KHRONOS GROUP

東京
SIGGRAPH
ASIA 2024
TOKYO

BRIEFING & NETWORKING RECEPTION

SIGGRAPH ASIA 2024 TOKYO 東京

KHRONOS GROUP

ANARI Kamaros OpenXR WebGL Vulkan glTF SYCL Slang OpenCL

Khronos Briefing and Reception

Neil Trevett

Khronos President and VP Developer Ecosystems at NVIDIA

Khronos Connects Software to Silicon

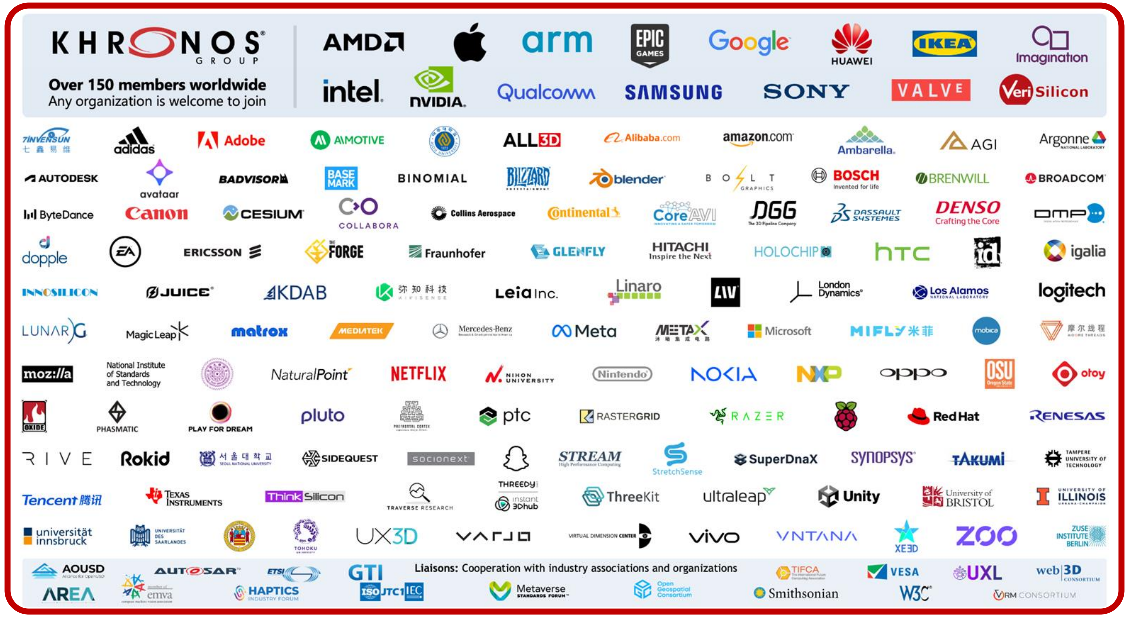

Founded in 2000

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

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

What is an Interoperability Standard?

INTEROPERABILITY standards define precise COMMUNICATION
E.g., software to hardware APIs, client to server networking protocols, common asset formats

Successful standards are thoughtful about level of abstraction
Enable and encourage *implementation* differentiation

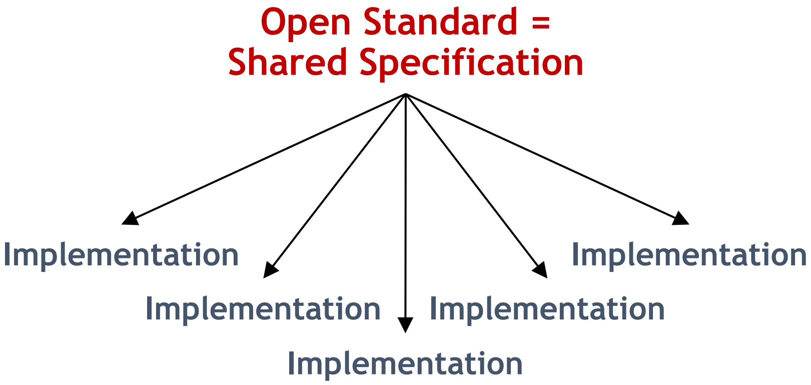

Open Standard = Shared Specification

Implementation
Implementation
Implementation
Implementation
Implementation
Implementation

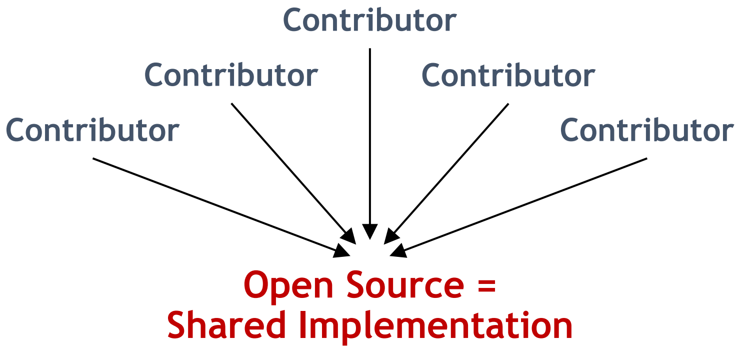

Contributor
Contributor
Contributor
Contributor
Contributor

Open Source = Shared Implementation

Open standards often need open source to grow and enable their ecosystems
Conformance testing, tooling, sample implementations

Sometimes open source is a better ecosystem solution than an open standard
E.g., kernel and shading language compilers

For Standards…

‘Open’ means…

Open to all who wish to participate in their creation
No restrictions on who can implement and adopt
Genuine multi-company governance is essential for industry confidence and adoption
Longevity not dependent on resources or current business focus of any one company

‘Free’ means…

No charge for access to specification documents
No charge to users of specifications
Royalty-free patent license to implementers from all involved in creating the specification
Well designed IP framework means more creator patents == MORE protection

Open Standards and Product Innovation

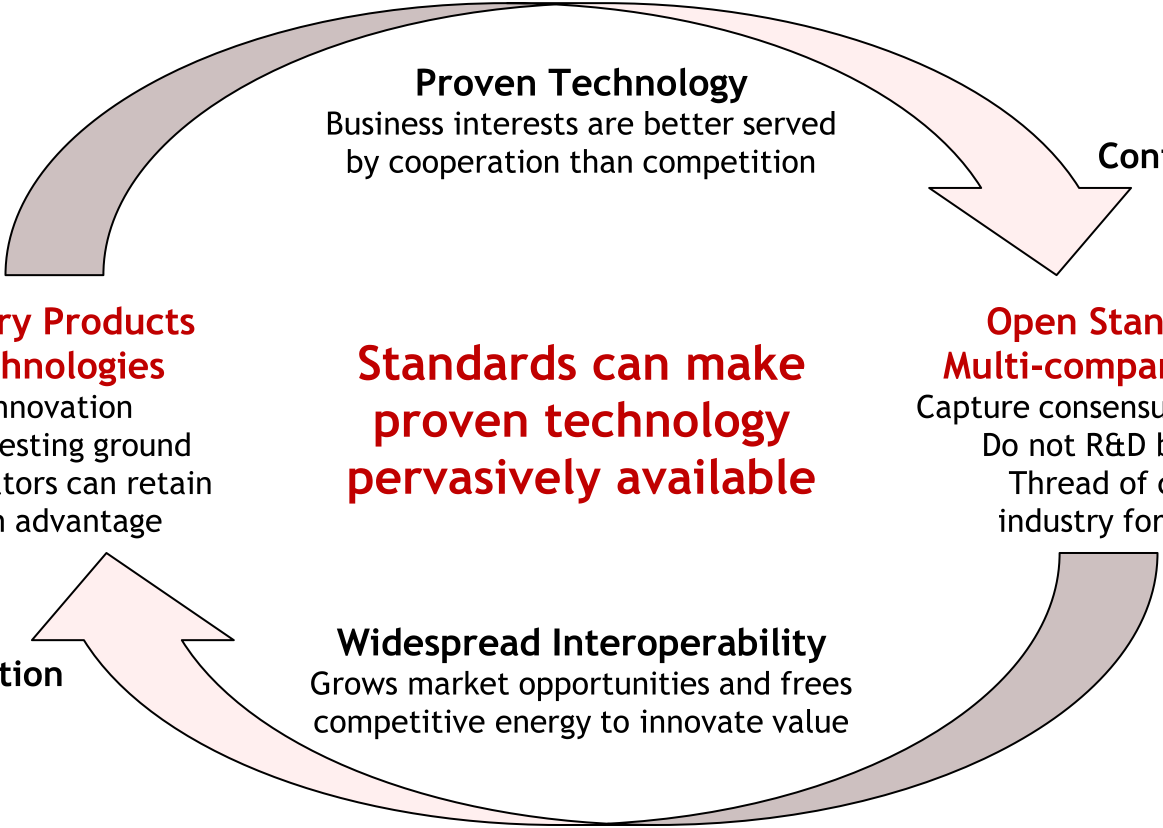

Proven Technology
Business interests are better served
by cooperation than competition

Contributions

Open Standards under Multi-company Governance
Capture consensus for wide adoption
Do not R&D by committee!!
Thread of continuity for
industry forward progress

Widespread Interoperability
Grows market opportunities and frees
competitive energy to innovate value

Adoption

Proprietary Products and Technologies
Rapid Innovation
Darwinian testing ground
Smart Innovators can retain
long-term advantage

Standards can make proven technology pervasively available

Open Standard Why’s, When's and How’s

Why?	**Grow Markets** Reduced consumer confusions and increased capabilities and usability	**Reduce Costs** Shared development between many companies and driving volume
	Speed Time to Market Well-proven functionality, testing and interoperability	**Enable Innovation** Free resources to compete on quality, performance, power etc.
When?	**When Technologies are Proven** Avoid research and development by standards committee	**Consensus Need** The downsides of not having a standard are obvious
How?	**Multi-company governance** Avoid single-company control or dependency	**Well-defined IPR Policy** Royalty-free standards have clearer path to wide adoption
	With Extensibility Enable a responsive innovation pipeline to meet customer and market needs	**Thoughtful Abstraction** Encourage innovation through freedom to innovate implementation details

Lesson 1: Timing is Everything

The optimal window for standardization is when:

A technology is just becoming proven – DON’T do R&D by standards committee!

There is strong consensus on the pain of NOT having a standard

glTF PBR

Incremental consolidation and meticulous specification of *accepted industry practice*

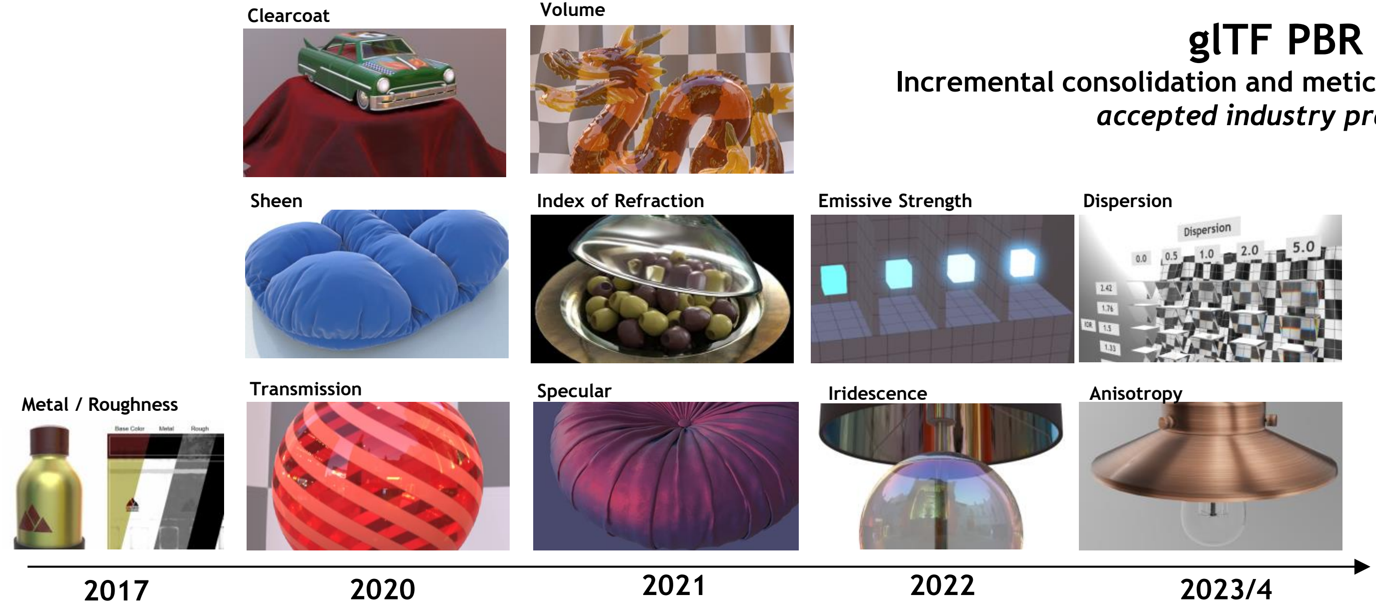

Clearcoat

Volume

Sheen

Index of Refraction

Emissive Strength

Dispersion

Metal / Roughness

Transmission

Specular

Iridescence

Anisotropy

2017

2020

2021

2022

2023/4

Lesson 2: Test, Test and Test

What isn’t tested is not working

Bugs and inconsistencies undermine the fundamental value of a standard

If ‘dialects’ evolve, industry and community collapse in trust in a standard can be swift and fatal

