KHRONOS GROUP

Khronos Fast Forward

Neil Trevett
Khronos President and VP Developer Ecosystems at NVIDIA

Speakers

Topic	Speaker	Length
Khronos, Kamaros, AI, ANARI, glTF, Web	Neil Trevett, Khronos	30 minutes
Vulkan	Ralph Potter, Samsung	5 minutes
Slang Shading Language	Shannon Woods, NVIDIA	5 minutes
OpenXR	Jian Zhang, PICO	5 minutes
Audience Q&A	All	15 minutes

Khronos Connects Software to Silicon

Founded in 2000
~ 150 Members | ~ 40% US, 30% Europe, 30% Asia

KHRONOS GROUP

Consortium creating open, royalty-free interoperability standards and open source to harness the power of GPU, XR and multiprocessor hardware

3D graphics, AI, augmented and virtual reality, parallel processing, and spatial computing

Non-profit, member-driven standards organization, open to any company

Proven multi-company governance and Intellectual Property Rights Framework

Khronos Active Standards

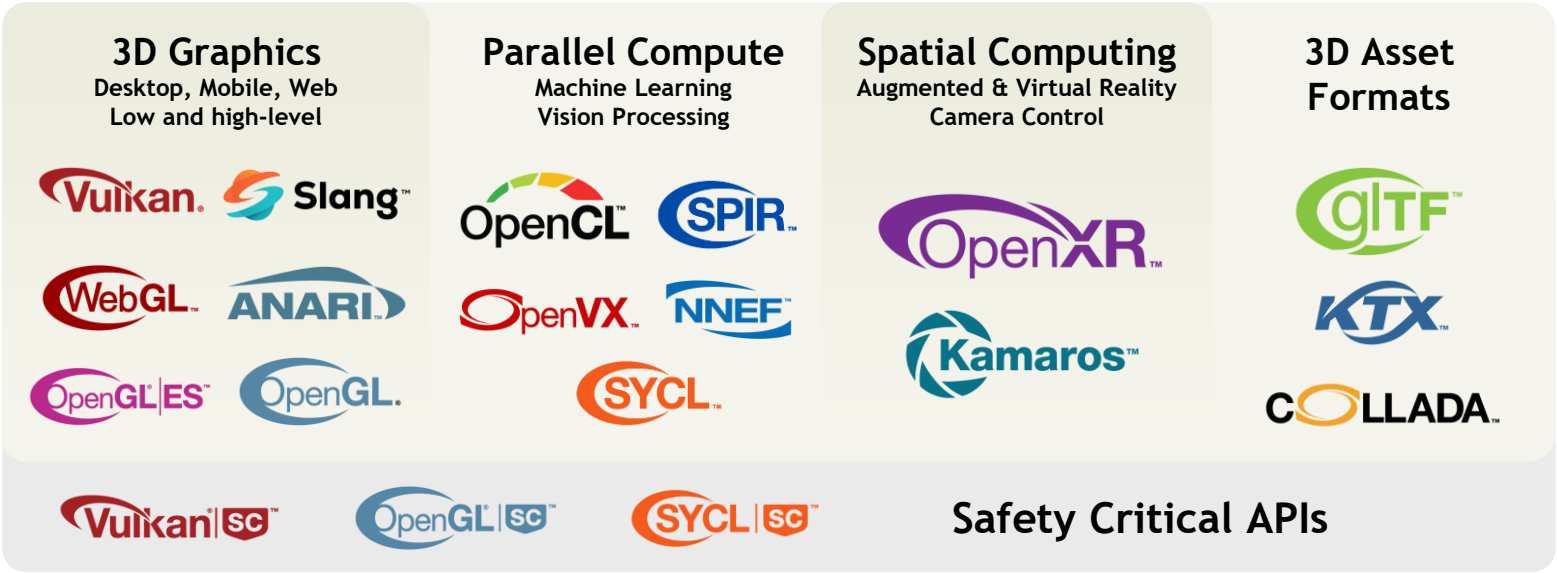

Khronos Standards for Spatial Computing

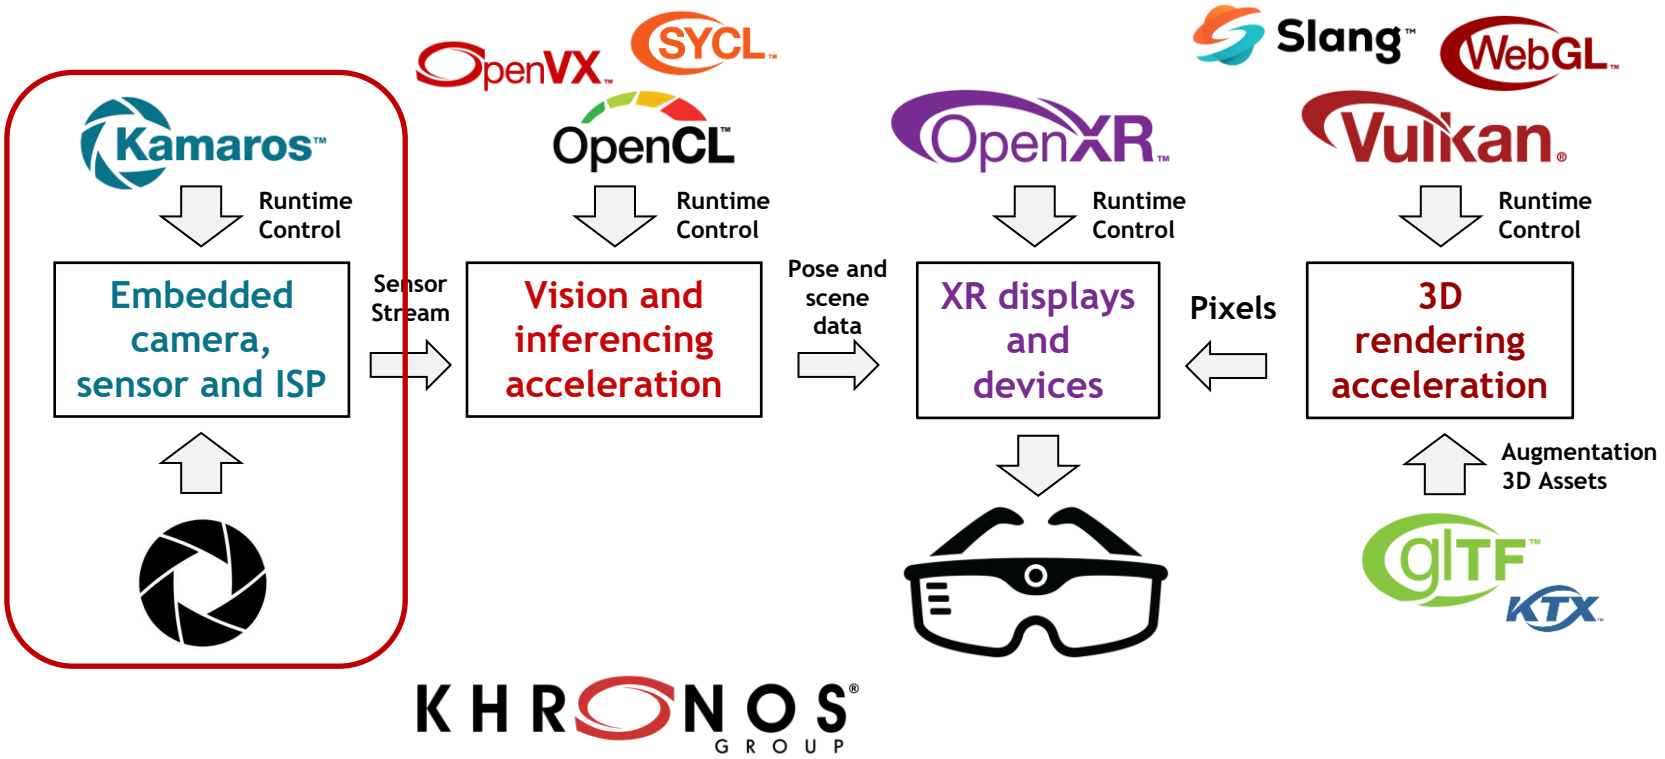

Needs for an Embedded Camera API Standard

Increasing Sensor Diversity
Including camera arrays and depth sensors such as Lidar

Multiple Sensors Per System
Synchronization and coordination become essential

Proprietary camera APIs often hide full camera capabilities to protect implementation IP

Cost and time to integrate and utilize sensors is major industry friction point

Tighter Accelerator Integration
Sensors feeding machine learning and traditional acceleration pipelines

Efficiency and Low Latency
Real-time interactive on power-constrained systems

Kamaros Scope

In Development!

