



GUIDELINES

Version 2.3 2021-01-25

VR Industry Forum

5177 Brandin Court
Fremont, CA 94538, USA
www.vr-if.org

THIS DELIVERABLE IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NON-INFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS 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 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 © 2020 VR Industry Forum. All rights reserved

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 production and 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.....	11
3.1	Definitions	11
3.2	Abbreviations.....	11
4	VR Technologies	13
4.1	Production Guidelines for VR Audio and Video Content.....	13
4.1.1	User Experience	13
4.1.2	Considerations for acquisition and creation	13
4.1.3	Master Format	23
4.2	Media Profiles	25
4.2.1	Introduction.....	25
4.2.2	Selected Media Profiles	25
4.3	Content Security.....	39
4.3.1	Scope.....	40
4.3.2	MovieLabs ECP 1.1 Deltas	41
4.3.3	DRM System Specifications	42
4.3.4	Platform Specifications	42
4.3.5	End-to-End Characteristics	42
4.3.6	Encrypted Media Extensions	43
4.4	Text & Fonts in VR.....	43
4.4.1	Scope.....	43
4.4.2	Guidelines for content creators	43
4.4.3	Font data formats and delivery mechanisms.....	45
4.4.4	Rendering and run-time environment considerations.....	48
4.4.5	Human factors and accessibility considerations	48
5	Vertical 1: OTT Download or Streaming of VR360 Content	56
5.1	Description of Vertical.....	56
5.2	Guiding Example Use Cases.....	56
5.3	Reference Architectures	57
5.3.1	Distribution Architecture.....	57
5.3.2	Client Architecture	59
5.4	Technical Enablers.....	62
5.4.1	Suitable Media Profiles.....	62
5.4.2	Suitable Presentation Profiles	62
5.4.3	Distribution Systems.....	63
5.5	Guidelines for Service Providers.....	68
5.5.1	Suitable Production Formats	68
5.5.2	Sphere-to-Texture Mapping and SEI Message Generation.....	72
5.5.3	Encoding and Content Preparation.....	75
5.5.4	Distribution.....	87
5.5.5	Security	90
5.6	Guidelines for Service Platform Developers	94
5.6.1	Overview	94
5.6.2	Rendering Process based on SEI messages.....	96
5.6.3	Distribution and Delivery.....	96

5.6.4	Security	96
5.7	Guidelines for App Developers	97
5.7.1	Distribution and delivery	97
5.7.2	Decoding and Rendering	97
5.7.3	APIs.....	98
6	Vertical 2: Live VR Streaming.....	101
6.1	Description of Vertical.....	101
6.2	Guiding Example Use Cases.....	101
6.2.1	Actors and set-up/equipment	101
6.2.2	Use case description (short)	101
6.2.3	Use case description (detailed)	101
6.2.4	Consumption of VR content.....	102
6.3	Reference architecture and workflow	103
6.4	Technical enablers.....	111
6.4.1	Production Systems for Live events.....	111
6.4.2	Suitable media profiles.....	120
6.4.3	Suitable presentation profiles	120
6.4.4	Distribution Systems.....	120
6.5	Guidelines for Producers.....	120
6.6	Guidelines for service providers.....	120
6.7	Guidelines for Services Platform developers.....	120
6.8	Guidelines for Application developers	121
Annex A	Video Master Format Metadata.....	122
A.1	Video Metadata.....	122
A.2	XML Schema for VR Video Master Format.....	124
Annex B	ISO BMFF Extractors (informative).....	129
Annex C	Example DASH Manifest for FOV enhanced tiling	131
Annex D	Live VR Deployment Information.....	135

Figures

Figure 1:	Example <i>Coverage</i> metadata for Partial Panorama	23
Figure 2:	Receiver Model for HEVC-based viewport-independent OMAF video profile.....	27
Figure 3:	OMAF-DASH Streaming Client model with interfaces	30
Figure 4:	OMAF-Download Client model with interfaces	31
Figure 5:	Depiction of FOV enhanced viewport dependent video profile.....	33
Figure 6:	LQ (left) and HQ (right)	34
Figure 7:	Example of textual and graphic annotations	45
Figure 8:	Example Architecture for simple VR Streaming Services	58
Figure 9:	Example Architecture for encrypted VR Streaming Services	59
Figure 10:	Client side processing on the example architecture	60

Figure 11: High-level VR software stack.....	60
Figure 11a: Khronos OpenXR Architecture	60
Figure 12: Example VR distribution system.....	64
Figure 13: Example protocol stack for VR content distribution	64
Figure 14: Exemplary DASH configuration setup	66
Figure 15: Latency sources contributing to M2HR latency [CDNOPT]	67
Figure 16: Content Preparation for DASH Distribution.....	77
Figure 17: Example of the packed picture and the respective projected picture of one of the 16 extractor tracks, for a viewing orientation above the equator.....	78
Figure 18: Video content preparation for DASH distribution with HEVC-based viewport-dependent OMAF video profile.....	79
Figure 18a: 6K stereo cubemap tile layout with 2x2 tiles per cube face. Both views are encoded with two qualities / bitrates illustrated with black and red font color.....	85
Figure 18b: Example tile layout of 8K mixed-quality ERP. The content is made available at two qualities / bitrates illustrated with black and red font color.....	86
Figure 18c: 8K effective stereo cubemap with 2x2 tiles per cube face.....	87
Figure 19: Logical Receiver Model	91
Figure 20: Unique A B marked video streams for each viewer	93
Figure 21: Logical Receiver Model	95
Figure 22: Rendering and viewport generation	96
Figure 23: OMAF-DASH Client model with interfaces.....	98
Figure 24: DASH Access Engine for HEVC-based viewport-dependent OMAF video profile.....	98
Figure 25: DASH Media Engine for HEVC-based viewport-dependent OMAF video profile.....	99
Figure 26: Inverse Projection/Renderer	99
Figure 27: Live workflow for Television augmented with VR.....	102
Figure 28: Linear TV and VR Live workflows.....	103
Figure 29: Linear TV and VR Live with - IP connected omnidirectional camera.....	103
Figure 30: Event Action Area: General differences in workflow between TV Linear Live and VR Live.....	104
Figure 31: Event Action Area for VR Live (IP delivered to Game Engine Model)	105
Figure 32: Event OB Facility.....	106
Figure 33: Event OB Facility: differences in workflow between TV Linear Live and VR Live.....	107
Figure 34: Host Facility HQ: differences in workflow between TV Linear Live and VR Live.....	108
Figure 35: Host Facility for VR Live (IP Delivered to Game Engine).....	109
Figure 36: Host Facility HQ	110
Figure 37: Consumer Environment: differences in workflow between TV Linear Live and VR Live.....	111
Figure 38: Two different camera locations: far from and close to the event.....	112
Figure 39: Continuous Live Streaming of VR360-Degrees 10K-Video with MPEG-OMAF and HEVC Tiles..	116
Figure 40: Fraunhofer HHI Omnicam-360	117
Figure 41: Example of a Cubemap projection.....	117
Figure 42: Segmentation of Cubemap projection tiles.....	118
Figure 43: Block diagram of HEVC Tile-Encoder.....	118

Figure 44: Block diagram of the OMAF packager and player.....	119
Figure 45: Single ISO BMFF File with one extractor track, N extractors and N MCTS tracks after subsegment concatenation.....	130

Tables

Table 1: Age related health and safety warnings	17
Table 2: Master File formats	24
Table 3: Maximum achievable resolution in the viewport using viewport-independent baseline media profile for HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with a FOV of 90° × 90° and a given content coverage	26
Table 4: Overview of OMAF media profiles for audio	38
Table 5: MPEG-H Audio MIME parameter according to RFC 6381 and ISO/IEC 23008-3.....	39
Table 6: VR Characteristics.....	40
Table 7: Mapping of SEI Message Information to OMAF Metadata	76
Table 8: Recommended tile layouts.....	77
Table 9: Master Format metadata	122

2 References

- [3DA] Information technology – High efficiency coding and media delivery in heterogeneous environments – Part 3: 3D Audio, Second Edition ([ISO/IEC 23008-3](#))
- [3DTECH] ANSES, [3D technologies and eyesight](#)
- [3GPPR15] [First 5G Release in 3GPP \(Phase 1\)](#)
- [3GPPR16] [Second 5G Release in 3GPP \(Phase 2\)](#)
- [80211AC] 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] 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 Enhancements for High Efficiency WLAN ([IEEE 802.11ax](#))
- [ACN] Ambisonic data exchange formats – Ambisonic Channel Number ([Wikipedia](#))
- [ADM] Audio Definition Model ([ITU-R BS.2076](#)), [BBC Software](#)
- [ADDSEI] HEVC Additional Supplemental Enhancement Information ([ICTVC AC1005](#))
- [AS11X1] MXF Program Contribution - DPP UHD ([AMWA AS-11 X1](#))
- [ATSC30] Advanced Television System Committee, Digital Television Broadcasting ([ATSC 3.0](#))
- [AOA-FAQ] American Optometric Association, [FAQ](#)
- [BS2088] Long-form file format for the international exchange of audio programme materials with metadata ([ITU-R BS.2088](#))
- [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](#))
- [BT709] Parameter values for the HDTV standards for production and international programme exchange ([ITU-R BT.709](#))
- [CDNOPT] TiledMedia, [CDN Optimization for VR Streaming](#)
- [CMAF] Information technology – Coding of audio-visual objects Information technology – Multimedia application format (MPEG-A) – Part 19: Common Media Application Format ([ISO/IEC 23000-19](#))
- [CMAFAMD1] SHVC media profile and additional audio media profiles ([ISO/IEC 23000-19:2018/Amd1:2018](#))
- [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](#))
- [DASHIFIOP] DASH-IF, [Interoperability Points v4.1](#)

- [DASHPUSH] Information technology – Dynamic adaptive streaming over HTTP (DASH) – Part 6: DASH with server push and WebSockets ([ISO/IEC 23009-6](#))
- [DAW1] Qualcomm, [3D Audio Tools](#)
- [DRC] Information technology – MPEG audio technologies – Part 4: Dynamic Range Control ([ISO/IEC 23003-4](#))
- [DSC] Video Electronics Standards Association (VESA) Display Stream Compression (DSC) v1.2
- [EBUR95] EBU, Safe Areas for 16:9 Television Production version 1.1 ([R95](#))
- [ECP] MovieLabs Specification for Enhanced Content Protection ([ECP V1.1](#))
- [EME] W3C, [Encrypted Media Extensions](#)
- [EN303560] ETSI, Digital Video Broadcasting (DVB) TTML subtitling systems ([EN 303560v1.1.1](#))
- [FLUS-GUIDE] 3GPP, Guidelines on the Framework for Live Uplink Streaming (FLUS) ([TR 26.939 v15.1.0 2018-09](#))
- [FLUS-PROT] 3GPP, Uplink Streaming ([TS 26.236 V15.1.0 2018-03](#))
- [FONTCOMP] Information technology – Coding of audio-visual objects -- Part 18: Font compression and streaming ([ISO/IEC 14496-18:2004](#))
- [HDMI] High Definition Multimedia Interface ([HDMI](#)) 2.1
- [IMF2E] Interoperable Master Format — Application #2E ([SMTPE 2067-21:2016](#))
- [HTTP2] Hypertext Transfer Protocol Version 2 ([IETF RFC7540](#))
- [ISOBMFF] Information technology – Coding of audio-visual objects – Part 12: ISO Base Media File Format ([ISO/IEC 14496-12](#))
- [JAUNT360] Jaunt Studios, [The Cinematic VR Field Guide: A Guide to Best Practices for Shooting 360°](#)
- [MP4FF] Information technology – Coding of audio-visual objects – Part 14: The MP4 File Format ([ISO/IEC 14496-14](#))
- [MBMS] Multimedia Broadcast/Multicast Service ([3GPP TS 23.246](#))
- [MVIEW1] I.D.D. Curcio, K. Kammachi-Sreedhar, S. Mate, “Multi-Viewpoint and Overlays in the MPEG OMAF Standard”, ITU Journal: ICT Discoveries, Vol. 3(1), 18 May 2020.
- [MVIEW2] M. M. Hannuksela, Y.-K. Wang, “An overview of Omnidirectional Media Format (OMAF)”, submitted to Proceedings of the IEEE.
- [N16584] MPEG-H 3D Audio Verification Test Report ([MPEG N16584](#))
- [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](#))
- [NORMS] Ambisonic data exchange formats – Normalization ([Wikipedia](#))
- [OFF] Information technology – Coding of audio-visual objects -- Part 22: Open Font Format ([ISO/IEC 14496-22:2015](#))
- [OMAF] Information technology – coded representation of immersive media – Part 2: Omnidirectional media format ([ISO/IEC 23090-2:2019](#))
- [OPENXR] [OpenXR 1.0 Specification](#)

- [OXRGDC] N. Whiting, “[Standardizing All the Realities: A Look at OpenXR](#)”
- [OV2110] SMPTE, Overview Document - Professional Media over Managed IP Networks Roadmap for the 2110 Document Suite ([OV 2110-0:2018](#))
- [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](#))
- [Pulkki] V. Pulkki, “Virtual sound source positioning using vector base amplitude panning” Audio Engineering Society Journal, [vol. 45, no. 6, pp. 456–466](#), 1997.
- [QUICdrafts] Quick UDP Internet Connections ([IETF datatracker](#))
- [R103] Video Signal Tolerance in Digital Television Systems ([EBU R103](#))
- [R128] Loudness Normalization and Permitted Maximum Level of Audio Signals ([EBU R128](#))
- [RFC6381] The 'Codecs' and 'Profiles' Parameters for "Bucket" Media Types ([IETF RFC 6381](#))
- [ST2110-10] SMPTE, Professional Media Over Managed IP Networks: System Timing and Definitions ([ST 2110-10:2017](#))
- [ST2110-20] SMPTE, Professional Media Over Managed IP Networks: Uncompressed Active Video ([ST 2110-20:2017](#))
- [ST2110-21] SMPTE, Professional Media Over Managed IP Networks: Traffic Shaping and Delivery Timing for Video ([ST 2110-21:2017](#))
- [ST2110-30] SMPTE, Professional Media Over Managed IP Networks: PCM Digital Audio ([ST 2110-30:2017](#))
- [ST2110-31] SMPTE, Professional Media Over Managed IP Networks: AES3 Transparent Transport ([ST 2110-31:2017](#))
- [ST2110-40] SMPTE, Professional Media Over Managed IP Networks: SMPTE ST 291-1 Ancillary Data ([ST 2110-40:2017](#))
- [TR26918] Virtual Reality (VR) media services over 3GPP ([3GPP TR26.918](#))
- [TS1015472] Plano-stereoscopic 3DTV; Part 2: Frame Compatible Plano-stereoscopic 3DTV ([ETSI TS 101 572-2](#))
- [TS103285] ETSI, Digital Video Broadcasting (DVB); MPEG-DASH Profile for Transport of ISO BMFF Based DVB Services over IP Based Networks ([TS 103 285 v1.2.1](#))
- [TS26234] Packet Switched Streaming Service ([3GPP TS 26.234](#))
- [TTML] W3C, [Timed Text Markup Language 2](#)
- [USB] Universal Serial Bus Revision 3.2 [Specification](#)
- [VDVS-QUAL] Subjective Quality Results of Viewport-Dependent Omnidirectional Video Streaming ([3GPP S4-170637](#))
- [WiGig] 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 3: Enhancements for Very High Throughput in the 60 GHz Band ([IEEE 802.11ad](#)), Dec. 2012.
- [WOFF1] W3C, [WOFF File Format 1.0](#)

[WOFF2] W3C, WOFF File Format 2.0

3 Conventions and Terminology

3.1 Definitions

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

Term	Definition
360° Content	Post produced or Live 360° video content with limited interactivity
3DOF	3 Degrees of Freedom – a rendering model whereby the viewing pose is only alterable through rotations on the x, y and z axes. These rotations represent roll, pitch and yaw respectively.
6DOF	6 Degrees of Freedom – an extension of the 3DOF rendering model in which translations along the x, y, and z axes are also permitted. These translations represent the forward-backward, left-right and up-down motions respectively.
Equirectangular	Projection of a spherical image in a rectilinear image frame
Haptic Feedback	Palpable feedback, usually servo-driven in human interfaces (hand controllers)
Nadir	Lowest point on a sphere antipodal to the Zenith. More generally, the point (or direction) represented by an elevation of -90° in any spherical coordinate system. Often used to refer to dead-spot in image below camera array. See also "Zenith".
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: Left/Right, 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
Vection	The visual peripheral information able to induce the illusion of self-motion.
Zenith	Highest point on a sphere, antipodal to the Nadir. More generally, the point (or direction) represented by an elevation of +90° in any spherical coordinate system. Often used to refer to dead-spot in image above camera array. See also "Nadir".

3.2 Abbreviations

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

Abbreviation	Definition
ϕ	Azimuth (longitude, increasing eastward)
θ	Elevation (latitude, increasing northward)
AES	Advanced Encryption Standard

BRIR	Binaural Room Impulse Response
CMP	Cube Map Projection
DAW	Digital Audio Workstation
DOF	Degrees of Freedom
ERP	Equirectangular Projection
FOA	First Order Ambisonics
FOV	Field of View
HDCP	High-bandwidth Digital Content Protection
HMD	Head Mounted Display
HOA	Higher Order Ambisonics
HRTF	Head Related Transfer Function
IPD	Interpupillary distance
M2HR	Motion to High Resolution
MCTS	Motion Constrained Tile Set
OB	Outside Broadcast
POV	Point of View
SBA	Scene-Based Audio
SEI	Supplemental Enhancement Information
SMP	Secure Media Pipeline
STB	Set Top Box
VOD	Video On Demand
VR	Virtual Reality

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 fully immersive Head Mounted Display (HMD) or non-immersive device such as a tablet or TV screen should allow the viewer to engage in a unique manner. 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”. 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).

It should be noted that due to the environmental isolation resulting from the unique and immersive nature of HMDs, viewers may experience discomfort when consuming 360° video experiences while in-transit. This discomfort is caused by either (a) the positional sensors in the HMD being affected by changes in the transit system resulting in unstimulated viewport change, or (b) the incoherence between thevection and the vestibular information.

4.1.2 Considerations for acquisition and creation

The following sub-sections provide guidance on certain factors which play a significant role in the creation of high quality VR content that can be viewed in a comfortable manner. Additional technical and artistic consideration can for example be found in [JAUNT360].

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 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 discomfort symptoms characteristic of the sensory sickness.

Sensory 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. During production attention has to be paid to reducing the need for head movement by keeping relevant content in the comfortable viewing area as well as providing some cues as to where any main action is occurring.

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 sensory sickness in the user, and a main source of “vection” is the camera motion.

Viewport changes that are the consequence of user movement (rotations in 3DOF and rotation and translation in 6DOF) could be sensory 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, viewport changes could potentially induce sensory sickness in particular when the vection suggests orientations and accelerations that are not coherent with that which is perceived by the vestibular system. Rotational gain can be used in some situations but should only be applied in a predictable or user specified manner to minimize the risk of sensory sickness.

To reduce the risk of sensory 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 (e.g., 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 too 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)

¹ Method implemented in the VR-game Eagle Flight© from Ubisoft

- Avoid repeated patterns and sharp lines like stripes, blocks, stairs, trees, and 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 an embodiment of the viewer's nose) could partially reduce sensory sickness probably due to a combination of perceptual effect (reduction of the vection due to static elements) and cognitive activity (identification of a 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.
- 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. However, if shooting at higher frame rates, consideration of the target display(s) characteristics should be made.

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. Significant changes in the scene orientation across cuts can cause the viewer to become lost in the experience and may result in physical discomfort while attempting to regain the area of action they were previously following.

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.

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 Sky sponsored survey of 300 people indicated 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

4.1.2.9.1 General

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 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.

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

³ Lightfield displays are expected to limit this issue.

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

- The long-term effects of repeated exposure to stereoscopic content on the development of the human visual systems are still under debate. As of today, no longitudinal studies have been conducted to evaluate such effects and the few statements produced by medical and governmental organizations are in part contradictory. The “French Agency for Food, Environmental and Occupational Health & Safety” for instance state that: “children under the age of 6 should not be exposed to 3D technologies” and “children under the age of 13 should only use 3D technologies in moderation, and that both they and their parents should be vigilant concerning any resulting symptoms” [3DTECH] while the “American Optometric Association” state that “By the age of 3 years most children will have binocular vision well enough established to enjoy viewing 3D television, movies or games” [AOA-FAQ].
- 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.

Stereo imagery is one of the key tools to enhance the experience of immersion.

- The experience should be 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.

4.1.2.9.2 Additional Health and Safety Issues

Most of Virtual Reality HMD producers include in their “health and safety warning” some statements concerning the risks of use of these devices by children. Examples of such warnings are presented in Table 1.

Table 1: Age related health and safety warnings

Device	Statement	Source
Playstation VR	The VR headset is not for use by children under age 12 .	https://www.playstation.com/en-us/network/legal/health-warnings/
Oculus Gear	The Gear VR should not be used by children under the age of 13 , as young children are in a critical period in visual development.	https://scontent-cdg2-1.xx.fbcdn.net/v/t39.2365-6/17640357_1698999383446748_1803373359325511680_n.pdf?oh=10d89b617e1c5f894bd8117bc3470ed8&oe=5A82F8CB

Oculus Rift	This product should not be used by children under the age of 13 , as the headset is not sized for children and improper sizing can lead to discomfort or health effects, and younger children are in a critical period in visual development	https://scontent-cdg2-1.xx.fbcdn.net/v/t39.2365-6/19896829_771660013013643_4087250127671001088_n.pdf?oh=642cb259720ceb661a181142929621d1&oe=5A8836EA
Google Daydream	Daydream View should not be used by children under the age of 13 .	https://support.google.com/daydream/answer/7185037?hl=uk
HTC vive	The product was not designed to be used by children. Do not leave the product within the reach of young children or allow them to use or play with it. They could hurt themselves or others, or could accidentally damage the product.	http://dl4.htc.com/vive/safty_guide/91H02887-08M%20Rev.A.PDF?_ga=2.37961125.787580076.1508146654-1752841129.1507818191

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 4.1.2.9
- Cuts should be consistent with the action and story.
 - Consideration of the viewer motion
 - Use of production techniques such as emphasized lighting or sound to direct the viewpoint
 - Consistent story telling
 - Audio consistency with the 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 combination of these (where two or more formats are used to represent the spherical sound field).

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 mono and stereo, horizontal-only (5.1) to immersive 7.1+4H and 22.2.

For faithful reproduction over headphones, a common methodology to use virtual loudspeakers and the corresponding set of HRTF/BRIRs that is relevant for a certain head position of the listener relative to the loudspeaker positions. When the head rotates, the sound field is rotated by updating the set of HRTF/BRIR. This approach requires an accurate and high-resolution set of HRTF/BRIRs available for all possible head locations relative to the loudspeaker positions as well as careful spatio-temporal interpolation when updating the HRTF/BRIRs 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 (mixed) into a summary element or stem while others (such as Dialog) can be carried as discrete objects.

For live capture and distribution of object-based audio, all audio objects must be tagged with the correct metadata (e.g. location, diegetic, diffusion, width, etc.) Dynamic audio objects require time-varying metadata and real-time tracking.

For post-production, an audio engineer typically pans and adjusts the objects in a 3D scene according to the video scene.

A pure object-based representation can require a multitude of individual audio tracks and their associated time-varying metadata. Typical cinematic content involves the use of several simultaneous objects. The bandwidth necessary for transmitting a sound scene depends on the number of simultaneous objects present at any point in time. Due to limitations in bandwidth for streaming or broadcast services, the high number of objects used in cinematic content typically must be reduced through object grouping (mixing).

Objects can be either individually binauralized using one discrete HRTF convolution process per object, or rendered (e.g., using Vector Base Amplitude Panning [Pulkki]) to a set of virtual loudspeakers which is then binauralized using one HRTF convolution process per virtual loudspeaker (as discussed in the channels section). While the first method results in the best possible rendering quality, the complexity increases

significantly with the number of objects. Alternative techniques trade off quality with complexity in different ways.

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, SBA is agnostic to the loudspeaker configuration. For loudspeaker playback, the renderer adjusts to 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 sound field is described by only the lowest four spherical harmonic coefficients. Higher Order Ambisonics provides a more accurate sound field representation by using additional spherical harmonic coefficients beyond the lowest four. As the number of spherical harmonic coefficients increases, so does the accuracy of the spatial audio representation. The number of coefficients for full 3D is $(N+1)^2$, where N is the ambisonics order.

SBA provides an efficient and accurate representation of the sound field with a limited number of coefficients. The accuracy of the sound field is dependent only on the Ambisonics order. Moreover, for a given Ambisonics order, the bandwidth is not a function of the number of sound sources in the scene.

The spherical harmonic based representation enables a matrix-based sound field rotation that is efficient and smooth, allowing for compelling VR experiences. The rotation operation does not increase complexity as it is typically integrated in the rendering operation.

Depending on the implementation technique, the computational complexity of the rendering increases with the HOA order. 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 sound fields can be achieved using compact microphone arrays.

A single scene-based audio representation can also represent any number of discrete audio elements without requiring additional metadata or increased bandwidth. In this approach, the audio scene is constructed or augmented by panning pre-recorded or live captured audio elements into the 3D space. The same benefits and limitations of Object-based audio panning, described above, applies here.

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 scene-based 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 sound field 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, sound field captured through microphones should be rotated in accordance with camera rotation.
- If the camera is to be translated (not just rotated), then attaching sound field 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 Object- and Channel-based production, unintentional audio crosstalk between (spot) microphones (e.g., due to proximity or reverberation) should be avoided.
- 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].
- 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 [R128] and the US CALM act. VR production requirements are subtly different in that it is not known which direction and position of the listener is not known at production time, meaning that these recommendations, while useful, may not have the same amount of benefit as traditional TV viewing.

- 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 sound field rotation with head tracking. A motion to sound latency lower than 30ms is recommended.
 - When mixing over headphones, an ideal binaural audio experience is achieved when the headphone feeds are created using Head Related Transfer Functions (HRTFs) and headphone equalization. It is recommended that HRTFs closely matching those of the mixing engineer be used. I
 - 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 reacts (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.

4.1.2.16 HDR content

High dynamic range (HDR) content is characterized by a wider color gamut and greater contrast ratio (between darkness and brightness). The richer colors and contrasts contribute to make the image more realistic and enhance the user experience.

While HDR contents could enhance the quality of the experience the following aspects have to be considered, particularly when the content is presented via HMD.

- The wider contrast and the enhanced peaks of luminosity that are induced in the HDR images could increase the flicker visibility in the HMD and consequently diminish the quality of the user experience. More information about flickering can be found in sections 4.1.2.3 and 4.1.2.4.
- Greater contrast and vivid images will induce a stronger stimulation of the human visual system and, depending on the content presented, a stronger vection with the associated risk of sensory sickness. More information about the effects of vection and the best practice to minimize the risk of sensory sickness were presented in the sections 4.1.2.1 and 4.1.2.3.

4.1.3 Master Format

4.1.3.1 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 [OMAF]
- 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 without padding, 2:1 Aspect Ratio with the following attributes:
 - No Zenith or Nadir Blind-spots are permitted.
- For partial 360° video
 - The **Coverage** metadata element as defined in Table 10 is included to describe the partial region.



Figure 1: Example **Coverage** metadata for Partial Panorama

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 Resolution

- Stereoscopic format: Separate Left/Right eye files.
- Metadata as depicted in Table 10.

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: $((\text{Coverage.AzimuthMax} - \text{Coverage.AzimuthMin}) / 360) \times 4096$
- Vertical minimum: $((\text{Coverage.ElevationMax} - \text{Coverage.ElevationMin}) / 180) \times 2048$

4.1.3.2.2 Video Metadata

Video Metadata to be included with video content is depicted in Table 10.

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

The monoscopic video frame rates are also acceptable for stereoscopic video, but preferable Progressive Frame Rates for stereoscopic video are:

- 50
- 60
- 75
- 90
- 100
- 120

A higher frame rate in the case of stereoscopic content is recommended to make the process of binocular fusion more comfortable and reduce visual fatigue.

4.1.3.2.4 File formats

Table 2: Master File formats

Bit Depth	10-bit
Color Sampling	4:2:2
Color Space	ITU Recommendation BT.709 [BT709] (gamut levels within the threshold defined by EBU R103 [R103]) ITU Recommendation BT.2100 [BT2100]
Scan	Progressive Frame
Delivery Format – Option 1	MXF Program Contribution (AMWA AS-11X1 [AS11X1] as per DPP specification)
Delivery Format – Option 2	IMF Application 2e [IMF2E] / JPEG2000 minimum data rate 150Mb/s
Delivery Format – Option 3	Avid DNxHR HQ
Delivery Format – Option 4	Apple ProRes422HQ

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] and 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 Profiles

Since early 2016, MPEG has worked on a project known as Omnidirectional Media Format (OMAF) which was published as an International Standard (IS) in January 2019 as ISO/IEC 23090-2 [OMAF]. OMAF includes two ways of representing an omnidirectional scene in video pictures: (i) a classical “equirectangular” projection like what has been used historically for maps of the globe, and (ii) a mapping of the scene onto the faces of a cube. It supports signaling of the metadata required for interoperable rendering of 360° monoscopic and stereoscopic audio-visual data, and provides a selection of audio-visual encoding formats for this application. It also includes technologies to arrange video pixel data in numerous ways to improve compression efficiency and reduce the size of video, a major bottleneck for VR applications and services.

4.2.1 Introduction

In the following sections, the VRIF media profiles for video and audio are presented. These media profiles are aimed to provide interoperability points for media codecs and associated metadata as well as media coding and encapsulation configurations that may be used for rendering, compression, streaming, and playback of the omnidirectional media content.

4.2.2 Selected Media Profiles

4.2.2.1 Video

4.2.2.1.1 Overview

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

1. HEVC-based viewport-independent OMAF video profile, further described in section 4.2.2.1.2
2. HEVC-based viewport-dependent OMAF video profile, further described in section 4.2.2.1.3

3. HEVC-based FOV enhanced video profile combining aspects of the HEVC based viewport-independent and viewport-dependent video profiles, further described in section 4.2.2.1.4
4. Unconstrained HEVC-based viewport-independent OMAF video profile, further described in section 4.2.2.1.5
5. Unconstrained HEVC-based simple tiling OMAF video profile, further described in section 4.2.2.1.6

For the HEVC viewport-independent OMAF video profile, the maximum achievable viewport resolution is constrained by the video decoder capabilities, specified by the elementary stream constraints in the video profiles with HEVC Main 10 Profile, Main Tier, Level 5.1. For the unconstrained HEVC viewport-independent OMAF video profile, there is no constraint on the HEVC bitstream level and therefore higher viewport resolutions can be achieved since this profile detaches the HEVC Main Level 5.1 codec requirement on the HEVC bitstream from the OMAF video profile definition. For instance, 8K content can be distributed using this profile with elementary streams conforming to HEVC Main 10 Profile, Main Tier, Levels 6, 6.1 or 6.2, assuming the corresponding video decoder capabilities are in place in the receiving device.

The following table shows the maximum viewport resolutions that can be achieved with the viewport independent baseline media profile for HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with an FOV of 90°×90° and a given content coverage.

Table 3: Maximum achievable resolution in the viewport using viewport-independent baseline media profile for HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with a 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

The HEVC-based viewport-dependent OMAF video profile allows for achieving higher resolutions of the viewport compared to the entries of Table 3 given the same capabilities (HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with an FOV of 90°×90°). These 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 download and streaming case, the HEVC-based viewport-dependent OMAF video profile can be used if resolutions higher than those achievable by the viewport independent baseline media profile are desired. In the streaming case, the HEVC-based viewport-dependent OMAF video profile can additionally be used to achieve bandwidth savings. However, viewport dependent streaming comes with additional latency requirements, as described in detail in section 5.4.3.

As described in section 5.4.3.1 in more detail, the download use case is considered to be very attractive for VR services. In such a case, on-demand VR content may be included in a single ISOBMFF file which is downloaded before playback. The structure of that file depends on the media profile used.

The HEVC-based FOV enhanced video profile allows for achieving a higher resolutions of the viewport compared to the entries of Table 3 given the same capabilities (HEVC Main 10 Profile, Main Tier, Level 5.1, 60fps, for a display with an FOV of 90°×90°). High quality tiles are used to construct the visual aspects of

the current field of view while a low quality panorama is always available when rapid changes are made to the viewport.

4.2.2.1.2 HEVC-based viewport-independent OMAF video profile

4.2.2.1.2.1 Introduction

This media profile is defined in ISO/IEC 23090-2 [OMAF] clause 10.1.2 and fulfils basic requirements to support omnidirectional video. Both monoscopic and stereoscopic spherical videos up to 360° are supported. The profile does neither require viewport dependent decoding nor viewpoint 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 [OMAF] 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 (i.e. the field of view of the rendering device) as well as possibly other metadata such as initial viewport information. Whereas 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.

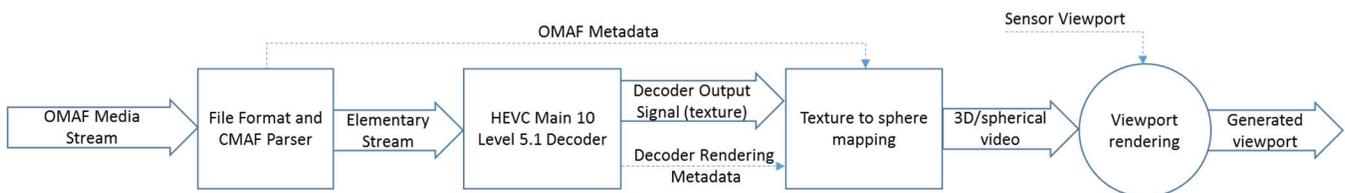


Figure 2: Receiver Model for HEVC-based viewport-independent OMAF video profile

The projection in this profile is restricted exclusively to EquiRectangular Projection (ERP), but permits delivery of less than full 360° spherical video by encoding predominantly only the pixels that are in coverage and therefore may allow to increase the spatial resolution of the covered area under the video decoders profile and level constraints.

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

For this media profile, video elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clause 10.1.2.2. In particular, SEI messages describing the omnidirectional video as defined in [ADDSEI] need to be present.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clauses 10.1.2.3 and 10.1.2.4. ISO BMFF files that contain such an encoded track are identified by the brand 'hevi'. The OMAF metadata is equivalent to the information that is present in the SEI messages for omnidirectional video as specified in ISO/IEC 23090-2 [OMAF] clause 10.1.2.3.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 section B.1.1. An Adaptation Set including Representations formatted according to this media profile is recommended to be signaled as

- @codecs='resv.podv+erpv.hvc1.1.6.L93.B0'
- @mimeType='video/mp4 profiles="hevi"'
- 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 could be played back. For additional information, the Supplemental Descriptor may be 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.

4.2.2.1.3 HEVC-based viewport-dependent OMAF video 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° 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 HEVC-based viewport-independent OMAF video profile. However, the profile requires commonly used content generation and distribution systems to be supplemented with fast stream switching techniques.

This profile allows streams to have a different quality or resolution for different areas/regions of the omnidirectional video each of them corresponding to a preferred viewport (the one to which the area(s)/region(s) with highest quality correspond to). When using this profile, different options are available:

1. Areas are encoded with higher or lower quality/fidelity: For example, the quantization step of transform coefficients is adapted spatially in such a way that the visual quality for the regions differ
2. Areas are encoded with different resolution: some areas are downscaled from their original resolution to a lower one.

For each of the two options listed above, there are two possible instantiations:

1. All areas of the omnidirectional video are offered in a single stream/track/Representation
2. Each of the areas of the omnidirectional video is offered using a separate stream/track/Representation

Note: Even for the latter case where each of the areas is offered as separate stream/track/Representation sufficient provisions are taken into account such that all areas corresponding to the whole omnidirectional video can be consumed using a single video decoder conforming to HEVC Main 10 Profile, Main tier, Level 5.1.

Regardless of which of the above-described instantiations is used, the following metadata is applicable:

1. Region-wise quality ranking [OMAF], enabling to indicate a relative quality order of regions. This metadata is applicable for both the single-resolution and multi-resolution viewport-dependent content.
2. Region-wise packing [OMAF], which provides a region-wise mapping between packed pictures and projected pictures. This metadata is typically unnecessary for single-resolution viewport-dependent content, but needed and essential when regions are coded with different resolutions.

In the configuration provided within these guidelines, only the second option is documented, i.e. for which each of the areas are offered in a separate track. . This is achieved by using HEVC tiles corresponding to different resolutions. These tiles of different resolutions can be combined such that the union of the decoded texture signals covers the entire omnidirectional video. Furthermore, the streams are provisioned such that the whole omnidirectional video can be decoded using a single video decoder conforming to HEVC Main 10 Profile, Main tier, Level 5.1 with a resolution/quality emphasis on the actual user viewport. This allows higher resolutions to be displayed in the viewport as compared to the viewport independent media profile.

With the configuration provided within these guidelines, when using 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 so they do not reference other tiles. In addition, this profile includes ISO Base Media File Format [ISOBMFF] tools that allows the receiver for easily generating a single HEVC stream that conforms to HEVC Main 10 Profile, Main tier, Level 5.1 and can therefore be decoded by a single HEVC Main 10 Profile, Main tier, Level 5.1 capable decoder. The resulting decoded texture signal needs to be dynamically processed in order to generate the actual viewport.

Minimum receiver capabilities required to support this profile:

- HEVC Main 10 Profile, Main tier, Level 5.1
- ISO BMFF extractors, as defined in ISO/IEC 14496-15 [NAL] specification
- Region Wise Packing, as defined in ISO/IEC 23090-2 [OMAF] specification

Note: The HEVC-based viewport-dependent OMAF video profile allows considerable freedom in the usage of Region Wise Packing with up to 255 separate regions. Likewise, the HEVC specification allows up to 110 tiles in Level 5.1 of HEVC Main 10 profile. These guidelines provide appropriate configurations that keep the number of regions for Region Wise Packing at a reasonable value.

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.