Content Developers
test and validate assets

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

Public comparative showcase between
multiple engines provides:
1) Early warning of inconsistencies
2) Incentive to implement functionality

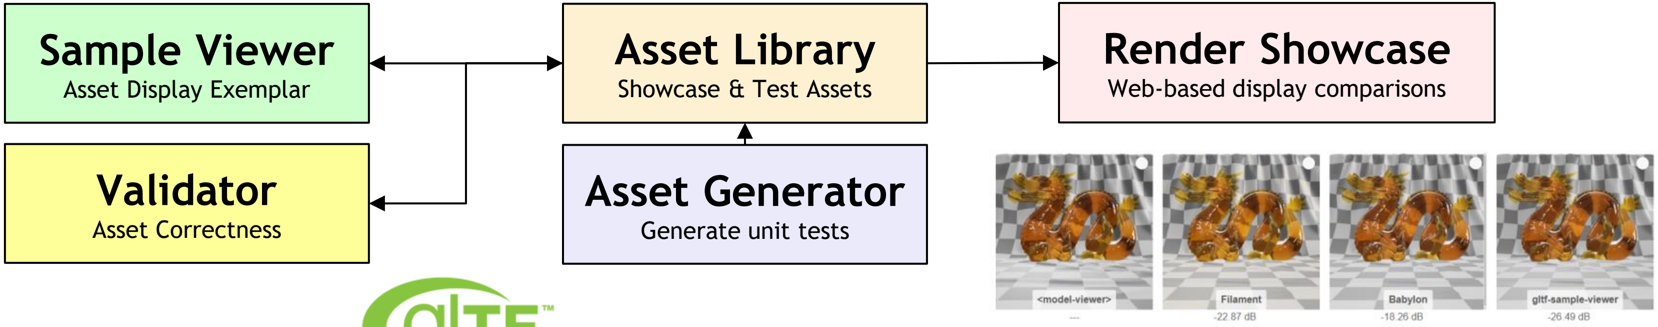

Sample Viewer
Asset Display Exemplar

Validator
Asset Correctness

Asset Library
Showcase & Test Assets

Asset Generator
Generate unit tests

Render Showcase
Web-based display comparisons

Lesson 3: Embrace Darwinism

The market knows best

Extensible standards enable community exploration of new functionality
Notice and react to what actually gets built and adopted and WHY
Strong industry Adoption is essential to sustain momentum and ongoing investment

These stages are aspirational - often spec evolution will not follow exact steps - that's OK	Proposal (optional)	Initial Draft	Review Draft	Release Candidate	Ratified	International Standard	Archived
Qualifying Criteria to enter stage	Draft with emphasis on motivation and 'philosophy'	Draft specification but not yet ready for wide review	Well-formed specification JSON schema Detailed design	Stable Draft incorporating feedback from primary stakeholders >1 implementations Notice for final Public Review	Only non-IP-breaking updates during Ratification Review	Stable and widely adopted	Advise against new usage Enables legacy assets
Work undertaken during stage			Spec incorporates initial input and feedback Initial Test Assets >0 implementations	Spec incorporates additional received feedback Full Test Asset suite Validator and Viewer Support Initial additional Tools support	Ratified PDF and HTML specs posted on Registry Expanded Tools support	Included in PAS Submission to JTC 1	
External Messaging	Alternative proposals or directional discussions?	High-level or directional feedback?	Detailed technical review requested	Last chance for feedback	Please report bugs and issues	Press Release	Archival Notice

Lesson 4: Don’t Be in a Bubble

Standardization is a cooperative endeavor

Reach out to other initiatives to seek synergies and avoid needless differences
Can strengthen understanding and focus on your core values to the industry

glTF™

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

ASWF /* ACADEMY SOFTWARE FOUNDATION

Multiple open-source projects including OpenPBR and MaterialX

OpenUSD

AOUSD Alliance for OpenUSD

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

Khronos Principles of Organization

Any company is welcome to join.
One company one vote

Strong IP Framework to enable
ROYALTY-FREE specifications:
all members agree not assert patents
against conformant implementations

Only invest where there is
strong industry momentum to
ensure industry relevance -
let Darwinism rule!

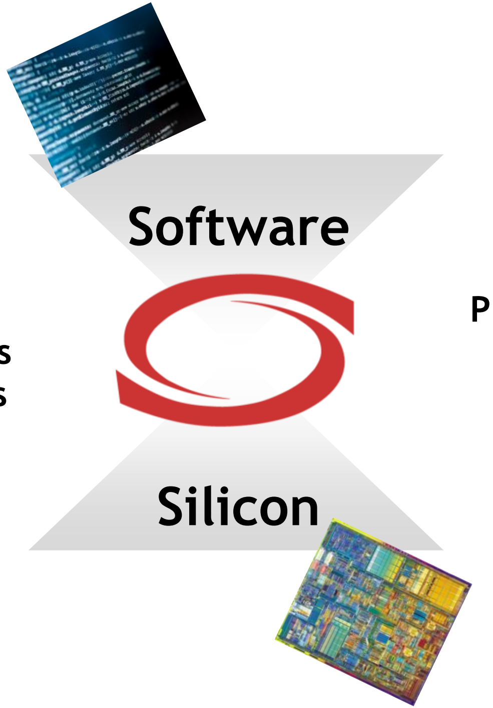

Software

Silicon

Defined process for members
and non-members to propose
new standards initiatives

Conformance Tests and Adopters
Programs for defining conformance,
specification integrity and cross-vendor portability

Non-profit organization -
All revenue is reinvested
into developing standards
that benefit the industry

Khronos Cooperative Framework

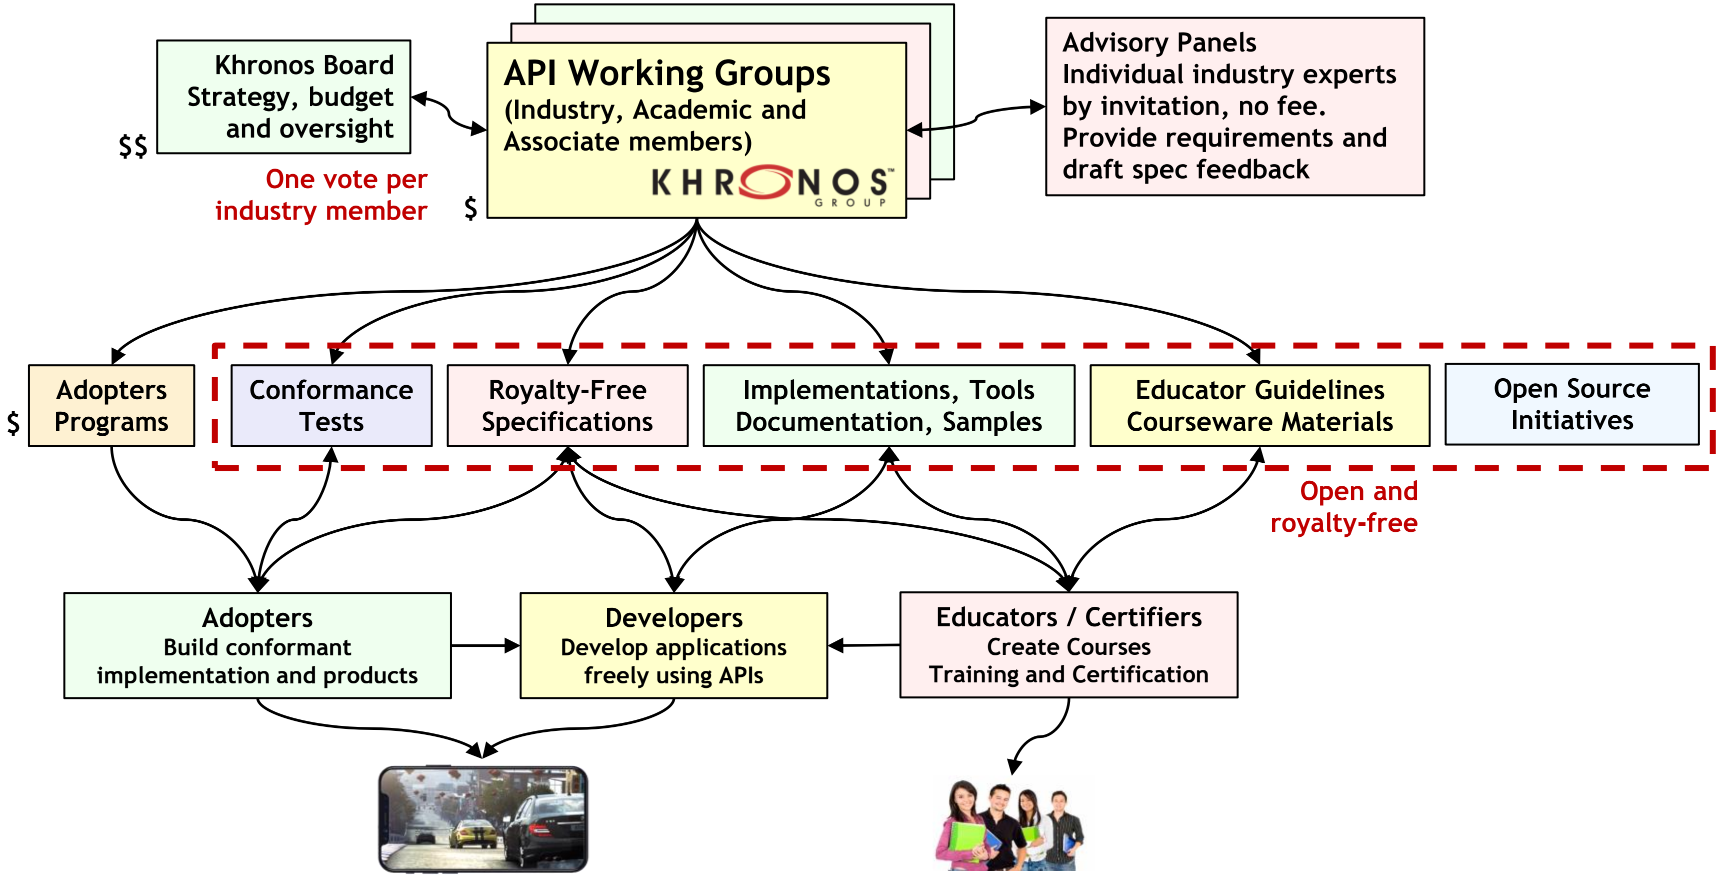

Khronos Board
Strategy, budget
and oversight

$$

API Working Groups
(Industry, Academic and
Associate members)

KHRONOS GROUP

$

Advisory Panels
Individual industry experts
by invitation, no fee.
Provide requirements and
draft spec feedback

