by the string representation of the value. The string representation of the value is created based on the type of the value following the instructions in Table 22.
- Concatenate the strings created in the previous step with an ampersand (&) character between each one. Eg name=value&name=value.
- Take the value of the @reportingUrl attribute, append a question mark (?) character and then append the string created in the previous step. This results in a string of the form `http://example.com/something?name=value&name=value`.
- Make an HTTP GET request to the URL contained within the string created in the previous step.

Table 22: Encoding of Reporting metric values

Type	How the value is encoded
Integer	Value converted to a decimal string containing no separators or leading zeroes.
String	URL encoded value
Real-Time	Date and time in UTC, formatted as a combined date and time according to ISO 8601 [30].
Media-Time	Time in milliseconds converted to a decimal string containing no separators or leading zeroes.
Enum	URL encoded value
List	Comma delimited objects
Object	Objects (which are essentially used as list entries) consist of the URL encoding of all the attributes of the object (as name=value pairs). The following characters shall additionally be encoded to allow the inclusion within a list: equals = percent % comma ,

The Player shall then make a GET request for the resulting URL. The Player may discard any body returned with the HTTP response.

If the Player is unable to make the report, for example because the @reportingUrl is invalid, the host cannot be reached, or an HTTP status code other than one in the 200 series is received, the Player shall cease being a reporting Player for the duration of the MPD.

10.12.4 Definition of the DVBErrors metric

Table 23 defines the metric for DVB Error events. The name "DVBErrors" shall be used to refer to the metric defined in the table.

Each entry in the DVBErrors list is an "error event". An error event shall be generated each time any of the following occur:

- an error as identified in clause 10.8.5 occurs;
- the Player changes BaseURL following the rules in clause 11.9.2;
- any other error with an errorcode assigned in Table 23, including SSL connection errors, unrecognized HTTP status codes or corrupt media, occurs.

Table 23: Identifiers which are to be substituted within an ErrorURL

Key		Type	Description	
DVBErrors		List	List of error events.	
	Entry			
	mpdurl	String	Absolute URL from which the MPD was originally retrieved (MPD updates will not change this value).	
	errorcode	String	The value of errorcode depends upon the type of error being reported. For an error listed in the ErrorType column below the value is as described in the Value column.	
			Error Type	Value
			HTTP error status code	HTTP status code
			Unknown HTTP status code	HTTP status code
			SSL connection failed	"SSL" followed by SSL alert value
			DNS resolution failed	"C00"
			Host unreachable	"C01"
			Connection refused	"C02"
			Connection error - Not otherwise specified	"C03"
			Corrupt media - ISO BMFF container cannot be parsed	"M00"
			Corrupt media - Not otherwise specified	"M01"
			Changing Base URL in use due to errors	"F00"
			Becoming an error reporting Player	"S00"
			terror	Real-Time
	url	String	Absolute URL from which data was being requested when this error occurred. If the error report is in relation to corrupt media or changing BaseURL, this may be a null string if the URL from which the media was obtained or which led to the change of BaseURL is no longer known.	
	ipaddress	String	IP Address which the host name in "url" resolved to. If the error report is in relation to corrupt media or changing BaseURL, this may be a null string if the URL from which the media was obtained or which led to the change of BaseURL is no longer known.	
	servicelocation	String	The value of the serviceLocation field in the BaseURL being used. In the event of this report indicating a change of BaseURL this is the value from the BaseURL being moved from.	

10.12.5 Example (informative)

This clause demonstrates the use of the DVB metric reporting mechanism to report the metric DVBErrors.

Suppose there is a following entry in the MPD:

```
<BaseURL dvb:priority="1" dvb:weight="1"
serviceLocation="CDN2">http://media.example.com/dash/</BaseURL>
<BaseURL dvb:priority="2" dvb:weight="1"
serviceLocation="CDN3">http://media2.example.com/dash/</BaseURL>
<Metrics metrics="DVBErrors">
  <Reporting schemeIdUri="urn:dvb:dash:reporting:2014" value="1"
dvb:reportingUrl="http://example.com/dash_error" dvb:probability="500"/>
</Metrics>
```

The Player would have a 0.5 probability of choosing to be a reporting Player. Assuming it did choose to be one, it would report becoming a reporting Player with a GET request from the following URL:

```
http://example.com/dash_error?mpdurl=http%3A%2F%2Fmedia.example.com%2Fdash%2Fmanifest.mpd&errorcode=
S00&terror=2014-01-14T20:00:00Z&url=&ipaddress=&servicelocation=CDN2
```

Now suppose the Player encountered an HTTP status of 503 when requesting a media segment and as a result of the error moves to a new BaseURL. Two error reports would be generated and hence two GET requests made, as follows:

```
http://example.com/dash_error?mpdurl=http%3A%2F%2Fmedia.example.com%2Fdash%2Fmanifest.mpd&errorcode=
503&terror=2014-01-14T20:18:00Z&url=
http%3A%2F%2Fmedia.example.com%2Fdash%2Fstream%2Fmedia%2F456&ipaddress=123.45.67.89&servicelocation=
CDN2
http://example.com/dash_error?mpdurl=http%3A%2F%2Fmedia.example.com%2Fdash%2Fmanifest.mpd&errorcode=
F00&terror=2014-01-14T20:18:00Z&url=&ipaddress=&servicelocation=CDN2
```

10.13 Reporting media duration and current play position to higher level applications (informative)

When a DVB Player is operating under the control of a separate application, there will typically be an API which exposes timing information to the application.

