

A dynamic splash of red liquid, possibly paint or ink, against a black background. The splash is centered and spreads outwards, creating a sense of motion and energy. The liquid forms various shapes, including droplets and larger, more complex splatters.

True Audio Next and Multimedia AMD APIs in game and VR application development

R A D E O N
SOFTWARE

Mikhail Mironov
Pavel Siniavine

March 2, 2017

Agenda

- ▶ Multimedia and games – use cases
 - ▶ TrueAudio Next (TAN)
 - ▶ Advanced Media Framework (AMF)
 - ▶ Game Engine Integration
 - ▶ 360 Video
-

Multimedia and Games

- ▶ Multimedia is an interactive application which combines multiple sources including video, audio, text and static images.
- ▶ Games use all of these content forms:
 - ▶ Video as animations or backgrounds
 - ▶ Audio
 - ▶ Text
 - ▶ Images as textures
 - ▶ Games are highly interactive
- ▶ Are games “multimedia”?

Video in Games and VR

Use cases

- ▶ 360 video as a background for a VR experience or a game
- ▶ Watching a movie on a virtual screen inside VR space
- ▶ Cameos (video inserts) as part of the game
- ▶ Live video chat with players
- ▶ Animated textures
- ▶ 360 video playback
- ▶ Live video streaming, eSports
- ▶ Remote displays
- ▶ Augmented reality



Audio in Games and VR

- ▶ Traditional 3D sound rendering
- ▶ Spatialization and Acoustic-based audio modelling
- ▶ Ambisonic audio in 360 video playback



Multimedia in Games and VR: Common Problems

- ▶ Performance... performance... performance...
 - ▶ Real-time environmental positional (3D) sound rendering for multiple sources requires a lot of computational power
 - ▶ Decoding 4K videos is a heavy CPU task
 - ▶ Even 2K HD encoding might take the most of your CPU
 - ▶ What about laptops?
- ▶ Solution: use hardware acceleration when available
 - ▶ Universal Video Decoder (UVD)
 - ▶ Video Compression Engine (VCE)
 - ▶ Compute queue for audio

TrueAudio Next

Why

Enabling Audio Presence

- ▶ Conventional gaming audio rendering is only satisfactory for flat screen games
- ▶ Virtual Reality demands full immersion
- ▶ People can instinctively hear the difference between real world and its rough approximations
- ▶ For immersive presence each sound source needs a dynamic update of impulse response based on user position and environment, applied as convolution in real-time
- ▶ CPU can process a limited number of concurrent audio streams





Computation of Conventional Gaming Audio

- ▶ Conventional gaming audio approximates reality audio with:
 - ▶ Simple low pass filters and sample rate conversion (occlusion and Doppler pitch shift)
 - ▶ Attenuation (far sounds are softer)
 - ▶ Head-Related Transfer Functions for 3D positioning (direct path only)
 - ▶ Global Reverbs (1 reverb for all sounds in the room plus 3 to 4 sources)
 - ▶ Mixer
- ▶ Conventional gaming audio is like a 3D mesh with no texture or lighting...




Creating Audio Presence Requires Acoustics-based Modeling

- ▶ When we hear a sound in the real world, we hear a dynamic superposition of a direct path, plus reflected, diffused and diffracted sound.
- ▶ Human brain has been trained to know instinctively if a sound is real by hearing these properties



Benefits of Acoustics-modeled Gaming Audio

- ▶ Can help enable delivery of immersive presence in Virtual Reality games
- ▶ Can simplify audio design process
 - ▶ Should not need to fabricate artificial environment filters in a studio:
 - ▶ The old way – tweak and re-tweak filters
 - ▶ The new way – adjust the physical room materials once for all sound sources in the room. Filters are auto-generated
 - ▶ Real time response is calculated as the user moves within a dynamic environment
- ▶ Can help enable support of infinite variation in sound characteristics, enhancing gaming experience



