



DRAFT GUIDELINES

Version 0.0draft010 2017-09-12

VR Industry Forum

5177 Brandin Court
Fremont, CA 94538, USA

www.vr-if.org

THIS DRAFT DELIVERABLE IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NON-INFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS DRAFT DELIVERABLE SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER VRIF, NOR ANY OF ITS MEMBERS OR CONTRIBUTORS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS DRAFT DELIVERABLE.

Disclaimer

The VR Industry Forum accepts no liability whatsoever for any use of this document.

Copyright Notification

No part may be reproduced except as authorized by written permission. Any form of reproduction and/or distribution of these works is prohibited.

Copyright © 2017 VR Industry Forum. All rights reserved

DRAFT FOR PUBLIC PREVIEW

About this Guidelines Preview

These Draft Guidelines are provided by the VR Industry Forum on 13 September 2017 for consideration, comment and contribution through public review. These guidelines should not be considered as complete, nor be used as any reference for implementations. It is anticipated that considerable updates will be made before the final version which is expected to be published in January 2018.

The following are some of the areas of ongoing work:

- The role of HDR in content presentation – what works and what does not
- Recommendations for stitching video content in the production domain
- Master format adaptations to permit distribution over constrained links
- Passing the desired viewport to the secure trust zone
- Content encryption and watermarking on viewport dependent media profile

Providing feedback

The VR Industry Forum welcomes feedback from the public and industry on this draft. Comments can be made either:

- On the VRIF website – <http://www.vr-if.org/guidelines/>, please provide contact information if you would like us to follow up on your comments
- On the VRIF GitHub Issue Tracker - <https://github.com/VR-IF/Guidelines/issues>

It is preferable for comments to be received by 31 October 2017 in order for us to fully evaluate and respond to them.

1 Introduction

The guidelines presented in this document cover all aspects of the distribution ecosystems, including compression, storage and delivery, in order to ensure high quality, comfortable consumer VR experiences. These guidelines are aimed at addressing best practices for VR content distribution as well as advocating interoperability and deployment guidelines based on common technical standards for VR content distribution, including promoting the use of common profiles across the industry.

The targeted audience includes content distributors, service providers, broadcasters, mobile operators, consumer electronics manufacturers, professional equipment manufacturers, software developers and technology companies that aim to enable deployment of VR content distribution services.

The scope of the guidelines presented herein includes:

- **Production:** Technical aspects of the media formats used in the interface between the content provider and the service provider along with human factors considerations for compelling and usable 360° video experiences.
- **Compression:** Media codecs for VR, i.e. encoding of different production formats and related media profiles for video, audio and possibly also other media types such as text, graphics, etc.. This includes decoding and rendering of the media based on an abstracted distribution data model.
- **Storage:** Media formats for VR content (e.g. file/segment encapsulation) for different distribution means, including but not limited to storage, download, adaptive bitrate streaming and broadcasting
- **Delivery:** Interfaces and protocols for Live, Linear and VOD delivery over streaming (unicast), and broadcast applications
- **Security:** VR specific threat identification and mitigation techniques as well as methods for implementing security and privacy protection functions.

Contents

1	Introduction	2
2	References.....	7
3	Conventions and Terminology.....	9
3.1	Definitions	9
3.2	Abbreviations.....	9
4	VR Technologies	11
4.1	Production Guidelines for VR Audio and Video Content.....	11
4.1.1	User Experience	11
4.1.2	Considerations for acquisition and creation	11
4.1.3	Master Format	19
4.2	Media and Presentation Profiles	21
4.2.1	Introduction.....	21
4.2.2	Selected Media Profiles	21
4.3	Content Security.....	29
4.3.1	Scope.....	30
4.3.2	MovieLabs ECP 1.1 Deltas	31
4.3.3	DRM System Specifications	32
4.3.4	Platform Specifications	32
4.3.5	End-to-End System Specifications.....	33
4.3.6	Encrypted Media Extensions	33
5	Vertical 1: OTT Download or Streaming of VR360 Content	34
5.1	Description of Vertical.....	34
5.2	Guiding Example Use Cases.....	34
5.3	Reference Architectures	35
5.3.1	Distribution Architecture	35
5.3.2	Client Architecture	37
5.4	Technical Enablers.....	37
5.4.1	Suitable Media Profiles.....	37
5.4.2	Suitable Presentation Profiles	37
5.4.3	Distribution Systems.....	38
5.5	Guidelines for Service Providers.....	43
5.5.1	Suitable Production Formats.....	43
5.5.2	Sphere-to-Texture Mapping and SEI Message Generation.....	44
5.5.3	Encoding and Content Preparation.....	48
5.5.4	Distribution.....	54
5.5.5	Security.....	56
5.6	Guidelines for Service Platform Developers	56
5.6.1	Overview	56
5.6.2	Rendering Process based on SEI messages.....	58
5.6.3	Distribution and Delivery.....	59
5.6.4	Decoding and Rendering	59
5.6.5	APIs.....	59
5.6.6	Security.....	59
5.7	Guidelines for App Developers	59
5.7.1	Distribution and delivery	60

5.7.2	Decoding and Rendering	60
5.7.3	APIs.....	60
5.7.4	Usage of Service Platform	63
Annex A Video Master Format Metadata.....		64
Annex B ISO BMFF Extractors (informative).....		67

Figures

Figure 1: Example of Partial Panorama metadata.....	19
Figure 2: Receiver Model for OMAF HEVC Viewport-Independent baseline media profile	23
Figure 3: OMAF-DASH Streaming Client model with interfaces	26
Figure 4: OMAF-Download Client model with interfaces	27
Figure 5: Example Architecture for simple VR Streaming Services.....	36
Figure 6: Example Architecture for encrypted VR Streaming Services.....	37
Figure 7: Example VR distribution system.....	39
Figure 8: Example protocol stack for VR content distribution	39
Figure 9: Exemplary DASH configuration setup	40
Figure 10: Latency sources contributing to M2HR latency [CDNOPT].....	42
Figure 11: Spherical surface coordinates ϕ , θ with yaw, pitch, and roll of the region of a sphere covered by the cropped output picture relative to the equator and 0 meridian	46
Figure 12: Content Preparation for DASH Distribution.....	50
Figure 13: Video content preparation for DASH distribution with OMAF HEVC Viewport-Dependent baseline profile.....	51
Figure 14: Logical Receiver Model.....	57
Figure 15: Rendering and viewport generation.....	58
Figure 17: OMAF-DASH Client model with interfaces.....	61
Figure 18: DASH Access Engine for OMAF HEVC Viewport-Dependent baseline profile	61
Figure 19: DASH Media Engine for OMAF HEVC Viewport-Dependent baseline profile	62
Figure 20: Inverse Projection/Renderer.....	62
Figure 21: Single ISO BMFF File with one extractor track, N extractors and N MCTS tracks after subsegment concatenation	68

Tables

Table 1: Master File formats	20
Table 2: Maximum achievable resolution in the viewport using viewport-independent baseline media profile for HEVC Main10 Profile, Main Tier, Level 5.1, 60fps, for a display with a FoV of 90°×90° and a given content coverage.....	22
Table 3: Overview of OMAF media profiles for audio.....	28
Table 4: MPEG-H Audio MIME parameter according to RFC 6381 and ISO/IEC 23008-3	29
Table 5: VR Characteristics.....	30
Table 6: Mapping of SEI Message Information to OMAF Metadata.....	49
Table 7: Master Format metadata.....	64

2 References

- [ECP] MovieLabs Specification for Enhanced Content Protection ([ECP V1.1](#))
- [EME] Encrypted Media Extensions ([W3C Proposed Recommendation](#))
- [CENC] Common Encryption in ISO Base Media Format Files ([ISO/IEC 23001-7:2016](#))
- [FRATES] [Les casques de réalité virtuelle et de jeux video](#)
- [ACN] Ambisonic data exchange formats – Ambisonic Channel Number ([Wikipedia](#))
- [ADM] Audio Definition Model ([ITU-R BS.2076](#)), [BBC Software](#)
- [NORMS] Ambisonic data exchange formats – Normalization ([Wikipedia](#))
- [TR26918] Virtual Reality (VR) media services over 3GPP ([3GPP TR26.918](#))
- [3DA] Information technology – High efficiency coding and media delivery in heterogeneous environments – Part 3: 3D Audio, Second Edition ([ISO/IEC 23008-3](#))
- [N16584] MPEG-H 3D Audio Verification Test Report ([MPEG N16584](#))
- [DRC] Information technology – MPEG audio technologies – Part 4: Dynamic Range Control ([ISO/IEC 23003-4](#))
- [DASH] Information technology – Dynamic adaptive streaming over HTTP (DASH) – Part 1: Media presentation description and segment formats, including ISO/IEC 23009-1:2014/Amd.4:2016 ([ISO/IEC 23009-1](#))
- [DASHPUSH] Information technology – Dynamic adaptive streaming over HTTP (DASH) – Part 6: DASH with server push and WebSockets ([ISO/IEC 23009-6](#))
- [ISOBMFF] Information technology – Coding of audio-visual objects – Part 12: ISO Base Media File Format ([ISO/IEC 14496-12](#))
- [ISOCMAF] Information technology – Coding of audio-visual objects Information technology – Multimedia application format (MPEG-A) – Part 19: Common Media Application Format ([ISO/IEC 23000-19](#))
- [MP4FF] Information technology – Coding of audio-visual objects – Part 14: The MP4 File Format ([ISO/IEC 14496-14](#))
- [MP4SYS] Information technology – Coding of audio-visual objects – Part 1: Systems ([ISO/IEC 14496-1](#))
- [VDVS-QUAL] Subjective Quality Results of Viewport-Dependent Omnidirectional Video Streaming ([3GPP S4-170637](#))
- [ADDSEI] Study text of ISO/IEC 23008-2:201x/DAM3 Additional supplemental enhancement information ([MPEG W17045](#))
- [UPDSEI] HEVC Additional Supplemental Enhancement Information ([ICTVC AB1005](#))

Editor's Note: [ADDSEI] and [UPDSEI] need to be rationalized

[OMAFSODIS] Study of ISO/IEC DIS 23000-20 (→ ISO/IEC 23090-2) Omnidirectional Media Format ([MPEG W16950](#))

[RFC6381] The 'Codecs' and 'Profiles' Parameters for "Bucket" Media Types ([IETF RFC 6381](#))

- [NAL] Information technology – Coding of audio-visual objects – Part 15: Carriage of network abstraction layer (NAL) unit structured video in the ISO base media file format (ISO/IEC 14496-15:2017)
- [BS2088] Long-form file format for the international exchange of audio programme materials with metadata ([ITU-R BS.2088](#))
- [BT709] Parameter values for the HDTV standards for production and international programme exchange ([ITU-R BT.709](#))
- [BT2020] Parameter values for ultra-high definition television systems for production and international programme exchange ([ITU-R BT.2020](#))
- [BT2100] Image parameter values for high dynamic range television for use in production and international programme exchange ([ITU-R BT.2100](#))
- [R103] Video Signal Tolerance in Digital Television Systems ([EBU R103](#))
- [Pulkki] V. Pulkki, “Virtual sound source positioning using vector base amplitude panning” Audio Engineering Society Journal, [vol. 45, no. 6, pp. 456–466](#), 1997.
- [CDNOPT] TiledMedia, CDN Optimization for VR Streaming
- [TS26234] Packet Switched Streaming Service ([3GPP TS 26.234](#))
- [80211AC] IEEE Standard for Information technology--Telecommunications and information exchange between systems—Local and metropolitan area networks--Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications—Amendment 4: Enhancements for Very High Throughput for Operation in Bands below 6 GHz. ([IEEE 802.11ac](#))
- [80211AX] High-Efficiency Wireless (IEEE 802.11ax)
- [DASHIFIOP] DASH-IF IOP Guidelines v4.1 ([DASH-IF IOP-v4.1](#))
- [TS1015472] Plano-stereoscopic 3DTV; Part 2: Frame Compatible Plano-stereoscopic 3DTV ([ETSI TS 101 572-2](#))
- [PROJCONV] Y. Ye, E. Alshina, J. Boyce (editors) Algorithm descriptions of projection format conversion and video quality metrics in 360Lib Version 4, (JVET G1003, [MPEG M41359](#))

3 Conventions and Terminology

3.1 Definitions

Additional terms are defined in the VRIF Lexicon, available at <http://www.vr-if.org/guidelines/>

Term	Definition
360 Content	Post produced or Live 360 Video content with limited interactivity
Cardboard	Google project to display 360 in cheap viewer using Gyro, accelerometer, compass of smart phones
Equirectangular	Projection of a spherical image in a rectilinear image frame
Haptic Feedback	Palpable feedback, usually servo-driven in human interfaces (hand controllers)
Hot Spot	Virtual object embedded in 360 content enabling additional functionality
Lat/Long	Projection of a spherical image in a rectilinear image frame
Nadir	Lowest point of a sphere – often used to refer to dead-spot in image below camera array
Nodal point	The point at which the optical centers of the camera array are cantered
Simulator Sickness	Uneasiness or disorientation when wearing a Head-Mounted Display
Spatial Audio	Audio which is perceived to have orientation in three possible axes: L/R, Up/Down, Back/Front
Spherical Video	Video content captured or displayed simultaneously in all directions
Stitch Lines	Visible junctures of blended or overlapped images from multiple camera arrays
Unity	(Game Engine) A cross-platform engine for game and VR design and implementation
Zenith	Highest point of a sphere – often used to refer to dead-spot in image above camera array

3.2 Abbreviations

Definitions of most of these are available in the VRIF Lexicon, <http://www.vr-if.org/guidelines/>

Abbreviation	Definition
DAW	Digital Audio Workstation
DOF	Degrees of Freedom
FOV	Field Of View

HDCP	High-bandwidth Digital Content Protection
HMD	Head Mounted Display
HRTF	Head Related Transfer Function
IPD	Interpupillary distance
PoV	Point of View
VR	Virtual Reality, historically referred to CGI -based 360 immersive experiences, now 360 video too

DRAFT FOR PUBLIC PREVIEW

4 VR Technologies

4.1 Production Guidelines for VR Audio and Video Content

4.1.1 User Experience

360 content created for display in a Head-mounted-Display (HMD) should allow the viewer to engage in a unique and immersive manner. The viewer may be separated from visual cues from the real environment and placed in a virtual environment or experience.

The uniqueness of the 360 video experiences is mostly due to two factors:

- The level of immersion induced by the wider field of view of the HMD
- The interactivity of the scene that respond to the user actions (at the lowest level, such interactivity is reduced to the ability to visually explore (look around) the scene like in the case of 3DOF)

These two elements imply some particular precautions to ensure that the content created should not lead to disorientation and unease.

From an audio perspective, content should be produced with an understanding that visual content is not primarily frontal as is assumed for linear TV consumption. Further, since audio is consumed over headphones, acoustic cues can be created from all directions including below ear level – something not usually done for loudspeaker consumption (since most loudspeaker configurations don't include loudspeakers below ear level).

4.1.2 Considerations for acquisition and creation

4.1.2.1 Field of view consideration for video

There is a significant difference between the human Field Of View (FOV) and the one covered by current commercial VR devices. The human horizontal FOV is estimated to be 190° (without eye rotation) and 220° (with eye rotation) while the vertical FOV is 120° (50° upwards and 70° downwards). The Head-mounted Displays (HMD) vary in the FOV that they can reproduce, but generally have a 60° to 90° per eye FoV range (for instance HTC Vive, Oculus Rift and Sony PlayStation VR claim to have a 110° total field of view, while smartphone based solution usually reach 95°). While these values are still far from the human FOV, they are notably wider than those of traditional devices like TV screens that, when viewed at recommended distances, cover approximately 40° of FOV.

The difference in the wideness of the FOV between HMD and traditional displays has various consequences for the production of the media. In particular a wider field of view implies a bigger amount of “vection” (visual peripheral information able to induce the illusion of self-motion) that could induce the discomfort symptoms characteristic of the Virtual Reality sickness.

Virtual Reality sickness is induced by the discrepancy between motion cues coming from the vestibular and the visual systems. It is important to notice that the vestibular system is able to give information concerning the orientation and acceleration (of the head, and consequently of the user) but is unable to distinguish between staying still and moving straight at a constant speed.

4.1.2.2 Content Position

While the human field of view is 190°, the perceptual capacities of the visual system are not the same over the whole visual field. The ability to recognize colors for instance is limited to a central area of about 120°. Shape is recognizable around 60° from the center and texts are readable when they are at less than 20°

from the center of the eye (fovea). As a consequence, humans tend to keep the most relevant visual content in the central area of their FOV, by constantly moving their eye gaze and head if necessary. This motion could be uncomfortable in particular when wearing an HMD. So, during production attention has to be taken to reduce the need of the head movement by keeping relevant content in the line of sight.

For example, content designed to be viewed in a seated position should be created such that the action in the scene requires that the viewer only moves their head through an arc of 90° to 120°. This will give an overall viewable arc of 180° to 210°, with significant action contained within a 90° range of head movement.

For content designed to be viewed in a standing position a wider arc may be considered.

4.1.2.3 Camera Motion

As previously stated, “vection” can induce VR sickness in the user, and a main source of “vection” is the camera motion.

Camera motion that is the consequence of user movement (rotations in 3DOF and rotation and translation in 6DOF) could be VR sickness free if the system is reactive (less than 20 milliseconds of Motion-to-Photon delay) and accurate (1:1 scale) because the vection is coherent with the vestibular information.

In all other cases camera motion could potentially induce VR sickness in particular when the vection suggest orientations and accelerations that are not perceived by the vestibular system.

To reduce the risk of VR sickness the main approach consists of reducing the amount of “vestibular incoherent” vection through different technique like:

- Avoid unnecessary change in speed and direction when performing camera motion, constant speed translations on a straight line are less disturbing.
- Avoid rotation of the camera that could influence the “horizontality of the horizon.”
- Whenever possible replace motion by adopting other metaphors (teleportation).
- The temporary reduction of the field of view sensibly reduce the strength of vection, peripheral vision is highly sensitive to motion cues¹.
- Avoid moving to close to objects and environmental elements that could induce high level of vection (walking close to a wall or very close to the ground in particular when they have repeated patterns)
- Avoid repeated patterns and sharp lines like stripes, blocks, stairs, trees, poles. Moving close to such patterns create a considerable amount of motion cues.
- Limit flickering and other artefacts that could generate sharp, high contrast, visual patterns that (like stripes) create a considerable amount of motion cues. For the same reason blurred images (and artificial blurring algorithms) could be used to reduce vection.
- The addition of static frames or elements around the view of the moving camera (like a view from inside a vehicle or a fake nose) could partially reduce VR sickness probably due to a combination of perceptual effect (reduction of the vection due to static elements) and cognitive activity (identification of an familiar situation that could help to interpret the incongruous motion cues)
- If the viewer is stationary and the scene moves independently from the viewer there is a perceived disconnect between the viewer's visual system and motion sensing system. This can lead to disorientation and in some cases nausea.

¹ Method implemented in the VR-game eagle fight from Ubisoft

- If motion is to be used as a creative intention then care should be taken. The following techniques can be useful:
 - Placing the viewer in a recognizable situation i.e. inside a vehicle or a flying bird's eye view.
 - If the motion is at ground level, avoid:
 - Accelerating too rapidly
 - Turns that are not relevant to the action
 - Yaw and pitch of the horizon
 - Vertical motion – low frequency - as in replication of human walking PoV
 - Where the action follows a relevant person or object, this can reduce motion effects.

4.1.2.4 Image capture rate and motion capture fidelity

- Content should be shot or created in a manner that reduces motion blur or motion stepping to comfortable levels as flickering images can lead to rapid fatigue or disorientation.
- Image capture at higher frame rates and designing content to work within the target display capabilities is recommended to reduce discomfort.
- When capturing moderately fast-paced action such as dance, sport and vehicles a minimum capture frame rate of 50/60fps should be considered.
- If shooting at higher frame rates, consideration of the target display(s) characteristics should be made.
- Frame rate metadata needs to be included with content.

4.1.2.5 Orientation

For 360 content, it is best practice to keep the horizon level at all times unless angled for a specific purpose. The viewer will always tend to match their head inclination to the perceived horizon – this can lead to discomfort and vertigo.

If the horizon oscillates this can prove to be very uncomfortable.

Panning and tilting the camera can also be very uncomfortable.

When transitioning between shots, care should be taken to orientate the action or intended direction of view of the subsequent shot to the outgoing shot.

4.1.2.6 Perceived Viewer eye height

Camera to scene height should always be relevant to the scene.

A small discrepancy in height can result in surprisingly off-putting perceptions of the relative sizes of subjects.

Ideally the camera height should be coherent with the final user position (seated vs standing up). For 3DOF systems, this could imply the user to select his viewing position (seated vs standing up)².