One vote per
industry member

$

Adopters
Programs

Conformance
Tests

Royalty-Free
Specifications

Implementations, Tools
Documentation, Samples

Educator Guidelines
Courseware Materials

Open Source
Initiatives

Open and
royalty-free

Adopters
Build conformant
implementation and products

Developers
Develop applications
freely using APIs

Educators / Certifiers
Create Courses
Training and Certification

Khronos New Initiative Process

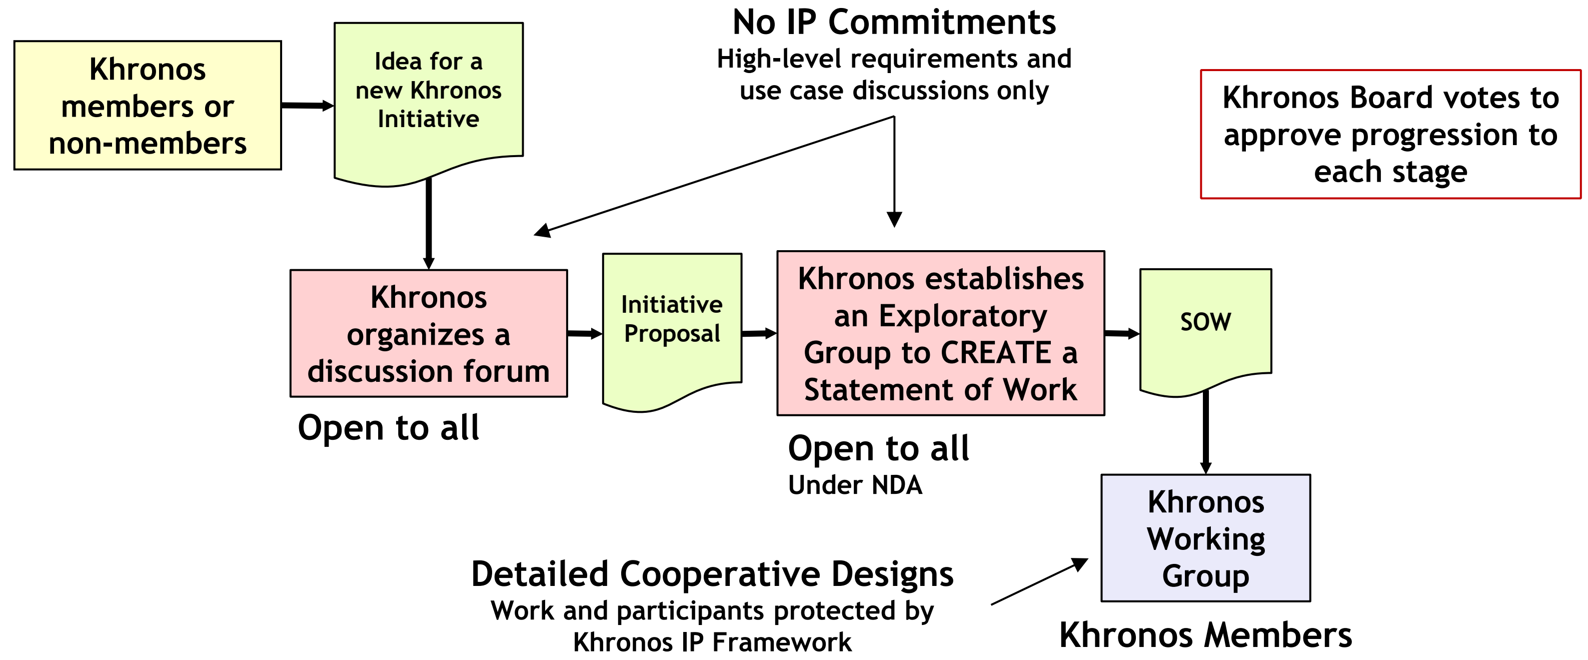

Khronos members or non-members
Idea for a new Khronos Initiative
No IP Commitments
High-level requirements and use case discussions only
Khronos Board votes to approve progression to each stage
Khronos organizes a discussion forum
Initiative Proposal
Khronos establishes an Exploratory Group to CREATE a Statement of Work
SOW
Open to all
Open to all
Under NDA
Khronos Working Group
Detailed Cooperative Designs
Work and participants protected by Khronos IP Framework
Khronos Members

The Exploratory Group Process enables Khronos, and its members, to explore whether to undertake a new initiative without committing any member to IP licensing obligations

The Value of Khronos Participation

Proven processes for building multi-company consensus for generation and governance of open interoperability standards

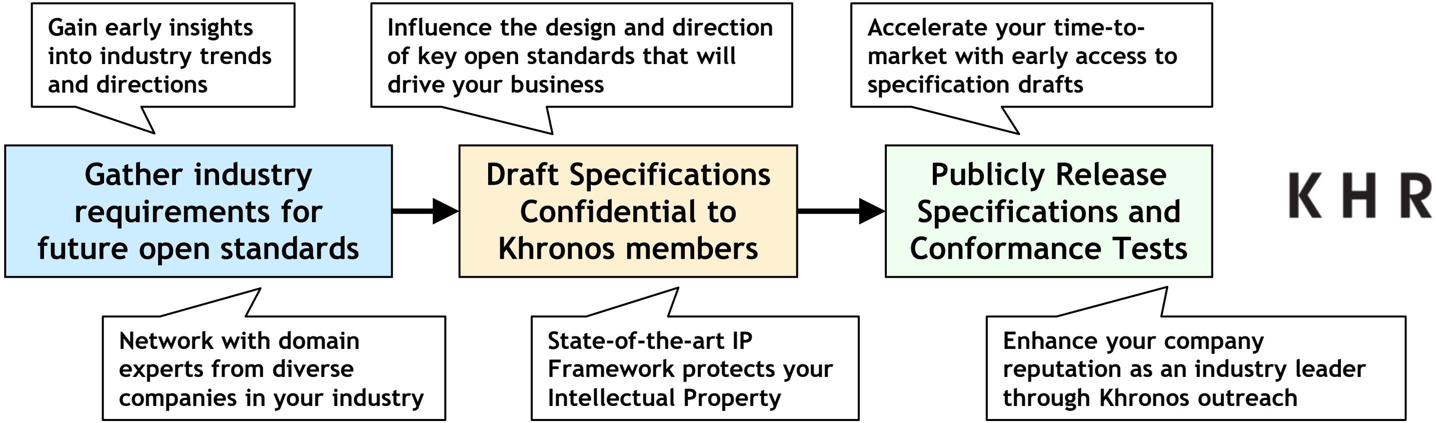

Gain early insights into industry trends and directions

Influence the design and direction of key open standards that will drive your business

Accelerate your time-to-market with early access to specification drafts

Gather industry requirements for future open standards

Draft Specifications Confidential to Khronos members

Publicly Release Specifications and Conformance Tests

Network with domain experts from diverse companies in your industry

State-of-the-art IP Framework protects your Intellectual Property

Enhance your company reputation as an industry leader through Khronos outreach

Khronos membership is open to any company for access to all standardization initiatives