TrueAudio Next: SDK

- ▶ A convenient SDK for integration with proprietary audio VR game engines
 - ▶ Real-time compute queue
 - ▶ Dedicated compute units
 - ▶ Uses OpenCL™ for computations
 - ▶ Convolutions
 - ▶ Helper functions
- ▶ Delivery
 - ▶ Headers
 - ▶ Runtime open source code
 - ▶ Documentation
 - ▶ Game engine integration TAN-based plug-in as open source code




TrueAudio Next: Integration

- ▶ Convolution acceleration for Wwise Audio engine and 3rd party advanced plug-ins
- ▶ AMD simple audio VR modelling + TAN convolution integrated into Unreal game engine




TrueAudio Next: Hardware Design Support Details

- ▶ Real Time(RT) compute and CU allocation support
 - ▶ Reserve a specified number of CUs
 - ▶ Dedicates a real-time priority queue to the reserved CUs
 - ▶ Applications must thoroughly characterize their workloads to support a wide range of systems
- ▶ Audio processing can be done on an APU, in the APU+dGPU case
- ▶ Audio processing is inherently scalable in quality and quantity by:
 - ▶ The number of streams
 - ▶ The length of convolution
 - ▶ The update rate
- ▶ True Audio Next SDK allows audio processing to be distributed between CPU and / or GPU