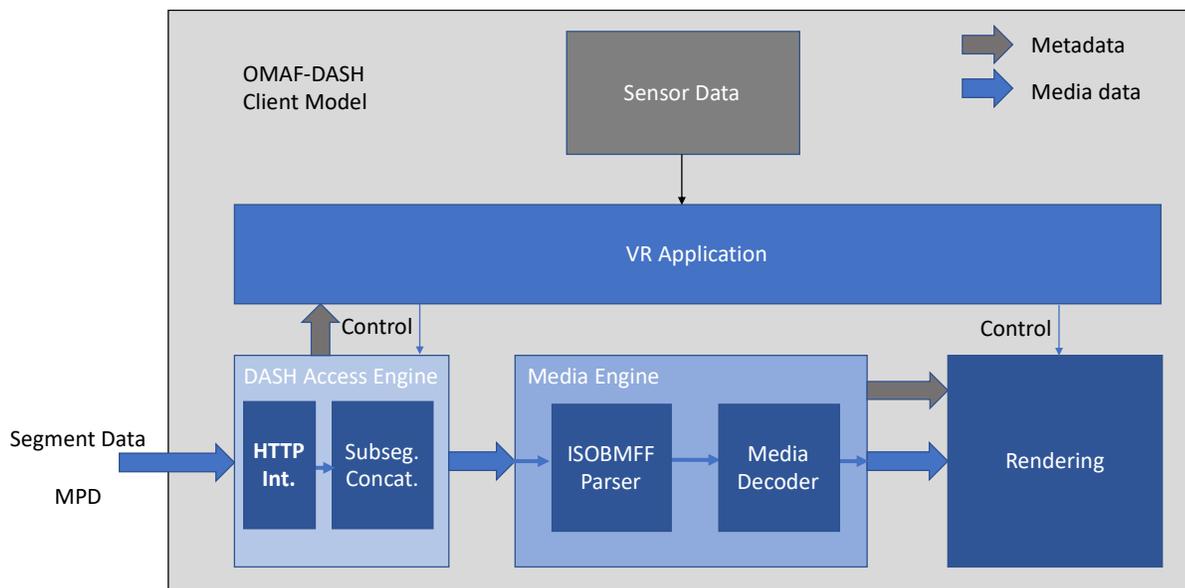


Figure 3: OMAF-DASH Streaming Client model with interfaces

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 which itself outputs decoded pictures and rendering metadata into to the Renderer.

Renderer: Renders the decoded video pictures to the HMD display taking into account sensor data as well as the field of view of the device.

Sensor Data: User positioning information (e.g. viewport position, direction, speed, etc.) taken e.g. from the HMD, e.g., by using OpenXR APIs.

Note: Generally, it is preferable to use hardware 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.

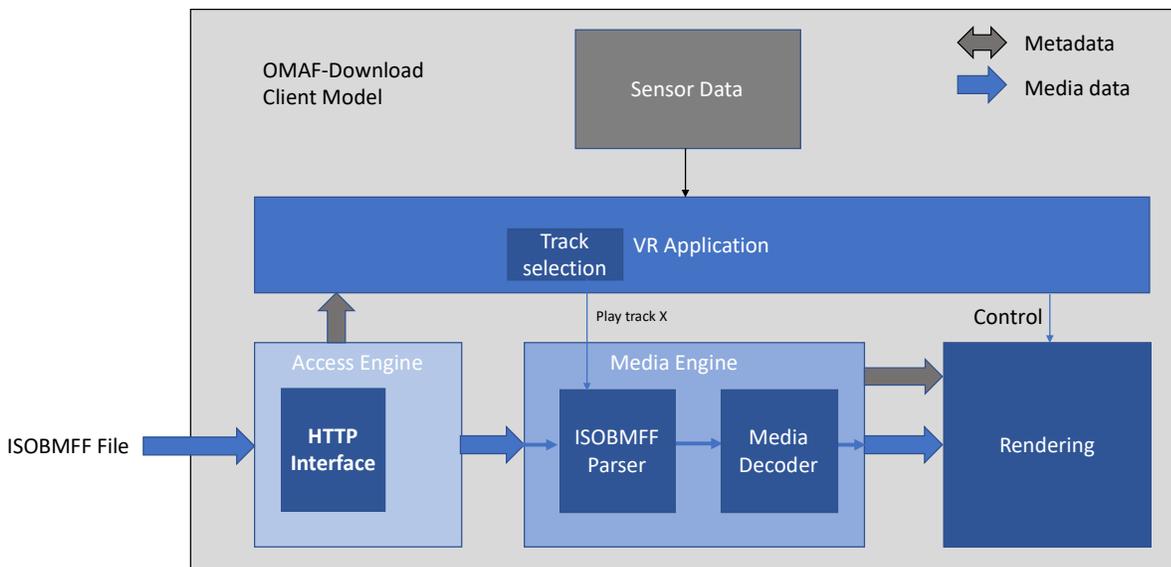


Figure 4: OMAF-Download Client model with interfaces

The main difference to the OMAF-DASH client for streaming as depicted in Figure 2 is the fact that the access engine does not 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 contain a track selection module that determines, based on the metadata provided by the access engine and the sensor data, what extractor track to decode (see control arrow labelled as “Play track X” in Figure 4). The extractor track is thereby dynamically selected 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 `RegionwisePackagingBox`, `SphereRegionQualityRankingBox` or the `2DRegionQualityRankingBox`, has to be parsed in the track selection module by the application in order to select the corresponding extractor track, based on the current user viewport. This requires the Access Engine to provide the VR Application with the relevant metadata, which means that the movie box ‘moov’ is passed to the VR Application or equivalent metadata extracted from it.

The entire ISO BMFF File is passed through to the ISO BMFF Parser, where only the selected extractor track is parsed and the resulting single HEVC Main10 Main Tier Level 5.1 elementary stream is passed to the Media Decoder. Then the stream is decoded and decoded pictures are forwarded to the renderer along with 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 [OMAF] clause 10.1.3.2.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clause 10.1.3.3. ISO BMFF files that contain such an encoded track are identified by the brand 'hevd'. The OMAF metadata is equivalent to the information that is present in the SEI messages for omnidirectional video as specified in ISO/IEC 23090-2 [OMAF] clause 10.1.2.3.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAF] annex B.1.2. An Adaptation Set including Representations formatted according to this media profile is recommended to be signaled as

- either (when corresponding to the Main Adaptation Set if Preselection is used or the Adaptation Set contains dependent Representation with @dependencyId)
 - @codecs='resv.podv+ercm.hvc2.1.6.L93.B0'
 - @mimeType='video/mp4 profiles="hevd"'
 - A Supplemental Descriptor or Essential Descriptor providing the frame packing arrangement may be used
- or (otherwise)
 - @codecs='resv.podv+erpv.hvc1.1.6.L93.B0' or 'resv.podv+ercm.hvc1.1.6.L93.B0'
 - @mimeType='video/mp4 profiles="hevd"'
 - A Supplemental Descriptor or Essential Descriptor providing the frame packing arrangement may be used

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 of [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.

4.2.2.1.4 HEVC-based FOV enhanced video profile

4.2.2.1.4.1 Introduction

This media profile fulfils the requirements to support the delivery of 360° video. This profile supports quality emphasis on the actual user viewport, which allows higher fidelity in the viewport and an overall reduced bitrate compared to the HEVC-based viewport-independent OMAF video profile.

This profile specifies two temporally aligned sets of video streams both of which represent the full spherical scene. One stream (referred to as “LQ”) is generally coded in a low quality format, such as 1080p or 2K, and provides the overall background. The second stream is coded in a higher quality format either at the source resolution or down-sampled to a resolution that more closely matches the best possible spherical view of the terminal device. This second stream (referred to as “HQ”) is also coded in a tiled format, where each tile represents a different rectangular area of the spherical view.

This profile typically makes use of the Equirectangular projection format, but could also be applied to a Cubemap based representation of the spherical view.

Spatial and temporal alignment of the two video streams is coordinated through signaling provided in the DASH manifest [DASH]. Figure 5 depicts the overall flow of the media elements from the live capture source to various terminals.

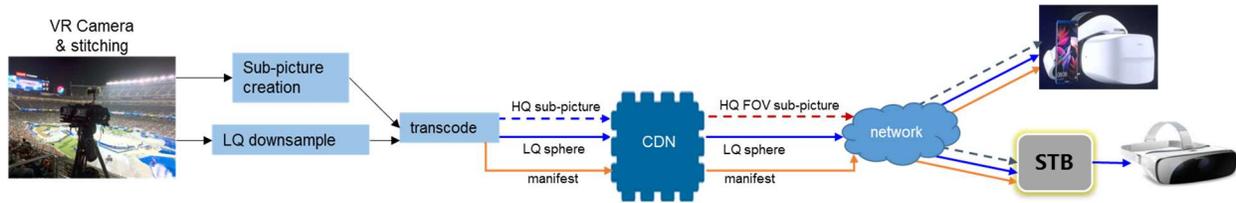


Figure 5: Depiction of FOV enhanced viewport dependent video profile

Each terminal receives the same LQ video stream and manifest, and selects the necessary HQ tiles in order to render the current viewport. In some scenarios, the LQ video stream could be distributed using multicast technologies in order to improve network traffic efficiency.

4.2.2.1.4.2 Content Preparation

The LQ spherical and HQ FOV sub-picture video streams are derived from a single content source and encoded as follows:

LQ Sphere

An equirectangular projection coded according to the HEVC-based viewport-independent OMAF video profile as defined in section 4.2.2.1.2.2.

HQ Tiles

A tile based equirectangular projection coded according to the HEVC-based viewport-dependent OMAF video profile as defined in section 4.2.2.1.3.3.

DASH Manifest

The LQ spherical video is described in an adaptation set that includes a Spatial Relationship Description (SRD) as defined in Annex H of ISO/IEC 23009-1 [DASH] and a Content Coverage (CC) descriptor as defined in ISO/IEC 23090-2 [OMAF] clause 8.3.4 indicating full 360×180° coverage, i.e.

```
<EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,0,1920,960,3840,1920" />
```

This property indicates that the actual resolution of the video (1920×960) should be scaled to a higher resolution (3840×1920).

```
<SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
  <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="0" centre_elevation="0" azimuth_range="23592960"
      elevation_range="11796480" />
  </omaf:cc>
</SupplementalProperty>
```

The content coverage property denotes that the coverage information is depicted as using two azimuth circles and two elevation circles, no view IDC is specified and that the representation is monoscopic.

The coverage info property indicates that the coverage for the adaptation set is 360° (azimuth_range="23592960") × 180° (elevation_range="11796480")

Each HQ tile is described in a separate adaptation set that includes an SRD indicating its area of coverage overlapping the LQ spherical video and a CC descriptor, i.e.

```
<EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,960,640,960,320,3840,1920,1" />
```

This property indicates the position of the upper left pixel (960, 640) and resolution of the tile (960×320) is part of a larger frame with a resolution is 3840×1920. Tiles have the same spatial_set_id (the last value in the property).

```
<SupplementalProperty schemeldUri="urn:mpeg:mpegI:omaf:2017:cc">
  <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="2449120" centre_elevation="983040"
      azimuth_range="5898240" elevation_range="1966080" />
  </omaf:cc>
</SupplementalProperty>
```

The content coverage property denotes that the coverage information is depicted as using two azimuth circles and two elevation circles, no view IDC is specified and that the representation is monoscopic.

The coverage info property indicates that the coverage for the adaptation set is 90° (azimuth_range="5898240") × 30° (elevation_range="1966080") and is centered 45° left and 15° up (these properties denote the tile labelled t04 in Figure 6).

Example

Consider a 4K (3840×1920) stitched equirectangular panorama which is expected to be streamed to a device with a per-eye resolution of 1280×1440 pixels and a field of view of 110×60°. This stream would be encoded as a single tile HD resolution video stream and 14 tiles as depicted in Figure X2.

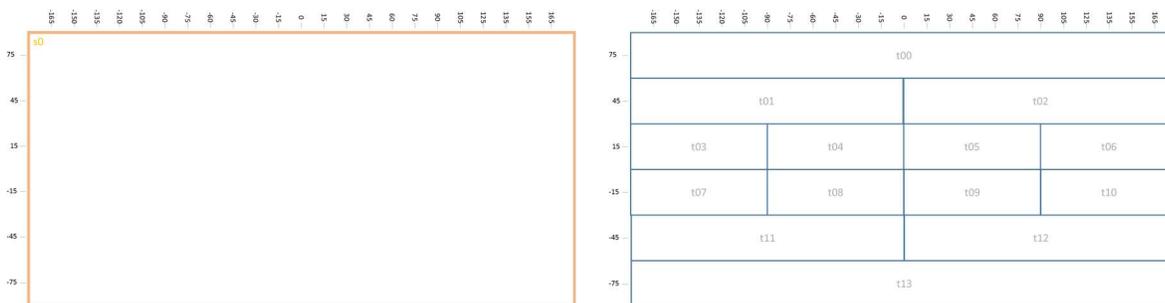


Figure 6: LQ (left) and HQ (right)

LQ has either a single tile (s0) of either (a) a resolution of 960x540 (qHD), (b) a resolution of 1920x1080 (Full HD) or (c) a resolution of 3840x1920 but with a quality that equates to a very low bitrate

Tiles in HQ (t00-t13) have resolutions of 3840x320, 1920x320 or 960x320. A complete DASH manifest for this example is provided in Annex C.

4.2.2.1.4.3 Streaming Client model

4.2.2.1.4.4 External specifications

4.2.2.1.4.5 Quality and Performance

In this profile, that part of the LQ video stream which is overwritten by the HQ FOV may be considered as redundant, this is an overhead that is traded off against any algorithmic and network complexity in selecting tiled areas for both the HQ viewport and LQ non-viewport areas. Motion-to-high-quality

performance can be controlled by a number of factors, including the segment and GOP sizes and forward retrieval of tiles based on current and predicted changes in the viewport.

4.2.2.1.5 Unconstrained HEVC-based viewport-independent OMAF video profile

4.2.2.1.5.1 Introduction

This profile provides the same set of features as the HEVC-based viewport-independent OMAF video profile described in 4.2.2.1.2. The only difference is that there is no constraint on the HEVC bitstream level. Therefore, service providers can use this profile with a variety of resolutions and not only with 4K resolutions enabling, for instance, content being prepared and decoded at a 8K resolution or larger, i.e., with HEVC Level 6, 6.1, or 6.2.

In particular, the unconstrained HEVC-based viewport-independent OMAF video profile as presented in this section may be used together with elementary streams conforming to the following HEVC codec profiles to distribute 8K content at various frame rates, assuming the corresponding video decoder capabilities are in place in the receiving device:

- HEVC Main 10 Profile, Main Tier, Level 6 for 8K spatial resolutions at 30fps
- HEVC Main 10 Profile, Main Tier, Level 6.1 for 8K spatial resolutions at 60fps
- HEVC Main 10 Profile, Main Tier, Level 6.2 for 8K spatial resolutions at 120fps

The content author signals the corresponding level of the bitstream by setting the appropriate value in the `general_level_idc` of the HEVC bitstream. This value of `general_level_idc` is additionally available in the ISO/BMFF encapsulation of the bitstream in the `HEVCDecoderConfigurationRecord` structure of `HEVCConfigurationBox` contained in the HEVC sample entry definition specified in ISO/IEC 14496-15. Lastly, the `@codecs` attribute in the DASH MPD exposes in a single string the level indication among other information.

Both monoscopic and stereoscopic spherical video up to 360° are supported.

4.2.2.1.5.2 External specifications

Video elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clauses 10.5.1.2.

ISO BMFF tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clauses 10.5.1.3 and 10.5.1.4. ISO BMFF files that contain such an encoded track are identified by the brand 'uhvi'. The OMAF metadata is equivalent to the information that is present in the SEI messages for omnidirectional video as specified in ISO/IEC 23090-2 [OMAF] clause 10.1.2.3 (HEVC-based viewport-independent OMAF video profile described in 4.2.2.1.2.).

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAF] section B.1.3. An Adaptation Set including Representations formatted according to this media profile is recommended to be signaled as

- `@codecs='resv.podv+erpv.hvc1.2.6'`
- `@mimeType='video/mp4 profiles="uhvi"'`
- A Supplemental Descriptor or Essential Descriptor providing the frame packing arrangement may be used.

As an example, if HEVC Level 6.1 is used for 8K resolution content, then may be set as follows:

- `@codecs='resv.podv+erpv.hvc1.2.6.L183.B0'`

4.2.2.1.5.3 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. At the time of these tests on the tested equipment, 3GPP TR26.918 [TR26918], 4K spatial resolution with ERP provided a fair quality but new VR equipment with better optics will highlight 4K resolution limits and show a clear advantage of the 8K resolution to bring a higher sense of reality.

4.2.2.1.6 Simple tiling OMAF video profile

4.2.2.1.6.1 Introduction

This media profile extends the capabilities of the HEVC-based viewport-dependent OMAF video profile as follows:

1. No HEVC Level limit is specified, which enables content to be prepared for devices having up to 8K decoding capacity (HEVC Level 6, 6.1, or 6.2).
2. Client-driven selection of OMAF tile tracks to be received, while enabling simple merging of the selected OMAF tile tracks to a single HEVC bitstream.
3. Possibility for requesting content at a granularity that is finer than (Sub)segments in order to reach a lower motion-to-high-quality latency.

Both monoscopic and stereoscopic spherical video up to 360° are supported.

This profile allows streams to have a different quality or resolution for different tiles of the omnidirectional video, wherein the term tile refers to a slice with a rectangular shape, which depends only on the collocated slice in reference pictures. When using this profile, different options are available:

1. Tiles are encoded with higher or lower quality/fidelity: For example, the quantization step of transform coefficients is adapted spatially in such a way that the visual quality for the regions differ
2. Tiles are encoded with different resolution: some areas are downsampled from their original resolution to a lower one.

Tile-based streaming enables players to select tiles adaptively and dynamically. In general, players select higher quality/resolution tiles to cover the viewport and lower quality/resolution tiles for the remaining regions of the spherical video. Consequently, tile-based streaming enables higher resolution in the viewport and/or reduced bitrate compared to viewport-independent spherical video with uniform quality and resolution.

Each coded tile sequence is stored as a separate track, referred to as an OMAF tile track, and offered as a separate Representation for streaming.

One or more OMAF base tracks are provided in addition to the OMAF tile tracks. An OMAF base track provides instructions to reconstruct a single video bitstream by merging samples of the referenced OMAF tile tracks. An OMAF base track can either be an HEVC tile base track or an extractor track. OMAF base tracks can be prepared for different device capabilities, such as for different decoding capacity and for monoscopic and stereoscopic displays.

Regardless of which of the above-described instantiations is used, the following metadata is applicable:

1. Region-wise quality ranking [OMAF], enabling to indicate a relative quality order of tiles. This metadata is applicable for both the single-resolution and multi-resolution viewport-dependent content.

2. Region-wise packing [OMAF], which provides a mapping of tiles to projected pictures. This metadata is typically unnecessary for single-resolution viewport-dependent content, but needed and essential when tiles are coded with different resolutions.

The possibility to achieve low motion-to-high-quality latency is enabled by separating the metadata and media data to different downloadable entities, referred to as Tile Index Segments and Tile Data Segments. Tile Index Segments contain the file metadata (most notably, MovieFragmentBoxes) for available OMAF tile tracks, and Tile Data Segments only contain the media data. By parsing a Tile Index Segment, an OMAF player can request media data at a granularity down to one picture interval. Thus, it is theoretically possible to react to viewing orientation changes even at one picture interval latency.

In addition to the receiver capabilities required HEVC-based viewport dependent OMAF video profile, the capability of deriving the region-wise packing of the OMAF base track from that of the client-selected OMAF tile tracks and the support for Tile Index and Tile Data Segment formats are required.

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

The simple tiling OMAF video profile enables two client models:

1. The client model described for the HEVC-based viewport-dependent OMAF video profile in section **Error! Reference source not found.**. This model enables fetching content at (Sub)segment granularity.
2. A client model where the content is fetched at a finer granularity than (Sub)segments.

In this model, an ISOBMFF Parser is a part of the DASH Access Engine or operationally coupled with the DASH Access Engine. The ISOBMFF Parser provides information on random access points, which can be used as switching points from one OMAF tile track to another, and hence as points to respond to viewing orientation changes. Moreover, the ISOBMFF Parser provides byte ranges within Tile Data Segments so that the DASH Access Engine can issue HTTP GET requests with byte ranges that could e.g. cover the minimum random access interval.

Furthermore, the DASH Access Engine is operationally connected with a bitstream rewriting functionality that merges the received tile data to an HEVC bitstream. An OMAF player can choose to implement either of the following

- i. An OMAF base track suitable for the client's purpose, such as decoding capacity, is selected and fetched. An ISOBMFF Parser parses an OMAF base track to create an HEVC bitstream from the received tile data.
- ii. An OMAF player may choose to operate without receiving and parsing any OMAF base track. In this case, some HEVC high-level syntax structures, such as sequence parameter sets and slice segment headers, are rewritten to create an HEVC bitstream from the received tile data.

When the OMAF tile tracks are encrypted, constraints as described in clause 5.5.5.1.2 apply.

4.2.2.1.6.3 External specifications

Video elementary streams are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clause 10.1.7.2.

ISO BMFF tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clause 10.1.7.3. ISO BMFF files that contain such an encoded track are identified by the brand 'siti'.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAF] annex B.1.4.

4.2.2.1.6.4 Quality and Performance

The quality and performance observations for the HEVC-based viewport-dependent OMAF video profile as described in section **Error! Reference source not found.** generally apply to the simple tiling OMAF video profile too, which however additionally also provides the following capability:

1. Greater flexibility to support up to 8K decoding capacity (HEVC Level 6, 6.1, or 6.2).
2. Possibility for requesting content at a granularity that is finer than (Sub)segments in order to reach a lower motion-to-high-quality latency.

4.2.2.2 Audio

4.2.2.2.1 Overview

This section defines media profiles for audio. Table 4 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 4: Overview of OMAF media profiles for audio

Media Profile	Codec	Profile	Level	Max. Sampling Rate	Brand	Section
3D Audio Baseline	MPEG-H 3D 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° 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 [OMAF] clause 10.2.2.2.

ISO BMFF Tracks are encoded following the requirements in ISO/IEC 23090-2 [OMAF] clause 10.2.2.3. The ISO BMFF track should be identified by the brand 'oab1'.

DASH Integration is provided following the requirements and recommendations in ISO/IEC 23090-2 [OMAF] clause 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 5.

Table 5: 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 3D Audio LC Profile Level 1	audio/mp4	mhm1.0x0B	'oab1'	ISO/IEC 23008-3
MPEG-H 3D Audio LC Profile Level 2	audio/mp4	mhm1.0x0C	'oab1'	ISO/IEC 23008-3
MPEG-H 3D Audio LC Profile Level 3	audio/mp4	mhm1.0x0D	'oab1'	ISO/IEC 23008-3
MPEG-H 3D Audio LC Profile Level 1, multi-stream	audio/mp4	mhm2.0x0B	'oab1'	ISO/IEC 23008-3
MPEG-H 3D Audio LC Profile Level 2, multi-stream	audio/mp4	mhm2.0x0C	'oab1'	ISO/IEC 23008-3
MPEG-H 3D 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 with the security requirements of the entire ecosystem is not practical.

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 6.

Table 6: 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

The current 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 high-value content (e.g. a movie or TV show) is embedded in a virtual theatre environment. It is a special case because (i) the video source may have its own DRM (e.g. streaming video or Blu-Ray disk), (ii) a compositing operation of high-value content and background is required and (iii) the high-value content is by definition attractive to pirates. It is expected that existing DRM technologies would be used to protect the high-value content separately from any items in the virtual theatre environment.

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

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

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

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.

4.3.2 MovieLabs ECP 1.1 Deltas

The MovieLabs Specification for Enhanced Content Protection (ECP) [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. 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 content, such reconstruction will be much more difficult because of the difficulty of getting all the tiles and dealing with padding nevertheless, Tiling itself is not a measure to avoid copying or pirating content.
- **Ripping Software (Interactive Applications):** Interactive VR application software is similar to 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 distribution environments include 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).
 - In future phases of the guidelines when interactive content is considered, considerations will be given as the VR content is not a simple decode, it may be a composite of 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 Characteristics

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 developed and published by the W3C [EME]. It provides a framework for web browsers to support playback of HTML5 video protected by DRM without the use of plug-ins. EME is the only standardized way to securely play back VR content in an HTML5-based browser but needs extensions to support the presentation of the desired viewport in the secure media pipeline. For non-browser playback (i.e. native applications), existing, widely deployed DRM systems should be used.

4.4 Text & Fonts in VR

4.4.1 Scope

VR and Augmented reality (AR) content may be created to serve a wide variety of users and application domains, including (but not limited to) such fields as entertainment and gaming, live broadcast and streaming services, immersive education, marketing and advertisement, industrial services (such as technical support and equipment maintenance, interactive user guides, etc.), and many more consumer applications. Significant portion of the content representing above-mentioned categories may require presenting textual information in one form or another – common examples of such use cases include subtitles in various languages, messaging applications within games, product descriptions / brand expressions in marketing and advertisement materials, and numerous text labels and instructions included as part of the content for educational and industrial applications. Considering a widely diverse (and, largely, unknown) language support requirements combined with creative needs of content authors – enabling the VR/AR ecosystem to support all world's languages and geographic regions, enabling multiple application domains, and ensuring authors' creative freedom to make their own content design decisions presents its own unique set of challenges that this section of the Guidelines aims to address.

4.4.2 Guidelines for content creators

4.4.2.1 General considerations for text in VR content

Creating an engaging and compelling content for VR/AR environment is a challenging undertaking, one that requires careful consideration of many different factors, including various limitations imposed by the viewing environments (head-mounted displays, flat-screen viewports into 360° video content, etc.), technical limitations of the display devices (varying screen sizes and resolutions, mediocre quality of HMD optical components), and varying platform capabilities. Content design choices made by the authors play a key role in making the users forget about technology limitations, and, when it comes to adoption of the new technology, content is truly the king – virtual reality applications would not have enjoyed wide level of adoption without high-quality VR games, entertainment programming, and educational content! Yet, it is vitally important to ensure that VR/AR ecosystem is also an attractive environment for presentation of advertisement and marketing content, that it is equally capable of supporting various industrial applications and tasks.

When it comes to designing user experience, graphical and textual components are vital parts of the overall content presentation, and we realize that the choice of resources to use as part of the content is solely authors' responsibility. Text is an important part of a presentation, and while a sequence of characters conveys the semantic meaning of the written word, fonts give it voice. Freedom to choose a custom font for content presentation is a highly desirable feature for authors, and their choice of fonts can be influenced by many factors, like the requirements of the content brand identity, typographic needs to ensure proper look and feel of the presentation, precise control over text layout and appearance by utilizing specific font functionality (such as e.g. using variable fonts to control text width and weight, enabling particular stylistic alternates of a selected typeface, or applying shaping features required by a particular writing system), and

supporting specific choice of languages (e.g. for subtitles) in various geographic regions where content is distributed.

However, freedom to choose a custom font for content presentation also imposes certain requirements and additional considerations. Custom fonts cannot be expected to be readily available on a device that would be used to access and view VR/AR content; therefore, the author-selected font resource has to be provided as part of the media, the same way that other components such as images and video/audio resources are made available for presentation on end-user devices. Careful consideration should be given to various technical solutions available for font embedding and delivery (discussed in details in the section 4.4.3), and, depending on the origin and availability of a particular font resource, decision to use a custom font may also require additional licensing considerations, to ensure that the chosen font is properly licensed for a particular content type and is free of any restrictions for a particular choice of technology used to enable font embedding and distribution as part of the content.

It is also important to consider certain limitations of rendering environments when making font choices. Text size, color contrast, typeface design, and a selection of a font style (among many other factors) can all play an important role and have an effect on text readability and legibility. However, in addition to traditional content design considerations, the VR/AR ecosystem imposes its own unique set of challenges – the content presentation environment is always dynamic, controlled by the viewer's positional and directional inputs. VR/AR content is presented in constant motion, creating an environment where authors' ability to control foreground / background objects and their color contrast is often limited. Rendering capabilities and device performance (discussed in more details in section 4.4.4) may also have an effect on content appearance, as dynamic video backgrounds (combined with even a subtle motion of a viewport) may necessitate constant presentation updates and require rendering of content be updated on a per-frame basis. Hardware-accelerated graphics rendering and video decoding in many cases alleviates performance-related concerns, but may require making specific design decisions when it comes to selecting font data formats and font delivery choices.

4.4.2.2 Content layout considerations

Content design for VR requires special attention be given to spatial relationships of different parts of content, their perceptual 3D positioning and size. In addition to traditional aspects of graphics design related to text content, where choice of fonts and text sizes may affect text readability and legibility, 3D content design with graphics overlays requires preserving spatial relationships between different parts. For example, 3D perspective has an effect on perceptual size of objects within a scene – text object using the same font size positioned further away from a viewer will be perceived as an object of much larger size than one placed closer to the foreground of the scene. These particular size relationships in perspective 3D scene have to be taken into account when graphics and text are introduced, and in order to preserve their relative sizes and positional relationships, the actual font size must be adjusted for rendering of dynamic content, especially if the content is moving due to changes in viewport positioning.

Even when VR content is presented on flat screens, the choice of text size may contribute to creation of volumetric effect, where larger size text objects introduced into VR/AR scene would appear be positioned closer to the viewer, and, conversely, smaller text sizes may be perceived as objects located farther away.



Figure 7: Example of textual and graphic annotations

Text legibility can also be affected to a significant degree by the choice of font weight and color contrast between foreground and background objects. To preserve higher levels of color contrast, the use of regular to medium heavy font weights is preferred, and in order to avoid visual dropouts (localized low contrast areas when static text objects are presented against constantly changing background imagery, e.g. a video background) – using an additional layer of separation such as drop shadows or semi-transparent back panel for text is often a good idea. This technique would also help to combat the side effects introduced by the optical components of many HMDs – things like color fringing and spatial distortions introduced by the lens of an HMD device, when content is positioned off center of the viewing axes, would have a much more profound effect on lightweight text at smaller sizes.

It is also worth noting that spatial relationships between synthetic parts of the VR/AR scene (such as text and graphics superimposed on top of the camera image) may play a vital role in reinforcing the virtual 3D perception of the scene when proper layout is achieved; on the other hand, it may be a major source of dissonance in 3D perception if the spatial relationships between text and visual objects are flawed – changes in perspective, incorrect depth positioning, and improper occlusions of the text and graphics components relative to the camera view of other natural objects within a scene (e.g. when a synthetically introduced graphics object appears to be placed behind another object within a scene, but is still visible) may have strong adverse effects on overall perception of VR scene, and, in certain cases, may completely ruin the 3D viewing experiences.

4.4.3 Font data formats and delivery mechanisms

Font data plays a critical role in enabling media presentation systems to support multiple languages of the world, and to enable authors produce content that would be free from any script- or language-specific considerations, and to benefit from the ability of presentation environments to display readable multilingual text utilizing standard OS services and software components.

Note: The ability of modern presentation systems to support any of the world languages and writing systems is based on a basic Unicode principle of separating the semantic load of a text message (i.e., the meaning of a sequence of characters encoded as a text string) from its presentation layer. In other words, a sequence of Unicode code-points encodes the units of text content, while the font encodes the units of text display and provides all necessary components to visualize the text (glyph outlines, shaping and layout rules, etc.). Presentation layer decisions often include the choice of glyphs to use to display text characters, how a text string should be shaped – its directionality (left-to-right or right-to-left), combining multiple characters into a single glyph (as it is often required for syllabic scripts such as Hangul or Devanagari), and many other decisions are based on the different clusters of data that are provided together in a single resource that is generally known as a font file.

It is, therefore, easy to see why font data, and, specifically, font data format used to encapsulate different components of a font in a single file is such an important subject to consider – a properly designed media presentation system would allow authors' freedom in making arbitrary font choices to achieve the desired look and feel of their content, to ensure certain guaranteed level of text legibility, brand identity of a message, as well as other aspects of presentation style, without being concerned whether font changes may affect semantics of a text message and its readability.

In general, there are different types of font data that are needed to support full-featured text presentation environment – among them, there are visual components (such as vector outlines to define font glyphs, or their pre-rendered images that have been traditionally supported by bitmap fonts), their corresponding metric information, and a set of rules for language-specific processing and text shaping. Certain font formats such as one defined by ISO/IEC 14496-22 "Open Font Format" (a.k.a. OpenType®) [OFF] are capable of delivering all types of font data as a single resource able to support language processing, text shaping and composition decisions, and glyph scaling and rendering at run time, while other formats such as bitmaps and/or their functional equivalents (e.g. Signed Distance Field (SDF) data sets) are only capable of supporting glyph rendering functionality and require all layout decisions (such as specific choice of glyphs to use, their exact locations and positioning, etc.) being done at the content creation stage.

The following brief overview of different available font data formats summarizes their features and capabilities with intent to help content authors and implementers to choose one that suites their particular needs best:

- ISO/IEC 14496-22 "Open Font Format" (OFF, also known as OpenType [OFF]) – binary-encoded font data that offers a feature-complete full support for all text shaping, layout and rendering implementations. This font data format is capable of supporting different type of vector outline formats, including TrueType and PostScript outlines, SVG-encoded glyph description for color font support, as well as embedded monochrome and color bitmaps. OFF also supports text shaping and layout decisions that are required by a variety of languages and writing systems. For implementations where supporting dynamic text inputs (such as subtitles in user-selected language, messaging within VR games, or aggregated text content for AR applications) is a requirement, OFF may be the only format capable of delivering full support for all required features mandated by a particular language and/or writing system.
- Web Open Font Format (versions 1 & 2, both developed by the W3C WebFonts Working Group) offer an efficient packaging mechanism for transporting compressed OFF data over IP networks. WOFF is a font-data-specific compression technique, with WOFF version 2 offering significantly higher compression rates. For both WOFF1 and WOFF2 – the payload must be provided in OFF format described above.
- Bitmap fonts offer relatively simple and convenient way of using pre-rendered glyph images (either monochrome or color) to be used for text visualization. While bitmap-based formats are limited in their capabilities to support scalable text and text layout decisions, they can be useful for certain implementations that expect text rendering be hardware accelerated using bitmap data encapsulated as textures. However, since many existing implementations offer support for full featured OFF fonts,

using dedicated bitmap font formats is rarely a necessity as textures can be rendered from OFF vector outlines and pre-loaded at runtime.

- SDF based font data can, to a certain degree, be considered as a variation of bitmap data. However, this format offers significant benefits for programmable, GPU-based implementations for hardware-accelerated text rendering. Main differentiating features of SDF-based fonts from their bitmap counterparts is that they offer support for scalable text rendering and perspective transformations, while also allowing easy integration of text objects into generic graphics rendering pipeline (thus allowing application of physics and ray-tracing engines, material effects, textures, etc.). Due to significantly higher CPU load that would be needed to generate SDF data at run time, certain implementations (where dynamic text rendering is not required) may choose to generate SDF data offline and preload it as a standalone graphics engine resource. However, it is also worth noting that there are existing implementations that are perfectly capable of generating SDF at runtime using OFF glyph outlines as input; these implementations would ultimately offer the best functionality covering all needs and use cases – supporting text layout and shaping needs during runtime for scene composition, and offering hardware-accelerated rendering capabilities for best in class text rendering.

Like we mentioned earlier, it is the prerogative of an author to select font resources that suit their content needs best, and the particular choice of fonts is their sole right and responsibility. Since the content is going to be delivered to an *a priori* unknown user device, preserving content's look and feel can only be done by delivering all needed resources and making them available for VR service platforms and applications. When it comes to font data delivery, multiple tools already exist to make it happen:

- For implementation where IP connections are available, WOFF encoding mechanism can be used to deliver fonts to VR device. Both WOFF1 [WOFF1] and WOFF2 [WOFF2] were developed by W3C as webfonts delivery mechanisms, and are widely adopted industry standards. According to the W3C Media Advisory issued in 2018 (<https://www.w3.org/2018/02/media-advisory-woff2-rec.html.en>), development and deployment of WOFF2 has led to an explosive growth of webfonts adoption, with over 70% major websites now using webfonts. WOFF-encoded font data files are not limited to web-only implementations, and can be delivered to any device on IP networks.
- Font data can also be delivered using ISO/IEC 14496-12 "ISO Base Media File Format" [ISOBMFF], as part of the font data stream specified by ISO/IEC 14496-18 "Font compression and streaming" [FONTCOMP]. Therefore, it allows encapsulating font data in the same file as other resources, including video and audio streams.
- For subtitles and other TTML-based streaming text implementations, font data can be embedded as part of the TTML data using a `font` element as defined in clause 9.1.4 of [TTML], which can be used to either define or refer to a font resource.
- Various other mechanisms have been developed to enable downloadable fonts in particular application domains. Some of them are based on the implementations of technologies described earlier in this section – for example, see "DVB: MPEG-DASH Profile for Transport of ISO BMFF Based DVB Services over IP Based Networks", [TS103285] clause 7.2 "Downloadable fonts", while others rely on previously deployed data delivery mechanisms, e.g. "DVB TTML subtitling systems" based on the DVB Transport Streams implementations, see clause 5.3 of [EN303560].

In conclusion, it should be noted that depending on a specific VR platform environment, any font delivery mechanism would work for as long as it can be reliably used to deliver a selected font resource to be used for rendering of text objects as a timed resource within a scene. Since the importance of enabling the usage of custom font resources and their timely delivery to an application platform are well understood, content authors should not be concerned with the details of various technical implementations, and should be encouraged to select font resources that are best for their content needs and requirements.

4.4.4 Rendering and run-time environment considerations

VR/AR content is intended to be viewed in highly dynamic environments where location of a viewer and his positional and directional inputs are used to define and render a particular view. Therefore, it is highly likely that due to constant motion and inputs of a content viewer, the particular scene of VR content would have to be updated on a frame by frame basis, making it a necessity to utilize hardware acceleration and available GPU rendering resources. Special consideration should be given to data formats used for hardware-accelerated graphics and text rendering, and it may also require taking into account specific design decisions made by rendering system implementers.

Prior experiences with hardware-accelerated 3D graphics rendering have shown that high-quality text rendering may be a challenging task, one that requires using custom tools and special-purpose data formats, such as e.g. glyph textures represented as signed distance fields (SDF) for GPU-based text rendering. For additional information on the subject matter, one may wish to consult the original SIGGRAPH 2007 Valve publication (https://steamcdn-a.akamaihd.net/apps/valve/2007/SIGGRAPH2007_AlphaTestedMagnification.pdf) and other resources including various publications and implementations (<https://onlinelibrary.wiley.com/doi/full/10.1111/cgf.13265>, <https://github.com/Chlumsky/msdfgen>).