The wide-angle distortion of some camera formats can exacerbate this, a low camera height can make subjects appear unusually tall – similarly a high camera height can make subjects appear unusually small.

² 6DOF systems (like the HTC vive) can adapt (at least in part) the video to be coherent with the user head height.

4.1.2.7 Proximity of objects in the scene

Objects close to camera have a disproportionate influence on the viewer and can unnecessarily dominate the scene. This effect can impact reaction in attempt to move away from the object.

Wide-angle perspective distortion of some camera formats may make objects approaching the camera appear to accelerate towards the viewer and so should be avoided.

If only 3DOF is supported, nearby objects create occlusion and the apparent loss of 3D effect due to inability to look around the object may break the illusion of immersion. In other words it makes it very clear to the user that there is no 6DOF and this might be disconcerting.

4.1.2.8 Duration of Content

Results from a survey (of 300 respondents) indicate that:

- Shorter content durations are more comfortable to the user
- Up to 20 minute durations are seen as acceptable if the content is not challenging.
- Content with significant action or motion should be shorter.

4.1.2.9 3D Stereo Content

The impression of depth obtained using stereoscopic contents is the consequence of the binocular disparity (the difference in the retinal projection of the 2 eyes) induced by the human inter pupillary distance IPD (on average 6 cm for the adults). The further the objects are from the user the smaller the difference of the retinal projection. In the case of objects that are more than 20 meters from the user, the binocular disparity plays a minor role in the perception of depth.

The vision of stereoscopic content in current HMDs³ induces a visual fatigue due to a phenomenon called the “vergence/accommodation conflict”. In fact, all the content presented in the HMD are at the same focal plane that is dictated by the lens (usually at 1.5 meters) while the binocular disparity imposes a vergence distance that varies as a function of the “stereoscopic distance (i.e. the horizontal disparity between the 2 retinal projections of a target)” of the object to keep in the line of sight (in the foveal area). The larger the difference between the focal and the vergence distance, the greater will be the fatigue.

The creation of stereoscopic content implies some specific Human Factor considerations:

- The human perception of depth evolves between the age of 6 and 12. For this reason stereoscopic content is forbidden to users younger than 6 years and is not recommended for users that are less than 13 years old.
- Differences in the luminosity, sharpness or colors of two images presented in a stereoscopic way will make it harder for the two images to be perceptually fused, with the most vivid image being predominant. Any lack of fusion may cause additional fatigue for the viewer.
- Stereoscopic depth information should be coherent with the other depth information. In particular, it is important to avoid that objects that appear in front of the “screen” (negative parallax) will move out from the HMD field of view. This problem known in literature as the “frame problem” induces visual discomfort due to the conflict between the stereoscopic depth cues that suggests that the object is close and the monocular (occlusion) information that suggests that the object is far.

³ Lightfield displays are supposed to limit this issue.

One of the most powerful cues which enables the perception of immersion is that of stereo imagery in the experience.

- The experience should have carefully planned and implemented stereo depth regime.
- Care should be taken to avoid too much inappropriate stereo disparity (uncomfortable stereo depth) as this will cause eye-strain and visual discomfort.
- Near-field objects should not be positioned or manipulated so as to cause adverse autonomous defensive reactions in the consumer.

4.1.2.10 Subtitles

Content can be produced with (non-positional) subtitles. These will be rendered by the HMDs with formats and positioned as selected by the end-user.

4.1.2.11 Video Live/Post Production

- Production should take into consideration the aspects described in sections 4.1.2.1 to 4.1.2.10
- Cuts should be consistent with the action and story.
 - Consideration of the viewer motion
 - Use of production to direct the viewpoint
 - Consistent story telling
 - Audio alignment with image content
 - Spatial alignment and temporal synchronization.

4.1.2.12 Audio for VR - overview

One of the essential novelties of VR is that a user is free to change their viewing gaze at will, allowing individual immersive experiences in any viewing direction at any given moment. Consequently, the methods of audio creation, transmission and reproduction applied for VR must be able to accompany the dynamically changing visual perspective.

This means that audio must be reproduced equally well in all directions, allowing the presentation of sounds from below or above the viewer with the same spatial accuracy as sounds from the front.

Further, for a realistic listening experience the audio presentation method must seamlessly adapt the spatial audio processing and recreate the sound scene coherently with respect to the dynamically changing viewer's gaze. Spatial audio is essential for compelling VR experiences. It can be used as a tool to immerse and guide the viewer in the VR scene.

Even when the video content has a limited FOV, (e.g. 180° or 270° content), full spherical audio is a requirement – since an artificial ‘silence’ from the back, in these cases, will result in the loss of immersiveness and/or the lack of suspension of reality.

4.1.2.13 Audio Formats

Audio for VR can be produced using three different formats. These are broadly known as channels-, objects- and scene-based audio formats. Audio for VR can use any one of these formats or a hybrid of these (where all three formats are used to represent the spherical soundfield).

Channels

Loudspeaker-based audio reproduction such as stereo 2.0 or surround 5.1 has been the de-facto standard for production and audio delivery to consumers for decades. To ensure the intended sound reproduction,

channel-based audio requires the same standardized loudspeaker placement at the production facility and the listener's reproduction location. Standardized loudspeaker configurations include simple mono and stereo, horizontal-only (5.1) to immersive 7.1+4H and 22.2.

For a faithful reproduction, the entire immersive virtual loudspeaker configuration must be virtually displaced by updating the HRTFs associated to the new direction of each speaker. This requires an accurate set of HRTFs available for all possible virtual loudspeaker positions as well as careful signal processing when updating the HRTFs in real time.

Objects

For Object-based audio the sound scene is composed of multiple individual sound sources (objects) along with metadata that describes its spatial characteristics (position, width, radiation pattern, room-reflection properties, etc). During playback, the audio scene is constructed or rendered using all the audio sources and the associated metadata. The format is thus agnostic of loudspeaker positions. For loudspeaker playback, the renderer considers the number and position of loudspeakers – and in the case of headphone playback such as VR – renders to headphones.

The use of object-based audio is quite effective for post-produced audio. It involves the use of DAWs to create the audio scene in which some objects might get grouped into a summary element or stem while others such as Dialog can be carried as a discrete object.

The capturing or live acquisition of audio using objects can be a challenging task because all sound sources must be individually captured and tagged with the correct location metadata. In post-production, this is typically achieved by panning objects while looking at the video. Leakage across the captured audio objects (e.g., due to proximity or reverberation) should be avoided.

A pure object-based representation can require a multitude of individual audio tracks and their associated time-varying metadata. This means that the bandwidth necessary for transmitting a sound scene depends on the number of simultaneous objects present at any point in time also known as the scene complexity. Typical cinematic content can involve the use of several simultaneous objects. For home delivery through streaming or broadcast services, objects typically require grouping to reduce scene complexity and transmission bandwidth requirements. In VR productions, similar methods are applicable

Objects can be either individually binauralized using one discrete HRTF convolution process per object, or rendered (e.g., using VBAP [Pulkki]) to a set of virtual loudspeakers which is then binauralized using one HRTF convolution process per virtual loudspeaker.

Scene-based Audio

Scene-based audio (SBA) represents the acoustic pressure field as a function of space and time using a set of coefficients that are the linear weights (or coefficients) of orthogonal spatial basis functions known as Spherical Harmonics. This is also known as Higher Order Ambisonics (HOA). Like Object-based audio, the soundfield representation is agnostic to loudspeakers. For loudspeaker playback, the renderer takes into account the number and position of loudspeakers.

First Order Ambisonics (FOA) (also commonly referred to as B format) is a basic form of scene-based audio in which the soundfield is described by only the lowest four spherical harmonic coefficients. Higher order Ambisonics (HOA) provides a concise and accurate representation of the entire sound field and overcomes the spatial resolution constraints of FOA.

SBA provides an efficient and accurate representation of the sound field with a limited number of coefficients, dependent only on the order of the representation that is chosen. Moreover, the SBA coefficients can be compressed to a fixed bandwidth irrespective of the complexity of the scene.

The spherical harmonic underpinnings of SBA allow for efficient and smooth rotation of the sound field – allowing for compelling VR experiences.

There are algorithms for efficient binauralization of HOA coefficients that are also independent of the complexity of the scene and the number of virtual speakers used in rendering. These computational advantages are invaluable in enabling head-tracked binauralization for VR on consumer devices.

Capturing/acquiring Ambisonics soundfields can be achieved by using compact microphone arrays.

A single scene-based audio representation can also carry any number of discrete objects without requiring additional metadata or increased bandwidth. Under this approach, the audio scene is constructed or augmented by panning pre-recorded or live captured audio objects into the 3D space. For live object panning, location tracking is required.

Hybrid formats

Hybrid scenes that involve the use multiple formats can be used for VR audio. Example scenarios include the combination of audio channel-beds and objects as well as the use of Ambisonics audio along with discrete audio objects (such as commentaries) and relevant metadata.

4.1.2.14 Audio Recording and Production

- Spatial audio from live scenes can be recorded using a range of microphones distributed in and around the recording scene.
- There are a number of microphone arrays that are available off the shelf that allow the capture of the spherical soundfield in HOA format. For 3DOF content, the positions of these microphones are considered to be relative to the camera position – which is generally considered to be at the origin.
- For 3DOF content, soundfield captured through microphones should be rotated in accordance with camera rotation.
- If the camera is to be translated (not just rotated), then attaching soundfield microphones to the camera is recommended.
- Microphones should be positioned with the following considerations:
 - Sound captured through stationary contact microphone modules provide a correct sampling of the audio at the position of the microphone. As such, it is relatively simple to ‘pan’ that recorded audio into 3D space by considering the position of the stationary microphone. This can be done through DAW software.
 - Sound captured through shotgun or parabolic microphones are not necessarily representative of the audio at the position of the microphones. To ‘pan’ that recorded audio into 3D space requires knowledge of the position of the audio object that the microphones are being directed at. If these positions can be estimated (either manually or automatically), the panning into 3D space can be achieved. The same is true for lapel microphones placed on a moving acoustic source. It is essential to ensure that a high degree spatial accuracy is maintained – such that the positional cues from the visual and acoustic scene are not contradictory.
- High spatial resolution can be achieved using audio channels, objects with metadata or Ambisonics audio or a combination thereof.
- Higher Order Ambisonics signals should be recorded and produced in at least full-spherical 4th order HOA. The higher the Ambisonics order, the higher the spatial resolution allowing the localization and tracking of acoustic features – a necessary condition for immersive experiences.

- Recordings from microphone-arrays can be combined with individually recorded and/or produced stems either live or during post-production.
- For a live and/or post production, the audio goes through a console or Digital Audio Workstation (DAW) that further processes the spatial characteristics of the sound field. This is on top of the basic audio production workflow. Example DAW software that allows the live capture as well as post-production of spatial audio include:
 - [DAW1] <https://developer.qualcomm.com/software/3d-audio-tools>.
- Consideration during production should account for the fact that spatial audio will be reproduced from all directions, including the lower hemisphere.
- The sampling-rate during recording should be at least 48 kHz.
- Diegetic/non-Diegetic Audio
 - Audio elements having a visual counterpart in the scene (Diegetic audio elements) must be spatially aligned and temporally synchronized with the video during head-motion.
 - Audio elements that do not have a visual counterpart in the scene (Non-Diegetic) may be produced so that their position is not compensated with the user's head-motion.
 - NOTE: An audio element may alternate over time between having a visual counterpart and not having a visual counterpart. An example is a narrator that is initially not present in the visual scene and then becomes part of the visual scene. Compensation for the user's head motion should only happen when the audio element has a visual counterpart.
- The audio export of the production should be loudness-normalized for consistent loudness across different content.
- Metadata for loudness of audio content should follow the appropriate regional recommendations for broadcast content delivery and exchange. Some examples of the recommendations include EBU R128 and the US CALM act.
- Presentation/Rendering for mixing and monitoring
 - Consideration during production should account for the fact that audio will be primarily experienced with headphones or in-ear buds.
 - Mixing using headphones should support low latency soundfield rotation with head tracking. A motion to sound latency lower than 30ms is recommended.
 - When mixing over headphones, an optimal binaural audio experience is achieved when the headphone feeds are created using personalized Head Related Transfer Functions (HRTFs) and headphone equalization. It is recommended that HRTFs closely matching those of the mixing engineer be used.
 - If the headphone feeds are created by an intermediate virtual loudspeaker renderer (before HRTF processing, for example), it is recommended that the renderer be indicated using appropriate metadata. Audio emission encoders (such as MPEG-H) often have the option of transmitting the renderer through the emission bitstream. This will allow for the use of the same renderer that was used for mixing - when played to consumers.
 - If mixing is done over a loudspeaker array, it is recommended that the renderer be indicated using appropriate metadata. Audio emission encoders (such as MPEG-H) often have the option of transmitting the renderer through the emission bitstream. This will allow for the use of the same renderer that was used for mixing - when played to consumers.

4.1.2.15 Interactivity and visual exploration

Compared to traditional media where the content presented in the display is mostly independent from user actions, content presented using VR devices react (at least in part) to user actions. At a minimal level (3DOF 360 videos for instance) the images displayed on the HMD are a function of the movement of the head of the user - which decides which part of the 360 environment to visually explore.

- This interactivity has strong implication for the creation of the content to grant an enjoyable and comfortable user experience.

4.1.3 Master Format

4.1.3.1.1 Projection and Aspect Ratio

The Master Video presentation format should have the following characteristics:

- Projection will be equirectangular projection.
- Co-ordinate system as described in [OMAFSODIS]
- The video may be generated from multiple camera arrays or composited imagery.
- The image should display no apparent stitch lines and occluded or missing picture information.
- Frame motion should not display motion artefacts such as blur or step motion.
- For fully 360° video, 2:1 Aspect Ratio with the following attributes:
 - No Zenith or Nadir Blind-spots are permitted.
- For partial 360° video
 - Metadata elements `RangeHorizontal` and/or `RangeTopBottom` as defined in Table 7 describing the partial region must be included.

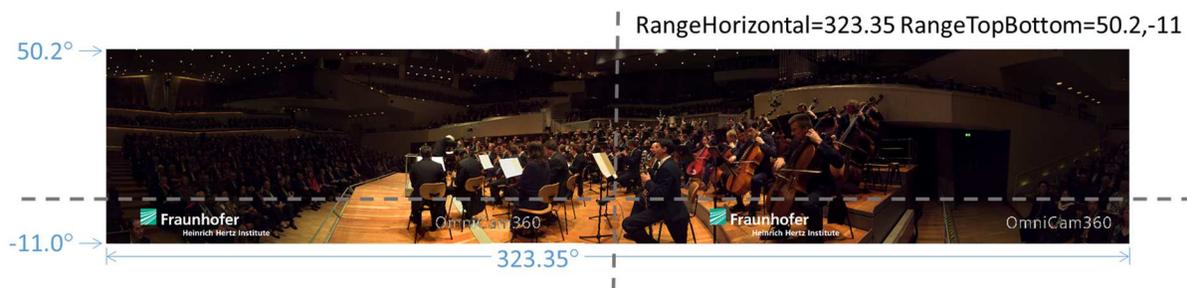


Figure 1: Example of Partial Panorama metadata

4.1.3.2 Video Master Format

A Master VR Video file from which transcodes for the various end-user platforms is desirable.

4.1.3.2.1 Video Metadata

Video Metadata to be included with video content is depicted in Table 7.

4.1.3.2.2 Resolution

- Stereoscopic format is: Separate Left/Right eye files.
- Metadata as depicted in Table 7.

For full 360° video

- Monoscopic: minimum 4096 H × 2048 V
- Stereoscopic: minimum 4096 H × 2048 V for each eye.

For partial 360° video

- 1:1 pixel ratio
- Horizontal minimum: $(\text{RangeHorizontal} / 360) \times 4096$
- Vertical minimum: $((\text{RangeTopBottom}[1] - \text{RangeTopBottom}[2]) / 180) \times 2048$

4.1.3.2.3 Frame rates

All media should be acquired and processed as Progressive frames. Acceptable progressive frame rates for monoscopic video are:

- 25 – Subject to content motion constraints
- 30 – Subject to content motion constraints
- 50
- 60

Acceptable Progressive Frame Rates for stereoscopic video are:

- 60
- 75
- 90
- 100
- 120

Comfortable frame rates for monoscopic is 50 and for stereoscopic is 100 [FRATES]. Other frame rates may be permitted by agreement.

4.1.3.2.4 File formats

Table 1: Master File formats

Bit Depth	10-bit
Color Sampling	4:2:2
Color Space	ITU Rec.709 (gamut levels within the threshold defined by EBU R103)
Scan	Progressive Frame
Delivery Format – Option 1	MXF Program Contribution (AMWA AS-11X1 as per DPP specification)
Delivery Format – Option 2	IMF Application 2e / JPEG2000 minimum data rate 150Mb/s
Delivery Format – Option 3	DnxHRHQ

4.1.3.3 Audio Master Format

4.1.3.3.1 File Format

An open production format such as the Audio Definition Model (ITU-R BS.2076-1) facilitates content exchange for contribution [ADM] should be used.

4.1.3.3.2 Audio Format

- For production in Scene-based Audio, full-spherical 4th order Ambisonics with or without additional audio objects for scene augmentation is recommended. The higher the Ambisonics order, the higher the spatial resolution and allows the better utilization of the capabilities of today's and future distribution platforms.
- The contribution should be in single tracks or in an interleaved track format with ACN ordering of the ambisonics signals [ACN]
- The normalization of the ambisonics signals should be N3D, or SN3D [NORMS]. N3D normalization is recommended for 32 bit resolution PCM floating point file formats. If legacy 16 bit PCM fixed-point file formats are used, SN3D normalization is recommended.

4.1.3.3.3 Bit depth and sampling rate

- A PCM file format supporting 32 bit floating point (e.g. ITU-R BS.2088 [BS2088]) is recommended. The minimum bit depth is 16 bits.
- The minimum sampling rate should be 48 kHz.

4.2 Media and Presentation Profiles

4.2.1 Introduction

tbd

4.2.2 Selected Media Profiles

4.2.2.1 Video

4.2.2.1.1 Overview

Editors Note: Download Optimization for (1) Quality and (2) File Size?

Editors Note: In case of download, what are the tradeoffs between viewport-independent vs. viewport-dependent? Consider quality, file size, decoder complexity, etc.

This section describes the selected media profiles for video, namely:

- (i) OMAF HEVC Viewport-Independent baseline media profile, further described in section 4.2.2.1.2
- (ii) OMAF HEVC Viewport-Dependent baseline media profile, further described in section 4.2.2.1.3

For the viewport-independent baseline media profile, the maximum achievable viewport resolution is constraint by the video decoder capabilities, specified by the elementary stream constraints in the video profiles with HEVC Main10 Profile, Main Tier, Level 5.1. The following table shows the maximum viewport resolutions that can be achieved with the viewport independent baseline media profile for HEVC Main10 Profile, Main Tier, Level 5.1, 60fps, for a display with an FOV of 90°×90° and a given content coverage.

Content coverage	Maximum viewport resolution
360°×180°	1K×1K
270°×180°	1.2K×1.2K
180°×180°	1.4K×1.4K
180°×120°	1.8K×1.8K

Table 2: Maximum achievable resolution in the viewport using viewport-independent baseline media profile for HEVC Main10 Profile, Main Tier, Level 5.1, 60fps, for a display with a FoV of 90°×90° and a given content coverage.

The OMAF HEVC Viewport-Dependent baseline media profile allows for achieving higher resolutions of the viewport compared to the entries of Table 2 given the same capabilities (HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with an FOV of 90°×90°). This guidelines describe how to generate content for this profile by mixing low and high-resolution tiles and thereby better leveraging the Max luma picture size as defined by the HEVC video profile and level definitions.

In the streaming case, the OMAF HEVC Viewport-Dependent baseline media profile can be used if resolutions higher than those achievable by the viewport independent baseline media profile are desired or for bandwidth optimization. However, viewport dependent streaming comes with additional latency requirements, as described in detail in section 5.4.3, that consist of sensor latency, network request delay, origin-to-edge delay (in case of cache miss), transmission delay (accounting for access network delay) and delays incurred in the client device due to buffering, decoding and rendering. These factors, might affect the time needed from the time instant the head movement happens until high quality/resolution content is shown to the user, when viewport dependent baseline media profile is used.

4.2.2.1.2 OMAF HEVC Viewport-Independent baseline media profile

4.2.2.1.2.1 Introduction

This media profile is defined in ISO/IEC 23090-2 [OMAFSODIS] section 10.1.2 and fulfils basic requirements to support omnidirectional video. Both monoscopic and stereoscopic spherical videos up to 360 degrees are supported. The profile does neither require viewport dependent decoding nor viewport dependent delivery. Regular HEVC encoders, DASH packagers, DASH clients, encryption technologies, file format parsers and HEVC decoder engines can be used for encoding, distribution and decoding. The profile also minimizes options to support basic interoperability.