Kamaros API provides controls for Camera Modules and close-to-sensor Image Signal Processing (ISP) hardware

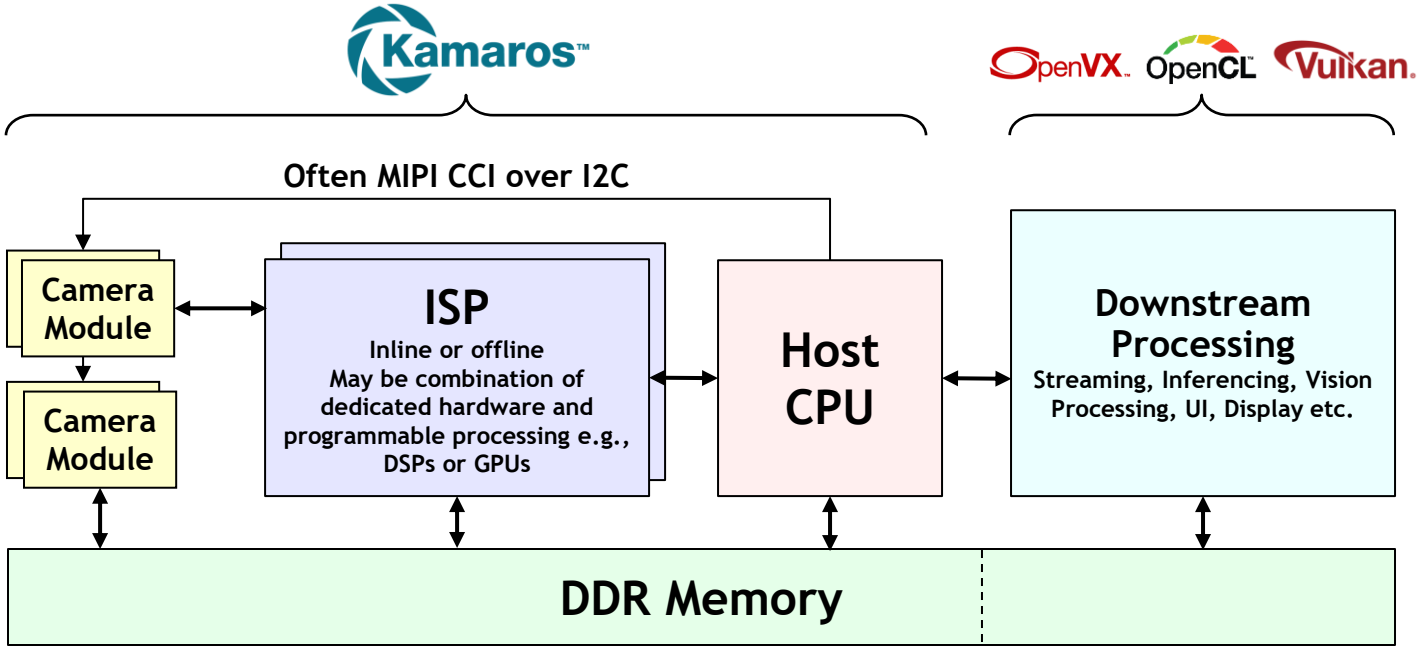

Khronos Standards for Spatial Computing

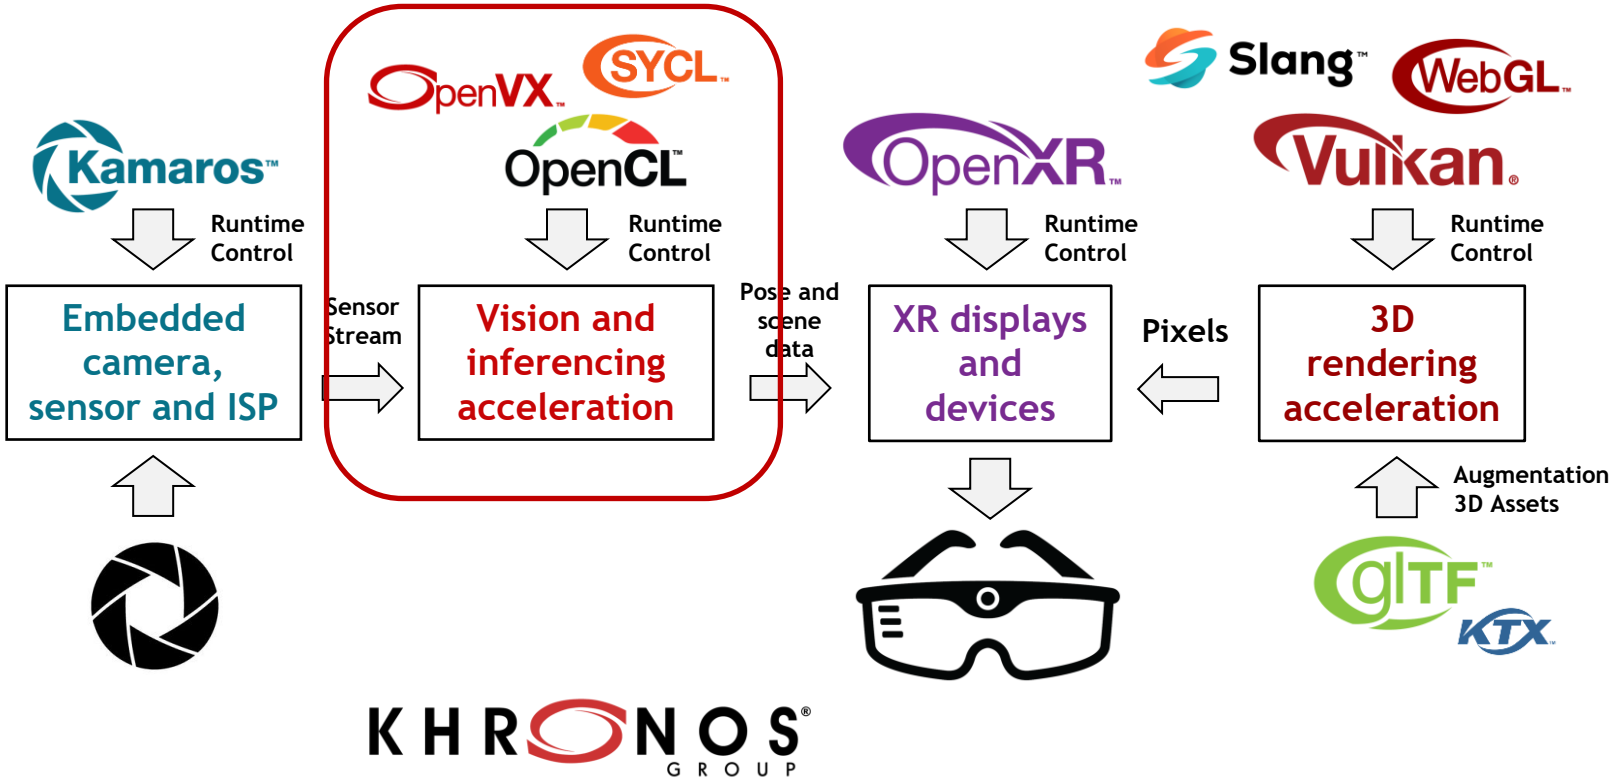

Machine Learning Acceleration APIs

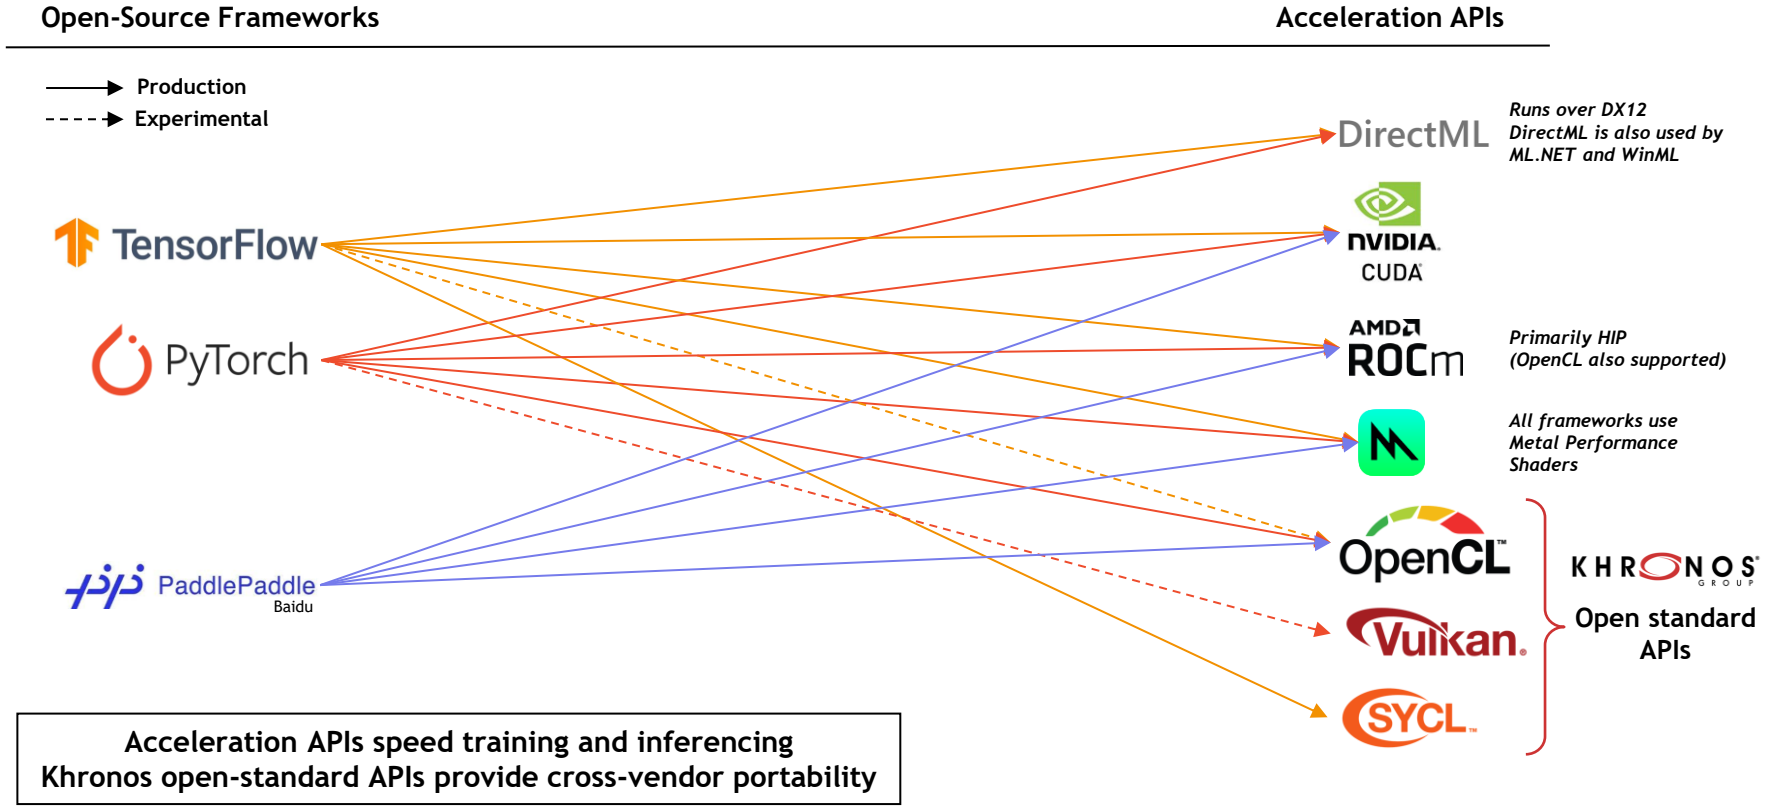

Acceleration APIs speed training and inferencing
Khronos open-standard APIs provide cross-vendor portability