Therefore, delivery of VR content may sometimes require supporting arbitrary custom data formats - in order to satisfy functional text rendering requirements (see details in the clause 4.4.3) both the original font data, which is required to support text shaping and layout decisions, and a special set of various resources (bitmap textures, SDF) may be needed for hardware-accelerated 3D text rendering. However, these additional resources [bitmaps and SDFs] may often be generated at runtime, doing so would eliminate the need to have them pre-generated and preloaded in advance. Content authors should, therefore, consider particular run-time environment capabilities (VR game engines, graphics rendering capabilities of a particular VR platform, etc.) when making their decisions on font data formats and data delivery requirements.

4.4.5 Human factors and accessibility considerations

4.4.5.1 Typographic recommendations for VR content

Creating content with textual elements for 3D viewing environment presents many additional challenges that authors need to consider when making design and style decisions. Quite often, when introducing type in 3D we tend to think of using novel, unusual 3D designs with extruded text, which, in practice, makes text much less legible and difficult to read. Except of few rare cases, such as representing specific logotype designs, or where text elements are integral parts of the VR environment (e.g., a title of VR game), using traditional 2D type makes text objects more legible and easier to read. However, special considerations should be given to composition and positioning of 2D text objects within virtual 3D scene, please refer to section 4.4.2.2 for details.

Designing typographic content for traditional applications in a physical world is a relatively well-known subject, where specific typography rules exist making it easier to establish a content structure and make text legible and readable. The same rules apply to the virtual 3D worlds – separation of content pieces using color contrast between foreground and background, creating clear content hierarchy using different typefaces, font styles and text size/weights helps to make text an effective communication tool that is legible and easy to read. At the same time, VR viewing environments do present additional challenges, something that authors need to consider. Specifically, in addition to certain aspects of VR content design already discussed in sections 4.4.2 and 4.4.4, considering the following list of recommendations would be helpful:

- Using typefaces and/or font styles to differentiate between different contexts. Examples of different contexts may include using distinct font choices for subtitles when separating background (behind the scene) audio from an audio stream attributed to a particular source or object within the visible VR content, using different typefaces to differentiate between text labels and descriptive parts of AR content, or using type to reinforce the perception of a particular brand or a product as part of marketing and advertisement. However, limiting the number of different typefaces and styles to select few, and using them consistently throughout the presentation helps improve user experience.
- Text placement in virtual 3D world needs to conform to the main scene composition and perspective. Choice of text size for a particular object has to take into account the depth positioning of that object within the scene, and should follow the perceptual relationships established by the main scene content. For example, text object of same physical size would be perceived as an object of much larger size, if it is positioned further away from a viewer. Making sure that text doesn't occlude any foreground objects is also a good idea; occlusions in virtual 3D content help reinforce the perception of VR worlds as 'real', but, at the same time, can easily ruin the 3D viewing experience if the occlusions and depth placements do not correlate with each other.
- Font weight and size also need to be chosen in correlation with each other. Due to specific limitations of virtual 3D content and its rendering environments, lightweight text of smaller size can be very illegible and extremely hard to read, especially if rendering of lightweight strokes of a glyph cause flickering and dropouts.

4.4.5.2 Designing for accessibility

Accessibility of content is an important factor for consideration that should be part of every decision made in various stages of content creation pipeline. When it comes to textual components, the following aspects and design guidelines should be considered:

- Color contrast – regardless of color properties assigned to a foreground text object, in a virtual 3D world the text object (or its portion) may at times be presented against a similarly colored background. This is particularly a concern for VR applications, where the movements within the scene background are controlled by the viewer's positional and directional inputs. We recommend using high-contrast opaque or semi-transparent backgrounds (such as drop shadows, text outlines, back plates) to separate the text against a physical background to improve user's focus and viewing experience.
- Font sizes and weights should be chosen to eliminate or significantly diminish the adverse rendering effects such as text flickering, stroke jitter, and dropouts. With proper contrast and positioning, use of larger text sizes set using regular or medium-weight fonts is recommended.
- Using font sizes and weights as tools to create a particular structure and establish hierarchy of content is highly recommended. The same goes for using different fonts as part of the structure for content separation, and is also recommended; however consider limiting different typeface and font style choices to select few (two to three at the most) for consistency.
- Choice of a typeface is likely to be the most critical decision affecting text legibility. Research has shown that modern, sans serif typefaces with larger x-heights, ample character spacing and open letterform designs (e.g. fonts that belong to humanist style sans serif designs such as "Frutiger") tend to significantly improve legibility of text on display screens.

4.4.5.3 Subtitles in VR content

Similar to widely established practices in traditional 2D applications, the subtitle streams can be used as optional content elements for translation of audio content, or to assist viewers with hearing impairments. The same general guidelines can be applied for use of subtitles services in VR applications. However, a few aspects should be considered when translating TV subtitle guidelines into VR environments

- Safe Area: Guidelines for safe graphic overlay areas for TV environments (such as e.g. EBU R95 [EBUR95]) should not be translated directly into a VR environment. Early user testing has shown that the safe area for textual content in VR should be smaller (i.e. to leave more margin to the viewport) compared to traditional TV format. This recommendation is based on two primary observations:
 1. TV guidelines are based on a fixed picture aspect ratio, where aspect ratio in VR can vary and depends on the rendering device; and
 2. In some end-user devices optical components often produce significant levels of distortion and chromatic aberrations, which cause perceptual image quality to fall off towards the edge of the viewport, which would decrease legibility in these parts of the image.

It is also important to note that VR devices with a very wide field of view should leave larger margins (to the left and right) since the viewport may extend the natural field of view of human vision system.

- Reading speed: Guidelines for reading speed may need to be adapted. However, at this point a specific recommendation cannot be provided.
- Speaker identification: Using different colors and other design choices to differentiate speakers should be considered an important part of the tool palette for subtitles in VR. Using different fonts to separate different contexts (e.g. someone speaking behind the scene vs. a speaker located in the current VR viewport) could also be an effective tool.

The audio can serve dual purpose in VR, and be used as a navigation tool to help guide the viewer through a VR scene tracking the action and offering additional cues for viewer orientation, especially when the action attributed to a particular audio stream is located outside of the current viewport of a VR scene. For people who use subtitles as a substitute for audio, these additional spatial cues are not available – the subtitle service should, therefore, make up for these missing cues as much as possible.

The existing guidelines for subtitles can be used in VR applications as a starting point; however, the different nature of a VR environment is tailored by additional aspects that authors should consider to make VR content more accessible:

- VR is an interactive environment where a viewer can look around a scene, and where viewer's inputs control the content of the VR viewport. Since audio elements can relate to a speaker or a sound that is located outside of viewer's particular viewport, subtitles can be used to help viewers navigate the VR scene by employing specific positional arrangements and/or including additional graphical elements that would help viewers obtain additional directional cues within the scene, and find a speaker or a sound-emitting object located outside of the viewport of a VR scene. It is worth noting that while these particular topics are currently being addressed by a number of researchers and user studies/experiments, particular methods for presenting subtitles in VR have not been established yet.
- Following the preceding item and observing current practices for rendering subtitles in VR, it can be noted that two basic strategies for subtitle rendering exist. Although we cannot give specific recommendations at this point, it may help to review advantages and disadvantages of each approach when designing subtitles for a specific content: One strategy is to attach subtitle rendering to the current viewport, such that subtitles follow the viewer's head movement. The second strategy is to attach subtitles to a video object. In that case each subtitle can also be multiplied and shown at several positions within the video object.

4.5 Multiple Viewpoints

Multiple viewpoints can be thought as a set of 360° cameras which, for example, may be scattered around a basketball field (see Figure 4.5.1). The OMAF 2nd edition specification enables a streaming format with

multiple viewpoints to allow, for example, switching from one viewpoint to another, as done by multi-camera directors for traditional video productions. This allows watching an event or object of interest from a different location and facilitates leveraging the well-established cinematic rules for multi-camera directors that make use of different shot types, such as wide-angles, mid-shots, close-ups, etc.

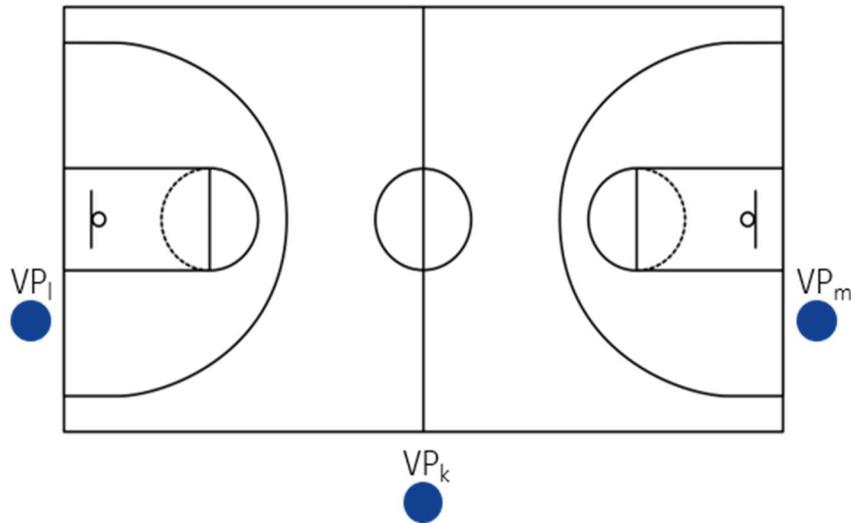


Figure 4.5.1: Example usage of multiple viewpoints in a basketball event [MVIEW1, MVIEW2].

Figure 4.5.1 illustrates an example of OMAF 2nd edition content with multiple viewpoints (VP_k, VP_l and VP_m). This allows the user to experience the action from different perspectives and facilitate interactive content consumption and creative storytelling.

OMAF 2nd edition provides means to signal multiple viewpoints including

- Cartesian coordinate (X, Y, Z) position of the viewpoint,
- GPS position of the viewpoint,
- Orientation of the coordinate system for the viewpoint in relation to the geomagnetic north
- Yaw, pitch, and roll rotation angles of X, Y, and Z axes, respectively, of the global coordinate system of the viewpoint relative to common reference coordinate system,
- Viewpoint group information,
- Viewpoint switching information, including number of possible switching transitions, viewing orientation in destination viewport modes, viewpoint transition effect types, and how timelines are handled when switching viewpoints,
- Viewpoint switch regions, that initiates viewpoint switching when selected by the user with types including a region on the viewport, a sphere region, an overlay,
- Viewpoint looping information,
- Viewpoint entity group information ‘vipo’ listing all tracks and image items for a given viewpoint, with each entity group identified by a certain grouping type and identifier

Viewpoint data can be static or dynamic. In case of dynamic viewpoints, OMAF defines a new timed metadata track to signal the viewpoint data dynamically changing over time. In addition, OMAF also defines an initial viewpoint timed metadata track, that indicates the initial viewpoint that should be used.

It is important to enable scene continuity across viewpoint transitions. In particular, scene continuity within the user’s viewport should be possible while switching from one viewpoint to a new viewpoint. OMAF contains signaling of object center point correspondence (OCPC) in order to enable this kind of

continuity. In particular, OMAF defines a new OCPC timed metadata track, that describes the object center points correspondence between viewpoints for important objects of interest. This enables to maintain the switching continuity between different viewpoints by maintaining the viewport oriented towards the selected object(s) of interest.

Toolset brand 'vwpt' in OMAF specifies viewpoint-related requirements on files and OMAF players, with the ISOBMFF-level mandate that the file contains two or more 'vipo' entity groups. At the player level, when an OMAF player conforming to this toolset brand is given a file containing the brand 'vwpt', the OMAF player is expected to handle viewpoints based on the expected behaviours as specified in the OMAF specification, including expected OMAF player behaviours based on signalling in DASH and those based on signalling in ISOBMFF, addressing viewpoint positioning, viewpoint switching and viewpoint looping.

4.6 Overlays

Overlays are a way to enhance the information content of 360 video. They allow superimposing another piece of content (e.g., a picture, another video with news, advertisements, text or other) to be rendered on top of the main (background) omnidirectional video. Overlays also allow the creation of interactivity points or areas.

OMAF 2nd edition defines four different overlay types, based on their spatial position (see Figure 4.6.1):

- Overlays may be displayed on a rectangular area at a specified position relative to the viewport on the users' viewing screen and always present on the users' viewport (viewport-relative or viewport-locked overlay);
- Overlays could be displayed on a plane at an indicated position within the unit sphere at a depth from the user viewing position (sphere-relative 2D overlay);
- Overlays may be displayed on a sphere surface at an indicated position within or on the unit sphere positioned over the background video without any gap between the two (sphere-relative projected omnidirectional overlay);
- Overlays may be displayed on indicated mesh elements of the 3D mesh at a given location within the unit sphere (3D mesh overlay).

Equator-level cross section

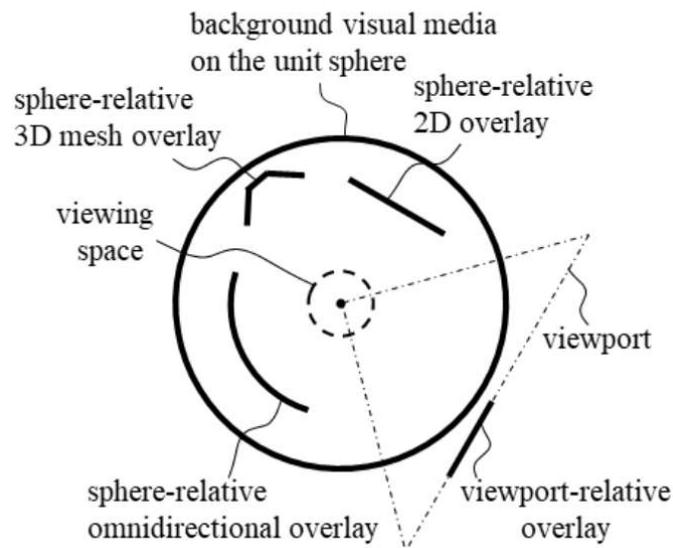


Figure 4.6.1: Different types of overlays defined in OMAF 2nd edition [MVIEW1, MVIEW2].

Below table describes various deployment scenarios for the first three types of overlays described above.

	Deployment scenarios
Viewport-relative 2D overlay	Advertisement Recommended viewport for the content Content owner's logo Legal notice, parental guidance Current score (in a sports event) Tickers (e.g. stock, news)
Sphere-relative 2D overlay	Annotation of 360 content (e.g. player statistics or player info "card") Selectable hotspot Advertisement Closeup or 2D camera view (for a region of interest)
Sphere-relative omnidirectional overlay	Selectable hotspot Semi-transparent "annotation map" overlays, e.g.: <ul style="list-style-type: none"> • highlighting/marking objects on the background • illustrated manufacturing/repairing instructions with phases at different overlays that the user can turn on/off

The following rendering properties may be indicated for an overlay:

- layering order, specifying the layering order among the overlays that are relative to the viewport, and separately among each set of overlays that have the same distance from the centre of the unit sphere
- opacity of the overlay that is applied for all pixels of the overlay source when rendering the overlay
- priority value, based on which a priority order of overlays can be derived and used in the case that an OMAF player does not have enough decoding capacity to decode all overlays
- alpha composition, specifying that the overlay is associated with an alpha plane used for determining pixel-wise opacity and blending when superimposing the overlay

As controls for user interaction with the overlays, OMAF 2nd edition defines signalling of the following flags:

- `change_position_flag`, when set to 1, specifies that users are allowed to move the overlay window to any location on the viewing sphere or the viewport.
- `change_depth_flag`, when set to 1, specifies that the depth of overlay can be chosen by user interaction. When both `change_position_flag` and `change_depth_flag` are set to 1 then the X,Y,Z position of the overlay can be freely chosen by user interaction.
- `switch_on_off_flag`, when set to 1, specifies that the user is allowed to switch ON/OFF the overlay.
- `change_opacity_flag`, when set to 1, specifies that the user is allowed to change the opacity of the overlay.
- `resize_flag`, when set to 1, specifies that the user is allowed to resize the overlay window. The field-of-view of the resized overlay window shall be same as that of original overlay window.
- `rotation_flag`, when set to 1, specifies that the user is allowed to rotate the overlay window to different directions. The field-of-view of the rotated overlay window shall be same as that of original overlay window.
- `crop_flag`, when set to 1 specifies that user is allowed to crop the overlay media window.
- `change_position_flag`, `change_depth_flag`, `switch_on_off_flag`, `change_opacity_flag`, `resize_flag`, `rotation_flag`, or `crop_flag` when set to 0, specifies that the user is disallowed to perform the respective operation on the overlay.

Overlay controls can be static or time-varying. Static overlay controls are contained in a sample entry in a track containing the overlay source or as an item property for an image item containing the overlay source. Time-varying overlay controls are included in a timed metadata track, which also specifies which overlays are active (i.e. turned on, unless otherwise controlled by a user interaction) or inactive (i.e. turned off) at a given time. Toolset brand 'ovly' in OMAF specifies overlay-related requirements on files and OMAF players. At the ISO BMFF level, the following constraints shall be obeyed in files containing the brand 'ovly' among compatible_brands of the FileTypeBox:

- The file shall contain one or more 'ovbg' entity groups.

Note: An 'ovbg' entity group associates overlays with a particular background visual media track or image item. An 'ovbg' entity group lists the tracks and image items that contain associated overlays or background visual media or contain timed metadata for the overlays or background visual media.

- When the `OverlayPriority` control structure is present for an overlay and `overlay_priority` in `OverlayPriority` is equal to 0, `overlay_control_essential_flag[i]` shall be absent or equal to 0 for all values of `i` greater than 13 in the `SingleOverlayStruct` of the overlay.

Note: When `overlay_priority` is equal to 0, OMAF players are required to display the overlay. OMAF 2nd edition specifies overlay controls up to index 13, inclusive. Consequently, this requirement states that whenever the 'ovly' brand is in use and the displaying of an overlay is mandatory, only the overlay controls that are specified in the OMAF 2nd edition are allowed for the overlay.

At the player level, when an OMAF player conforming to this toolset brand is given a file containing the brand 'ovly', the OMAF player is expected to render the background visual media and the overlays of an 'ovbg' entity group as based on the expected behaviours as specified in the OMAF spec, including expected OMAF player behaviours on processing of overlay control structure, sphere-relative overlay rendering procedure, viewport-relative overlay rendering procedure, overlay source region access procedure and handling of user interactions with overlays.

Any encryption applied to the overlay tracks may be different and performed separately from the encryption to the tracks carrying main 360-degree video. Moreover, it is also possible that the main 360-degree video is encrypted while the overlays are not, and vice versa.

Here is an example workflow on advertisements using overlays in regular 360 video:

1. The content creator has a VR asset where there are specific locations (2D areas) where an advertisement can be placed.
2. The VR asset is processed offline to be delivered in DASH format using OMAF.
3. The sponsor of the content creator have 2D content (images, videos, gif, etc)
 - This content is integrated into VR asset using OMAF as a sphere-relative 2D overlay and prepared for DASH streaming.
4. When the overlay is in the viewport, the advertisement content starts to play.
5. Whenever possible: with eye tracking, the sponsor may check if and when the user is viewing the advertisement.
6. Analytics
 - How much time is within Viewport?
 - How much time is it assumed being watch?
 - Click to advertisement?

Here is another example workflow on dynamic insertion of advertisements using overlays in regular 360 video:

1. The content creator has a VR asset where it is desired to place an advertisement dynamically.
2. The VR asset is processed offline to be delivered in DASH format using OMAF.
3. The sponsor of the content creator have 2D content (images, videos, gif, etc)
 - This content is integrated into VR asset using OMAF as a viewport-relative 2D overlay and prepared for DASH streaming.
4. User start to stream the VR asset content and watch.
5. The DASH server reacts based on an activation (e.g., a context aware trigger) to play or place the advertisement content
6. The advertisement starts playing as a viewport-relative 2D overlay while the user is consuming the VR asset.
7. Whenever possible: with eye tracking, the content provider may check if and when the user is viewing the advertisement.
8. Analytics
 - How much time is within Viewport?
 - How much time is it assumed being watch?
 - Click to advertisement?

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 VR 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 × 2k (ERP), 10bit, BT.709, as low as 30fps
 - High-quality: up to 8k × 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 8 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 expected to be synchronized in time and aligned in the 3D 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. Interfaces Da and Dv 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. Operational timed metadata may be added, for example to provide information on a director's cut. Content may be made available such that it is optimized for 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.

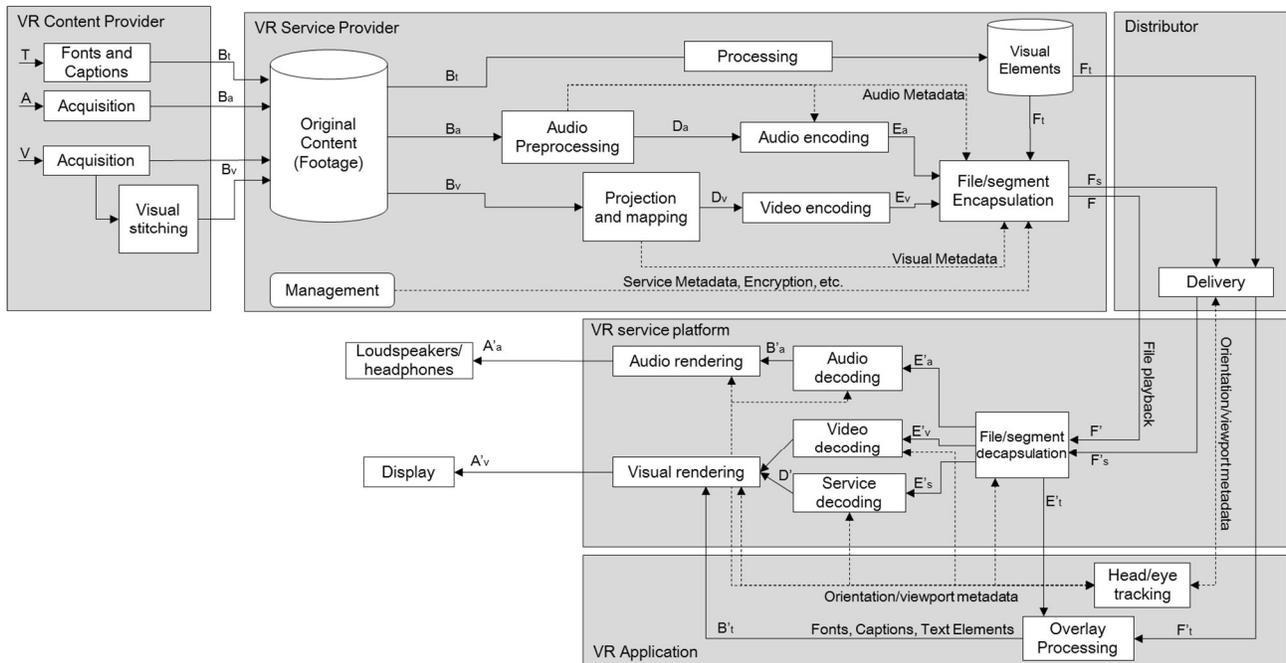


Figure 8: Example Architecture for simple VR Streaming Services

Note that certain functionalities (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 9 provides an extension to the architecture for an encrypted service. The secure media pipeline, see section 4.3.4, exists to provide a means to decrypt, decode and render content without allowing applications or other parts of the system to read or copy the content. Where inputs are required to modify or influence the rendering process, for example to specify or modify the portion of the current frame of the content to be displayed in the users view, these inputs need to be well defined and the operations they perform must not return data or result in transformations which allow the original content to be extracted or deduced.

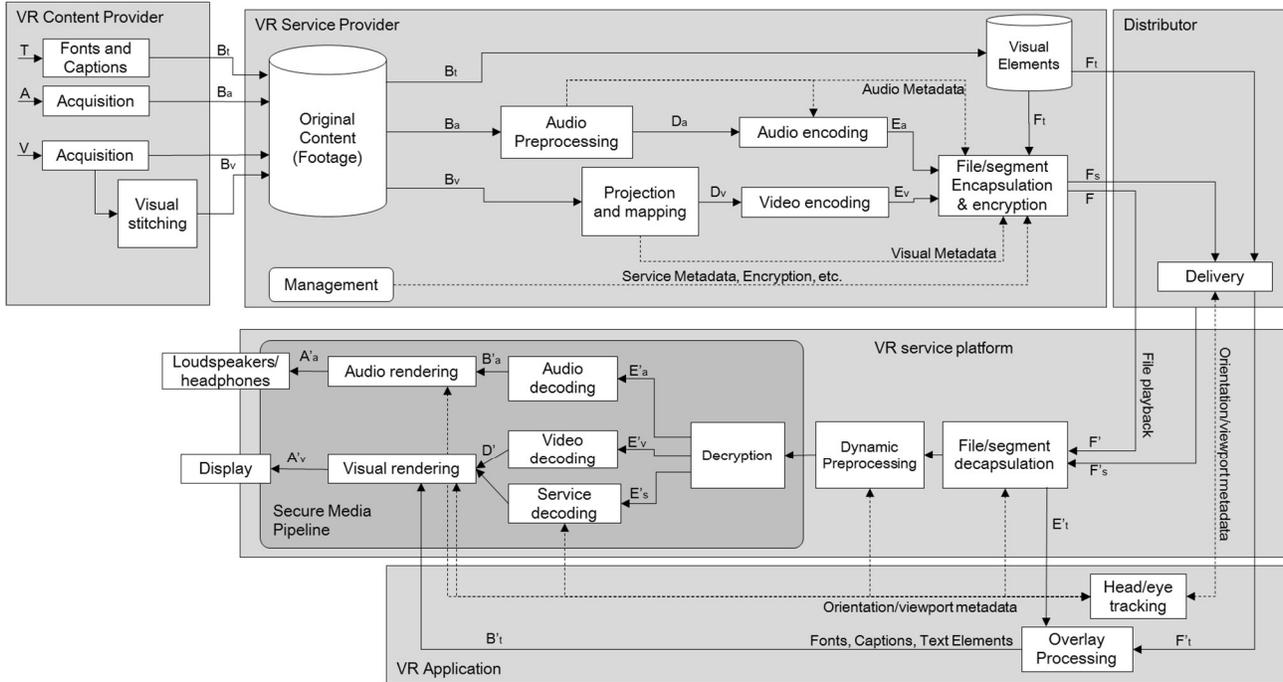


Figure 9: Example Architecture for encrypted VR Streaming Services

5.3.2 Client Architecture

In the distribution architectures depicted in Figure 8 and Figure 9, a VR compute platform (e.g., PC, console, tablet, smartphone, etc.) receives the delivered VR content, performs processing toward de-capsulation, decoding and rendering of the VR content, and forwards the processed input to the displayed at the VR device (e.g., HMD). VR device includes loudspeakers/headphones, display and sensors for head/eye tracking. Available communication technologies to realize the interface between the VR compute platform and VR device include USB for relaying sensory input data and HDMI for delivering the viewport image. This is depicted in Figure 10. In particular, the latest USB 3.2 specification [USB] can deliver speeds up to 20 Gbps and latest HDMI 2.1 specification [HDMI] can deliver speeds up to 48 Gbps with support for 8K and 10K resolutions and frame rates up to 120 fps, with the latter supporting VESA’s Display Stream Compression (DSC) standard [DSC]. It is also expected that wireless technologies can also be used for communication between the VR compute platform and the VR device. For example, such wireless connectivity may be realized via the WiGig standard [WIGIG] interface that uses the 60 GHz band for short-range high bandwidth wireless connectivity, with data rates up to 7 Gbps and less than 7 ms latency.

Note: It should be noted that an exception to the client-side depicted in Figure 10 is the integrated HMD platform where VR compute platform is part of the VR device.

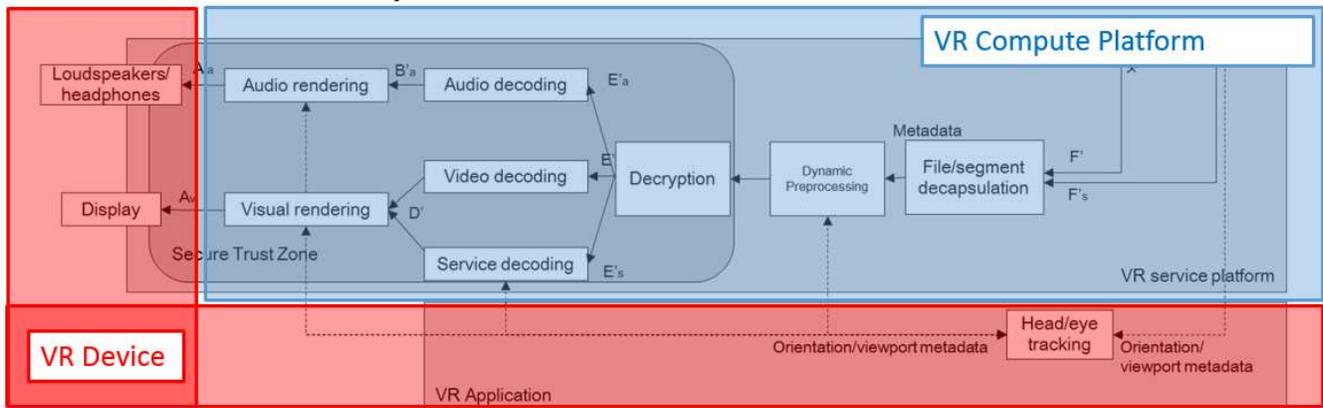


Figure 10: Client side processing on the example architecture

An example high-level VR software stack is depicted in Figure 11, showing the relationship between VR application, VR runtime and VR hardware / device drivers on the VR compute platform. Here VR hardware for instance includes HW capabilities (e.g., CPU/GPU) on the VR platform for decoding and rendering. Device drivers for instance include drivers for display and head/eye tracking sensors, and thus provides access to devices.

A bi-directional application interface (i.e., APIs) between the VR application and VR runtime carries information both ways. Another directional driver (HW) interface between the VR runtime and VR hardware / device drivers carries information both ways. Data mostly flows upward in this diagram from the drivers to the application, but some requests (like haptics on controllers) flows back down to the driver and is then passed on to the actual device.

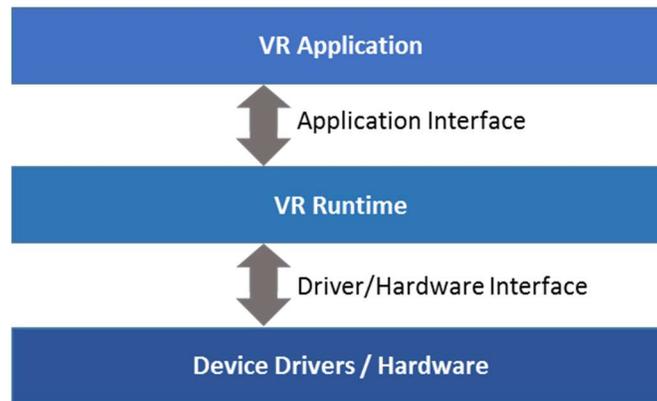


Figure 11: High-level VR software stack

The VR runtime interacts with both the VR application and device drivers via the APIs through the application and driver (HW) interfaces, and contains functions such as server and compositor functions. As such, VR applications do not have to interact with the devices directly, but instead the interaction occurs through the logical abstractions as enabled by the VR runtime. The VR runtime may also manage simultaneous interactions with multiple physical devices (e.g., display, controller, sensors, etc.) and may activate only a subset of devices to optimally use the VR platform resources.

The VR runtime includes functions such as predictive tracking to estimate and project the head pose based on sensor data. Moreover, the VR runtime performs functions such as Asynchronous Time Warp (ATW) and/or Asynchronous Space Warp (ASW) to display smooth motion even if the application is unable deliver new frames on time, by processing previously rendered frames based on predicted changes in user's head orientation. Other operations performed by the VR runtime includes taking pre-distorted images from the

VR application and applying of barrel distortion to correct for pincushion distortion caused by the lenses. Moreover, the VR runtime further can interact with the CPU/GPU on the platform to perform other operations such as decoding, decryption and graphics rendering.

For the application interface, API examples on the information from the VR application to the VR runtime includes pre-distortion image to display and haptics information. API examples on the information from the VR runtime to the VR application include predicted poses and controller/peripheral states.

For the driver (HW) interface, API examples on the information flowing from the VR runtime to the VR device drivers include the output of the audio rendering and visual rendering, and haptics information. API examples on the information flowing from the VR device to the VR runtime include head/eye tracking information, controller/peripheral state, and sensor data, e.g., hand/foot sensor information.

Standardization of APIs for the above mentioned two interfaces, i.e., application interface and driver (HW) interfaces, is the charter of the Khronos OpenXR Working Group. Without a cross-platform standard, VR applications, games and engines must port to each vendors' APIs. In turn, this means that each VR device can only run that apps that have been ported to its SDK. The result is high-development costs and confused customers – limiting market growth. To address this gap, the Khronos OpenXR working group is currently standardizing the APIs for these two interfaces with the goal of harmonizing the market [OXRGDC].

In particular, OpenXR 1.0 [OPENXR] has been released in July 2019 which provides the initial set of APIs for the application interface above in Figure 11. The driver-HW interface, referred to as OpenXR device plug-in extension in the OpenXR spec [OPENXR], is still under development and is also an optional feature of the specification. The OpenXR architecture is depicted in Figure 11a – it should be noted that this is in line with the high-level software stack depicted in Figure 11 above, but also extends the VR runtime to an XR runtime with the intention to not only address VR but also more broadly Extended Reality (XR) applications (including Augmented Reality (AR)). The specification currently supports two form factors, namely handheld displays (e.g., smartphone) with mono view and head mounted displays (HMDs) with stereoscopic view (one view per eye). Physical devices that the application chooses to use (e.g., HMD and controllers) can be grouped into logical systems of related devices that may support display, input, tracking etc. Session initiation, management and termination are supported to allow an application to indicate to the XR runtime how it wants to render and output VR/AR frames. Input and haptics support on the XR runtime is provided through interaction profiles, which are actually aligned with the capabilities of existing controllers in the market, e.g., those for Google DayDream, HTC Vive, Oculus Go, etc.

VR streaming and download applications may rely upon OpenXR APIs to obtain sensor data in relation to viewport and pose information. For instance, in the OMAF-DASH Streaming Client model in Figure 3 and OMAF-Download Client model in Figure 4, the VR application may obtain the sensor data on viewport and pose information from the VR runtime by calling the OpenXR APIs. This information could then be passed on to the rendering engine, as well as to the DASH access client for viewport-dependent delivery.

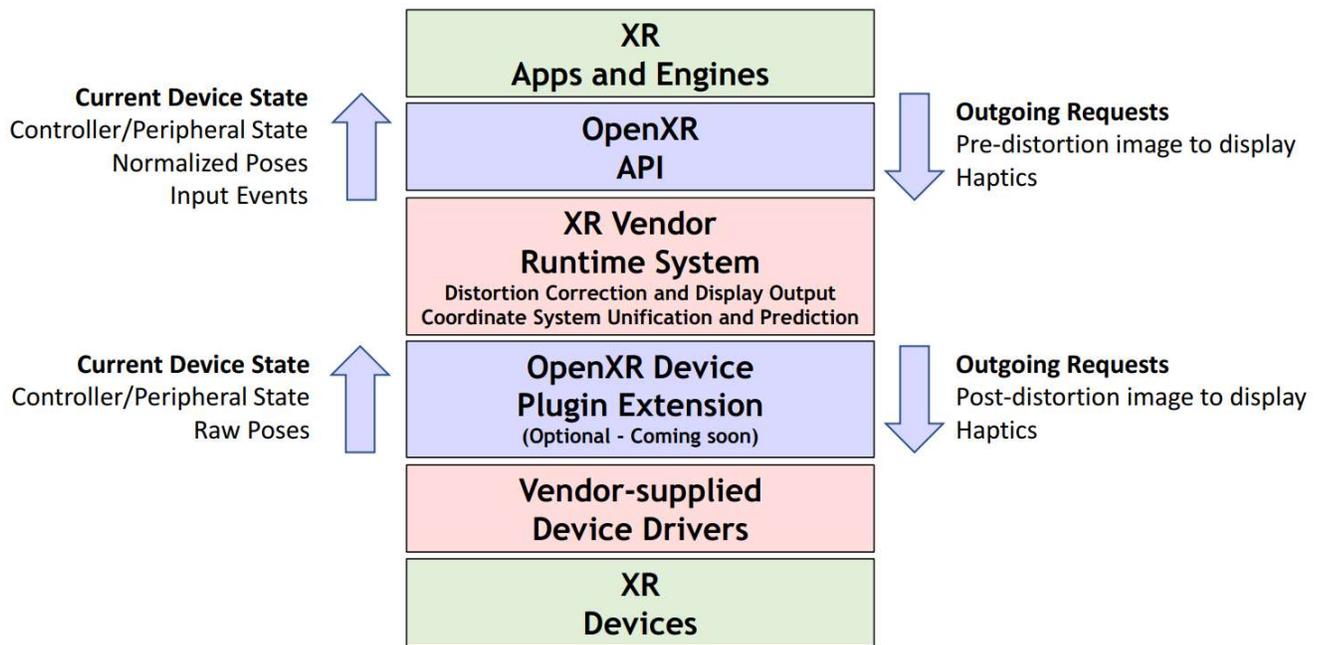


Figure 11a – Khronos OpenXR architecture

5.4 Technical Enablers

5.4.1 Suitable Media Profiles

5.4.1.1 Video

The HEVC-based viewport-independent OMAF video profile as presented in section 4.2.2.1.2 may be used to fulfil the use case and provide broad interoperability.

The HEVC-based viewport-dependent OMAF video profile as presented in section 4.2.2.1.3, may be used to fulfil the use case and provide broad interoperability, while achieving a higher resolution at the viewport than that of the HEVC-based viewport-independent OMAF video profile, but requiring fast stream switching support in distribution and decoding in order to meet the immersive experience expectations.

The unconstrained HEVC-based viewport-independent OMAF video profile as presented in section 4.2.2.1.5 may be used together with profile with elementary streams conforming to HEVC Main 10 Profile, Main Tier, Level 6, 6.1 or 6.2 to fulfil the use case and provide interoperability with 8K capable devices, while achieving a higher resolution at the viewport than that of the HEVC-based viewport-independent OMAF video profile.