This media profile is expected to be supported by HMDs and other devices rendering omnidirectional video powered by platforms released in 2015 and onwards. The key support is the availability of HEVC Main-10 level 5.1 decoders to process 4k by 2k frames at frame rates up to 60 fps as well as GPU-based rendering. The profile permits improved immersive experiences beyond the basic capabilities.

Figure 2 provides an overview of a possible receiver architecture that recovers the spherical video. Note that this figure does not represent an actual implementation, but a logical set of receiver functions. More implementation aspects are covered later. Based on an OMAF media stream, the receiver parses, possibly decrypts and moves the elementary stream to the HEVC decoder. Either the OMAF Metadata as defined in [OMAFSODIS] or the Decoder Rendering Metadata (SEI messages) may be used by the Texture-to-Sphere Mapping function to generate a spherical video based on the decoded output signal, also known as “texture”. The viewport is then generated from the spherical video signal by taking into account viewport position information from sensors, display characteristics as well as possibly other metadata such as initial viewport information. Whereas the decryption and decoding is typically done in hardware on devices, the OMAF restricted scheme permits to use existing texture mapping and rendering functionalities on GPUs to generate the viewport.

Editors Note: Secure rendering is for further study.

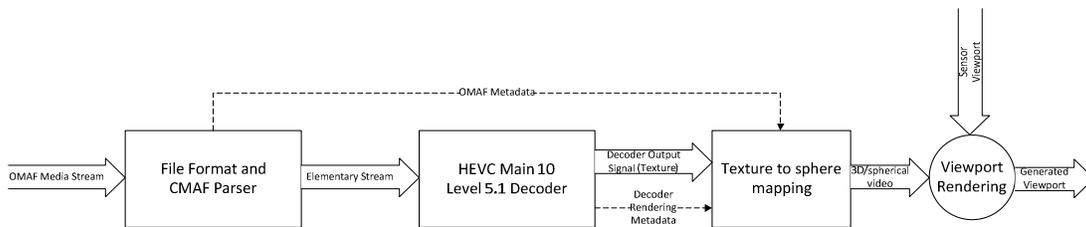


Figure 2: Receiver Model for OMAF HEVC Viewport-Independent baseline media profile

The projection in this profile is restricted exclusively to EquiRectangular Projection (ERP), but permits delivery of less than full 360 degrees spherical video.

The texture signal is restricted to 4096 samples in horizontal and vertical direction.

Details on definitions and coordinate systems can be found in [UPDSEI]. Note that the following terms are used in the remainder of this document following the definitions in this document:

- **coverage sphere region:** *sphere region* that is covered by a *cropped decoded picture*.
- **global coordinate axes:** coordinate axes associated with *omnidirectional video* that are associated with an externally referenceable position and orientation.
- **local coordinate axes:** coordinate axes having a specified rotation relationship relative to the *global coordinate axes*.
- **omnidirectional video:** video content in a format that enables rendering according to the user's viewing orientation, e.g., if viewed using a head-mounted device, or according to a user's desired *viewport*, reflecting a potentially rotated viewing position
- **packed region:** region in a *region-wise packed picture* that is mapped to a *projected region* according to a *region-wise packing*
- **projected picture:** picture that uses a *projection format* for *omnidirectional video*.
- **projected region:** region in a *projected picture* that is mapped to a *packed region* according to a *region-wise packing*.
- **projection:** specified correspondence between the color samples of a *projected picture* and azimuth and elevation positions on a sphere.
- **region-wise packed picture:** decoded picture that contains one or more *packed regions*.

- **region-wise packing:** transformation, resizing, and relocation of *packed regions* of a *region-wise packed picture* to remap the *packed regions* to *projected regions* of a *projected picture*.
- **sphere coordinates:** azimuth and elevation angles identifying a location of a point on a sphere.
- **sphere region:** region on a sphere, specified either by four *great circles* or by two *azimuth circles* and two *elevation circles*, or such a region on a rotated sphere after applying yaw, pitch, and roll rotations.
- **viewport:** region of *omnidirectional video* content suitable for display and viewing by the user.

4.2.2.1.2.2 External Specification

Video elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.3.2.2. In particular, SEI messages describing the omnidirectional video as defined in [UPDSEI] need to be present.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.3.2.3. ISO BMFF files that contain such an encoded track are identified by the brand 'ovid'. The OMAF metadata is equivalent to the information that is present in the SEI messages for omnidirectional video.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 section 10.1.2.6. An Adaptation Set including Representations formatted according to this media profile is recommended to be signaled as

- @codecs='resv.podv.hvc1.1.6.L93.B0' or @codecs='resv.erpv.hvc1.1.6.L93.B0'
- @mimeType='video/mp4 profiles="ovid"'
- A Supplemental Descriptor or Essential Descriptor providing the frame packing arrangement may be used.

Note: By the use of the restricted video scheme and the @profiles referring to this media profile, the DASH client has all information to identify if this media profile can be played back. For additional information, the Supplemental Descriptor is used to provide some details on the configuration of the contained Representations.

4.2.2.1.2.3 Quality and Performance

3GPP TR26.918 [TR26918] contains subjective test experiments that investigate the dependency of perceived visual quality on spatial resolution of the omnidirectional video. According to the results reported in 3GPP TR26.918 [TR26918], perceived quality increases with the video resolution. On the tested equipment, 3GPP TR26.918 [TR26918], 4K spatial resolution with ERP provides good quality.

While the tests in 3GPP are limited and not comprehensive, they are the only ones available up to date that can be publicly accessed. It is relevant to understand that no general conclusions can be drawn from the tests. Readers are encouraged to look at the detailed test setup and results before drawing any conclusions.

According to the TR, the coded video bitrate is expected to be in the range of 5-20 Mbit/s, depending on the content and good quality being achieved in the upper range of the bit rate range.

Editors Note: SOME MORE INFORMATION ON THIS SUBJECT WOULD BE BENEFICIAL

4.2.2.1.3 OMAF HEVC Viewport-Dependent baseline media profile

4.2.2.1.3.1 Introduction

This media profile fulfils basic and advanced requirements to support omnidirectional video. Both, monoscopic and stereoscopic spherical video up to 360 degrees is supported. This profile supports quality emphasis on the actual user viewport, which allows higher resolution in the viewport and/or reduced bitrate compared to the OMAF HEVC Viewport-Independent baseline media profile.

Editors Note: The MPEG OMAF HEVC Viewport-Dependent baseline media profile not only enables tiling but also additional capabilities to realize viewport dependency, such as region-wise packing. Currently, VRIF's OMAF HEVC Viewport-Dependent baseline media profile only considers tiling, and relevance of additional aspects of the OMAF HEVC Viewport-Dependent media profile are for further investigation. Moreover, VRIF will investigate the feasibility and benefit of defining a simpler viewport dependent media profile that does not require tiling.

This profile allows using tiles of different resolutions, while using a single video decoder conforming to HEVC Main 10 Profile, Main tier, Level 5.1 at the receiver side. When using the presented media profile, tiles of different resolutions can be combined to cover the entire omnidirectional video. I.e. the whole omnidirectional video is decoded with a resolution/quality emphasis on the actual user viewport, which cannot be done with viewport independent media profile and allows for higher resolutions to be displayed in the viewport as compared to the viewport independent media profile.

Within this profile one or more HEVC streams of the omnidirectional video are offered at different qualities/resolutions and are encoded comprising HEVC tiles that are encoded as Motion-Constrained Tile Sets (MTCS), i.e. tiles are encoded in such a way that does not reference other tiles. In addition, this profile includes ISO Base Media File Format (ISO BMFF) tools that allow for generating a single HEVC stream that can be decoded by a single HEVC Main 10 Profile, Main tier, Level 5.1 capable decoder and presented.

Minimum receiver capabilities required to support this profile:

- HEVC Main 10 Profile, Main tier, Level 5.1
- ISO BMFF extractors, as defined in MPEG ISO BMFF specification
- Region Wise Packing, as defined in MPEG OMAF specification

Note: The OMAF HEVC Viewport-Dependent baseline media profile allows considerable freedom in the usage of Region Wise Packing with up to 256 separate regions. Likewise, the HEVC specification allows up to 110 tiles in Level 5.1 of HEVC Main10 profile. This might lead to undesirable implementation complexity. Beyond simple tiling examples, the current draft of this document does not introduce detailed guidelines for tiling configurations or further restrictions of the above mentioned limits. Please note that such work is still under progress in VRIF and such guidelines are expected to be provided in the final version of the VRIF guidelines.

4.2.2.1.3.2 OMAF-DASH Viewport-Dependent Streaming and Download Client model

This section provides an overview of the OMAF-DASH streaming client model as well as of the OMAF Download client model and briefly describes their components. A detailed description of these components and the corresponding interfaces can be found in section 5.7.3.

Figure 3 shows a high-level structure of the OMAF-DASH client model with interfaces for streaming. It consists of 5 sub modules and illustrates the interfaces between them.

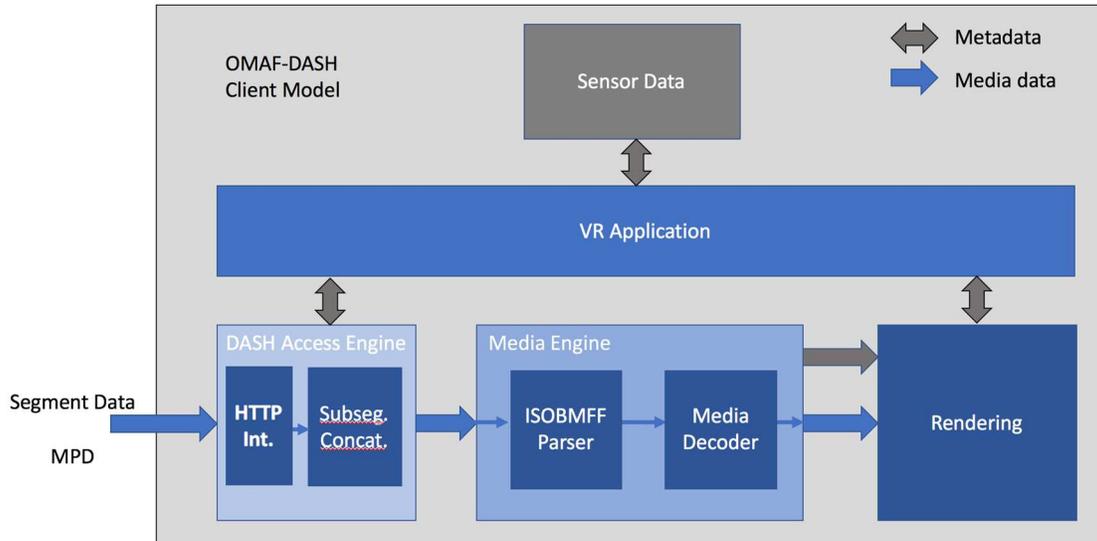


Figure 3: OMAF-DASH Streaming Client model with interfaces

DASH Access Engine: Downloads all OMAF Media streams and generates a single ISO BMFF file by concatenating subsegments as indicated in the figure.

VR application: Determines which OMAF Media streams should be downloaded by the DASH Access Engine. Controls the rendering depending on sensor data and HMD capabilities.

Media Engine: Plays the ISO BMFF file (i.e. plays the extractor track) and outputs a single NAL unit video stream into the Media Decoder that outputs decoded pictures and metadata into to the Renderer.

Renderer: Renders the decoded video pictures to the HMD display.

Sensor Data: User metadata (e.g. viewport position, direction, speed, etc.) taken e.g. from the HMD.

Note: Generally, it is preferable to use HW supported functionalities to optimize speed, latency, power consumptions and overall performance. Each of those above functions may be accessed with APIs. Specific APIs, possibly supported on SDKs and media frameworks are currently under development for example in Khronos or W3C.

A detailed description of each sub-module and its interfaces is given in section 5.7.3.

Figure 4 shows a high-level structure of the OMAF-Download client model with its interfaces.

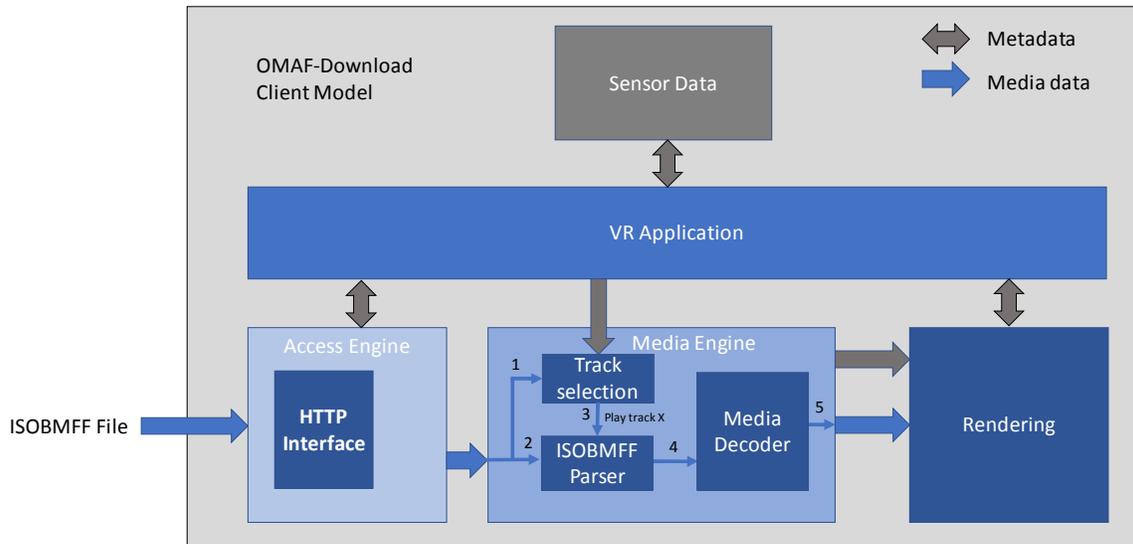


Figure 4: OMAF-Download Client model with interfaces

The main difference to the OMAF-DASH client for streaming as depicted in Figure 2 is the fact that the access engine does carry out any subsegment concatenation, since the download file contains all tracks corresponding to the tiles and all extractor tracks corresponding to the different viewports. The other main difference is that the VR applications has an interface to the track selection box to choose what extractor track to play back (3) based on the sensor data (i.e. based on the current viewport of the user). The extractor track is thereby selected dynamically based on the current user viewport and may change as coding configuration and client implementation allow.

Therefore, OMAF metadata corresponding to the “main” viewport of the tracks (i.e. viewport corresponding to the regions with higher quality/resolution): namely `RegionwisePackinBox`, `SphereRegionQualityRankingBox` or the `2DRegionQualityRankingBox` has to be parsed in the track selection box (1) in order to select the corresponding extractor track (3) based on the current user viewport. The entire ISO BMFF File is passed through to the ISO BMFF Parser (2) where, only the selected extractor track is parsed and a single corresponding HEVC Level 5.1 elementary stream (4) is passed to the Media Decoder. Then the stream is decoded and decoded pictures (5) are forwarded to the renderer along the necessary metadata for rendering.

4.2.2.1.3.3 External specifications

Video elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.1.3.2.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.1.3.3.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAFSODIS] section 10.1.3.4.

Note: Signaling of other parameters such as applied color transform and transfer characteristics is under discussion in MPEG and may be added.

4.2.2.1.3.4 Quality and Performance

Firstly, according to the results presented in 3GPP SA4 in 3GPP TR 26.918 [TR26918] in section 7.3, e.g. Table 7.3 [VDVS-QUAL], significant bitrate reduction can be achieved by tiling the omnidirectional video and mixing different resolution as enabled through this viewport dependent baseline media profile using HEVC Tiles.

Secondly, and as pointed out in section 2 of [VDVS-QUAL] with examples in section 10.1, compared to viewport-independent distribution methods, higher resolution can be obtained in the viewport by used viewport-dependent distribution methods such as this viewport dependent baseline media profile using HEVC Tiles.

4.2.2.2 Audio

4.2.2.2.1 Overview

This section defines media profiles for audio. Table 3 provides an overview on the supported features, but is not considered to be comprehensive. The detailed specification is subsequently provided in the referred section.

Table 3: Overview of OMAF media profiles for audio

Media Profile	Codec	Profile	Level	Max. Sampling Rate	Brand	Section
3D Audio Baseline	MPEG-H Audio	Low Complexity	1, 2 or 3	48 kHz	'oab1'	10.2.2

4.2.2.2.2 OMAF 3D Audio Baseline Media Profile

4.2.2.2.2.1 Introduction

This media profile fulfils the requirements to support omnidirectional audio. Channels, objects and Higher-Order Ambisonics (HOA) are supported, as well as combinations of those. The profile is based on MPEG-H 3D Audio [3DA].

MPEG-H 3D Audio [3DA] specifies coding of immersive audio material and the storage of the coded representation in an ISO Base Media File Format (ISO BMFF) track. The MPEG-H 3D Audio decoder has a constant latency, see Table 1 – “MPEG-H 3DA functional blocks, internal processing domain and delay numbers” of ISO/IEC 23008-3 [3DA]. With this information, content authors can synchronize audio and video portions of a media presentation, e.g. ensuring lip-synch. When orientation sensor inputs (i.e. pitch, yaw, roll) of an MPEG-H 3D Audio decoder change, there will be some algorithmic and implementation latency (perhaps tens of milliseconds) between user head movement and the desired sound field orientation. This latency will not impact audio/visual synchronization (i.e. lip synch), but only represents the lag of the rendered sound field with respect to the user head orientation.

MPEG-H 3D Audio specifies methods for binauralizing the presentation of immersive content for playback via headphones, as is needed for 360 Media VR presentations. MPEG-H 3D Audio specifies a normative interface for the user’s orientation, as Pitch, Yaw, Roll, and 3D Audio technology permits low-complexity, low-latency rendering of the audio scene to any user orientation.

4.2.2.2.2 External Specification

Audio elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.2.2.2.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAFSODIS] section 10.2.2.3. The ISO BMFF track should be identified by the brand ‘oab1’.

Note: The Media Profile specified in MPEG does not contain yet a compatibility brand. The ‘oab1’ brand is used for compatibility with CMAF Media Profile for OMAF 3D Audio Baseline Media Profile. It is envisioned that this will be corrected at the next MPEG meeting, and the ‘oab1’ brand will be used for the OMAF 3D Audio Baseline Media Profile while the ‘cab1’ brand will be used for compatibility with CMAF Media Profile for OMAF 3D Audio Baseline Media Profile, as specified in current document.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAFSODIS] section 10.1.3.4. An Adaptation Set including Representations formatted according to this media profile should provide the following signaling according to [RFC6381] and ISO/IEC 23008 3 [3DA] section 21 as shown in Table 4.

Table 4: MPEG-H Audio MIME parameter according to RFC 6381 and ISO/IEC 23008-3

Codec	Media type	codecs parameter	profiles	ISO BMFF encapsulation
MPEG-H Audio LC Profile Level 1	audio/mp4	mhm1.0x0B	‘oab1’	ISO/IEC 23008-3
MPEG-H Audio LC Profile Level 2	audio/mp4	mhm1.0x0C	‘oab1’	ISO/IEC 23008-3
MPEG-H Audio LC Profile Level 3	audio/mp4	mhm1.0x0D	‘oab1’	ISO/IEC 23008-3
MPEG-H Audio LC Profile Level 1, multi-stream	audio/mp4	mhm2.0x0B	‘oab1’	ISO/IEC 23008-3
MPEG-H Audio LC Profile Level 2, multi-stream	audio/mp4	mhm2.0x0C	‘oab1’	ISO/IEC 23008-3
MPEG-H Audio LC Profile Level 3, multi-stream	audio/mp4	mhm2.0x0D	‘oab1’	ISO/IEC 23008-3

4.2.2.2.3 Quality and Performance

MPEG-H Audio LC Profile provides excellent sound quality (as assessed per ITU-R BS.1534) for 2D and 3D program material as shown in 3D Audio Verification Test Report [N16584].

4.3 Content Security

This section presents guidelines for protecting Virtual Reality content. The word “content” is used advisedly: there are non-content VR assets that are not in the scope of VR content security. In these guidelines, VR content is defined as a new audio-visual media type that shares many characteristics of traditional linear audio-visual content but differs in significant ways (mainly rendering and display).

Accordingly, these guidelines highlight those differences from traditional video that have a material impact on VR content protection. This approach is based on the assumption that existing Digital Rights Management systems will be used as a baseline. The format will indicate the limitations of using existing DRM on VR content, allowing informed judgment to be made as to where additional protection mechanisms may be required.