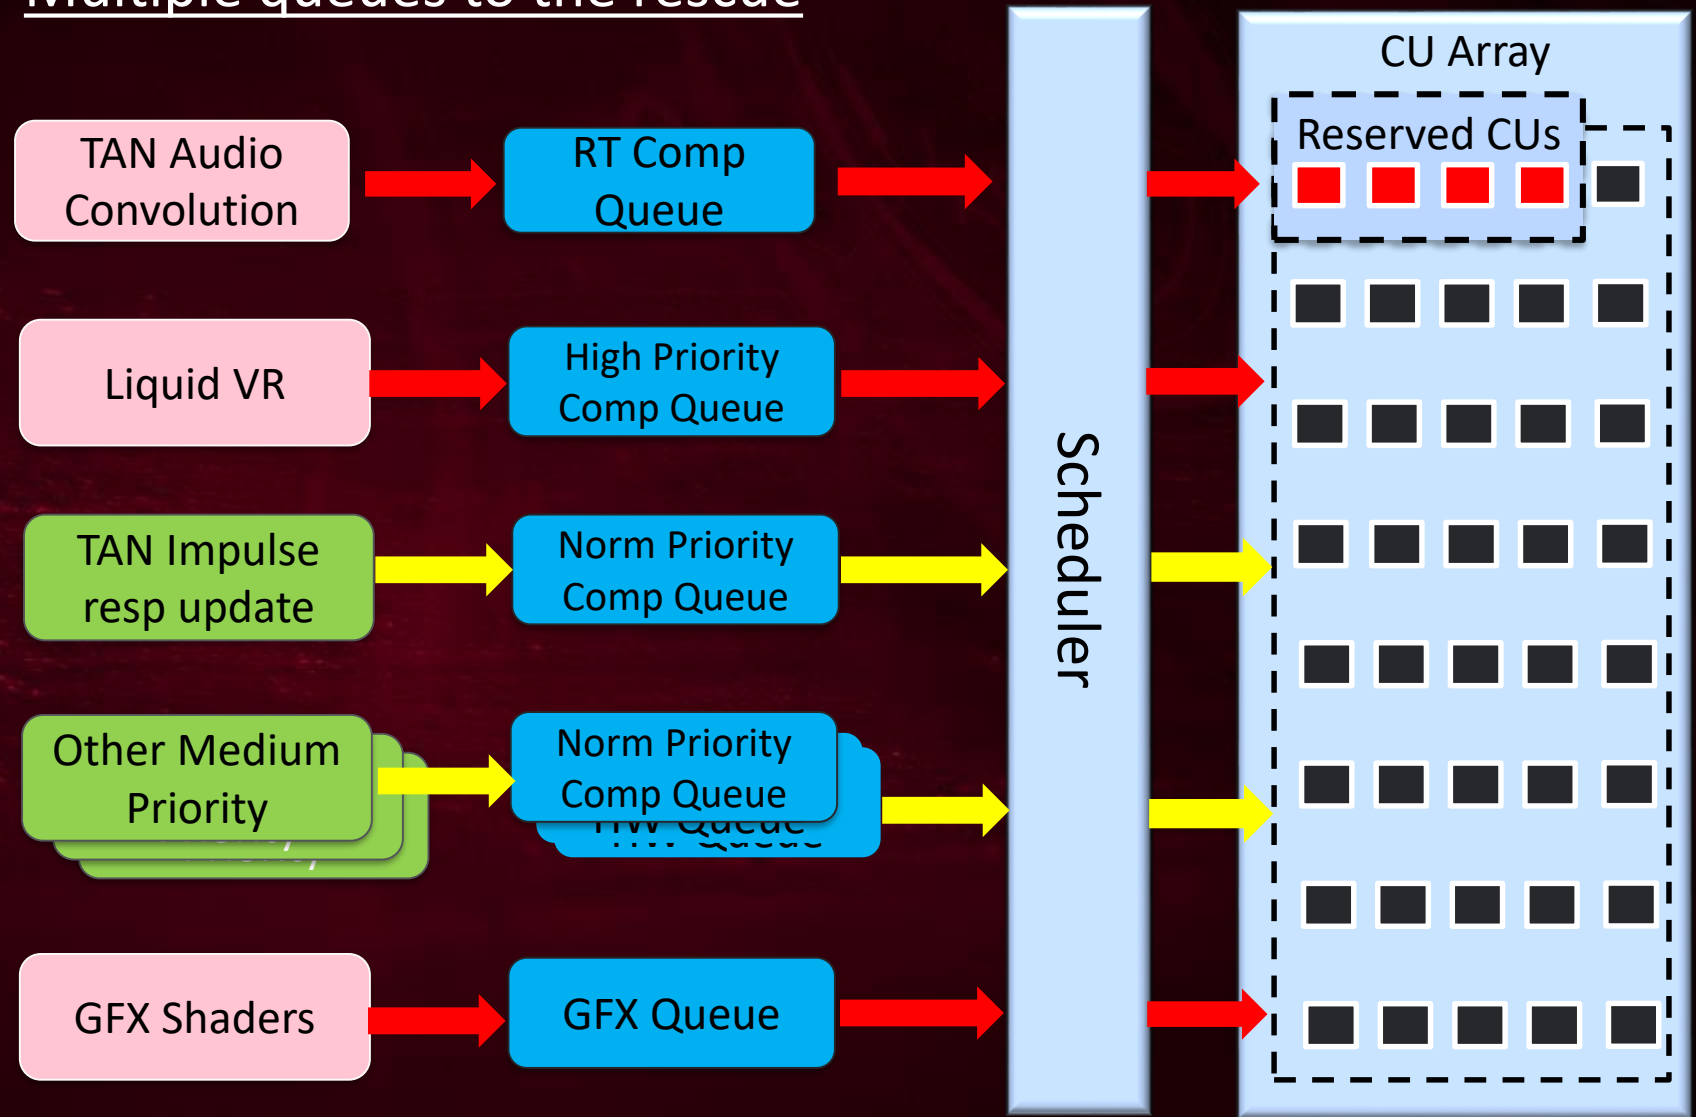
TrueAudio Next: Performance

- ▶ For a 2-second convolution:
 - ▶ A CPU-based solution using one dedicated core can deliver up to 8 stereo audio streams. Adding more cores cannot guarantee quality of service.
 - ▶ 4 Compute GPU units can process up to 32 stereo audio streams within QoS




Rendering Audio and Other Low-latency Workloads on GPU

- ▶ TrueAudio Next uses the Real Time (RT) queue and a dynamically reserved set of CUs
- ▶ Liquid VR uses own high-priority queue
- ▶ TAN uses a normal priority queue to update impulse responses