5.4.1.2 Audio

The OMAF 3D Audio Baseline Media Profile as presented in section 4.2.2.2.2 may be used to fulfil the use case and provide broad interoperability.

5.4.2 Suitable Presentation Profiles

The presentation profiles defined in clause 11 of [OMAF] are applicable to this use case. Any other media information that is added to the presentation should utilize the same coordinate system as the video and audio formats to ensure proper rendering and presentation.

5.4.3 Distribution Systems

5.4.3.1 OTT Distribution using HTTP

This section primarily addresses interface F on distribution in Figure 8. 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 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 [MBMS] or ATSC3.0 [ATSC30]. In addition, HTTP delivery 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 against its playback capabilities before downloading it. For this purpose, 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 12. 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.

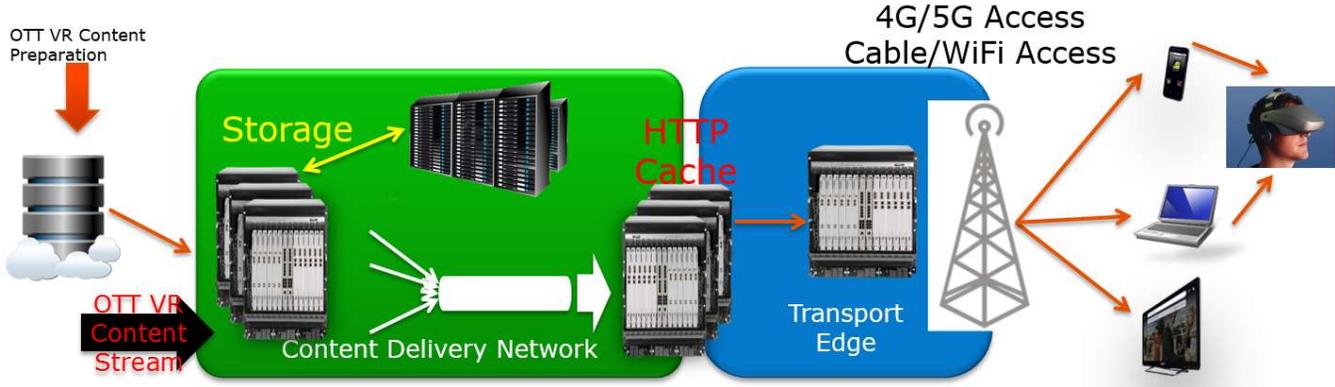


Figure 12: Example VR distribution system

Figure 13 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 13: Example protocol stack for VR content distribution

5.4.3.2 Download

If the HEVC-based viewport-independent OMAF video 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 HEVC-based viewport-dependent OMAF video 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 media profile in section 4.2.2.1.3.

If the unconstrained HEVC-based viewport-independent OMAF video 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 4.2.2.1.5.

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.

5.4.3.3 DASH Distribution

If the HEVC-based viewport-independent OMAF video 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 unconstrained HEVC-based viewport-independent OMAF video 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 4.2.2.1.5.

If the HEVC-based viewport-dependent OMAF video 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 4.2.2.1.3. In addition, one Adaptation Set per potential viewing direction is included in each period. Each of these Adaptation Sets either:

- contains a Preselection Supplemental descriptor that indicates which Adaptation Sets contribute to the Preselection corresponding to a viewing direction
- or contains a Representation with a @dependencyId attribute that indicates the Representations from other Adaptation Sets that are required for a given viewing direction.

Figure 14 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.

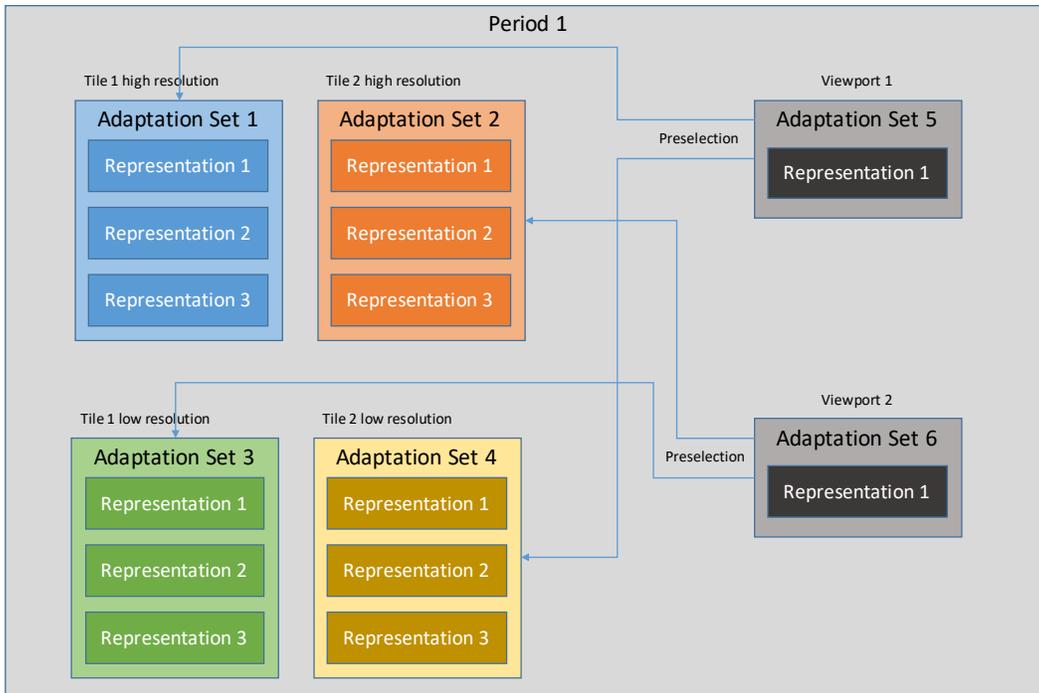


Figure 14: Exemplary DASH configuration setup

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 segment 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 sensory sickness. CDNs 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 VRIF 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 15 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.

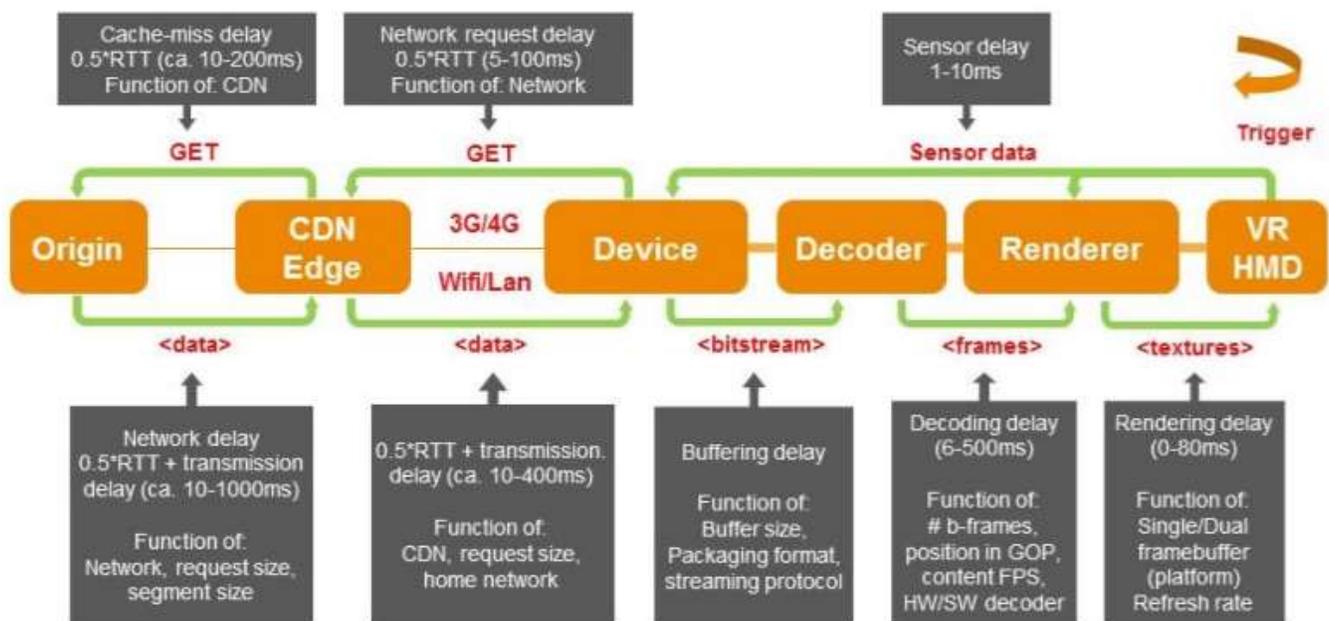


Figure 15: Latency sources contributing to M2HR latency [CDNOPT]

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 [HTTP2] enable features to reduce latency such as header compression and request bundling. In combination with QUIC [QUICdrafts] this can improve throughput by reducing dependence on RTT and the effects of TCP windowing.
- Access Network Considerations
- 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 improving on those available for LTE. Feasibility of VR services in 3GPP networks has been studied in 3GPP TR 26.918 [TR26918].

5.4.3.4.1 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 HEVC-based viewport-independent OMAF video profile

The video format is expected to conform to the video master format in clause 4.1.3.2 and video metadata is provided along with the video format.

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 *Coverage* parameter.

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

The original video format must be chroma subsampled from 4:2:2 to 4:2:0. After chroma subsampling, if the subsampled original video signal format is in the constraints of HEVC Main 10 Main Tier Level 5.1 and the scheme constraints, then the HEVC-based viewport-independent OMAF video profile may be used directly on the subsampled source signal to generate elementary streams following the media profile constraints as described in clause 10.1.2.2 of [OMAF] and the encoding and content preparation in section 5.5.2.

The relevant HEVC Main 10 Main Tier Level 5.1 are:

- Max luma picture size is 8,912,896
- Max luma sample rate (samples/s) is 534,773,760
- The maximum bitrate (kbit/s) is 40,000

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

- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 8,912,896$
- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 534,773,760$

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

Stereoscopic signals frame packed with side-by-side or top-bottom that can be distributed with this profile need to fulfill the following requirements:

- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 4,456,448$
- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 267,386,880$

Stereoscopic signals frame packed using temporal interleaving that can be distributed with this profile need to fulfill the following requirements:

- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 8,912,896$
- $\text{ceil}_{16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil}_{16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 267,386,880$

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

- Monoscopic:
 - 4096 H × 2048 V, 4:2:0, at 25, 30, 50 and 60fps with full content coverage
- Stereoscopic:
 - Each view with 4096 H × 2048 V, 4:2:0 at 25 and 30fps, if frame-packed using temporal interleaving.
 - Each view with 2048 H × 2048 V, 4:2:0, at 25, 30, 50 and 60fps, if framed packed using side-by-side.
 - Each view with 4096 H × 1024 V, 4:2:0, at 25, 30, 50 and 60fps, if framed packed using top-bottom.
 - Each view with 2944 H × 1472 V, 4:2:0, at 25, 30, 50, and 60fps, if frame packed using top-bottom.

Note: Picture sizes do not correspond necessarily to the full 360° reference if Coverage does not indicate the whole 360° sphere. They correspond to the content covered by Coverage and Cropping, i.e. sizes correspond to $(\text{FullWidthPixel} - (\text{Cropping.Left} + \text{Cropping.Right})) * (\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom}))$. Therefore, *FullWidthPixel* and *FullHeightPixel* might be greater than 4096 and 2048 respectively.

If the original video signal after chroma subsampling 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 HEVC-based viewport-independent OMAF video profile may be used after preprocessing of the original video content such that the constraints are fulfilled. Examples for videos that require preprocessing are

- Monoscopic:
 - 6144 H × 3072 V, 4:2:0
 - 8192 H × 4096 V, 4:2:0
- Stereoscopic:
 - For each 4096 H × 2048 V, 4:2:0, at 50 and 60fps, i.e. 8192 H × 2048 V for side-by-side and 4096 H × 4096 V for top-bottom
 - 6144 H × 3072 V, 4:2:0, i.e. 12288 H × 3072 V for side-by-side and 6144 H × 6144 V for top-bottom
 - 8192 H × 4096 V, 4:2:0, i.e. 16384 H × 4096 V for side-by-side and 8192 H × 8192 V for top-bottom

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 HEVC-based viewport-dependent OMAF video profile

For the HEVC-based viewport-dependent OMAF video 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.

5.5.1.4 Unconstrained HEVC-based viewport-independent OMAF video profile

As described in clause 4.2.2.1.5, this profile can be used to distribute 8K content using elementary streams conforming to HEVC Main 10 Profile, Main Tier, Level 6, 6.1 or 6.2, assuming the corresponding video decoder capabilities are in place in the receiving device. This section presents suitable production formats using this profile for 8K content distribution formats based on elementary streams conforming to HEVC Main 10 Profile, Main Tier, Level 6.1. The corresponding 8K video decoding capabilities are expected to be supported by HMDs and other devices rendering omnidirectional video powered by platforms released in 2020 and onwards. The key support is the availability of HEVC Main 10 Level 6.1 decoders to process 8k by 4k frames at frame rates up to 60 fps as well as GPU-based rendering.

The video format is expected to conform to the video master format in clause 4.1.3.2 and video metadata is provided along with the video format. 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 Coverage parameter. 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

The original video format must be chroma subsampled from 4:2:2 to 4:2:0. After chroma subsampling, if the subsampled original video signal format is in the constraints of HEVC Main 10 Main Tier Level 6.1 and the scheme constraints, then the unconstrained HEVC-based viewport-independent OMAF video profile may be used directly on the subsampled source signal to generate elementary streams and the encoding and content preparation in section 5.5.2. The relevant HEVC Main 10 Main Tier Level 6.1 are:

- Max luma picture size is 35,651,584
- Max luma sample rate (samples/s) is 2,139,095,040
- The maximum bitrate (kbit/s) is 120,000

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

- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 35,651,584$

- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 2,139,095,040$ with $\text{ceil16}(x)$ the smallest integer that is greater or equal than x and a multiplier of 16.

Stereoscopic signals frame packed with side-by-side or top-bottom that can be distributed with this profile need to fulfill the following requirements:

- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 17,825,792$
- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 1,069,547,520$

Stereoscopic signals frame packed using temporal interleaving that can be distributed with this profile need to fulfill the following requirements:

- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) \leq 35,651,584$
- $\text{ceil16}(\text{FullWidthPixel} + \text{Padding} - (\text{Cropping.Left} + \text{Cropping.Right})) * \text{ceil16}(\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom})) * \text{FrameRate} \leq 1,069,547,520$

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

- Monoscopic:
 - 8192 H × 4096 V, 4:2:0, at 25, 30, 50 and 60fps with full content coverage
- Stereoscopic:
 - Each view with 8192 H × 4096 V, 4:2:0 at 25 and 30fps, if frame-packed using temporal interleaving.
 - Each view with 4096 H × 4096 V, 4:2:0, at 25, 30, 50 and 60fps, if framed packed using side-by-side.
 - Each view with 8192 H × 2048 V, 4:2:0, at 25, 30, 50 and 60fps, if framed packed using top-bottom.

Note: Picture sizes do not correspond necessarily to the full 360° reference if Coverage does not indicate the whole 360° sphere. They correspond to the content covered by Coverage and Cropping, i.e. sizes correspond to $(\text{FullWidthPixel} - (\text{Cropping.Left} + \text{Cropping.Right})) * (\text{FullHeightPixel} - (\text{Cropping.Top} + \text{Cropping.Bottom}))$. Therefore, FullWidthPixel and FullHeightPixel might be greater than 8192 and 4096 respectively.

If the original video signal after chroma subsampling is beyond the constraints of HEVC Main 10 Level 6.1 constraints, then the unconstrained HEVC based viewport-independent OMAF video profile may be used after preprocessing of the original video content such that the constraints are fulfilled.

Examples for videos that require preprocessing are

- Monoscopic:
 - > 8192 H × 4096 V, 4:2:0
- Stereoscopic:
 - 8192 H × 4096 V, 4:2:0, i.e. 16384 H × 4096 V for side-by-side and 8192 H × 8192 V for top-bottom

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

5.5.2 Sphere-to-Texture Mapping and SEI Message Generation

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 in clause 5.1 of [OMAF].

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
- 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 monoscopic or stereoscopic. If stereoscopic, 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 -180 degrees inclusive to 180 degrees exclusive.
 - **AzimuthMax** specifies the maximum azimuth value of the coverage sphere region in the range of -180 degrees inclusive to 180 degrees exclusive. 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 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 the full signal is not provided but a cropped version of it is, 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 center point of a sample location (i, j) is derived following clause 5.2.2 of [OMAF] for monoscopic or each of the views for **StereoMode** separately, invoked with **FullWidthPixel**, **FullHeightPixel**, **Cropping.Left** + i and **Cropping.Top** + j as inputs.

If the **Rotation** parameter is not present, then the global projected sphere coordinates (ϕ', θ') for the sample location for the center 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 center point of a sample location (i, j) are derived based on its the local sphere coordinates (ϕ, θ) .

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 [ADDSEI]) 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 following general rules apply for SEI message generation:

- An SEI message with payload type 150 (equirectangular projection) 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 **Padding** parameter is present, then the `erp_padding_flag` is set to 0 and no region-wise packing SEI message is present.

When the video provides full 360° coverage and **Padding** parameter is present, then the following applies:

- `erp_padding_flag` is set to 0 in the equirectangular projection SEI message.
- region-wise packing SEI messages (as defined in sections D.2.41.4 and D.3.41.4 of [ADDSEI]) is generated in order to maximize the visible information in the encoded 2D image using the **Padding** 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**
 - `proj_region_height[0]` is set to **FullHeightPixel**
 - `proj_region_top[0]` is set to 0

- `proj_region_left[0]` is set to 0
- `transform_type[0]` is set to 0
- `packed_region_width[0]` is set to **FullWidthPixel**
- `packed_region_height[0]` is set to **FullHeightPixel**
- `packed_region_top[0]` is set to 0
- `packed_region_left[0]` is set to **Padding/2**

No guidance is given for parameters which are not listed above.

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

- region-wise packing SEI messages (as defined in sections D.2.41.4 and D.3.41.4 of [ADDSEI]) is 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**

No guidance is given for parameters which are not listed above.

When the video is stereoscopic, then the frame packing needs to be generated and an appropriate frame packing arrangement SEI message (as defined in [ADDSEI] section D.3.16) 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.

5.5.3 Encoding and Content Preparation

5.5.3.1 HEVC-based viewport-independent OMAF video profile

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

If the original chroma-subsampled video signal is beyond the constraints of HEVC Main 10 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 may include temporal and/or spatial subsampling.

Specifically, for stereoscopic content at 4096 H × 2048 V, 4:2:0 at 25 and 30fps using temporal interleaving frame-packing is the most suitable format for distribution.

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 stereoscopic at 50 and 60fps may be preprocessed to
 - 2880 H × 1440 V, 4:2:0, per eye at 50 or 60 fps, if frame-packed top-and-bottom
 - 2048 H × 2048 V, 4:2:0, per eye at 50 and 60fps, if frame-packed side-by-side
 - 4096 H × 1024 V, 4:2:0, per eye at 50 and 60fps, if frame-packed top-and-bottom
 - 4096 H × 2048 V, 4:2:0 at 25 and 30fps using temporal interleaving
- 6144 H × 3072 V, 4:2:0, monoscopic 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, stereoscopic with full content coverage
 - To any of the formats that can be directly distributed.
- 8192 H × 4096 V, 4:2:0, monoscopic or stereoscopic with full content coverage
 - To any of the formats that can be directly distributed.

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 **Coverage** parameter.

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 monoscopic:
 - 4096 H × 2048 V, 4:2:0, monoscopic at 25, 30, 50 and 60fps with full content coverage or more than 180 degree coverage
- For stereoscopic:

- 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

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 arrangement 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 [CMAF] 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-IF IOP [DASHIF IOP] or ISO/IEC 23000-19 annex D [CMAF]). The number of Representations per Adaptation Set as well as their encoding bitrates depend on different factors, such as encoder performance, content complexity, and distribution parameters. In the absence of other information, a first idea on suitable bitrates and their performance, the 3GPP TR26.918 [TR26918] provides some ideas (further details are provided 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, it is 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 7, more details are provided in [OMAF].

Table 7: Mapping of SEI Message Information to OMAF Metadata

SEI Message	OMAF Metadata
<code>equirectangular_projection</code>	<code>ProjectedOmniVideoBox</code> <code>RegionWisePackingBox (Padding)</code>
<code>region_wise_packing</code>	<code>RegionWisePackingBox</code> <code>CoverageInformationBox (Optional)</code>
<code>frame_packing_arrangement</code>	<code>StereoVideoBox</code>

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

Figure 16 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.

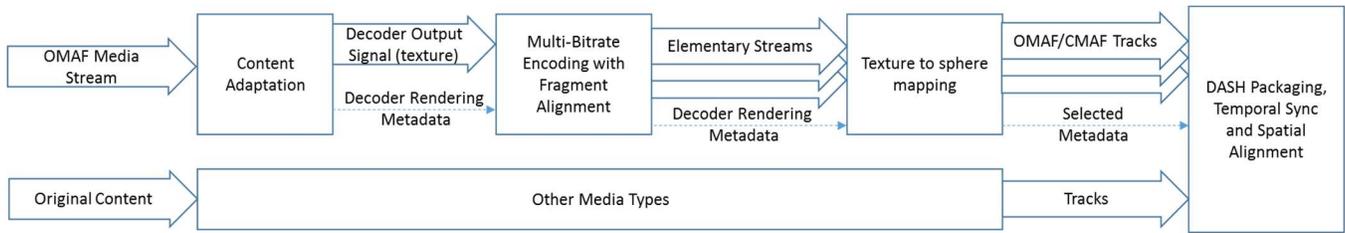


Figure 16: Content Preparation for DASH Distribution

5.5.3.2 HEVC-based viewport-dependent OMAF video profile

If the HEVC-based viewport-dependent OMAF video 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 Main 10 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 HEVC-based viewport-dependent OMAF video profile as explained below.

Four 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
- Number of Motion Constraint Tile Set for each resolution: numMCTS(r)
- 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

Table 8 summarizes the recommended tile layout patterns based on the characteristics mentioned above.

Table 8: Recommended tile layouts

Title for the tile layout scheme	Target FOV	Resolution of ERP before pre-processing	Target display resolution	Definition of scheme
6K effective ERP	Approx. 120° or less	6144 H × 3072 V or greater	2560 H × 1440 V	Clause A.6.3 of [OMAF]

6K effective cubemap	Approx. 135° or less	Greater than 6144 H × 3072 V	2560 H × 1440 V	Clause A.6.4 of [OMAF]
----------------------	----------------------	------------------------------	-----------------	------------------------

While the recommended schemes are defined in [OMAF], a brief summary is provided below.

For ERP, either provided by the source signal or after adaptation and pre-processing, the content is expected to be in the following versions:

- For monoscopic
 - 6144 H × 3072 V, 4:2:0, monoscopic up to 60fps
 - 3072 H × 1536 V, 4:2:0, monoscopic up to 60fps
 - 1536 H × 768 V, 4:2:0, monoscopic up to 60fps

In the 6K effective ERP scheme the content for the viewport originates from an ERP sequence of 6K resolution (6144×3072), while other parts of the content originate from either a 3K (3072×1536) version or 1.5K (1536×768) version. The polar stripes (with a value for elevation higher than 60° or lower than -60°) are encoded at resolutions 1.5K and 3K, while the central part (with a value for elevation between -60° and 60°) that covers an elevation range of 120° is encoded at resolutions 3K and 6K. Motion-constrained tile sets (MCTSs) are used in the encoding. The encoded MCTS sequences are combined with extractor tracks to packed pictures for 16 distinct viewing orientations, each corresponding to a selection of four spherically adjacent MCTSs from the 6K bitstream and a viewing orientation either above or below the equator. Region-wise packing metadata is included in the extractor tracks to indicate the mapping of the packed regions to the respective projected regions. Figure 17 illustrates an example for a viewing orientation above the equator. Each colored rectangle of a particular color indicates a packed region and the respective projected region. The picture size of the bitstream resolved from the extractor track is 3840×2304, which conforms to HEVC Main10 Main Tier Level 5.1.

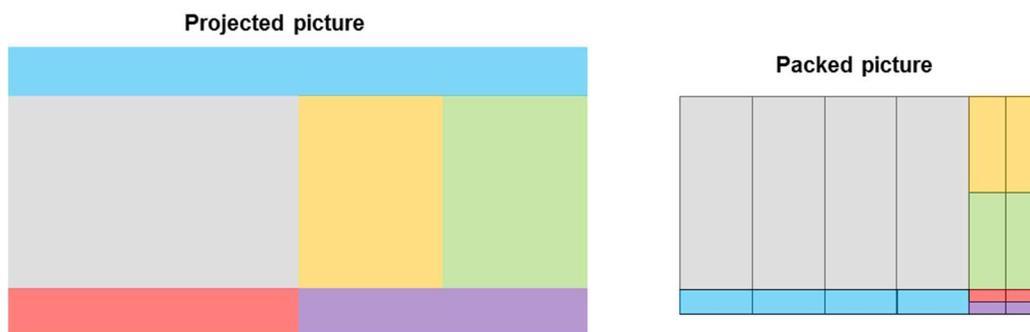


Figure 17: Example of the packed picture and the respective projected picture of one of the 16 extractor tracks, for a viewing orientation above the equator.

For CMP, the content, after adaptation and pre-processing, is expected to be in the following versions in CMP format:

- For monoscopic
 - 4608 H × 3072 V, 4:2:0, monoscopic up to 60fps
 - 2304 H × 1536 V, 4:2:0, monoscopic up to 60fps

This requires the source signal in ERP is expected to be in the following version or at higher resolution:

- For monoscopic
 - > 6144 H × 3072 V, 4:2:0, monoscopic up to 60fps

The effective 6K cubemap arrangement codes the viewport with cube faces of 1536×1536 samples, which could be considered to approximately equivalent to 6K ERP in terms of sampling density. In the arrangement, 12 tiles (encoded as MCTS) originate from the high-resolution version, and the remaining tiles (encoded as MCTS) are extracted from a cubemap having a quarter resolution compared to the high-resolution bitstream. 24 extractor tracks are created, each for different viewing orientation. The bitstreams resolved from the extractor tracks have resolution 1920×4608, which conforms to HEVC Main10 Main Tier Level 5.1.

Note: The same tiling granularity could be used if lower target display resolutions are considered with lower resolutions than the ones listed above for source and content versions.

Figure 18 illustrates the video content preparation for DASH distribution using the HEVC-based viewport-dependent OMAF video profile.

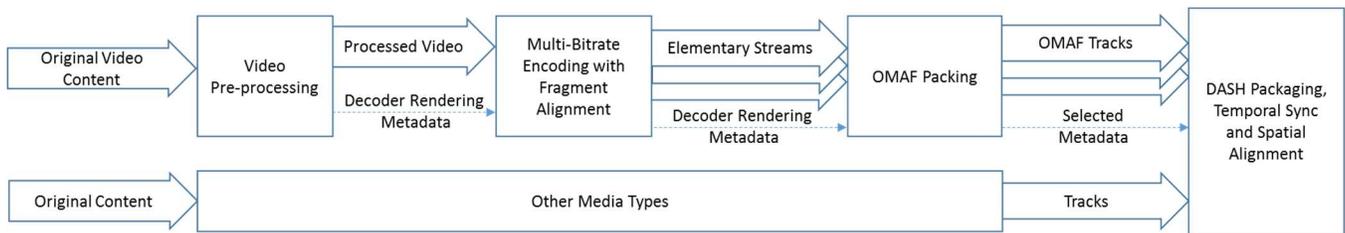


Figure 18: Video content preparation for DASH distribution with HEVC-based viewport-dependent OMAF video profile

Video pre-processing

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 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×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, i.e. same number of samples and same Decoding Time (DT) and Composition Time (CT) for all samples.

Note: In cases where adjacent tiles are always used together the constraint of NxM tiles being encoded as MCTS can be relaxed to achieve a better coding efficiency (see for instance the polar stripes for the ERP configuration in clause D.6.3 of [OMAF]).

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 HEVC-based viewport-dependent OMAF video profile. Therefore, it is recommended to have frequent random access points in the encoding, for example every second or less.

For DASH streaming cases, it is also possible 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. Thus, different elementary streams are offered with different switch point intervals. As for the different bitrate versions, elementary streams with different switch point intervals should be encoded with aligned fragments, i.e. same number of samples and same Decoding Time (DT) and Composition Time (CT) for all samples (or in other words same GOP structures). In addition, segments and subsegments of the Representations generated by these elementary should be aligned (have the same amount of the described fragments).

OMAF Packaging

After encoding, each bitstream of MCTS 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 MCTS 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 MCTS is assigned a unique `mctsID` (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 MCTS is encapsulated into a single ISO BMFF file containing a single MCTS track with the following considerations as described in [OMAF]:

- 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
- `'rwpk'`: indicating one region for the MCTS with its respective dimension and position within the region-wise packed frame and corresponding dimension and position in the projected frame.

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 projection SEI message or the 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 could be calculated using following equation:

$$trackID = mctsID + \sum_{i=0}^r numMCTS(i)$$

with $r \in [0, \dots, R - 1]$ indicating which resolution level they belong to, and $mctsID > 0$, 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 HEVC-based viewport-dependent OMAF video 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:

$$trackID = c + 1 + \sum_{i=0}^{R-1} numMCTS(i)$$

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 as many extractors as track_IDs of TL_c for each sample within the extractor track; each of which is generated using an inline constructor (optional for the first extractor) 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. (optional for the first extractor)

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

The sample constructor references the dependent track of the corresponding MCTS 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 NAL unit length and 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 MCTS with their respective dimensions and position within the region-wise packed frame and corresponding dimension and position in the projected frame.

As for the files containing the MCTS tracks, these files include 'hevd' 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 equirectangular projection SEI message or the cubemap projection SEI message 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_local_flag is set equal to 1
- quality_type is set equal to 1

- `quality_ranking` is set to the value $r + 1$, with $r \in [0, \dots, R - 1]$ indicating which resolution level the region (corresponding MCTS) belongs to
- `orig_width` is set equal to the width of the content resolution used for encoding the region
- `orig_height` is set equal to the height of the content resolution used for encoding the region

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_local_flag` is set equal to 1
- `quality_type` is set equal to 1
- `quality_ranking` is set to the value $r + 1$, with $r \in [0, \dots, R - 1]$ indicating which resolution level the region (corresponding MCTS) belongs to.
- `orig_width` is set equal to the width of the content resolution used for encoding the region.
- `orig_height` is set equal to the height of the content resolution used for encoding the region.

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 [CMAFAMD1], Annex J should be taken into account. Multiple bitrates may be generated by adapting the bitrate of the audio.

5.5.3.4 Unconstrained HEVC-based viewport-independent OMAF video profile

As described in clause 4.2.2.1.5, this profile can be used to distribute 8K content using elementary streams conforming to HEVC Main 10 Profile, Main Tier, Level 6, 6.1 or 6.2, assuming the corresponding video decoder capabilities are in place in the receiving device. This section presents encoding and content preparation guidelines using this profile for 8K content distribution formats based on elementary streams conforming to HEVC Main 10 Profile, Main Tier, Level 6.1.

If the original chroma-subsampled video is in the constraints of the encoder, then it may be distributed directly. If the original chroma-subsampled video signal is beyond the constraints of HEVC Main 10 Level 6.1, then the original video content needs to be adapted to be encoded properly with an HEVC Main 10 Level 6.1 encoder and to meet the profile level constraints. Adaptation may include temporal and/or spatial subsampling. Specifically, for stereoscopic content at $8192 \text{ H} \times 4096 \text{ V}$, 4:2:0 at 25 and 30fps using temporal interleaving frame-packing is the most suitable format for distribution. As examples, the signals beyond the limits of the HEVC profile level constraints documented in section 5.5.1 may be preprocessed as follows:

- $8192 \text{ H} \times 4096 \text{ V}$, 4:2:0, per eye stereoscopic at 50 and 60fps may be preprocessed to
 - $4096 \text{ H} \times 4096 \text{ V}$, 4:2:0, per eye at 50 and 60fps, if frame-packed side-by-side

- 8192 H × 2048 V, 4:2:0, per eye at 50 and 60fps, if frame-packed top-and-bottom
- 8192 H × 4096 V, 4:2:0 at 25 and 30fps using temporal interleaving
- > 8192 H × 4096 V, 4:2:0, monoscopic or stereoscopic with full content coverage
 - To any of the formats that can be directly distributed.

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 Coverage parameter. 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 monoscopic:
 - 8192 H × 4096 V, 4:2:0, monoscopic at 25, 30, 50 and 60fps with full content coverage or more than 180 degree coverage
- For stereoscopic:
 - 8192 H × 4096 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
 - 8192 H × 4096 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

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 arrangement 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 [CMAF] 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-IF IOP [DASHIFIOP] or ISO/IEC 23000-19 annex D [CMAF]). The number of Representations per Adaptation Set as well as their encoding bitrates depend on different factors, such as encoder performance, content complexity, and distribution parameters. In the absence of other information, a first idea on suitable bitrates and their performance, the 3GPP TR26.918 [TR26918] provides some ideas (further details are provided 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, it is 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 7,

more details are provided in [OMAF]. Otherwise, no specific aspects for VR content need to be taken into account.

5.5.3.5 Simple tiling OMAF video profile

The simple tiling OMAF video profile can utilize the tile layout schemes that are designed for the HEVC-based viewport-dependent OMAF video profile. Thus, the tile layout schemes presented in section 5.5.3.2 are applicable. Moreover, most of the main principles for selecting a tile layout configuration and most of the tools for processing content are the same.

The simple tiling OMAF video profile provides capability for taking advantage of HEVC Level 6, 6.1 or 6.2 decoding capability. This section presents tile layout patterns applicable for such decoding capability.

The following characteristics may be taken into account when preparing content and/or when an OMAF player selects an OMAF base track and/or OMAF tile tracks:

- Whether the content is monoscopic or stereoscopic
- Whether the targeted display is monoscopic or stereoscopic
- Target FOV
- Available decoding capacity, which may be expressed in terms HEVC Level.
 - Levels 5, 5.1, and 5.2 correspond approximately to 4K resolution at 30, 60, and 120 Hz picture rate, respectively.
 - Levels 6, 6.1, and 6.2 correspond approximately to 8K resolution at 30, 60, and 120 Hz picture rate, respectively.
- Resolution of ERP or CMP video before pre-processing
- Target display resolution

Table 9 summarizes the example tile layout patterns based on the characteristics mentioned above.

Table 9: Example tile layouts for the simple tiling OMAF video profile

Title for the tile layout scheme	Mono/ stereo	HEVC Level	Target FOV	Per-view ERP resolution for pre-processing	Target display resolution (both views)	Definition of scheme
6K effective ERP	Mono	5/5.1/ 5.2	Approx. 120° or less	6144 H × 3072 V or greater	2560 H × 1440 V	Clause D.6.3 of [OMAF]
6K effective cubemap	Mono	5/5.1/ 5.2	Approx. 135° or less	Greater than 6144 H × 3072 V	2560 H × 1440 V	Clause D.6.4 of [OMAF]
6K stereo cubemap	Stereo	6/6.1/ 6.2	Any	6144 H × 3072 V or greater	2560 H × 1440 V	This section
8K mixed-quality ERP	Mono	6./6.1 /6.2	Any	8192 H × 4096 V or greater	4096 H × 2048 V	This section
8K effective stereo cubemap	Stereo	6/6.1/ 6.2	Approx. 135° or less	8192 H × 4096 V or greater	4096 H × 2048 V	This section

Figure 198a presents the tile layout for 6K stereo cubemap with 2x2 tiles per cube face. In this example, both views are encoded with two picture qualities or bitrates, but generally any number of qualities or bitrates can be used. Moreover, in this example 2x2 tiling per cube face is in use, but generally any tile partitioning can be used. The tile partitioning affects the granularity of selecting the viewport and the number of HTTP GET requests. OMAF tile tracks containing collocated tiles with different qualities are grouped into an 'alte' track group. Each tile layout position in an OMAF base track refers to the respective 'alte' track group. An OMAF player resolves each reference to an 'alte' track group to either of the collocated OMAF tile tracks in the track group in a manner that the viewport is covered by the higher quality of the available two qualities.