It is important that applications are presented with accurate and consistent timing information so that applications can interact with the viewer in the context of the current media timeline. Therefore, where timing information is exposed, the representation of *current play position* and *media duration* should be as described here.

Subject to the additional guidelines below on Period removal and xlink, *current play position* should be calculated as the total duration of media described by the MPD up to the current play position. That is the start time of the current Period plus the sum of all media segment durations before the current segment in the current Period plus the sum of all sample durations in the current segment prior to the sample currently being presented. The start time of the current Period is determined from its Period@start attribute or from the sum of previous Period@duration attributes going back to the most recent Period with a Period@start attribute, as described by clause 5.3.2 of ISO/IEC 23009-1:2013 [1].

For static MPDs which contain no xlink references, *media duration* should be the total duration of the MPD, as indicated by the MPD@mediaPresentationDuration attribute. For MPDs which contain items included by xlink the MPD@mediaPresentationDuration attribute may be incorrect after xlink resolution.

For dynamic MPDs, if the MPD@mediaPresentationDuration attribute is present this should be used as the *media duration*, otherwise the *media duration* should be reported in one of the following ways:

- If the API supports reporting of an indeterminate duration, this should be used.
- If the API does not support this, the *media duration* should signal to the application the extent of media currently available so that the presentation ends at the last segment currently available. The *media duration* will therefore grow as time elapses. Subject to the additional guidelines below on Period removal and xlink, this means the reported *media duration* is the start time of the current period plus the sum of all media segment durations up to the last available segment.

For dynamic MPDs describing live services, Periods may be removed from the beginning of the MPD and others added at the end as time advances. Removal of Periods can be detected using the Period@id attribute which is mandatory for a dynamic MPD. Players should retain knowledge of the total duration of Periods that have been removed from the MPD since playback began and continue to report current play position and duration with respect to the datum determined when the MPD was first retrieved.

If, due to an MPD update, the duration of a Period in a dynamic MPD changes or a new Period is inserted, the *media duration* and *current play position* should be recalculated and the new values reported to the application.

APIs should surface the following additional information which is necessary for applications to handle live streams correctly:

- The value of MPD@availabilityStartTime when the MPD was first retrieved. This provides an absolute time reference for the relative play position values in a live presentation.
- The region of the media timeline which is available for seeking. For a live stream, this is determined by MPD@timeshiftBufferDepth.

It is recommended that APIs exposing times are defined with a suitable range and precision to ensure that individual frames can be identified and to handle the very long durations which are typically associated with live services. In order to indicate time with millisecond precision for a stream that began more than a couple of months in the past, an integer with more than 32 bits is required.

Where a more advanced timing model is exposed through an API, for example indicating the Period structure, it may be appropriate to provide the current play position using the fields defined for the MPD Anchor defined in clause C.4 of ISO/IEC 23009-1:2013 [1].

11 Content Provider Guidelines (Informative)

11.1 Use of Profiles

All Representations that are intended to be decoded and presented by a DVB conformant Player should be such that they will be inferred to have an @profiles attribute that includes the profile name defined in clause 4.1, as well as either the one defined in clause 4.2.5 or the one defined in clause 4.2.8.

11.2 Video Resolutions and Frame Rates in Adaptation Sets

11.2.0 Introduction

A content provider can choose any number of representations using any video resolution and frames per second within an adaptation set. This clause provides some guidance to help the content provider choose which would be preferable sets of video resolution and frames per second, and the tradeoffs.

Adaptive Bit Rate preparation usually consists of two steps, though they may both be within the same product: Encoding and Packaging. The encoding is very similar to traditional broadcast encoding, taking the playout stream in, and producing a single encoded stream out quite often in the form of a standard MPEG-2 Transport Stream. The stream is then demultiplexed, scaled and segmented into several aligned streams, each of decreasing bandwidth, that are then stored on a web server referenced by URLs in the MPEG-DASH manifest.

The choice of video resolutions and frame-rates are more art than science, given the many trade-offs possible. DVB has chosen the resolutions and frame rates from the experience of our members, but they may not suite your chosen service.

11.2.1 Video Resolutions

It is recommended that the resolutions for progressive representations are taken from the list in Table 24. The Table is aimed at several types of devices and connectivity, so you may want to restrict the lower horizontal and vertical resolutions if you are delivering to a Connected TV connected by high speed broadband.

Table 24 shows the source resolution, which may be scaled from the original resolution by the MPEG-DASH Packager. This is different from the display resolution of the IRD, which is as described in ETSI TS 101 154 [3], which is the resolution decided by the Player.

Table 24: 16:9 Video Resolutions for Progressive Representations

Horizontal @maxwidth	Vertical @maxheight
3 840	2 160
3 200	1 800
2 560	1 440
1 920	1 080
1 600	900
1 280	720
960	540
768	432
640	360
480	270
384	216
320	180
192	108

11.2.2 Frame Rates

A Content Provider should check the capabilities of the device population as the capabilities of phones and tablets may not be able to display high resolution at high frame rates. This means a Content provider should either provide:

- Two representations, one at the artistically chosen fps, and the other for devices that can only accommodate the lower fps.
- One representation at the lowest common denominator, usually the lower fps.

Service Continuity considerations may also require lowering the frame rate for lower resolutions, such as for resolutions below 768×432, to save on bandwidth.

