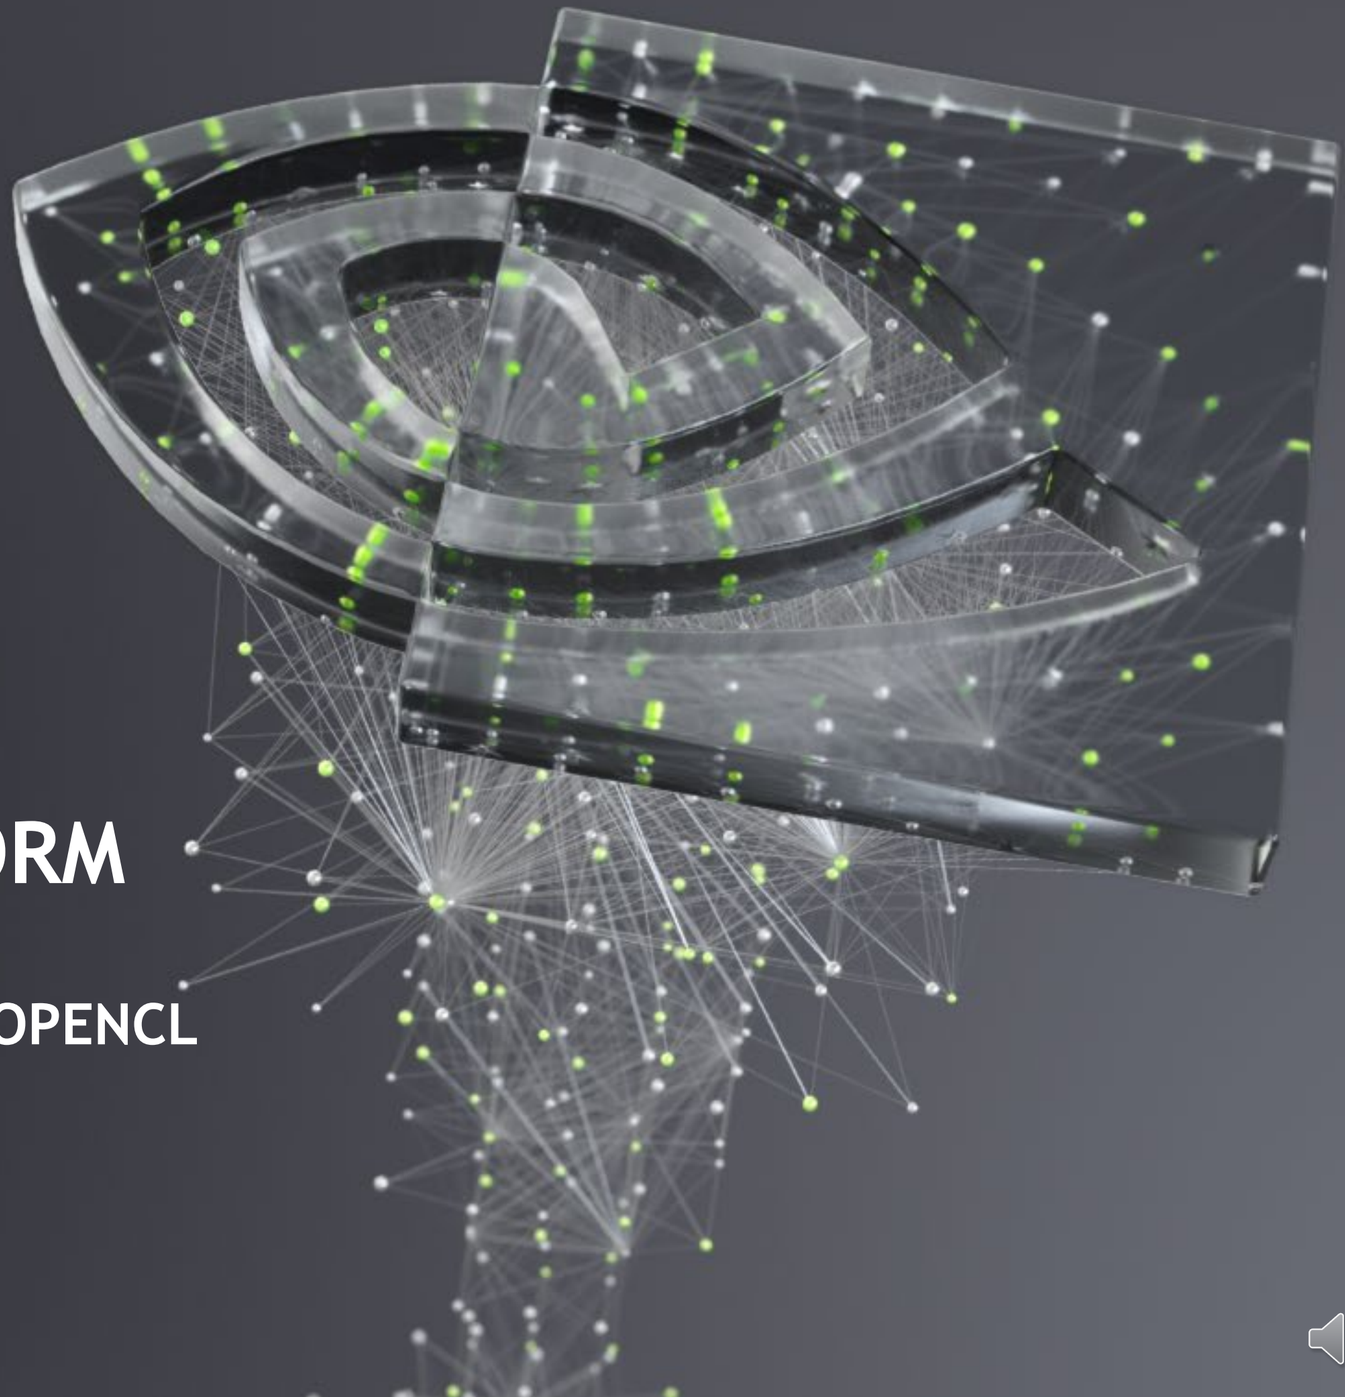

KHRONOS CROSS-PLATFORM STANDARDS UPDATE

VULKAN, ANARI, OPENXR, GLTF AND OPENCL

Neil Trevett, GTC, March 2020
VP NVIDIA, Khronos President

AGENDA

Latest updates on key Khronos open standards

That are of most interest to the GTC audience
Vulkan, ANARI, SPIR-V, OpenXR, glTF and OpenCL

On Khronos format slides

NVIDIA and Khronos Standards

How NVIDIA is supporting and deploying
Khronos Open Standards

On NVIDIA format slides

Khronos Connects Software to Silicon

Open interoperability standards to enable software to effectively harness the power of multiprocessors and accelerator silicon