Differences notwithstanding, the goals of protecting traditional and VR content are fundamentally the same. That is to protect the rights of the copyright holders and distributors. Of particular concern are unauthorized distribution (by whatever means), unauthorized modification (delivering an inferior experience), and violation of user privacy. There are also new usage rights VR content owners may want to control, e.g. entitlements for resolution, output control, or 3DoF vs 6DoF navigation.

4.3.1 Scope

VR experiences come in many forms. There is a diversity of presentation formats (e.g. augmented, immersive), display devices (e.g. HMDs, tablets), and distribution methods (e.g. streamed vs. downloaded). The landscape is constantly changing, and keeping up the security requirements of the entire ecosystem is not practical.

The targets of this section’s guidelines are those designing or specifying commercial systems for viewing VR content. Guidelines are therefore confined to the use cases identified by the VRIF’s Requirements and Guidelines committees (see the appendix for relevant documents). They are further restricted to those parts of the ecosystems developed enough to be commercially realizable, and based on standards.

A small set of characteristics are used to distinguish the systems discussed in the security guidelines. They are content type, distribution system, and interaction model, as shown in Table 5.

Table 5: VR Characteristics

Category	Video Format	Audio Format	Distribution Format	Duration	Live	Client Type
360° Video	Spherical	Spatial	DASH	Fixed	No	Player
	Spherical	Spatial	DASH	Open-ended	Yes	Player
	Spherical	Spatial	Downloaded ISO BMFF	Fixed	No	Player
Theatrical VR	Spherical	Spatial	n/a	Fixed	Either	Player
Interactive VR	Mixed ⁴	Spatial	Downloaded	Scripted	No	Game Engine ⁵

⁴ Visualizable media types besides spherical video include CGI Models, textures, point clouds, etc.

⁵ Or some other stand-alone application to handle interactivity.

The primary focus is 360° video that is streamed to clients equipped with a player (including web browsers extended to support view-dependent projection). The bulk of this content will be streamed, but downloading content in its entirety before playback will also have use cases.

Live 360° content is similar but will always be streamed. The major difference is that potentially time-consuming operations such as encoding or watermarking must be done in near real-time (introducing no more than a few seconds of latency).

Theatrical VR is a special case of 360° video where a traditional piece of traditional high-value content (e.g. a movie or TV show) is embedded in a virtual theatre environment. It is a special case because 1) the video source may have its own DRM (e.g. streaming video or Blu-Ray disk; 2) a compositing operation of high-value content and background is required; 3) the high-value content is by definition attractive to pirates.

Interactive VR refers to content intended for high-end VR systems such as the Oculus Rift, although it will increasingly be found on mobile devices as well. The key difference is that instead of only a fixed 360° video sequence, the images sent to the display are under the control of software, and may be a composite of multiple media types including CGI imagery. These non-video media types may have a value which the owner or creator wants to protect or track. Hence, they may have their own security requirements and corresponding solutions which may or may not be distinct solutions from the video element's security solutions.

Any of this content may be viewed with a HMD (fully immersive), or on a non-immersive device such as a tablet. In the former case the viewport is determined by the orientation of the HMD. In non-immersive displays, orientation may also be determined by sensors, or may be controlled directly by the user (e.g. finger swipes). This sort of display is sometimes referred to as a “magic window.”

4.3.2 MovieLabs ECP 1.1 Deltas

The MovieLabs Specification for Enhanced Content Protection (ECP) describes a set of high-level requirements for securing content. They are intended to be general enough to be applicable to any content distribution system (including future ones), but specific enough to serve as a template for evaluating a specific instance of such a system. No document can substitute for a certification process, but the ECP permits rapid identification of potential problem areas.

Although the ECP is a good starting point for defining guidelines for systems distributing VR content, additional work is required because VR is a fundamentally different media type than traditional audio-visual content. Some of these differences require new mechanisms (or modifications of existing ones) to secure the VR content. In this section those differences are noted as deltas to the ECP. They are enumerated following the format of the ECP, which is divided into four major sections: 1) threats to content; 2) requirements for DRM systems; 3) Platform Specifications; and 4) End-to-End System Specifications.

4.3.2.1 Problems/Threats

The ECP enumerates several threats to protected traditional content. These are also applicable to VR content, with the following notes:

- **Ripping Software (360° video):** For 360° video, the most common distribution format will be adaptive streaming, which is not strictly speaking ripping. Circumvention efforts will focus on these new formats.
- **Availability of Rips (360° video):** If the video is not tiled, transcoding it to a standard format and putting it on file-sharing and torrents is the logical path. If tiles are used to distribute, such

reconstruction will be much more difficult because of the difficulty of getting all the tiles and dealing with padding.

- **Ripping Software (Interactive Applications):** These are similar to AAA games, and similar attacks can be expected. These are focused on reverse-engineering the application well enough to remove the section involved with authentications.
- **Availability of Rips (Interactive Applications):** As with AAA games, hacked versions with authentication logic removed will be created and distributed.
- **Output Capture:** Intercepting the output of VR content may take more work to make useful because only the current viewport is displayed. There is an exception in the case of live sports and other events: others may be willing to accept another person's view of a premium event such as a championship game.

4.3.2.2 New Problems/Threats

VR has some issues that are new and not covered by the ECP. These are listed below:

- **Return Path Data:** Users engaging in VR experiences generate several types of return path data, including the HMD tracking data and controller input. This data is a form of "digital exhaust": it must be protected from interception and properly anonymized, otherwise it could reveal information about the user.
- **Degradation of Experience:** A user in an immersive experience is more vulnerable to disruption caused by targeted (e.g. spoofing input data) or generic (DDoS) attacks.

4.3.3 DRM System Specifications

The ECP lists requirements for DRM systems hosting protected content. The same requirements apply to systems hosting VR content, with the following notes:

- **Connection:** Not exclusive to VR, but note that broadcast environment includes scenarios that may require different content protections techniques, such as:
 - Online full 2-way communication
 - Online 1-way (broadcast)
 - Intermittent
 - Completely offline
- **Outputs & Link Protection:** DRM on the output (e.g. HMD view) needs to be selectable to account for use cases where re-transmission of a user's output (e.g. Twitch) is desirable. Also, many existing HMDs use chipsets that only support HDCP 1.4.

4.3.4 Platform Specifications

The ECP lists requirements for hardware platforms hosting protected content. The same requirements apply to platforms hosting VR content, with the following notes:

- **Encryption:** VR content has the same requirements for encryption as traditional content.
- **Secure Media Pipeline:** there are several issues that must be taken into account:
 - VR content that requires special hardware such as a GPU should be part of the SMP (note that this may not be possible on some current platforms).
 - Interactive VR content is not a simple decode, it may composite different media types.
- **Secure Computation Environment:** VR content has the same requirements for a Secure Computation Environment as traditional content.

- **Hardware Root of Trust:** VR content has the same requirements for a Hardware Root of Trust as traditional content.

4.3.5 End-to-End System Specifications

This section captures requirements whose implementation extends across multiple system components (e.g. client and server), with the following notes:

- **Forensic Watermarking:** VR audio-visual content is created and rendered differently from traditional content so current algorithms will not work, for example:
 - The Field of View in the HMD may only be a small percentage of the full 360° frame. This violates the assumption that a full frame is displayed.
 - The various VR display transforms (equirectangular projection, spherical lens warping, correcting for chromatic distortion, foveated rendering) may make the watermark difficult to recover.

This is an ongoing area of investigation; forensic watermarks are required to respond to breaches.

- **Playback Control Watermark:** recovering audio watermarks may be hindered by spatial audio processing. It is not clear this is relevant to VR content.
- **Breach Response:** VR content has the same requirements for breach responses as traditional content.

4.3.6 Encrypted Media Extensions

Encrypted Media Extensions (EME) is a recommendation under development by the W3C. It provides a framework for web browsers to support playback of HTML5 video protected by DRM without the use of plug-ins. Although controversial because it enables proprietary technology to be used in the otherwise open browser ecosystem, the W3C announced in July 2017 its intentions to publish EME as a web standard. EME is recommended for securely displaying VR content in an HTML5-based browser. For non-browser playback (i.e. native applications), traditional DRM systems should be deployed.

5 Vertical 1: OTT Download or Streaming of VR360 Content

5.1 Description of Vertical

This vertical primarily addresses the economically feasible distribution of VR360 content to emerging devices.

5.2 Guiding Example Use Cases

A service provider (content aggregator) offers a library of 360 A/V content. The library is a mixture of content formats from user generated content, professionally generated studio content, VR documentaries, promotional videos, as well as highlights of sports events. The content enables to change the field-of-view based on user interaction.

The service provider wants to create a portal to distribute the content to a multitude of devices that support 360-A/V and VR processing and rendering. This device may implement functions in hardware for reduced power and battery consumption, optimized processing, minimal thermal impacts and minimized latencies. Some solutions may be embodied in software (such as apps). Typically, VR applications make use of well-defined interfaces to hardware functionality, notably decoders.

The service provider wants to target two types of applications:

- Primarily, view in a HMD with head motion tracking.
- As a by-product, the content provider may permit viewing on a screen with the field-of-view for the content adjusted by manual interaction (e.g. mouse input or finger swipe)

The service provider expects different types of consumption and rendering devices with different capabilities in terms of decoding and rendering. However, it wants to target devices that fulfil a certain quality threshold expressed by decoder and rendering capabilities.

The service provider has access to the original footage of the content and is permitted to encode and transcode to appropriate distribution formats.

The footage includes different types of 360 A/V VR content, such as

- For video:
 - One of
 - Pre-stitched monoscopic video, i.e. a (360 and possibly less than 360) spherical video without depth perception, with Equirectangular Projection (ERP).
 - Pre-stitched stereoscopic video, i.e. a spherical video using a separate input for each eye, typically with ERP.
 - Original content
 - Original content, either in on original uncompressed domain or in a high-quality mezzanine format.
 - Basic VR content: as low as 4k x 2k (ERP), 8 or 10bit, BT.709, as low as 30fps
 - High-quality: up to 8k x 4k (ERP), 10 bit, possibly advanced transfer characteristics and color transforms, sufficiently high frame rates, etc.
 - Sufficient metadata is provided to appropriately describe the A/V content
- For audio

- Spatial audio content for immersive experiences, provided in the following formats:
 - Channel-based audio
 - Object-based audio
 - Scene-based audio
 - Or a combination of the above
- Sufficient metadata for encoding, decoding and rendering the spatial audio scene permitting dynamic interaction with the content. The metadata may include additional metadata that is also used in regular TV applications, such as for loudness management.
- Diegetic and non-diegetic audio content.

The service provider is responsible for monetizing the content and fulfilling necessary accessibility requirements. Subtitles are considered to be important and need to be supported in a standardized way.

The service provider is also responsible for securing the content, if required by the content provider, including DRM systems.

The service provider wants to

- reach as many devices as possible
- minimize the number of different formats that need to be produced and distributed
- ensure that the content is presented in highest quality on the different devices.

The service provider provides an application (e.g. browser-based, native app) or makes use of an installed third party application, and may rely on the decoding and rendering capabilities of the device, typically in hardware or by pre-installed or downloaded software decoders.

The service provider wants to reach devices that are already in the market or are expected to be in the market by end of 2017.

The service provider wants to avoid testing each device, but rather prefers simple interoperability, e.g. standardized interfaces.

The service provider wants to enable that some of the library items can be downloaded to devices, primarily through HTTP, and is played back on the device after downloading. The service provider wants to ensure that a device downloads only content that it can decode and render while providing the best user experience for the device capabilities.

For certain library items, the service provider wants to ensure that content is rendered instantaneously after selection, so a DASH-based streaming is considered. The service provider wants to ensure that a device accesses only content that it can decode and render while providing the best user experience for the device capabilities. The service provider also wants to ensure that the available bandwidth for the user is used such that the rendered content for the user is shown in the highest quality possible.

5.3 Reference Architectures

5.3.1 Distribution Architecture

The architecture introduced in this section addresses service scenarios for the distribution of VR content in file or segment based download and streaming services, including DASH-based services.

The role of the VR Content Provider, the VR Service provider, the distributor, the application and the service platform are differentiated.

Figure 5 considers a functional architecture for such scenarios. VR Content is captured by a VR Content provider and split in audio B_a and video in B_v on the interfaces. Both media come with metadata and are synchronized in time and space. The content is uploaded to a VR Service Provider Portal which stores the original footage. Then the content is prepared for distribution by pre-processing, encoding and file format/DASH encapsulation. Interface D_a and D_v provide formats that enable encoding by existing media encoders. After media encoding, the content is made available to file format encapsulation engine as elementary streams E and the file format may generate a complete file for delivery or segmented content in individual tracks for DASH delivery over interface F . Metadata may be added. Content may be made available in different viewpoints, so the same content may be encoded in multiple versions. Content may also be encrypted.

At the receiving end, there is an expectation for the availability of a VR application that communicates with the different functional blocks in the receiver's VR service platform, namely, the delivery client, the file format decapsulation, the media decoding, the rendering environment and the viewport sensors. The reverse operations are performed. The communication is expected to be dynamic, especially taking into account the dynamics of sensor metadata in the different stages of the receiver. The delivery client communicates with the file format engine, and different media receivers decode the information and provide also information to the rendering.

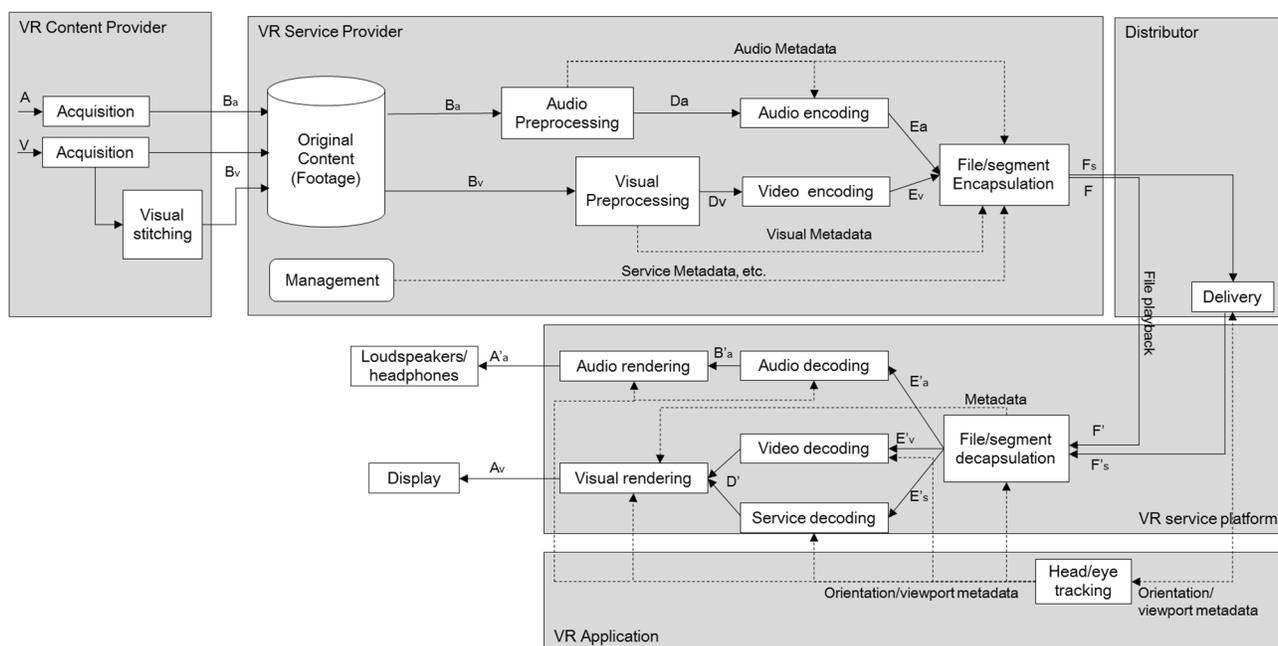


Figure 5: Example Architecture for simple VR Streaming Services

Note that certain functionality (such as audio decoding and audio rendering) depicted in the VR Service Platform box above may in certain circumstances take place within the VR Application box, depending on the VR Service Provider's needs and platform capabilities. However, there are benefits in enabling VR Applications to use a native VR Service platform for decoding and rendering to minimize latency, thermal impact, processing power and power consumption.

Figure 6 provides an attempt for an encrypted system. It is important that the interfaces to the secure decoding and rendering platform are limited and therefore require that well defined conforming bitstreams are provided to the trust zone.

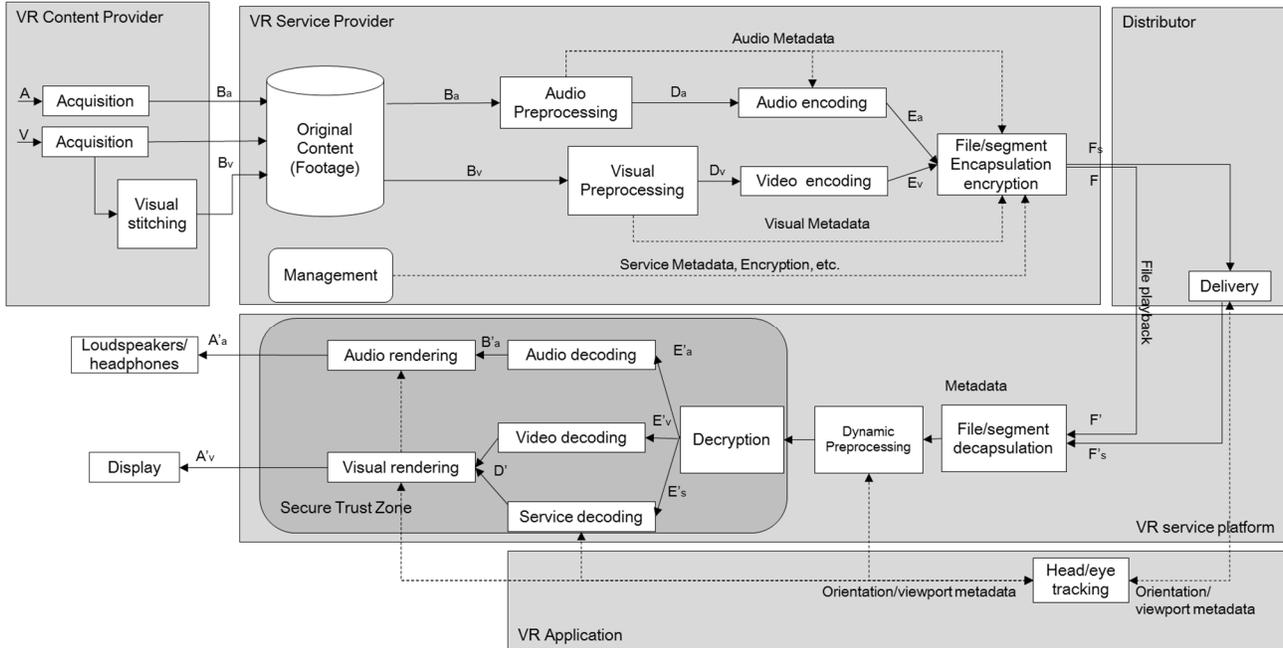


Figure 6: Example Architecture for encrypted VR Streaming Services

5.3.2 Client Architecture

Editors Note: This section is to be completed

5.4 Technical Enablers

5.4.1 Suitable Media Profiles

5.4.1.1 Video

The OMAF HEVC Viewport-Independent baseline media profile as presented in section 4.2.2.1.2 may be used to fulfil the use case and provide broad interoperability.

The OMAF HEVC Viewport-Dependent baseline profile as introduced, may be used to fulfil the use case and provide broad interoperability, while achieving a higher resolution at the viewport than that of the OMAF HEVC Viewport-Independent baseline media profile.

5.4.1.2 Audio

The OMAF 3D Audio Baseline Media Profile may be used to fulfil the use case and provide broad interoperability.

5.4.2 Suitable Presentation Profiles

Editors Note: Viewport independent baseline media profile

5.4.3 Distribution Systems

5.4.3.1 OTT Distribution using HTTP

This section primarily addresses interface F on distribution in Figure 5. In the context of this use case, streaming and download are considered. In recent years the use of HTTP-CDNs for distribution of content over the open Internet has gained more and more attraction and is nowadays the premium OTT distribution means. The popularity of HTTP is manifold, but primarily the scalability of CDNs with distributed caching architectures, the ability to pass firewalls and NATs as well as the ability of HTTP protocol stacks on many different devices as well as in browser endpoints makes HTTP the most far-reaching distribution protocol. The downsides of TCP/IP and HTTP (such as variable bitrates, object-based delivery, download latencies, etc.) have been compensated by smart formats and the usage of the formats in different applications. Predominantly download and adaptive bitrate (ABR) streaming is used for distributing new media services. Furthermore, the formats and the HTTP-based APIS also more and more have found support in not only unmanaged distribution, but also in IP multicast as well as in broadcast, e.g. in 3GPP MBMS or ATSC3.0. In addition, HTTP architectures and CDNs are continuously improved adding new protocols such as HTTP/2.0 or combine unicast and multicast protocols.