Annual membership fees start at $4,000 for smaller companies

https://www.khronos.org/members/

Khronos Active Standards

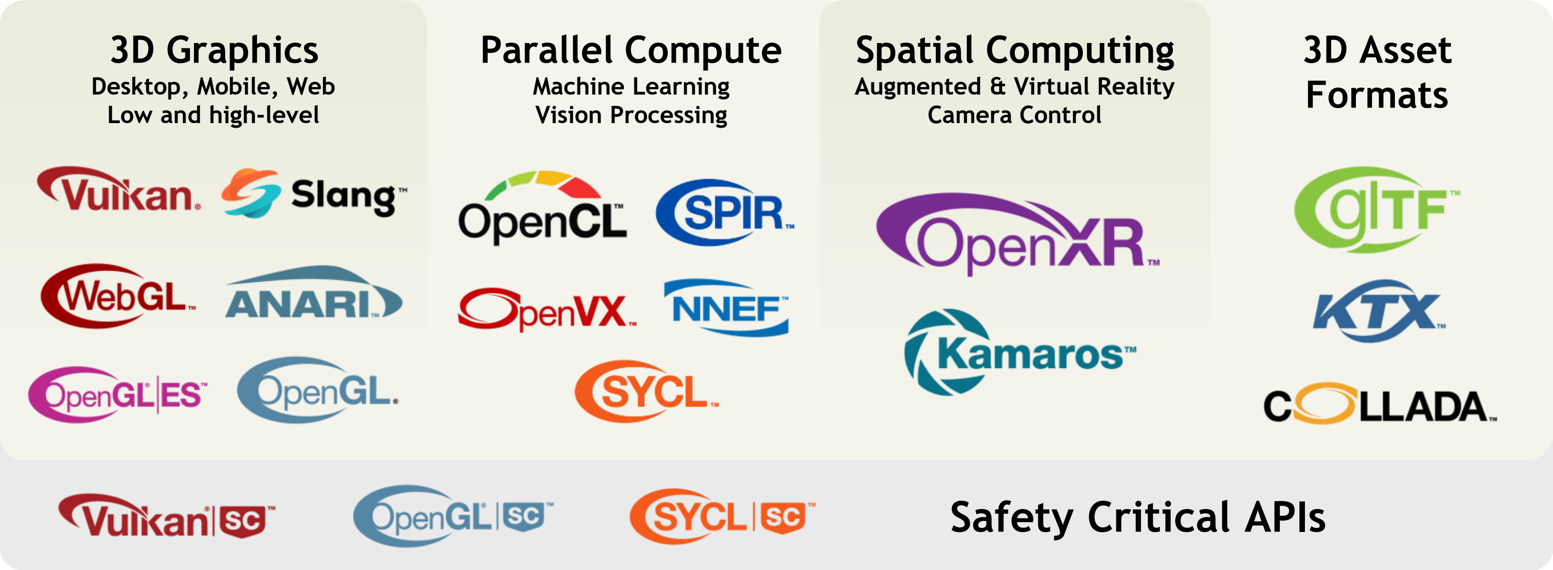

3D Graphics

Desktop, Mobile, Web
Low and high-level

Parallel Compute

Machine Learning
Vision Processing

Spatial Computing

Augmented & Virtual Reality
Camera Control

3D Asset Formats

Safety Critical APIs

Khronos Standards for Spatial Computing

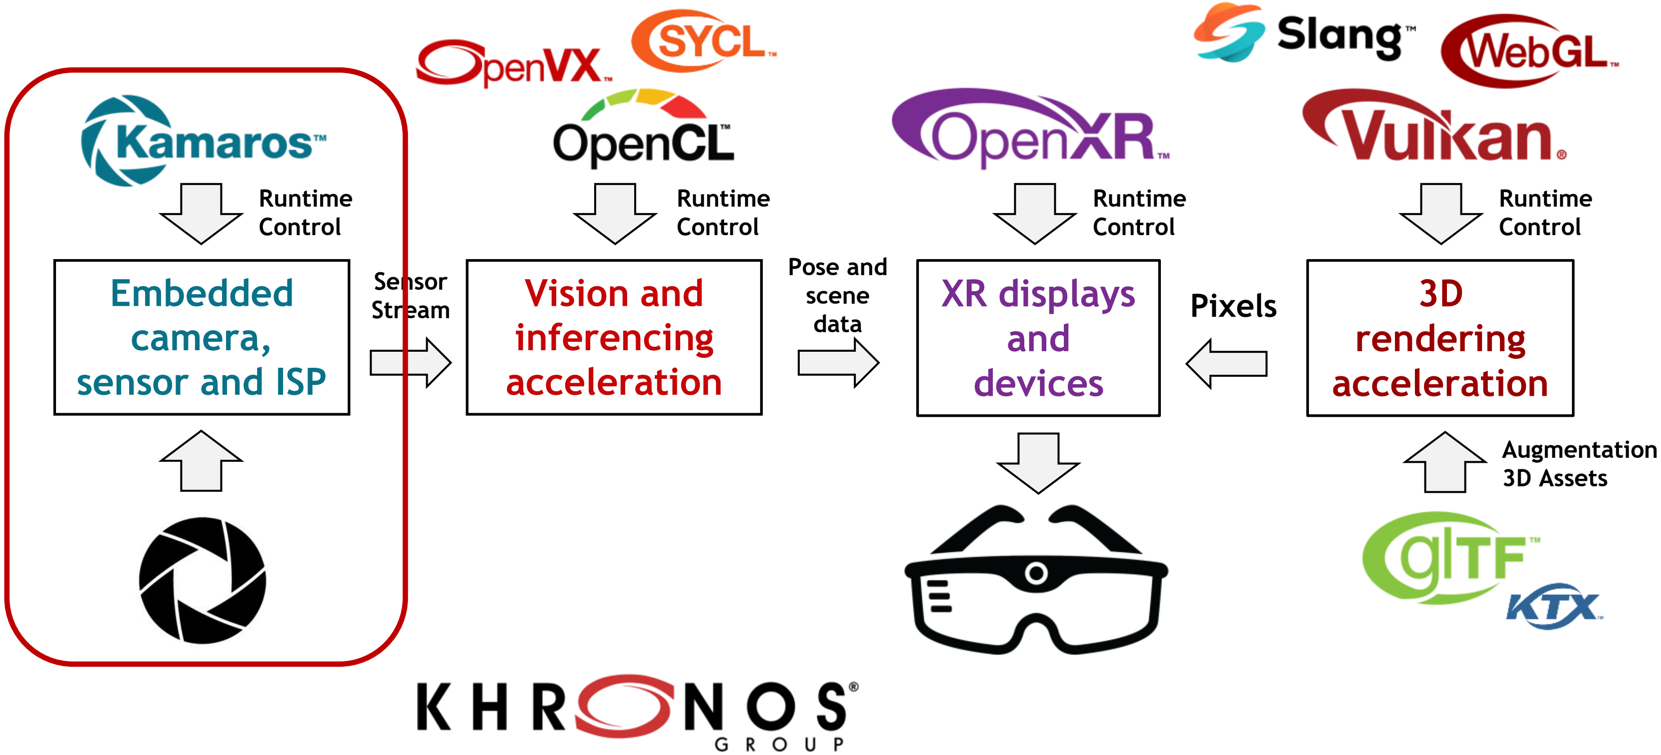

Runtime Control

Embedded camera, sensor and ISP

Sensor Stream

Runtime Control

Vision and inferencing acceleration

Pose and scene data

Runtime Control

XR displays and devices

Pixels

Runtime Control

3D rendering acceleration

Augmentation 3D Assets

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

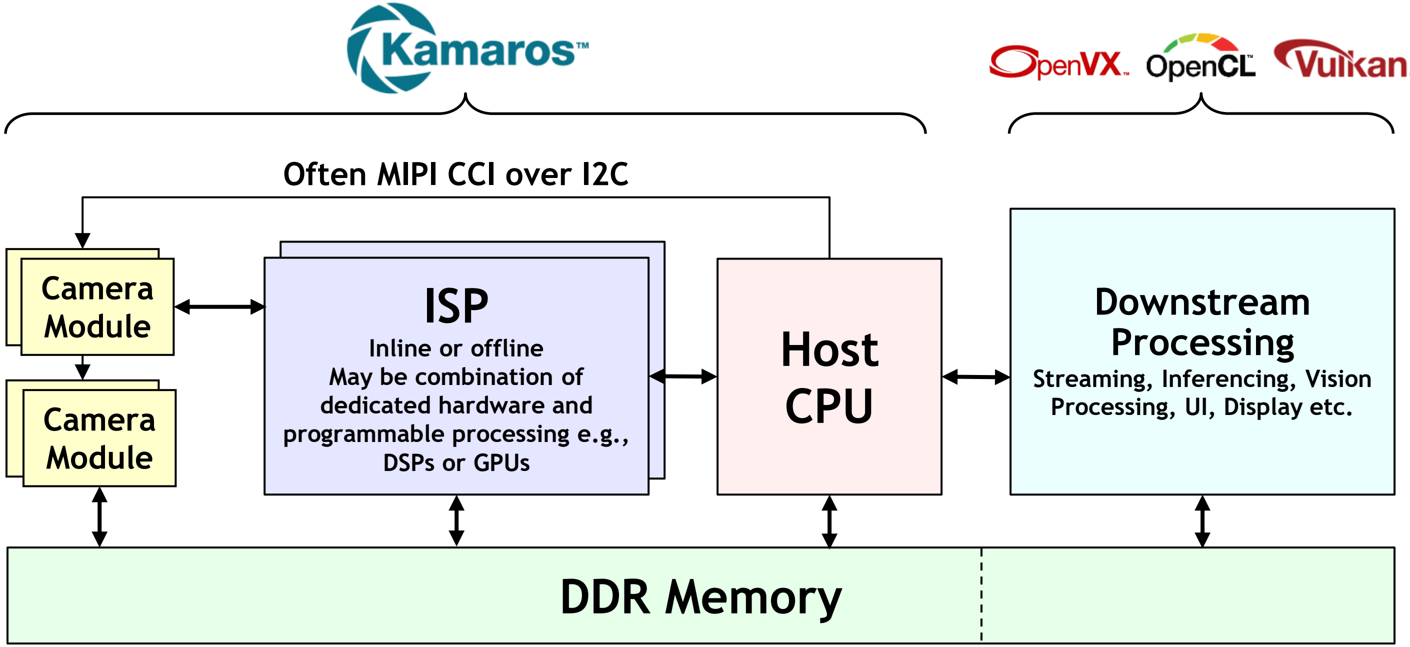
Kamaros™
OpenVX
OpenCL
Vulkan
Often MIPI CCI over I2C
Camera Module
Camera Module
ISP
Inline or offline
May be combination of dedicated hardware and programmable processing e.g., DSPs or GPUs
Host CPU
Downstream Processing
Streaming, Inferencing, Vision Processing, UI, Display etc.
DDR Memory

Khronos Standards for Spatial Computing

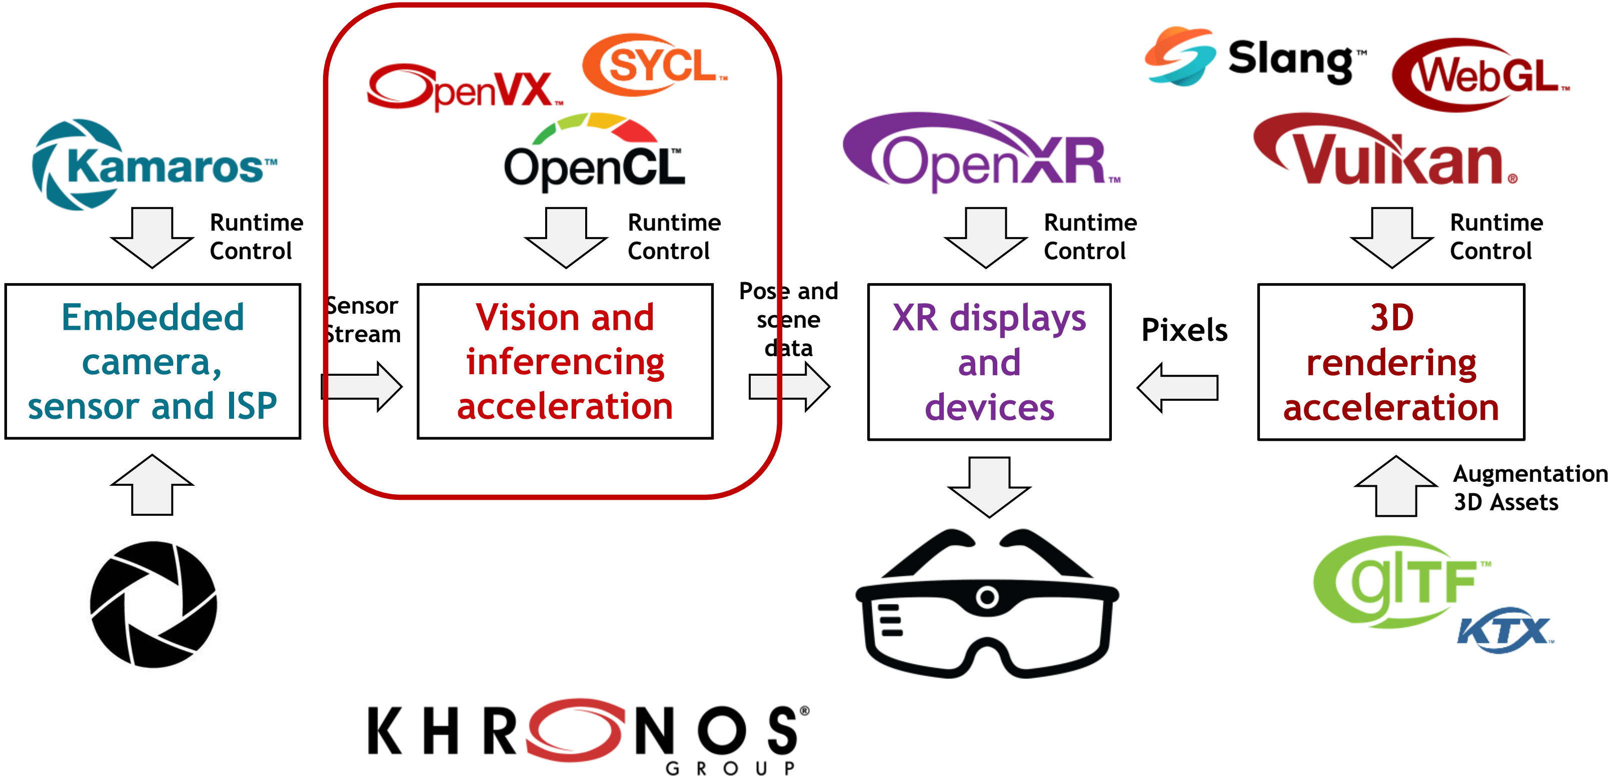

Machine Learning Acceleration APIs

Open-Source Frameworks | Acceleration APIs

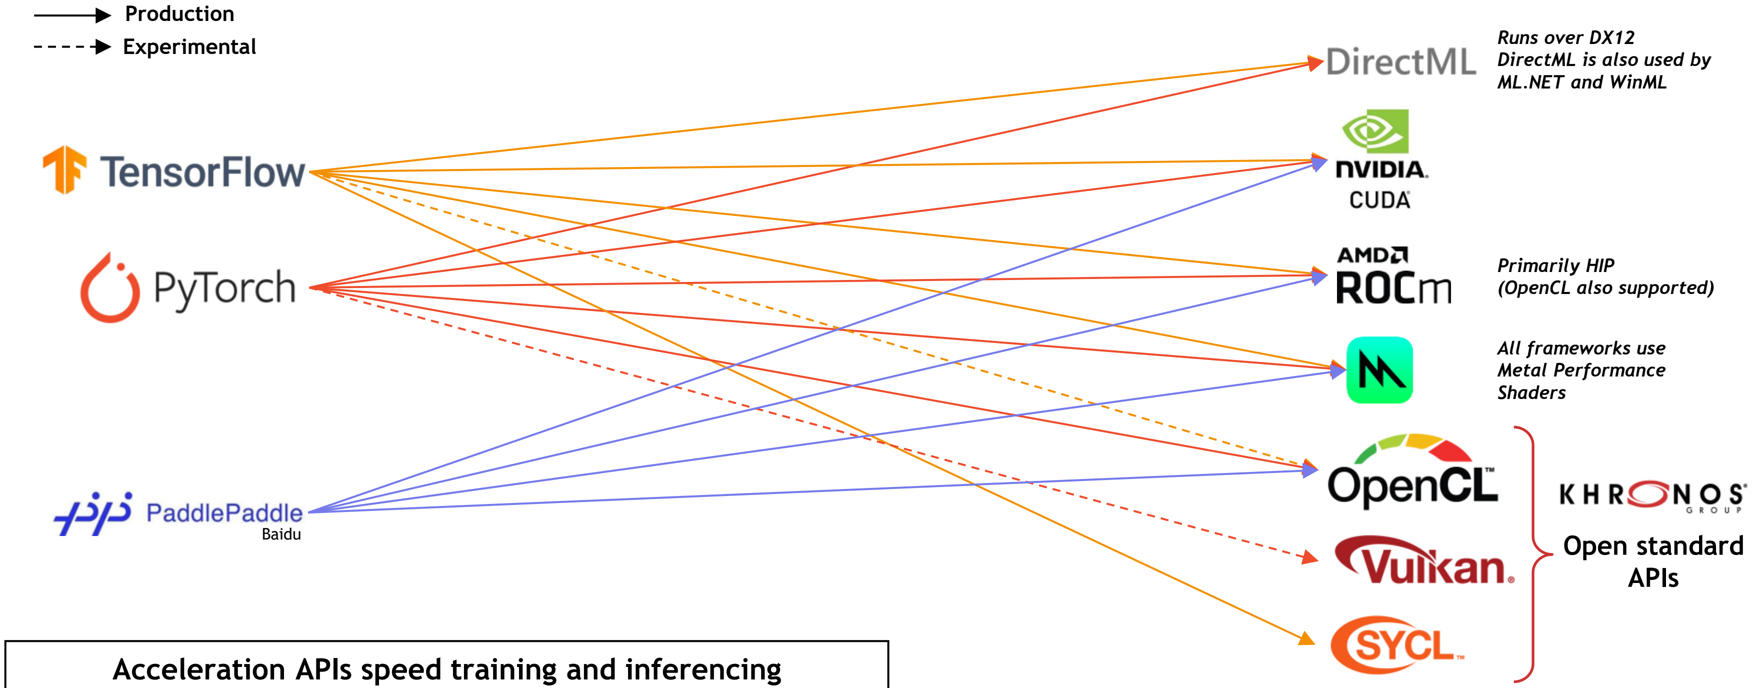

Runs over DX12
DirectML is also used by ML.NET and WinML

Primarily HIP
(OpenCL also supported)

All frameworks use Metal Performance Shaders

Open standard APIs

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

Machine Learning Acceleration APIs

Open-Source Frameworks	Compilers, Runtimes and Libraries (some in open source)	Acceleration APIs