>150 Members ~ 40% US, 30% Europe, 30% Asia

3D graphics, XR, parallel programming, vision acceleration and machine learning

Non-profit, member-driven standards-defining industry consortium

Open to any interested company

All Khronos standards are royalty-free

Well-defined IP Framework protects participant's intellectual property

Khronos Active Initiatives

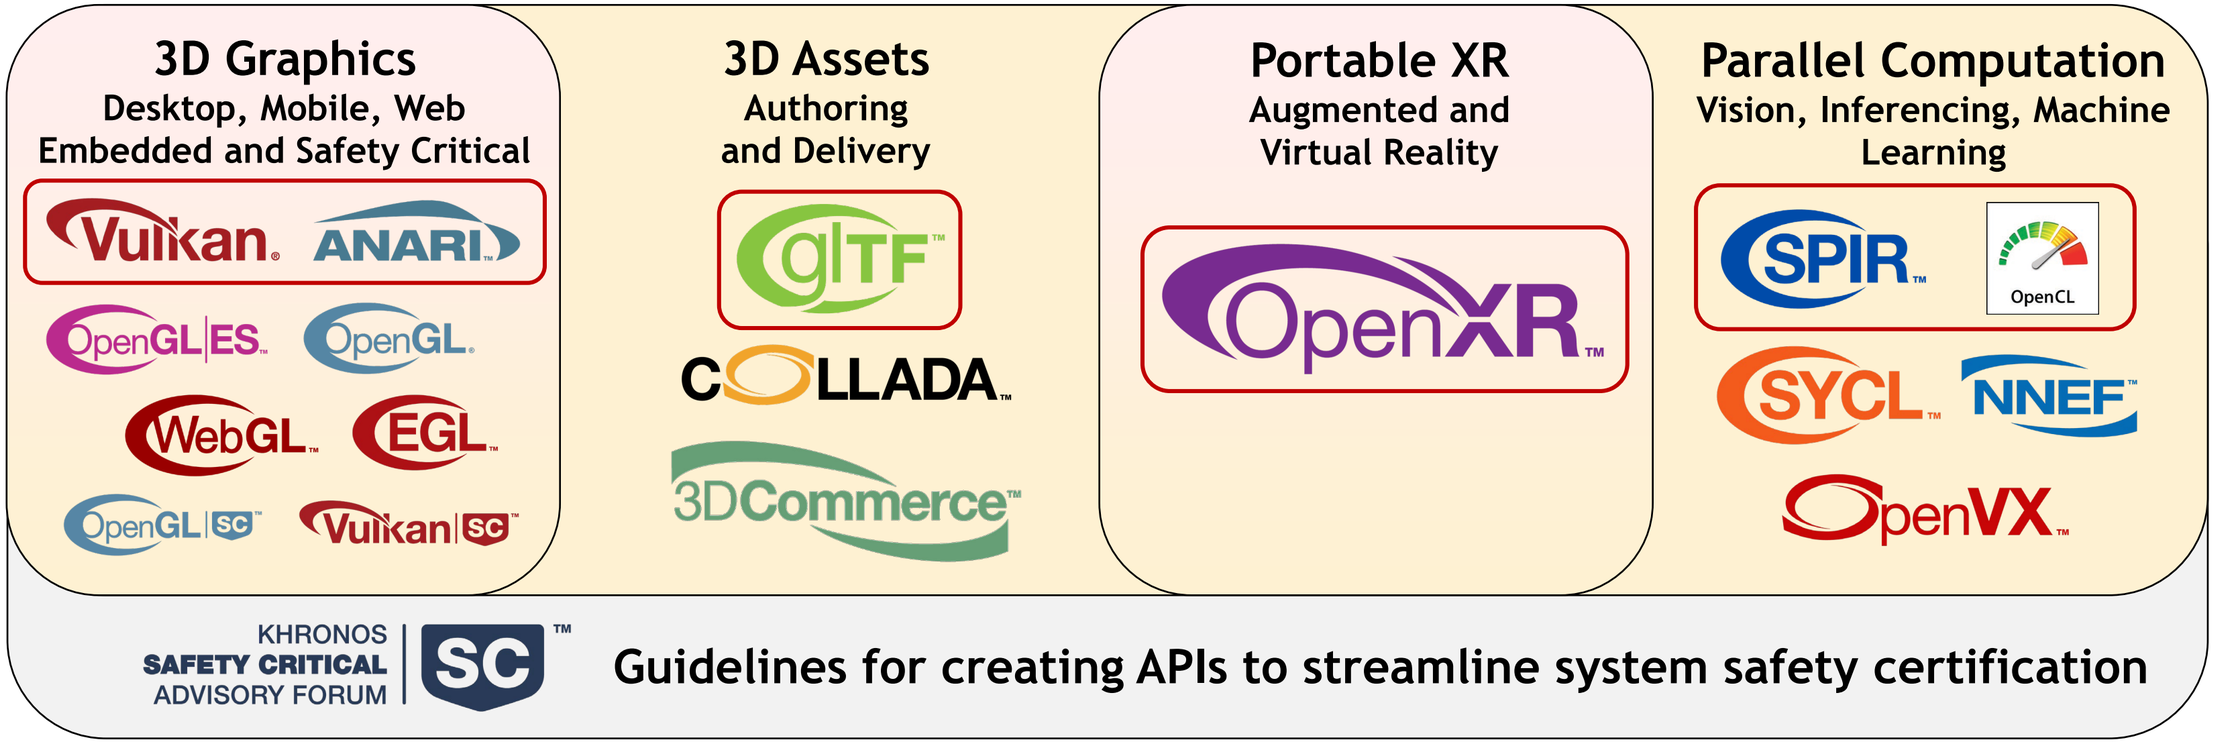

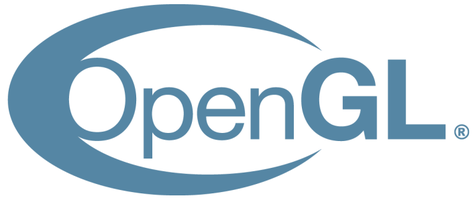

NVIDIA is fully committed
to continuing to support OpenGL
Exploring expanded interop with Vulkan
New functionality is primarily Vulkan-focused

Pervasive Vulkan

Desktop and Mobile GPUs

AMD arm BROADCOM Imagination intel NVIDIA QUALCOMM VeriSilicon

http://vulkan.gpuinfo.org/

Platforms

Desktop

Android
(Android 7.0+)
(Vulkan 1.1 required on Android Q)

iOS macOS

Apple
(via porting layers)

Media Players

Consoles

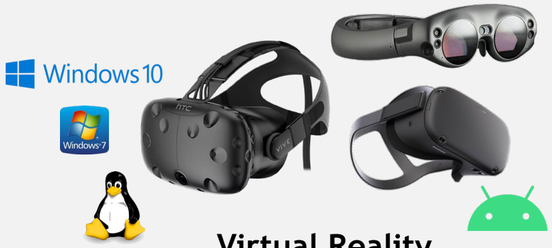

Virtual Reality

Cloud Services

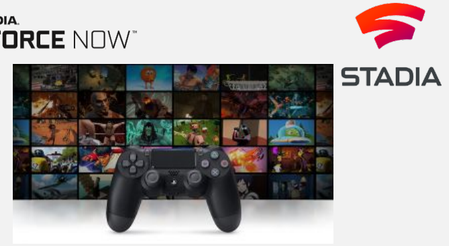

Game Streaming

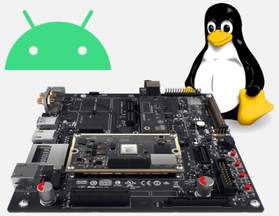

Embedded

Engines

CIDER CINDER XENKO GODOT CRYENGINE unity source id UNREAL ENGINE Croteam Serious Engine UX3D NetEase Games Passion of gamers