While for regular TV applications, downloading an entire content is considered less and less appealing (except for recent services that enable offline access to content libraries, for example in air planes), for VR services and content download may still be a very attractive option, in particular if the content is short, the content is not live and/or the real-time access to high-quality content is not possible due to bitrate restrictions. Hence, download will remain to be an attractive option and with the use of HTTP-based download, CDN capabilities can be fully exploited. It is beneficial that the application can check the content before downloading it and HTTP headers and capability exchange can be used for this purpose together with the formats well defined Internet media type using the content-type HTTP header.

The most popular standardized protocols for streaming are HTTP Progressive Download and Dynamic Adaptive Streaming over HTTP (DASH). Both enable to provide services to deliver on-demand VR content over HTTP protocols, including the metadata and media data composing the on-demand VR service. As such, again standard HTTP servers and standard HTTP caches can be used for hosting and distributing on-demand VR content. Note that in the context of this distribution model, it is expected that the delivery network is typically unaware of the content, whether it is VR or any other media content. This makes HTTP-based delivery attractive for launching services.

An example VR distribution system is depicted in Figure 7. On-demand VR content including media content and metadata may be stored on one or more media origin servers, along with the DASH media presentation description (MPD). The MPD contains the relevant information on the different encoded versions of the DASH VR content, including VR-related content information such as those on available viewports, projection and region-wise packing metadata, along with the traditional MPD parameters on codecs, bitrates and resolutions. The media origin server is typically an HTTP server such that the MPD and media segments related to on-demand VR content can be requested via clients and be delivered via HTTP. A DASH client in the user terminal obtains a current viewing orientation or viewport e.g. from the HMD that detects the head orientation and possibly also eye orientation. By parsing metadata from the MPD, the DASH client concludes which Adaptation Set and Representation cover the current viewing orientation at the highest quality and at a bitrate that can be afforded by the prevailing estimated network throughput. The DASH client issues (Sub)Segment requests accordingly. In case of HTTP progressive download, the on-demand VR content may be included in an ISO base media file format (ISO BMFF) file as one track and the entire ISO BMFF file may be offered on an HTTP server or on a CDN for downloading.

The massively scalable distribution of on-demand VR content is typically enabled via content delivery networks (CDNs) that consist of a geographically distributed set of HTTP proxy caches, with the goal of enabling content access to end users with high availability, proximity and high performance. At the network edge, VR content may be delivered through the client devices via different access networks, such as 4G/5G access and cable/WiFi access. Client compute devices (e.g., PC, tablet, smartphone, etc.) typically perform processing toward de-capsulation, decoding and rendering of the VR content, and forward the processed input to be displayed at the HMD.

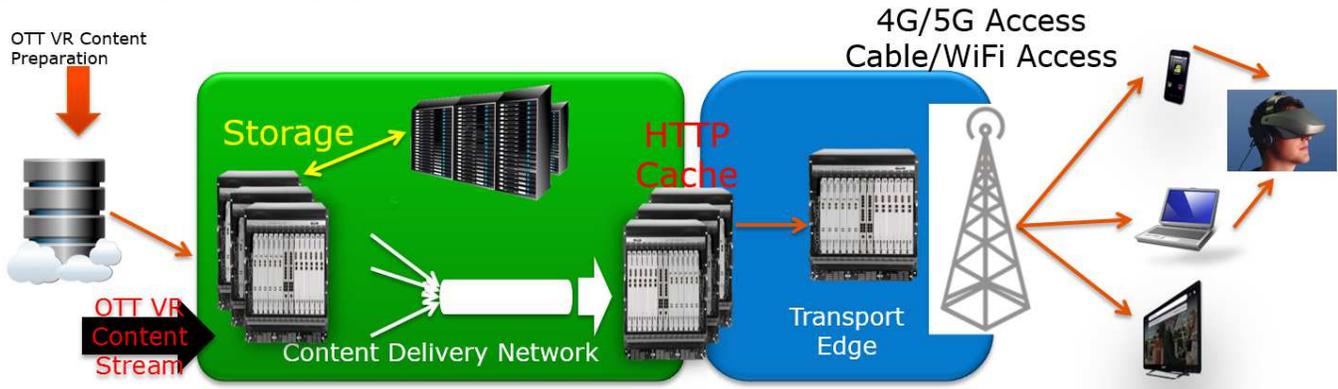


Figure 7: Example VR distribution system

Figure 8 shows an example protocol stack for on-demand VR content delivery services. OMAF-based media content, associated metadata, DASH-formatted MPD and media segments can be made accessible through HTTP.

Video Audio	DASH Media Presentation Description
OMAF File Format	
HTTP	
TCP	
IP	

Figure 8: Example protocol stack for VR content distribution

5.4.3.2 Download

If the OMAF HEVC Viewport-Independent baseline media profile is used for download, one track is included in the ISO BMFF file that follows the requirements and recommendations of the media profile in section 4.2.2.1.2.

If the OMAF HEVC Viewport-Dependent baseline media profile is used for downloading tiles, one 'hvc1' track per tile per resolution is included in the ISO BMFF file and one 'hvc2' track per potential viewing direction is included in the ISO BMFF that follow the requirements and recommendations of the OMAF media profile in section 10.1.3.3 in [OMAFSODIS].

If the OMAF 3D Audio Baseline Media Profile is used for download, one track is included in the ISO BMFF file that follows the requirements and recommendations of the media profile in section 4.2.2.2.2, with the first sample of the movie as a Stream Access Point (SAP) of type 1 (i.e. sync sample).

5.4.3.3 DASH Distribution

If the OMAF HEVC Viewport-Independent baseline media profile is used for DASH-based streaming, one Adaptation Set is included in each Period based on the requirements and recommendations of the media profile in section 4.2.2.1.2.

If the OMAF HEVC Viewport-Dependent baseline media profile is used for tile-based DASH streaming, one Adaptation Set per tile per resolution is included in each Period based on the requirements and recommendations of the OMAF media profile in section 10.1.3.4 in [OMAFSODIS]. In addition, one Adaptation Set per tile per potential viewing direction is included in each period that contains a @dependencyId attribute or a Preselection as a Supplemental descriptor.

Figure 9 shows an exemplary DASH configuration where two tiles (Tile 1 and Tile 2) with two different resolutions (high and low resolution) are grouped into separate Adaptation Sets (1, 2, 3 and 4). Each of those Adaptation Sets contains three Representations where each MCTS track is encoded with a different bitrate. In addition, two Adaptation Sets (5 and 6), representing two different viewport orientations, are included in the same Period pointing to the correct Adaptation Sets containing one of the tiles in high and one of the tiles in low resolution.

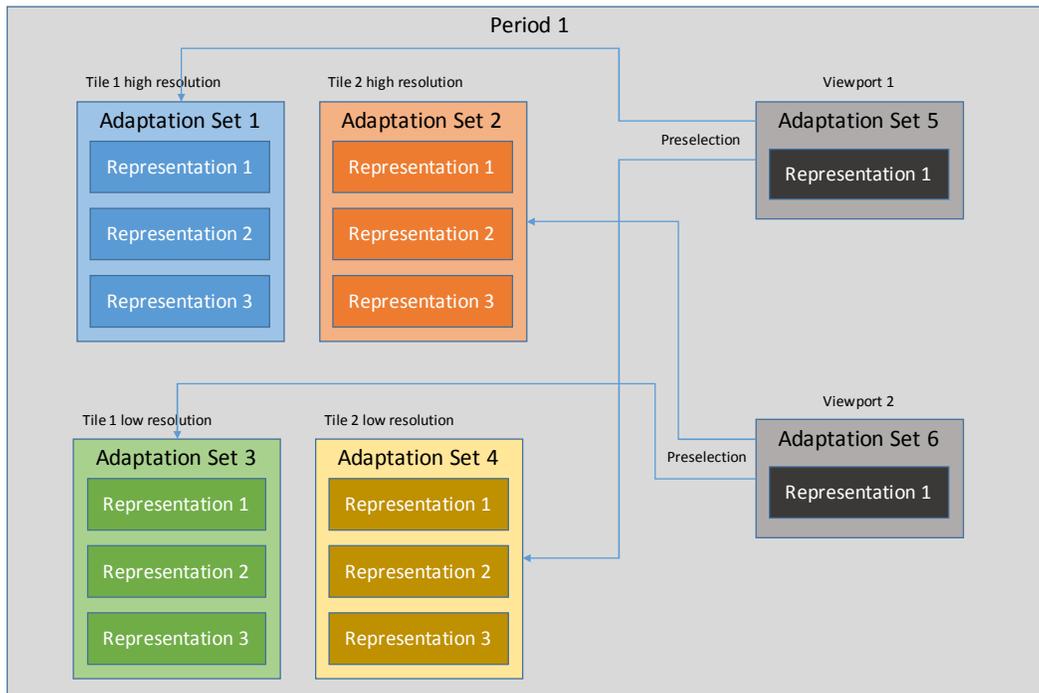


Figure 9: Exemplary DASH configuration setup

Editors Note: This expresses the preselection implementation using the descriptor rather than the concept. The OMAF HEVC Viewport-Dependent baseline media profile only supports use of Preselection as a descriptor and not as an element.

If the OMAF 3D Audio Baseline Media Profile is used for DASH-based streaming, one Adaptation Set is included in each Period based on the requirements and recommendations of the media profile in section 4.2.2.2.2, with the first sample of every fragment as a Stream Access Point (SAP) of type 1 (i.e. sync sample).

5.4.3.4 CDN Considerations

The role of a CDN is critical for enabling access to on-demand VR content with high availability and high performance. On top of the traditional adaptive bitrate (ABR) / DASH content delivery, on-demand VR content delivery requires further degree of client adaptation as it involves continuous change of the user's viewports (e.g., as the user wearing the HMD changes his/her head orientation), which makes it important that CDNs enable means to reduce download latencies toward enabling interactive VR experiences with high degree of responsiveness. High latency in delivering the user viewport may cause poor user experience, including motion sickness. CDN may benefit from edge computing techniques deployed in proximity to the users in the service provider or operator network to allow the content served close to the clients, and thereby minimizing latency and optimizing network bandwidth efficiency.

CDN strategies toward improving OTT VR content delivery may depend on the specific OMAF media profile in use. In particular, CDN considerations could differ between viewport independent vs. viewport dependent media profiles. In case of viewport independent VR content format, the viewport agnostic nature of the VR content delivery implies that it is sufficient for the CDN to apply the existing ABR / DASH content delivery enhancement mechanisms (and regular DASH clients may be used), but also ensure that the delivered throughput performance meets the bandwidth requirements of VR content delivery (e.g., as reported in sections 7.2 and 7.3 of 3GPP TR 26.918 [TR26918] for viewport-independent and viewport-dependent media profiles, respectively), since streaming the entire 360° panorama requires significant bandwidth. In case of viewport dependent VR content format, there are further optimizations possible at the CDN level. In this case, the DASH VR content may be made available for different viewports, so the same content and associated media segments may be stored in multiple versions. For example, the CDN may cache different DASH adaptation sets corresponding to different viewport versions (e.g., tiles) of the content to provide the client the ability to interactively switch across different viewports (e.g., in response to user's change in head orientation). In case of viewport dependent delivery via use of tiled streaming, this allows that high quality / resolution tiles corresponding to the new viewport can be fetched very quickly by the client from the CDN. Thus, CDN's VR content caching ability corresponding to different viewports may help towards reducing download delays and also potentially improving the network bandwidth utilization efficiency.

A relevant latency metric to consider for tiled streaming is motion to high resolution (M2HR) latency, defined as the latency from the change in head orientation to display of the first high resolution frame based on the new field of view. The sources contributing to M2HR latency as depicted in Figure 10 and include sensor latency, network request delay, origin-to-edge delay (in case of cache miss), transmission delay (accounting for access network delay) and delays incurred in the client device due to buffering, decoding and rendering. All of these latency sources are relevant for viewport-dependent streaming while only the rendering latency is relevant for viewport-independent streaming.

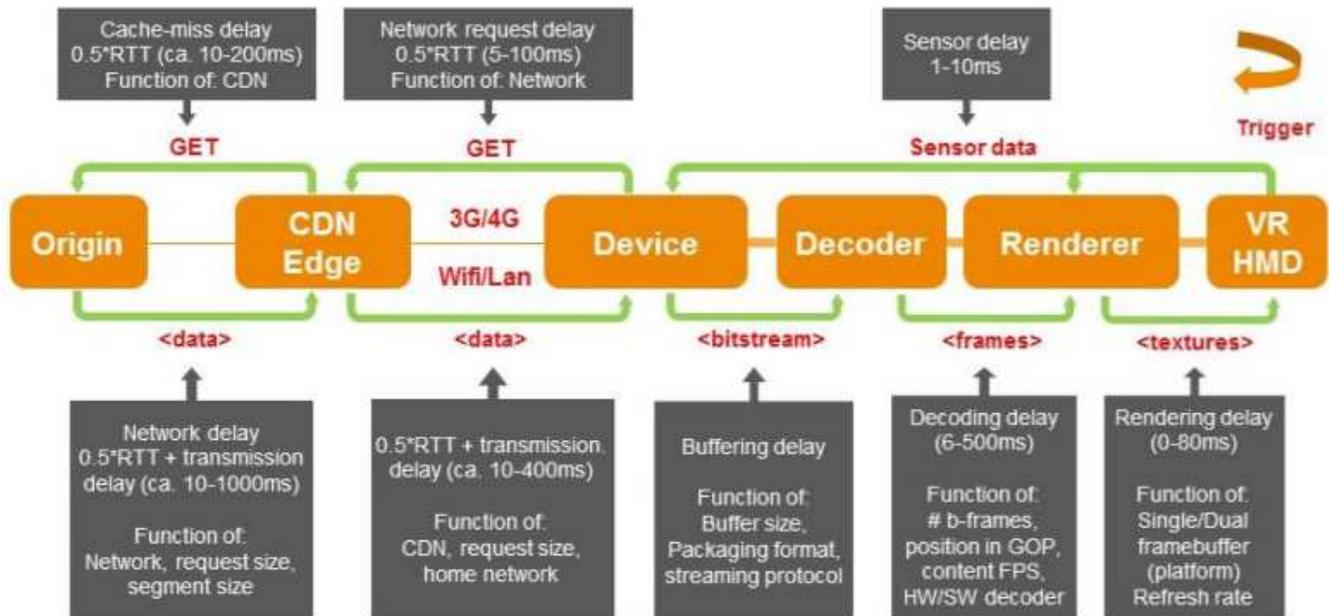


Figure 10: Latency sources contributing to M2HR latency [CDNOPT]

As reported in [CDNOPT], the M2HR latency cost of tile-based streaming can be reduced considerably by various CDN and streaming protocol optimizations including:

- Intelligent caching of tiles at the edge to avoid cache misses and consequent origin-to-edge latency
- Transport level improvements such as the use of HTTP/2.0 and QUIC

5.4.3.5 Access Network Considerations

5.4.3.5.1 3GPP/4G/5G

3GPP Packet Switched Streaming Service [TS26234] enables on-demand VR progressive download and DASH content delivery over HTTP protocols, using 3GPP-based core network architectures (e.g., evolved packet core (EPC)) and radio access air interfaces (e.g., Long Term Evolution (LTE)). 3GPP is currently developing a new 5G air interface that is expected to deliver data rates much higher than those of LTE. Feasibility of VR services in 3GPP networks has been studied in 3GPP TR 26.918 [TR26918].

5.4.3.5.2 Cable/Fiber/Copper and WiFi to Home

WiFi and cable/fiber/copper technologies provide access network connectivity to the home. WiFi is developing its next generation standards to deliver better data rates to support high-bandwidth applications such as VR, with data rates up to 2-4 Gbps (see [80211AC] and [80211AX]). Currently deployed fiber access technologies (G-PON and EPON) deliver speeds up to 3Gbps and next generation (XG-PON and 10G-PON) will deliver data rates up to 10 Gbps.

5.5 Guidelines for Service Providers

5.5.1 Suitable Production Formats

5.5.1.1 OMAF HEVC Viewport-Independent baseline media profile

The original video signal may be a full 360 sphere content.

The original video signal may also be restricted in coverage, i.e. only cover a subset of the full 360 sphere as indicated in the parameters `RangeHorizontal` and/or `RangeTopBottom`.

The original video source may be monoscopic or stereoscopic as indicated by the `stereoMode` metadata element.

The following parameters are expected to be constant over the sequence of the content:

- Spatial resolution
- Frame rate
- Coverage

If the original video signal format is in the constraints of HEVC Main-10 Main-Tier level 5.1 and the scheme constraints, then the OMAF HEVC Viewport-Independent baseline media profile may be used directly on the source signal to generate elementary streams following the media profile constraints as described in section 10.1.2.2 of [OMAFSODIS] and the encoding and content preparation in section 5.5.2.

The relevant HEVC Main-10 Main Tier level 5.1 are:

- Max luma sample rate (samples/s) is 534,773,760
- The maximum bitrate (kbit/s) is 40,000

Mono signals that can be distributed with this profile need to fulfill the following requirements:

$$\text{ceil16}(\text{FullWidthPixel} - (\text{Cropping.Left} + \text{Cropping.Right})) *$$

$$\text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) *$$

$$\text{FrameRate} \leq 534,773,760$$

With $\text{ceil16}(x)$ the smallest integer that is greater or equal than x and a multiplier of 16.

Stereo signals that can be distributed with this profile need to fulfill the following requirements:

$$\text{ceil16}(\text{FullWidthPixel} - (\text{Cropping.Left} + \text{Cropping.Right})) *$$

$$\text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) *$$

$$\text{FrameRate} \leq 267,386,8800$$

Based on this, examples for production formats that can directly be distributed are:

- 4096 H × 2048 V, 4:2:0, mono at 25, 30, 50 and 60fps with full content coverage
- 6144 H × 3072 V, 4:2:0, mono at 25 and 30fps with full content coverage
- 4096 H × 2048 V, 4:2:0, per eye at 25 and 30fps, if frame-packed in top and bottom and sample aspect ratio 1:1 with full content coverage
- 4096 H × 2048 V, 4:2:0, per eye at 50 and 60fps with up to 180 degree content coverage, if frame-packed in top and bottom, sample aspect ratio 1:1
- 6144 H × 3072 V, 4:2:0, mono at 25, 30, 50 and 60fps with 180 degree content coverage

If the original video signal is beyond the constraints of HEVC Main-10 level 5.1 constraints or the constraints dictated by the restricted scheme applied for this profile, then the OMAF HEVC Viewport-Independent baseline media profile may be used after preprocessing of the original video content such that the constraints are fulfilled. Examples for videos that require preprocessing are

- 4096 H × 2048 V, 4:2:0, per eye stereo at 50 and 60fps may be preprocessed to
- 6144 H × 3072 V, 4:2:0, mono at 50 and 60fps with full content coverage
- 6144 H × 3072 V, 4:2:0, stereo with full content coverage
- 8192 H × 4096 V, 4:2:0, mono or stereo with full content coverage

The original signal is then pre-processed, encoded and distributed following the constraints for this media profile as described in section 5.5.3.

5.5.1.2 OMAF HEVC Viewport-Dependent baseline media profile

For the OMAF HEVC Viewport-Dependent baseline media profile, any original video using the Equirectangular Projection (ERP) can be used as defined in section 4.1.3.2.

5.5.1.3 OMAF 3D Audio Baseline media profile

If the original audio signal is in the constraints of MPEG-H 3D Audio, LC profile, Level 3, then the OMAF 3D Audio Baseline Media Profile should be used directly on the source signal to generate elementary streams following the media profile constraints.

If the original audio signal is beyond the constraints of MPEG-H 3D Audio, LC profile, Level 3, then the OMAF 3D Audio Baseline Media Profile may be used after pre-processing of the original audio content such that the constraints are fulfilled. The pre-processed signal is then encoded and distributed the constraints for this media profile.

5.5.2 Sphere-to-Texture Mapping and SEI Message Generation

Editors Note: The text in this section is primarily for the viewport-independent media profile.

A key issue is the mapping of the spherical video to a 2D texture at the content generation and the reverse operation at the receiver. Based on the system diagram in section 4.2.2.1.2, SEI messages are added to describe the source content and the expected resulted processing from the 2D texture to a spherical video.

The mapping of the color samples of 2D texture images onto a spherical coordinate space in angular coordinates (ϕ , θ) for use in omnidirectional video applications for which the viewing perspective is from the origin looking outward toward the inside of the sphere. The spherical coordinates are defined so that ϕ is the azimuth (longitude, increasing eastward) and θ is the elevation (latitude, increasing northward) as depicted in Figure 11.