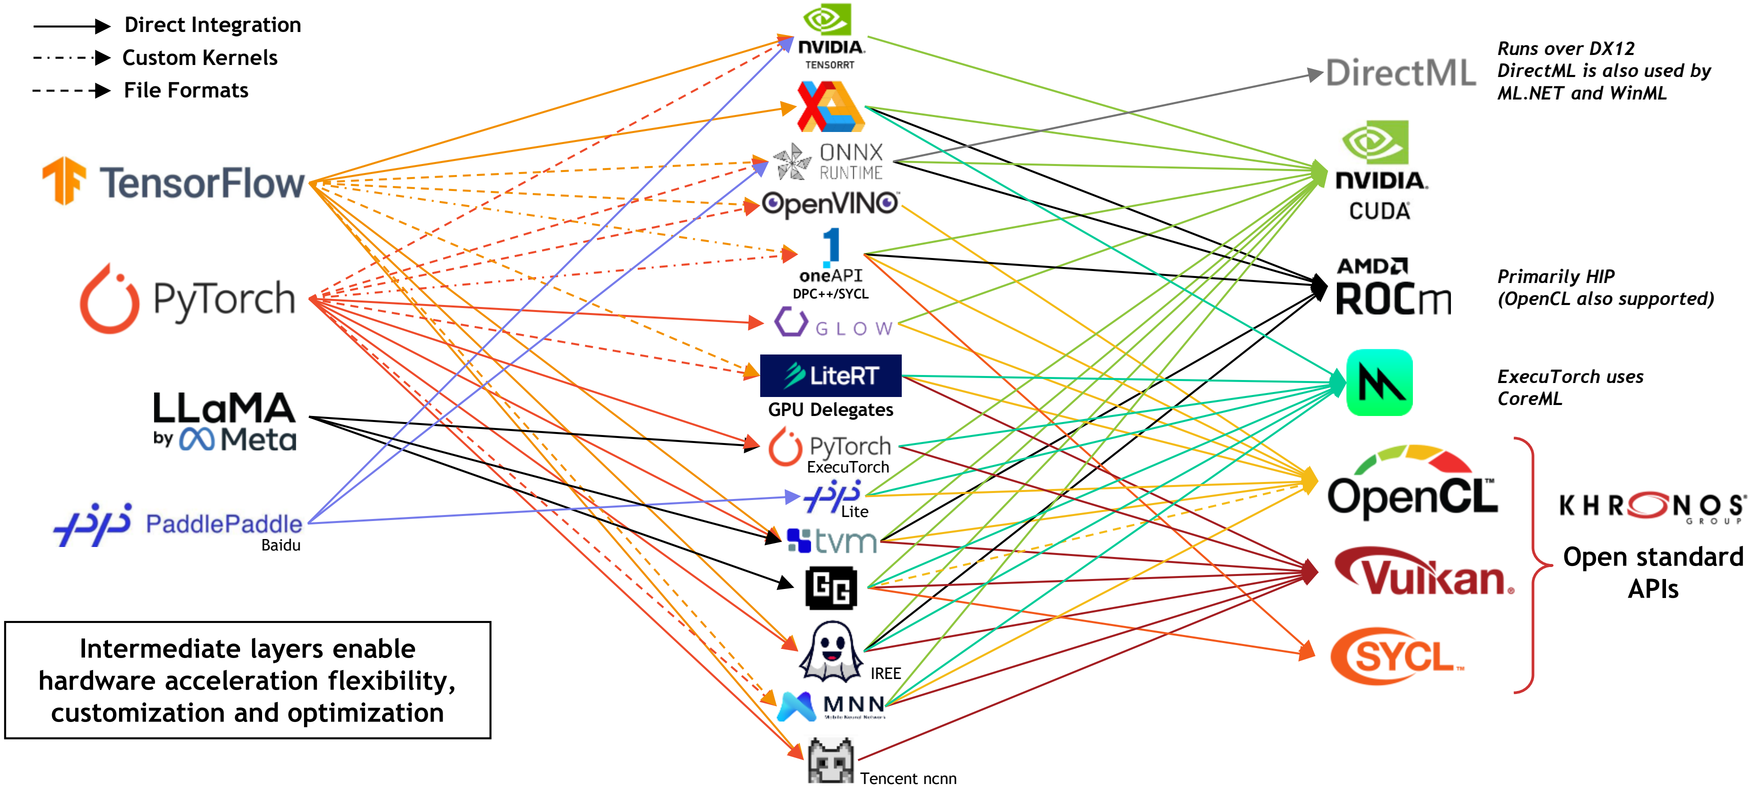

Runs over DX12 DirectML is also used by ML.NET and WinML

Primarily HIP (OpenCL also supported)

ExecuTorch uses CoreML

Open standard APIs

Intermediate layers enable hardware acceleration flexibility, customization and optimization

Machine Learning Acceleration APIs

Open-Source Frameworks | Compilers, Runtimes and Libraries (some in open source) | Acceleration APIs

⟶ Direct Integration

-·-·-▶ Custom Kernels

- - - -▶ File Formats

Using open standard APIs to access cross-vendor acceleration can reduce ecosystem complexity and development costs for machine learning developers

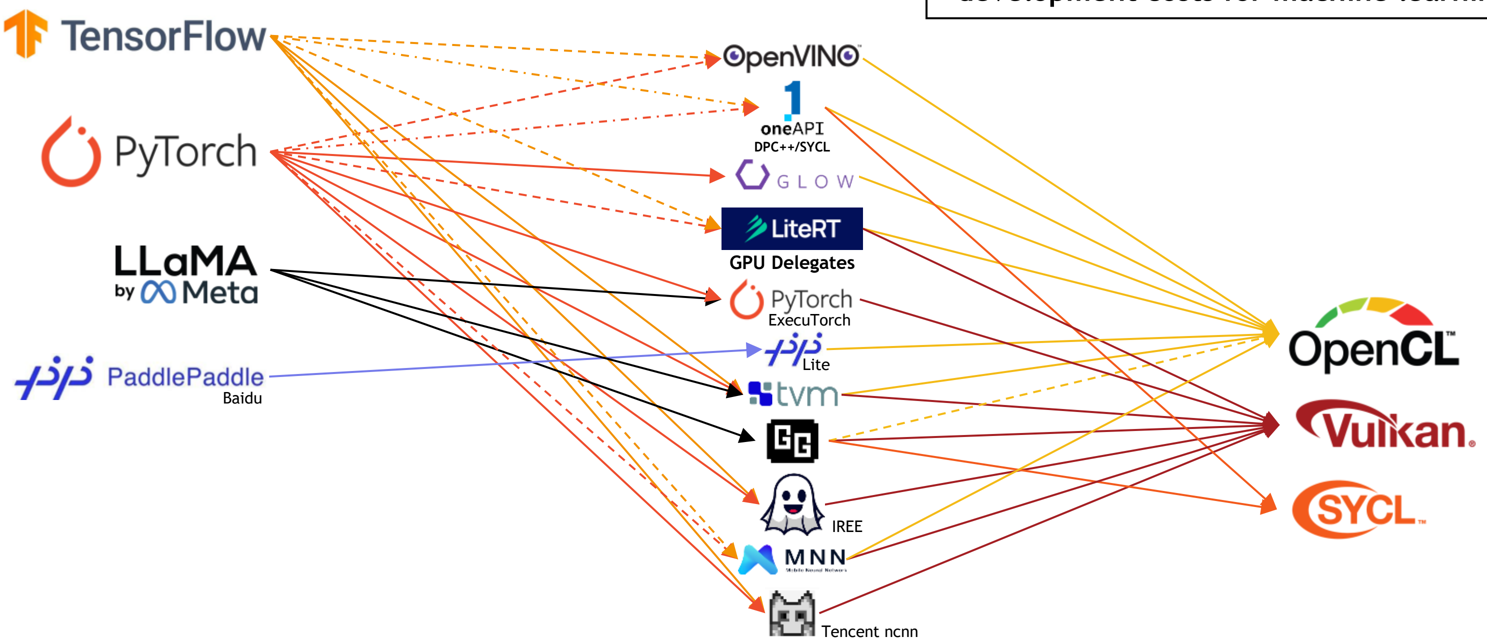

TensorFlow

PyTorch

LLaMA by Meta

PaddlePaddle Baidu

OpenVINO

oneAPI DPC++/SYCL

GLOW

LiteRT GPU Delegates

PyTorch ExecuTorch

Lite

tvm

IREE

MNN

Tencent ncnn

OpenCL

Vulkan

SYCL

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) 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 Acuity ML/AI Software Foundation	Runtime		OpenCL OpenVX
Xiaomi Mace MACE Mobile AI Compute Engine	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?

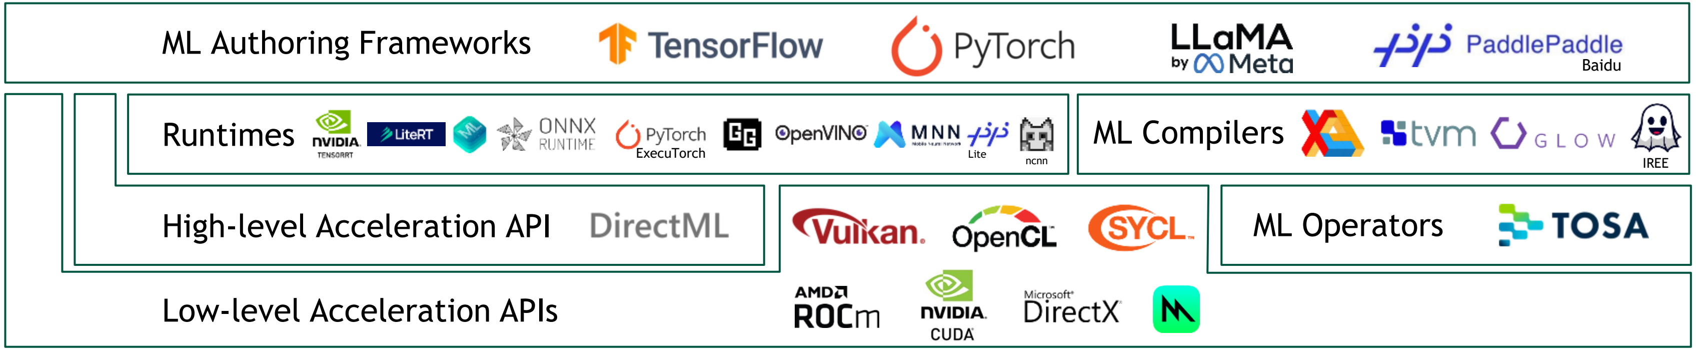

ML Authoring Frameworks: TensorFlow, PyTorch, LLaMA by Meta, PaddlePaddle Baidu

Runtimes: NVIDIA TensorRT, LiteRT, ONNX RUNTIME, PyTorch ExecuTorch, OpenVINO, MNN, Lite, ncnn