Note: The version of Vulkan available will depend on platform and vendor

Vulkan AAA Content

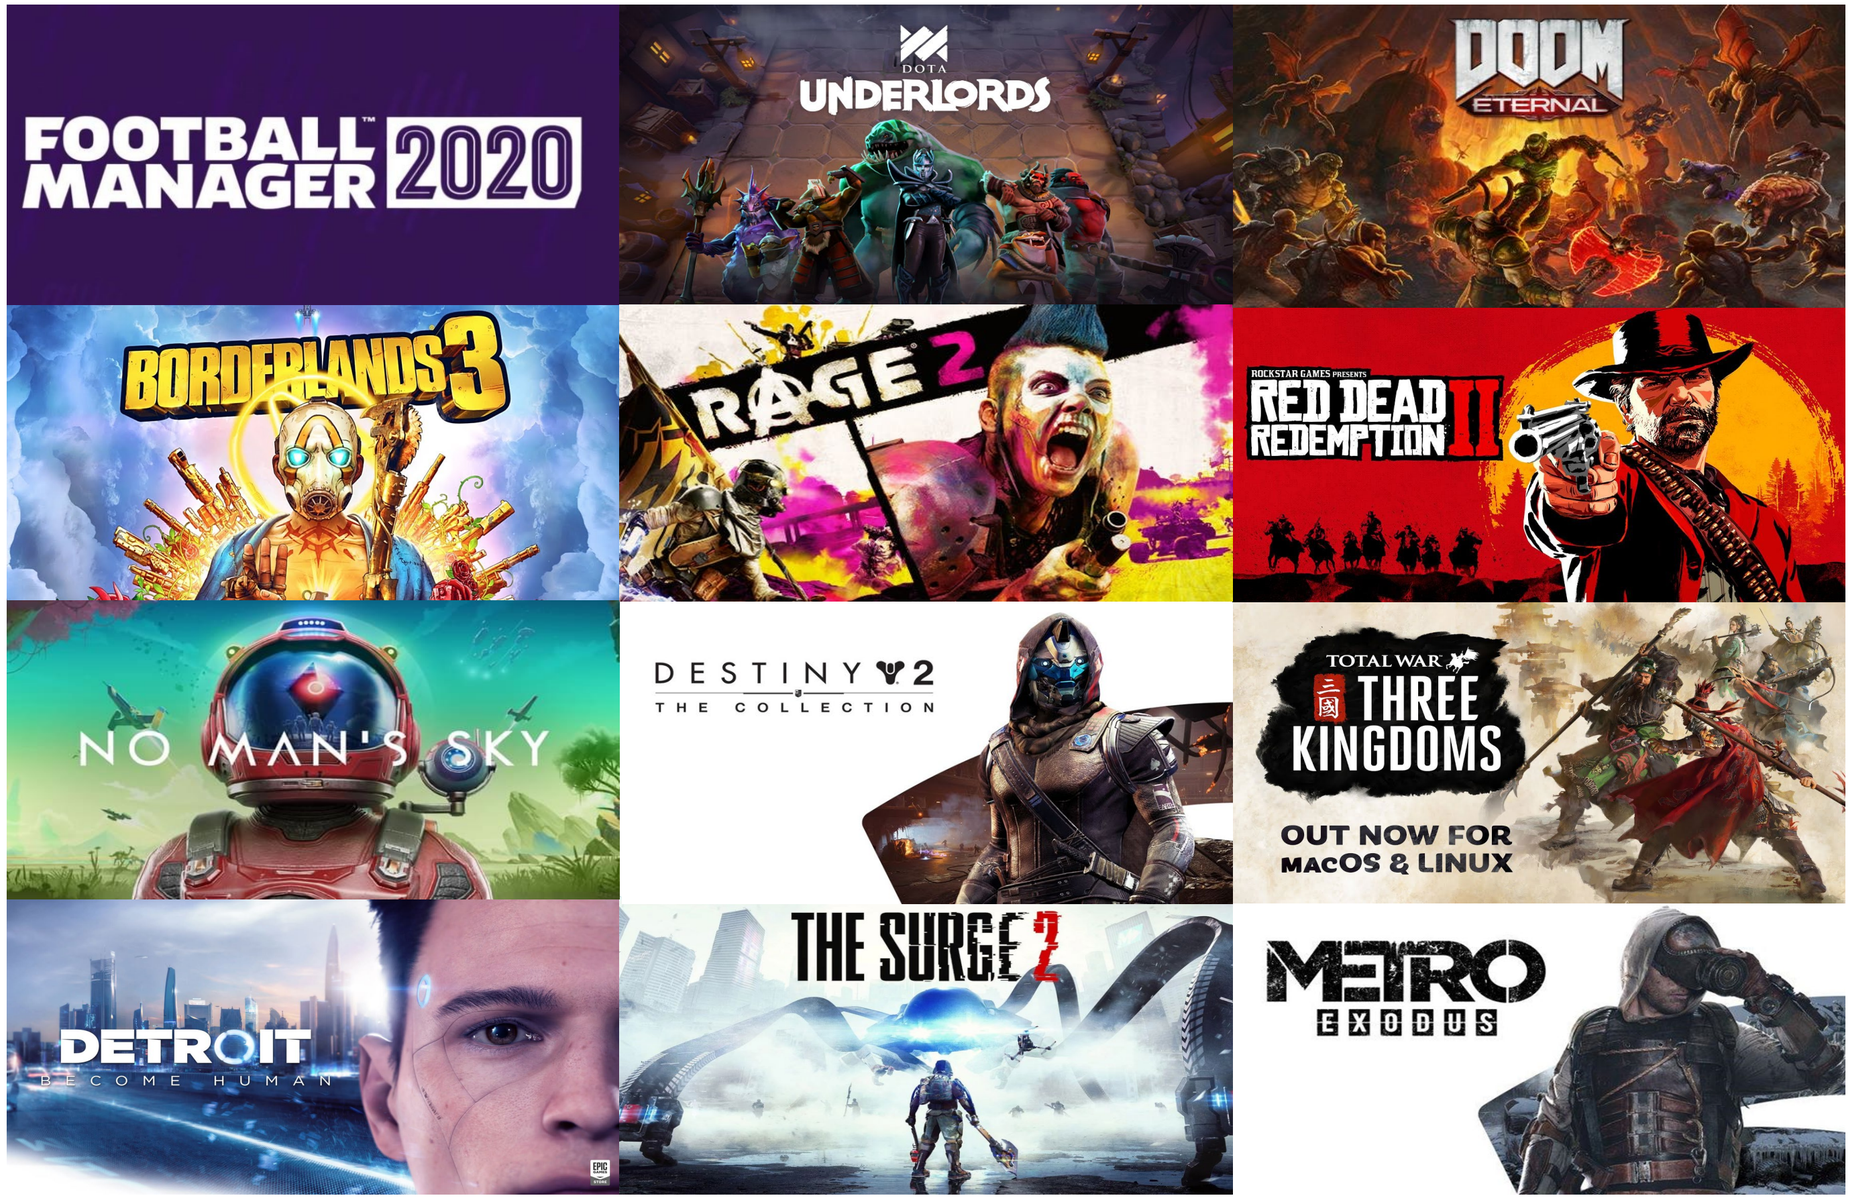

Windows 10

Windows 7

STADIA

NVIDIA GEFORCE NOW

macOS

Vulkan Roadmap

Vulkan 1.1 Extensions

Maintenance updates plus additional functionality

Timeline semaphores
DX/HLSL compatibility
Bindless resources
Reduced precision arithmetic
Formal memory model
Buffer references
SPIR-V 1.5

January 2020

Roadmap Discussions

Ray Tracing
Variable Rate Shading
Accelerated Video Encode/Decode
Machine Learning Primitives
Mesh Shaders

NVIDIA

DLSS 2.0 smart scaling can be used with Vulkan applications

Does not need per application training

Vulkan Ray Tracing

Set of Extensions to Vulkan, GLSL and SPIR-V

Seamlessly integrates ray tracing into Vulkan 1.X

Familiar Ray Tracing Pipeline Architecture

Straightforward porting between Vulkan Ray Tracing and DXR
- including re-use of ray tracing shaders written in HLSL

	Vulkan Ray Tracing	DX12 / DXR
Ray Tracing Pipelines	Yes	Yes
Ray Queries	Optional	DXR 1.1 Inline raytracing
Language for Ray Tracing Shaders	GLSL or HLSL	HLSL
Pipeline Libraries	Yes	Yes
Build Acceleration Structure on Host	Optional	No
Deferred Host Operations	Optional	No
Capture/Replay Support for Tools (e.g. RenderDoc)	Optional	No

Vulkan®

The industry's first open, cross-vendor, cross-platform standard for ray tracing acceleration

Can be accelerated on existing GPU compute and dedicated ray tracing cores

Extensions are Provisional
Launched 17th March 2020

Open to developer feedback before finalization

https://khr.io/vkrayprovfeedback

NVIDIA®

Straightforward port from NVIDIA VKRay vendor extension to KHR extensions

Shipping beta drivers today

Example code coming soon

https://www.khronos.org/registry/vulkan/

KHRONOS® GROUP

Vulkan Ray Tracing and Shading Languages

HLSL and Vulkan with DXC

Microsoft's DXC HLSL compiler was open sourced in Jan 2017
Google and others have added SPIR-V code generation to DXC
with Microsoft's knowledge and approval
Vulkan developers can now choose between GLSL and HLSL!

Vulkan Ray Tracing includes GLSL and SPIR-V Extensions

Enabling compiled GLSL/SPIR-V shaders to operate in a
Ray Tracing Pipeline - similar to HLSL features used in Direct3D's DXR

HLSL for Vulkan Ray Tracing

NVIDIA added code generation to DXC to generate SPIR-V for the
NVIDIA VKRay ray tracing vendor extension from HLSL
Vulkan Ray Tracing Extensions supported in HLSL soon

Developers can port HLSL shaders
with minimal changes between Vulkan Ray Tracing and DXR

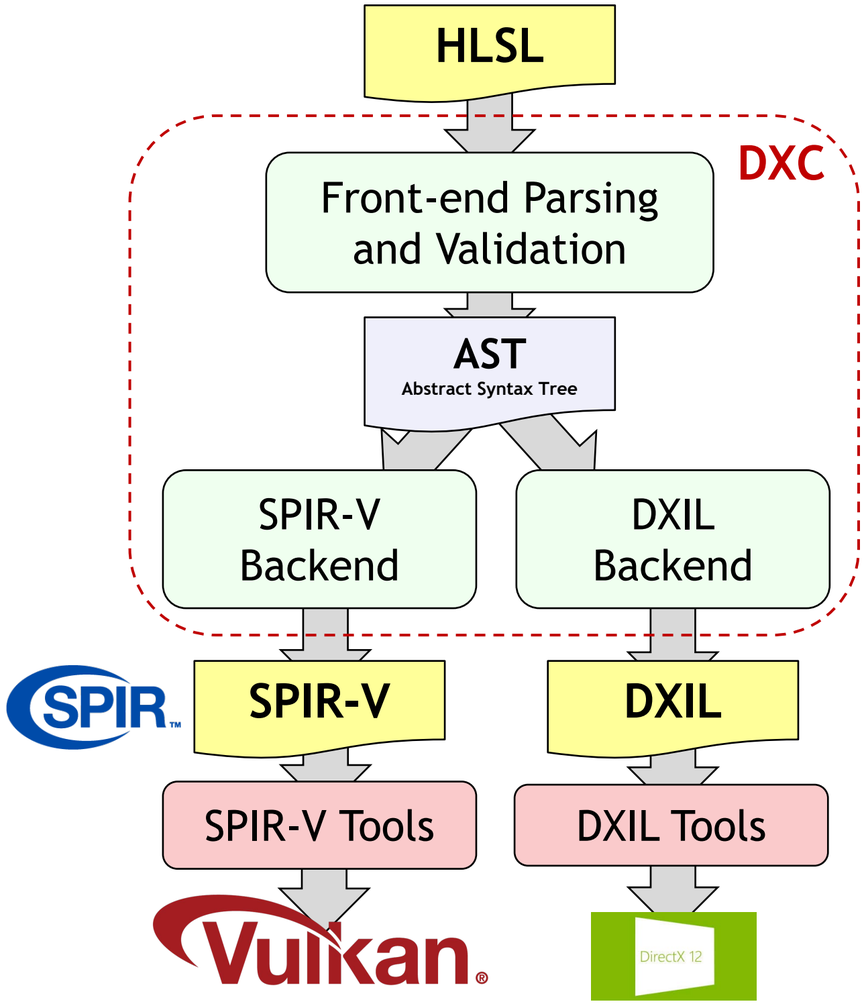

NVIDIA AND VULKAN

NVIDIA deeply engaged in Vulkan and driving extend/consolidate cycle

- NVIDIA shipped Vulkan 1.0, 1.1 and 1.2 on day of spec releases
- Shipped beta Vulkan Ray Tracing extensions on day of spec release
- Increased Vulkan support in NSIGHT 2020.2 development tool – close parity now with DX12

NVIDIA chairing multiple Vulkan initiatives at Khronos

- Ray Tracing: contributed NVIDIA VKRay to catalyze Vulkan Ray Tracing KHR extensions
- Vulkan Portability: bringing layered Vulkan to Apple, WebGPU, silicon without native drivers etc.
- Machine Learning: low-level inferencing primitives

Cloud – Windows and Linux

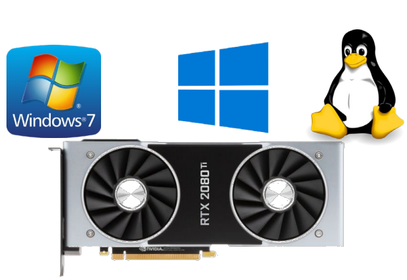

Desktop – Kepler and later

Nintendo Switch

Shield Android TV

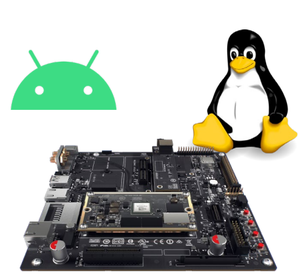

Jetson

Vulkan®

Open Source Layering Projects

Fighting Platform Fragmentation

Layers Over	Vulkan	OpenGL	OpenCL	OpenGL ES	DX12	DX9-11
Vulkan		Zink	clspv clvk	GLOVE Angle	vkd3d	DXVK WineD3D
OpenGL	gfx-rs Ashes			Angle		WineD3D
DX12	gfx-rs	Microsoft 'GLOn12'	Microsoft 'CLOn12'			Microsoft D3D11On12
DX9-11	gfx-rs Ashes			Angle		
Metal	MoltenVK gfx-rs		*clspv over MoltenVK?*	MoltenGL Angle		

Vulkan is effective porting layer for API portability and stack simplification

'Vulkan everywhere'! Even if no native drivers on platform

Working towards 'OpenCL Everywhere'!