TrueAudio Next Demo



TrueAudio Next: Summary

- ▶ TAN provides a practical platform for acoustic-modelled game audio
- ▶ Allows real time audio to leverage the OpenCL™ runtime and co-exist with graphics
 - ▶ To AMD knowledge this has never been done before
- ▶ Game developers can budget and allocate compute resources dedicated to audio based on the needs of their game
 - ▶ To AMD knowledge nothing like this available in CPU-based implementations

Advanced Media Framework



AMF SDK

AMF SDK is an API providing:

- ▶ Easy access to AMD dGPU/APU hardware blocks:
 - ▶ Video Compression Engine (VCE) - encoder
 - ▶ Unified Video Decoder (UVD) - decoder
- ▶ Easy interoperability with the relevant GPU APIs:
 - ▶ DX9, DX11, DX12
 - ▶ Vulkan
 - ▶ OpenCL
 - ▶ OpenGL
- ▶ Common building blocks for multimedia applications:
 - ▶ Color space conversion
 - ▶ Scaling
 - ▶ Composition
 - ▶ Compute
- ▶ Sample implementations of various GPU-accelerated pipelines



AMF SDK: Use Cases

▶ Gaming and VR


- ▶ Remote gaming
- ▶ Cloud gaming
- ▶ Video textures in games
- ▶ Wireless VR
- ▶ 360 video playback
- ▶ 360 video stitching

▶ Traditional

- ▶ Video playback
- ▶ Transcoding
- ▶ Video editing
- ▶ Remote desktop
- ▶ Video conferencing
- ▶ Medical imaging