Where frame rate needs to be lowered, it is recommended that this be done as a pre-process before encoding, every other frame in the original 50p or 59,94p being deleted (and additional/made-up/synthesised "motion blur" added) giving a 25p or 29,97 "judder-version" to the DASH packager.

The frame rates used have to be kept within the same families to enable seamless switching, as described in clause 10.4, and should be integer multiples of each other.

11.3 Considerations for Audio in Adaptation Sets

11.3.0 Recommendation for Audio as a Percentage of Total Stream Bandwidth

If the service being delivered is a video service, then audio should be 20 % or less of the total stream bandwidth, exceptions to this recommendation can occur at very low video bit rates (e.g. in a mobile environment) or when then audio is the primary program feature.

11.3.1 Considerations for Advanced Audio Coding

This clause provides additional guidelines for Advanced Audio Coding (AAC) including the selection of AAC Profiles and typical bit rates. It also covers considerations on backward compatibility. It should be noted, that those guidelines are informal recommendations which can be adapted to fit the need of a particular DVB service.

Backwards Compatibility of AAC Profiles: There are three AAC Profiles that are considered in the following: AAC-LC, HE-AAC, and HE-AACv2. Those profiles are backwards compatible in the sense, that HE-AACv2 decoders can always decode AAC-LC and HE-AAC bit streams. However, for the encoding of a particular bit stream (or Representation) the service provider has to decide for one of those profiles. Recommendations for this choice are given below.

HE-AAC: The HE-AAC Profile can be seen as the default AAC Profile for DASH. The HE-AAC Profile can be used over a wide range of bit rates and can reach excellent audio quality for mono, stereo, and surround. For stereo audio, it can provide good quality down to 24 kbps and can improve quality consistently by adding more bit rate - up to 160 kbps at which point the audio quality is excellent even for critical items. For 5.1 surround good audio quality can be maintained down to 64 kbps and broadcast quality can be reached at 160 kbps. Higher surround bit rates help to improve quality for critical items. Hence, the complete range from lowest bit rate to highest quality can be covered with a single Adaptation Set allowing seamless bit rate switching and service continuity.

HE-AACv2: The HE-AACv2 Profile allows even lower bit rates by using the Parametric Stereo (PS) tool. Compared to HE-AAC there is an advantage for bit rates below 32 kbps. However, it does not scale to excellent audio quality when increasing the bit rate beyond 48 kbps. Therefore HE-AACv2 is mainly advantageous for audio-only stereo services where a difference of 16-32 kbps matters. For video services and channels with sufficient bandwidth, HE-AAC is preferred for its greater flexibility and availability. Hence, HE-AACv2 is mainly recommended for audio-only services if the average available channel bit rate is below 32 kbps.

AAC-LC: The AAC-LC Profile can be seen as a legacy profile that is replaced by the HE-AAC Profile for most streaming services. However, for stereo bit rates above 128 kbps AAC-LC can achieve transparent audio quality, which offers a certain advantage over HE-AAC. This may e.g. be of interest to a music streaming services addressing the audiophile user community, because transparent audio quality can be assured also for critical items. Hence, the AAC-LC Profile is a valid option if sufficient bit rate can be assured and transparent audio quality is a key requirement. Another reason for using AAC-LC in DASH may be backwards compatibility to devices not yet supporting HE-AAC as explained in more detail below.

Multiple Adaptation Sets: If a broad range of devices and bit rates will be covered it is possible to offer multiple Adaptation Sets, each containing a single AAC Profile with constant configuration. The MPD will then include several Adaptation Sets with the @codecs parameter as given in clause 6.2. A DASH Player may then e.g. select the HE-AACv2 Adaptation Set when connected via 3G while using the HE-AAC Adaptation Set for WiFi and DSL. However, it will have to stick to this selection for the duration of the current Period. At Period boundaries a switch of Adaptation Set is of course possible again. Though the use of multiple Adaptation Sets is possible and increases flexibility in system design, it has to be done with care as the selection of the "correct" Adaptation Set by the Player needs to be well defined. In general it is therefore recommended to use a single HE-AAC Adaptation Set unless there is good reason to do otherwise.

Typical Bitrates: The tables below show typical stereo and 5.1 surround bit rates for AAC-LC, HE-AAC, and HE-AACv2 Adaptation Sets. Note that each profile has a sweet spot for normal operation (bold) but can also be operated at lower and higher bit rates. The lower bit rates are fall back modes which should only be used temporarily during severe network congestion. The higher bit rates are a saturation mode which may not yield significant gains in quality except for critical content. The bit rates are recommendations only and may change for different service requirements.

**Table 25: Typical stereo bit rates [kbps] for AAC Adaptation Sets
(normal operation shown inside [])**

Profile	AOT	@codecs	bit rate [kbps] for 44,1/48 kHz							
HE-AACv2	2+5+29	mp4a.40.29	18	[24]	[32]	48				
HE-AAC	2+5	mp4a.40.5		24	[32]	[48]	[64]	[96]	128	
AAC-LC	2	mp4a.40.2					64	[96]	[128]	[160] 256

**Table 26: Typical 5.1 surround bit rates [kbps] for AAC Adaptation Sets
(normal operation shown inside [])**

Profile	AOT	@codecs	bit rate [kbps] for 44,1/48 kHz							
HE-AAC	2+5	mp4a.40.5	64	96	[128]	[160]	[192]	256	320	
AAC-LC	2	mp4a.40.2				160	192	[256]	[320]	[384] 448

Backwards Compatibility: In order to support Players which e.g. only support AAC-LC the content creator has to provide the corresponding Adaptation Set. For example, he may offer an HE-AAC Adaptation Set as the default but needs to add an AAC-LC Adaptation Set for backwards compatibility reasons. The legacy player would then select the AAC-LC Adaptation Set because it is the only Adaptation Set he is able to decode. A player supporting the HE-AACv2 Profile does have the choice of selection either of the two Adaptation Sets. However, it is recommended that the player selects the HE-AAC Adaptation Set because it can support lower bit rates and therefore increased stability.