Vulkan Portability on macOS and iOS

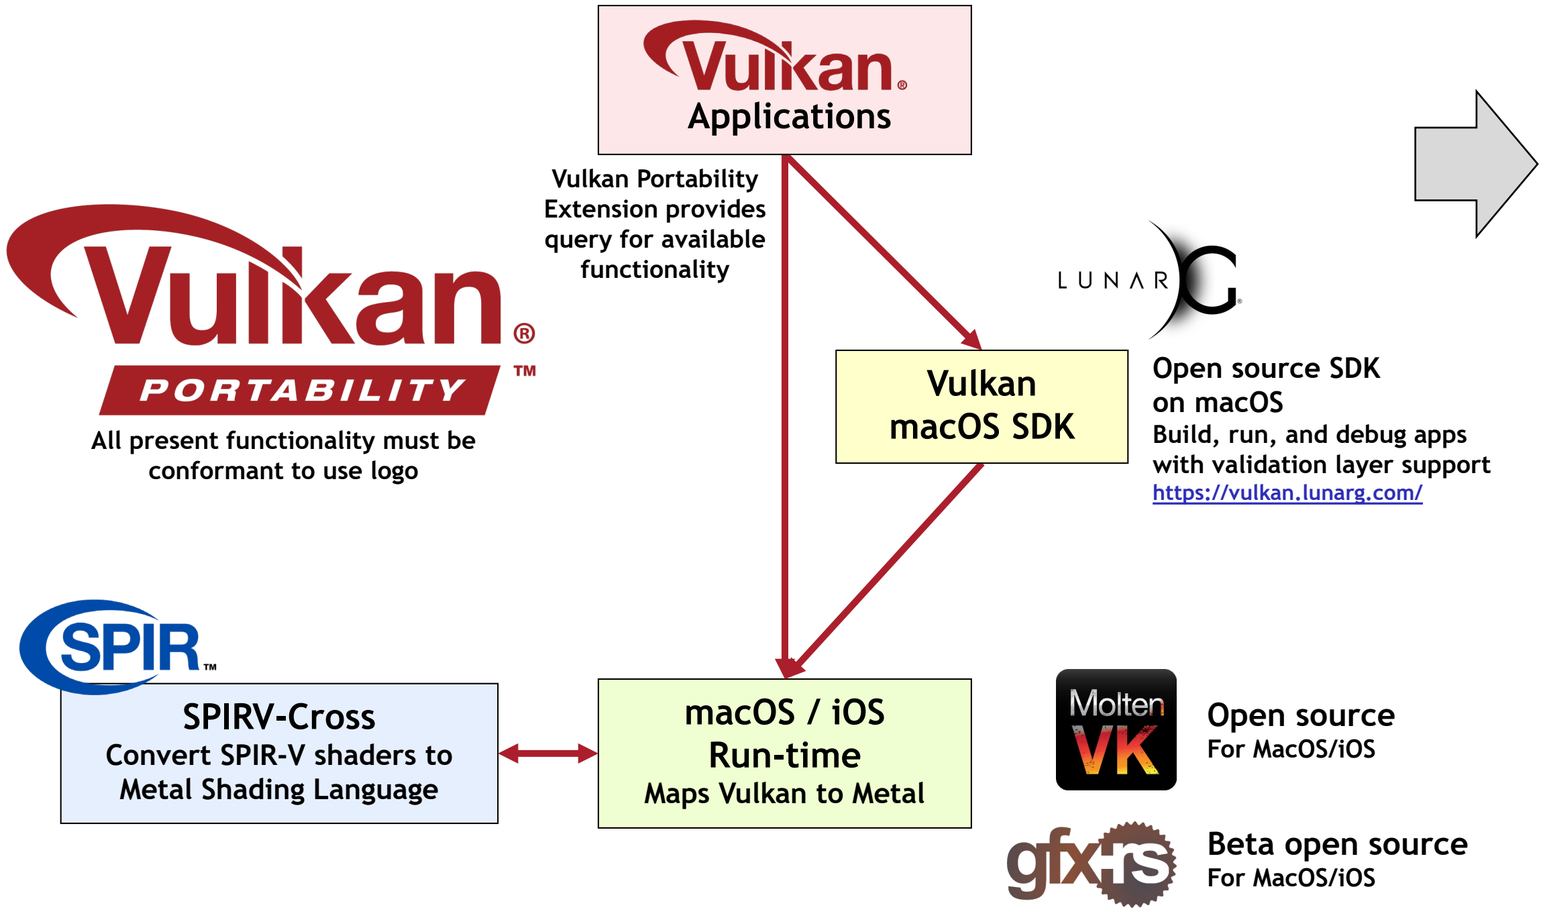

Vulkan applications shipping today on macOS or iOS using MoltenVK

KHRONOS GROUP

ANARI - Analytic Rendering API

Working Group Announced 3rd March 2020

Scientific Visualization Portability

APIs to describe objects in a scene
- the renderer does the rest

Ray tracing was key motivation
- but can drive any renderer

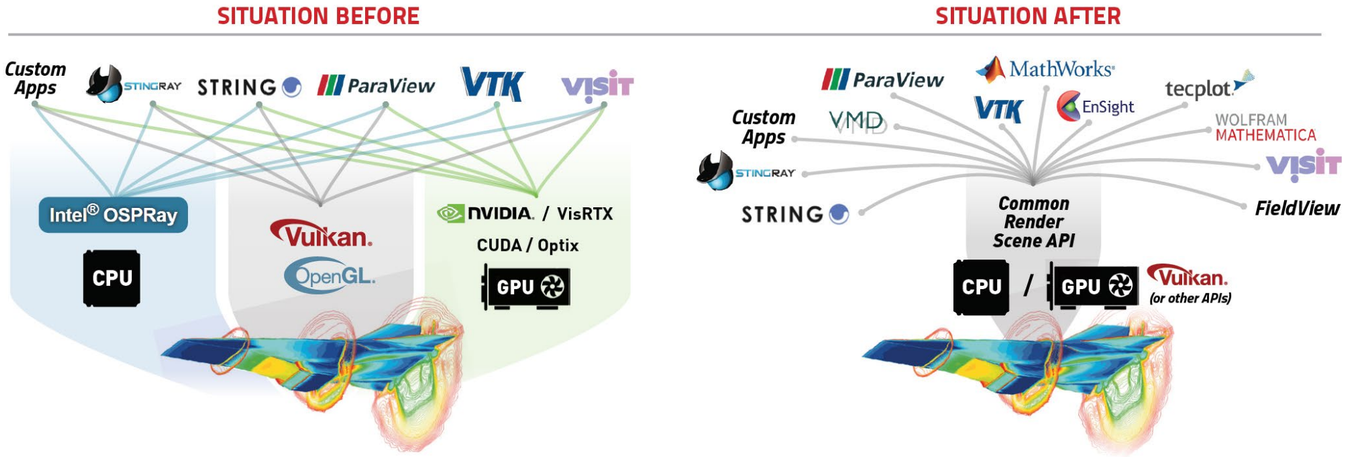

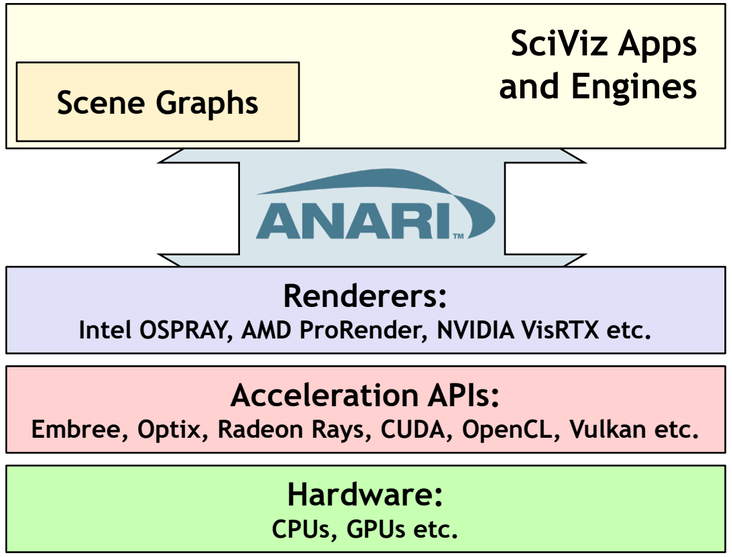

Industry Support

OpenXR - Cross-Platform Portable AR/VR

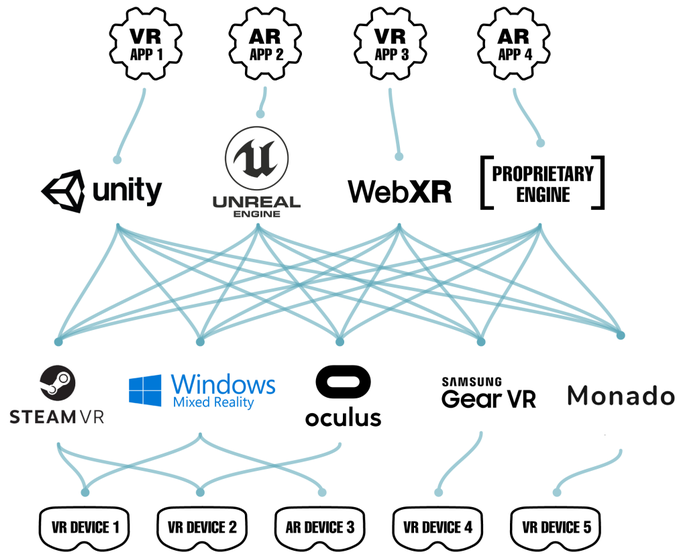

Before OpenXR

XR Market Fragmentation

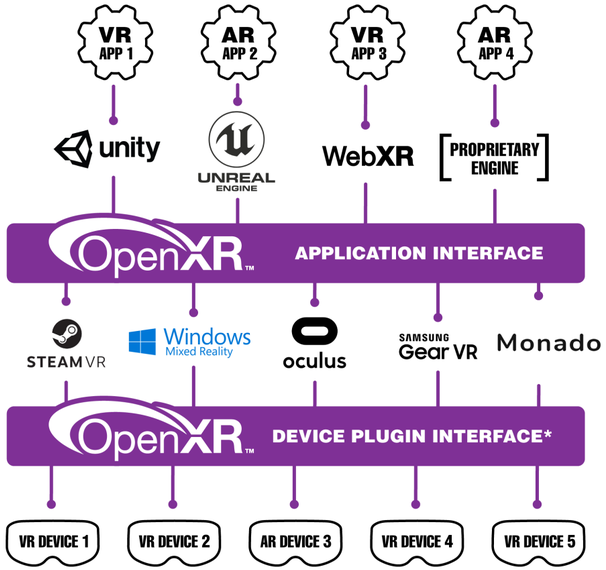

After OpenXR

Wide interoperability of XR apps and devices

Working Group Participants

Games Engines

Desktop/Console VR

AR

Mobile VR

GPUs

UI Devices

OpenXR is a collaborative design

Integrating many lessons from proprietary 'first-generation' XR API designs

* OpenXR 1.0 is focused on enabling cross-platform applications. Optional device plugin interface will be supported post V1.0