Machine Learning Acceleration APIs

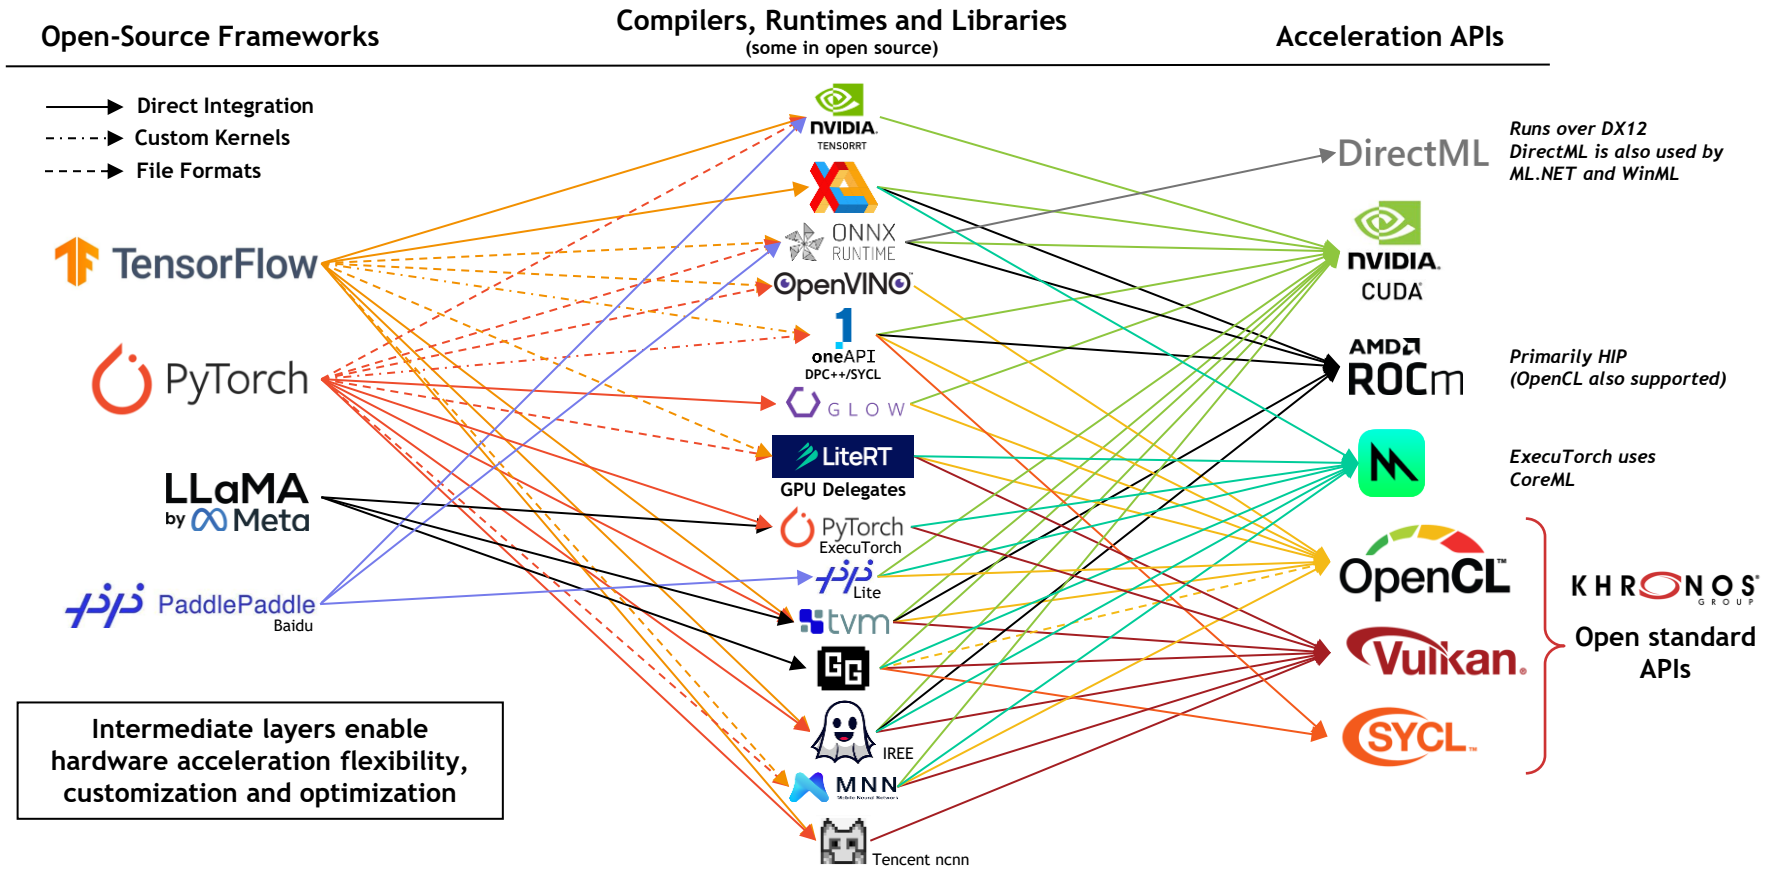

Machine Learning Acceleration APIs

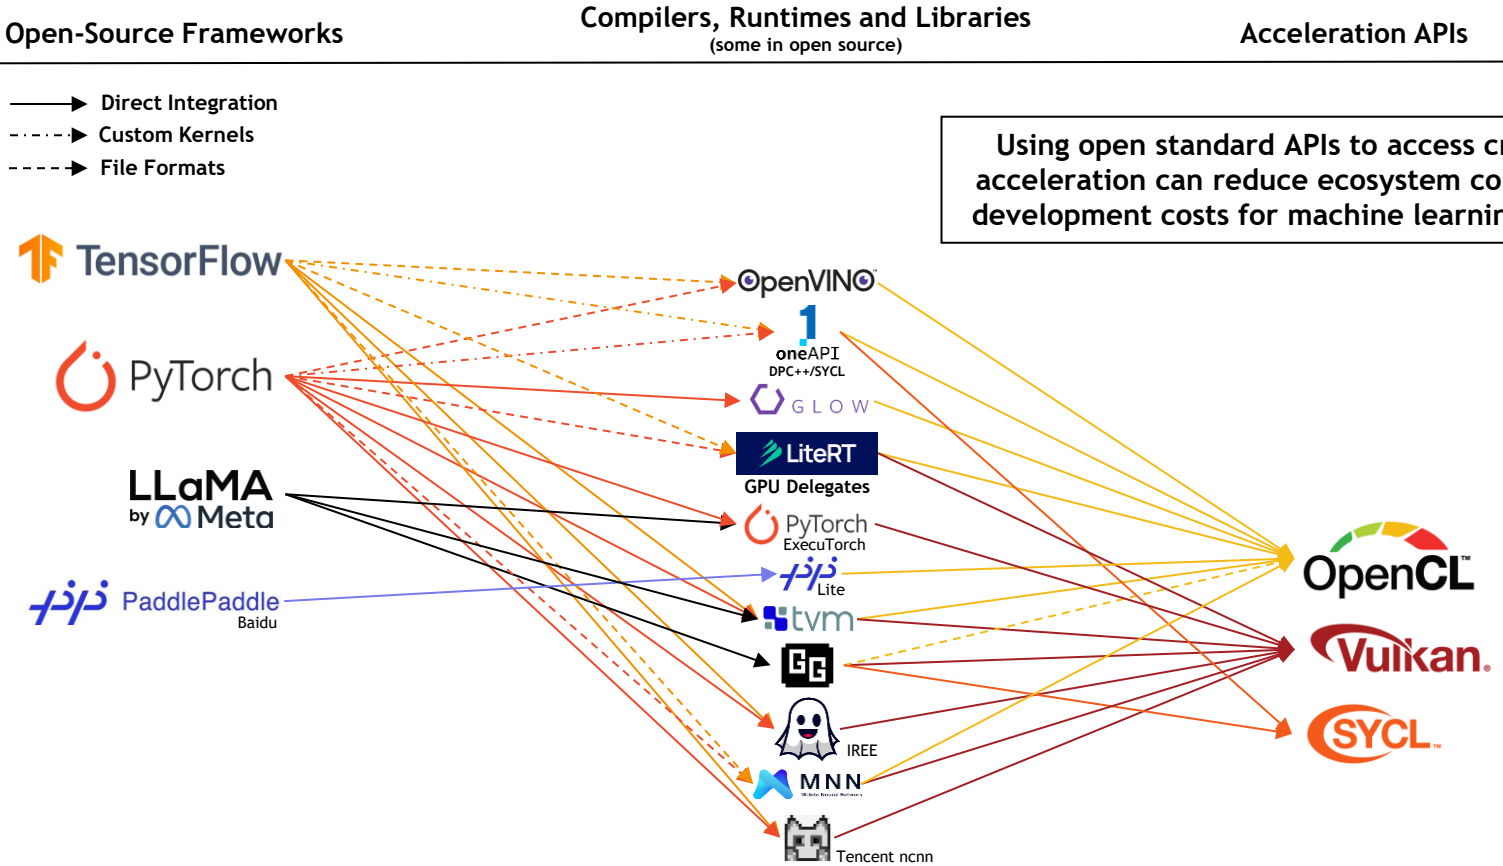

Embedded Machine Learning

Name	Type	Authoring Frameworks	Acceleration APIs
Cadence Xtensa Neural Network Compiler (XNNC)	Compiler	TensorFlow, PyTorch, ONNX	OpenCL, OpenVX
CEVA Deep Neural Network compiler (CDNN)	Compiler	TensorFlow, PyTorch, Caffe, ONNX	OpenCL, OpenVX
Synopsis MetaWare EV	Runtime	TensorFlow, PyTorch, Caffe, ONNX	OpenCL, OpenVX
Texas Instruments DL Library (TIDL)	Runtime	TensorFlow, PyTorch, Caffe, ONNX	OpenCL, OpenVX
VeriSilicon Acuity	Runtime		OpenCL, OpenVX
Xiaomi Mace	Runtime	TensorFlow, Caffe, ONNX	OpenCL, Vulkan
Xilinx Vitis AI	Runtime	TensorFlow, PyTorch	OpenCL, Native

OpenCL and OpenVX are the open standard APIs of choice for inferencing acceleration in embedded (and often mobile) devices

Accelerated ML Industry Discussions

- Do intermediate runtimes or compilers deliver best performance?
- Do high-level or low-level acceleration APIs deliver the best performance?
- What functionality should APIs provide for effective tensor acceleration?
- What is the most effective way to balance inferencing and other loads on a GPU?
- How can APIs provide acceleration across diverse hardware - such as GPUs and NPUs?
- Should the industry agree on a standard tensor operator set such as Arm's TOSA?