Service Continuity for Surround: The requirement on keeping the channel configuration constant within an Adaptation Set means that the Player cannot switch from e.g. 5.1 surround to stereo within a Period. This is sometimes seen as a limitation for bit rate adaptation, i.e. adapting from high bit rate surround to very low bit rate mono. However, there are several reasons why a constant channel configuration is preferred: First, a switch from e.g. 5.1 to stereo is not perceived as seamless but will cause a discomfort for the listener, especially if the configuration is switching back and forth more frequently. Hence, service providers typically want to have control over this behaviour and not leave it up to the adaptation logic of the Player. Second, a switch of the channel configuration may cause a re-configuration of the output device and therefore a discontinuity. Finally, HE-AAC can provide 5.1 surround audio at 64 kbps, and therefore eliminate the need to switch to a stereo codec. Note that the HE-AAC encoder can allocate the bit rate to the stereo channels if considered advantageous and thus seamlessly switch to a stereo configuration internally. Hence, there is no need to enforce this switch through an external channel configuration. The encoder can do this internally in the most efficient way while keeping the external channel configuration constant.

11.3.2 Considerations for Dolby Enhanced AC-3 and AC-4

Enhanced AC-3 decoders will decode all bitstreams at valid bit rates and with channel configurations as constrained by ETSI TS 102 366 [10]. AC-4 decoders will decode all bitstreams at valid bit rates and with channel configurations as constrained by ETSI TS 103 190 [23]. The number of available playback channels will be determined by the implementation included in the rendering device.

Note also that Enhanced AC-3 and AC-4 bitstreams are not limited to stereo or 5.1 channel audio; typical bit rates for stereo and 5.1 channel presentations are offered as examples.

Table 27: Typical stereo bit rates [kbps] for Enhanced AC-3 and AC-4 Adaptation Sets (normal operation shown inside [])

Profile	codecs	bit rate [kbps]					
Enhanced AC-3	ec-3			64	[96]	[128]	160
AC-4	ac-4	32	48	64	89		

Table 28: Typical 5.1 surround bit rates [kbps] for Enhanced AC-3 and AC-4 Adaptation Sets (normal operation shown inside [])

Profile	codecs	bit rate [kbps]							
Enhanced AC-3	ec-3		128	160	[192]	[224]	[256]	320	384
AC-4	ac-4	96	128	160	192				

Service continuity for surround: An Enhanced AC-3 encoder may allocate available bitrate to the front stereo pair of channels of a multichannel bit stream if considered advantageous, thus maintaining a constant coding mode and removing the possibility that a downstream rendering device (e.g. audio/video receiver) may cause a discontinuity during a coding mode change.

Stream switching and codec latency considerations: Enhanced AC-3 and AC-4 use sophisticated methods in order to achieve highest coding efficiency and at the same time provide important capabilities to enable seamless switching on the system layer.

Enhanced AC-3 streams come with reliable in-band configuration information and constant bit rates. Both Enhanced AC-3 and AC-4 come each with a characteristic but constant coding delay across coding configurations. Therefore switching between different Enhanced AC-3 streams or different AC-4 streams with different audio coding configurations, e.g. on period boundaries, does not require additional considerations for compensation of coding latencies. This can drastically simplify content creation when using different encoder entities to create content for different adaptation sets or different periods (e.g. for different programs or program components or other chapters).

Keeping A/V in sync and at the same time maintaining non-overlapping or gap-free A/V play out at switching points (e.g. ad-insertion) is a challenge. The main reason for this is that audio and video frames usually have different durations and that the end of both audio and video is typically not temporally aligned accurately at the end of segments and periods.

AC-4 comes with the capability to adapt the audio frame rate to commonly used video frame rates.

The A/V alignment feature of AC-4 in conjunction with the static configuration-independent latency of AC-4 can be used to simplify or even completely avoid some of the complex system level challenges introduced by such switching scenarios without introducing artefacts through audio gaps or additional, overlapping audio.

Content encoded with AC-4 should be encoded so that the audio frames and corresponding video frames are temporally aligned.

11.3.3 Considerations for DTS Audio Formats

This clause provides additional guidelines for the DTS-HD family of codecs. The DTS family of codecs includes a number of constant bit rate (CBR) coding modes, a "core + lossless" mode and lossless only (VBR) coding modes. It also supports 96 KHz sampling frequency in both lossy and lossless modes, and 192 KHz audio in lossless coding mode.

Table 29 provides a general relationship between the codec, channel configuration and bit rate for the CBR configurations.

Table 29: DTS Profiles for Digital Delivery

Profile	@codecs	audio sampling frequency (KHz)	Channel Layout	Typical bit rates (kbits/sec)
DTS LBR	dtse	48	stereo	64 to 255
			5.1	160 to 510
			7.1	447 to 768
DTS core	dtsc	48	stereo	192 to 768
			5.1	384 to 1524
DTS HD (core + extension substream)	dtsh	48 or 96	up to 7.1	2 046 to 5 760
DTS-HD (core + lossless extension)	dtsh	48, 96 or 192	up to 7.1	CBR core bit rates from 768 to 1 509, with a VBR lossless extension
DTS-HD (lossless extension only)	dtsl	48, 96 or 192	up to 7.1	Variable bit rate

11.4 Considerations for Interlaced and Progressive in Adaptation Sets

Interlaced and Progressive representations may both be included within the same adaptation set. It is strongly recommended that this only be used when 1 080i is mixed with other progressive representations within the same adaptation set.