OpenXR 1.0 Availability

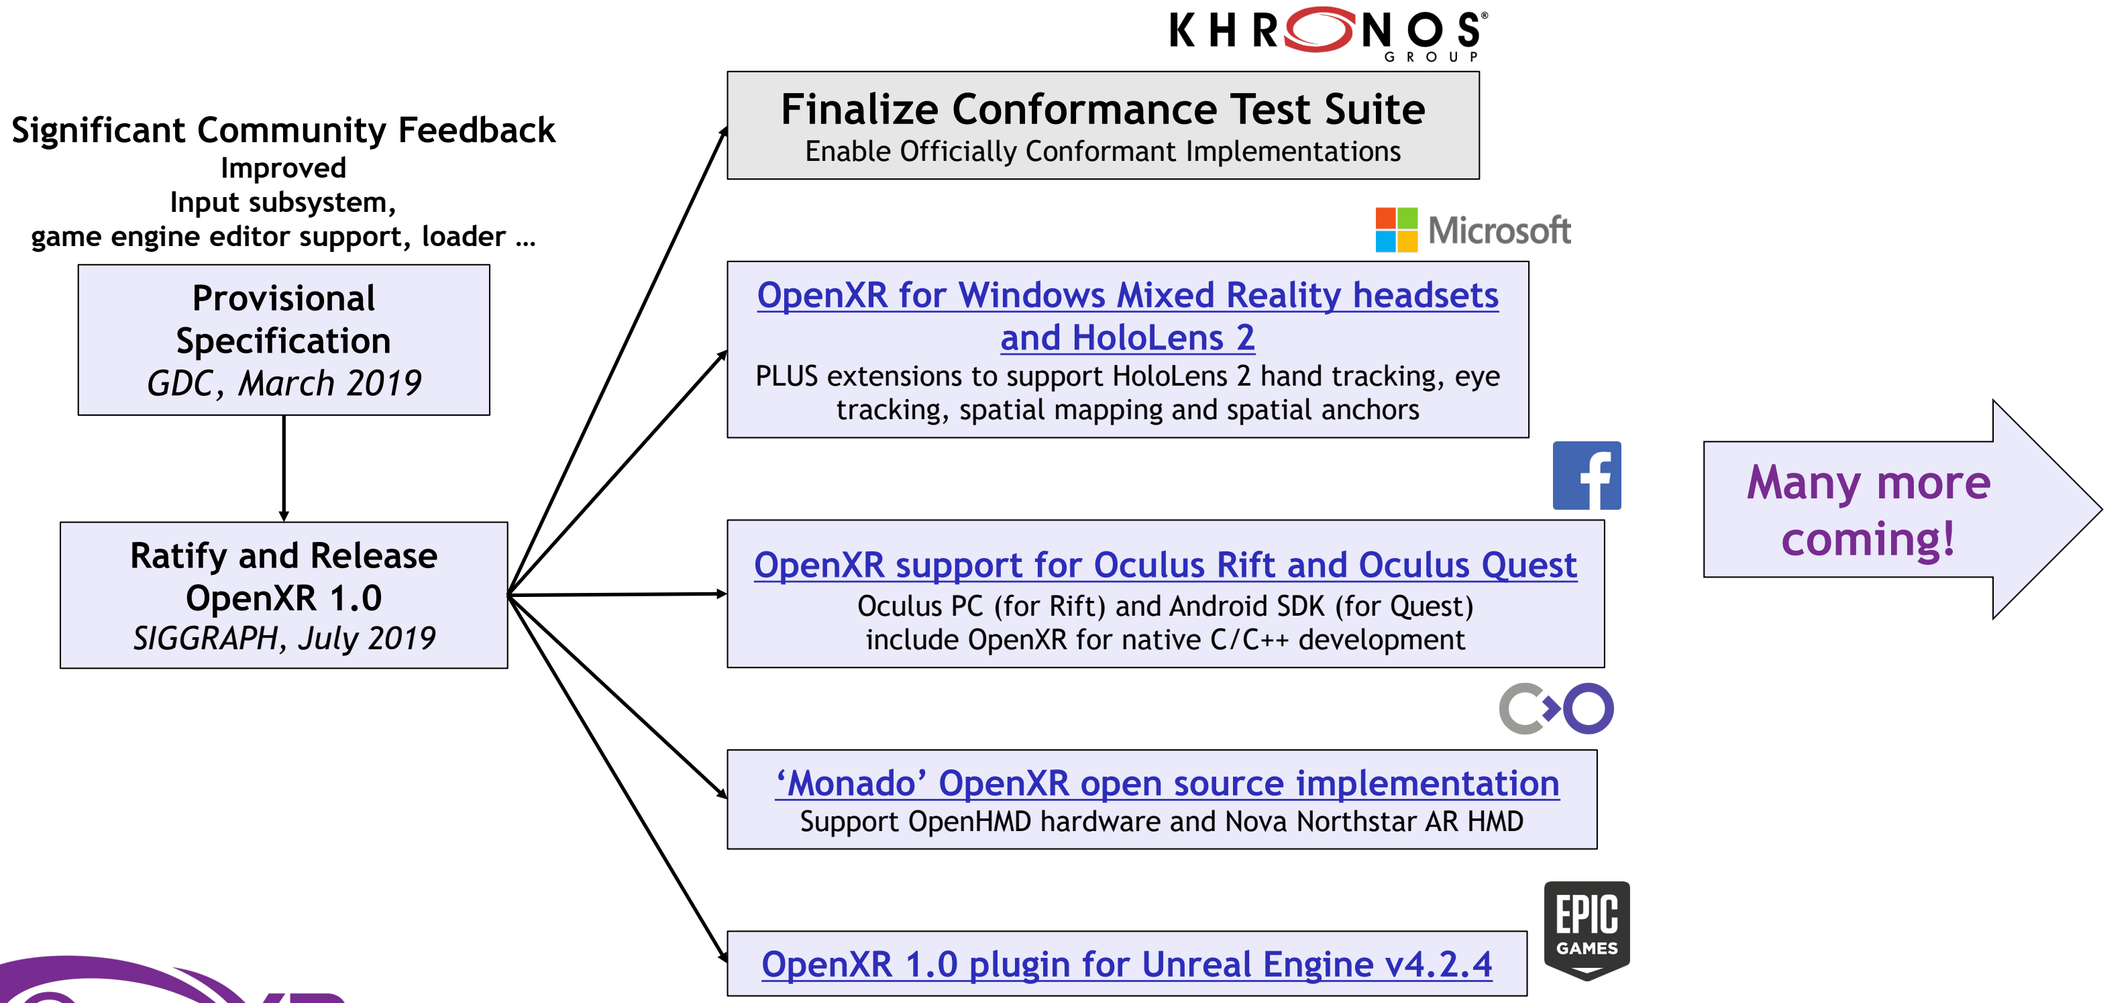

OpenXR is used with a 3D API

Application or Engine

High-performance, low-latency 3D rendering and composition*
- Multiview
- Context priority
- Front buffer rendering
- Tiled rendering (beam racing)
- Variable rate rendering

Display, composition and optical correction parameters

Cross-platform access to XR HMDs and sensors
- XR application lifecycle
- Input device discovery and events
- Sensor tracking and pose calculation
- Frame timing and display composition
- Haptics Control

Vulkan®

OpenXR™

* OpenXR can be used with other 3D APIs such as Direct3D, OpenGL and OpenGL ES

OPENXR AND EDGE SERVER APPS

Low latency
Sensor Data

5G

OpenXR APIs hide the 5G round trip from applications

Sensor handling

OpenXR

Display composition

Wireless mobile device with display and sensors

MEC (Multi-access Edge Computing) Server

1. Processes sensor data, including machine learning for environmental lighting, occlusion, scene semantics, object reconstruction and UI
2. Generates imagery from 3D models, including stereo, foveal rendering, ray-tracing, optics pre-distortion, varifocal processing