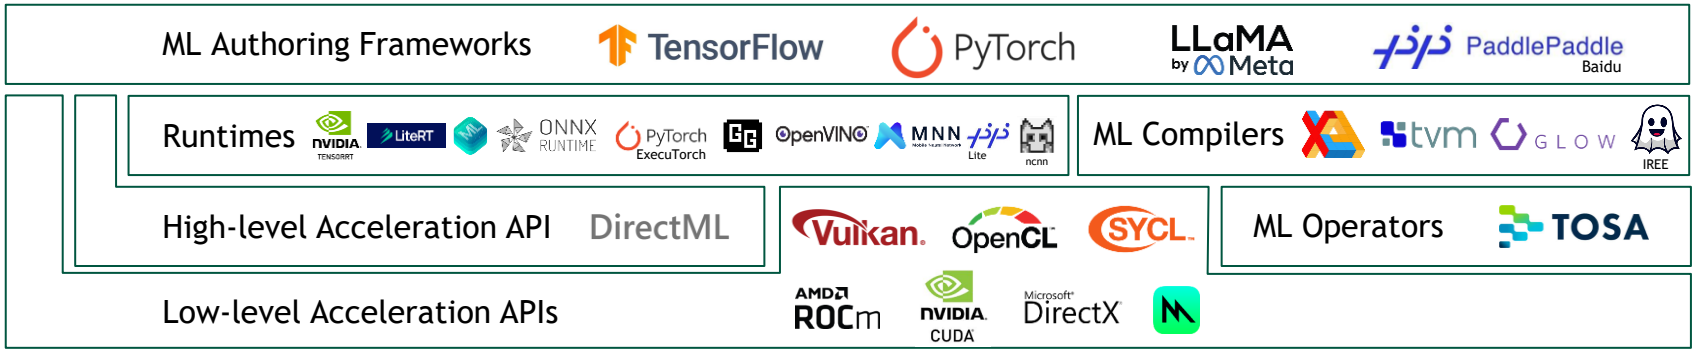

Native machine learning stack

Similar discussions happening in the JavaScript stack for machine learning in the Web

ANARI

Rendering Engine API

Neil Trevett, NVIDIA

Scientific Visualization Before and After ANARI

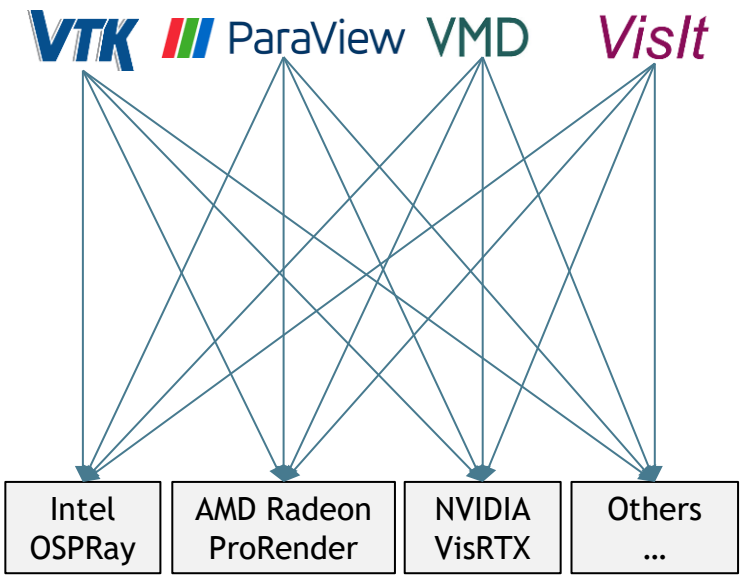

Before ANARI

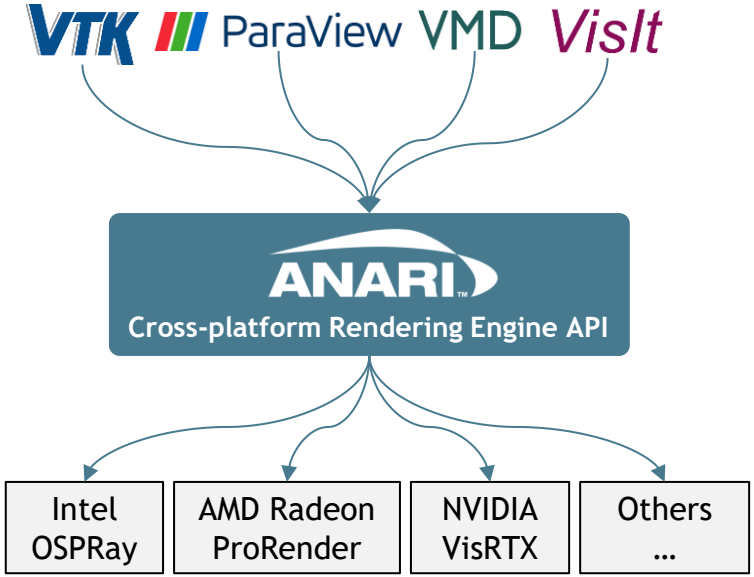

After ANARI

ANARI applications are portable to any engine supporting the ANARI API
Independently of vendor, platform or ecosystem

ANARI Scene Representation

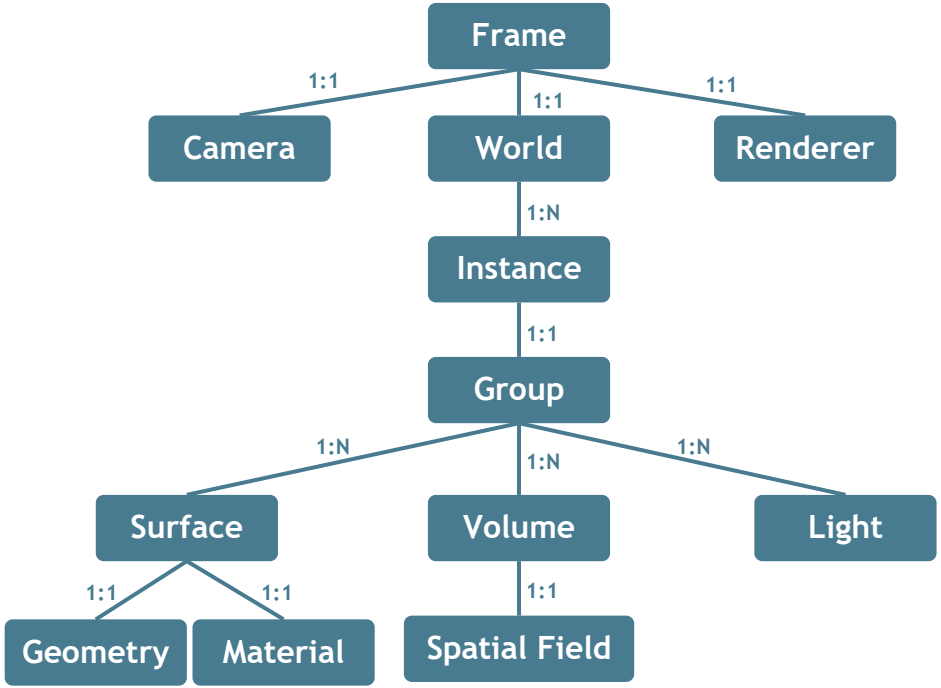

ANARI API used to build in-memory scene representation

Hierarchical object tree that expresses the complete scene for a single frame

Sections of the tree can be re-used to optimize resource utilization

Scene representation can be used to drive any rendering backend - rasterization techniques are NOT prescribed

ANARI only defines "what" and "when" not "how"

ANARI - Simplified and Portable Rendering

ANARI 1.0 Specification shipped in 2023

ANARI can be used by any type of application for portable access to diverse rendering engines including those using ray tracing and global illumination

Shipping implementations

AMD ProRender
Intel OSPRay
NVIDIA Omniverse
NVIDIA VisRTX
Helide open-source sample implementation

New implementations coming online

Visionaray (CPU/CUDA
Barney (OptiX + MPI)
Cycles (prototype)
...

ANARI bindings for C99, C++, Python ...