11.5 Mixing HD and SD in Adaptation Sets

It is recommended that if the master is in HD with BT.709 [21] colour then this is used to make a lower "SD" resolution set of representations distributed via MPEG-DASH. An SD only capable display device will then take the "SD" resolution version and display it as SD. This allows a single adaptation set to work for a range of devices and a single content protection license for both.

If content protection or other reasons mean that you need to separate HD and SD licenses, then two adaptation sets: one for SD and the other for HD need to be used. The HD can still adapt down into the SD space, if needed for service continuity, however this will add to CDN costs.

11.6 Mixing H.264/AVC and HEVC in Adaptation Sets

The mixing of HEVC and H.264/AVC in the same adaptation set is not allowed (see clause 4.3). The Content Author can include separate Adaptation Sets for AVC and HEVC in the same MPD to allow support of devices that have HEVC, H.264/AVC or both.

11.7 Subtitles

Subtitles carried according to the present document are carried within EBU-TT-D documents encapsulated within ISO BMFF containers. This provides a mechanism through which both live and on-demand subtitles can be accommodated.

Content providers may wish to use media segments for subtitles which have a duration greater than the audio or video segments. For on-demand content the only restriction which applies is the maximum subtitle media segment size in clause 4.5. This limit is likely to be around 3 hours - if the segment size limit is likely to be exceeded then more than one segment needs to be used. For live content the most significant restriction is from the requirement in clause 4.3 for subtitles to be available at or before the time at which the corresponding media is available. This means either the subtitle segments need to be the same duration as the other segments, or if they are longer they need to have an earlier availability time signalled using the @availabilityTimeOffset attribute. This is only likely to be possible if the subtitle production is faster than the video encoding.

When splitting subtitles, into samples, each sample should only contain subtitles which are for display at some time during that sample. Times do not need to be truncated at the start and end of the sample as the Player will carry that out based on the sample's start time and duration, however including subtitles which are not going to be displayed increases the Player's work unnecessarily. It is envisaged that each subtitle media segment will only contain one sample. As each sample is a complete EBU-TT-D document in itself, the use of more samples increases the overhead from document headers and styling information.

HTTP servers supplying subtitle files may find the use of gzip compression has a significant effect on the amount of data delivered.

11.8 Audio/Video Switching Across Period Boundaries

Content Providers should prepare content to ensure that the adaptation set on both sides of the periods have the correct properties if they want seamless switching can happen.

11.9 Signalling multiple serving locations using BaseURLs

11.9.1 Introduction

The present document contains a number of requirements on Players in clause 10.8.2 to enable content providers to control the way Players make use of BaseURLs, including a provision for load balancing, and to require automatic failover between different BaseURLs where the signalling indicates this is available.

These features are intended to be used to provide high level load balancing, for example between different CDNs or datacentres, and redundancy at a similar level - allowing failover in the event of a failure of an entire hosting location, or link from one ISP to a CDN. They are not intended to provide load balancing over large numbers of individual web servers - this is expected to be carried out within hosting centres or CDNs in a manner transparent to the Players.

This clause in the present document provides guidance on the way the signalling was expected to be used in a typical hosting scenario.

11.9.2 Prioritizing BaseURLs

BaseURLs containing absolute URLs can be assigned a priority using the @priority attribute. The value of this is a positive integer, with the lowest value indicating the highest priority. This allows additional BaseURLs to be specified to improve the reliability of a service, but on the basis that some of the BaseURLs are only to be used if higher priority ones cannot be used successfully.

The distance between values of @priority has no effect, so two BaseURLs with priorities of 1 and 2 would be the same as if they had values of 1 and 4.

It is possible for more than one BaseURL to have a given priority value, which allows load balancing to happen between the BaseURLs with that priority, as described in the following clauses.

As an example, if a content provider uses three CDNs, but they charge differently, with cheap-cdn being the cheapest and costly-cdn being the most expensive, that content provider may wish to use cheap-cdn most of the time, but make use of the others if there are problems. In that case they would create three BaseURLs, with the @priority attribute being used to signal the order in which they are used.

```
<BaseURL dvb:priority="1">http://cheap-cdn.example.com/</BaseURL>
<BaseURL dvb:priority="2">http://moderate-cdn.example.com/</BaseURL>
<BaseURL dvb:priority="3">http://costly-cdn.example.com/</BaseURL>
```

This would mean that Players used cheap-cdn first, falling back to moderate-cdn if there is a problem with cheap-cdn. Costly-cdn would only be used if both cheap-cdn and moderate-cdn had problems.

11.9.3 Load balancing

Load balancing can often be carried out using either equipment in the network, or DNS load balancing. However such mechanisms may not be ideal for balancing between different hosting locations or CDNs, or where the load needs to be distributed in non-equal ratios.

A facility for load balancing is provided in clause 10.8.2, and a @weight attribute is added to the BaseURL element to allow the relative weighting of different BaseURLs to be signalled. The values of @weight are only significant relative to the values of other BaseURLs with the same @priority value. The BaseURL to use is chosen at random by the Player, with the weight of any given BaseURL being its @weight value divided by the sum of all @weight values.