AMF SDK: Helps Enable

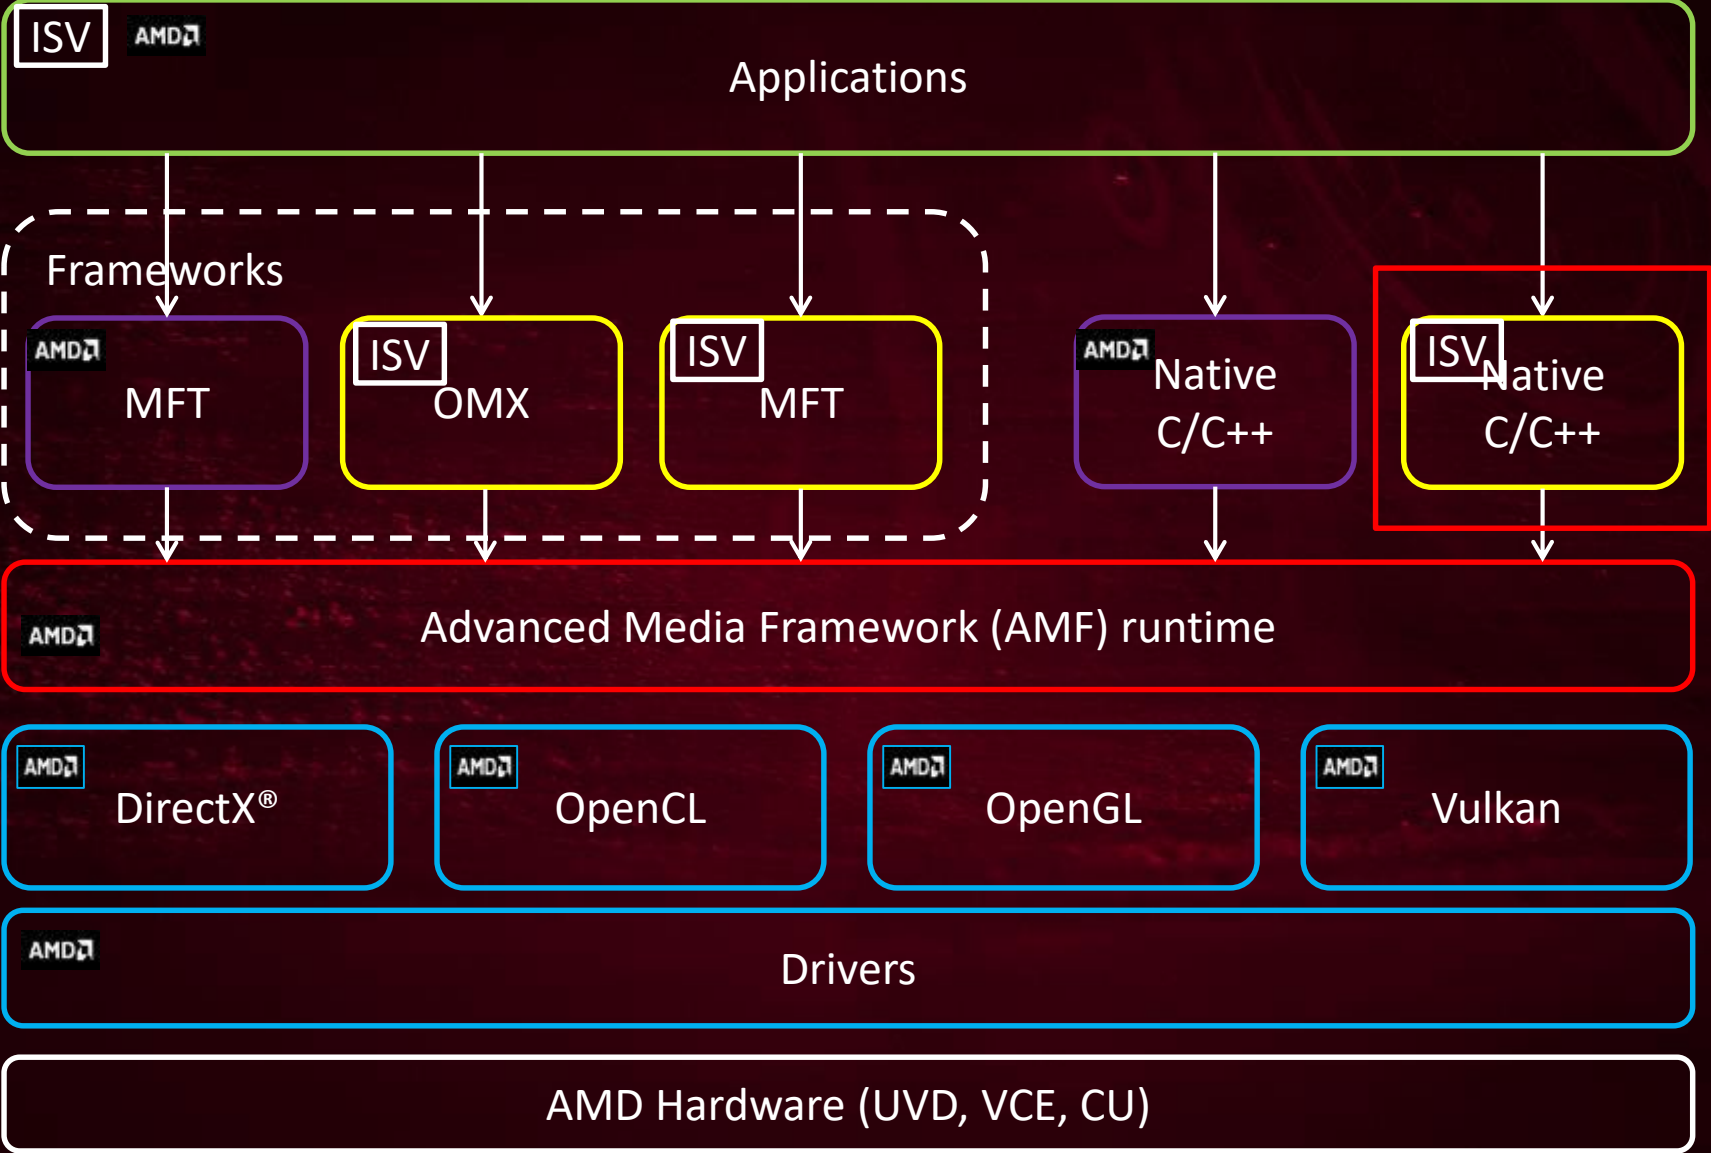
- ▶ Video Decoder: H.264, VC1, WMV9, MPEG2, MPEG4 P2, H.264 MVC, MJPEG; HEVC(H.265)
- ▶ Video Encoder: H.264, AVC SVC, HEVC(H.265)
- ▶ Video scaler and color space converter
- ▶ Capability Management
- ▶ Simple Compute API