Engines use in-memory scene representation to drive rendering operations

Low-level APIs provide explicit control over hardware resources and operations

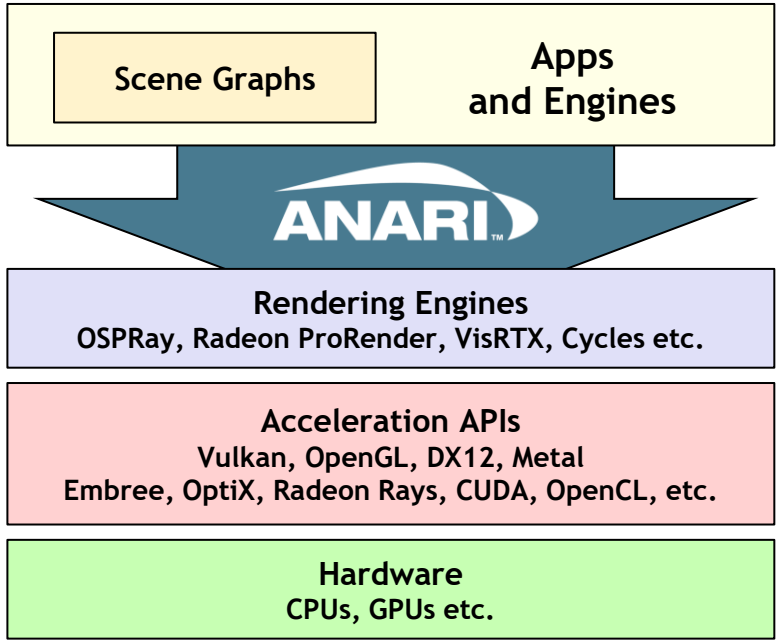

ANARI Virtual Hackathon 2024

- **Developers from around the globe gathered for three days**
  - To push the boundaries of 3D rendering with ANARI
- **Results included:**
  - Overhaul to VTK-ANARI integration
  - New VTK-m implementation
  - Volume rendering with the Blender Cycles renderer

For presentations, results and news of future events go to [ANARI Virtual Hackathon 2024]

ANARI Updates

- **Adopters Program Released**
  - Improved Conformance Test Suite (with more on the way!)
- **Many SDK improvements and additions**
  - Initial version of an OpenUSD Hydra plugin ‘hdAnari’ now available
  - New Blender add-on - contributions welcome!
  - Application debug layer can be enabled without code changes
  - Helium now provides a generic host-side array implementation

ANARI SDK is in open source on GitHub

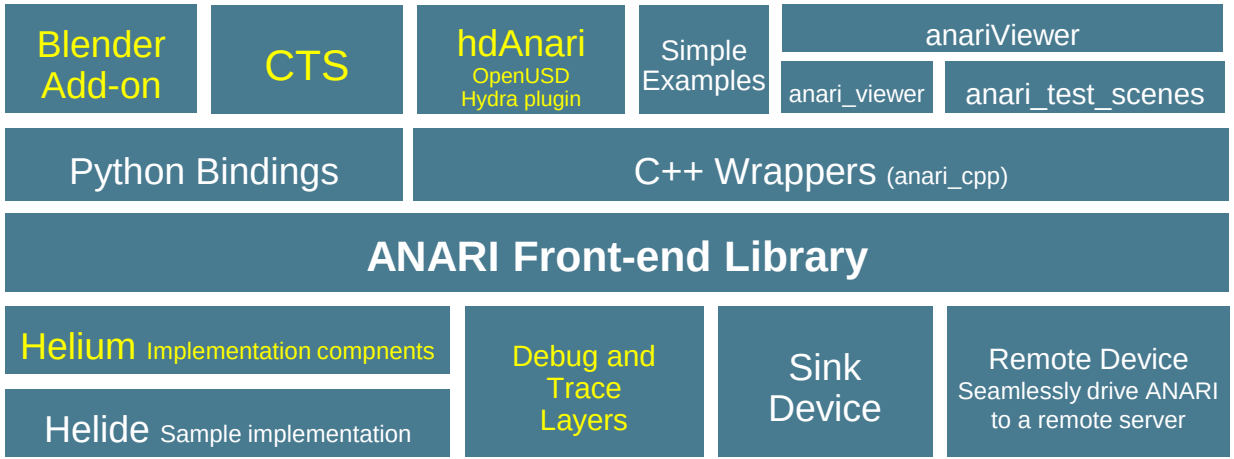

glTF & 3D Commerce

Neil Trevett

3D Authoring Tools

VR / AR Authoring Tools

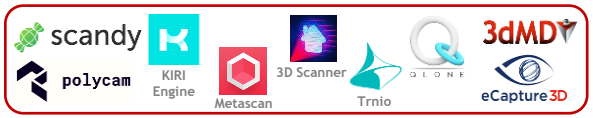

3D Scanning Tools

Convertors, Optimizers and Loaders

Validation and Reference Tools

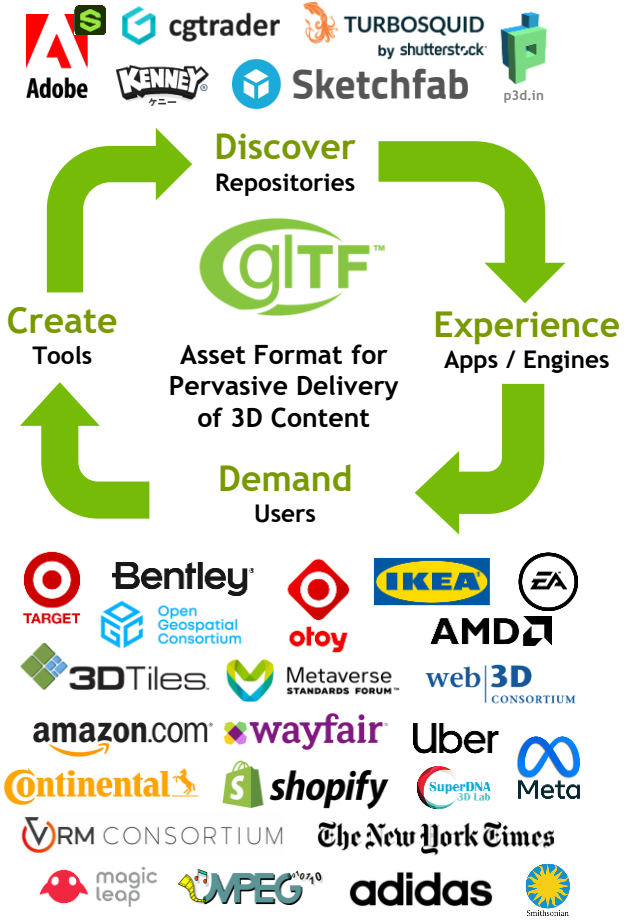

Game Engines

Web Engines

Apps and Engines

VR / AR Apps and Engines

Productivity and Social Apps

glTF PBR Materials Roadmap

Incremental consolidation and meticulous specification of *proven and accepted industry practice*

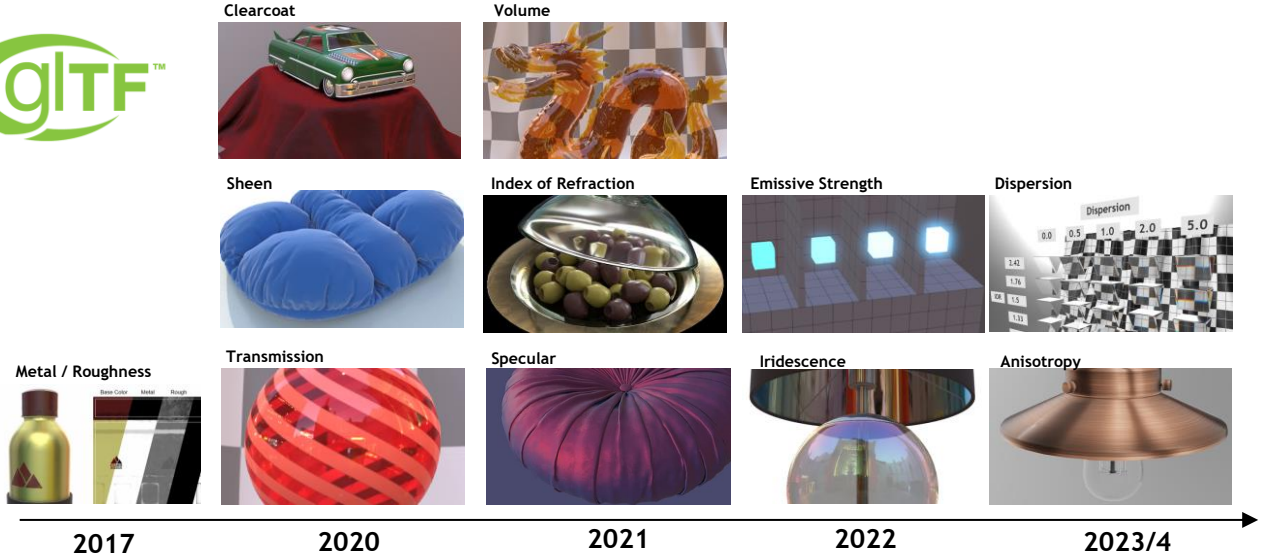

Subsurface
In development

MaterialX
1) Export glTF PBR as MaterialX node graph
2) Use MaterialX to drive procedural texture inputs into glTF PBR

OpenPBR
Working to align glTF and OpenPBR

Industry Collaboration

glTF Spatial Computing Roadmap

Interactivity

Node-based graph handling of user actions or events
Distillation of widespread accepted practice
Flexible computed scene state updates and animations

Compose complex scenes from referenced glTF assets
Efficiency and flexibility in transmission/delivery use cases
Placement, configuration, cache reuse, personalization, deferred loading, LODs, mesh variants etc.