As a simple example, consider a content provider uses two CDNs, cdn1 and cdn2 and wishes the load to be equally spread across them. This can be achieved by setting the @weight value on each to be the same, e.g.:

```
<BaseURL dvb:priority="1" dvb:weight="1">http://cdn1.example.com/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="1">http://cdn2.example.com/</BaseURL>
```

In a more complex example, consider a content provider using three CDNs, with one of them having a lower purchased capacity. The content provider needs to ensure that cdn3 only has 10 % of the load, with the remaining 90% being shared equally by cdn1 and cdn2. They could do this by using the following entries:

```
<BaseURL dvb:priority="1" dvb:weight="9">http://cdn1.example.com/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="9">http://cdn2.example.com/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="2">http://cdn3.example.com/</BaseURL>
```

11.9.4 Combining priority and load balancing

The failover options described in clause 11.9.2 and the load balancing described in clause 11.9.3 can be combined to provide a fully specified set of load balancing and failover rules. To assist with this the @serviceLocation attribute is used to indicate when BaseURLs share the same hosting platform and should not be used if another with the same value of @serviceLocation was found to be unusable.

When creating an MPD which combines these features it is important to note the rules which Players follow to select a BaseURL to use after the failure of the BaseURL they are using. In particular when a BaseURL fails, its @serviceLocation is added to a blacklist and no BaseURL using that @serviceLocation will be used again for the current session. Additionally, and perhaps at first unexpectedly, no BaseURLs with a @priority value the same as a BaseURL with a blacklisted @serviceLocation will be used again. This behaviour is to ensure that a BaseURL which has limited capacity is not unintentionally exposed to a larger load due to the failure of a different BaseURL which it was load balanced with. In the second example in clause 11.9.3, it would be inappropriate for the failure of cdn2 to lead to cdn3 having nearly twice as much load as it was intended to have.

Where the same BaseURLs are to be load balanced in normal use, but can also be used with or without load balancing then there need to be additional entries with the same @serviceLocation value but different @priority values. Typically this will lead to the lowest @priority value being used for a set of BaseURLs which are to be load balanced across as the primary location. Then there may be additional BaseURLs with increasing @priority values indicating BaseURLs which can be used as part of a smaller load balancing group or on their own. Finally there may be some BaseURLs with still higher values of @priority which are not normally used, but can be in the event of failure of all the normal ones.

As an example, consider the second example in clause 11.9.3, where there are three CDNs: cdn1, cdn2 and cdn3, and also the costly-cdn from clause 11.9.2.

For normal operation the content provider uses cdn1, cdn2, and cdn3 with @priority=1. For dealing with a failure of either cdn1 or cdn2 it provides three separate groups which load balance amongst the other two in the case of one failing. Finally it provides cdn1 and cdn2 separately in case both cdn3 and one of the other two fail. Then as a final fallback costly-cdn is provided.

```
<!--These are used first -->
<BaseURL dvb:priority="1" dvb:weight="9" serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="9" serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>
<BaseURL dvb:priority="1" dvb:weight="2" serviceLocation="cdn3">http://cdn3.example.com/</BaseURL>

<!--This would be used if cdn3 was in use and failed -->
<BaseURL dvb:priority="2" dvb:weight="1" serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>
<BaseURL dvb:priority="2" dvb:weight="1" serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>

<!--This would be used if cdn2 was in use and failed -->
<BaseURL dvb:priority="3" dvb:weight="9" serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>
<BaseURL dvb:priority="3" dvb:weight="1" serviceLocation="cdn3">http://cdn3.example.com/</BaseURL>

<!--This would be used if cdn1 was in use and failed -->
<BaseURL dvb:priority="4" dvb:weight="9" serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>
<BaseURL dvb:priority="4" dvb:weight="1" serviceLocation="cdn3">http://cdn3.example.com/</BaseURL>
```

```

<!--This would be used if cdn2 and cdn3 both failed -->
<BaseURL dvb:priority="5" dvb:weight="1" serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>

<!--This would be used if cdn1 and cdn3 both failed -->
<BaseURL dvb:priority="6" dvb:weight="1" serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>

<!--This would be used if cdn1, cdn2 and cdn3 all failed -->
<BaseURL dvb:priority="7" dvb:weight="1" serviceLocation="costly-cdn">http://costly-
cdn.example.com/</BaseURL>

```

Figure 1 shows the various routes through all of these options based on the percentage chance of each option being fixed and assuming that each option eventually fails. This is for illustrative purposes only - the failure of all BaseURLs should not be a common occurrence.