Rotation angles yaw (α), pitch (β), and roll (γ) are also used in the specification of these semantics.

Relative to an (x , y , z) Cartesian coordinate system, yaw expresses a rotation around the z (vertical, up) axis, pitch rotates around the y (lateral, side-to-side) axis, and roll rotates around the x (back-to-front) axis. Rotations are extrinsic, i.e., around x , y , and z fixed reference axes. The angles increase clockwise when looking from the origin towards the positive end of an axis.

Assume a signal with the following parameters is provided:

- Projection is ERP with the coordinate system and mapping according to Figure 11.
- The frame rate of the signal is provided as **FrameRate**

- The full reference 360 video has spatial resolution **FullWidthPixel** times **FullHeightPixel** with picture aspect ratio 2:1
- The signal may follow the mono or stereo. If stereo, the signal is provided separately per eye. The type is expressed in the **StereoMode** parameter.
- The signal may have a restricted coverage expressed in the **Coverage** Parameter, if present, in the spherical domain expressed as follows:
 - **AzimuthMin** specifies the minimum azimuth value of the coverage sphere region in the range of -360 to 360 degrees.
 - **AzimuthMax** specifies the maximum azimuth value of the coverage sphere region in the range of -360 to 360 degrees. This value is greater than **AzimuthMin**.
 - **ElevationMin** specifies the minimum elevation value of the coverage sphere region in the range of -90 to 90 degrees.
 - **ElevationMax** specifies the maximum elevation value of the coverage sphere region, in in the range of -90 to 90 degrees.
- The signal may have prerotation expressed in the **Rotation** parameter, if present, in the spherical domain expressed as follows
 - **RotationYaw** specifies the value of the yaw rotation angle in the range of -180 to 180 degrees. When not present, the value is inferred to be equal to 0.
 - **RotationPitch** specifies the value of the pitch rotation angle in the range of -90 to 90 degrees. When not present, the value is inferred to be equal to 0.
 - **RotationRoll** specifies the value of the roll rotation angle in the range of -180 to 180 degrees. When not present, the value is inferred to be equal to 0.
- If not the full signal is provided but a cropped version of it, then this is expressed by the **Cropping** Parameter with the four following values
 - **Top**: the number of pixel cropped by on the top compared to the full pixel height.
 - **Right**: the number of pixel cropped by on the right compared to the full pixel height.
 - **Bottom**: the number of pixel cropped by on the bottom compared to the full pixel height.
 - **Left**: the number of pixel cropped by on the left compared to the full pixel height.
- The provided image sequence therefore has a luma component with
 - Width being **FullWidthPixel - (Cropping.Left + Cropping.Right)**
 - Height being **FullHeightPixel - (Cropping.Top + Cropping.Bottom)**
 - Note that the Cropping parameter should be chosen such that all pixels that are in coverage are included in the image.

The local projected sphere coordinates (ϕ, θ) for the sample location for the centre point of a sample location (i, j) is derived as follows:

- $\phi = (\text{AzimuthMin} + (0.5 - i \div \text{FullWidthPixel}) * (\text{AzimuthMax} - \text{AzimuthMin}))$
- $\theta = (\text{ElevationMin} + (0.5 - j \div \text{FullHeightPixel}) * (\text{ElevationMax} - \text{ElevationMin}))$

If the **Rotation** parameter is not present, then the global projected sphere coordinates (ϕ', θ') for the sample location for the centre point of a sample location (i, j) are identical to the local sphere coordinates (ϕ, θ) .

If the **Rotation** parameter is present with parameters **RotationYaw** (α), **RotationPitch** (β), **RotationRoll** (γ) - all in units of degrees - then the global projected sphere coordinates (ϕ', θ') for the sample location for the centre point of a sample location (i, j) are derived based on its the local sphere coordinates (ϕ, θ) as follows

- $x1 = \text{Cos}(\phi) * \text{Cos}(\theta)$
- $y1 = \text{Sin}(\phi) * \text{Cos}(\theta)$
- $z1 = \text{Sin}(\theta)$
- $x2 = \text{Cos}(\beta) * \text{Cos}(\gamma) * x1 - \text{Cos}(\beta) * \text{Sin}(\gamma) * y1 + \text{Sin}(\beta) * z1$
- $y2 = (\text{Cos}(\alpha) * \text{Sin}(\gamma) + \text{Sin}(\alpha) * \text{Sin}(\beta) * \text{Cos}(\gamma)) * x1 + (\text{Cos}(\alpha) * \text{Cos}(\gamma) - \text{Sin}(\alpha) * \text{Sin}(\beta) * \text{Sin}(\gamma)) * y1 - \text{Sin}(\alpha) * \text{Cos}(\beta) * z1$
- $z2 = (\text{Sin}(\alpha) * \text{Sin}(\gamma) - \text{Cos}(\alpha) * \text{Sin}(\beta) * \text{Cos}(\gamma)) * x1 + (\text{Sin}(\alpha) * \text{Cos}(\gamma) + \text{Cos}(\alpha) * \text{Sin}(\beta) * \text{Sin}(\gamma)) * y1 + \text{Cos}(\alpha) * \text{Cos}(\beta) * z1$
- $\phi' = \text{Atan2}(y2, x2) * 180 \div \pi$
- $\theta' = \text{Asin}(z2) * 180 \div \pi$

Editors Note: THIS DIAGRAM SHOULD BE USED IN THE INTERFACE BETWEEN CONTENT PRODUCTION AND DISTRIBUTION

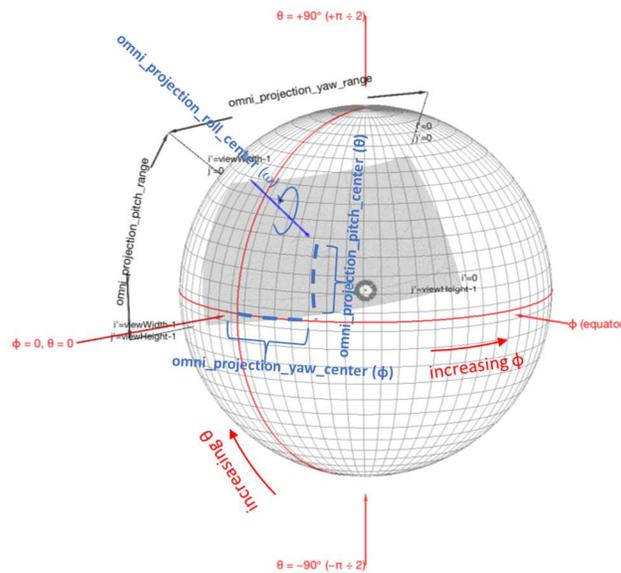


Figure 11: Spherical surface coordinates ϕ , θ with yaw, pitch, and roll of the region of a sphere covered by the cropped output picture relative to the equator and 0 meridian

Editors Note: THIS DIAGRAM MUST BE UPDATED TO DESCRIBE THE CORRECT CONTENT MODEL

The above content parameters may be mapped directly to the encoded signal or a preprocessing needs to be applied such that the above parameters are adjusted. Without loss of generality we assume that the above parameters are now directly mapped to the relevant SEI messages.

The equirectangular projection SEI message (as defined in, sections D.2.41.1 and D.3.41.1 of [UPDSEI]) provides information to enable remapping of the color samples of the output decoded pictures onto a spherical coordinate space in angular coordinates (ϕ , θ) for use in omnidirectional video applications for which the viewing perspective is from the origin looking outward toward the inside of the sphere. The

spherical coordinates are defined so that ϕ is the azimuth (longitude, increasing eastward) and θ is the elevation (latitude, increasing northward) as depicted in Figure 11.

The following general rules apply for SEI message generation:

- An SEI message with payload type 150 is generated
- The `erp_cancel_flag` is set to 0
- The `erp_persistence_flag` is set to 1

When the video provides full 360 coverage and no **Coverage** parameter is present, then the `erp_explicit_coverage_range_flag` is set to 0.

When the video does not provide full 360 coverage as indicated by the **Coverage** parameter), then

- the `erp_explicit_coverage_range_flag` is set to 1
- the `erp_azimuth_min`, `erp_azimuth_max`, `erp_elevation_min` and `erp_elevation_max` are set accordingly using the **Coverage** parameter values and the mapping defined in D.3.41.5 of [UPDSEI]
- the region-wise packing SEI messages (as defined in sections D.2.41.3 and D.3.41.3 of [UPDSEI]) should be generated in order to maximize the visible information in the encoded 2D image using the **Cropping** information parameters as follows:
 - The `rwp_cancel_flag` is set to 0
 - The `rwp_persistence_flag` is set to 1
 - `num_packed_regions` is set to 1
 - `proj_picture_width` is set to **FullWidthPixel**
 - `proj_picture_height` is set to **FullHeightPixel**
 - `packing_type[0]` is set to 0
 - `proj_region_width[0]` is set to **FullWidthPixel - (Cropping.Left + Cropping.Right)**
 - `proj_region_height[0]` is set to **FullHeightPixel - (Cropping.Top + Cropping.Bottom)**
 - `proj_region_top[0]` is set to **Cropping.Top**
 - `proj_region_left[0]` is set to **Cropping.Left**
 - `transform_type[0]` is set to 0
 - `packed_region_width[0]` is set to **FullWidthPixel - (Cropping.Left + Cropping.Right)**
 - `packed_region_height[0]` is set to **FullHeightPixel - (Cropping.Top + Cropping.Bottom)**
 - `packed_region_top[0]` is set to **Cropping.Top**
 - `packed_region_left[0]` is set to **Cropping.Left**

When the video is stereoscopic, then the frame packing needs to be generated and an appropriate frame packing arrangement SEI message (as defined in [UPDSEI] section D.3.41) needs to be generated as follows

- An SEI message with payload type 45 is generated
- The `frame_packing_arrangement_cancel_flag` is set to 1
- The `frame_packing_arrangement_type` is set to one of the following values: 3 or 4. For more details on the choice of one of the formats, see below.
- The `quincunx_sampling_flag` is set to 0

Using frame-compatible plano-stereoscopic video formats means that the left-eye and right-eye images are arranged in a spatial multiplex which results in a composite image that can be treated like a conventional 2D image. Annex A of TS 101 547-2 [TS1015472] provides an informative overview of the frame

compatible video formats and how a single 2D image can be generated if `frame_packing_arrangement_type` with a value of 3 or 4 is in use.

Editors Note: More details to be added

5.5.3 Encoding and Content Preparation

5.5.3.1 OMAF HEVC Viewport-Independent baseline media profile

If the original video is in the constraints of the encoder, then it may be distributed directly.

If the original video signal is beyond the constraints of HEVC Main10 level 5.1, then the original video content needs to be adapted to be encoded properly with an HEVC Main-10 level 5.1 encoder and to meet the profile level constraints. Adaptation includes

- Temporal subsampling

The preferred order and guidelines on what preprocessing should be used is for further study. In the absence of further study the above order also expresses a preference.

As examples, the signals beyond the limits of the HEVC profile level constraints documented in section 5.5.1 may be preprocessed as follows:

- 4096 H × 2048 V, 4:2:0, per eye stereo at 50 and 60fps may be preprocessed to
 - 4096 H × 2048 V, 4:2:0, per eye at 50 and 60fps, if frame-packed and sample aspect ratio 2:1 or 1:2
 - 3264 H × 1632 V, 4:2:0, per eye at 50 fps, if frame-packed in top and bottom and sample aspect ratio 1:1
 - 2880 H × 1440 V, 4:2:0, per eye at 50 or 60 fps, if frame-packed in top and bottom and sample aspect ratio 1:1
- 6144 H × 3072 V, 4:2:0, mono at 50 and 60fps with full content coverage
 - To any of the formats that can be directly distributed.
- 6144 H × 3072 V, 4:2:0, stereo with full content coverage
 - To any of the formats that can be directly distributed.
- 8192 H × 4096 V, 4:2:0, mono or stereo with full content coverage

The original or preprocessed video signal may also be restricted in coverage, i.e. only cover a subset of the full 360 sphere as indicated in the parameters `RangeHorizontal` and/or `RangeTopBottom`.

In this case it is recommended that:

- The signal is properly rotated such that the covered area is centric
- The signaled is properly cropped such that a minimum of the non-covered area is included in the original signal.

Further types of adaptation may be applied for efficient encoding such as content pre-rotation. This is typically applied for increasing coding efficiency and can be achieved by moving the content specific high-motion regions into content regions (typically ERP equator) where motion is less distorted (compared to ERP poles) through global rotation. Rotation, if applied, is static for the entire stream and cannot be applied dynamically.

Either provided by the source signal or after adaptation and pre-processing, the content is expected to be in the constraints of the following parameters:

- For mono:
 - 4096 H × 2048 V, 4:2:0, at 25, 30, 50 and 60fps with full content coverage or more than 180 degree coverage
 - 6144 H × 3072 V, 4:2:0, mono at 25 and 30fps with full content coverage or more than 180 degree coverage
 - 6144 H × 3072 V, 4:2:0, mono at at 25, 30, 50 and 60fps with less than 180 degree content coverage
- For stereo
 - 4096 H × 2048 V, 4:2:0, per eye at 25 and 30fps, if frame-packed in top and bottom and sample aspect ratio 1:1 with full content coverage
 - 4096 H × 2048 V, 4:2:0, per eye at 50 and 60fps with up to 180 degree content coverage, if frame-packed in top and bottom, sample aspect ratio 1:1

Editors Note: Additional items to be determined and added

According to the requirements of the OMAF profiles, SEI messages are added based on the original or pre-processed sequence to signal the used projection format, pre-rotation, region-wise packing and Frame-packing SEI message are added as well, if the processing applies. For details on the SEI messages, see section 5.5.2.

Depending on the applications, the content provider should take into account regular random access points in the encoding, for example every 2 seconds.

If the content is prepared for adaptive bitrate streaming, then also the constraints from the HEVC CMAF Video Track as defined in [ISOCMAF] annex B.1 should be taken into account. Multiple quality representations may be generated by adapting the bitrate of the video. Note that each Representation is required to have the same OMAF metadata and SEI messages in order to ensure consistency when bitrate switching is applied.

Regular DASH or ABR recommendations for content encoding may be used (see DASH-IP IOP [DASHIFIOP] or CMAF Annex D [ISOCMAF]). The number of Representations per Adaptation Set as well as their encoding bitrates depend on different factors, such as encoder performance, content complexity, distribution parameters and so on. In the absence of other information, a first idea on suitable bitrates and their performance, the 3GPP TR26.918 [TR26918] provides some ideas. (cf-further details in section 4.2.2.1.2.3). Generated Representations should be checked for perceptual quality and it is recommended to check the lowest bitrate Representations, if they still meet the perceptual quality expectations. If they don't, those may preferably not be offered to regular DASH clients as valid alternatives.

However, important to note that the profile prohibits spatial sub-sampling of Representations in one Adaptation Set to ensure that the rendering metadata is identical for all Representations in one Adaptation Set. A basic mapping is provided in Table 6, more details are **tbd**.

Editors Note: WE NEED TO SAY HOW TO GENERATE THE OMAF METADATA FROM THE SEI MESSAGES AND HOW ALL SAMPLE ENTRIES ARE SET

Table 6: Mapping of SEI Message Information to OMAF Metadata

SEI Message	OMAF Metadata
equiangular_projection	ProjectedOmniVideoBox

	CoverageInformationBox
region_wise_packing	RegionWisePackingBox
frame_packing_arrangement	StereoVideoBox

Otherwise, no specific aspects for VR content need to be taken into account.

Figure 12 provides an overview on the encoding process such that the spatially aligned and time-synchronized content can be prepared for distribution for this media profile.

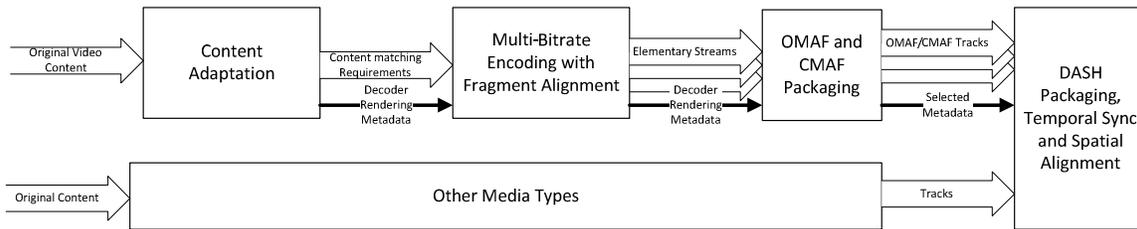


Figure 12: Content Preparation for DASH Distribution

5.5.3.2 OMAF HEVC Viewport-Dependent Profile

If the OMAF HEVC Viewport-Dependent baseline profile is used for distribution or download, video to be encoded may use Equirectangular Projection (ERP) or Cubemap Projection (CMP) as a projection. The ERP or CMP video may be encoded by using a Motion Constraint Tile Set (MCTS) capable HEVC Main10 encoder.

Since the production format is only defined for ERP, if CMP is used the content needs to be converted from ERP to CMP. For guidance on the conversion between projection formats, the reader is referred to [ADDSEI] and [PRO]CONV]. A related software package is available at:

https://jvet.hhi.fraunhofer.de/svn/svn_360Lib/

Preprocessing of the original video content may also be required to generate different versions (e.g. resolutions) of the original video that are then encoded and offered fulfilling the constraints of the OMAF HEVC Viewport-Dependent baseline media profile as explained below.

Three main configuration parameters need to be chosen:

- Tiling granularity: number of tile columns and rows ($N \times M$; N =Number of horizontal tiles; M =Number of vertical tiles)
- Available resolutions: number of resolutions (R =Number of resolutions) and ratios between them
- Representations with a preferred viewing direction: combinations of tiles with different resolutions corresponding to a different preferred viewing direction with the combinations mixing tiles of the available resolutions (C =Number of Viewing Directions)

In order to determine the three main configuration parameters, following characteristics may be taken into account:

- Target FoV
- Resolution of ERP or CMP video before pre-processing
- Target display resolution

Editors Note: Detailed guidelines on tiling layout is still work in progress. For instance, examples are to be provided on tiling layout for ERP and CMP and how this impacts content quality. Relevant input parameters to decide on the tiling configuration are also to be completed.

Figure 13 illustrates the Video content preparation for DASH distribution using the OMAF HEVC Viewport-Dependent baseline media profile.

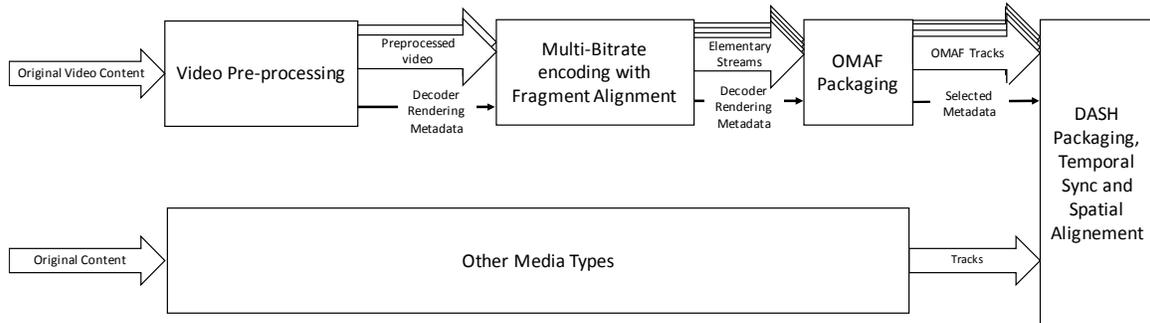


Figure 13: Video content preparation for DASH distribution with OMAF HEVC Viewport-Dependent baseline profile

Video pre-processings

Depending on the original video content format, various pre-processing steps may be performed before feeding the input video into the encoder:

- Projection: The input output video of the pre-processing module needs is required to be ERP or CMP projected. If the original video uses a different projection format a transformation to the targeted projection has to be applied. For further information on projection formats and conversion the reader is referred to [ADDSEI] and [PROJCONV].
- Subsampling: The input video may need to be spatially sampled to generate lower resolution variants of the input video when it is desired to offer mixed resolution video.
- Pre-rotation: For better compression efficiency, it might be beneficial to pre-rotate the input video before encoding.

The output of the video pre-processing is one video or multiple videos (of the same source) at different resolutions, which are used in the Multi-Bitrate encoding with Fragment Alignment function in order to create elementary streams with MCTS for each tile.

Multi-Bitrate and encoding with Fragment Alignment

As a result, the R pre-processed videos at different resolutions are encoded into $N \times M$ tiles (with N =number of horizontal tiles and M =number of vertical tiles), each tile being an MCTS. Each of the R pre-processed videos at different resolutions might be encoded at different bitrates, with all streams having aligned fragments.