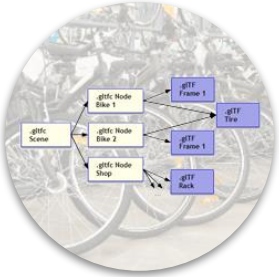

Complex Scenes

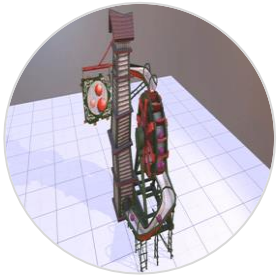

Physics

Describes physical properties of assets
Distillation of widespread accepted practice
Rigid Bodies: motions, collisions, materials , joints filters

Triggered and controlled from interactivity node graph
3D spatialized audio with 6DoF source/listener capabilities,
Play, stop, pause, loop, and speed controls
Splitting, merging, up/down-mixing, reverb, filtering

Audio

Market Segment Industry Collaboration

Khronos welcomes working collaboratively to leverage glTF extensibility

Market-specific extensions and use of glTF controlled by partner standards organization
Accelerates development of market segment functionality
Avoid needless duplication and fragmentation

VRM	Open Geospatial Consortium 3DTiles	MPEG	X3D	PDF	Additional Market Segments
Avatar Format *.vrm extension*	*.b3dm and .i3dm extensions*	ISO 23090-14:2023 MPEG-I *.mp4, miv, ivr*	ISO 19775-1:2023 X3D4 *.x3d extension*	ISO/TS 32007 glTF in PDF *.PDF extension*	

glTF™

Foundation of
Core specification and
glTF working group extensions

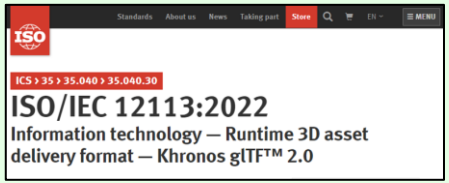

KHRONOS GROUP

Khronos 3D Commerce

Making 3D Pervasive - in the Real World

Build Once, Use Everywhere

Developing tools and techniques for 3D assets to be reliably and consistently used and displayed across diverse platforms and engines

Multiple Projects Underway

Render Showcase - evolve and expand Render Fidelity Site

Tone Mapping (PBR Neutral), exposure and lighting

Apparel: Skeletal & Facial Anchoring, Virtual Try-On, Stitching / detailing, Simulation

Content Developers test and validate assets

The equivalent of an API Conformance Test Suite (Vulkan has >3M tests)

Public comparison of multiple engines
1) Early warning of inconsistencies
2) Incentive to implement functionality

Sample Viewer
Asset Display Exemplar

Validator
Asset Correctness

Asset Library
Showcase & Test Assets

Asset Generator
Generate unit tests

Render Showcase
Web-based display comparisons

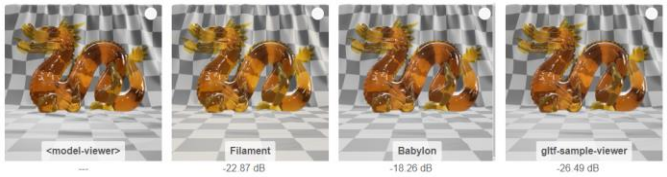

Khronos PBR Neutral Tone Mapper

- **True-to-Life Color Rendering of 3D Products**
  - Released in May 2024
  - Specification and sample implementation
- **1:1 match for colors up to a certain maximum value**
  - The remainder of color space used as headroom for compressed highlights
- **Wide adoption and support by 3D tools and engines**
  - <model-viewer>, Autodesk, Babylon.js, Blender, Dassault, Filament
  - London Dynamics, Phasmatic, Three.js, and ThreeKit

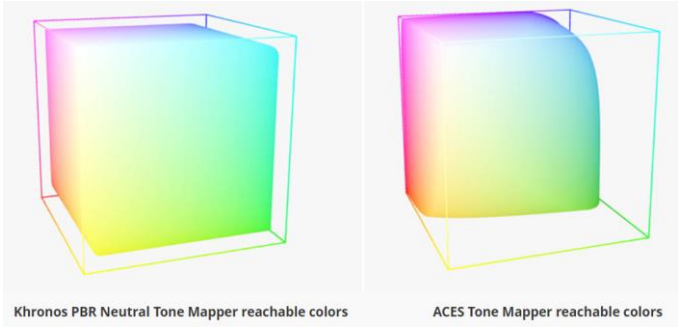

glTF Roadmap Overview

	Baseline Today	Short Term **Roadmap (23-24)**	Roadmap **Discussions**
Geometry	Draco Mesh Compression meshopt Compression		Quads, SubDiv Vector Displacement Implicit spheres and strands? Nerfs, Gaussian Splats
External References		Placement, Configuration, Cache Reuse, Personalization, Deferred Loading, LODs, Mesh Variants	
Textures and Materials	KTX 2.0 textures w Basis Universal Material Variants PBR Core + PBR Extension Wave 1-4	PBR Wave 5: Subsurface Scattering MaterialX Node graph update HDR Universal Textures Video Textures, Procedural Textures	PBR: Diffuse Transmission, Material X Procedural Textures
Animations	Keyframe/Skinned	Blender-compatible animation	Multi-track animation/blending Skeletons, Rigs and Anchors Animation Compression
Lights	Punctual Point, spot, and directional		IES, Rectangular Area Dome/Image
Interactivity		Node-based Behavior Graph	
Physics		Collisions, Rigid Bodies, Joints	Deformable Bodies
Audio		Playback (e.g., play, stop, loop), Spatial audio, Signal processing (gain, delay, pitch, reverb, filtering), multiple channels with splitting, merging Animation control and dynamic update of node properties	

glTF/USD 3D Asset Interoperability

The Metaverse Standards Forum

A neutral venue for pre- and post-standardization cooperation

Enable standards organizations to leverage each other's work and minimize overlaps/fragmentation

The Forum's glTF/USD 3D Asset Interoperability Working Group is enabling communication and cooperation between the glTF and USD communities

KHRONOS GROUP

Asset format to enable 3D content to be pervasively delivered and displayed on a wide diversity of native and web viewers, applications and engines

Metaverse STANDARDS FORUM™

Cooperation between glTF and USD ecosystems is a significant industry benefit

Alignment over requirements and roadmaps

Pragmatic projects to address immediate, real-world interoperability pain points

Extensible framework and ecosystem for describing, composing, simulating, and collaboratively navigating and constructing 3D scenes

KHRONOS GROUP

glTF BOF Session!

SIGGRAPH Asia Khronos BOFs

glTF: Transforming 3D Content Delivery for Real-Time Graphics

Friday 6th

11:00-12:00 noon, G408

Topics

glTF update and roadmap

3D Commerce Working Group

VRM Avatar Format and Khronos cooperation

Metaverse Standards Forum and OpenUSD cooperation

Metaverse STANDARDS FORUM™

KHRONOS GROUP

Forging the Immersive Web BOF

Neil Trevett

WebGL Update

- **Khronos is fully supporting development of WebGPU at W3C**
  - Working for a smooth transition for developers between WebGL and WebGPU
  - WebGPU brings GPU Compute to the Web using Vulkan/DX12/Metal backends
- **WebGL is pervasive and will be used by many applications for many years**
  - Khronos is evolving the WebGL specification and supporting multiple implementations
  - New extensions: Pixel Local Storage and more OpenGL ES functionality

WebGL 2.0 is available on >96% of browsers

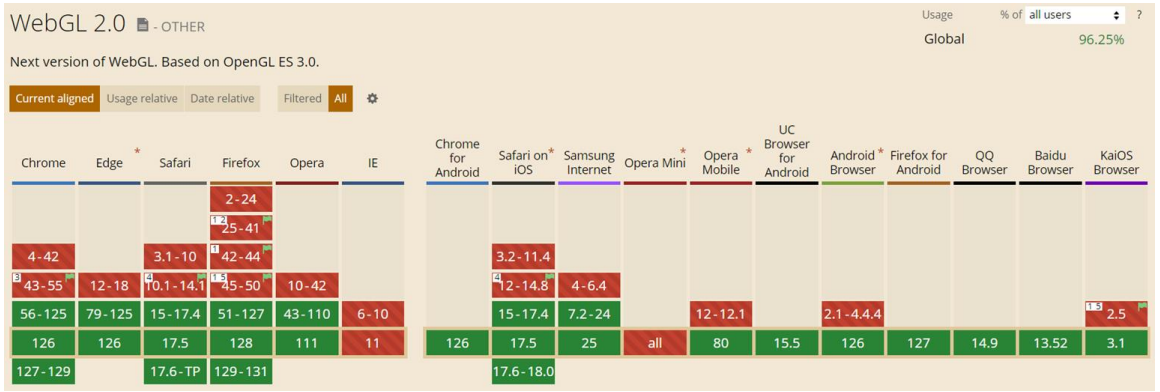

Khronos and W3C - XR Cooperation

XR Applications and Engines have access to native and JavaScript APIs with aligned functionality

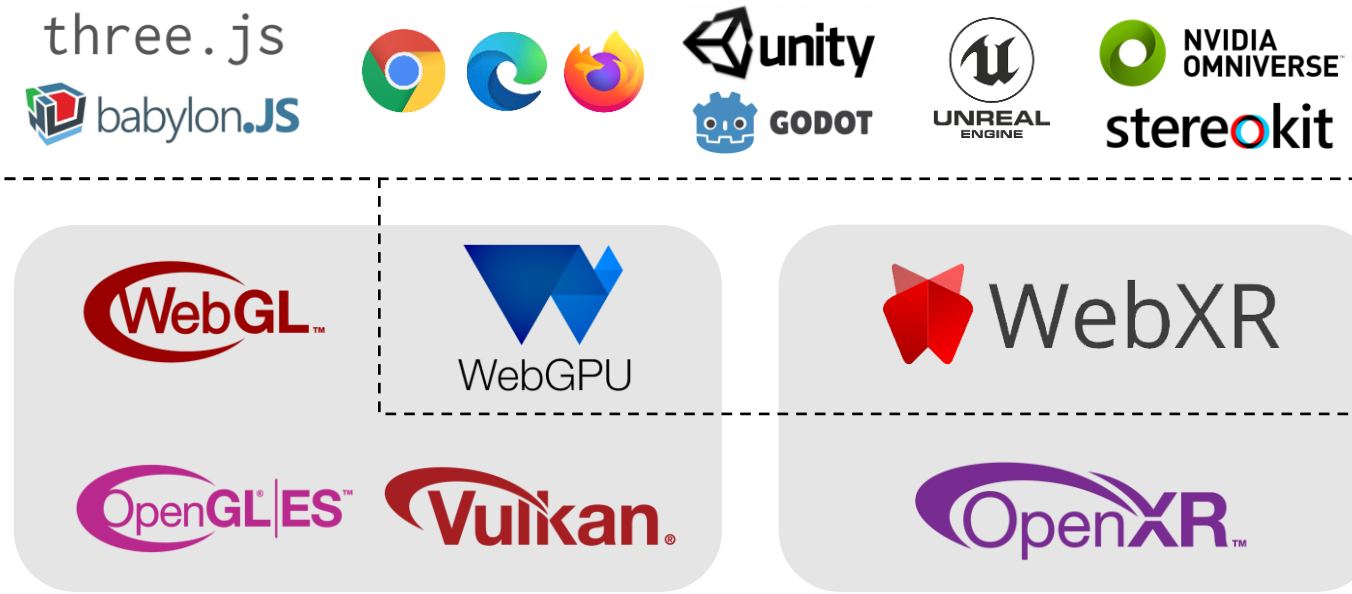

Immersive Web: Standards & Open Source

Creating the Immersive Web will need and leverage the work of many standards organizations, consortia and open-source projects

Standardization is a cooperative endeavor!