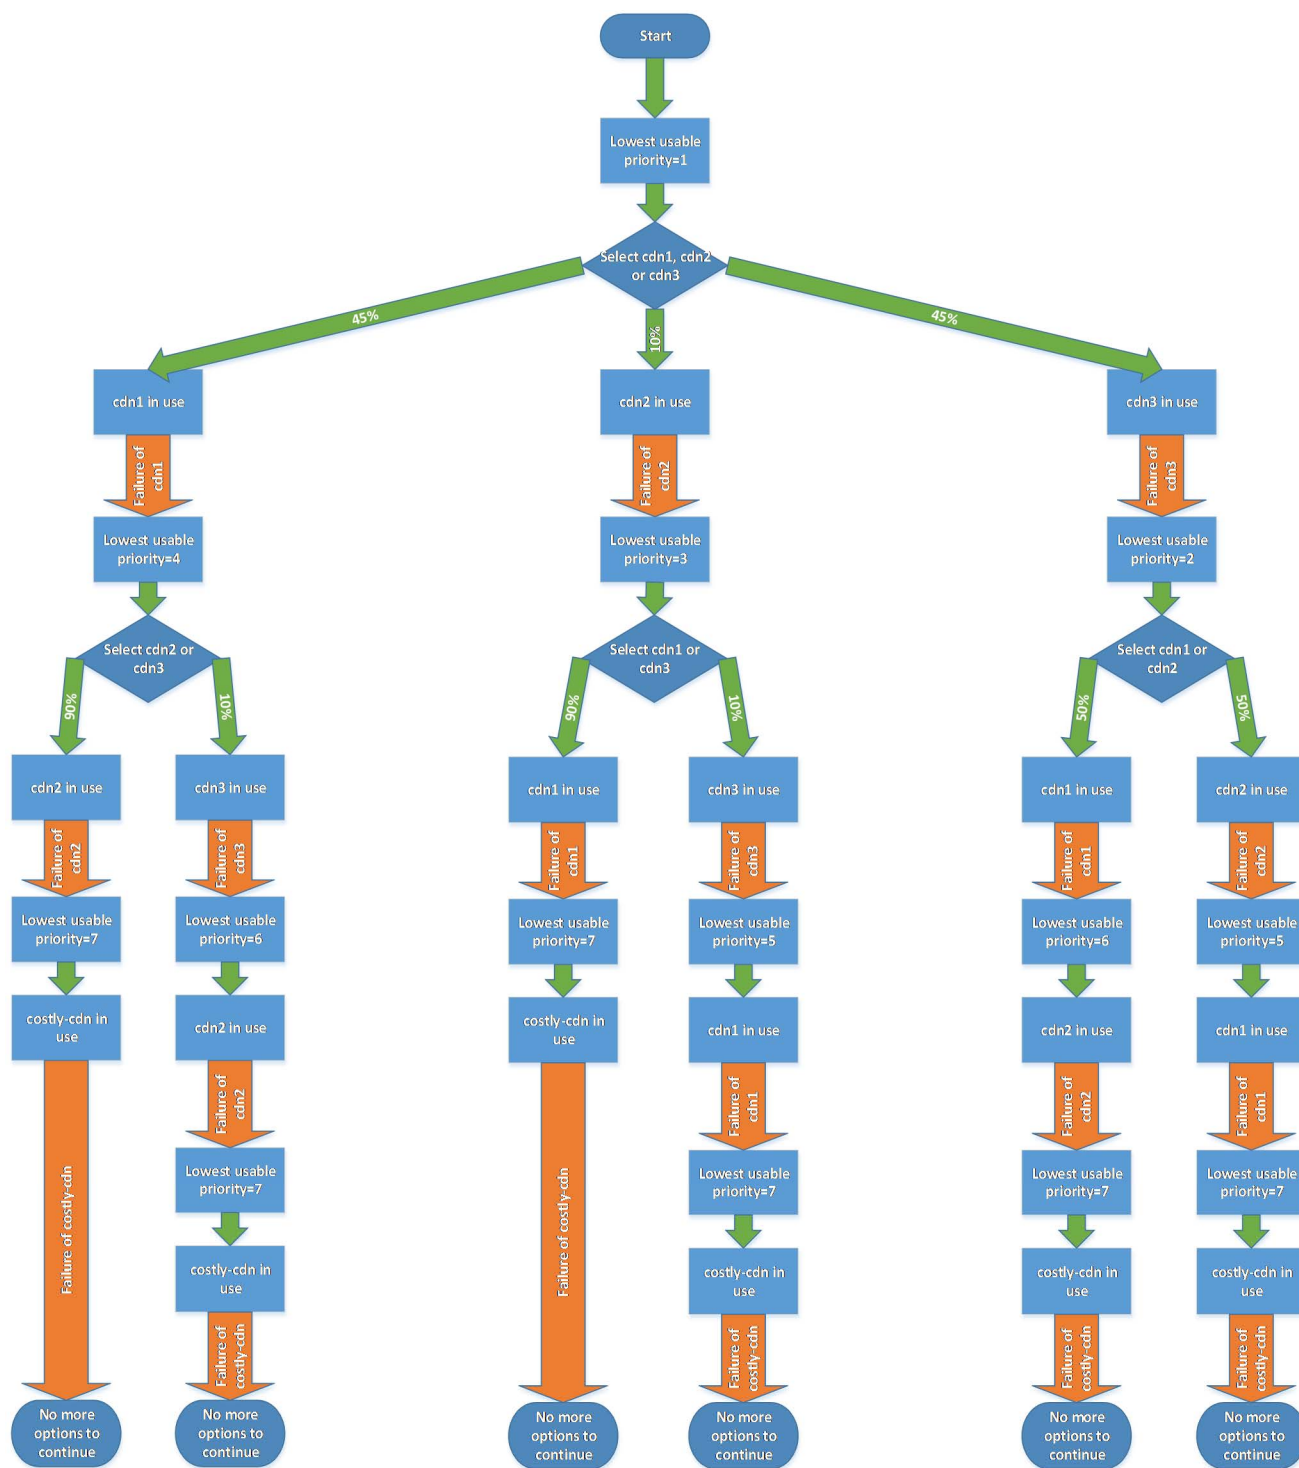


Figure 1: The various routes through the BaseURLs given in the example

11.9.5 Relative and Absolute BaseURLs

According to ISO/IEC23009-1 [1], BaseURLs can be carried at a number of levels within an MPD, and may contain both absolute and relative URLs, with relative URLs being resolved either against an absolute URL contained within a BaseURL at a higher level within an MPD or the base URI of the MPD itself.