ML Compilers: tvm, GLOW, IREE

High-level Acceleration API: DirectML

Vulkan, OpenCL, SYCL

ML Operators: TOSA

Low-level Acceleration APIs: AMD ROCm, NVIDIA CUDA, Microsoft DirectX

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

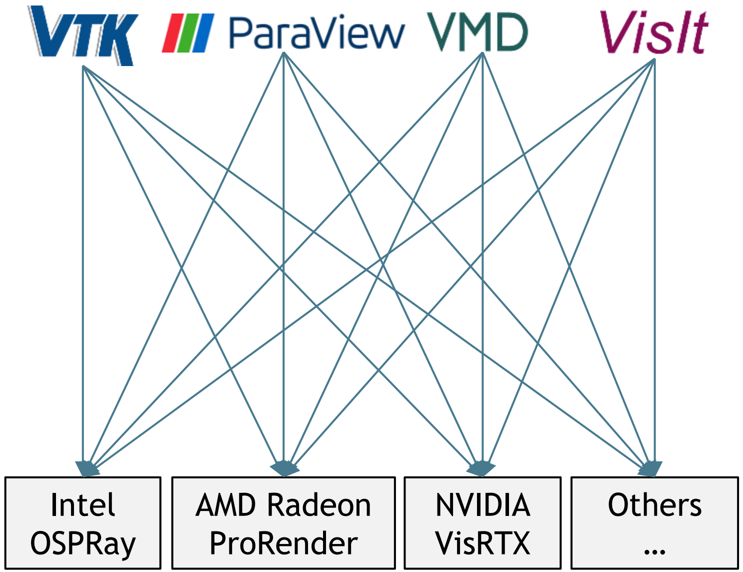

Before ANARI

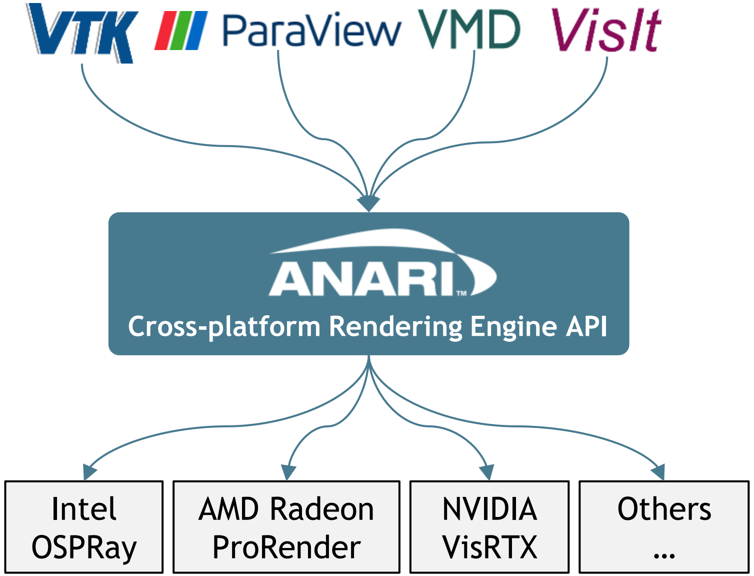

After ANARI

ANARI applications are portable to any engine supporting the ANARI API

Independently of vendor, platform or ecosystem

ANARI Scene Representation

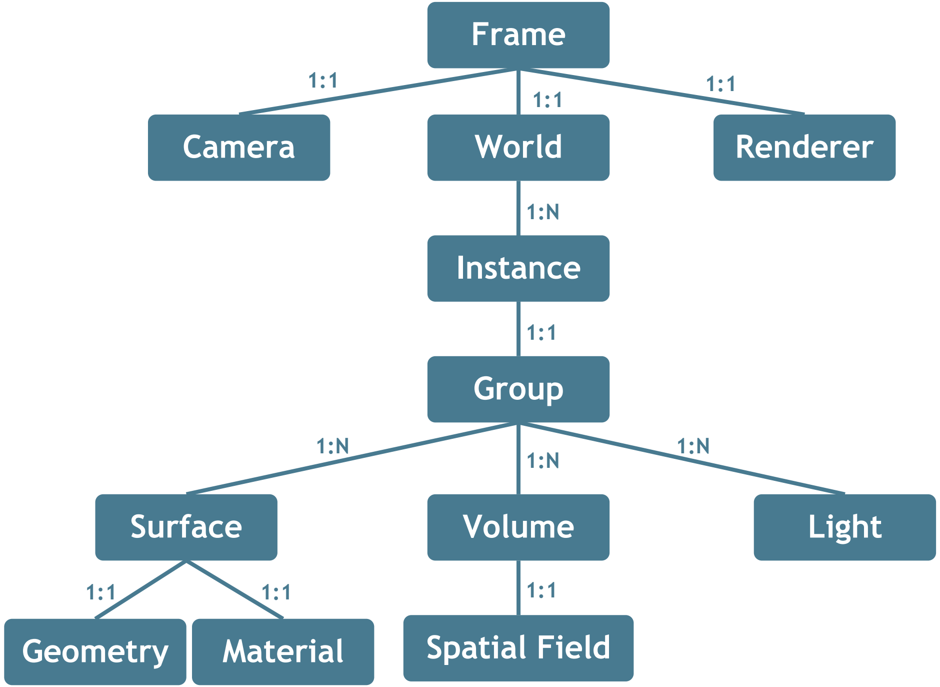

ANARI API used to build in-memory scene representation

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

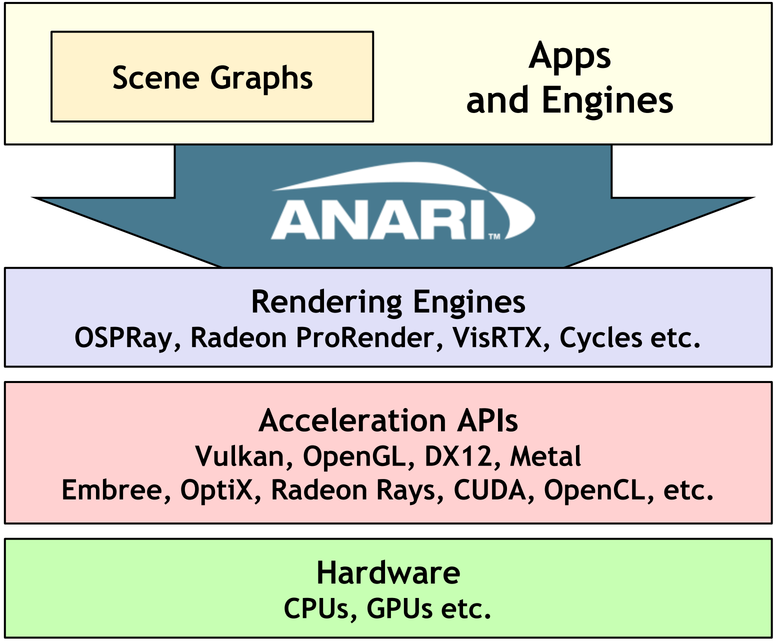

Scene Graphs

Apps and Engines

Rendering Engines
OSPRay, Radeon ProRender, VisRTX, Cycles etc.

Acceleration APIs
Vulkan, OpenGL, DX12, Metal
Embree, OptiX, Radeon Rays, CUDA, OpenCL, etc.

Hardware
CPUs, GPUs etc.

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

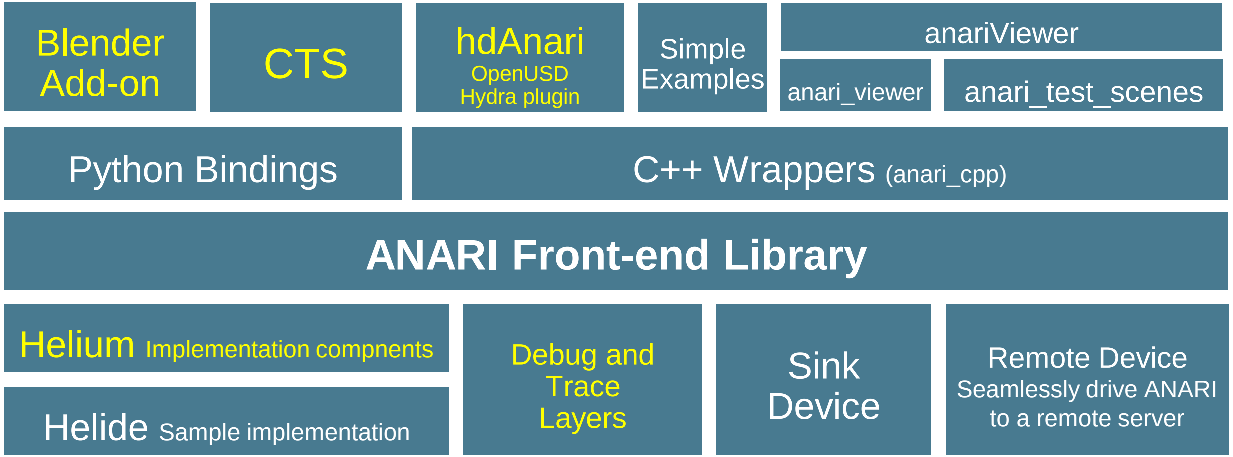

Blender Add-on

CTS

hdAnari OpenUSD Hydra plugin

Simple Examples

anariViewer

anari_viewer

anari_test_scenes

Python Bindings

C++ Wrappers (anari_cpp)

ANARI Front-end Library

Helium Implementation compnents

Helide Sample implementation

Debug and Trace Layers

Sink Device

Remote Device Seamlessly drive ANARI to a remote server

glTF & 3D Commerce

Neil Trevett

AUTODESK
3 MAX, M AYA, R RVT, F 360, I PRO
blender
SOLIDWORKS
SketchUp
3DEXCITE
Flutter
3DCoat
Rhinoceros
UU3D v3 (64-bit)
Houdini
Blockbench
RealityServer
NOMAD SCULPT
SelfCAD
Clara.io
SIEMENS NX
MODO
MARMOSET TOOLBAG
READY PLAYER ME
3DSLASH
EF EVE
HYPAR
SculptGL
MagicaVoxel
CINEMA 4D
VoxEdit
Creature
Daz3D
chaos V-Ray
Dust3D
Sg
Pt
KeyShot by Luxion
Cheetah3d model • animate • render
Titania
Dn
An
Adobe
Twinmotion
Archilogic

3D Authoring Tools

NIANTIC
VECTARY
EVERGINE
Meta Spatial Editor
Ar
Adobe
ptc vuforia studio
VNTANA

VR / AR Authoring Tools

scandy
KIRI Engine
polycam
Metascan
3D Scanner
Trnio
QLONE
3dMD
eCapture3D

3D Scanning Tools

SIMPLYGON
RapidCompact by DGG
gltfpack
Reality Convertor
Assimp Open Asset Import Library
PiXYZ SOFTWARE
SAFE SOFTWARE
cgltf
CrossManager
Collada2gltf
OBJ2GLTF
FBX2glTF

Convertors, Optimizers and Loaders

gltf-vscode
AGI
KHRONOS
glTF Sample Viewer
glTF-validator
glTF-Toolkit
Microsoft
glTF-asset-generator

Validation and Reference Tools

Adobe
cgtrader
TURBOSQUID by shutterstock
KENNEY
Sketchfab
p3d.in

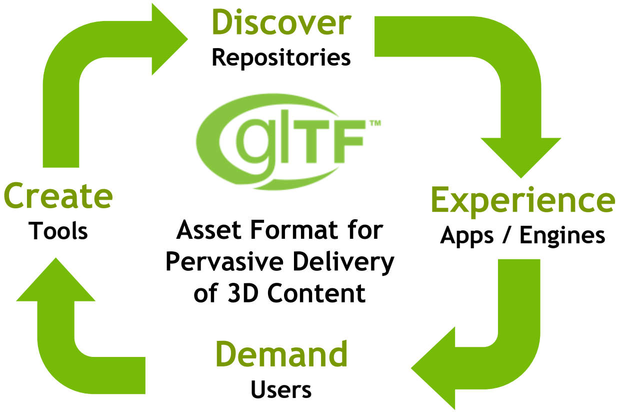

TARGET
Bentley
Open Geospatial Consortium
otoy
IKEA
EA
AMD
3DTiles
Metaverse STANDARDS FORUM
web3D CONSORTIUM
amazon.com
wayfair
Uber
Meta
Continental
shopify
SuperDNA 3D Lab
VRM CONSORTIUM
The New York Times
magic leap
MPEG
adidas
Smithsonian

Unity
UNREAL ENGINE
THE FORGE
O3DE OPEN 3D ENGINE
GODOT Game engine
ARMORY3D
OGRE
JMONKEYENGINE
Diligent Graphics
The G3D Innovation Engine

Game Engines

babylon.js
three.js
PLAYCANVAS
xeogl
PHASMATIC
model-viewer
Threedium

Web Engines

ROBLOX
MINECRAFT
CESIUM
Microsoft Flight Simulator
T I N K E R C A D
AUTODESK Platform Services
STK
KDAB
VTK
NVIDIA OMNIVERSE
xeogl
Qt
ArcGIS CityEngine
COMSOL
instant3Dhub
GESTALTOR
threekit
VENTUZ
seek
emersya
Filament
ParaView

Apps and Engines

WebXR
horizon Worlds
VRCHAT
SECOND LIFE
arvizio
THE WILD
METAVRSE
MESHROOM3D
worldviz
ARCore
NIANTIC 8thWall
AUCTA
ETHAR The XR Experience Platform

VR / AR Apps and Engines

Office
WORDPRESS
PDF

Productivity and Social Apps

KHRONOS GROUP

glTF PBR Materials Roadmap

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

Clearcoat

Volume

Sheen

Index of Refraction

Emissive Strength

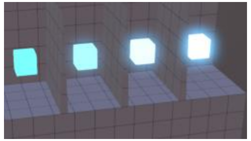

Dispersion

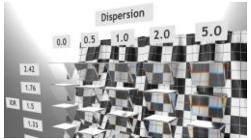

Subsurface
In development

Metal / Roughness

Transmission

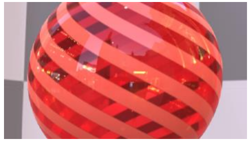

Specular

Iridescence

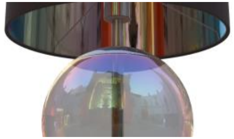

Anisotropy

2017 2020 2021 2022 2023/4

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

ASWF /* ACADEMY SOFTWARE FOUNDATION

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.