Distribution and Samples

- ▶ AMF Runtime is distributed with the Crimson public driver
 - ▶ Easy updates
 - ▶ Application should not require distribution of AMF binaries
 - ▶ Avoids DLL conflicts
 - ▶ Designed to support current hardware
- ▶ AMF SDK includes:
 - ▶ Header files
 - ▶ Sample source code
 - ▶ Helper functions and classes
- ▶ AMF SDK uses Ffmpeg for file certain format support and audio: binaries, libraries and headers (included in the package)
- ▶ AMF SDK is published on GPUOpen / GitHub
- ▶ Released under the MIT license and complies with the FFmpeg's LGPL v2.1 license

AMF Software Stack



AMF Projects



AMF: Game Engine Integration

▶ Unreal

- ▶ Default implementation based on Media Foundation (MF) framework with a sample grabber that works via system memory. It is hard to control and adjust timing in MF-based pipelines

▶ Unity

- ▶ Default implementation uses a CPU-based video playback pipeline based on QuickTime.

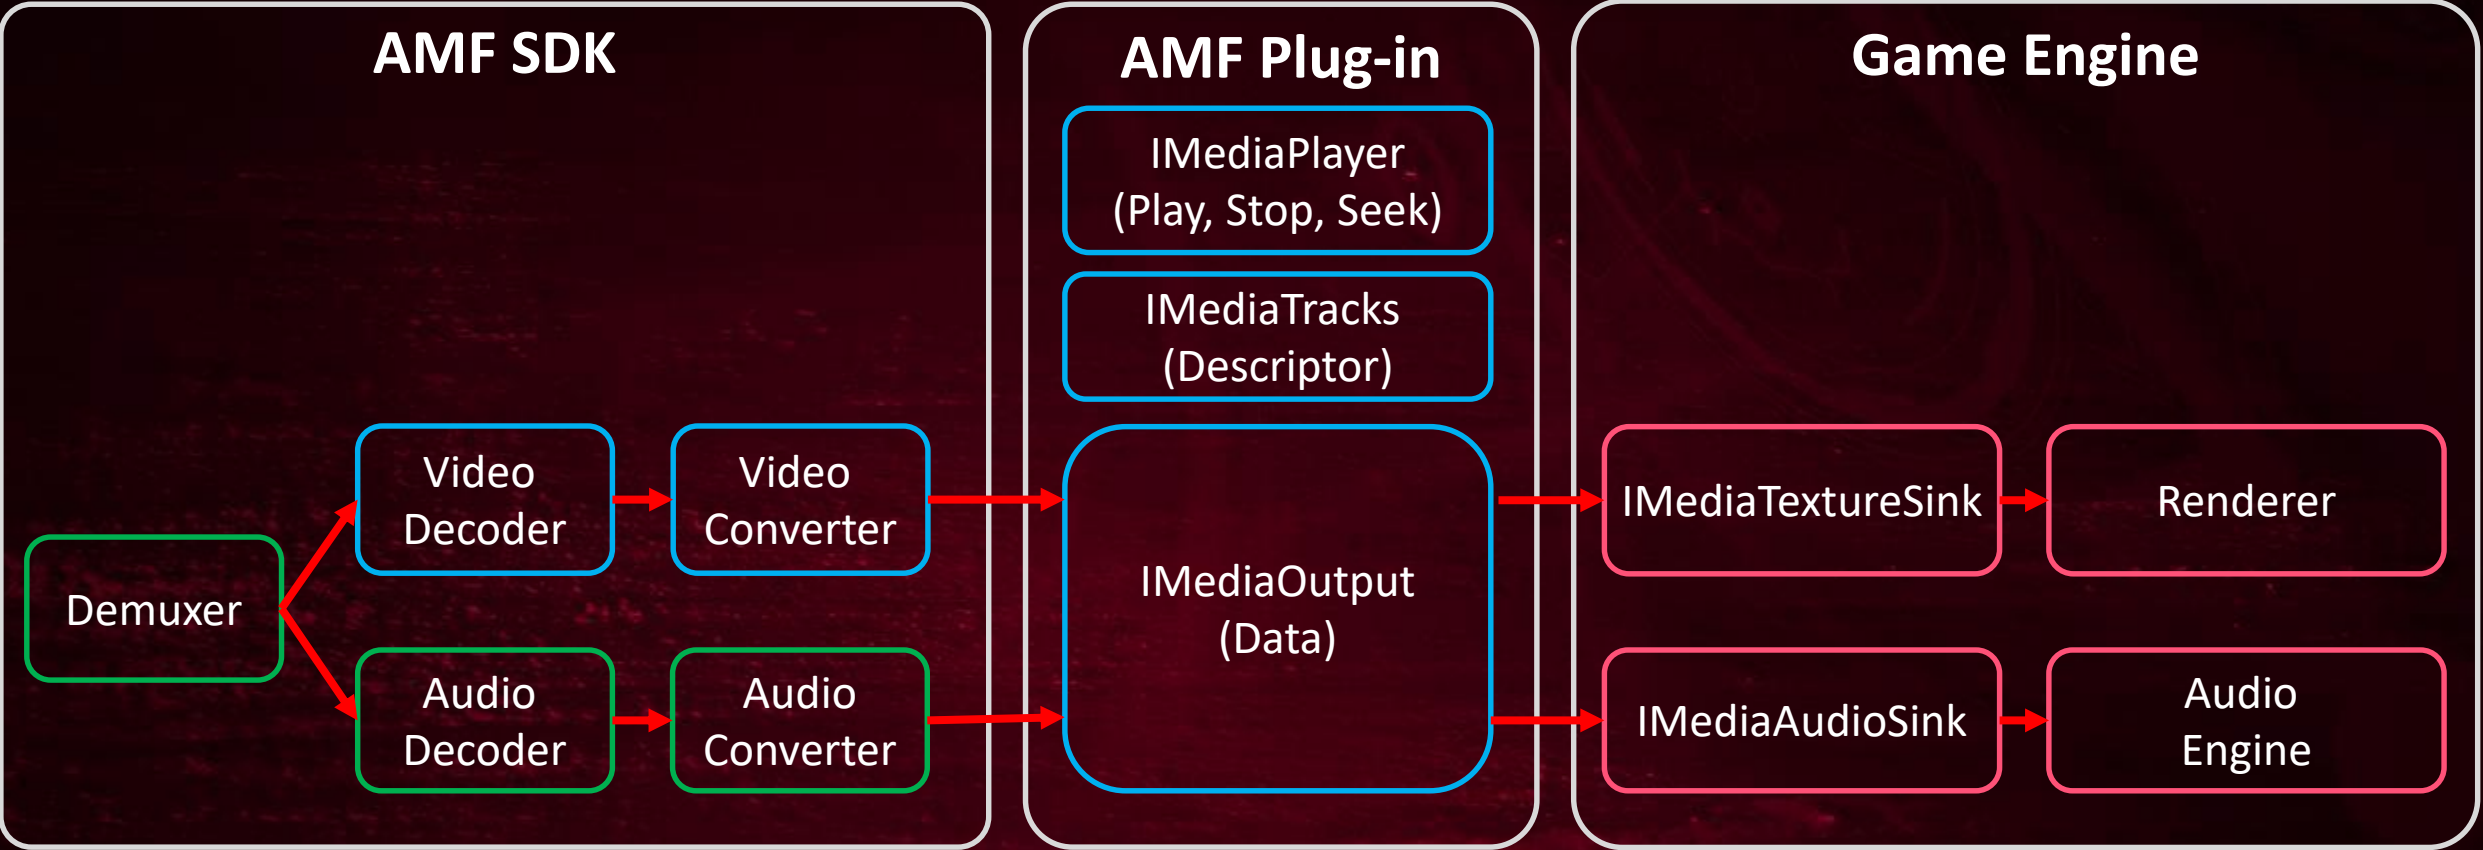
▶ Demand

- ▶ Game developers want to integrate high quality video streams into games without sacrificing performance
- ▶ VR games and experiences could benefit from the use of 360 videos as backgrounds

AMF: Game Engine Integration - AMD Offers:

- ▶ AMF-based open source video plug-ins replacing the default implementations when the AMD hardware is present