In order to improve clarity it is recommended that:

- BaseURLs containing absolute URLs are placed at one level throughout an MPD, eg at the MPD level or the Period level.
- No level contains BaseURLs containing both absolute and relative URLs
- BaseURLs which will be ignored are not placed in the MPD at all (a base URL will be ignored if it is not needed in URL resolution).
- Multiple BaseURLs containing relative URLs are not placed within the same element (this would effectively indicate a choice about which to use without any guidance).

Where BaseURLs contain relative URLs, these should not include the @serviceLocation, @priority or @weight attributes, as these attributes will be ignored by the Player.

Ideally an MPD will contain a set of BaseURLs containing absolute URLs either at the MPD level, or if different periods are served from different locations at the Period level. Then single BaseURLs containing a relative URL, placed as necessary at the Period, Adaptation Set and Representation levels.

For example an MPD structure (showing only the BaseURLs and their containing elements) might look as follows:

```
<MPD>
  <BaseURL dvb:priority="1" dvb:weight="9"
serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>
  <BaseURL dvb:priority="1" dvb:weight="9"
serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>
  <BaseURL dvb:priority="1" dvb:weight="2"
serviceLocation="cdn3">http://cdn3.example.com/</BaseURL>
  <BaseURL dvb:priority="2" dvb:weight="1"
serviceLocation="cdn1">http://cdn1.example.com/</BaseURL>
  <BaseURL dvb:priority="3" dvb:weight="1"
serviceLocation="cdn2">http://cdn2.example.com/</BaseURL>
  <Period id="p1">
    <BaseURL>media/period1/</BaseURL>
    <AdaptationSet>
      <BaseURL>video/</BaseURL>
    </AdaptationSet>
    <AdaptationSet>
      <BaseURL>audio/</BaseURL>
    </AdaptationSet>
  </Period>
  <Period id="p2">
    <BaseURL>media/period2/</BaseURL>
    <AdaptationSet>
      <!-- In this adaptation set there is no BaseURL as the SegmentTemplate includes the path
within the @media and @initialization attributes -->
    </AdaptationSet>
    <AdaptationSet>
      <!-- In this adaptation set there is no BaseURL as the SegmentTemplate includes the path
within the @media and @initialization attributes -->
    </AdaptationSet>
  </Period>
</MPD>
```

11.10 Using Events

It is recommended that Content Providers do not use the MPEG-DASH MPD Patch event.

11.11 Asset Identifiers

AssetIdentifier descriptors identify the asset to which a period belongs and may be used for implementation of Player functionality that depends on distinguishing between adverts and main content. An **AssetIdentifier**, may be used, and if used has to be unique per asset within an MPD.

NOTE 1: This includes both the content and the advertising material.

The value of the @value remains undefined, however, it has to be the same for all parts of the same asset.

Periods with the same **AssetIdentifier** should have the same representations and same initialization segments.

NOTE 2: This may allow for a Player optimization that allows some sharing of initialization data across periods for the same asset.

Content authors should offer an MPD with *period-continuous* Adaptation Sets if the MPD contains Periods with identical Asset Identifiers.

11.12 Considerations for Running 24/7 services

11.12.1 A/V synchronization

In a long running live stream the alignment between audio and video streams cannot be achieved by starting both at 0 when an encoder starts as they may need to be restarted during the stream. Techniques similar to those needed to support multiple encoders may be needed, as described in the following clause.

11.12.2 Dual (or more) redundant encoding and hosting infrastructure

Content providers often wish to ensure their services are reliable and usually have backup systems in place to prevent single failures causing the service to go 'off-air'. When doing this with DASH services, redundant systems can be used to provide content through different Internet end points. This has the benefit that single points of failure can be further removed by using the Player behaviour described in clause 10.8 to manage the Player switching between the systems.

For this approach to work without interruption, there needs to be a system in place to ensure the different paths are time aligned, such that media segments produced through one half of a redundant system can be used interchangeably with segments from the other half. To enable this, the following need to be the same on the output of each path:

- segment names
- track fragment decode times
- fragment sequence numbers

Encoders and packagers may need to synchronize to a timecode in the media stream to allow them to set these values deterministically, and to have rules to decide where segment boundaries go, to ensure that segment boundaries are aligned, even after one of the paths is restarted after a failure or maintenance. In particular, care needs to be taken with regards to the audio access unit timing relative to the video when switching from one redundant audio encoder to another.

11.12.3 Encryption

When encrypting long running services, consideration needs to be given to whether key rotation is required. The following may affect this:

- Amount of traffic to be encrypted with one key
- Avoiding the reuse of IVs with the same key
- Allowing entitlement changes over time

11.12.4 Further information

Further information on running live MPEG-DASH services can be found in clause 4 of DASH-IF: Guidelines for Implementation: DASH-IF Interoperability Points [i.1].

11.13 Cache Control

Be aware of HTTP caching directives when serving dynamic MPDs, in order to ensure that the Player can acquire appropriate updated versions when required to do so by the MPD@minimumUpdatePeriod attribute or MPD validity expiration events.

History

Document history		
V1.1.1	May 2015	Publication