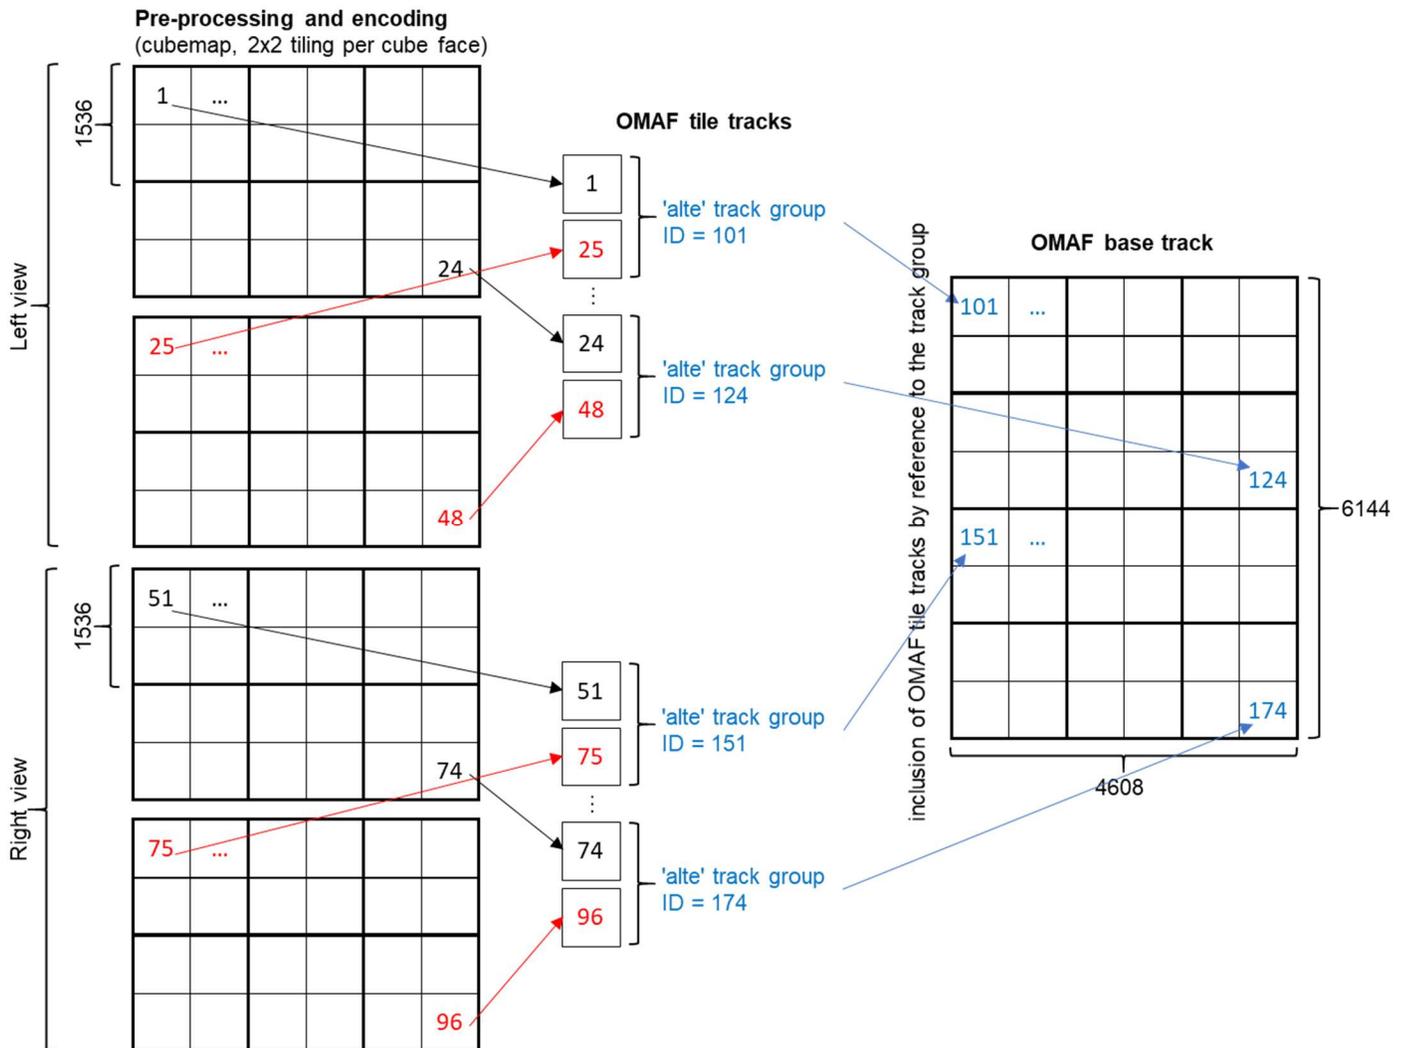


Figure 198a: 6K stereo cubemap tile layout with 2x2 tiles per cube face.
Both views are encoded with two qualities / bitrates illustrated with black and red font color.

Figure 18b presents an example tile layout for 8K mixed-quality ERP. In this example, the same original video is coded with two picture qualities or bitrates, illustrated with different font colors. but generally any number of qualities or bitrates can be used. Moreover, in this example the four "equator" tiles have 90° azimuth range and 135° elevation range, and the two "polar" cover the entire azimuth range. In a typical viewing situation and display FOV, reception of two "equator" tiles at high quality and the remaining tiles at lower quality would be streamed. Generally, the 8K mixed-quality ERP could use any tile partitioning, and the tile division in Figure 18b is one example among many viable options. The tile partitioning affects the granularity of

selecting the viewport and the number of HTTP GET requests. OMAF tile tracks containing collocated tiles with different qualities are grouped into an 'alte' track group. Each tile layout position in an OMAF base track refers to the respective 'alte' track group. An OMAF player resolves each reference to an 'alte' track group to either of the collocated OMAF tile tracks in the track group in a manner that the viewport is covered by the higher quality of the available two qualities.

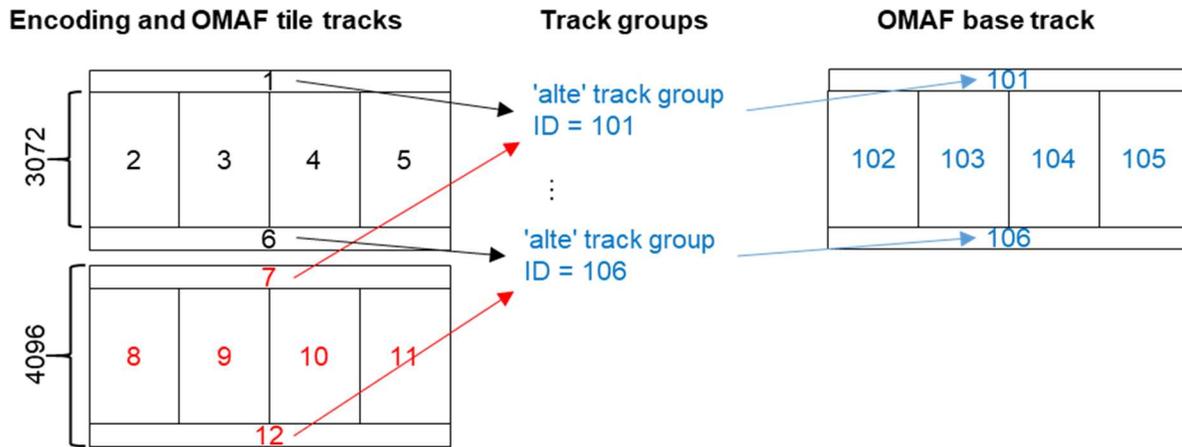


Figure 18b: Example tile layout of 8K mixed-quality ERP.

The content is made available at two qualities / bitrates illustrated with black and red font color.

Figure 18c presents the tile layout for 8K effective stereo cubemap with 2x2 tiles per cube face. In this example, both views are encoded with two resolutions 2048x2048 per cube face (corresponding to 8K resolution at the equator) and 1024x1024 per cube face. In addition to different resolutions, versions of different qualities or bitrates can be encoded for finer bitrate adjustment. Moreover, in this example 2x2 tiling per cube face is in use, but generally any tile partitioning can be used. The tile partitioning affects the granularity of selecting the viewport and the number of HTTP GET requests. OMAF tile tracks of a particular view and of a particular size are grouped into an 'alte' track group, and thus four 'alte' track groups are formed: high-resolution left view tiles, low-resolution left view tiles, high-resolution right view tiles, and low-resolution right view tiles. The OMAF base track contains 12 tile positions for each tile group. An OMAF player resolves each reference to an 'alte' track group to one of the OMAF tile tracks in the referenced tile group in a manner that the viewport is covered by the high-resolution tiles.

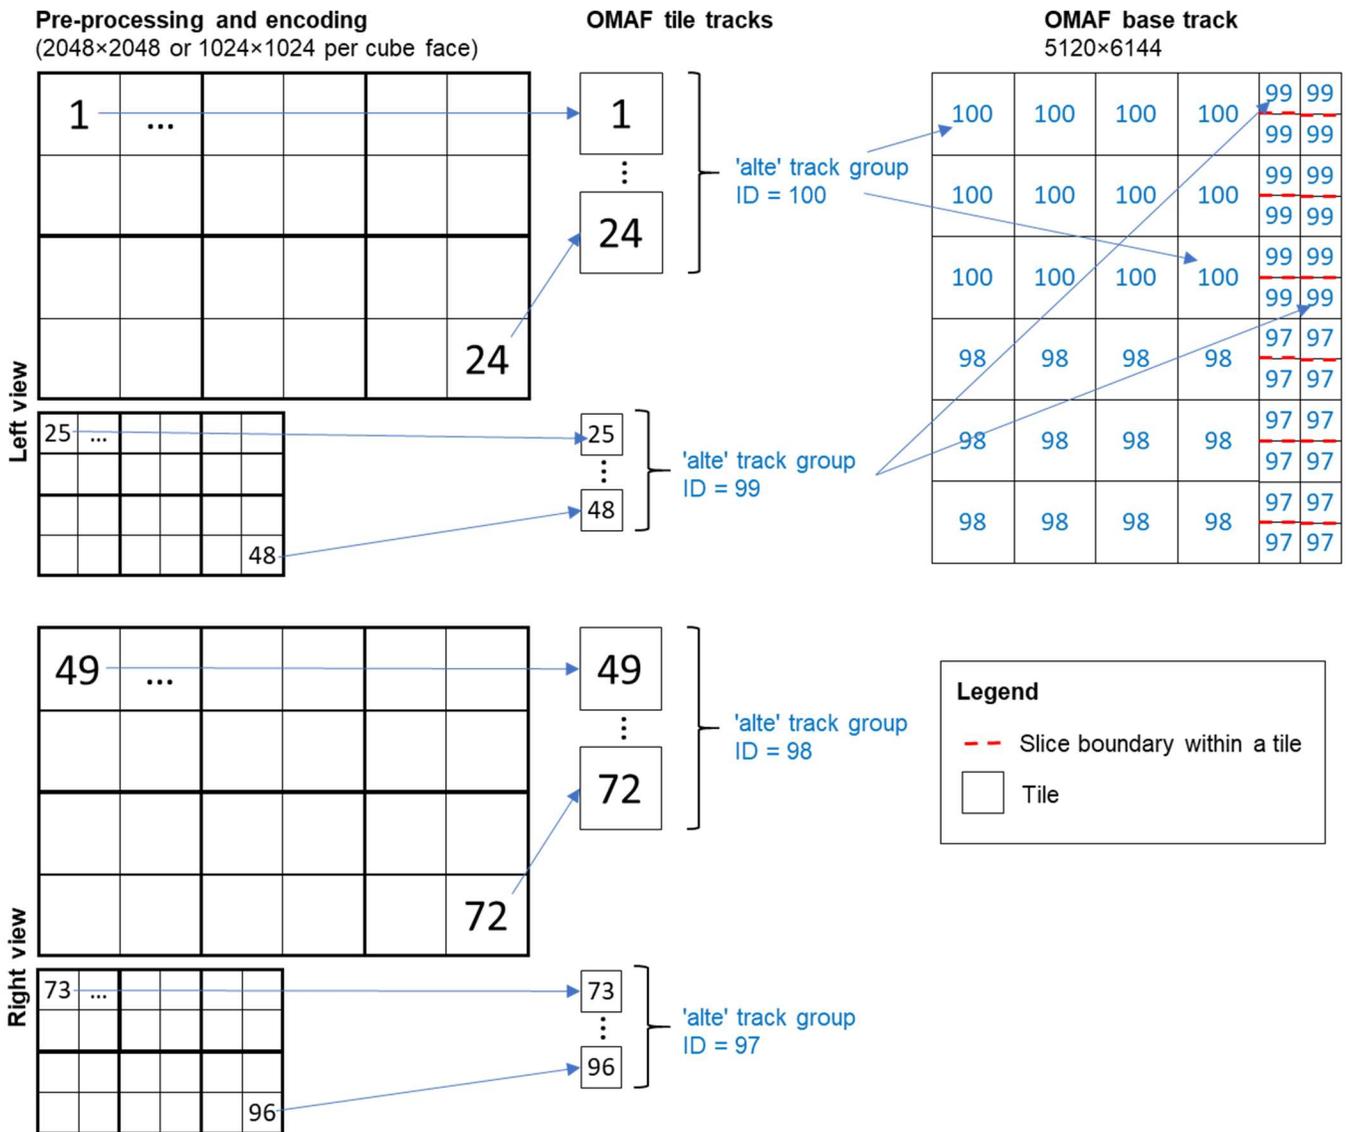


Figure 18c: 8K effective stereo cubemap with 2x2 tiles per cube face.

5.5.4 Distribution

5.5.4.1 Download

5.5.4.1.1 HEVC-based viewport-independent video profiles

If the HEVC-based viewport-independent OMAF video profile or unconstrained HEVC-based viewport-independent OMAF video 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='hevi' codecs='resv.pov+erpv.hvc1.1.6.L93.B0'. If the visual component

is added for a download/storage application and is generated according to the viewport-independent profile, then the entire visual component resides in a single track of the ISO file format package.

Encoding should be done in a way such that the quality is sufficiently high to minimize any visual artefacts due to encoding, under the profile/level constraints of the viewport-independent profiles.

The encoding should prioritize quality rather than bitrate or size of the file. In the encoding configuration, typical offline encoding aspects may be considered, such as variable-bitrate encoding, multi-path encoding, content-dependent encoding, sync samples to enable seek and other trick modes, etc.

Such sync samples may be added in regular distance, or if more appropriate also content dependent, for example at scene-change boundaries.

5.5.4.1.2 HEVC-based viewport-dependent OMAF video profile

If the HEVC-based viewport-dependent OMAF video 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 clause 10.1.3.3 of [OMAF].

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. The `Content-Type` in the HTTP header should be set as

```
video/mp4 profiles='hevd' codecs='resv.podv+erpv.hvc1:1.1.6.L93.B0'  
or 'resv.podv+ercm.hvc1:1.1.6.L93.B0, resv.podv+ercm.hvc2:1.1.6.L93.B0'
```

5.5.4.2 DASH Streaming

5.5.4.2.1 HEVC-based viewport-independent OMAF video profiles

If the HEVC-based viewport-independent OMAF video profile or unconstrained HEVC-based viewport-independent OMAF video 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 [DASHIF IOP].

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 HEVC-based viewport-dependent OMAF video profile

If the HEVC-based viewport-dependent OMAF video profile is used for tile-based DASH Streaming, an MPD file is generated with $\sum_{i=0}^{R-1} numMCTS(i) + C$ Adaptation Sets: one for each tile at each resolution $\sum_{i=0}^{R-1} numMCTS(i)$ and one for each of the C extractor tracks. This implies that each Adaptation Set contains Representations with the same resolution. Each of the $\sum_{i=0}^{R-1} numMCTS(i)$ Adaptation Sets (corresponding to each tile at each resolution) can contain a Content Coverage (CC) SupplementalProperty element as defined in [OMAF] 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.

All $\sum_{i=0}^{R-1} numMCTS(i) + C$ Adaptation Sets may contain a quality ranking for each region; but at least the C Adaptation Sets with Representations corresponding to the extractor tracks contain the quality ranking indication. Either the spherical region-wise quality ranking (SRQR) SupplementalProperty element as defined in [OMAF] is used or the 2D region-wise quality ranking (2DQR) SupplementalProperty element as defined in [OMAF] is used in order to signal which tiles have higher quality in respect to other tiles. 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 SRQR or 2DQR descriptor is set according to the SphereRegionQualityRankingBox or 2DRegionQualityRankingBox as described in section 5.5.3 respectively. If SphereRegionQualityRankingBox is present in the ISOBMFF file, the value of shape_type in the CC descriptor of each of the $\sum_{i=0}^{R-1} numMCTS(i)$ Adaptation Sets for each tile at each resolution 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 CC 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 CC 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 clause 5.2.1 of [OMAF].

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 $\sum_{i=0}^{R-1} numMCTS(i)$ 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 $\sum_{i=0}^{R-1} numMCTS(i)$ + 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 HEVC-based viewport-dependent OMAF video 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 SwitchingInterval 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.

5.5.5.1 Encryption

Encrypting media prevents it from being read by a user or client without the correct decryption key. It does not prevent the media from being copied or redistributed but it makes those activities pointless unless the recipient also has access to the decryption key.

5.5.5.1.1 Viewport Independent Baseline Media Profile

The Viewport Independent Media Profile is compatible with commonly deployed DRM functionalities and encryption work flows - current DRM and encryption technology allow a Viewport Independent Media Profile video to be encrypted, transported, decrypted and made available for rendering. The rendering process may differ across players and platforms. Example guidelines for the usage of DRM and security in DASH are provided in the DASH-IF interoperability guidelines [DASHIF10P], clause 7. This provides a good overview of widely deployed adaptive streaming DRM and encryption systems.

5.5.5.1.2 Viewport Dependent Delivery

This section applies to the HEVC-based viewport-dependent OMAF video profile and the simple tiling OMAF video profile.

When the DASH Access engine in the VR Service Platform performs DASH sub-segment concatenation or, when using the simple tiling OMAF video profile, concatenates samples from OMAF tile tracks, it will construct a single ISOBMFF file. This file will contain encrypted data from individual DASH streams for each tile that will make up the frame, concatenated into the single ISOBMFF file.

Each sample in the frame will contain encrypted data received from the DASH stream for each tile included in that the sample.

ISOBMFF supports the definition of encryption metadata at the granularity of a sample, but not different encryption metadata within a single sample. For this reason it is necessary that the key-id and the initialization vector for each part of the sample is the same. The decryption function will decode the byte

ranges indicated by the subsample information stored in the Sample Auxiliary Information within the single ISOBMFF file.

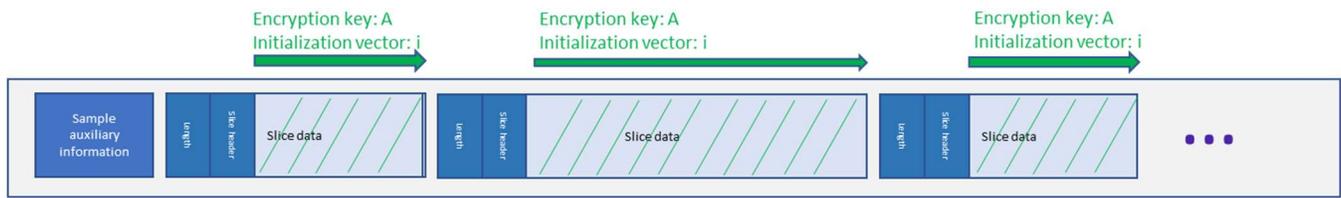


Figure 20: Logical Receiver Model

This restricts the AES encryption mode that can be used – `ctr` and `cbc1` cannot be used. For this reason the recommended encryption mode for viewport dependent media profile at this time is `cbc3`.

Note: This section of the guidelines remains a work in progress and further investigation is required to verify the operation of VR Players and to analyze the performance implications of this approach.

There exists proposals to support the definition of different encryption metadata for different parts of a single sample within an ISOBMFF file. This may allow the support of AES `ctr` and `cbc1` modes also in the future, again with the caveat that player support and performance implications would need to be understood.

5.5.5.2 Watermarking

Watermarking media provides a means to identify the source of the media, or to identify a transition point in a media’s journey. It does not physically prevent sharing of media, nor does it by itself include any mechanisms to restrict playback. Rather, it allows data to be inserted into the media that may be used to trace the source of the media, or more accurately to trace the point where the watermark was inserted. From a media security point of view, this enables the media owner to identify the source of copied or pirated content and either stop the supply at source or take action to recover losses from the source point, or both.

It is not the intent to provide a detailed overview of watermarking in this document, rather to focus on any special considerations when applying watermarking to VR video. A high-level overview of some of the key types of watermarking is provided below, but readers wanting to get a more detailed view can refer to publicly available documents from the Ultra HD Forum (<https://ultrahdforum.org/wp-content/uploads/Ultra-HD-Forum-Guidelines-v1.4-final-for-release.pdf>) and the Streaming Video Alliance (<https://www.streamingvideoalliance.org/download/27623/>).

5.5.5.2.1 Watermarking drivers

Watermarking has overhead associated with it, as with most technologies – it may require extra processing, extra storage or simply extra effort to implement so it is important to look at why one would watermark content, and to understand some of the drivers.

There may, in fact be multiple drivers – for example

- Detecting the source of pirated material.
- Supporting analytics.
- Tracing separate streams and sources.

Additionally, it is worth considering what exactly we want to watermark – this may depend on the driver:

- All parts of all content.
- Only high value portions of content.

The techniques to insert and detect watermarks may not be affected by these considerations, but when they are applied, and what they are applied to, may be.

5.5.5.2.2 Types of watermarking

5.5.5.2.2.1 Visible vs invisible watermarking

For some applications a visible watermark may meet the requirements. A visible watermark may be, for example, a number or logo which appears either throughout the video or at intervals during the video. This number can be associated with a particular copy of the video, for example the copy shared with a certain airline for inflight movies, or the copy broadcast to a particular room for hotel movies.

The two primary disadvantages of visible watermarks are that the visible marks may impact the viewer's enjoyment of the video and that a pirate may simply modify the video to remove the watermark if they can see the watermark clearly.

For the purposes of these guidelines we will focus only on invisible watermarks. Note that although the watermarks are generally invisible to the naked eye, they do of necessity change the image in certain frames of the video and a trained observer can detect small video artefacts in certain circumstances. In fact the 'invisibility' of the watermark is usually an engineering trade-off that can be balanced against the content type, the robustness of the watermark and the required speed of detection. Different applications and use cases may have different weightings for these factors.

5.5.5.2.3 Single stage vs two stage watermarking

Single stage watermarking inserts the invisible watermark into a frame or frames of the video in a single step, as the name suggests. The watermarking detection point reads the watermark or watermarks in the displayed frames.

It is usually deployed on the client side just before or during playback.

Two stage watermarking typically first creates two versions of the media, inserting an 'A' watermark into each chunk or segment of one version and a 'B' watermark into the corresponding chunk or segment of the other version. A second switching function then delivers a unique combination of A and B chunks to each end point, for example a subscriber or distribution point. A watermarking detection function reads in the chunks of video displayed and retrieves the unique combination allowing it to be mapped to the given end point.

Two stage watermarking is typically implemented in the head end or 'server side'

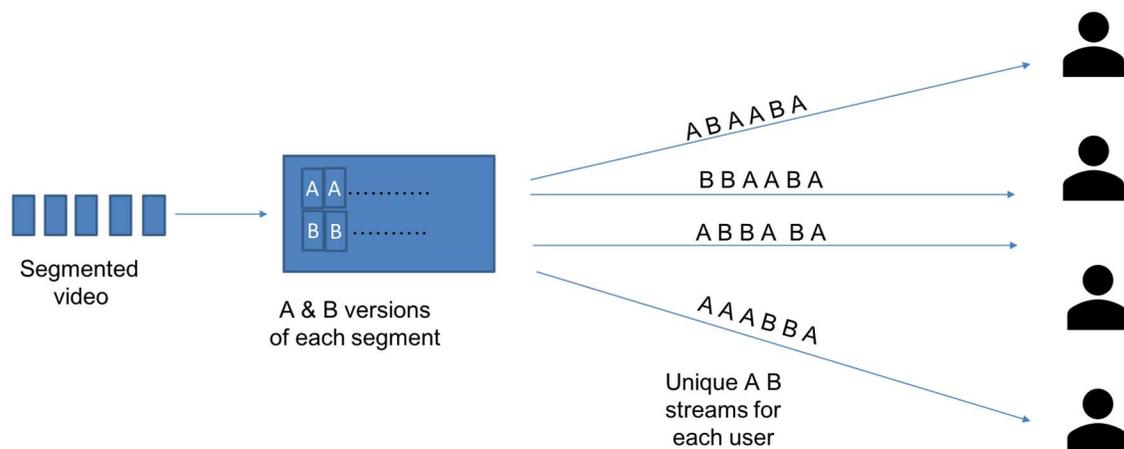


Figure 21: Unique A B marked video streams for each viewer

5.5.5.3 Watermarking in the VR domain – different use cases

It is important to identify the type of piracy that an operator or content owner wants to prevent when selecting the appropriate watermarking approach.

In fact, it is important to be sure that the content itself is content that needs to be protected in this way – as mentioned above, and like most technologies there is an overhead associated with implementing a watermarking solution.

Some examples of scenarios where watermarking may be appropriate in the VR domain:

- A flat 2D movie being shown on a screen in a virtual theatre or living room within a VR environment. The movie itself is not a VR 360 video, but the theatre or living room is. Typically in this scenario it is the movie itself that need to be protected against piracy, not the theatre or room.
- A VR 360 video being shown on a flat screen or in an HMD. From a watermarking point of view, the requirement may be just to watermark and trace the current view port only, or it may be to watermark and trace the entire 360 video.

At this time, these guidelines will focus on the Home Theatre ‘flat’ video VR use case, from the point of view of watermarking.

5.5.5.4 Watermarking Home Theatre ‘flat’ video in VR environments

When considering this use case, it is important to remember that the video itself is not a VR or 360 video, but a regular ‘flat’ video that is being displayed on a surface or ‘screen’ in a VR environment.

5.5.5.4.1 How the content will be pirated

It is useful to consider how a video may be pirated in the VR environment when considering how watermarking can be used to combat the piracy.

- Raw video stream copied

The flat video displayed in a VR environment is delivered as one of more regular streams, e.g. MPEG-DASH, which the client will display on a vertical surface or plane in the VR world. Because the video is delivered as a distinct stream it is reasonable to assume that it might be intercepted and played on a regular non-VR client and screen. In this scenario, the pirating approach, e.g. screen capture or recording the screen, is likely to be the same as that used for non VR video pirating.

- Viewport copied

In this scenario the pirate copies just what is being viewed in the viewport, for example by placing a camera in the same position as a user’s eyes would be in an HMD.

- Flat screen VR on tablet screen copied

This is similar to the viewport capture as the flat screen will typically only be displaying the current viewport. The screen may be captured as a whole or the person pirating the video may point a camera at the viewport display on the flat display screen.

5.5.5.4.2 Guidelines for watermarking Home Theatre ‘flat’ video in VR environments

In general, current ‘regular’ video watermarking techniques should be utilized for VR home theatre flat video.

This means that existing one stage, typically client side, and two stage, typically server side, watermarking techniques should be used to insert and detect forensic watermarks in this type of media.

Insertion technologies in particular, although impacted by some of the considerations in the following sections, should from an architectural point of view be as today's solutions.

Detection technologies, should also remain largely the same, although the interfaces to capture and analyze pirated content may need to allow for the format that the pirated video is provided in – for example if the full VR360 environment is the pirated 'entity' then the detection technology must be able to accept this as an input and to isolate the VR home theatre flat video within it.

5.5.5.4.3 Special watermarking considerations for Home Theatre 'flat' video in VR environments

5.5.5.4.3.1 Resolution

For users using Head Mounted Devices to view home theatre flat video, the resolution of the video will be required to be higher than standard device's displays as the user will be viewing the screen at very close quarters and via magnifying lens.

With lower resolutions, this proximity and magnification will cause the viewer to perceive distinct pixels in the display. Up scaling may allow a partial solution, but ultimately higher resolution videos are expected to be common.

These higher resolution videos, while providing more area in which to insert a watermark in, also provide a challenge to watermarking insertion and detection.

5.5.5.4.3.2 Processing

Higher resolution video will typically require more processing power to insert watermarks into, as each frame is much larger and simply decoding and encoding a frame, even before watermark insertion is added, will require greater processing.

5.5.5.4.3.3 Delay

Related closely to processing, any extra delay added into a stream may be an important consideration, especially for live streams. In simple terms, greater processing power may be required to support real-time watermark insertion with greater resolution.

5.5.5.4.3.4 Viewing

Users of HMD's watching home theatre flat video will typically be viewing the screen at a similar angle to a movie theatre – i.e. looking up slightly at the screen. This is typically a different viewing angle than many home television and computer displays. This different viewing angle will require further study to understand whether there are any impacts on the visibility of hidden watermarks.

5.6 Guidelines for Service Platform Developers

5.6.1 Overview

Figure 22 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.

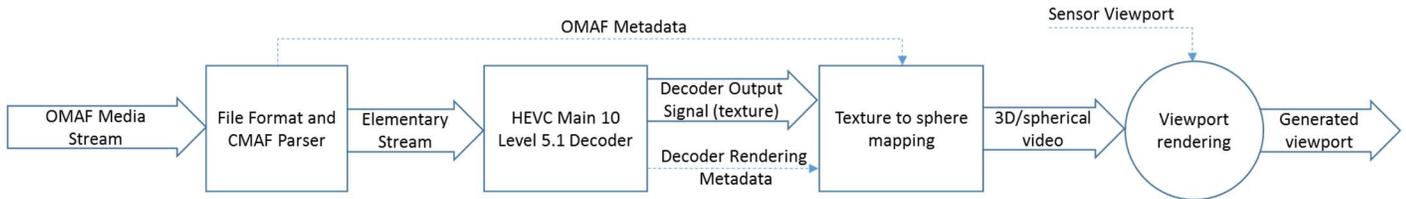


Figure 22: Logical Receiver Model

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 23. The available metadata in the OMAF stream or the SEI messages is summarized as follows:

- Region-wise packing information as constraint in clause 7.3.1.2 of [OMAF]
- 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. Additional information regarding rendering the spherical image using SEI messages can be found in both [ADDSEI] and [OMAF]

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.

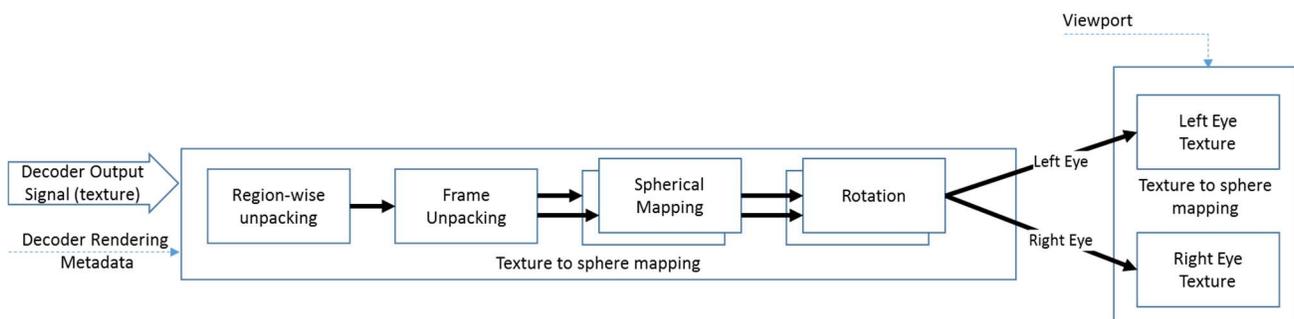


Figure 23: Rendering and viewport generation

5.6.2 Rendering Process based on SEI messages

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 [ADDSEI] section D.3.41.6, the sample location remapping process is documented.

For the restricted scheme based for the viewport-independent media profile, to remap color sample locations of a region-wise packed picture to a unit sphere, the following ordered (also indicated in Figure 23) 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 [ADDSEI] section D.3.41.6.4. 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 [ADDSEI] section D.3.41.6.6. 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 [ADDSEI] section D.3.41.6.2.

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 [ADDSEI] section D.3.41.6.3. Otherwise, the global coordinate axes are identical to the local coordinate axes.

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 [DASH-PUSH]. If both are supported at CDNs and servers a more efficient network performance can be achieved.

5.6.4 Security

It should be understood that for encrypted media there is no access to any decoded 360° video pixels outside of the secure media pipeline and the video bit stream should contain all the information required to recover a 360° video. Additionally, the graphics subsystem must be capable of receiving any external inputs required to produce the final display, for example orientation sensor inputs.

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.
- 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.

5.7.1 Distribution and delivery

5.7.2 Decoding and Rendering

If the HEVC-based viewport-independent OMAF video 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 HEVC-based viewport-dependent OMAF video 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.

If the unconstrained HEVC-based viewport-independent OMAF video profile is used to distribute 8K@60fps content, an HEVC Main 10 profile, Main tier, Level 6.1 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.

5.7.3 APIs

Figure 24 shows an OMAF-DASH Client model for illustration.

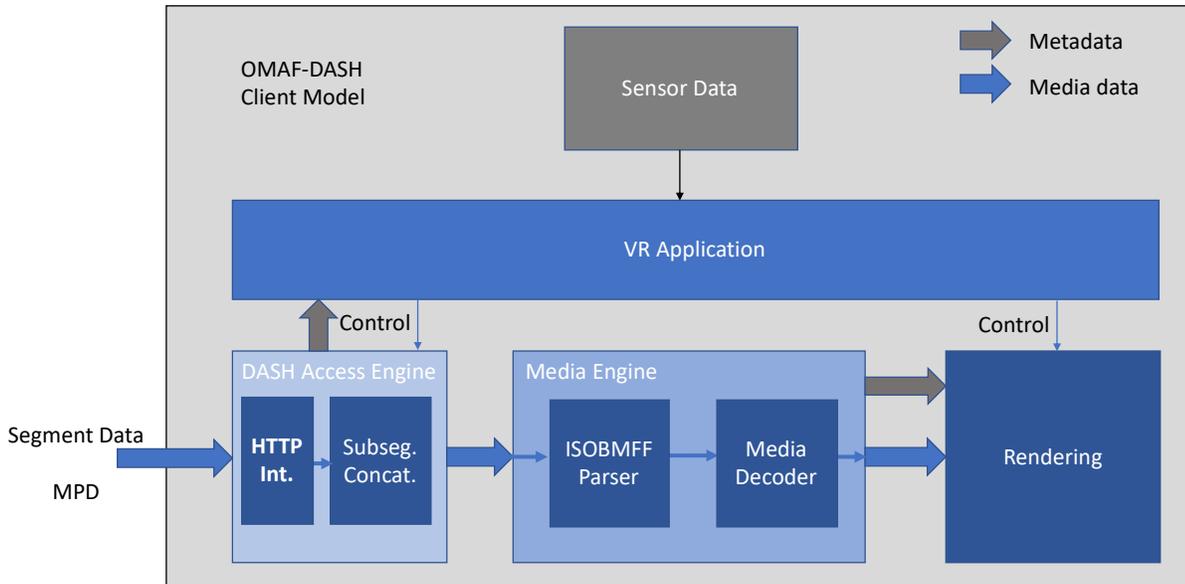


Figure 24: OMAF-DASH Client model with interfaces

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

DASH Access Engine

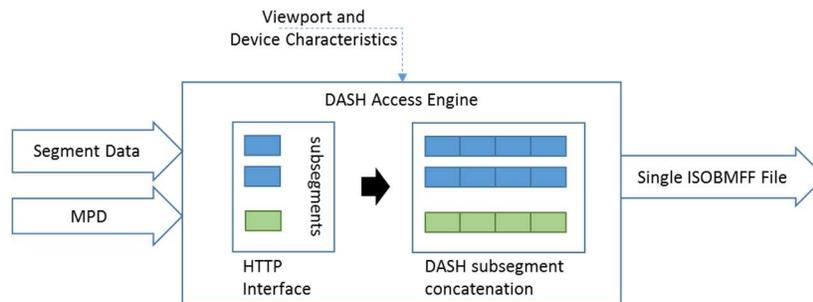


Figure 25: DASH Access Engine for HEVC-based viewport-dependent OMAF video profile

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 [DASH/PUSH] 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. In order to do so, the DASH Access Engine provides the relevant metadata to the VR Application, which might include the MPD and Initialization segments or equivalent metadata extracted from them.

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 clause 10.1.3.3 of [OMAF] which correspond to the HEVC-based viewport-dependent OMAF video profile.

Since each of the frames consist of multiple tiles that together depict the whole covered scene (with a 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

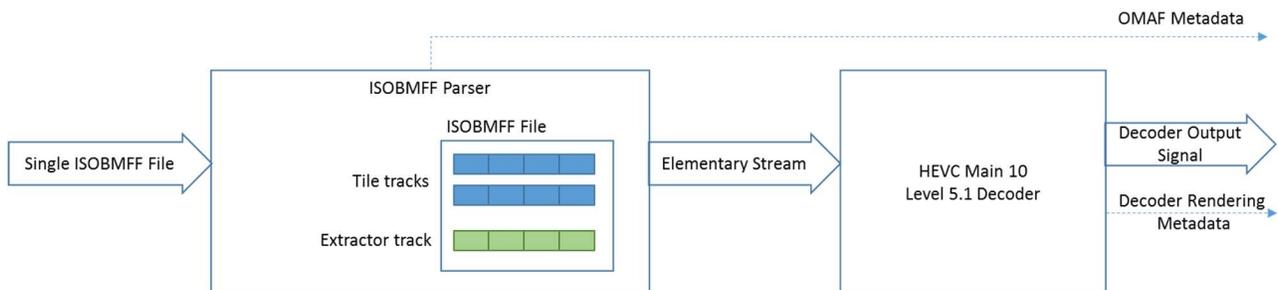


Figure 26: DASH Media Engine for HEVC-based viewport-dependent OMAF video profile

The input to the Media Engine is a single ISO BMFF file, which is the output of the DASH access engine (see Figure 25). 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

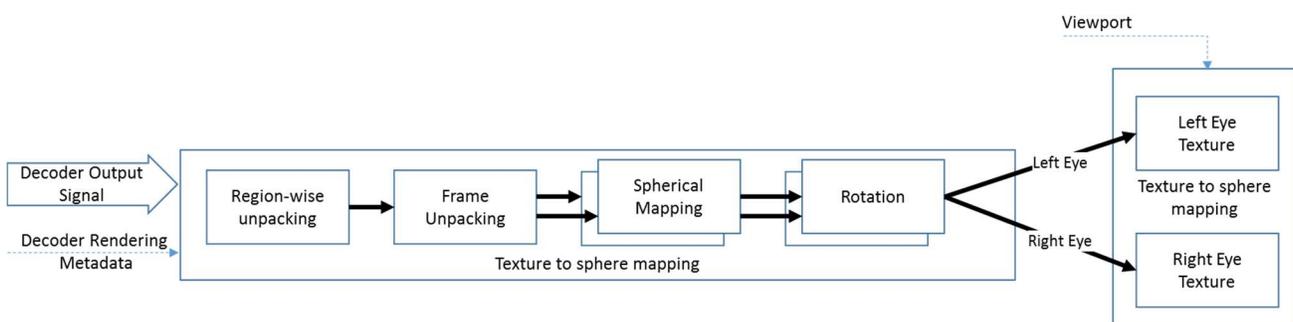


Figure 27: Inverse Projection/Renderer

A more detailed flowchart of the logical steps of rendering chain is provided in Figure 27. 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`
- Padding (“Indicates whether there is padding in the packed frame”)
 - Decoder Rendering Metadata:
 - Equirectangular projection SEI (only for ERP and simple padding) or Region-wise packing SEI
 - OMAF Metadata:
 - `RegionwisePackingBox`
- Stereoscopic 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:
 - Sphere rotation SEI
 - OMAF Metadata:
 - `Rotation Box`
- Coverage restriction (“Indicates content coverage”)
 - Decoder Rendering Metadata:
 - Region-wise packing SEI
 - OMAF Metadata:
 - `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 27, 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, as provided by the VR Application as control data in Figure 24.