NVIDIA

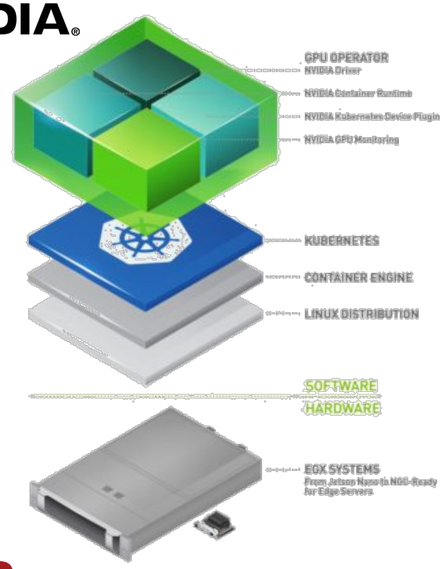

Vulkan

NVIDIA EGX

Location-aware Content Requests

Needed assets loaded to edge server

Apps and 3D Assets

glTF

5G

Generated Augmentations & Scenes

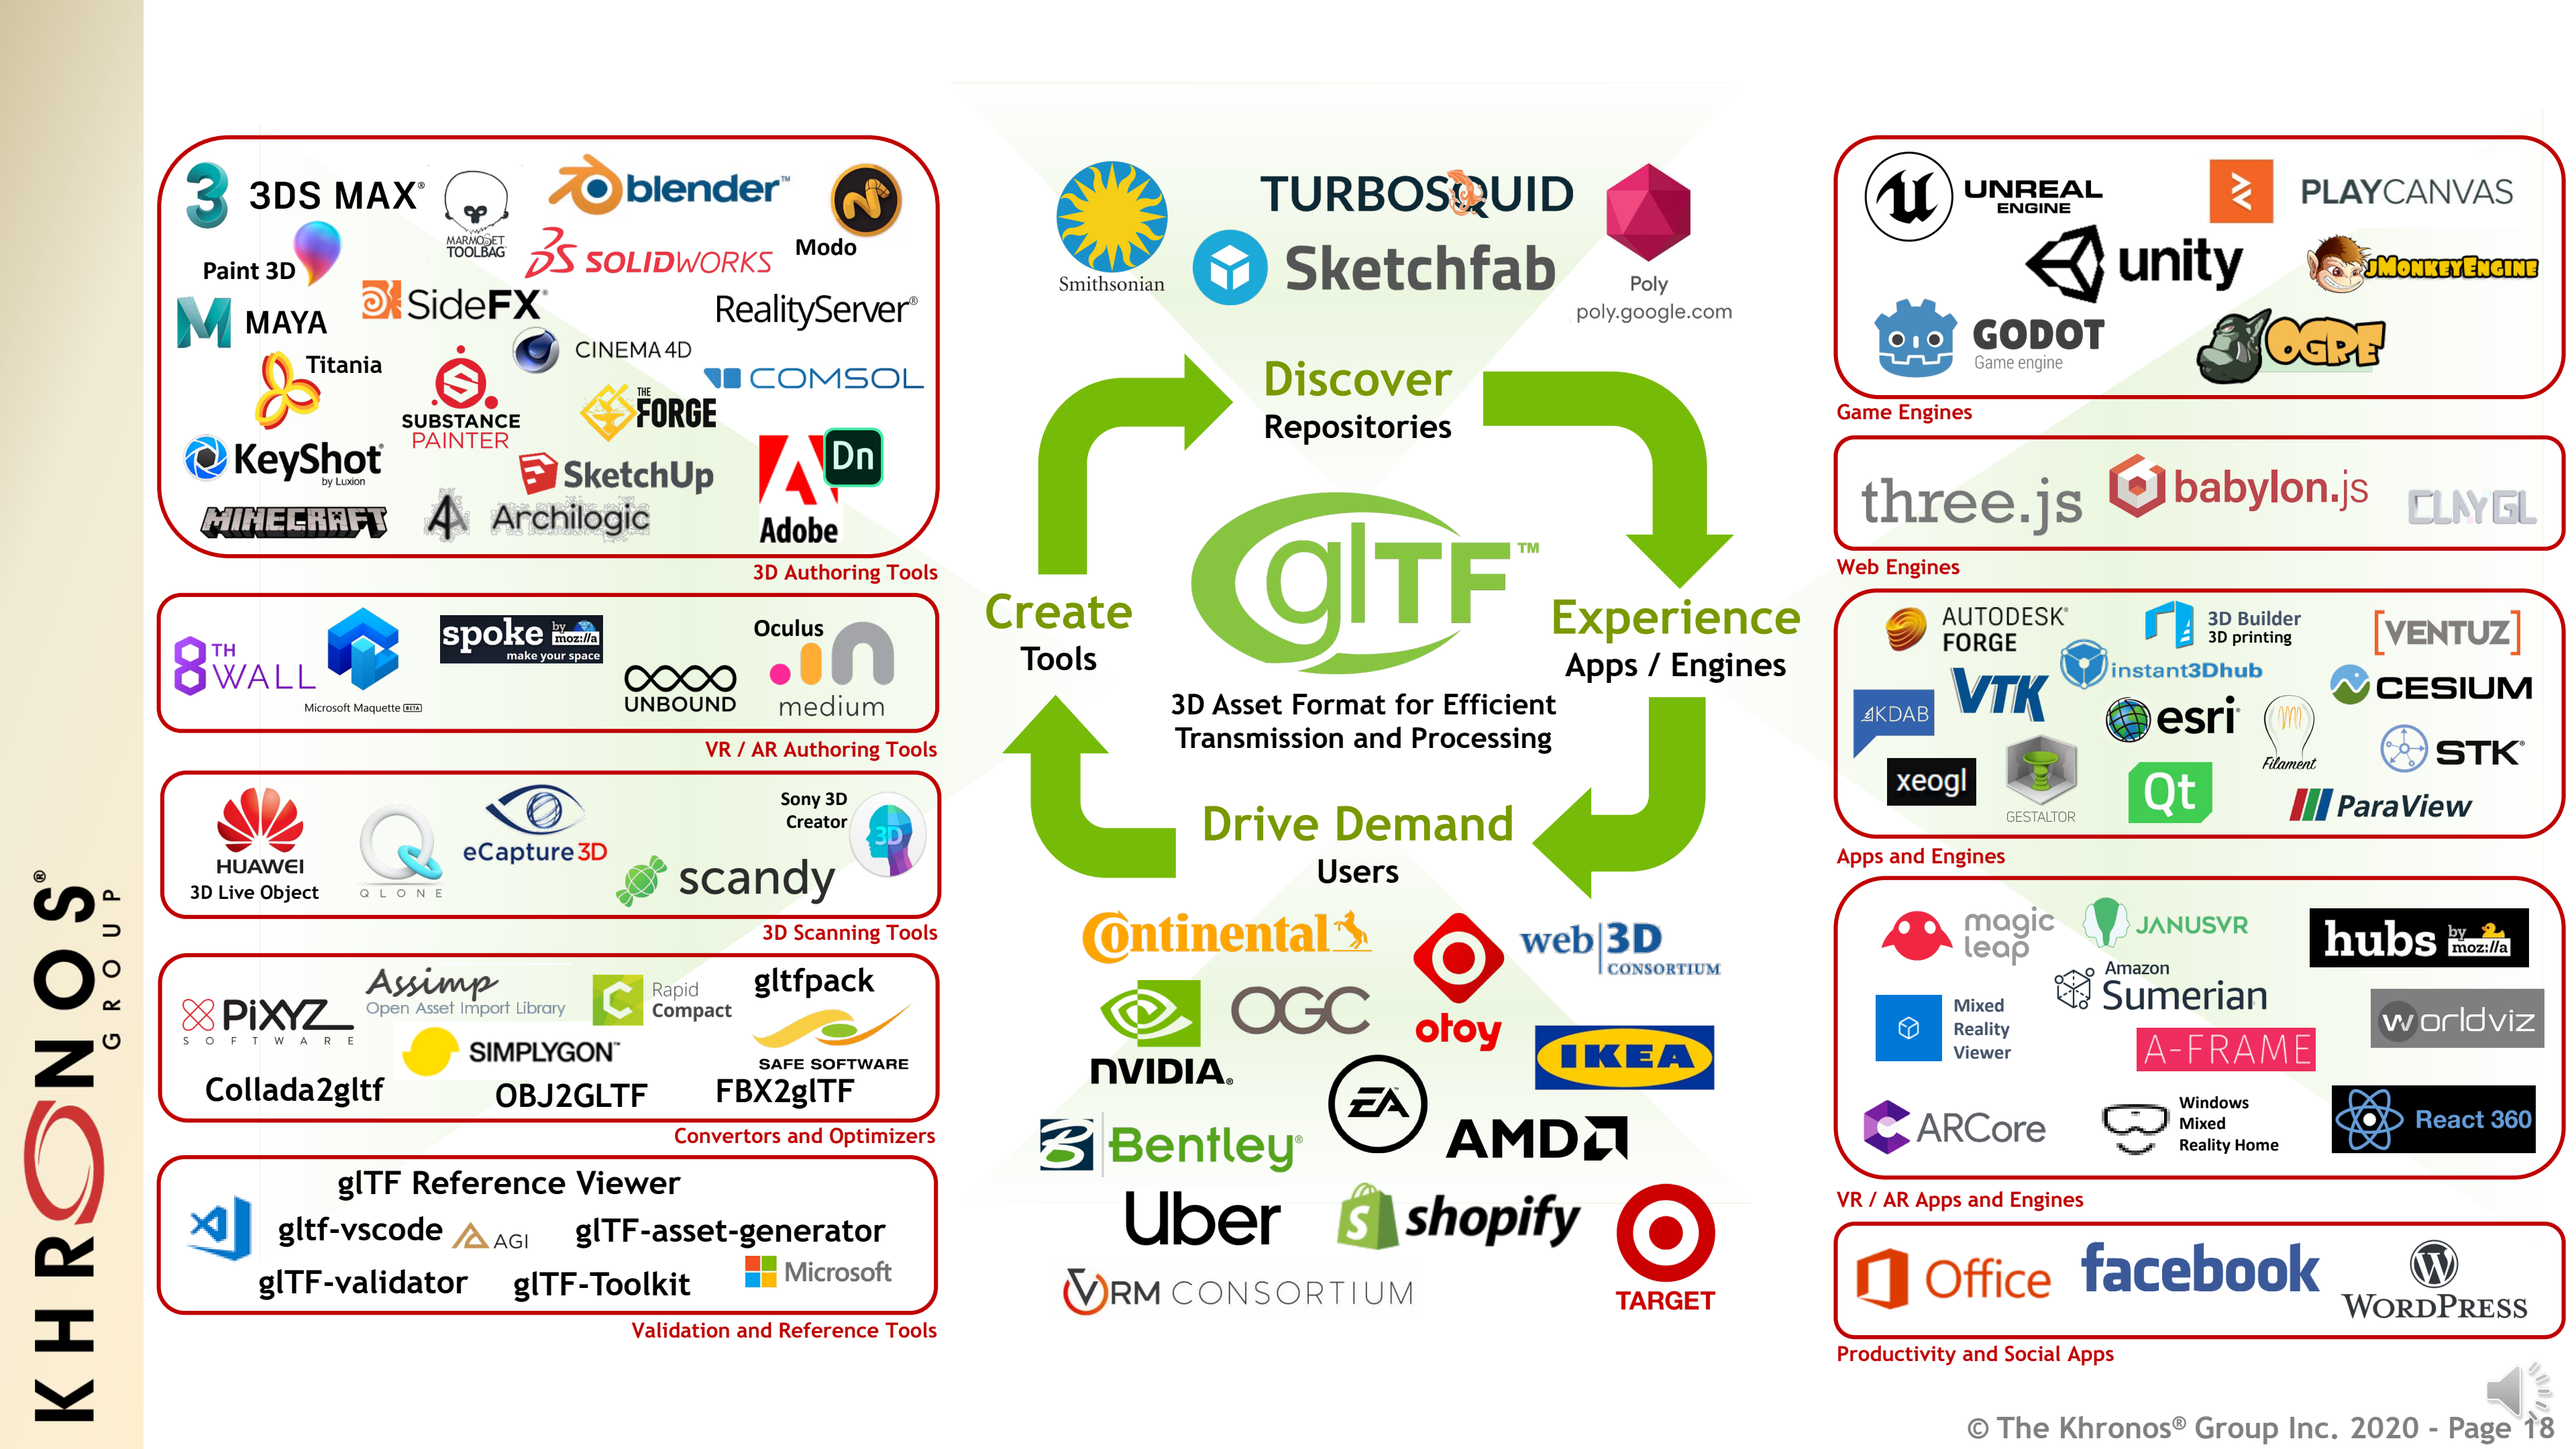
3DS MAX
blender
Modo
Paint 3D
MARMOSET TOOLBAG
SOLIDWORKS
MAYA
SideFX
RealityServer
CINEMA 4D
Titania
COMSOL
SUBSTANCE PAINTER
THE FORGE
KeyShot by Luxion
SketchUp
Dn
MINECRAFT
Archilogic
Adobe
3D Authoring Tools
8TH WALL
Microsoft Maquette BETA
spoke by moz://a make your space
UNBOUND
Oculus
medium
VR / AR Authoring Tools
HUAWEI 3D Live Object
QLONE
eCapture 3D
Sony 3D Creator
scandy
3D Scanning Tools
PiXYZ SOFTWARE
Assimp Open Asset Import Library
Rapid Compact
gltfpack
SIMPLYGON
SAFE SOFTWARE
Collada2gltf
OBJ2GLTF
FBX2glTF
Convertors and Optimizers
glTF Reference Viewer
gltf-vscode
AGI
glTF-asset-generator
glTF-validator
glTF-Toolkit
Microsoft
Validation and Reference Tools
Smithsonian
TURBOSQUID
Sketchfab
Poly
poly.google.com
Discover
Repositories
Create
Tools
glTF™
Experience
Apps / Engines
3D Asset Format for Efficient Transmission and Processing
Drive Demand
Users
Continental
web 3D CONSORTIUM
NVIDIA
OGC
otoy
IKEA
Bentley
AMD
Uber
shopify
VRM CONSORTIUM
TARGET
UNREAL ENGINE
PLAYCANVAS
unity
JMonkeyEngine
GODOT Game engine
OGPF
Game Engines
three.js
babylon.js
CLAYGL
Web Engines
AUTODESK FORGE
3D Builder 3D printing
VENTUZ
instant3Dhub
VTK
KDAB
esri
CESIUM
Filament
STK
xeogl
GESTALTOR
Qt
ParaView
Apps and Engines
magic leap
JANUSVR
hubs by moz://a
Amazon Sumerian
Mixed Reality Viewer
worldviz
A-FRAME
ARCore
Windows Mixed Reality Home
React 360
VR / AR Apps and Engines
Office
facebook
WordPress
Productivity and Social Apps