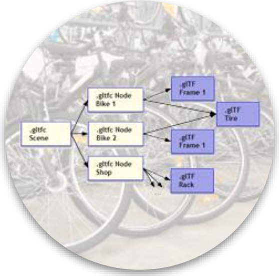

Complex Scenes

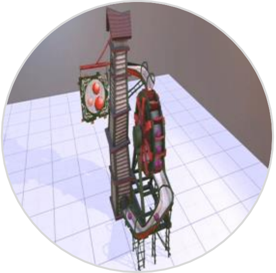

Physics

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

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

Additional Market Segments

**Foundation of
Core specification and
glTF working group extensions**

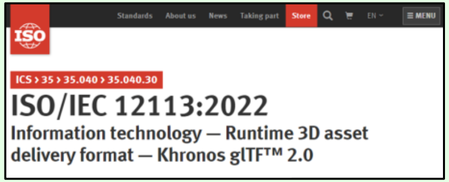

ICS > 35 > 35.040 > 35.040.30

ISO/IEC 12113:2022

Information technology — Runtime 3D asset delivery format — Khronos glTF™ 2.0

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

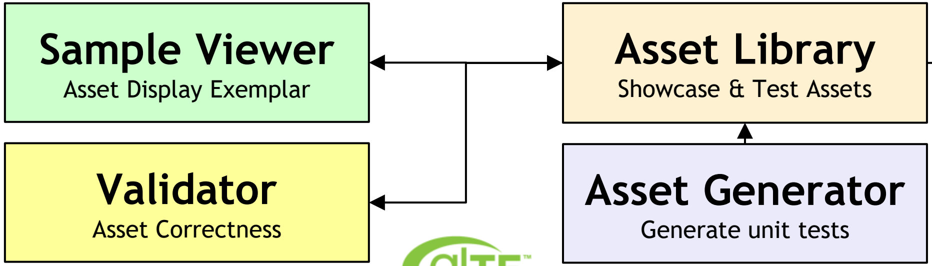

Sample Viewer
Asset Display Exemplar

Validator
Asset Correctness

Asset Library
Showcase & Test Assets

Asset Generator
Generate unit tests

Render Showcase
Web-based display comparisons

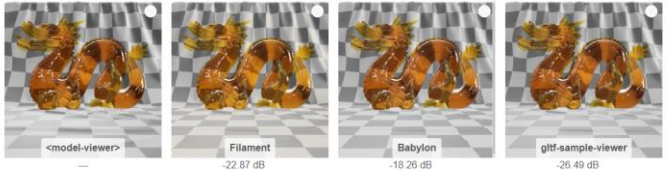

Khronos PBR Neutral Tone Mapper

- **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

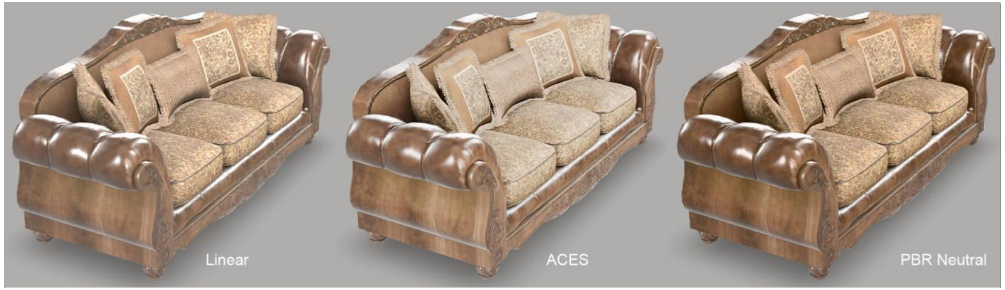

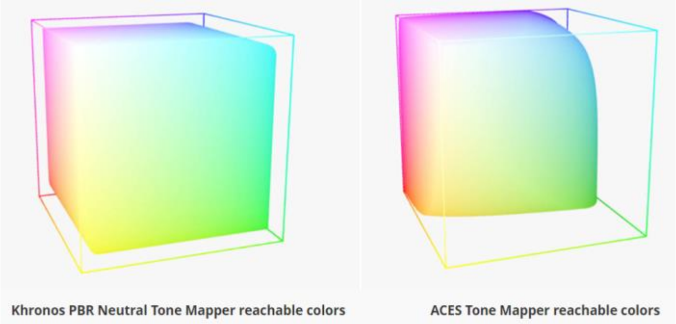

glTF Roadmap Overview

	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

glTF™

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

OpenUSD

AOUSD Alliance for OpenUSD

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

SIGGRAPH ASIA 2024 TOKYO 東京

KHRONOS GROUP

W3C® WebGL™ WebGPU glTF™ WebXR

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

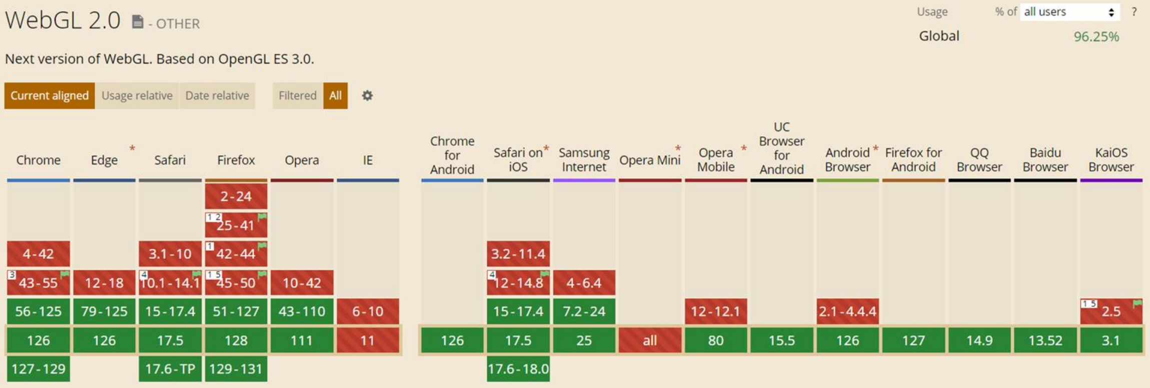

WebGL 2.0 - OTHER

Usage % of all users

Global 96.25%

Next version of WebGL. Based on OpenGL ES 3.0.

Chrome	Edge *	Safari	Firefox	Opera	IE	Chrome for Android	Safari on iOS *	Samsung Internet	Opera Mini *	Opera Mobile *	UC Browser for Android	Android Browser *	Firefox for Android	QQ Browser	Baidu Browser	KaiOS Browser
			2-24													
			[1,2] 25-41													
4-42		3.1-10	[1] 42-44				3.2-11.4									
[3] 43-55	12-18	[4] 10.1-14.1	[1,5] 45-50	10-42			[4] 12-14.8	4-6.4								
56-125	79-125	15-17.4	51-127	43-110	6-10		15-17.4	7.2-24		12-12.1		2.1-4.4.4				[1,5] 2.5
126	126	17.5	128	111	11	126	17.5	25	all	80	15.5	126	127	14.9	13.52	3.1
127-129		17.6-TP	129-131				17.6-18.0									

Khronos and W3C - XR Cooperation

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

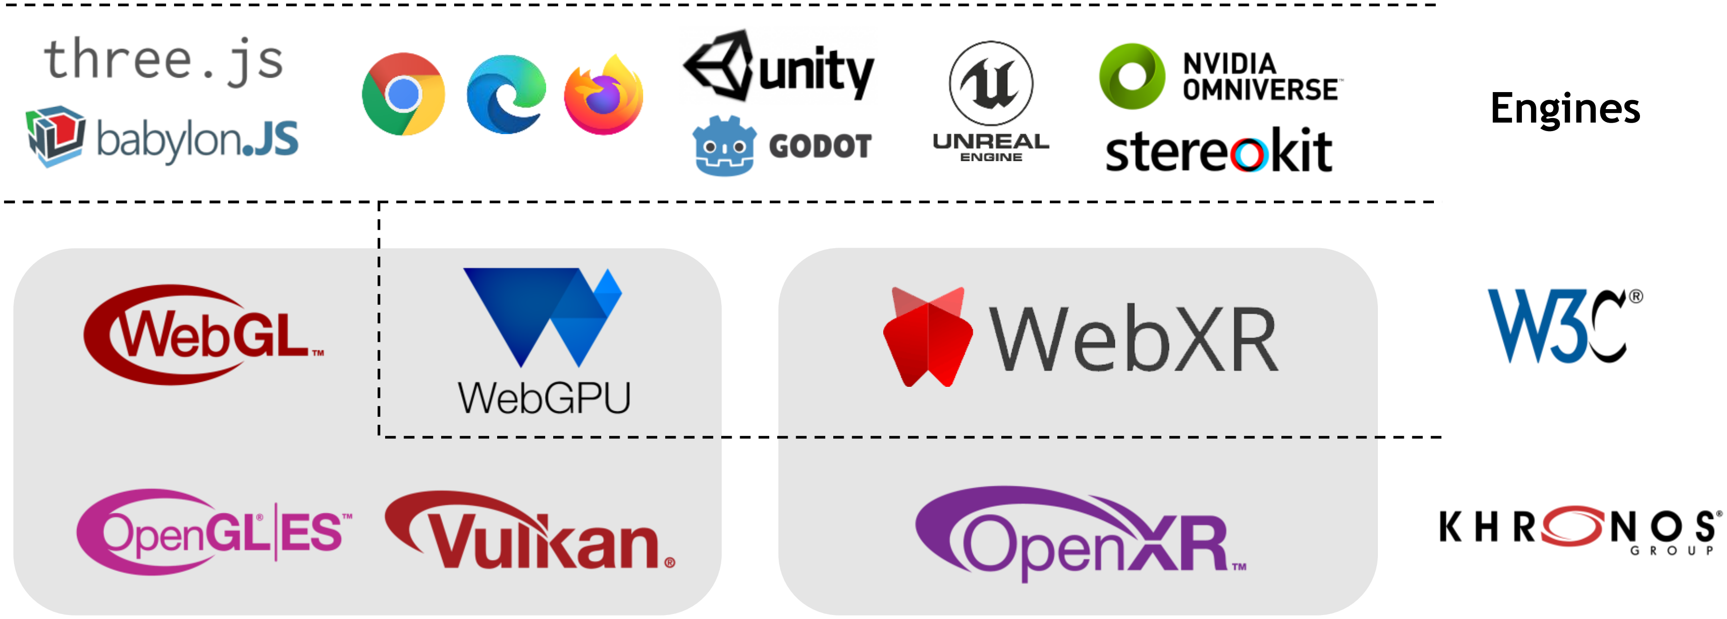

Engines

3D Stack
Driving GPUs to render scenes and augmentations

XR Stack
Handling XR Devices for creating UI

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

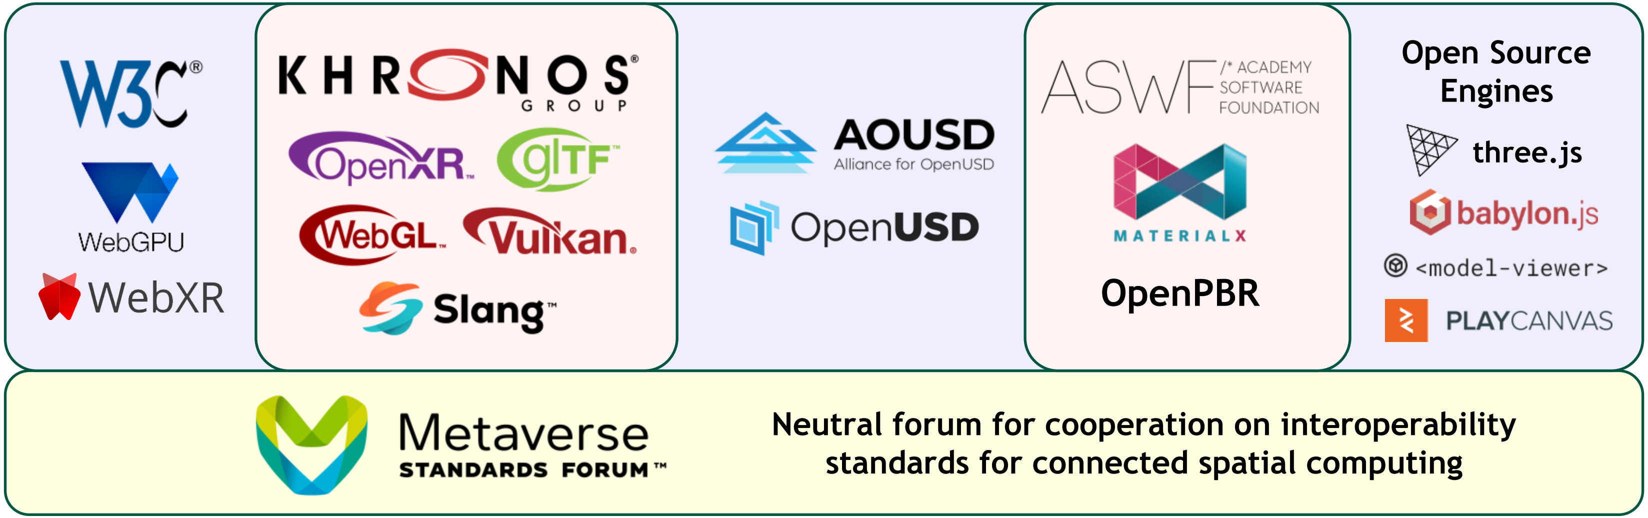

OpenPBR

Open Source Engines

Neutral forum for cooperation on interoperability standards for connected spatial computing

Vulkan: Forging Ahead!