Furthermore, implementation may combine several of the logical steps to reduce memory consumption and processing time compared to a naive 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.

6 Vertical 2: Live VR Streaming

6.1 Description of Vertical

Live VR provides an immersive experience that certainly comes close to “being there”. For many people it is not possible to experience some live sports events or artistic performances, either because the “best seats in the house” are already sold or due to geographic constraints. A well planned Live VR event places omnidirectional cameras the prime viewing locations allowing the same viewing experience to be sold multiple times.

This vertical primarily addresses the live acquisition and distribution of VR360 content to emerging devices.

Annex D provides information on live VR deployments in the industry. In particular, Annexes D.1 and D.2 describe live VR deployments from TiledMedia and Intel, respectively.

6.2 Guiding Example Use Cases

6.2.1 Actors and set-up/equipment

José and María are ready to watch a live sport event in their living room. They turn on the TV and they realize that the event also is broadcasting a secondary Live VR sport feed. José takes his tablet and via an application they start to see the secondary VR360 video content synchronized with the main TV event. Their son, Pedro, joins them in the living room and he decides to go fully immersive to watch the VR content. He uses his mobile phone in a VR headset using a viewer provided by the phones manufacturer.

For the main sport event (the non-VR content) Broadcasters can re-use legacy TV receivers (HD, UHD-1 Phase 1 and UHD-1 Phase 2). Any VR headsets can be used if they support the video format used by the broadcasters and if they support Internet connection. Depending on José and María’s Internet subscription, the tablet can be connected to Internet via local WIFI or cellular connection. The VR service can be over the top (OTT) or by other managed network to reduce the latency.

6.2.2 Use case description (short)

The content acquisition of the Live VR feed is captured using additional cameras dedicated for this purpose (in addition to the traditional cameras). In the venue, there might be several locations where the VR content is captured. The director may choose which Live VR feeds are distributed to home delivery, either by producing one single VR feed, or by giving the choice to the end user to select which one to watch (in the secondary screen), or a by offering a combination of a produced feed and individual cameras.

The broadcaster has support for hybrid delivery of live TV content, that is, the main Live TV event is broadcasted using “traditional means” (such as terrestrial, satellite, cable, etc.) while the VR content is delivered via unicast-streaming service in parallel with the broadcast content⁶. To reduce the transmission bandwidth, a viewport aware delivery method is used and thus, the VR headset is required to transmit information to the delivery server. Note that broadcasting of the VR content (i.e. with no bidirectional communications path) is not considered in this use case.

6.2.3 Use case description (detailed)

In this use case, it is assumed that there are one or more supplementary live VR feeds for home delivery in addition to the main live feed that it is broadcasted by “traditional means” such as terrestrial, satellite,

⁶ This means that both feeds are not necessarily in sync. The current Live VR scenario does not consider technologies to synchronise both feeds

broadband or cable delivery (see Figure 28). The Live VR feed is captured using additional cameras dedicated for this purpose. The output of these cameras maybe already stitched or may require stitching in a secondary process. The stitching process may be supervised by director at the outside broadcast (OB) truck or at central studio at the main broadcast facility. In the venue, there might be several locations where the VR content is captured. The director may choose which Live VR feeds are distributed to home delivery, by producing one single VR feed, by giving the choice to the end user to select which camera, or by offering both a produced feed and individual camera positions.

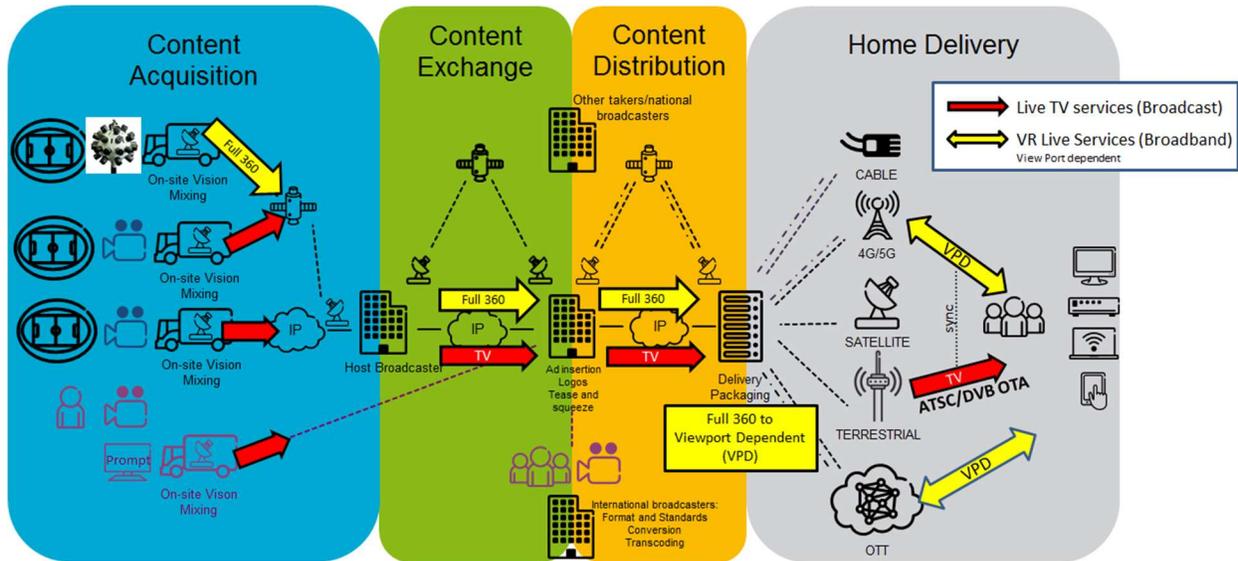


Figure 28: Live workflow for Television augmented with VR⁷.

As in traditional Live TV, the content may be exchanged between other affiliates and even distributed to other regions including international destinations. This exchange occurs using the full 360° video content and is normally at very high quality, at least 8K and possibly even higher). Finally, a local broadcaster (or service provider) delivers the Live VR content to the home with additional bandwidth optimization, such as viewport dependent delivery.

In this use case, it is assumed that the broadcaster has support for hybrid delivery of Live TV content, that is, the main Live TV event is broadcasted using “traditional means” with the VR content delivered in parallel via unicast streaming. It is also assumed that viewport-aware delivery relies on the client device to request the appropriate content from the server or CDN. This means that VR content broadcasting (i.e. without a bidirectional connection) is not considered in this use case.

6.2.4 Consumption of VR content

In order to facilitate the adoption of VR services to a wider audience, this use case also presumes the distribution of VR content to 2D displays, such as tablets, smartphones, and possibly set top boxes or even smart TVs. In the case of tablets and smartphones, the viewport is generally controlled by touching the screen or moving the device, and an STB or TV by a remote control. Other scenarios may utilize a connected smartphone or tablet to function as the swiping input mechanism for the STB or smart TV.

⁷ VPD: Viewport dependent: a portion of the 360 video is delivered to reduce bandwidth and complexity of decoding at the receiver (e.g. instead of decoding 8K for full 360, VPD requires lower decoding capabilities in the device).

6.3 Reference architecture and workflow

The following diagrams depict the parallel content workflows where the live event is captured using traditional cameras (in the upper part of the diagrams) and omnidirectional cameras (in the lower part of the diagrams)

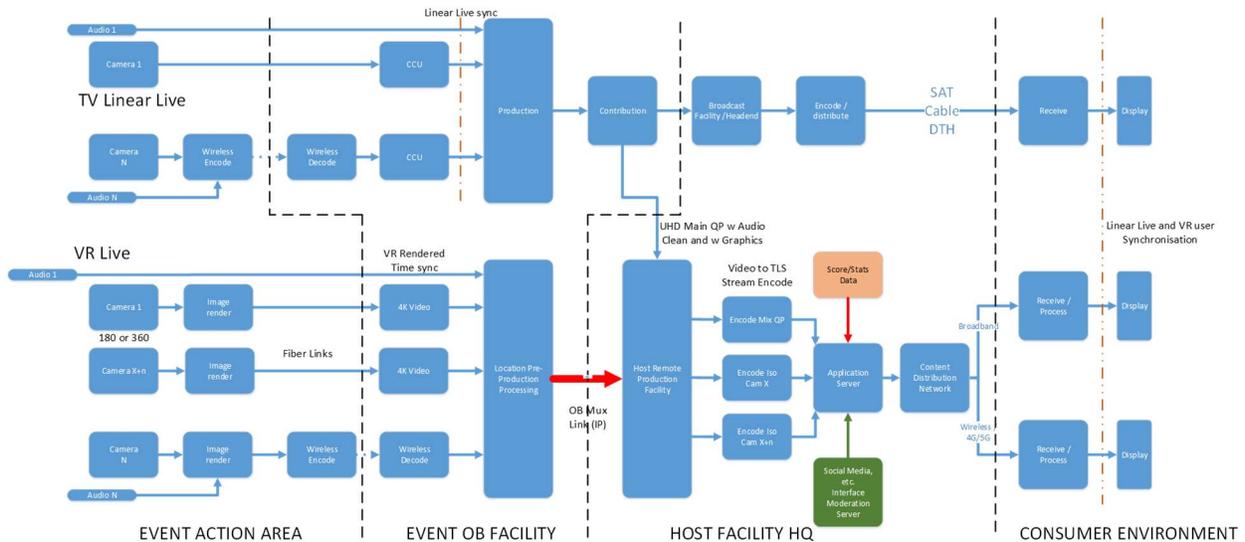


Figure 29: Linear TV and VR Live workflows

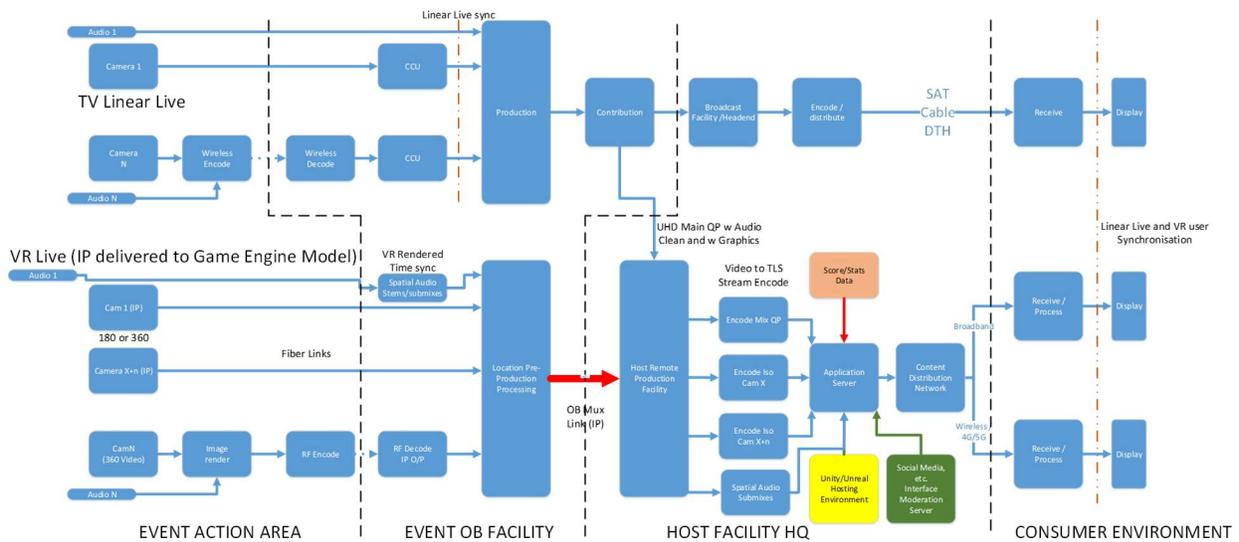


Figure 30: Linear TV and VR Live with - IP connected omnidirectional camera

It can be observed from Figure 29 and Figure 30 that similar workflows are utilized for traditional Live TV and Live VR. This should not be unexpected as the same operations of acquisition, production, processing and delivery apply for all entertainment experiences although the nature of the functions may differ.

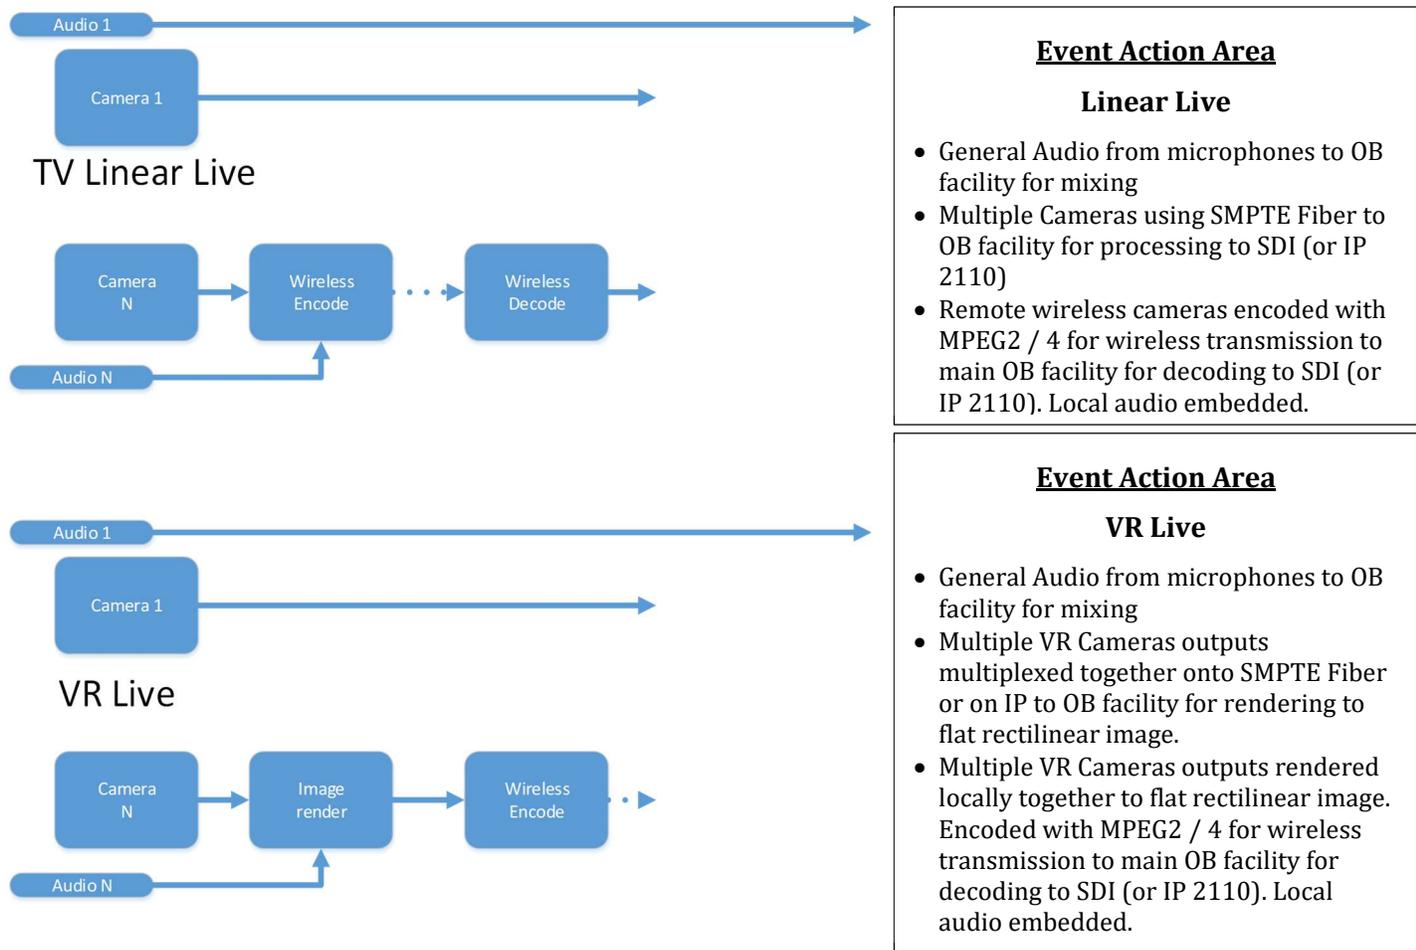
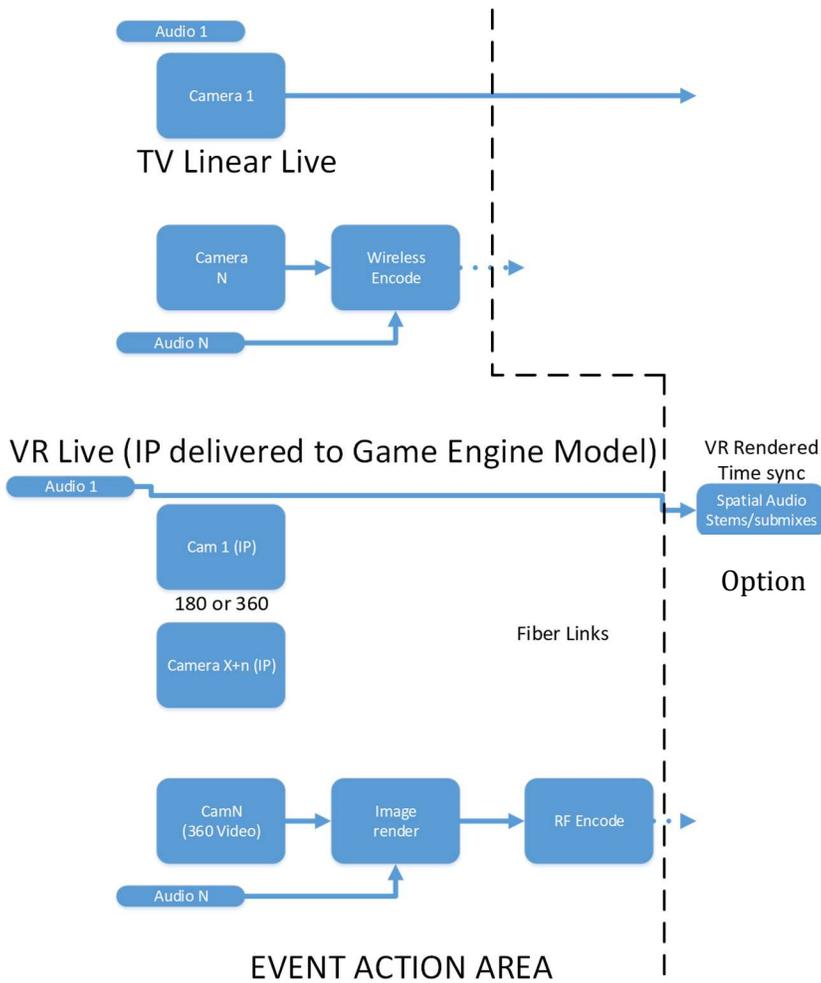


Figure 31: Event Action Area: General differences in workflow between TV Linear Live and VR Live

Current broadcasting production workflows can be used for VR production by making use of an equirectangular projection in a vision mixer, for instance a 4K broadcast infrastructure will work without problems if the full 360 VR video is captured in 4K. If VR is captured in 8K video or even higher, then VR and traditional TV Linear live workflows will start to differ. It is important to note that in order to take full advantage of the OMAF format with view port dependent profile (as defined in sections 4.2.2.1.3 and 4.2.2.1.4), it would require to capture a 360 video with at least 8K resolution.



Event Action Area

Linear Live

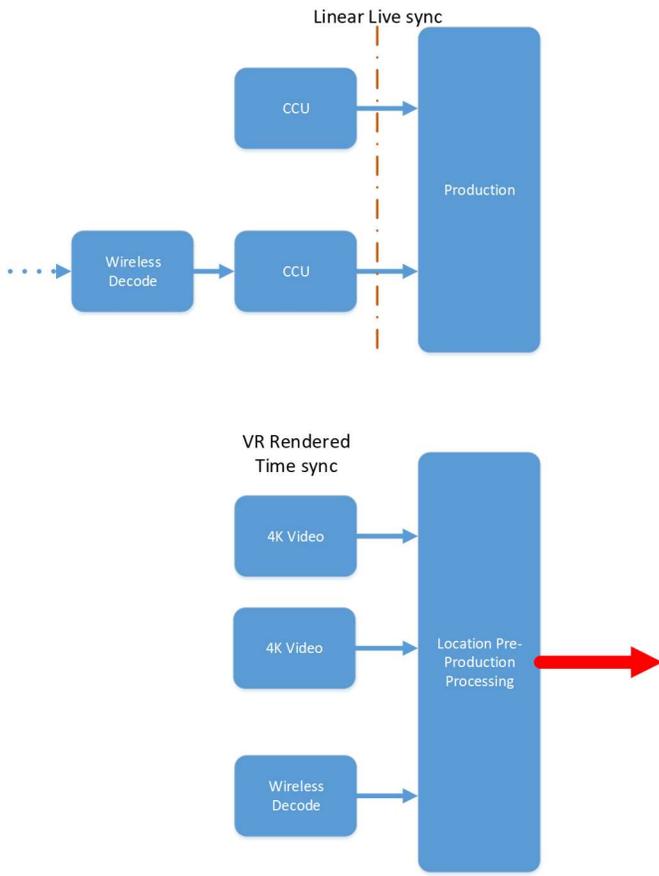
- General Audio from microphones to OB facility for mixing
- Multiple Cameras using SMPTE Fibre to OB facility for processing to SDI (or IP 2110)
- Remote wireless cameras encoded with MPEG2 / 4 for wireless transmission to main OB facility for decoding to SDI (or IP 2110). Local audio embedded.

Event Action Area

VR Live

- General Audio from microphones to OB facility for mixing
Note: if conventional Broadcaster is present – feeds may be taken here and Delayed
- Spatial Audio Sub-mixer may be derived for Per-Camera view Audio feeds
- Multiple VR Cameras outputs multiplexed together onto SMPTE Fiber or on IP to OB facility for rendering to flat rectilinear image.
Note: on most current “Live” camera systems the render is on-board
- Multiple VR Cameras outputs rendered locally together to flat rectilinear image. Encoded with MPEG2 / 4 for wireless transmission to main OB facility for decoding to SDI (or IP 2110). Local audio embedded.

Figure 32: Event Action Area for VR Live (IP delivered to Game Engine Model)



EVENT OB FACILITY

Event OB Facility

VR Live

- It would be feasible to implement remote production for VR Live
- Camera feeds from a small Head-end Location Unit would be fed via ASI/IP to a remote Gallery for Processing
- OB Parking space is generally at a premium
- Hardware would solely comprise fiber camera interface and ASI interface to link to Broadcast.

Figure 33: Event OB Facility

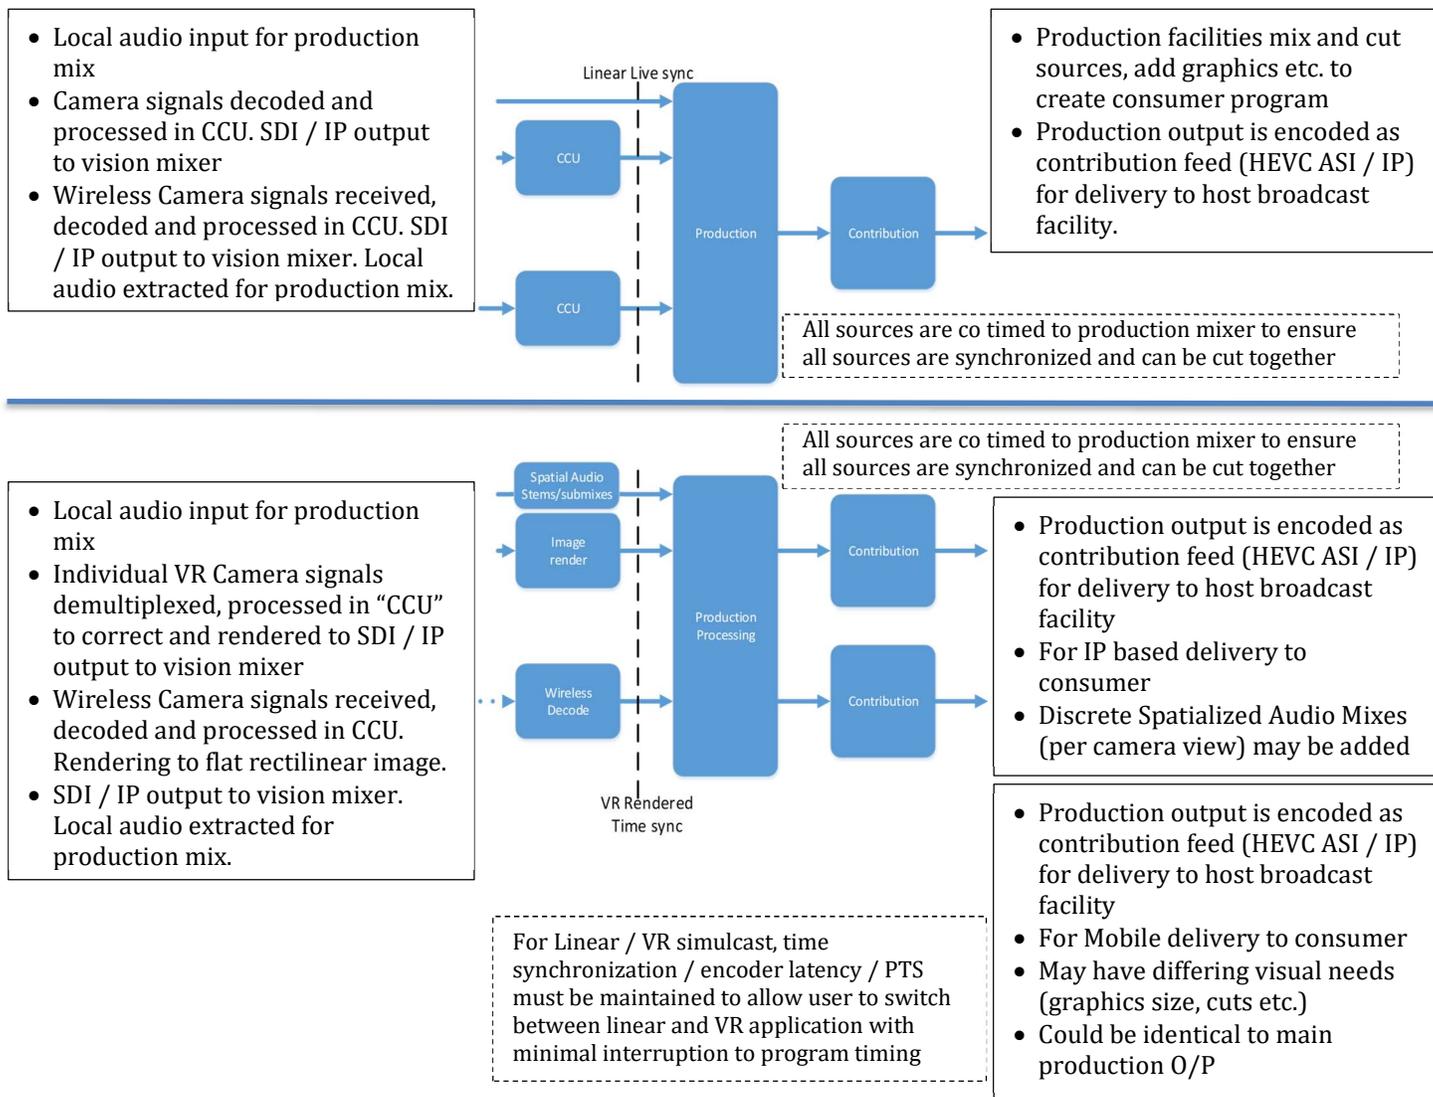


Figure 34: Event OB Facility: differences in workflow between TV Linear Live and VR Live

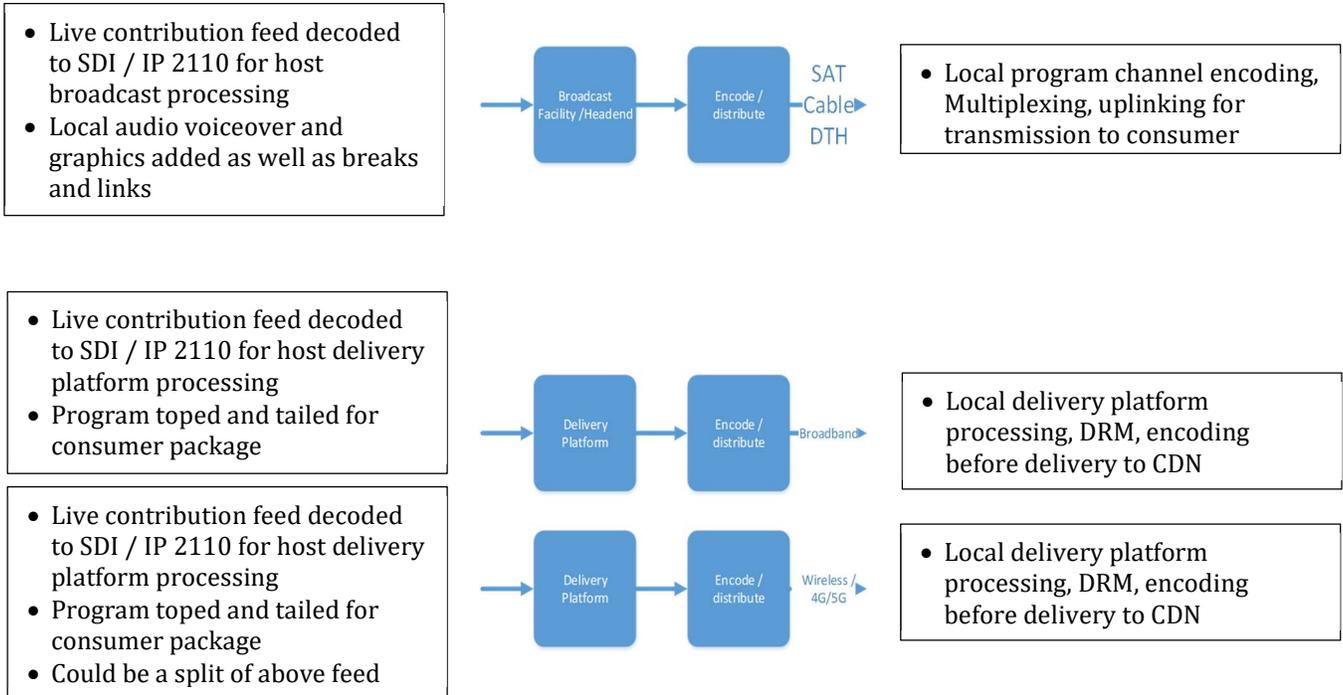
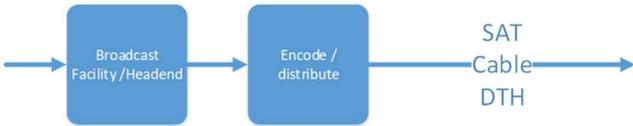


Figure 35: Host Facility HQ: differences in workflow between TV Linear Live and VR Live

- Live contribution feed decoded to SDI / IP 2110 for host broadcast processing
- Local audio voiceover and graphics added as well as breaks and links



- Live contribution feed decoded to SDI / IP 2110 for host delivery platform processing
- Programme topped and tailed for consumer package

- Live contribution feeds (IP) for platform processing
- Hosted in Game Engine environment
- Spatial Audio submixes (per Camera View)
- Picture feed optimisation for best Game Engine performance

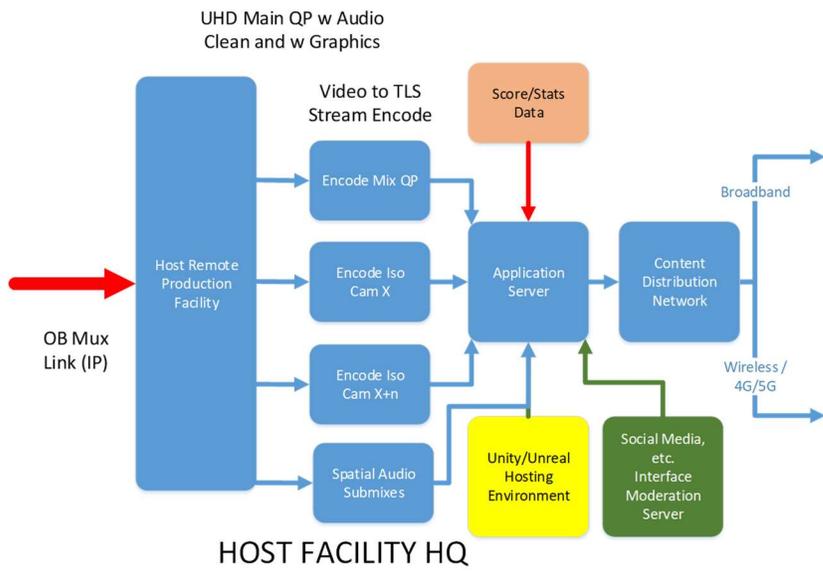
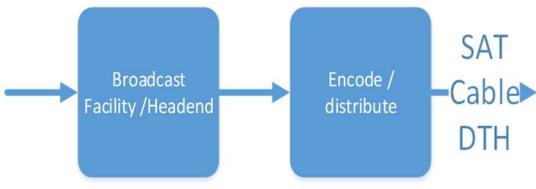
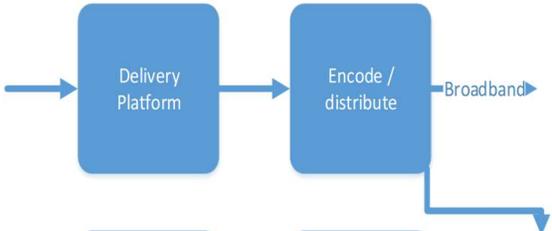


Figure 36: Host Facility for VR Live (IP Delivered to Game Engine)

- Live contribution feed decoded to SDI / IP 2110 for host broadcast processing
- Local audio voice over and graphics added as well as breaks and links



- SDI Live contribution feed decoded to SDI / IP 2110 for host delivery platform processing
- Programme topped and tailed for consumer package



- SDI Live contribution feed decoded to SDI / IP 2110 for host delivery platform processing
- Programme topped and tailed for consumer package
- Could be a split of above feed

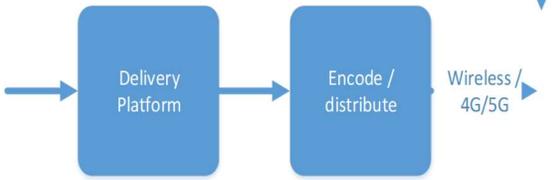


Figure 37: Host Facility HQ

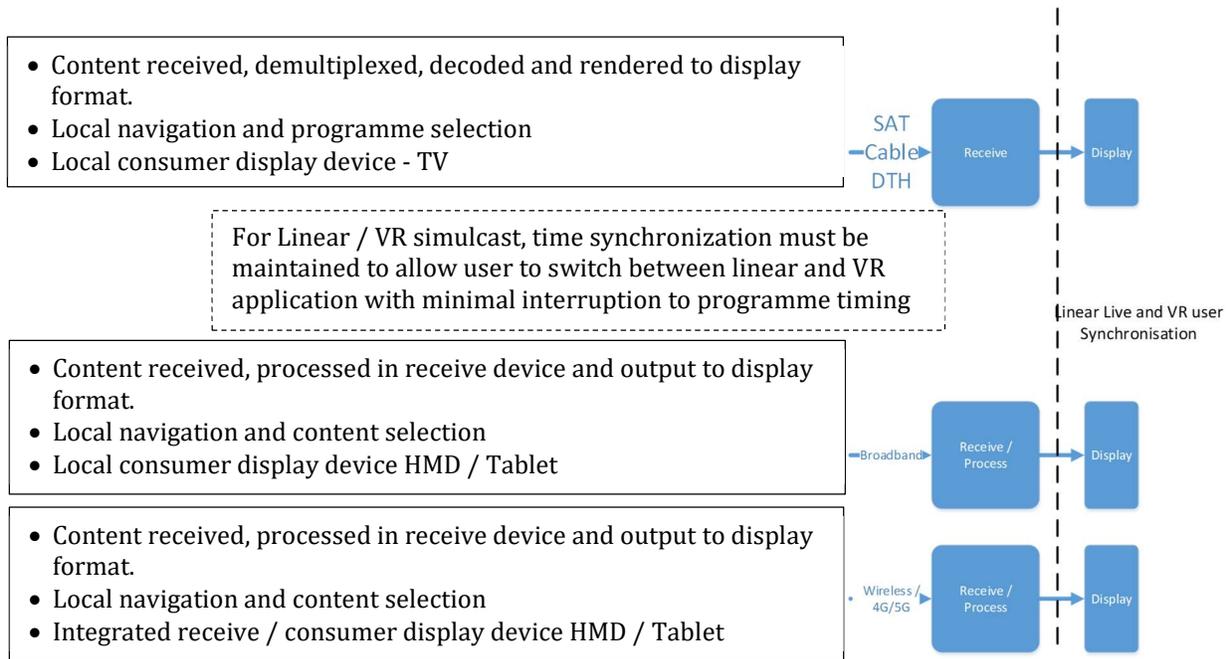


Figure 38: Consumer Environment: differences in workflow between TV Linear Live and VR Live

6.4 Technical enablers

6.4.1 Production Systems for Live events

6.4.1.1 Considerations on production infrastructure

6.4.1.1.1 Best practices for camera positioning for Live 360 Video

VR is a new medium and live 360 video production is perhaps the most intense, stressful, and gratifying way to create VR content. As always, careful planning reduces risk and helps ensure a smooth production.

As identified above, many aspects of live 360 production are shared with standard live video production; these common aspects are out of scope for this document. This section focus on the unique aspects of production for 360 video.

Camera placement

“Camera represents the best/optimal viewing position”

Camera placement is everything. Compare the two camera locations in Figure 39.



Figure 39: Two different camera locations: far from and close to the event

Closer is always better. Always fight for camera placement. 360 video works when the camera is close to the subject; it becomes boring and disengaging when the camera is far away. It is the single most important thing you can do to make a compelling production. Remember that you are giving your audience an opportunity to see what they could not otherwise experience. With digital pan⁸, you can re-center your area of interest but you cannot zoom, so get as close as the production will let you. The use of digital pan needs to follow the same criteria as a moving/rotating camera or 360 scene, if you move/pan/rotate too quickly the viewer will become spatially lost in the scene.