glTF Roadmap

FlightHelmet_baseColor
2048 x 2048, RGB

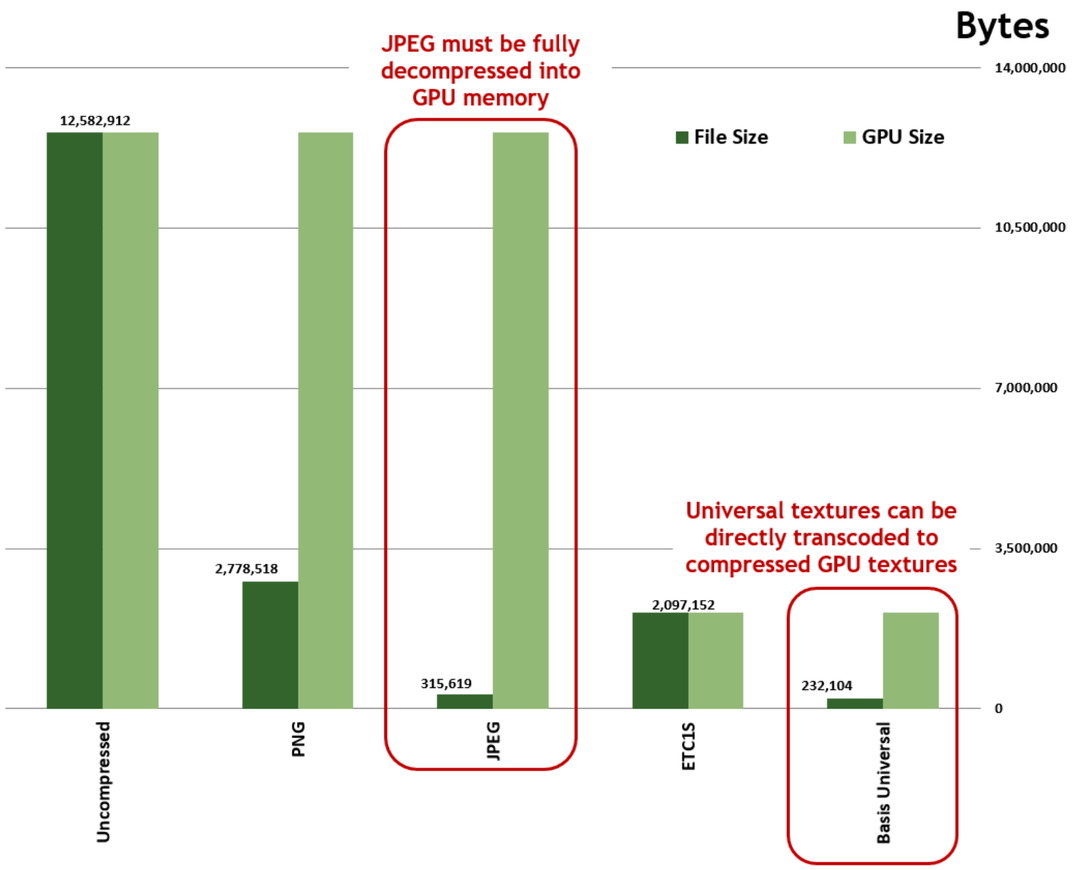

glTF Universal Textures

(imminent)

Basis Universal encoding/transcoding
KTX2 Container

Working Group is constantly balancing feature requests against the 'glTF Prime Directive' - remain a universal and easy to process delivery format

Second Generation Physically-Based Rendering (PBR)

Set of coherent extensions
Clear coat (imminent)
Absorption/attenuation
Subsurface scattering
Anisotropy

Inspiration from Dassault Systèmes Enterprise PBR Shading Model (DSPBR) and MDL

Wide industry cooperation

Seeking Requirements

Subdivision surfaces
Advanced Animation
LOD and Streaming
Compressed Point Clouds
Cross-asset linking
Enhanced Metadata
Composability
Instancing
CAD/BIM model support
Encryption and security
3D Printing

GPU COMPUTE APIS AT NVIDIA

Developer Choice

- Heterogeneous compute devices
- Cross-vendor open standard
- Simpler to program than rendering APIs

- CUDA-X libraries and tools
- Integrated HW/SW roadmap for rapid innovation
- BUT GPU Only

- Widely available across multiple platforms
- Integrated rendering, data movement and compute
- Explicit hardware control
- BUT GPU Only

Increasing layering, tooling and interop functionality enabling enhanced development and deployment flexibility

OpenCL is Widely Deployed and Used

Desktop Creative Apps

Parallel Languages

OpenACC DIRECTIVES FOR ACCELERATORS

SYCL

aparapi

PyOpenCL

Linear Algebra Libraries

SYCL-BLAS

ViennaCL

CLBlast

Machine Learning Libraries and Frameworks

Intel OPENVINO™

Xiaomi MACE Mobile AI Compute Engine

Intel clDNN

NNAPI

Caffe

SYCL-DNN

Acuity ML/AI Software Foundation VeriSilicon

Qualcomm Neural Processing SDK for AI

Synopsis MetaWare EV

TI DL Library (TIDL)