An exemplary configuration could be a CMP video at 2 different resolutions with each cube face encoded using 2×2 tiles, i.e. 24 tiles per resolution. A viewport-dependent combination of these encoded tiles could consist of 8 high-resolution tiles and 16 low-resolution tiles.

SEI messages need to be added to signal the used projection format and possibly the use of frame-packing, if stereoscopic video is distributed.

The content provider should take into account that fast switching capabilities are desired for the OMAF HEVC Viewport-Dependent baseline media profile. Therefore, it is recommended to have frequent random access points in the encoding, for example every second or less.

DRAFT FOR PUBLIC PREVIEW

For DASH streaming cases, since very frequent random access points might decrease the coding efficiency, it could be desirable to have an accompanying media stream with less frequent random access point (e.g. every 2 seconds), so that users that are not changing their viewport at a point in time can download a video bit stream that has a better efficiency.

Editors Note: When elementary streams are offered that have a different switch point interval, certain considerations need to be taken into account on alignment of the elementary streams. For instance, same GOP structures are chosen for the elementary streams and the switch points of the elementary stream with a longer switching point period are aligned to those of the elementary stream with the shorter switching point period, i.e. if the AU with a given Decoding Time (DT) is an IDR in the bistream with longer switching point period the AU with same DT in the short switching point period is an IDR as well. Additional information will appear in following version of the document.

OMAF Packaging

After encoding, each bitstream of each tile is encapsulated into separate ISO BMFF file that contains a single track referred to as the MCTS track. In addition, several ISO BMFF files containing the extractor tracks are generated. In the following, more details of this process are provided:

Generation of ISO BMFF files with MCTS tracks

First, each of the $N \times M$ tiles is converted into an HEVC conforming bitstream following the MCTS sub-bitstream extraction process described in [ADDSEI]. Thereby, each MCTS track is HEVC conformant. A different `track_ID` is assigned to each MCTS track. Each tile has a unique `tileID` (derived as `MCTS index + 1` following [ADDSEI]) assigned to it which is greater than zero and is used later for calculation of the `track_IDs`. Then, each tile is encapsulated into a single ISO BMFF file containing a single MCTS track with the following considerations as described in [OMAFSODIS]:

- The `track_not_intended_for_presentation_alone` flag of the `TrackHeaderBox` may be used to indicate that a track is not intended to be presented alone.
- ‘povd’: indicating the projection used (e.g. ERP or CMP)
- ‘covi’: indicating the coverage of the track

These files include ‘hevd’ in the compatible brands in the ‘ftyp’ box. Each of the MCTS tracks covering the tiles converted into an HEVC conforming bitstream is offered with a restricted sample entry ‘resv’ with an original format box ‘frma’ indicating the ‘hvc1’. The `HEVCDecoderConfiguration` contains the Equirectangular or Cubemap Projection SEI message that was obtained as a result of the MCTS sub-bitstream extraction process or generated at this stage. The `track_ID` of the file is calculated using following equation:

$$trackID = tileID + r \cdot N \cdot M$$

with $r \in [0, \dots, R - 1]$ indicating which resolution level they belong to, and $tileID > 0$, so that the all `track_IDs` are different as required in the ISO BMFF specification.

Generation of ISO BMFF files with extractor tracks

Second, C tile combinations are chosen that result in a rectangular shape fulfilling the media constraints of the OMAF HEVC Viewport-Dependent baseline media profile, each of which has a different preferred viewing direction.

Using the chosen C combinations, the files containing the extractor tracks are generated. There are C extractor track files which define different viewport configurations. Each extractor track file contains exactly one extractor track and a set of MCTS tracks TL_c , which correspond to tiles at their respective

resolutions for the selected viewport combination. All tracks are added to the 'moov' box of the extractor track file. The `track_ID` of the track containing the extractors is set to:

$$trackID = c + 1 + R \cdot N \cdot M$$

with $c \in [0, \dots, C - 1]$ as a number of a specific viewport configuration. This track contains all the `track_IDS` of TL_c in the 'tref' box (also included in the 'moov' as mentioned).

There are $N \times M$ extractors for each sample within the extractor track; each of which is generated using an inline constructor and a sample constructor as defined in [NAL].

- **Inline Constructor:** Data carried in the extractor track (NAL and slice headers)

The inline constructor contains the NALU and slice header of the original bitstreams with an adjusted `slice_segment_address` correctly reflecting the spatial position of the each tile within the combined video bitstream.

- **Sample Constructor:** Data referenced by the extractor track (slice payloads)

The sample constructor references the dependent track of the corresponding tile and contains a `data_offset` field that allows skipping the data up to the first byte of the `slice_segment_data()`, i.e. skipping the slice header within the MCTS track. In addition, the `data_length` within the sample constructor is set to the maximum value.

These files further contain:

- 'povd': indicating the projection used (e.g. ERP or CMP)
- 'rwpk': indicating the regions for each of the tiles with their respective dimensions and position within the region-wise packed frame.

As for the files containing the MCTS tracks, these files include 'hevvd' in the compatible brands in the 'ftyp' box. The extractor tracks are offered with a restricted sample entry 'resv' with an original format box 'frma' indicating the 'hvc2'. The `HEVCDecoderConfiguration` contains the Omnidirectional Projection SEI as for the whole bitstreams before the extraction process, i.e. with full coverage. In addition, the `HEVCDecoderConfiguration` contains the Region-wise packing SEI message. Note that the region-wise packing SEI needs to be included into the extractor tracks by the OMAF Packaging module to signal the region-wise packed picture result of combining tiles with different resolutions into a rectangular picture. The region-wise packing SEI contains the same information as the 'rwpk' box and is included to the `HEVCDecoderConfiguration`.

Either `SphereRegionQualityRankingBox` or `2DRegionQualityRankingBox` should be added to the ISO BMFF file containing extractor tracks.

In case of including `SphereRegionQualityRankingBox`, the following applies:

- `region_definition_type` is equal to 1 if the projection format is ERP or equal to 0 if the projection format is equal to CMP.
- `num_regions` has the same value as in 'rwpk'
- `quality_ranking` is set to the value $r + 1$, with $r \in [0, \dots, R - 1]$ indicating which resolution level the region (corresponding tile) belongs to

In case of including `2DRegionQualityRankingBox`, the following applies:

- `num_regions` has the same value as in 'rwpk'

- regions defined by `left_offset`, `right_offset`, `top_offset` and `bottom_offset` are aligned to the regions defined in `'rwpk'`
- `quality_ranking` is set to the value $r + 1$, with $r \in [0, \dots, R - 1]$ indicating which resolution level the region (corresponding tile) belongs to.

DASH Packaging

See section 5.5.4.2.

5.5.3.3 OMAF 3D Audio Baseline media profile

If the OMAF 3D Audio Baseline Media Profile is used for distribution or download, the audio elementary stream may be generated by using a regular MPEG-H 3D Audio LC Profile, Level 3 encoder.

The content provider should take into account regular random access points in the encoding, for example every 2 seconds. Otherwise, no specific aspects for VR need to be taken into account.

If the content is prepared for adaptive bitrate streaming, then also the constraints from the MPEG-H Audio Track as defined in [ISOCMAF], Annex J should be taken into account. Multiple bitrates may be generated by adapting the bitrate of the audio.

5.5.4 Distribution

5.5.4.1 Download

5.5.4.1.1 OMAF HEVC Viewport-Independent baseline media profile

If the OMAF HEVC Viewport-Independent baseline media profile is used for distribution for download, the generated bitstream is included in an ISO BMFF file as one track.

The entire ISO BMFF file may be offered on an HTTP server or a CDN for download. The HTTP server should expose the capabilities of the file. No specific requirements need to be taken into account for the file format encapsulation. The `Content-Type` in the HTTP header should be set as

```
video/mp4 profiles='ovid' codecs='resv.erpv.hvc1.1.6.L93.B0'
```

5.5.4.1.2 OMAF HEVC Tile Based Viewport-Dependent Profile

If the OMAF HEVC Viewport-Dependent baseline media profile is used for download, one 'hvc1' track per tile per resolution is included in the ISO BMFF file and one 'hvc2' track per potential viewing direction is included in the ISO BMFF that follow the requirements and recommendations of the media profile in section 10.1.3.3 of [OMAFSODIS].

The entire ISO BMFF file may be offered on an HTTP server or a CDN for download. The HTTP server should expose the required capabilities to process the file.

5.5.4.2 DASH Streaming

5.5.4.2.1 OMAF HEVC Viewport-Independent baseline media profile

If the OMAF HEVC Viewport-Independent baseline media profile is used for distribution for DASH Streaming, regular DASH distribution means can be used. For details of DASH distribution methods, please refer to DASH-IF IOP Guidelines [DASHIFIOP].

Different types of DASH client architectures may be considered:

- 1) A native DASH client on an existing platform is used. In this case the MPD is handed to the playback and the DASH client is expected to handle sufficiently well downloading and playback of the content. This example follows for example the model in smart TVs for HbbTV today or playback of a DASH Media Presentation in a video element. This model is typically referred to as type 1.
- 2) In a variant of this, the application still uses a native client, but has additional control interfaces to influence DASH client decisions. This may control the playback of the media (rendering, stop and resume, etc.). Such a model is provided in today in browsers with interfaces on the video element to control certain features. This model is typically referred to as type 2.
- 3) The DASH client is part of the application and the application optimizes the operation for this profile. Such cases require that the application has a full DASH client library. Examples for such approaches are dash.js or Shaka Player for which the DASH client part of the web page. This model requires more knowledge for the app provider on DASH operations, but also provides flexibility and optimization potentials for the app provider. In this case the APIs to the service platform are typically on codec and elementary stream level. This model is typically referred to as type 3.

Prior to more knowledge on factors impacting the performance of streaming VR content, it is recommended that a conservative approach is taken with a focus on quality of the distributed video rather than on factors such as reduced latency, fast start-up and so on.

Among others, the following is recommended:

- At startup, not necessarily the lowest bitrate should be chosen, but a bitrate that is expected to provide sufficiently good quality.
- Buffer sizes should rather be kept longer, for example in the range of several seconds to 30 seconds. This reduces the necessity for rebuffering or requiring the need to down-switch to Representations with too low quality.
- If the content is short, then an approach similar to progressive download, e.g. filling the buffer with half of the content before playback, may be preferable. Buffer sizes up to 30 seconds may be considered by the DASH client to ensure high-quality playback. While this may result in longer startup delays, it is expected that this is beneficial for the user experience. Other application means may be used to enable long buffering, or content may for example be generated with low complexity in the beginning to ensure that sufficiently long buffers can be built.

5.5.4.2.2 OMAF HEVC Tile Based Viewport-Dependent Profile

If the OMAF HEVC Viewport-Dependent baseline media profile is used for tile-based DASH Streaming, an MPD file is generated with $N \times M \times R + C$ Adaptation Sets: one for each tile at each resolution ($N \times M \times R$) and one for each of the C extractor tracks. This implies that each Adaptation Set contains Representations with the same resolution. Each of the $N \times M \times R$ Adaptation Sets (corresponding to each tile at each resolution) can contain a Spherical Spatial Relationship (SSR) supplementalProperty element as defined in [OMAFSODIS] to signal which portion on the sphere is covered by the corresponding tile. Several Representations might be available within each Adaptation Set if each tile at a given resolution is encoded at different bitrates. These, Adaptation Sets may contain a Preselection descriptor as an Essential Property descriptor to indicate to which Preselection they belong to.

Each of the C Adaptation Sets (corresponding to each of the C extractor tracks) contain a region-wise quality ranking (RWQR) supplementalProperty element as defined in [OMAFSODIS] in order to signal which tiles have higher quality in respect to other tiles addressed by the corresponding extractor track. This information helps the client to identify the correct Adaptation Set (with high quality inside the desired viewport) depending on the viewing orientation of the client within an omnidirectional video. The RWQR

descriptor is set according to the `SphereRegionQualityRankingBox` or `2DRegionQualityRankingBox` as described in section 5.5.3. If `SphereRegionQualityRankingBox` is present in the ISOBMFF file, the value of `shape_type` in the SSR descriptor is set to the same value as `region_definition_type` of the `SphereRegionQualityRankingBox`. Values of `center_azimuth`, `center_elevation`, `center_tilt`, `hor_range`, and `ver_range` of the SSR descriptor are set to the values of `center_azimuth`, `center_elevation`, `center_tilt`, `hor_range`, and `ver_range` in `SphereRegionStruct` of the `SphereRegionQualityRankingBox`. If only the `2DRegionQualityRankingBox` is present in the ISOBMFF file, the value of `shape_type` should be set to 1 if the projection format is ERP or equal to 0 if the projection format is equal to CMP. Besides, `center_azimuth`, `center_elevation`, `center_tilt`, `hor_range`, and `ver_range` of the SSR descriptor are computed based on `left_offset`, `right_offset`, `top_offset` and `bottom_offset` values of the `2DRegionQualityRankingBox`, the `RegionwisePackingBox` and the sample location derivation as defined in section 5.2.1 of [OMAFSODIS].

In addition, each of the C Adaptation Sets may contain a `Preselection` descriptor as a Supplemental Property descriptor in order to signal which of the $N \times M \times R$ Adaptation Sets are linked to the corresponding extractor track (main media component). In addition, the C Adaptation Sets contain an Essential Property descriptor indicating that the video is packed and contains region-wise packing information.

Since the generated streams have aligned segments and subsegments, all $N \times M \times R + C$ Adaptation Sets contain the same unsigned integer value for `@segmentAlignment` and `@subsegmentAlignment`.

For VoD services, it is recommended that the content is offered at the MPD using the ISO Base Media File Format On-Demand profile: `urn:mpeg:dash:profile:isoff-on-demand:2011` profile.

If low latency considerations are considered and encodings are performed with various random-access points configurations, it is recommended that the content is offered at the MPD using the ISO-Base Media File Format Broadcast TV profile: `urn:mpeg:dash:profile:isoff-broadcast:2015` profile.

NOTE: The OMAF HEVC Viewport-Dependent baseline media profile typically requires a low delay operation and fast switching. This requires frequent stream access points (e.g., lower than 1 second interval) to be available, which can be achieved by providing different representations with different `switching@interval` values or with 'sidx' boxes having different `starts_with_SAP` values for each of the subsegments.

5.5.5 Security

The guidelines presented in section 4.3 apply in their entirety to Service Providers in the delivery of VR360 content.

[Ed. (DP) relates to AP3 of VRIF2017.154.00: To be updated with input from Security TF. All tiles within a frame are encrypted with the same key using the 'CBCS' subsample encryption scheme as defined in section 9.5 in CENC specification [CENC].

5.6 Guidelines for Service Platform Developers

5.6.1 Overview

Figure 14 shows the basic receiver for one media component in OMAF. In case of DASH streaming, it is considered that the DASH client is part of the application and a conforming OMAF media stream is handed to the file format parser, potentially as a result from an adaptive streaming process and from a concatenation of DASH Segments/Subsegments or CMAF Fragments.

The OMAF conforming media stream is processed by the file format parser. Rendering metadata that is present in the media stream as defined for this profile is extracted and forwarded to the texture to sphere mapping. The function generates the described 3D/spherical signal. The sensor viewpoint information is then used to generate the actually rendered view. OMAF primarily describes the metadata to translate from the decoder output texture information to a 3D/spherical video.

The elementary media stream is decoded by the media decoder. The elementary stream contains the equivalent rendering metadata and may be used instead of the file format metadata as the information is available on both layers.

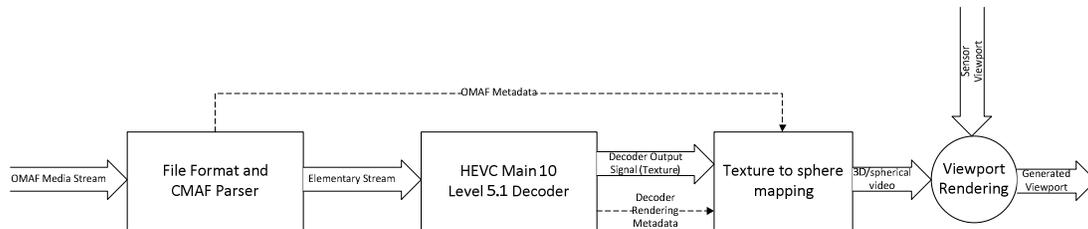


Figure 14: Logical Receiver Model

OMAF media profiles define requirements for the possible presence or absence of such rendering information in the file format and/or the in the elementary streams as well as the required receiver capabilities to process the elementary stream.

A more detailed flowchart of the logical steps of rendering chain is provided in Figure 15. The available metadata in the OMAF stream or the SEI messages is summarized as follows:

- Region-wise packing information as constraint in 7.3.1.2 in OMAF [OMAFSODIS]
- Frame-packing information
- Projection format parameters with content rotation

Based on this information a spherical content can be recovered for each eye with the information of the viewport the viewport can be generated dynamically.

Editors Note: HOW TO USE THE SEI MESSAGE METADATA FOR RENDERING

Note that basically all processing is independent of the viewport. At the same time, implementations may take into account the viewport for optimized performance to reduce processing load and power consumption.

Furthermore, implementation may combine several of the logical steps to reduce memory consumption and processing time compared to a naïve implementation of the rendering chain. For instance, it is viable to generate a polygon mesh according to the rotation of the pre-rotated content and set texture coordinates for each polygon according to the region-wise and frame-packed content at once.

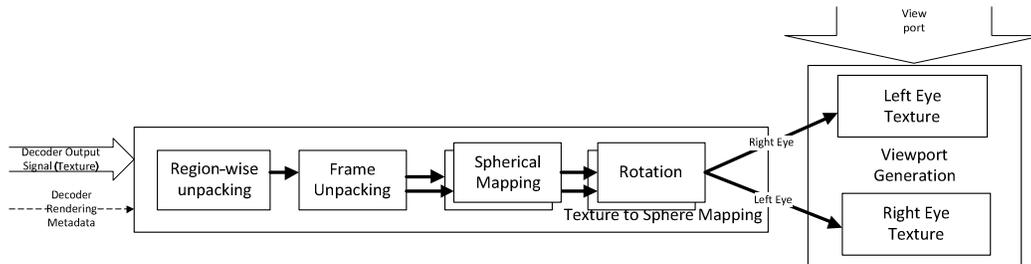


Figure 15: Rendering and viewport generation

5.6.2 Rendering Process based on SEI messages

Editors Note: add more details

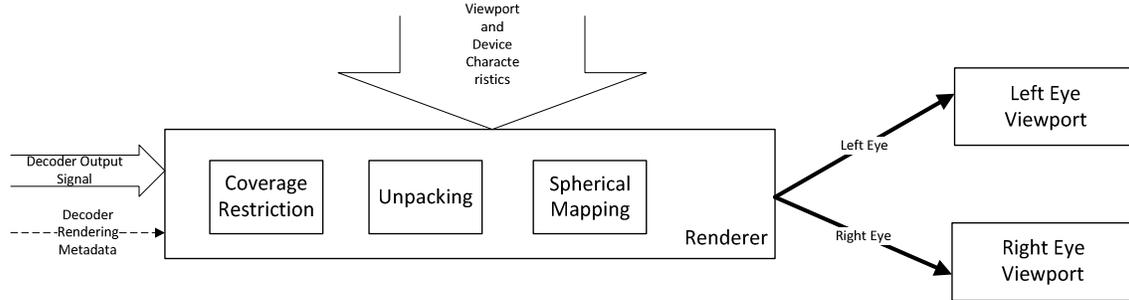
Based on the service provider guidelines in section 5.5, the elementary stream may include SEI message that permit the 2D texture mapping to the spherical coordinate - spatially aligned – for left and right eye. In [UPDSEI] section 3.41.5, the sample location remapping process is documented.

For the restricted scheme based for the viewport-independent media profile, to remap colour sample locations of a region-wise packed picture to a unit sphere, the following ordered (also indicated in Figure 15) steps are applied:

- If a `region_wise_packing` SEI message is present a region-wise packed picture is obtained as the cropped output picture by decoding a coded picture. For purposes of interpretation of chroma samples, the input to the indicated remapping process is the set of decoded sample values after applying an (unspecified) upsampling conversion process to the 4:4:4 color sampling format as necessary when `chroma_format_idc` is equal to 1 (4:2:0 chroma format) or 2 (4:2:2 chroma format). This (unspecified) upsampling process should account for the relative positioning relationship between the luma and chroma samples as indicated by `chroma_sample_loc_type_top_field` and `chroma_sample_loc_type_bottom_field`, when present.
- Furthermore, the sample locations of the region-wise packed picture are mapped to sample locations of the respective projected picture as specified in [UPDSEI] section D.3.41.5.4. Note that this is a 1:1 mapping for the viewport-independent profile
- If frame packing is indicated, the sample locations of the projected picture are converted to sample locations of the respective constituent picture of the projected picture, as specified in [UPDSEI] section D.3.41.5.6. Otherwise, the constituent picture of the projected picture is identical to the projected picture.
- The sample locations of a constituent picture the projected picture are converted to sphere coordinates relative to the local coordinate axes, as specified in [UPDSEI] section D.3.41.5.2.