The Immersive Web BOF will provide latest updates and how the industry is working together

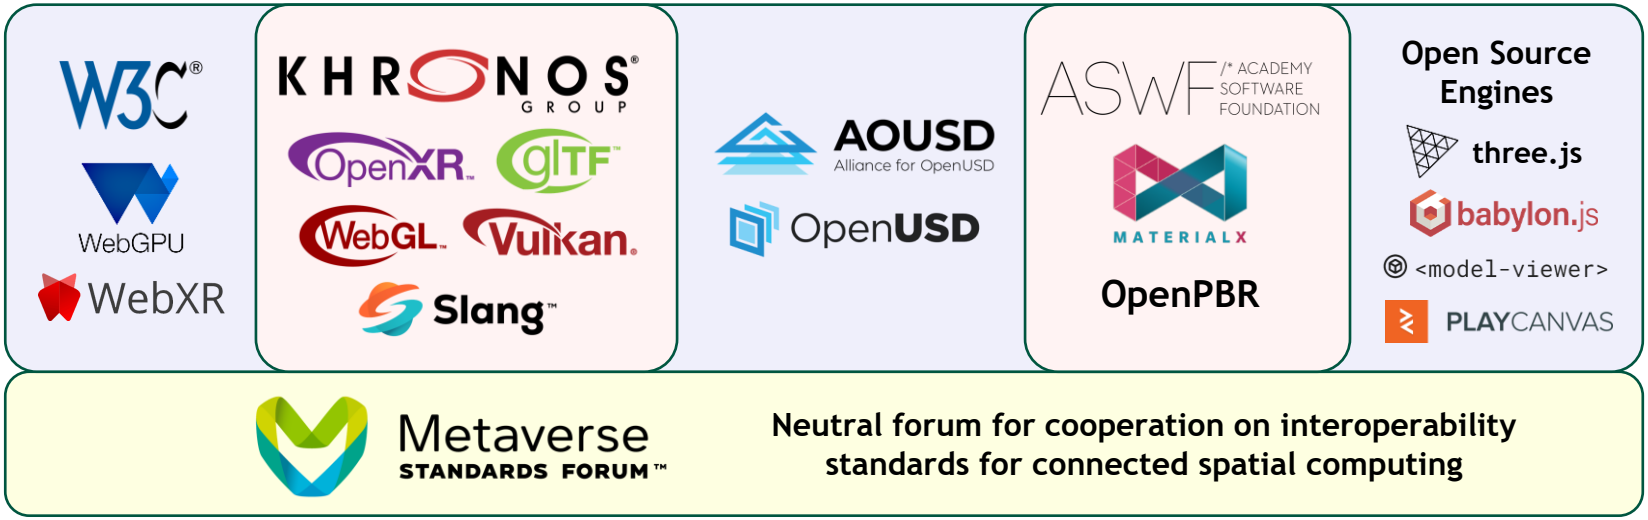

Immersive Web BOF Session!

WebGL

glTF™ WebXR

WebGPU

SIGGRAPH Asia Khronos BOFs

Forging the Immersive Web through Real-Time Graphics

Wednesday 4th

3:30-4:30PM, G408

Topics

glTF and WebGL updates

Khronos and W3C cooperation

Immersive Web and WebXR update

three.js update (mrdoob!)

WebGPU update

OpenPBR and the web

ASWF /* ACADEMY SOFTWARE FOUNDATION W3C® three.js KHRONOS GROUP

KHRONOS GROUP

Vulkan: Forging Ahead!

Ralph Potter (Samsung), Vulkan Working Group chair

Releasing Today!

https://khr.io/vulkan14

Vulkan

An explicit API for graphics and compute on GPUs

- Radically cross-platform, from embedded to desktop
- Focus on high performance and user control

Driving the future evolution of graphics hardware

- Setting requirements for new hardware
- Ensuring compatibility with current hardware
- Focus on solving issues raised by industry experts

Developed collaboratively by industry experts

- Input considered from a wide range of sources

Vulkan Roadmap: Setting industry direction

The roadmap sets forward-looking targets for new hardware, with milestones established for new mid-to-high-end GPUs in the year of their release

Roadmap **2022**

Roadmap **2024**

Roadmap **2026**

Core Specifications: Support for Current GPUs

Features that can be supported on current GPUs are brought into core

Maintenance Updates

Vulkan **Core 1.0**

Vulkan **Core 1.1**

Vulkan **Core 1.2**

Vulkan **Core 1.3**

Vulkan **Core 1.4**

Vulkan **Core 1.5**

2016 2018 2020 2022 2024

Roadmap Sets Direction, Core Solidifies It

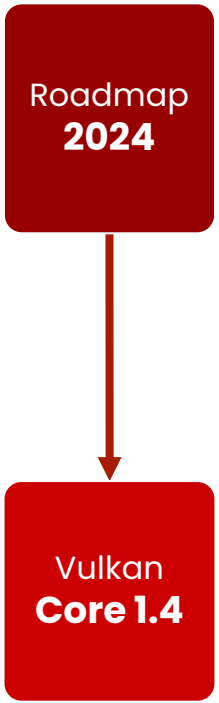

- **Vulkan 1.4 is the first core version derived from the roadmap**
  - Notable benefits both in design and development
  - Enabled huge increase in supported features
- **Most of the tough questions for 1.4 largely already answered**
  - Future direction already set with the roadmap
  - Features already designed, shipped, and implemented
  - Vendors already knew which hardware could support what
- **“Just” had to put the pieces together**
  - Much easier development cycle than previous cores
  - Allowed us to focus on future roadmap items

Vulkan 1.4 Core Specification

Integrates significant requested functionality proven as extensions

Mandated support for new functionality ensures availability on all Vulkan 1.4 implementations

Dynamic rendering local read bringing subpass support to the dynamic rendering API

Streaming transfers via host image copy or mandatory async transfer queue support

Fine-grained control of floating point optimization behavior

Mandating previously optional features such as scalar block layout and 8/16 integer support

Maintenance extensions up to VK_KHR_maintenance6

Several limit increases, including 8K rendering with up to eight separate render targets

And more…
16 extensions promoted
28 optional features become mandatory
31 minimum hardware limits raised

Vulkan BOF Session!

Vulkan: Forging Ahead
Thursday 5th
3:30-5:30PM, G407

Topics
Vulkan 1.4 launch!
Vulkan SDK Update
HDR support in Vulkan
GPU-driven Rendering in Vulkan
Vulkan Safety Critical

Khronos BOFs

KHRONOS GROUP

Slang Fast Forward

Shannon Woods
NVIDIA, Slang Working Group Chair

Open-Source, Cross-Platform Compiler

Slang compiler ingests Slang shaders AND provides onramps for HLSL and GLSL codebases

Slang
HLSL 2020
GLSL

Open source
https://github.com/shader-slang

SPIR-V → Vulkan
HLSL → Microsoft DirectX
GLSL → OpenGL
MSL
WGSL → WebGPU
PyTorch
C (CPU)
CUDA
Optix

Diverse backend targets enable shader code to be written once to run almost anywhere

Khronos now hosts the Slang open-source shading language and compiler
Builds on 15 years development at NVIDIA
Modular code, portable deployment, and neural computation

Slang is Now Under Khronos Open Governance

- **3D developers need a modern, responsive domain-specific shading language**
  - To fit an ecosystem that has grown in scope and complexity
- **Slang leverages 15 years of R&D and development experience**
- **Transition to Khronos to foster industry-wide collaboration and innovation**

For developers, by developers!

- Community structure built from OSS best practices
- **Any company or individual** is welcome to become a contributor, not just Khronos members
- Decision-making and development in the open – you can join technical conversations today on Discord, or propose features directly to the repository.
- Slang developers make the decisions about what goes into the language, and you can become one

The Neural Rendering Revolution

- **Intuitive use of machine learning in real-time shaders**
  - Generate geometry/texture/material LODs, compression, approximations, parameter tuning
  - Embed training inside the renderer e.g., accelerate learned Neural Materials
- **Deeply integrated, first-class automatic differentiation**
  - Eliminates need for additional, separate differential version of every shader

Compression
16x more Texels than D3D/Vulkan textures in same footprint

Neural Materials
Movie-quality photoreal materials in real-time with appearance-based LOD

Neural Shapes
New content creation paradigms and LOD solution

Neural Lighting
Neural Radiance
Neural Path Sampling

Thank you for your interest in Slang!