Arm Compute Library

Vision and Imaging Libraries

OpenCV

ImageMagick

VisionCpp

Halide

OpenVX

Machine Learning Compilers

plaidML

tvm

Glow

Math and Physics Libraries

BULLET PHYSICS LIBRARY

Wolfram Mathematica

GNU Octave

ArrayFire C, C++, Fortran

Matlab

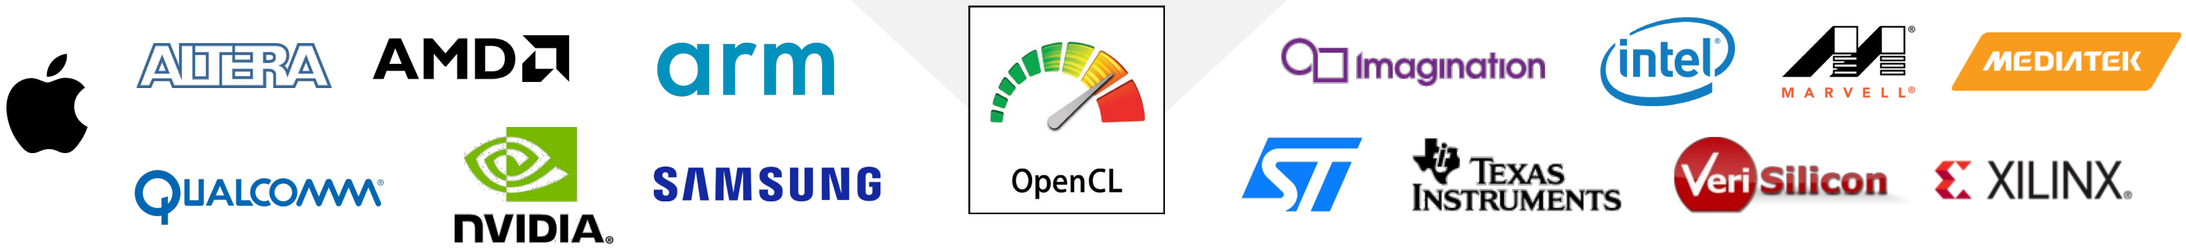

Hardware Implementations

KHRONOS GROUP

OpenCL Applications over Vulkan

- **Clspv - Google's open source OpenCL kernel to Vulkan SPIR-V compiler**
  - Tracks top-of-tree LLVM and clang, not a fork
- **Clvk - prototype open source OpenCL to Vulkan run-time API translator**
- **Used for shipping production apps and engines on Android**
  - Adobe Premiere Rush video editor - 200K lines of OpenCL C kernel code
  - Butterfly Network iQ Ultrasound on Android
  - Xiaomi MACE inferencing engine

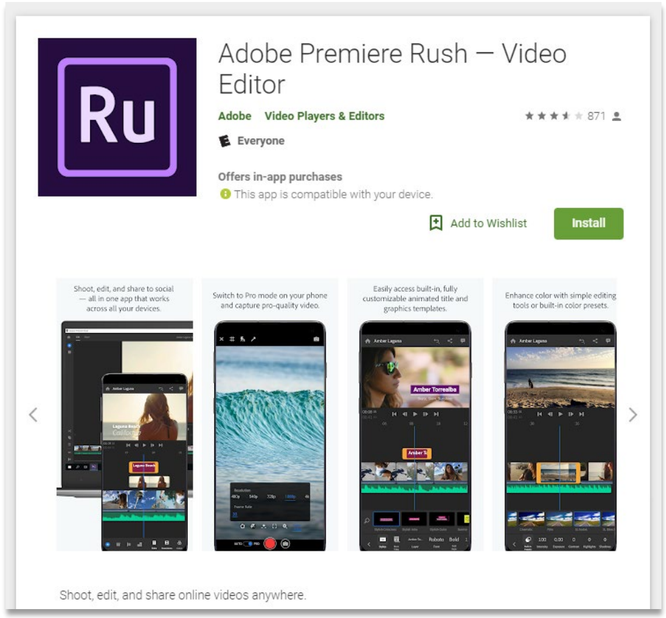

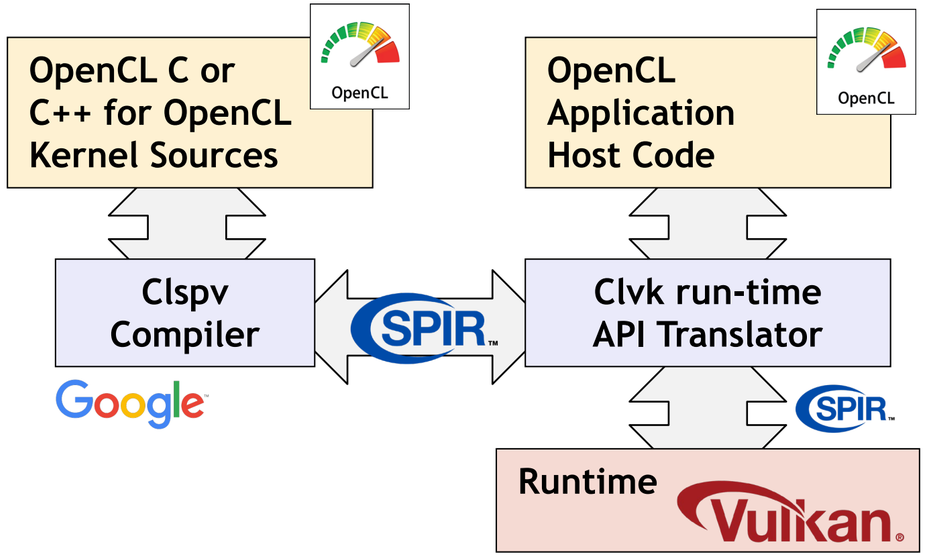

https://github.com/kpet/clvk
https://github.com/google/clspv

NVIDIA AND OPENCL

Active Investment and Support

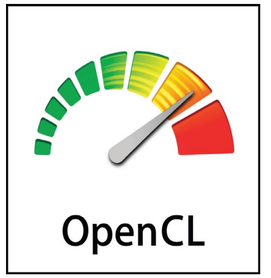

NVIDIA Chairs the Khronos OpenCL Working Group

Production-class OpenCL 1.2 on Linux and Windows

Active, ongoing improvements in performance and power efficiency for new architectures

Multi-GPU optimizations, multi-command-queue use case tuning

Optimized data transfers and GPU memory allocation

New Functionality

Half & Half2, arithmetic and conversions - in development

Expanding OpenCL Interop options

OpenGL and D3D9/10/11 – today. DX12 being planned

OpenCL/Vulkan interop – aiming for 2H20 - using standard Vulkan external memory interop hooks
NVIDIA driving Vulkan/OpenCL interop extension at Khronos

Providing guidance for deploying OpenCL apps over Vulkan

Leveraging the open source SPIR-V compiler

NVIDIA AND KHRONOS API STANDARDS

Cloud – Windows and Linux

Desktop – Kepler and later

Nintendo Switch

Shield Android TV

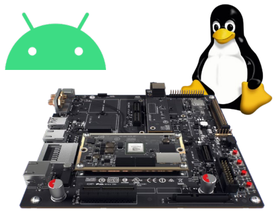

Jetson

NVIDIA has shipped Vulkan 1.2 and Ray Tracing drivers on day of specification release
NVIDIA chairing multiple Vulkan initiatives at Khronos: Ray Tracing, Machine-Learning, Vulkan Portability etc.
Increased Vulkan support in NSIGHT graphics development tools
https://developer.nvidia.com/Vulkan

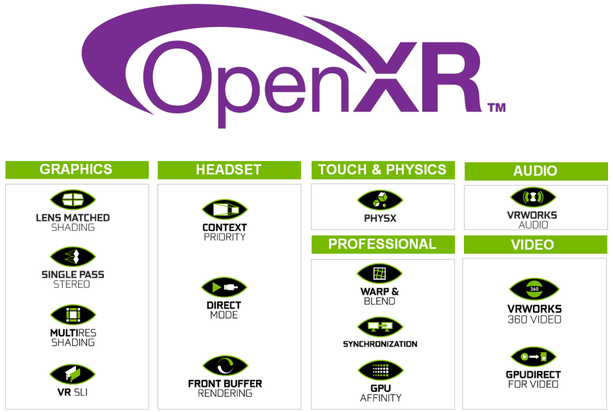

Actively working to help evolve OpenXR to complement VRWorks and enable EGX Edge Server AR
https://developer.nvidia.com/vrworks

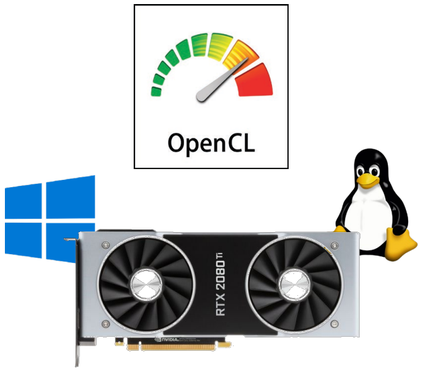

Production-class OpenCL 1.2 on Linux and Windows
Active improvements in efficiency and performance
Vulkan Interop in development
https://developer.nvidia.com/opencl

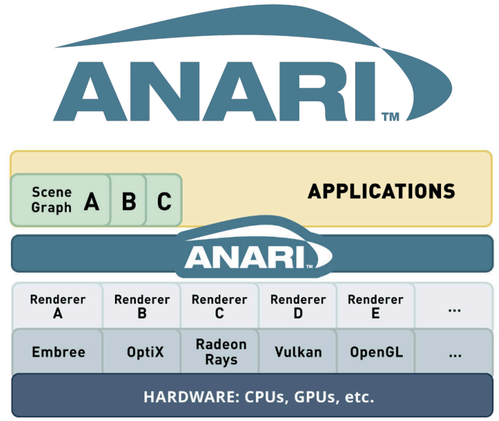

NVIDIA initiated industry cooperation
Portable Analytic Rendering API
NVIDIA chairing the working group
https://www.khronos.org/anari

How To Get Involved!

- **More information at other GTC recorded sessions**
  - [S21770] NVIDIA Vulkan Features Update – including Vulkan 1.2 and Ray Tracing
  - [S22694] Ray Traced Reflections in Wolfenstein Youngblood
- **Any company or organization is welcome to join Khronos!**
  - For a voice and a vote in any of these standards – membership starts at $3,500
- **OR request an invite to Vulkan, OpenCL, OpenXR Advisory Panels**
  - No fee, execute Khronos NDA and IP Framework, provide requirements and spec feedback
- **We welcome your feedback at NVIDIA or Khronos**
  - Khronos Forums: https://forums.khronos.org/
  - Khronos Slack Channels: https://khronosdevs.slack.com/messages
  - Khronos open source GitHub repositories: https://github.khronos.org/
- **Contact Neil Trevett**
  - ntrevett@nvidia.com | @neilt3d | www.khronos.org