If rotation is indicated, the sphere coordinates relative to the local coordinate axes are converted to sphere coordinates relative to the global coordinate axes, as specified in [UPDSEI] section D.3.41.5.3. Otherwise, the global coordinate axes are identical to the local coordinate axes.

A more detailed description of the rendering process is **for further study**.



5.6.3 Distribution and Delivery

5.6.3.1 CDN Considerations

If the OMAF HEVC Tile Based Viewport-Dependent Profile is used, the amount of resources that are accessed via HTTP transactions is increased. Therefore, it is recommended to use HTTP/2.0 and ISO/IEC 23009-6 [DASHPUSH]. If both are supported at CDNs and servers a more efficient network performance can be achieved.

5.6.4 Decoding and Rendering

Editors Note: This section is to be completed

5.6.5 APIs

Editors Note: This section is to be completed

5.6.6 Security

Editors Note: This section is to be completed

5.7 Guidelines for App Developers

App developers that attempt to playback content provided according to this profile on service platform are expected to have an OMAF metadata functionality included as well as APIs available to service platform. The OMAF metadata and the scheme restrictions may either be handled by the application by parsing and processing the OMAF metadata, or the application instructs the service platform to use the included SEI messages for proper rendering.

The app developer has two options:

- It checks if the rendering platform supports the usage of the SEI message. If the case, the rendering may be deferred to the rendering platform. A proper capability mechanism is **tbd**.
- If the rendering platform does not support the functionalities, the app developer may interpret the OMAF metadata to map the 2D texture output to the sphere.

Typically, the app developer needs the following functions, either from the service platform with proper APIs, or integrated into the app.

- A DASH client, unless the DASH client is part of the application, (type 3), possibly with configuration APIs to supported optimized playback and rendering.
- OMAF metadata functionality to parse and extract the relevant information or at the minimum to instruct playback in the media pipeline.
- File format parsing functionality for video playback
- Decryption module, if the content is encrypted
- HEVC video decoder to decode the video content
- Rendering and GPU functionalities to generate viewports
- Sensors for viewport tracking

Generally, it is preferable to use HW supported functionalities to optimize speed, latency, power consumptions and overall performance. Each of those above functions may be accessed with APIs. Specific APIs, possibly supported on SDKs and media frameworks are currently under development for example in Khronos, CTA WAVE or W3C.

More detailed API documentation is [for further study](#).

5.7.1 Distribution and delivery

5.7.2 Decoding and Rendering

If the OMAF HEVC Viewport-Dependent baseline profile is used, an HEVC Main 10 profile, Main tier, Level 5.1 capable decoder is needed. In addition, specific Metadata needs to be present to perform the inverse projection/rendering function on the receiver side. This metadata is either carried as OMAF Metadata in Fileformat signaling or as Decoder Rendering Metadata as SEI messages within the elementary stream.

If the OMAF HEVC Viewport-Dependent baseline media profile is used, an HEVC Main 10 profile, Main tier, Level 5.1 capable decoder is needed. In addition, specific Metadata needs to be present to perform the inverse projection/rendering function on the receiver side. This metadata is carried as OMAF Metadata in Fileformat signaling and as Decoder Rendering Metadata as SEI messages within the elementary stream as detailed below in 5.7.3. Depending on the implementation of the renderer either one or the other metadata can be used.

5.7.3 APIs

Figure 16 shows an OMAF-DASH Client model for illustration.

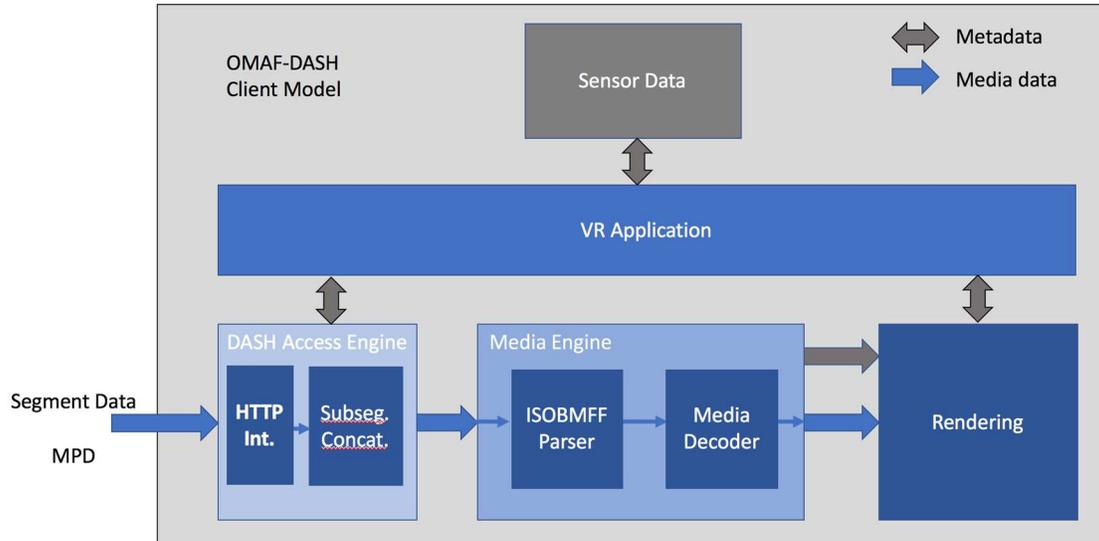


Figure 16: OMAF-DASH Client model with interfaces

In the following each sub module and the associated interfaces are described:

DASH Access Engine

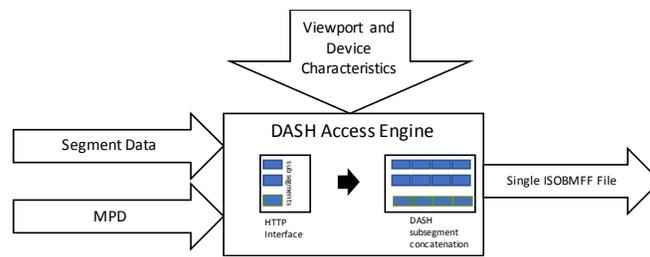


Figure 17: DASH Access Engine for OMAF HEVC Viewport-Dependent baseline profile

The DASH Access Engine is responsible for downloading of all OMAF media streams as well as for the fetching of the DASH manifest file. All network related aspects with respect to HTTP streaming (HTTP version, ISO/IEC 23009-6 [DASHPUSH] aspects, rate adaptation, buffering, etc.) are considered inside this module.

DASH Access Engine is connected to a VR Application module, which constantly provides the information on the selected viewport-dependent adaptation sets for the next requests.

When @dependencyId is used, the initialization segment of the Representation corresponding to 'hvc2' and subsegments of the dependent Representations and complementary Representations in the order as indicated by @dependencyId and increasing presentation order are concatenated (i.e. as specified in section 5.3.5.1 of ISO/IEC 23009-1 [DASH]). When Preselection is used, the initialization segment of the Representation corresponding to 'hvc2' can be concatenated with subsegments of the component of the Preselection in any order. The results leads to an ISOBMFF file conforming to the constraints defined in section 10.1.3.3 of OMAF [OMAFSODIS] which correspond to the OMAF HEVC Viewport-Dependent baseline media profile.

Since each of the frames consist of multiple tiles that together depict the whole covered scene (with an quality/resolution emphasis on a selectable viewport), obviously, all tiles need to be received before the frame can be decoded. Therefore, the subsegments of all tiles and subsegments corresponding of an

extractor track are concatenated. It is also important to mention that the bitrate of each of the tiles is much lower than that of the entire video stream, so downloading a segment for each of the tiles is comparable to downloading the equivalent sub portion of a segment when using the independent profile.

Media Engine

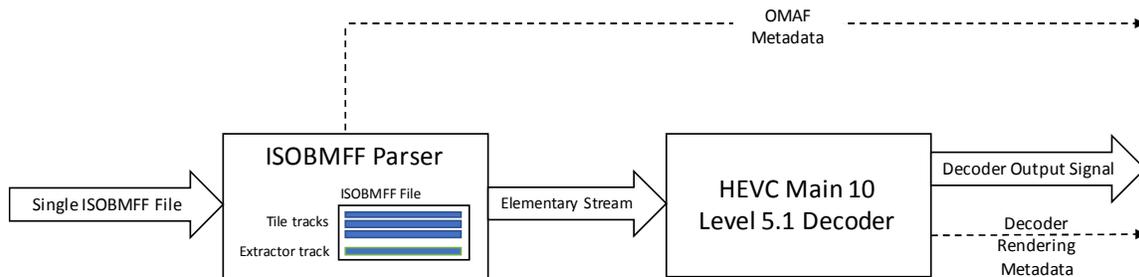


Figure 18: DASH Media Engine for OMAF HEVC Viewport-Dependent baseline profile

The input to the Media Engine is a single ISO BMFF file, which is the output of the DASH access engine (see Figure 17). By usage of HEVC tiles and ISO BMFF Extractor functionality, the ISO BMFF Parser generates a single HEVC Main-10 level 5.1 elementary stream compliant to video decoders available in the market. The ISO BMFF Extractor is specified in [NAL] and contains all aggregation rules for reconstruction of a single HEVC Main 10 level 5.1 elementary stream from multiple MCTS tracks. A detailed overview of the Extractor functionality specified in [NAL] and how they are resolved is given in Annex B.

The output signal, i.e. the decoded picture, is then rendered using either the metadata contained in the File format container or using the SEI messages contained in the video elementary streams.

Renderer

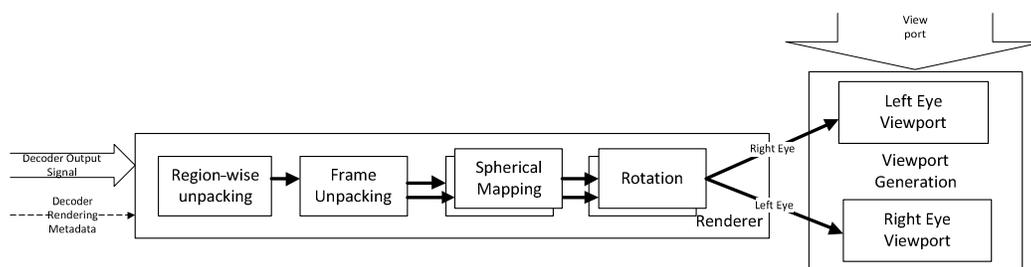


Figure 19: Inverse Projection/Renderer

A more detailed flowchart of the logical steps of rendering chain is provided in Figure 19. The available metadata in the OMAF stream or the SEI messages is summarized as follows:

- Region-wise packing (“Indicates the packing format of the content”):
 - Decoder Rendering Metadata:
 - Region-wise packing SEI
 - OMAF Metadata:

- RegionwisePackingBox
- Stereo frame packing (“Indicates the packing format for stereoscopic content”)
 - Decoder Rendering Metadata:
 - Frame packing arrangement SEI
 - OMAF Metadata:
 - StereoVideoBox
- Projection mapping (“Indicates the projection format”)
 - Decoder Rendering Metadata:
 - Equirectangular Projection SEI
 - Cubemap Projection SEI
 - OMAF Metadata:
 - ProjectionFormatBox
- Content pre-rotation (“Indicates if content is pre-rotated”)
 - Decoder Rendering Metadata:
 - Equirectangular Projection SEI
 - Cubemap Projection SEI
 - OMAF Metadata:
 - Projection orientation box
- Coverage restriction (“Indicates content coverage”)
 - Decoder Rendering Metadata:
 - (Equirectangular Projection SEI or Cubemap Projection SEI) and RegionWisePacking SEI
 - OMAF Metadata:
 - Global coverage information box and RegionwisePackingBox

Based on this information a spherical content can be generated dynamically for each eye with the information about the desired viewport.

Note that basically all processing at the renderer function as defined in Figure 19, is independent of the viewport. At the same time, implementations may take into account the viewport for optimized performance to reduce processing load and power consumption. In this case, the sensor data or viewport information needs to be available in the Renderer.

Furthermore, implementation may combine several of the logical steps to reduce memory consumption and processing time compared to a naïve implementation of the rendering chain. For instance, it is viable to generate a polygon mesh according to the rotation of the pre-rotated content and set texture coordinates for each polygon according to the region-wise and frame-packed content at once.

5.7.4 Usage of Service Platform

Editors Note: This section is to be completed

Annex A Video Master Format Metadata

Editors Note: RangeTopBottom needs to be reconsidered now that new SEI messages are available

Table 7: Master Format metadata

Name	Description	Data Type	Parameter Space	Comments
Spherical Video				
<i>Spherical</i>	Flag indicating if the video is a spherical video	Boolean	yes	
<i>ProjectionType</i>	Projection type used in the video frames	Enum	equirectangular	
<i>StereoMode</i>	Description of stereoscopic 3D layout	Enum	mono, stereo-left, stereo-right	
<i>Coverage</i>	Coverage Parameters of the video	CoverageType See section 5.5.2		
<i>Rotation</i>	Rotation Parameters	RotationType See section 5.5.2		
<i>InitialView</i>	Initial View	InitialViewType		
2D Video Parameters				
<i>FullWidthPixels</i>	Width of the reference video frame in pixels.	Integer	See section 5.5.2	
<i>FullHeightPixels</i>	Height of the full reference video frame in pixels.	Integer	See section 5.5.2	
<i>FrameRateInFPS</i>	Frame Rate in frames per second	Integer	25, 30, 48, 50, 60, 90, 120	

DRAFT FOR PUBLIC PREVIEW

Name	Description	Data Type	Parameter Space	Comments
<i>CroppingInformation</i>	Cropping Information top, right, bottom, left	Integer, Integer, Integer, Integer	can be any information in pixel See section 5.5.2	
<i>PictureAspectRatio</i>	Picture aspect ratio	Integer : Integer	2:1	
<i>ChromaFormat</i>	Chroma format	Enum	YCbCr	
<i>ColourSampling</i>	Colour sampling format	Enum	4:2:0	
<i>SampleAspectRatio</i>	Sample aspect ratio	Integer : Integer	1:1	
<i>BitDepth</i>	Bit depth	Integer	8, 10	
<i>ColourPrimaries</i>	Colour primaries	Integer, Integer	1, 1: ITU-R BT.709: colour_primaries=1, matrix_coefficients=1 9, 9: ITU-R BT.2100 colour primaries and non-constant luminance matrix coefficients, i.e. colour_primaries=9, matrix_coefficients	
<i>TransferFunction</i>	Transfer function	Integer	1: BT.709 14: SDR BT.2020 16: HDR PQ10	
<i>RangeHorizontal</i>	Horizontal range of the Equirectangular video, centered at the midpoint	Float	0 - 360	Assumed to equal 360 if omitted
<i>RangeTopBottom</i>	Elevation angles representing the upper and lower	Float	(-90 - 90, -90 - 90)	Assumed to equal (90, -90) if omitted
Production Metadata				

Name	Description	Data Type	Parameter Space	Comments
<i>Duration</i>	Duration of the content	xs:duration		
<i>TimeStamp</i>	Epoch Time stamp when first frame was recorded	xs:utc		
<i>DirectorCut</i>	Directors cut data			
<i>Heatmap</i>				
<i>Annotation</i>				
<i>StitchingSoftware</i>	Name and version of stitching software			
<i>Copyright</i>	Copyright information			
<i>License</i>	License information			

DRAFT FOR PUBLIC PREVIEW

Annex B ISO BMFF Extractors (informative)

The single ISO BMFF file contains several tracks (one per tile) with an original format equal to 'hvc1' and a track with original format equal to 'hvc2', which corresponds to the tracks with extractors.

The ISO BMFF parser has to play the track with original format equal to 'hvc2'. When parsing the 'moov' box, the parser finds a non-empty 'tref' box within the 'trak' box with original format 'hvc2'. The 'tref' box contains a list of `track_IDs` that indicate the tracks which the 'hvc2' track depends on.

Then, the ISO BMFF parser parses (single track) fragments until it finds the ones corresponding to the 'hvc2' track. Then, it gets the samples of the 'hvc2' track by accessing the bytes indicated in the 'trun' box within the fragments. A sample of the 'hvc2' track contains extractors, i.e. NAL units that have `nal_unit_type` equal to 49.

When the ISO BMFF parser finds such a NAL unit it needs to resolve the extractor, i.e. it parses the body of the extractor and replaces it with the corresponding data. The data that is replaced can be either data encapsulated within the extractor, i.e. within a construct called inline constructor, or data from other tracks, i.e. data that is referenced at a construct called sample constructor.

When parsing an extractor, a parser might find zero or more inline constructor and zero or more sample constructors. An inline constructor consists of a length field and a data field (of size indicated by the length field). The parser, simply takes the data in the data field and extracts it.

A sample constructor consists of an index pointing to one of the `track_IDs` in the 'tref' box, a sample offset, data offset and a data length field. The sample offset indicates the difference to the decoding time of the sample in the tracks containing the extractors and the sample in the referenced track. This value is zero since in the OMAF HEVC Viewport-Dependent baseline media profile extractors are used for tile aggregation and all have the same decoding time. The data offset indicated how many bytes of the sample of the referenced track are skipped from the start of the sample and data length indicates the number of bytes that are copied after skipping data offset bytes. Note that if the indicated data length + data offset is bigger than the length derived from the 'trun' box of the referenced track, data length is clipped to the end of the sample as indicated in the 'trun' box.

In summary, when a sample constructor is found, the 'trun' box of the track with `track_ID` corresponding to the track reference index is parsed and the sample with same decoding time is searched, then the bytes from the sample following data offset up to data offset + data length (clipping if bigger than the sample size) are extracted to the extracted data.

After resolving all extractors, a valid sample is obtained.

Figure 20 shows how ISO BMFF Extractors are used for a sample. As can be seen in the figure there are N Extractors in an extractor track. The first NALU in the aggregated stream is not modified and therefore the first extractor only contains a sample constructor referencing the whole NALU with the preceding length field in the referenced track. The rest of extractors contain an inline constructor responsible of prepending an appropriate length field and slice header (for the aggregated stream) and a sample constructor that fetches the slice payload from the referenced data.

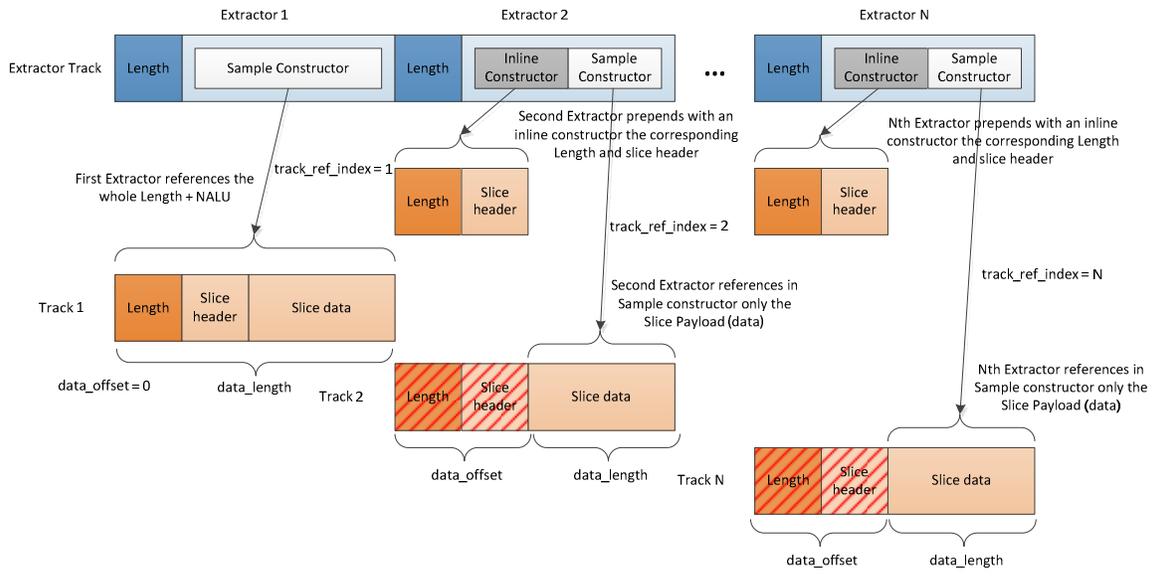


Figure 20: Single ISO BMFF File with one extractor track, N extractors and N MCTS tracks after subsegment concatenation

DRAFT FOR PUBLIC PREVIEW