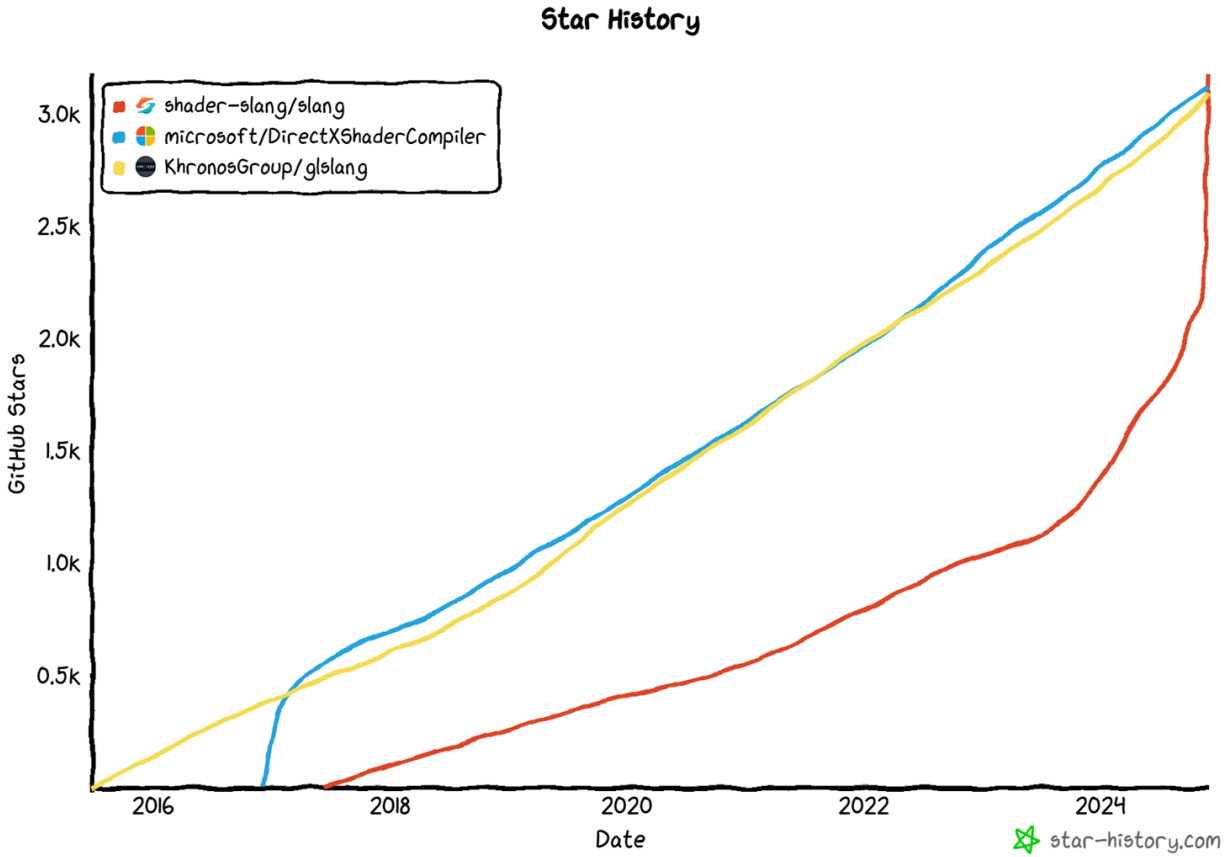

Slang BOF Session!

Slang and the 3D Shading Language Landscape

Wednesday 4th
1:00-2:00PM, G407

Khronos BOFs

KHRONOS GROUP

KHRONOS GROUP

OpenXR™

OpenXR Fast Forward

Jian Zhang
Head of XR Foundation Engineering, PICO

OpenXR Cross-Platform Portability

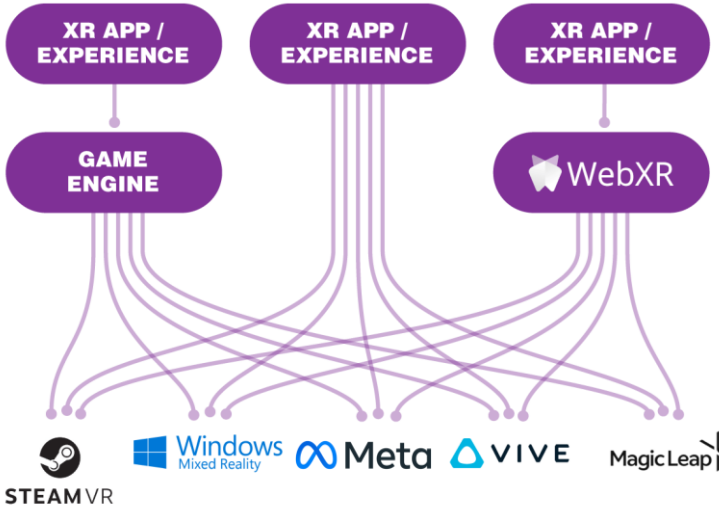

Before OpenXR: Applications and engines needed separate proprietary code for each device on the market.

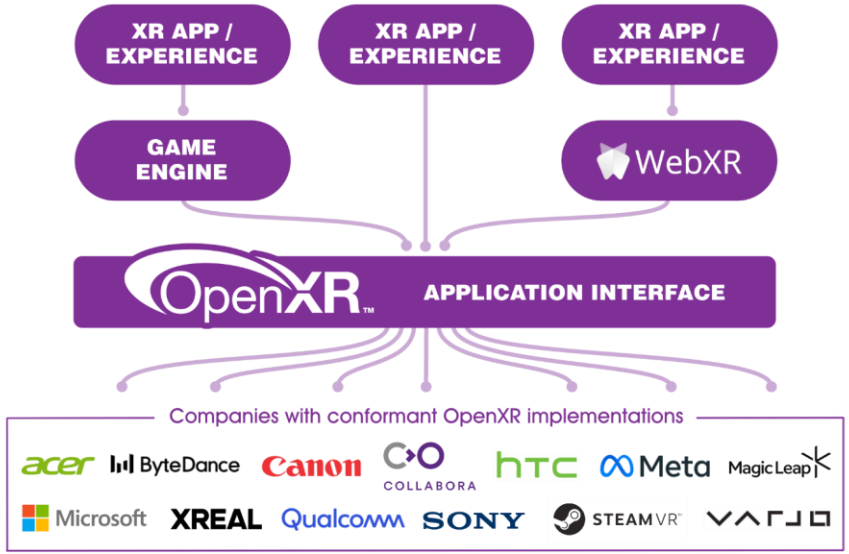

OpenXR provides a single cross-platform, high-performance API between applications and all conformant devices.

Applications and engines can portably access any OpenXR-conformant hardware

Conformant OpenXR Devices

 Microsoft	 Meta	 htc
HoloLens and Mixed Reality Headsets. Hand and eye tracking extensions	Rift S, Quest 3, Quest 2 and Quest Pro Meta Deprecated own API for OpenXR	Vive Focus 3, Vive Cosmos, Vive XR Elite, Vive Wave Runtime
 VALVE	 VARJO	 Canon
Valve Index Valve Deprecated OpenVR APIs for OpenXR	All Varjo Headsets are fully compliant XR-3, XR-4	MREAL X1
 Magic Leap	 XREAL	 Snapdragon spaces
Magic Leap 2	XREAL Air 2, Air 2 Pro, Air 2 Ultra	Qualcomm Snapdragon Spaces XR Development Platform
 acer	 PICO	 SONY
Spatial Labs Display Series	Neo 3, PICO 4, PICO 4 Pro, PICO 4 Ultra	Spatial Reality Displays

The OpenXR Story So Far…

Empowering Cross-platform Immersive Experiences

OpenXR 1.1
Consolidates multiple extensions to streamline application development and reduce fragmentation
Adds new functionality with spec improvements

Vendor Proprietary API fragmentation
Clear industry demand need for a cross-platform XR open standard

Establishing baseline XR functionality
Though industry consensus and contributed designs

OpenXR 1.0 specification drafted

OpenXR achieves wide industry adoption

OpenXR is foundation for experimentation
New functionality introduced through extensions

Increased focus on regular core spec updates
Balancing the need to ship new functionality *AND* consolidate widely proven technology

Leverage OpenXR's flexible design to explore new use cases
e.g., body tracking and advanced spatial computing

OpenXR Working Group Formed — 2017

OpenXR 1.0 Released — 2019

New! OpenXR 1.1 Released — April 2024

OpenXR 1.1 Key Extensions Promoted to Core

- **Local Floor Reference Space**
  - Gravity-aligned world-locked origin for standing-scale content
  - Estimated floor height built in
  - Recenter to current user position at the press of a button without a calibration procedure
- **Grip Surface**
  - Anchors visual content relative to the user's physical hand
  - Can be tracked directly or inferred from a physical controller’s position and orientation
- **Stereo with Foveated Rendering for XR headsets**
  - Runtimes MAY optionally expose eye-tracked or fixed foveated rendering
  - Portable across multiple graphics rendering APIs
- **Additional enhancements**
  - Interaction Profile improvements
  - Spec language cleanup and clarifications

Coming Soon to OpenXR

- **Extending hand tracking**
  - To include full body tracking
- **Enhanced handling of spatial entities**
  - Standardized methods to interact with the user's environment
  - Support for advanced spatial computing applications
- **Expanded haptics support**
  - Support immersive experiences through PCM, vibrotractiles, and transients
- **Controller render models (glTF)**
  - Showing and animating a model of the user's actual controller

OpenXR Specification

OpenXR SDK GitHub

OpenXR BOF Direction Discussions

- **Empowering Cross-Platform Immersive Experiences**
  - The long-term goal of OpenXR
- **Enabling Multi-Application Support in XR**
  - Extend the XR use cases beyond Fully Immersive experiences: Multi-App
  - Rendering Architecture: Unified Rendering vs Self Rendering
  - *Jian Zhang, Praveen Babu J D (PICO)*
- **Secured and private access to XR device cameras / Sensors**
  - Framework to run your own ML model on XR devices
  - Developer driven mixed reality effects
  - *Jimmy Alamparambil, Jane Tian (PICO)*

OpenXR BOF Session!

OpenXR: Transforming the Future of Cross-Platform XR Development

Thursday 5th

2:15-3:15PM, G407

Khronos BOFs

KHRONOS GROUP

Khronos BOFs at SIGGRAPH Asia

Day	Time / Room	Session Title	Standards and Projects
Tuesday 3rd	1:00-2:00PM, G408	Khronos Fast Forward	Vulkan, OpenXR, Slang, ANARI, glTF
Wednesday 4th	1:00-2:00PM, G407	Slang Shading Language	Slang
Wednesday 4th	3:30-4:30PM, G407	Immersive Web with Khronos and W3C	WebGL, WebXR, WebGPU, three.js
Thursday 5th	2:15-3:15PM, G407	OpenXR Update and Roadmap	OpenXR
Thursday 5th	3:30-5:30PM, G407	Vulkan Update and Ecosystem	Vulkan, Vulkan SC, Slang
Friday 6th	11:00-12PM, G408	glTF 3D Transmission Format	glTF, VRM Avatar Format

All BOF slides and videos will be uploaded to the
Khronos SIGGRAPH event page

SIGGRAPH ASIA 東京
2024 TOKYO 3–6 DEC

Khronos BOFs

KHRONOS GROUP

Khronos Information
www.khronos.org
memberservices@khronosgroup.org