 - ▶ Unreal
 - ▶ Unity
 - ▶ Other engines planned to follow
- ▶ Live streams are possible
- ▶ In the future the existing integration with Twitch could be replaced with hardware-accelerated encoding

AMF Game Engine Integration: Unreal Engine 4 Pipeline



AMF 360 Video Player: Challenges



- ▶ These days 360 video files / streams have up to 4Kx2K resolution in equirectangular format. With HMD field of view of about 100 degree the visible portion of the image is 1Kx0.5K. HMDs have higher resolution than video input - not enough pixels.
- ▶ 360 Video is encoded to 30 (less frequently to 60) frames per second, while many HMDs run at 90 frames per second. Not enough frames.
- ▶ Transmission of 4Kx2K videos requires a lot of network bandwidth. Many providers limit video stream bitrate, degrading quality. Not enough details.
- ▶ Majority of 360 players use software decoding or open source back-ends based on DXVA (DX9) for decoding while none of the HMD APIs supports DX9. Performance issues.

AMF: 360 Video Player Sample



- ▶ A 360 video player in the form of a sample source code based on DX11
- ▶ Integrated with major HMD vendors
- ▶ Demonstrates the use of full hardware acceleration and timing
- ▶ Renders ambisonic audio
- ▶ Showcases advanced video processing algorithms to compensate for low video quality and low frame rate

Call for Action

- ▶ Get TAN SDK from GPUOpen and GitHub:
 - ▶ <http://gpuopen.com/gaming-product/true-audio-next/>
 - ▶ <http://gpuopen.com/gaming-product/true-audio-next/>
- ▶ Get AMF SDK from GPUOpen and GitHub:
 - ▶ <http://gpuopen.com/gaming-product/advanced-media-framework/>
 - ▶ <https://github.com/GPUOpen-LibrariesAndSDKs/AMF>
- ▶ AMF plug-in for Unreal Engine 4 planned to be available soon
- ▶ 360 Player sample planned to be available soon
- ▶ Try to integrate the SDKs into your engine
- ▶ Stay tuned for new features
- ▶ Contact:
 - ▶ Mikhail.Mironov@amd.com
 - ▶ Pavel.Siniavine@amd.com
- ▶ Special thanks to our AMD colleagues:
 - ▶ Carl Wakeland
 - ▶ Gabor Sines

Q&A



Disclaimer & Attribution

DISCLAIMER

- ▶ The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.
- ▶ The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.
- ▶ AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.
- ▶ AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions.

Other names are for informational purposes only and may be trademarks of their respective owners.

DirectX is a registered trademark of Microsoft Corporation" and "OpenCL" is a trademark of Apple Inc. used by permission by Khronos.

© 2017 Advanced Micro Devices, Inc. All rights reserved.