Wide lenses used in 360 cameras tend to push the action “further away” and loses the consumer association with the experience. Higher quality or higher performing stitching is needed for closer objects, and avoiding the action moving across a stitching point is desirable. The use of wide angle lenses brings the advantage of reducing the number of stitching lines in the captured video however cameras should be positioned in such a way that the close up action does not cross the stitching areas.

Moving Cameras

Despite what you might have heard about the impact of moving a 360 camera, it can be used to enhance certain aspects of the production. Camera motion can have a nauseating effect on many viewers, information on the effects of 360 camera motion can be found in section 4.1.2.3.

As long as you keep your horizon level, and the motion is even and slow, motion adds enormous depth and interest to your productions. Combined with motion parallax, stereo capabilities provide a compelling sense of immersion. With a digital pan feature, you can maintain the user's viewing direction even when using jibs and booms to move in and out from the action. Experiment with jibs, booms, dollies, and other motion platforms and develop your own technique.

Multi-camera Shoots

Like multi-camera television shows, 360 video productions are much more compelling with a skilled director managing the viewers' attention across multiple camera viewpoints. Use the following multi-camera best practices as a starting point, but experiment and draw your own conclusions based in the type of event and the primary target audience.

- Hard cuts are usually disconcerting for users with an HMD. Start with a ½ second crossfade and adjust according to the action in the scene and drama that is desired.

⁸ Digital Pan – rotation of the image around the Y/vertical axis in the production process before it is distributed to the viewer. Keeps the primary view “front and centre”

- Use digital pan to set the most salient point of a camera's feed at the center of the frame before you switch to it. This feature is extremely effective when used to match viewpoints before switching cameras.
- The composition of a VR360 panorama typically may impose an extra delay on what the director sees. He or she needs to anticipate this delay in commands to jib operators, who of course operate in real time.
- Confidence monitors for each omnidirectional camera are nice, but you don't strictly need them - after all, any camera feed shows you the entire scene in equirectangular view. With a little practice it's possible to direct from the program feed alone.

LED Light Flicker

Indoor scenes for sports and concerts are often lit using LED lighting. In some cases, the refresh rates on these lights will cause flickering of the captured images. This is a common problem across all digital sensor technology.

To reduce or eliminate this effect, try reducing the camera shutter speed as low as you can – this will allow more LED refreshes to occur during a single shutter, and can reduce or eliminate flickering. Temporal noise reduction may help to reduce this problem but can have an impact on sharpness.

As a very last resort, you can overexpose the image and compensate with color correction during post production, however this will reduce scene definition, especially in brighter areas.

Direct Sunlight

Direct sunlight can hit individual sensors causing lens flares and color variations. Use per-lens color correction capability to compensate for these anomalies. Color compensation for a single sensor can be accommodated using differential correction in relation to other sensors.

Note: Do not use the auto white- and black-balance features if you're color correcting individual lenses.

6.4.1.1.2 SMPTE 2110 (IP based) vs SDI (legacy infrastructure)

Early workflows for VR360 content made use of customized applications of devices that were repurposed from regular 2D production. This included the use of SDI (Serial Digital Interface) for transporting uncompressed video from camera rigs to processing devices and between processing devices. The Nokia OZO camera, with its 8 wide angle lenses and 8 microphones, made use of a single link 1.5G HD-SDI interface to deliver proprietary compressed sensor data to a powerful computer running the OZO Live suite which performed the stitching, rectification and other content processing and formatting functions. Along with several GPUs, the computer required dedicated PCIe SDI boards. While SDI provides a suitable interface for “close” connections, inherently fixed nature lacks flexibility and dynamic configurability.

The physical limitations of coaxial cables used in SDI along with the added hardware for converting SDI signals into data streams that can be processed in software environments, led to the development of the set of protocols in the SMPTE ST 2110 family which utilize RTP over UDP as the underlying transport mechanism. At the same time, the IP switching and routing fabrics being used in IT environments were becoming more robust and capable of reliably transporting multi-gigabit services.

The general ambition for SMPTE ST 2110 is as a replacement of coaxial SDI interfaces in the production suite, including the venue or studio and any outside broadcast facilities. Since ST 2110 is intended for the carriage of uncompressed audio and video, it is currently not well suited as a contribution format to cloud-based or remote production facilities.

The following documents are available in the SMPTE digital library

- SMPTE ST 2110-10/-20/-30 — addressing system concerns and uncompressed video and audio streams
- SMPTE ST 2110-21 — specifying traffic shaping and delivery timing of uncompressed video
- SMPTE ST 2110-31 — specifies the real-time, RTP payload based transport of SMPTE ST 291-1 Ancillary (ANC) Data packets related to digital video data streams, over IP networks, referenced to a common reference clock.
- SMPTE ST 2110-40 — maps ancillary data packets (as defined in SMPTE ST 291-1) into Real-Time Transport Protocol (RTP) packets that are transported via User Data Protocol/Internet Protocol (UDP/IP) and enables those packets to be moved synchronously with associated video and audio essence streams

The SMPTE ST 2110 Overview [OV2110] provides an introduction to the ST 2110 suite of standards.

As we look towards the role that ST 2110 can play in the VR domain, the follow usage scenarios are feasible.

Omnidirectional 360 camera with inbuilt stitching

A single omnidirectional camera rig with either full or partial spherical coverage looks to use ST 2110 for delivering the panoramic scene via IP to additional processing systems. The camera rig includes built in stitching software which spatially composites each sensors bit stream and outputs in “raw” format. The camera rig also includes multiple microphones to capture the “sound field” which is processed by an inbuilt processor and output as a bit stream that may or may not be temporally aligned with the stitched video output.

Omnidirectional 360 camera without inbuilt stitching

A single omnidirectional camera rig with either full or partial spherical coverage looks to use ST 2110 for delivering the captured scene via IP to additional processing systems. The camera rig outputs each sensors bit stream as a separate “track” in an uncompressed format. Any microphone sources are also output as separate tracks. In order to reconstruct an omnidirectional visual scene and its corresponding sound field, orientation information (often referred to as intrinsic/extrinsic metadata) should be provided for each track.

Multiple Omnidirectional 360 cameras with/without inbuilt stitching

At a live event, multiple omnidirectional camera rigs are deployed at strategic locations. These cameras may or may not include inbuilt processing of the video and audio data (according to the “Omnidirectional 360 camera with inbuilt stitching” and “Omnidirectional 360 camera without inbuilt stitching” scenarios presented above). All video and audio bit streams are sent over an IP contribution link to a cloud based processing system which creates the necessary composite view. In order to allow the positions of each camera rig to be shown within a single omnidirectional view, accurate positioning information (either global or relative) is required.

Personal Point-of-View

At a sporting or artistic performance, several fixed wide angle IP connected cameras are positioned at various locations (for example, in the goal mouth, in front of the band) and the content is uplinked to a cloud based processing center where a virtual PTZ-style view can be created for each viewer based on their pose along with the ability to hop between camera positions.

Follow my favorite

This scenario uses fixed wide angle cameras such as that described in “Personal Point-of-View” but the cloud based processing system performs object/actor/player recognition across all the IP uplinked camera feeds and creates a view tailored to the desires of the viewer.

Volumetric Acquisition

Volumetric capture systems use large numbers of fixed, calibrated inward facing cameras to record still or motion objects in real time. The feed from each camera needs to be uploaded to a local or cloud service such that processing, including the final production of a fixed or dynamic 3D model (using point clouds or texture-and-mesh) can be developed.

6.4.1.1.3 Wireless connectivity from camera to production using cellular networks

3GPP Release 15 [3GPPR15] supports a protocol called Framework for Live Uplink Streaming (FLUS) to stream point-to-point uplink live feeds (e.g. media from wireless cameras) over the cellular network to a second party. In TR 26.939 [FLUS-GUIDE] and subsequent updates the main use cases have been described, and include:

- capture of 360° video through a VR camera and its sharing to a social network
- live uplink 360° video from a drone or any other moving vehicle (or even a stationary device)
- Immersive media conversations where the 360° video is not tightly synchronized with the bi-directional speech streams.
- Traditional rectangular media production.

The Technical Specification 26.238 [FLUS-PROT] specifies the actual protocol suite for FLUS. RTP is used as transport protocol for media in order to guarantee low-delays in transmission. The considered media are 360° video and UHD video. At the time of writing of this document, 3GPP is working on Enhancements to the FLUS framework, known as E-FLUS. This work item is expected to be part of 3GPP Release 16 [3GPPR16] and be ready in September 2019.

New features will include

- Mechanisms for an increased QoS flexibility (e.g., lower delays for drone video upstreaming, compared to a typical conversational video service)
- Interfaces for 3rd party content providers to request QoS support;
- Support for network-based media processing (e.g., for video stitching operations, media transcoding, content reformatting)
- Definition of 3GPP media codec profile(s) for FLUS.

The work will eventually be documented in TS 26.238⁹ and TR 26.939¹⁰, but most recent work at the time of writing of this document can be found in S4-190179.zip¹¹.

⁹ <http://www.3gpp.org/DynaReport/26238.htm>

¹⁰ <http://www.3gpp.org/DynaReport/26939.htm>

¹¹ http://www.3gpp.org/ftp/TSG_SA/WG4_CODECS/TSGS4_102_Bruges/Docs/S4-190179.zip

The requirement working group in 3GPP SA1 is working on a more generic set of requirements for Audio Video Production. It is expected that when the work is finished, the result likely to go to 3GPP SA4 for specification development.

6.4.1.1.4 Live VR production events

Live Streaming demonstration of Ultra-High Resolution 360° Video with MPEG-OMAF and HEVC Tiles

This section describes a demonstration of 8K resolution 360° video captured with a Fraunhofer HHI OmniCam-360¹² being streamed live to mobile devices using the MPEG-OMAF standard in combination with HEVC compression as it is described in section 4.2.2.1.3 of these guidelines.

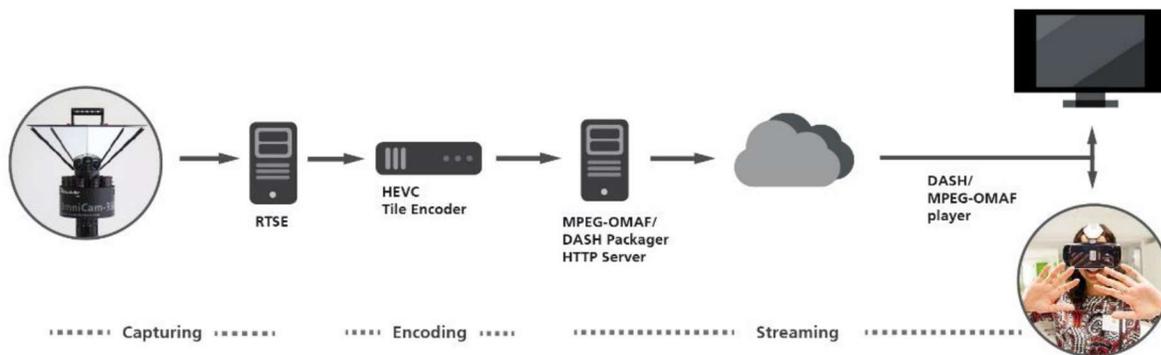


Figure 40: Continuous Live Streaming of VR360-Degrees 10K-Video with MPEG-OMAF and HEVC Tiles

The demonstration used the viewport-dependent profile of MPEG-OMAF standard (see section 5.5.3.2) to stream 6K CMP (produced from an 8K ERP) in real time to a Samsung Gear VR device capable of decoding HEVC Main 10 Level 5.1. The demonstration showed the implementation of OMAF in a continuous real-time transmission chain. Figure 40 illustrates the complete live-chain with the following key components:

Capturing

The OmniCam-360 from Fraunhofer HHI produced an 8K ERP using the Real-Time Stitching Engine (RTSE) whereby the eleven individual camera segments were stitched in real time to form the 8K ERP.

¹² <https://www.hhi.fraunhofer.de/en/departments/vit/technologies-and-solutions/capture/panoramic-uhd-video/omnicam-360.html>



Figure 41: Fraunhofer HHI Omnicam-360

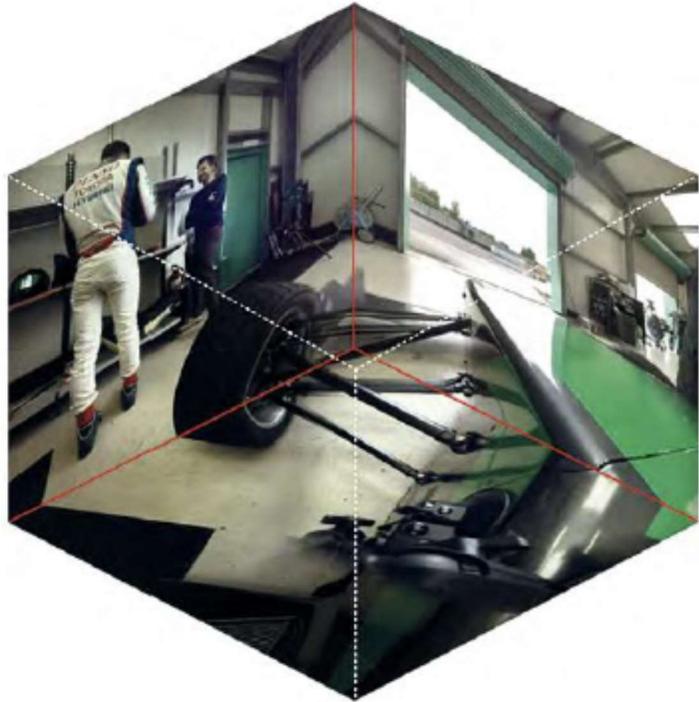


Figure 42: Example of a Cubemap projection

Figure 41 shows the current Omnicam-360. Ten HD cameras and one UHD camera are used to take pictures with a panorama resolution of approx. 10,000 x 3,300 pixels at an opening angle of 360 horizontal x 120 vertical degrees. The cameras used here support 25, 30, 50 and 60 frames per second and have a global shutter and genlock to ensure perfect stitching of the individual camera segments even under critical conditions. The camera signals are processed into a seamless panorama in real time with a latency of five frames.

The resulting panorama was converted into Cubemap projection format. In the Cubemap projection, as shown in Figure 42, the 360° view is projected onto the six inner sides of a cube, each with a resolution of 2,048 x 2,048 pixels.

Encoding

The software-based live encoder from Fraunhofer HHI integrates HEVC tiles according to section 5.5.3.2.



Figure 43: Segmentation of Cubemap projection tiles

The RTSE of the OmniCam-360 delivers the 360° video as a Cubemap projection in YUV 4:2:0 planar format over 10 Gbit/s Ethernet. Each of the six sides of the cube corresponds to one of the fields of view: Front, Left, Right, Back, Top and Bottom. A three times two arrangement of the six sides (2048 x 2048) results in a resolution of 6,144 x 4,096 pixels. Before encoding, these cube sides are split into four 1,024 x 1,024 pixel tiles, resulting in a total of 24 tiles as shown in Figure 43. Each of the 24 tiles is encoded live in two different resolution levels, once in the original and once in reduced resolution of 512 x 512 pixels.

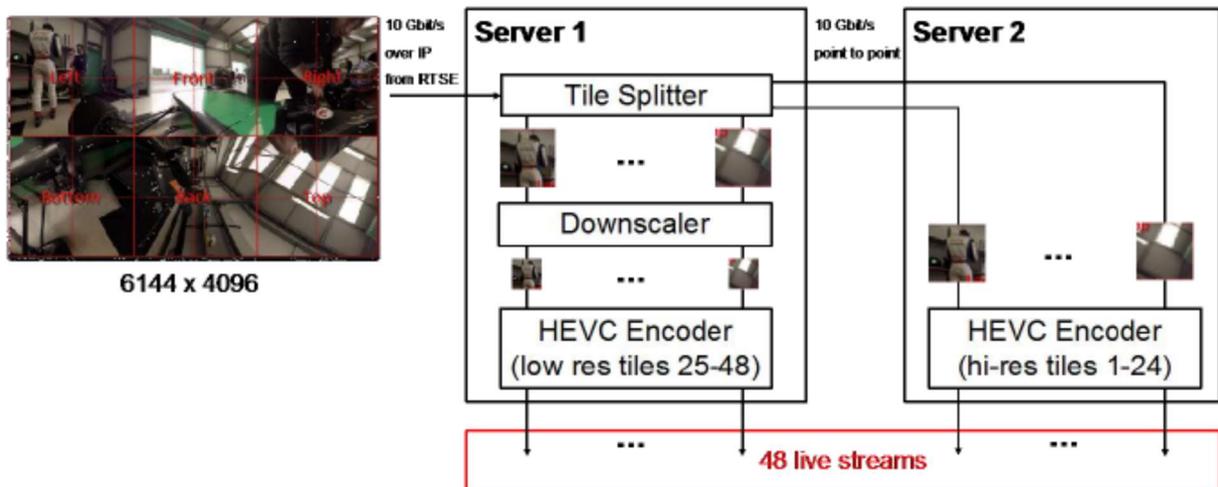


Figure 44: Block diagram of HEVC Tile-Encoder

Two servers, each with two Intel Xeon E5-2697v3 14-core CPUs, are used for HEVC encoding. Figure 44 shows the block diagram of the HEVC tile encoder. Server 1 segments the Cubemap images into 24 tiles,

scales down the resolution to 512 x 512 pixels and then encodes these low-resolution tiles. Server 2 encodes the 24 tiles of the highest quality with the original resolution of 1,024 x 1,024 pixels per tile. The delay of the encoder is 2 seconds.

Streaming

The Viewport-Dependent Media Profile of the MPEG-OMAF standard is used for streaming with 24 different viewpoints. The DASH Packager receives the 48 live encoded HEVC streams, then synchronizes and cuts them to the key frames. Then each segment is packaged individually and made available as a DASH segment on a simple HTTP server, creating 24 “Extractor-Track” segments (one for each view port) in addition to the 48 tile stream segments. The Extractors describe for each viewport which tile stream segments belong to which specific viewport. In addition, the initialization segments and a DASH manifest with all necessary OMAF metadata are generated once in the beginning. The upper part of Figure 45 illustrates how the OMAF/DASH packager works.

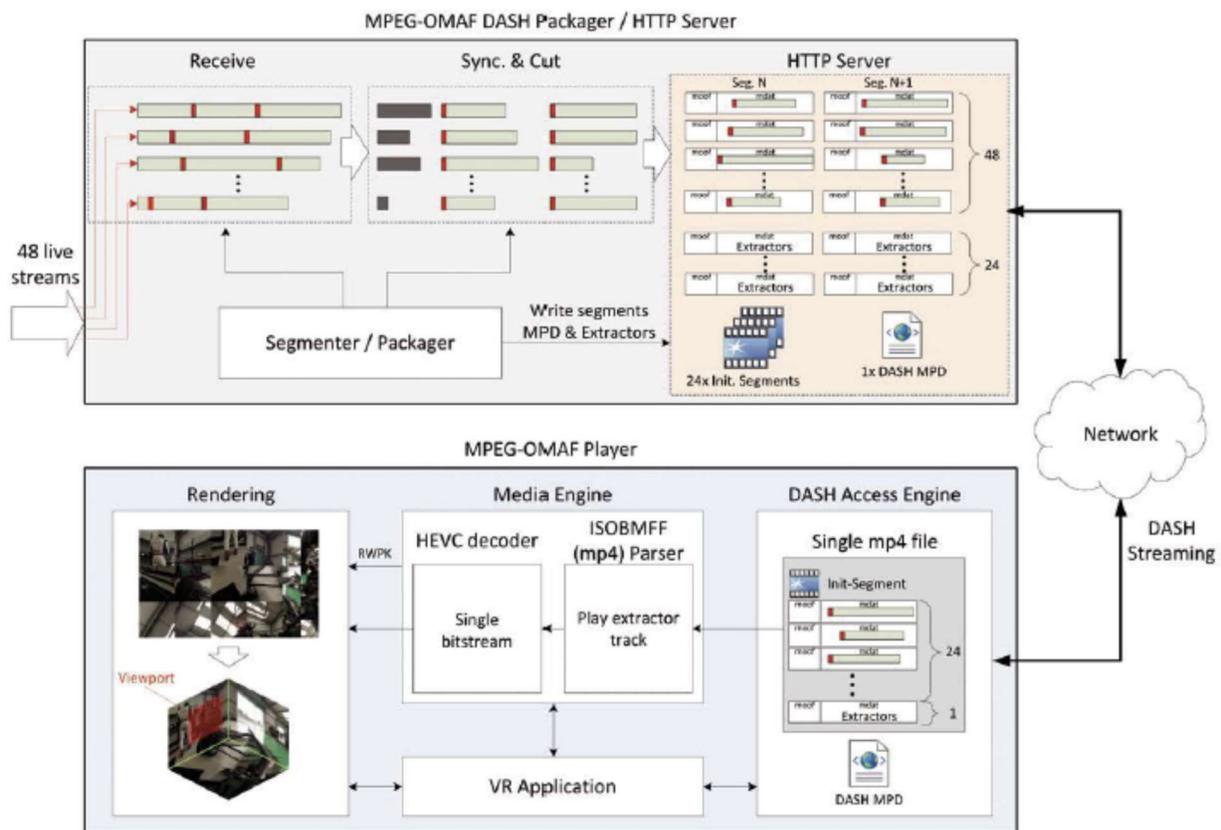


Figure 45: Block diagram of the OMAF packager and player

Since all separately encoded HEVC tiles are now made available on the server in two resolutions in the form of DASH segments, each receiver, such as VR glasses or television sets, can download the suitable HEVC tiles with the high resolution for the current viewport position. Outside the viewport, the HEVC tiles are consequently downloaded in a low resolution. OMAF metadata are analyzed so that each receiver can independently determine which tiles exactly are to be downloaded in high resolution. "Sphere region quality ranking" signals the region-dependent quality in 3D space and its assignment to respective DASH segments (tiles). The exact function of an OMAF player is illustrated in the lower part of Figure 45 and is briefly explained below.

After the respective 24 DASH segments have been downloaded for the entire 360° image, these are aggregated together with an Extractor Track segment and the corresponding Initialization segment to a file and passed to the mp4 parser. This parser extracts a single HEVC bit stream from the file, as specified in the extractor track, which is passed to the decoder. The decoded image is then projected onto the individual regions of the cube using OMAF region-wise packing and projection format metadata. The high-resolution tiles land on the cube surfaces in the viewport and the low-resolution tiles are scaled up and distributed on the remaining surfaces. This achieves a higher resolution in the viewport, which would not be possible with a traditional method.

To ensure that users can always see the high-resolution image even when the viewport moves quickly, the segments are prepared in such a way that a switching process can be carried out with every new frame (0.36 seconds). Depending on the current viewport, the bit rate of the transmitted content is 5 to 10 Mbps.

6.4.2 Suitable media profiles

6.4.2.1 Contribution media profiles

The media profile specified in section 5.5.1.2 is applicable to Live VR services.

6.4.2.2 Distribution media profiles

The video profiles specified in section 5.4.3.3 are applicable to the delivery of Live VR.

6.4.3 Suitable presentation profiles

The presentation profiles defined in clause 11 of [OMAF] are applicable to this use case. Any other media information that is added to the presentation should utilize the same coordinate system as the video and audio formats to ensure proper rendering and presentation.

6.4.4 Distribution Systems

The distribution systems specified in section 5.4.3 are also applicable to Live VR, with the exception that the download method specified in section 5.4.3.2 is generally not applicable.

6.5 Guidelines for Producers

Aside from the omnidirectional camera based aspects in section 6.4.1.1.1, the production of VR content closely copies traditional live production. Additional end-to-end (or glass-to-glass) latency is likely to be observed due to the media processing functions (image rectification, stitching, tiling) that needs to be performed to provide a high quality experience. Under optimal circumstances this latency should not exceed 10 seconds, but productions should anticipate up to 15 seconds of overall latency

6.6 Guidelines for service providers

Live VR experiences follow the same service provider guidelines as for On Demand content described in section 5.5. In addition, it is expected that the Live VR experience is “streamed” rather than “downloaded” as the later implies that the viewing time is not closely aligned with the actual event.

Nothing in these guidelines prevents a live event from being delivered immediately to viewers, but also separately process for future use. Such future use may use a download model for the HEVC-based viewport-independent OMAF video profile which is likely to be less resource intensive during playback.

6.7 Guidelines for Services Platform developers

The receiver model as that described in section 5.6 for on-demand VR content is applicable to Live VR.

6.8 Guidelines for Application developers

As the delivery of VR content to the end users follows the same workflow as on-demand VR content, the guidelines for the playback of Live VR content are the same as those provided in section 5.7.

Annex A Video Master Format Metadata

A.1 Video Metadata

Table 10: Master Format metadata

Name	Description	Data Type	Parameter Space	Defaults
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	mono
<i>Coverage</i>	Coverage parameters of the video	CoverageType See section 5.5.2		Full spherical
<i>Rotation</i>	Rotation parameters	RotationType See section 5.5.2		none
<i>InitialView</i>	Initial viewing point of the projected view	InitialViewType		If omitted, elevation=0, azimuth=0, tilt=0 is assumed
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>FrameRate</i>	Frame Rate in frames per second	Integer	25, 30, 50, 60, 75, 90, 100, 120	

Name	Description	Data Type	Parameter Space	Defaults
<i>Cropping</i>	Cropping information expressed as top, right, bottom, left	Integer, Integer, Integer, Integer	See section 5.5.2	none
<i>Padding</i>	Number of horizontal pixels expressing overlapping region on the left side of the frame. Used when the frame overlaps when wrapping around the equator.	Integer	See section 5.5.2	none
<i>PictureAspectRatio</i>	Picture aspect ratio	Integer : Integer	2:1	
<i>ChromaFormat</i>	Chroma format	Enum	YCbCr	
<i>ColourSampling</i>	Color sampling format	Enum	4:2:2	
<i>SampleAspectRatio</i>	Sample aspect ratio	Integer : Integer	1:1	
<i>BitDepth</i>	Bit depth	Integer	10	
<i>ColourPrimaries</i>	Color primaries	Enum	ITU-R BT.709: colour_primaries=1, matrix_coefficients=1 ITU-R BT.2100 : color primaries and non- constant luminance matrix coefficients, i.e. colour_primaries=9, matrix_coefficients=9	
<i>TransferFunction</i>	Transfer function	Integer	1: BT.709 14: SDR BT.2020 16: HDR PQ10 18: HDR HLG	

Name	Description	Data Type	Parameter Space	Defaults
Production Metadata				
<i>Duration</i>	Duration of the content	xs:duration		
<i>TimeStamp</i>	Epoch Time stamp when first frame was recorded	xs:dateTime		
<i>DirectorCut</i>	Directors cut data	DirectorCutType	Named sequence of time stamped center and tilt angles points	
Annotation				
<i>Contact</i>	Name, phone, email of person or organization to contact	string		
<i>StitchingSoftware</i>	Name and version of stitching software	string		
<i>Copyright</i>	Copyright information	string		
<i>License</i>	License information	string		

A.2 XML Schema for VR Video Master Format

XML documents containing VR Video format properties should be instantiated against the following schema. Descriptions for the use of the elements can be found above and in the main sections of these guidelines.

```
<?xml version="1.0" encoding="utf-8"?>
<schema targetNamespace="http://vr-if.org/VRVideoMetadata/1"
  elementFormDefault="qualified"
  xmlns:this="http://vr-if.org/VRVideoMetadata/1"
  xmlns="http://www.w3.org/2001/XMLSchema">

  <element name="VRVideoMetadata" type="this:VRVideoMetadata_type"/>
  <complexType name="VRVideoMetadata_type">
    <sequence>
      <!-- Spherical Video parameters-->
      <element name="Spherical" type="boolean" default="true" minOccurs="0"/>
      <element name="ProjectionType" type="this:AllowedProjections_type"
        minOccurs="0" default="equirectangular"/>
    </sequence>
  </complexType>
</schema>
```

```

<element name="StereoMode" type="this:StereoMode_type" minOccurs="0" default="mono"/>
<element name="Coverage" type="this:Coverage_type" minOccurs="0"/>
<element name="Rotation" type="this:Rotation_type" minOccurs="0"/>
<element name="InitialView" type="this:InitialView_type"/>
<!-- 2D Video parameters -->
<element name="FullWidthPixels" type="unsignedInt"/>
<element name="FullHeightPixels" type="unsignedInt"/>
<element name="FrameRate" type="this:FrameRate_type"/>
<element name="Cropping" type="this:Cropping_type" minOccurs="0"/>
<element name="Padding" type="unsignedInt" minOccurs="0" default="0"/>
<element name="PictureAspectRatio" type="this:PictureAspectRatio_type"
  minOccurs="0"/>
<element name="ChromaFormat" type="this:ChromaFormat_type" default="YCbCr"
  minOccurs="0"/>
<element name="ColourSampling" type="this:ColourSampling_type" default="4:2:2"
  minOccurs="0"/>
<element name="SampleAspectRatio" type="this:SampleAspectRatio_type" minOccurs="0"/>
<element name="BitDepth" type="this:BitDepth_type" minOccurs="0" default="10"/>
<element name="ColourPrimaries" type="this:ColourPrimaries_type"/>
<element name="TransferFunction" type="this:TransferFunction_type"/>
<!-- Production metadata -->
<element name="Duration" type="duration"/>
<element name="TimeStamp" type="dateTime"/>
<element name="DirectorCut" type="this:DirectorCut_type" minOccurs="0"/>
<!-- Annotation -->
<element name="Contact" type="this>Contact_type" minOccurs="0" maxOccurs="5"/>
<element name="StitchingSoftware" type="string" minOccurs="0"/>
<element name="Copyright" type="string" minOccurs="0"/>
<element name="License" type="string"/>
</sequence>
</complexType>

<simpleType name="AllowedProjections_type">
  <restriction base="string">
    <enumeration value="equirectangular"/>
  </restriction>
</simpleType>
<simpleType name="BitDepth_type">
  <restriction base="unsignedInt">
    <enumeration value="10"/>
  </restriction>
</simpleType>
<simpleType name="ChromaFormat_type">
  <restriction base="string">
    <enumeration value="YCbCr"/>
  </restriction>
</simpleType>
<simpleType name="ColourPrimaries_type">
  <restriction base="string">
    <enumeration value="ITU-R BT.709"/>
    <enumeration value="ITU-R BT.2100"/>
  </restriction>
</simpleType>
<simpleType name="ColourSampling_type">
  <restriction base="string">

```

```

        <enumeration value="4:2:2"/>
    </restriction>
</simpleType>
<complexType name="Contact_type">
    <sequence>
        <element name="Name" type="string" minOccurs="0"/>
        <element name="Email" type="string" minOccurs="0"/>
        <element name="Phone" type="string" minOccurs="0"/>
    </sequence>
</complexType>
<complexType name="Coverage_type">
    <sequence>
        <element name="AzimuthMin" type="this:OMAFMinusPlus180_type" minOccurs="0" />
        <element name="AzimuthMax" type="this:OMAFMinusPlus180_type" minOccurs="0" />
        <element name="ElevationMin" type="this:OMAFMinusPlus90_type" minOccurs="0" />
        <element name="ElevationMax" type="this:OMAFMinusPlus90_type" minOccurs="0" />
    </sequence>
</complexType>
<complexType name="Cropping_type">
    <sequence>
        <element name="Top" type="unsignedInt" minOccurs="0"/>
        <element name="Right" type="unsignedInt" minOccurs="0"/>
        <element name="Bottom" type="unsignedInt" minOccurs="0"/>
        <element name="Left" type="unsignedInt" minOccurs="0"/>
    </sequence>
</complexType>
<complexType name="DirectorCut_type">
    <sequence>
        <element name="description" type="string" minOccurs="0" />
        <element name="gp" type="this:DirectorCut_entry" minOccurs="0"
            maxOccurs="unbounded" />
    </sequence>
</complexType>
<complexType name="DirectorCut_entry">
    <attribute name="ts" type="duration" use="required"/>
    <attribute name="azi" type="this:OMAFMinusPlus180_type"/>
    <attribute name="ele" type="this:OMAFMinusPlus90_type"/>
    <attribute name="tilt" type="this:OMAFMinusPlus180_type"/>
</complexType>
<simpleType name="FrameRate_type">
    <restriction base="unsignedInt">
        <enumeration value="25"/>
        <enumeration value="30"/>
        <enumeration value="50"/>
        <enumeration value="60"/>
        <enumeration value="75"/>
        <enumeration value="90"/>
        <enumeration value="100"/>
        <enumeration value="120"/>
    </restriction>
</simpleType>
<complexType name="InitialView_type">
    <sequence>
        <element name="InitialAzimuth" type="this:OMAFMinusPlus180_type" default="0"
            minOccurs="0" />
    </sequence>

```

```

    <element name="InitialElevation" type="this:OMAFMinusPlus90_type" default="0"
      minOccurs="0" />
    <element name="InitialTilt" type="this:OMAFMinusPlus180_type" default="0"
      minOccurs="0" />
  </sequence>
</complexType>
<complexType name="PictureAspectRatio_type">
  <sequence>
    <element name="Width" default="2" minOccurs="0">
      <simpleType>
        <restriction base="unsignedInt">
          <enumeration value="2"/>
        </restriction>
      </simpleType>
    </element>
    <element name="Height" default="1" minOccurs="0">
      <simpleType>
        <restriction base="unsignedInt">
          <enumeration value="1"/>
        </restriction>
      </simpleType>
    </element>
  </sequence>
</complexType>
<complexType name="Rotation_type">
  <sequence>
    <element name="RotationYaw" type="this:OMAFMinusPlus180_type" default="0"
      minOccurs="0"/>
    <element name="RotationPitch" type="this:OMAFMinusPlus90_type" default="0"
      minOccurs="0"/>
    <element name="RotationRoll" type="this:OMAFMinusPlus180_type" default="0"
      minOccurs="0"/>
  </sequence>
</complexType>
<complexType name="SampleAspectRatio_type">
  <sequence>
    <element name="Width" default="1" minOccurs="0">
      <simpleType>
        <restriction base="unsignedInt">
          <enumeration value="1"/>
        </restriction>
      </simpleType>
    </element>
    <element name="Height" default="1" minOccurs="0">
      <simpleType>
        <restriction base="unsignedInt">
          <enumeration value="1"/>
        </restriction>
      </simpleType>
    </element>
  </sequence>
</complexType>
<simpleType name="StereoMode_type">
  <restriction base="string">
    <enumeration value="mono"/>
  </restriction>
</simpleType>

```

```

        <enumeration value="stereo-left"/>
        <enumeration value="stereo-right"/>
    </restriction>
</simpleType>
<simpleType name="TransferFunction_type">
    <restriction base="unsignedInt">
        <enumeration value="1">
            <annotation>
                <documentation>BT.709</documentation>
            </annotation>
        </enumeration>
        <enumeration value="14">
            <annotation>
                <documentation>SDR BT.2020</documentation>
            </annotation>
        </enumeration>
        <enumeration value="16">
            <annotation>
                <documentation>HDR PQ10</documentation>
            </annotation>
        </enumeration>
        <enumeration value="18">
            <annotation>
                <documentation>HDR HLG</documentation>
            </annotation>
        </enumeration>
    </restriction>
</simpleType>

<simpleType name="OMAFMinusPlus180_type">
    <restriction base="float">
        <minInclusive value="-180"/>
        <maxExclusive value="180"/>
    </restriction>
</simpleType>
<simpleType name="OMAFMinusPlus90_type">
    <restriction base="float">
        <minInclusive value="-90"/>
        <maxInclusive value="90"/>
    </restriction>
</simpleType>
</schema>

```

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 HEVC-based viewport-dependent OMAF video 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 46 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.

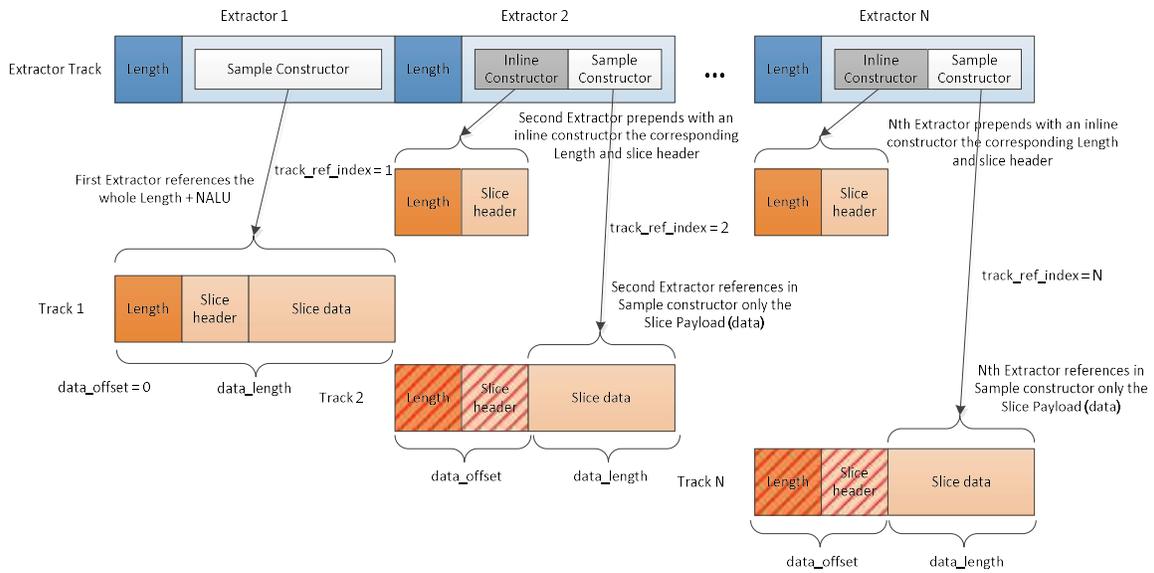


Figure 46: Single ISO BMFF File with one extractor track, N extractors and N MCTS tracks after subsegment concatenation

Annex C Example DASH Manifest for FOV enhanced tiling

```
<?xml version="1.0" encoding="UTF-8" ?>
<MPD xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xmlns="urn:mpeg:DASH:schema:MPD:2011"
xmlns:omaf="urn:mpeg:mpeg1:omaf:2017" xsi:schemaLocation="urn:mpeg:dash:schema:mpd:2011 DASH-MPD.xsd
urn:mpeg:mpeg1:omaf:2017 OMAF.xsd" type="dynamic" minimumUpdatePeriod="PT5S" timeShiftBufferDepth="PT30M"
availabilityStartTime="2018-02-23T17:10:40" minBufferTime="PT4S" profiles="urn:mpeg:dash:profile:isoff-live:2011">
  <Period start="PT0S">
    <AdaptationSet id="0" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
      <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,0,3840,1920,3840,1920" />
      <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
        <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
          <omaf:coverageInfo centre_azimuth="0" centre_elevation="0" azimuth_range="23592960" elevation_range="11796480" />
        </omaf:cc>
      </SupplementalProperty>
      <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
      <Representation id="s0" width="3840" height="1920" bandwidth="9664816" />
    </AdaptationSet>
    <AdaptationSet id="1" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
      <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,0,1280,320,3840,1920,1" />
      <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
        <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
          <omaf:coverageInfo centre_azimuth="0" centre_elevation="4915200" azimuth_range="23592960" elevation_range="1966080"
/>
        </omaf:cc>
      </SupplementalProperty>
      <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
      <Representation id="t00" width="320" height="320" bandwidth="918200" />
    </AdaptationSet>
    <AdaptationSet id="2" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
      <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,320,1920,320,3840,1920,1" />
      <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
        <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
          <omaf:coverageInfo centre_azimuth="5898240" centre_elevation="2949120" azimuth_range="11796480"
elevation_range="1966080" />
        </omaf:cc>
      </SupplementalProperty>
      <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
      <Representation id="t01" width="320" height="320" bandwidth="656528" />
    </AdaptationSet>
    <AdaptationSet id="3" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
      <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="01920,320,1920,320,3840,1920,1" />
      <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
        <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
          <omaf:coverageInfo centre_azimuth="-5898240" centre_elevation="2949120" azimuth_range="11796480"
elevation_range="1966080" />
        </omaf:cc>
      </SupplementalProperty>
      <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
```

```

    <Representation id="t02" width="320" height="320" bandwidth="1306472" />
  </AdaptationSet>
  <AdaptationSet id="4" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,640,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
      <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
        <omaf:coverageInfo centre_azimuth="8847360" centre_elevation="983040" azimuth_range="5898240"
elevation_range="1966080" />
      </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t03" width="320" height="320" bandwidth="36696" />
  </AdaptationSet>
  <AdaptationSet id="5" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,960,640,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
      <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
        <omaf:coverageInfo centre_azimuth="2449120" centre_elevation="983040" azimuth_range="5898240"
elevation_range="1966080" />
      </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t04" width="320" height="320" bandwidth="108640" />
  </AdaptationSet>
  <AdaptationSet id="6" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,1920,640,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
      <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
        <omaf:coverageInfo centre_azimuth="-2449120" centre_elevation="983040" azimuth_range="5898240"
elevation_range="1966080" />
      </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t05" width="320" height="320" bandwidth="34384" />
  </AdaptationSet>
  <AdaptationSet id="7" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,2880,640,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
      <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
        <omaf:coverageInfo centre_azimuth="-8847360" centre_elevation="983040" azimuth_range="5898240"
elevation_range="1966080" />
      </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t06" width="320" height="320" bandwidth="30832" />
  </AdaptationSet>
  <AdaptationSet id="8" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,960,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
      <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">

```

```

        <omaf:coverageInfo centre_azimuth="8847360" centre_elevation="-983040" azimuth_range="5898240"
elevation_range="1966080" />
    </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t07" width="320" height="320" bandwidth="39416" />
</AdaptationSet>
    <AdaptationSet id="9" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,960,960,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="2449120" centre_elevation="-983040" azimuth_range="5898240"
elevation_range="1966080" />
    </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t08" width="320" height="320" bandwidth="35288" />
</AdaptationSet>
    <AdaptationSet id="10" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,1920,960,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="-2449120" centre_elevation="-983040" azimuth_range="5898240"
elevation_range="1966080" />
    </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t09" width="320" height="320" bandwidth="66192" />
</AdaptationSet>
    <AdaptationSet id="11" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,2880,960,960,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="-8847360" centre_elevation="-983040" azimuth_range="5898240"
elevation_range="1966080" />
    </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t10" width="320" height="320" bandwidth="160384" />
</AdaptationSet>
    <AdaptationSet id="12" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
    <EssentialProperty schemeldUri="urn:mpeg:dash:srd:2014" value="0,0,1280,1920,320,3840,1920,1" />
    <SupplementalProperty schemeldUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
    <omaf:coverageInfo centre_azimuth="5898240" centre_elevation="-2949120" azimuth_range="11796480"
elevation_range="1966080" />
    </omaf:cc>
    </SupplementalProperty>
    <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
    <Representation id="t11" width="320" height="320" bandwidth="193464" />

```

```

</AdaptationSet>
<AdaptationSet id="13" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
  <EssentialProperty schemeIdUri="urn:mpeg:dash:srd:2014" value="0,1920,1280,1920,320,3840,1920,1" />
  <SupplementalProperty schemeIdUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
      <omaf:coverageInfo centre_azimuth="-5898240" centre_elevation="-2949120" azimuth_range="11796480"
elevation_range="1966080" />
    </omaf:cc>
  </SupplementalProperty>
  <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
  <Representation id="t12" width="320" height="320" bandwidth="39392" />
</AdaptationSet>
<AdaptationSet id="14" contentType="video" mimeType="video/mp4" codecs="hev1.010101" frameRate="25"
segmentAlignment="true" startWithSAP="1">
  <EssentialProperty schemeIdUri="urn:mpeg:dash:srd:2014" value="0,0,1600,3840,320,3840,1920,1" />
  <SupplementalProperty schemeIdUri="urn:mpeg:mpeg1:omaf:2017:cc">
    <omaf:cc shape_type="1" view_idc_presence_flag="0" default_view_idc="0">
      <omaf:coverageInfo centre_azimuth="0" centre_elevation="-4915200" azimuth_range="23592960"
elevation_range="1966080" />
    </omaf:cc>
  </SupplementalProperty>
  <SegmentTemplate timescale="90000" duration="90000" initialization="$RepresentationID$/init.mp4"
media="$RepresentationID$/Number$.m4s" startNumber="1" />
  <Representation id="t13" width="320" height="320" bandwidth="66544" />
</AdaptationSet>
</Period>
</MPD>

```

Annex D Live VR Deployment Information

D.1 Tiledmedia

D.1.1 IBC Demonstration

Together with a team of world-leading companies, the Intel Visual Cloud Conference at the 2019 IBC was delivered live to a global audience in 8K 360VR. This visual experience (i.e. *IBC 360 Live*) was made available to a variety of mobile and head mounted devices, showcasing what the ecosystem can deliver over today's networks.

Besides Intel and Tiledmedia, the demonstration partnership included Akamai, Google, IBC, Iconic Engine, KPN, Oculus, and Voysys. This joint effort enabled the first-ever live 8K VR distribution with worldwide availability, stringing together hardware, software and services that are all commercially available today.

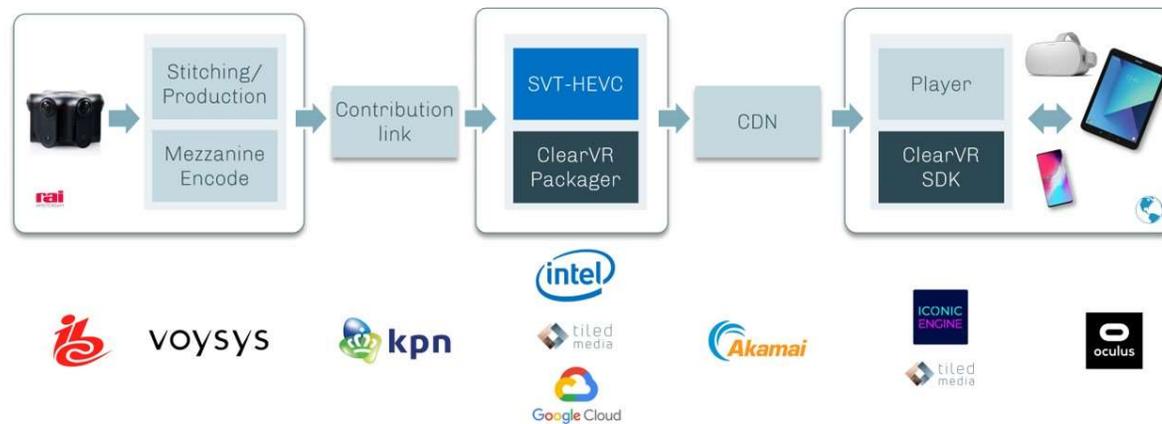


Figure D.1.1.1 IBC demonstration overview

D.1.2 Bringing High-Quality VR to a Mass Market

Virtual and Augmented Reality services have a great future ahead of them. A well-designed VR experience can be transformative: the user will feel completely transported to a new environment, away from home to a new and exciting location. While VR adoption has been growing steadily, mainstream adoption has been impeded by ecosystem and technology development. This is not a surprise, given the complexities of delivering these very rich live experiences. For VR360 service deployment the end user Quality of Experience (QoE) is key

to accelerating adoption. To enable a great QoE, it is necessary to be able to produce amazing content in high quality, and to distribute it to large audiences while maintaining that high quality. This is exactly what the IBC cooperation shows.

D.1.3 The Challenge: Good Quality VR Distribution

To understand why this project was unique, there needs to be an explanation of what makes distribution at 8K quality hard. A “VR camera” is usually a set of linked conventional video cameras that all record a piece of the environment. Stitching software then combines these video images into a single spherical video in the form of an “equirectangular projection” or “ERP” (a flat map of the earth is also an ERP). A VR user only sees a small part of that sphere at any one point in time in their VR headset: about 1/8th of the complete picture. Since the video is played right in front of the user’s eye, magnified by special lenses, the resolution needs to be very good. The industry standard in VR is to use 4K video for the entire sphere. 4K in VR is 4 096 x 2 048 pixels, but the user will only see, approximately, a mere 1K x 1K per eye, right in front of them – not a good experience! If 8K video (8 192 x 4 096 pixels) is utilized, the image becomes considerably better. But to get that quality to the user, one needs to send a huge number of pixels: the equivalent of about 16 high definition channels, requiring some 60 – 100 Mbit/s second. Even if it is possible to get that bandwidth across, there are only very few devices that could decode it.

D.1.4 Solution: Smart VR Delivery Through Tiles

The *IBC 360 Live* distribution system gives consumers an 8K experience while transmitting and decoding much less information; the platform sends *only* what a user actually sees. To make this possible, the image is cut up into around 100 high-quality tiles and only the actual tiles in view are sent, decoded and displayed. The real time nature of the experience, including the need to instantly respond to a user’s head motion, makes this approach particularly challenging. When the user turns their head, the system retrieves new high-resolution tiles, decodes and displays them – all within a tenth of a second. This happens so fast that users will hardly notice the low-resolution background layer that is always present to prevent black areas from appearing in their field of view.

Tiled streaming is at the heart of the end-to-end *IBC 360 Live* experience which was powered by the following components:

- **Cameras:** High-resolution VR360 cameras were used to capture the event. Kandao™ Obsidian R cameras were employed as they are compact and affordable. Each of the VR cameras produces 3 individual feeds, which need to be stitched in real-time into a single 360 sphere. Intel’s Visual Cloud Conference at **IBC** was recorded with two such cameras, letting the user freely switch between these in the application.
- **Production Processing:** Voysys VR Producer software performed the stitching of the three individual camera signals into a single “Equirectangular Projection” with a resolution of 8 092 x 4 046 pixels. (An ERP is like a flat image of the globe as used in maps). Voysys VR Producer was also used to mix a crisp image of the conference speaker slides as a virtual screen into the scene, and there was also the option to mix in the “director’s cut” as produced by IBC TV. The ERP was then converted into a “cubemap” (a globe projected onto the six faces of a cube), which is a format more suitable for tiled distribution – see below.

- **Processing:** Two servers with **Intel® Xeon® Scalable** processors and graphics cards were used to run the Voysys software and perform an HEVC mezzanine encode for each of the cubemaps, at 150 Mbit/s. The mezzanine encoder lightly compresses the camera feeds for contribution to the cloud data center where the feeds are processed for tiled distribution.
- **Data Links:** KPN supplied a dedicated, 400 Mbit/s fiber link to the Internet, on which the two mezzanine-compressed ERPs were sent to the Eemshaven **Google Cloud Platform (GCP)** instance in Groningen, The Netherlands. This was the GCP closest to the RAI in Amsterdam.
- **Delivery Processing: Tiledmedia's** ClearVR Cloud Live software running in Google Cloud first decoded the cubemaps and divided the decoded video in 96 tiles, 16 per cube face. For each of the tiles, two versions were encoded using Intel's open source SVT-HEVC encoder, which supports the encoding constraints that are required for tiled distribution. There are two encoded versions of each tile because one needs to enable both fast switching (a short GOP) and high efficiency (a longer GOP). Six larger tiles are also encoded, one for each cube face – these provide a lower-resolution background that is always present at the decoder. The HEVC-encoded tiles were then packaged into mp4 files for distribution. The processes to create the cubemap, and then tile, encode and package the full 8K content run on multiple hundreds of Intel Xeon cores, in parallel processes managed by the ClearVR Orchestrator. ClearVR Orchestrator ensures all tiles are ready in time for ingest into the CDN and provides for redundancy and fail-over in the encoding and packaging. The total bitrate of the packaged files amounts to about 120 Mbit/s. Note that these numbers all apply to one single 8K VR feed; and there were two parallel camera feeds, so the numbers doubled.
- **CDN:** Next, the packaged MP4 files were ingested into **Akamai's** CDN using 6 parallel Media Services Live (MSL) ingest points per camera, where each of these ingests can handle 45 Mbit/s. It is required that these parallel ingests to accommodate the exceptionally high 120 Mbit/s bandwidth mentioned in the previous bullet, and allow for some variation in the individual ingests. Note that non-VR streams do not come close to Akamai's current 45 Mbit/s ingest limit, but that VR streams easily exceed them. The multiple-ingest configuration was developed by Akamai and Tiledmedia in 2018 to accommodate live 6K and 8K VR events. As the *IBC 360 Live* production used two parallel 8K cameras, twelve (2 x 6) MSL ingests were used in total. The ingested files were then recombined and uploaded onto Akamai's origin for its EU region and, from there, distributed across all regions globally.
- **Format Optimisation:** To distribute VR content, live or on-demand, the Akamai network uses a VR-specific configuration that optimizes the likelihood that requested tiles are available at the edge cache. This is important in giving the optimal response when users move their head (or swipe on a flat device – see below). This configuration, again developed by Akamai and Tiledmedia, has been available since 2017. It relies on the http/2 or QUIC (http/3) protocol using with multipart byterange requests and smart pre-fetching of tiles to the edge cache. With tiled streaming, a short request/response delay is much more important than for regular on-demand services. In our configuration, the “motion-to-high-resolution” latency is typically 3 frames or less at 30 frames per second: after head motion, over 85% of high-resolution tiles are in the user's viewport within these 3 frames. Given then way the human visual system works, this is virtually instant to the viewer. The final distributed resolution was 8 192 x 4 096, slightly higher than the ingested resolution because tiling requires certain multiples of 16 to be used.
- **Delivery:** The final part of the delivery is handled by the application. **Iconic Engine**, an end-to-end provider of XR solutions, developed a special *IBC 360 Live* application for the Oculus platform, available on the **Oculus Go** and Gear VR. Another version of the

same app was available in Apple's App Store and the Google Play Store for iOS and Android tablets and phones; this allowed playback of the VR streams on "flat" devices and in so-called cardboard viewers. Iconic Engine's application platform integrates Tiledmedia's ClearVR library, which enabled Iconic Engine to develop the special IBC app in a matter of weeks. In the *IBC 360 Live* app, ClearVR which determines what the users sees and retrieves the required tiles for that viewport using http retrieval as described above. The SDK next reassembles the bitstream snippets for these individual tiles into one single legal bitstream, which is decoded by the hardware HEVC decoder in the user device. The last step is putting all the tiles in the right place on the rendered sphere. All of this happens at the last possible moment, in real time, giving the user the highest possible quality on existing headsets and flat screens. The final end-user bitrate hovered between 12 and 15 Mbit/s; and it was possible to tune this bitrate in real-time by operating the ClearVR Cloud Live platform. The end-to-end latency was about 30 seconds. Using the latest chunk-based streaming techniques, the partners expect to bring this down to the single-digit domain before the end of the year.

D.1.5 From Idea to Production in Less than Two Months

The production of the IBC 360 Live event was conceived and realized in less than 2 months, which was only possible because no new elements needed to be developed; all partners contributed elements that they have commercially available today. The production was still unique: never had a 5-day event been broadcast non-stop in 8K VR, and never was such a distribution truly global. The encoded and distributed video was compliant with the HEVC standard which is implemented in hardware in all modern mobile devices. Tiledmedia is currently working in MPEG to contribute its innovations to the upcoming *Omnidirectional Media Format (OMAF)* version 2 standard, planned for approval by mid-2020.

Because the tiled streaming relies on regular http requests and does not require any per-user (edge) processing, the distribution system used for this event was as scalable as the CDN itself, suitable for mass distribution of the next major sports or music event. The 8K quality makes users feel more present than ever before.

D.1.6 IBC Solution Summary

Number of cameras: *two, user-switched*
Event Duration: *5 days; 8+ hours per day*
Capture resolution: *8 092 x 4 046*
Contribution bitrate: *150 Mbit/s per camera (HEVC)*
Distribution resolution *8 192 x 4 096*
Distribution Encoder: *SVT-HEVC*
Distribution Format *MP4-based ClearVR packaging*
Distribution bitrate: *12 – 15 Mbit/s*
Streaming Protocol..... *standard http/2 with multipart byterange requests*
User device decoder:..... *HEVC Main level 5.1*
Glass-to-glass latency:..... *~30 seconds*
Supported devices: *Oculus headsets,
iOS and Android tablets and phones*
Cloud processing: *well over 1 000 Intel cores in Google Cloud Platform for the two streams (dynamically managed)*

D.2 Intel Sports

D.2.1 Immersive Media Experiences in Sports Guiding Example Use Case

D.2.1.1 Background and Set-up/Equipment for Use Case

Intel Sports has partnerships with some of the world's top sports leagues, teams, and broadcasters. The Intel immersive media platform provides partners an end-to-end solution for processing immersive media content and delivering immersive media experiences to fans worldwide. Intel Sports offers services to partners that enable them to personalize immersive media experiences based on the way their fans prefer to experience the game. During the recording of a sporting event, the capture systems, Intel® True View and Intel® True VR work independently or side-by-side at the stadium.

Intel® True VR is a transportable stereoscopic camera solution that can be set up in stadiums on game day. The stereoscopic camera pods are placed close to the action, and often coincide with the placements of traditional broadcast cameras. The cameras are placed in specific locations to enhance immersive media experiences. Intel® True VR outputs panoramic video from the stadium and sends it through the immersive media processing pipeline.

Intel® True View is comprised of a camera array that is built into the perimeter of the stadium. The high-resolution cameras are angled to capture the entire field of play. In parallel or independently, Intel® True View outputs volumetric video through the same immersive media processing pipeline as Intel® True VR.

Once the content from either capture system is processed, the distribution pipeline enables delivering a catered stream for each supported device. The Intel Sports immersive media platform supports distributing both live and non-live content. The immersive media platform also supports 2D screen viewing experiences like mobile phones, traditional broadcast television, and web players; as well as, VR/AR/MR HMDs and other immersive media streaming platforms.

D.2.1.2 Use Case Description Short

On the day of a live event the portable stereoscopic Intel® True VR cameras are brought in and set up inside the sports venue. In partnership with the various teams, leagues, and broadcasters, one to eight camera pods, each with as many as 12 cameras, are set up closest to the action. Fans watch the live stream through Intel Sports partners' immersive media viewing applications, on one of the many streaming platforms supported during the event.

The streams are consumed in applications built by our partners or built by the cross-platform Intel Sports immersive media platform. Applications built with the immersive media platform offer fans a more personalized experience through produced third-party audio and video feeds, camera switching, and live game stats. Intel's production team captures game footage from many angles, and can provide a produced feed of the best vantage points of the game. Based on the type of event, a sideline reporter or a host will be added to the production team for full end-to-end coverage. In addition to the produced feed, the user with the camera switching feature can select other isolated views of the various camera pods.

Stadiums installed with Intel® True View can independently or in parallel with Intel® True VR create content. Working with the creative producers from the teams and broadcasters, the production team can create volumetric content focused on the most action-packed or analysis worthy parts of the game. The panoramic and volumetric video created by the capture systems is then pushed to the processing and distribution pipeline. Current distribution of the True View content feeds that various outputs as a highlight. The outputs are: in-house jumbotron, national broadcasters, and social media channels.

During a live event, the immersive media platform distributes a catered experience through the unified immersive media processing and distribution pipeline for all forms of immersive content.

D.2.1.3 Use Case Description Detailed

In this use case, there are simultaneous immersive media capturing methods occurring in the stadium. During a live event there is a traditional broadcast feed, portable stereoscopic camera pods at preset locations, and a stationary volumetric capture system built into the stadium itself. The portable stereoscopic camera pods are placed at the vantage points with the best static viewing angle for the sporting event. As seen in *Figure D.2.1.3.1*, cameras are placed around the basketball court and mounted behind the backboard. The white chevrons in *Figure 1* represents how some of the camera pods are actually laid out during a basketball game. Some camera pods are close to the court or field and others are placed throughout the stadium.



Figure D.2.1.3.1: Intel® True VR for Stereoscopic Capture at a Stadium

For the stadium as shown in *Figure D.2.1.3.2*, a camera array of many high-definition 5K cameras are installed around the venue. The cameras are placed in a way that enables them to capture the entire field.

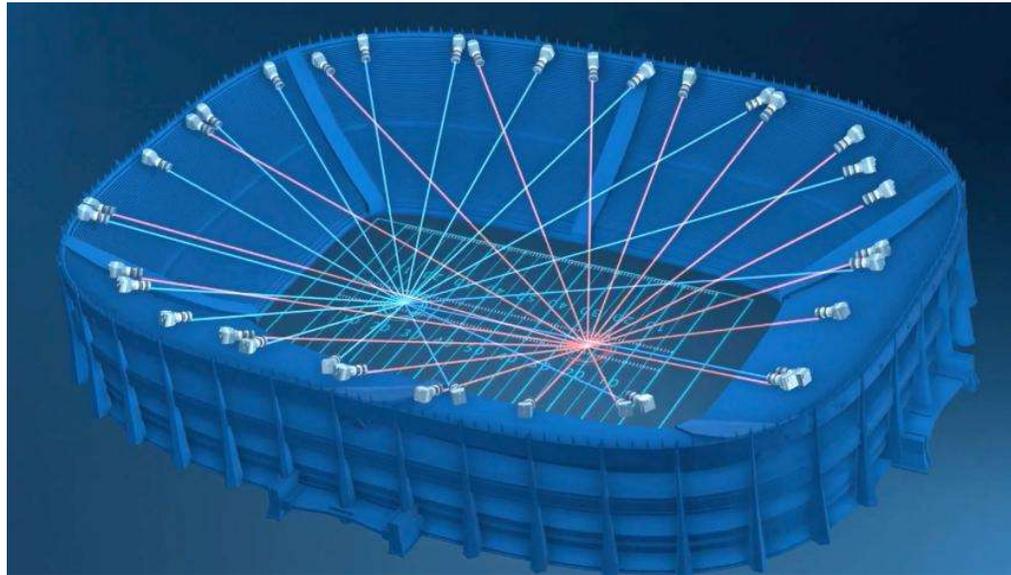


Figure D.2.1.3.2: Intel® True View for Volumetric Capture at a Stadium

This capture system generates massive amounts of volumetric data in the form of voxels. These voxels contain data describing height, width, *depth*, and other attributes used to reconstruct the volumetric capture. The camera array communicates to dedicated on-site servers through a fiber-optic network.

During a live event, Intel technologies work together to deliver end users a more personalized immersive media experience, by providing them a mix of static and dynamic viewing angles that differ from those in traditional broadcast feeds. Depicted in *Figure D.2.1.3.3* are the user's in-app camera selection offerings, which include the produced cast that contain the overlay of the televised broadcast stream.



Figure D.2.1.3.3: In-App Camera Selection for Viewing Angles and Produced Stream

Fans have the ability to access more camera angles than just the broadcast feed. Those viewing angles are vetted through Intel’s content creators. The content creators provide fans new perspectives that are not possible with traditional broadcast television. The live streams are delivered to immersive viewing platforms, where a user can also consume past content, pre-curated content like interviews, and access to volumetric replays as they become available. The use of storytelling and personalization through the user interface flow of the application creates a unique live immersive media experience that sets it apart from traditional sports viewing experiences like live broadcast television.

To reach a wide audience, Intel Sports immersive media platform is designed to support a large amount of immersive media viewing platforms and devices. This use case is mostly geared towards HMD and non-headset 2D screen viewing experiences.

D.2.2 Immersive Media Platform Reference Architecture and Workflow

The following diagram, *Figure D.2.2.1*, showcases the Intel Sports immersive media platform that is powered by Intel® True VR and Intel® True View. This workflow highlights the video capture steps that occur at the stadiums, the immersive media processing and distribution pipeline happening in the cloud, and the enablers for distributing experiences on the client-side.

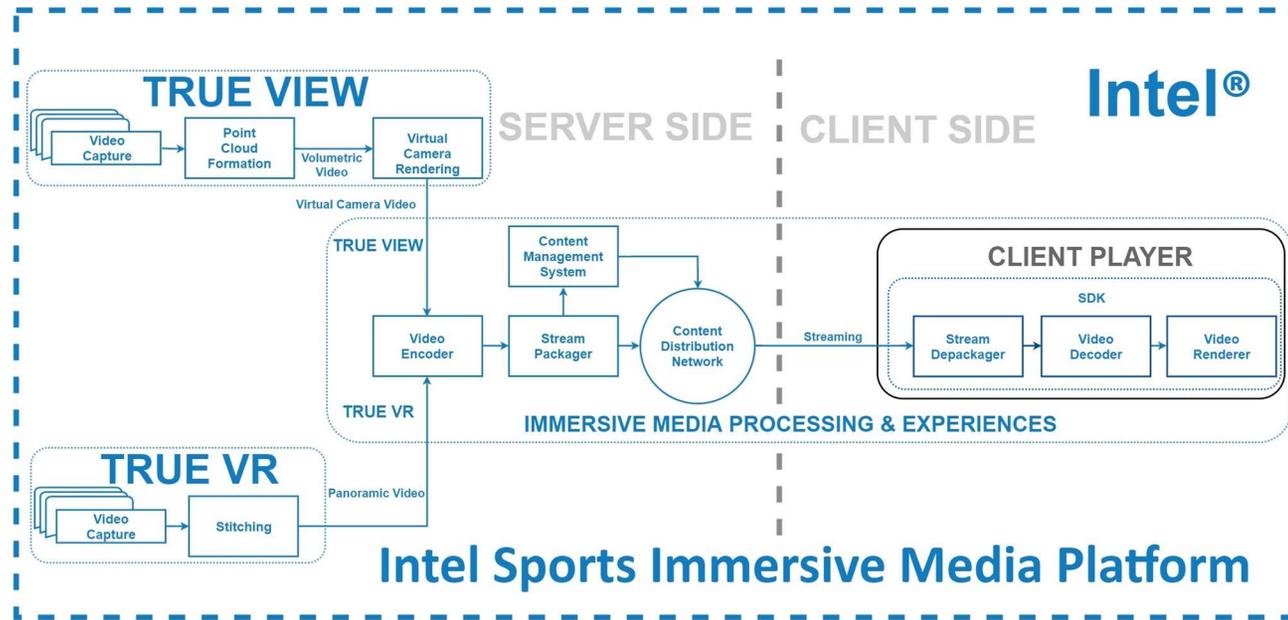


Figure D.2.2.1: Immersive Media Platform Architecture and Workflow

D.2.3 Technical Enablers

D.2.3.1 Production Systems for Immersive Content

D.2.3.1.1 Intel® True View

During a live event, video is captured from each of the high-definition 5K cameras that are installed in the venue. The video captured at the stadium generates massive amounts of volumetric data in the form of voxels. The voxels capture height, width, *depth*, and relative attributes that are needed for the point cloud formation step of the Intel® True View capture system. The camera array is connected by fiber to dedicated on-site Intel servers. The data is stored, synchronized, analyzed, and processed in the cloud. The volumetric video created by the point cloud sends terabytes of volumetric data that are rendered into high-fidelity 3D video in the form of virtual camera videos. Current challenges are with matching the color output to broadcast standards and variability of the weather can also be a factor in the segmentation and calibration. The cameras shown in *Figure D.2.3.1.1.1* are Intel® True View cameras.

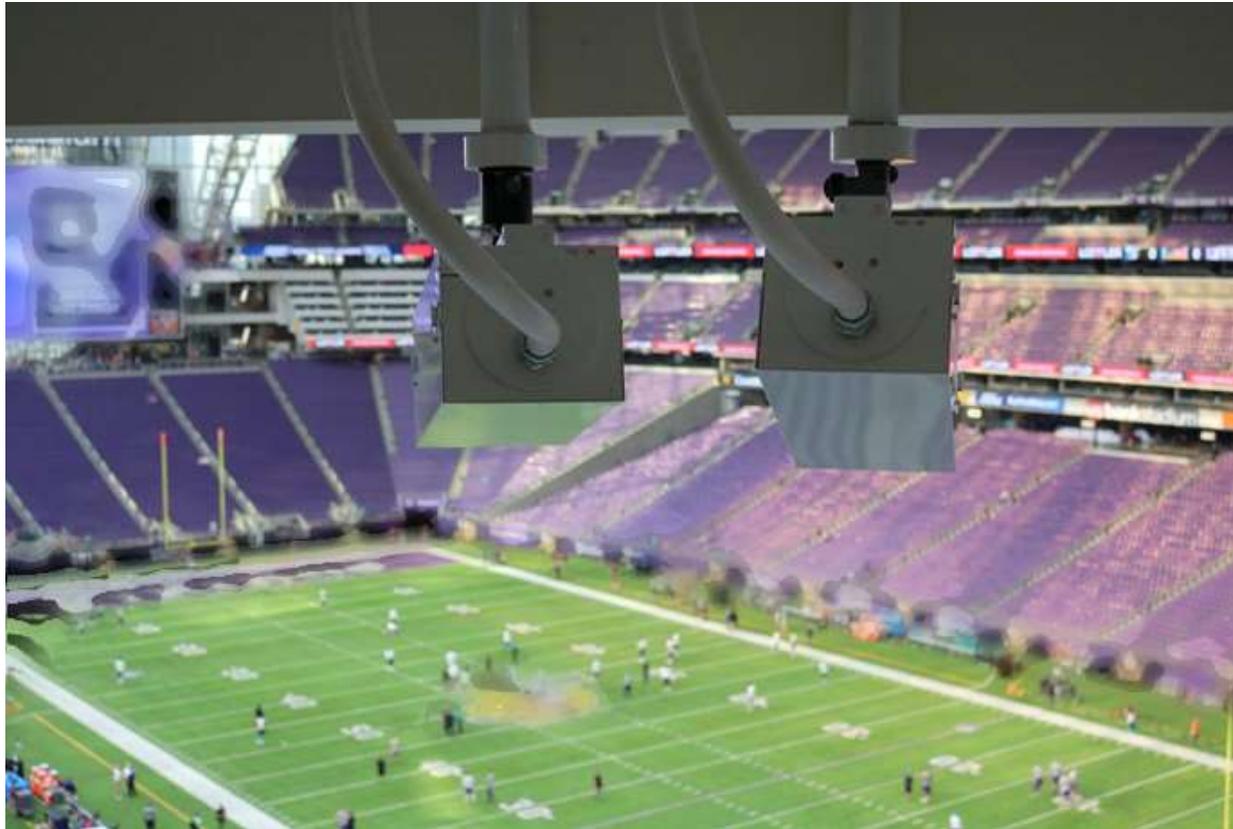


Figure D.2.3.1.1.1: In Stadium Intel® True View Cameras

D.2.3.1.2 Intel® True VR

In parallel with the Intel® True View capture system, the Intel® True VR capture system also captures video from multiple stereoscopic camera pods. The camera pods are Intel proprietary paired-lens stereoscopic camera pods that are setup in the stadium for the live event. An event can have as many as one to eight camera pods, each pod with as many as 12 cameras that are stitched together in the cloud. The stitched camera pod views are then turned into panoramic video that is sent to the common media processing and distribution pipeline. The camera pod shown in *Figure D.2.3.1.2.1* is part of Intel® True VR.



Figure D.2.3.1.2.1: Intel® True VR Camera Pod

D.2.3.1.3 Intel Sports Immersive Media Processing & Experiences Pipeline

Intel® True VR and Intel® True View share a common video processing and distribution pipeline that is hosted in the cloud. The video processing component of the pipeline takes the raw and uncompressed virtual camera videos from the Intel® True View capture and the panoramic video from the Intel® True VR capture then encodes the video sources. The video encoder software converts the uncompressed videos, into a compressed digital video format. The system is flexible to support industry standard codecs as H.264 (AVC), H.265 (HEVC), and M-JPEG to support a wide range of platforms and devices. The stream packager takes the encoded videos and converts the memory bits into consumable bitstreams. The content distribution network (CDN) enables the distribution component of the pipeline to take stream packages then stream the content to client players.

On the client side, the stream de-packager takes and reads from the consumable bitstreams and converts it into a format that the decoder can understand. The client player's video decoder takes the compressed video, then decompresses the video into a series of images. The video renderer takes the series of images and renders them sequentially into what is seen on an end users' device.

D.2.3.1.4 Intel Sports Production Team

During events, just as sports teams work together on the field, the production team covering the event also excels through teamwork. The team consists of producers, technical and production managers, audio and video technicians, graphics and replay operators, commentators, and more. Like an orchestra, these unique skillsets combine together to gather, curate and produce the content in the most compelling way possible for the sport. The quality control of the curated content goes through a master control engineer filter who reviews the content for consistency and app readiness. The content is then distributed by the CDN. There is a post-production team that takes the archived content and edits the full replay into highlights or segments. Shown in *Figure D.2.3.1.4.1* is the production team looking over the multiple camera streams in preparation for a live event.



Figure D.2.3.1.4.1: Intel Sports Production Team

D.2.3.1.5 Intel Sports Content Management System

The Intel Sports immersive media platform consists of a content management system (CMS) that enables our business partners to control which video streams are outputted to immersive media streaming applications. The CMS stores the location of the CDN streams and manages the output streams shown to end users.

D.2.3.2 Suitable Media Profiles

Intel Sports immersive media client players support wide platform choices, such as HMDs, mobile devices (Android and iOS), and web browsers in order to reach many audiences. The settings of stream types, codecs, compression quality, picture resolutions, and bitrates for immersive media broadcasting are highly dependent on the target end user experiences and expected network conditions.

The deployment uses the most common and standards compliant media codecs and streaming protocols in our products: H.264 for video codec, AAC for audio codec, MP4 for file format, and HLS (HTTP Live Streaming) for distribution.

The configuration of bitrates targets a wide range to accommodate varying network conditions. Typically, a good rate control, such as VBR (variable bitrate) with cap, is required to ensure the encoding bitrate falls within target bandwidth range. Also see the 'alternate quality streams' section in the live immersive media deployment guidelines.

Intel Sports deployment does not have any interoperability needs as of today, because the content format of our distribution is conforming to the typical audio and video formats for online media content streaming. However, there will be an interoperability need for 360°, 3DoF+, and 6DoF VR distribution in the future because the generated content of Intel® True View and Intel® True VR includes full panoramic video and point-cloud video data.

D.2.3 Immersive Media Deployment Guidelines

D.2.3.1 Alternate Quality Streams

In order to accommodate for varying last-mile conditions, alternate stream' capabilities of streaming standards such as HLS (HTTP Live Streaming) and DASH (Dynamic Adaptive Streaming over HTTP) are utilized to provide high, medium and low quality of media streams. The client player then chooses the best suitable quality stream to latch onto during run-time. It is essential for us to configure encoder settings in such a way that there is a good spread of compression ratios in terms of display resolution and target bitrates.

D.2.3.2 Specific Endpoint for Device Type

Given the variations in displays and processing capabilities in different devices, it is not unusual for a provider to offer different end points catering to different types of client devices if the client capabilities vary a lot. Each end point can have media streams conforming to a different set of profile and level sequence level information.

D.2.3.3 Seamless Multi-Camera Transitions through Continuous Audio

In order to not break immersion when switching between camera viewing feeds during a live event, it's important to have a consistent audio stream. An example would be how during a live event at a concert or sporting game, there is one unified audio stream that's independent of the viewing camera angle the end user is currently viewing.

D.2.3.4 Time-Sync between Audio and Visuals

During the production of immersive content, it's important to have time synced between the audio and visuals. This becomes extremely important during time sensitive sports, like the final buzzer in basketball. A disconnect between audio and visuals will break users' sense of presence.

D.2.3.5 Identifying Areas of Greatest Interest

The live event type will depict where the cameras should be placed for the best viewing angles. Every sport has designated points of interest where the most action takes place. The static cameras should be positioned at these locations, so that they are as close to the action as possible. The dynamic cameras of systems such as volumetric capture, should focus their attention on the highest point of interest during the live event. In sports such as basketball, most of the time the area with the greatest point of interest is the ball.

D.2.3.6 Curating a Personalized Live Immersive Experience

Many immersive media streaming platforms and app development tools already have analytics built in. During a live event, analytics can be used to determine the best stream for the device type of the user. Because live events are open to real-time troubleshooting, their production mimics more that of live broadcast television than produced on-demand content. When a camera pod encounters issues, it's important that the end user is quickly navigated toward a working camera pod and if possible, their audio should not cut out.

D.2.3.7 Maintaining High Quality Streams

When deploying a live immersive media experience, it's important to deliver the highest quality streams to fans. Poor video quality and buffering frustrates end users, deterring them from continuing to view the experience. Ensuring the fans of your live immersive media experience maintain a high-bandwidth internet connection, is key to delivering a high quality experience. Giving a user the ability to pause or rewind during a live event gives users control over their viewing experience.

D.2.3.8 Increasing User Engagement

The use of storytelling and easy to navigate user interface capabilities are important for user retention in an app. Fans need to feel immersed, so they stay within the viewing experience. During a live immersive experience, audio and visual cues should guide users toward the action.