Ralph Potter (Samsung), Vulkan Working Group chair

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 is Everywhere

Desktop and Mobile GPUs and SOCs

AMD BROADCOM HUAWEI Imagination intel NVIDIA MEDIATEK QUALCOMM arm Google SAMSUNG VeriSilicon

http://vulkan.gpuinfo.org/

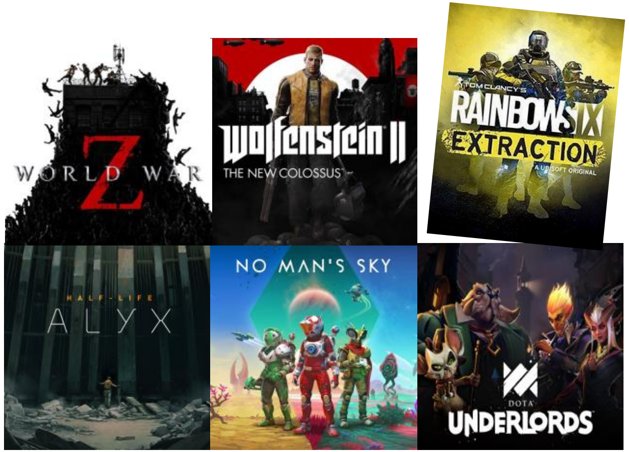

Desktop Games

Mobile Games

AUTODESK® FUSION 360®

Cross-platform post-processing and display of simulation results

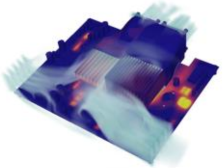

Adobe Sg

Substance 3D Stager
Cross-platform ray tracing

Applications

Engines

CIDER CINDER GODOT XENKO 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 is Unique

Vulkan

Vulkan is the *only open standard* modern GPU API

Under multi-company governance

Supported by all major GPU vendors

Cross-platform support reduces developer porting effort and costs

Used extensively by games and applications

Windows 10

Windows 11

Windows and Linux Desktops and Cloud

android

Mobile

NVIDIA GEFORCE NOW

Game Streaming Platforms

SteamOS

NINTENDO SWITCH

Gaming Platforms

iOS

macOS

Apple Platforms
(via translation layer)

KHRONOS GROUP

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

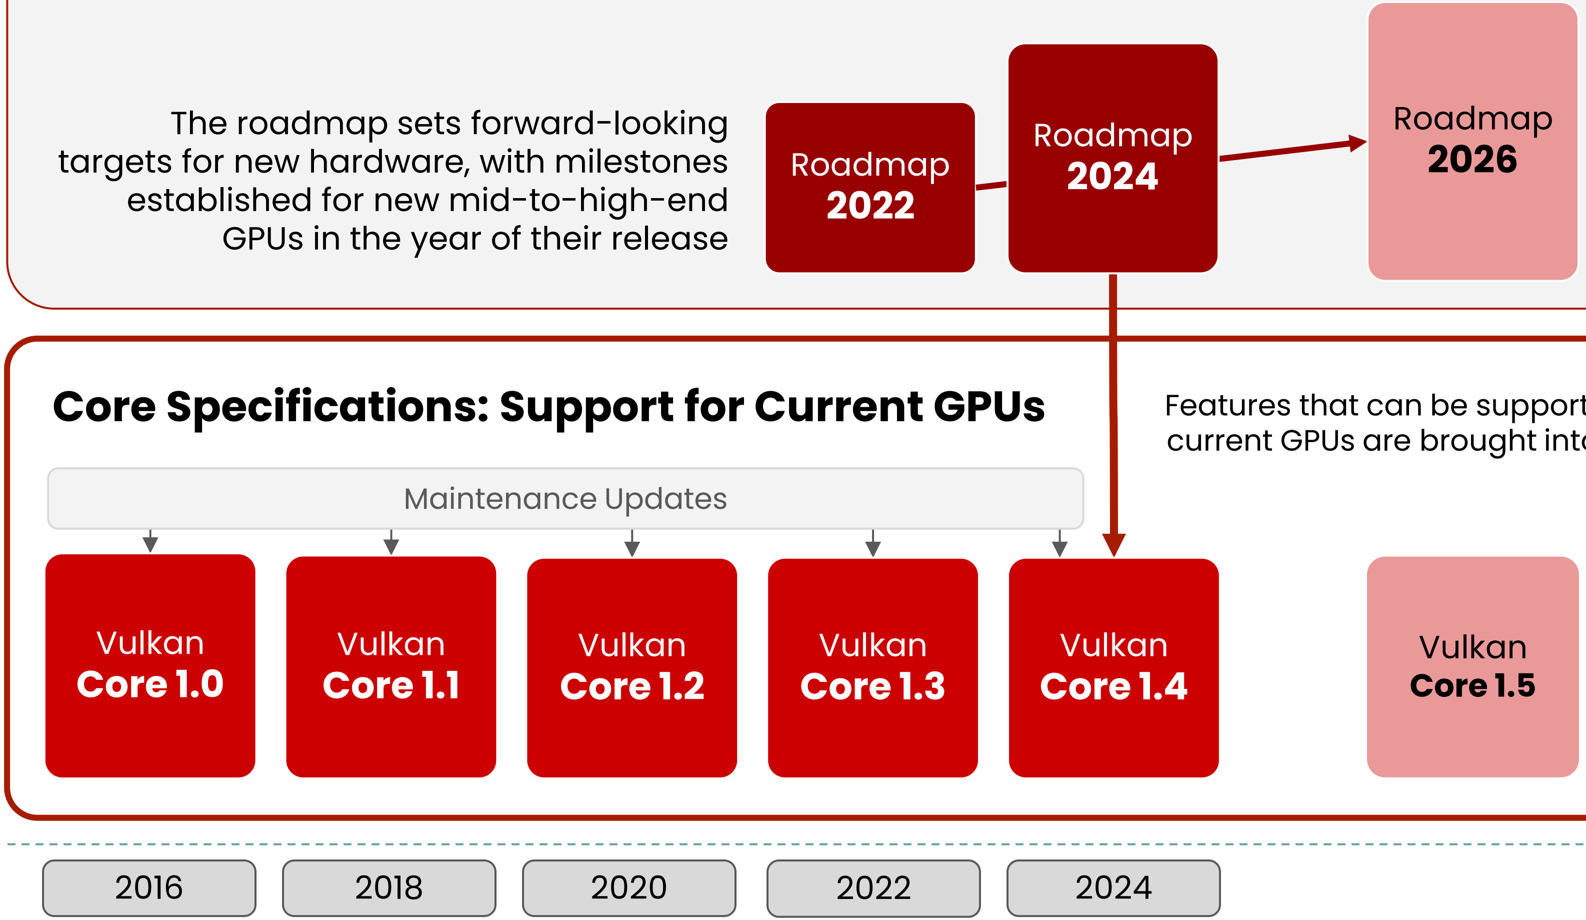

Roadmap **2022**

Roadmap **2024**

Roadmap **2026**

Core Specifications: Support for Current GPUs

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

Roadmap Sets Direction, Core Solidifies It

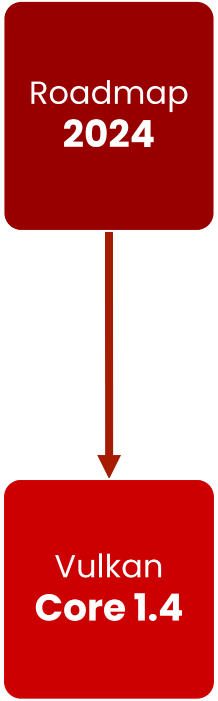

Roadmap **2024**

Vulkan **Core 1.4**

- **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 in total

Raising the Bar

Vulkan 1.0 was designed to run on GLES 3.1-class GPUs (circa 2014)

Core versions need to run on the broadest set of devices

Vulkan 1.4 raises minimum hardware requirements

Optional functionality and artificially low limits increase complexity for developers

Makes 28 previously optional features mandatory, including scalar block layout and 8/16 bit integer support in shaders

Raises minimum limits on 31 properties

Provides reliable access to functionality across all supported platforms

Streaming Transfers

Streaming image resources without interrupting rendering

Previously required copies on GPU timeline

VK_EXT_host_image_copy (optionally) promoted to core

Enables CPU-side image copies

If host copy is not supported, then a dedicated asynchronous transfer queue is mandatory

https://www.khronos.org/blog/copying-images-on-the-host-in-vulkan

Dynamic Rendering Local Read

Vulkan 1.3 promoted dynamic rendering to core…

- Removed the need for render pass and framebuffer objects
- Greatly simplified the programming model

…but the original extension didn’t address input attachments or subpasses

- Critical for performance on tile-based GPUs

Vulkan 1.4 includes local reads for color attachments/storage resources

- Closes the gap versus legacy render passes
- Local reads for depth/stencil/multisampled attachments are optional

https://www.khronos.org/blog/streamlining-subpasses

Vulkan 1.4 Release Schedule

- **Specification - Available Now**
  - https://docs.vulkan.org
- **API Headers - Available Now**
  - https://github.com/KhronosGroup/Vulkan-Headers
- **Conformance Tests - Available Now**
  - https://github.com/KhronosGroup/VK-GL-CTS
- **Vulkan Loader - Available Now**
  - https://github.com/KhronosGroup/Vulkan-Loader
- **Vulkan Validation Layers - MR in review**
  - https://github.com/KhronosGroup/Vulkan-ValidationLayers/pull/8955
- **Complete Vulkan SDK release - January 2025**
  - https://www.lunarg.com/vulkan-sdk
- **Implementations**
  - AMD, Arm, Imagination, Intel, NVIDIA, Qualcomm, and Samsung passing conformance today
  - NVIDIA and Mesa drivers publicly available today
  - Other vendors coming soon

Releasing Today!

https://khr.io/vulkan14

Slang Fast Forward

Shannon Woods
NVIDIA, Slang Working Group Chair

The Shading Language Landscape Today

- Shader codebases have become incredibly large & complex
- Developers need to deploy to many platforms
- Shader combinatorial explosion
- New graphics techniques & neural graphics discontinuity
- GLSL no longer innovating new language features

Open-Source, Cross-Platform Compiler

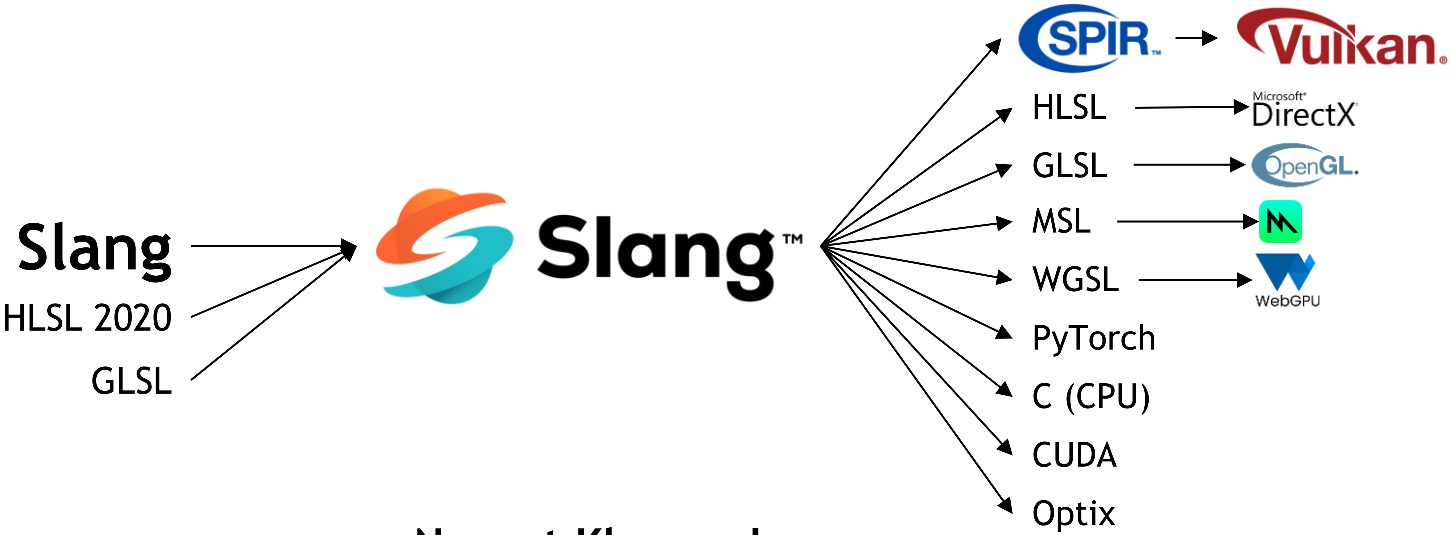

Slang

HLSL 2020

GLSL

Slang™

HLSL

GLSL

MSL

WGSL

PyTorch

C (CPU)

CUDA

Optix

Now at Khronos!

Slang + Khronos = Developers Win

- Shading language diversity means more competition & innovation
- No single company controls the language, so it can evolve as developers need

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

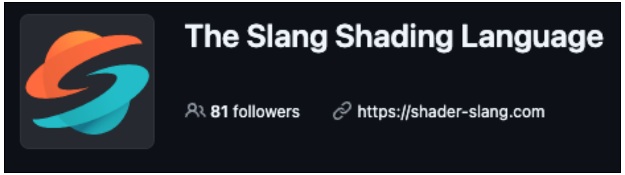

The Slang Shading Language

81 followers https://shader-slang.com

Why Another Shading Language?

	GLSL	MSL	WGSL	HLSL	Slang™
Actively Evolving	NO	YES	YES	YES	YES
Modular Code Management	NO	NO	NO	NO	YES
Converging with C++	NO	YES	NO	YES	NO*
Auto-diff / Neural Shading	NO	NO	NO	NO	YES
Diverse Backend Targets	NO	NO	NO	DXIL and SPIR-V	YES
Open-Source Compiler(s)	YES	NO	YES	YES	YES
Open Governance	YES	NO	YES	NO	YES

* Slang and HLSL are taking complementary evolutionary paths
HLSL will remain and evolve as a critically important shading language for many developers
Language diversity and choice is good for the graphics ecosystem!

Language Evolution

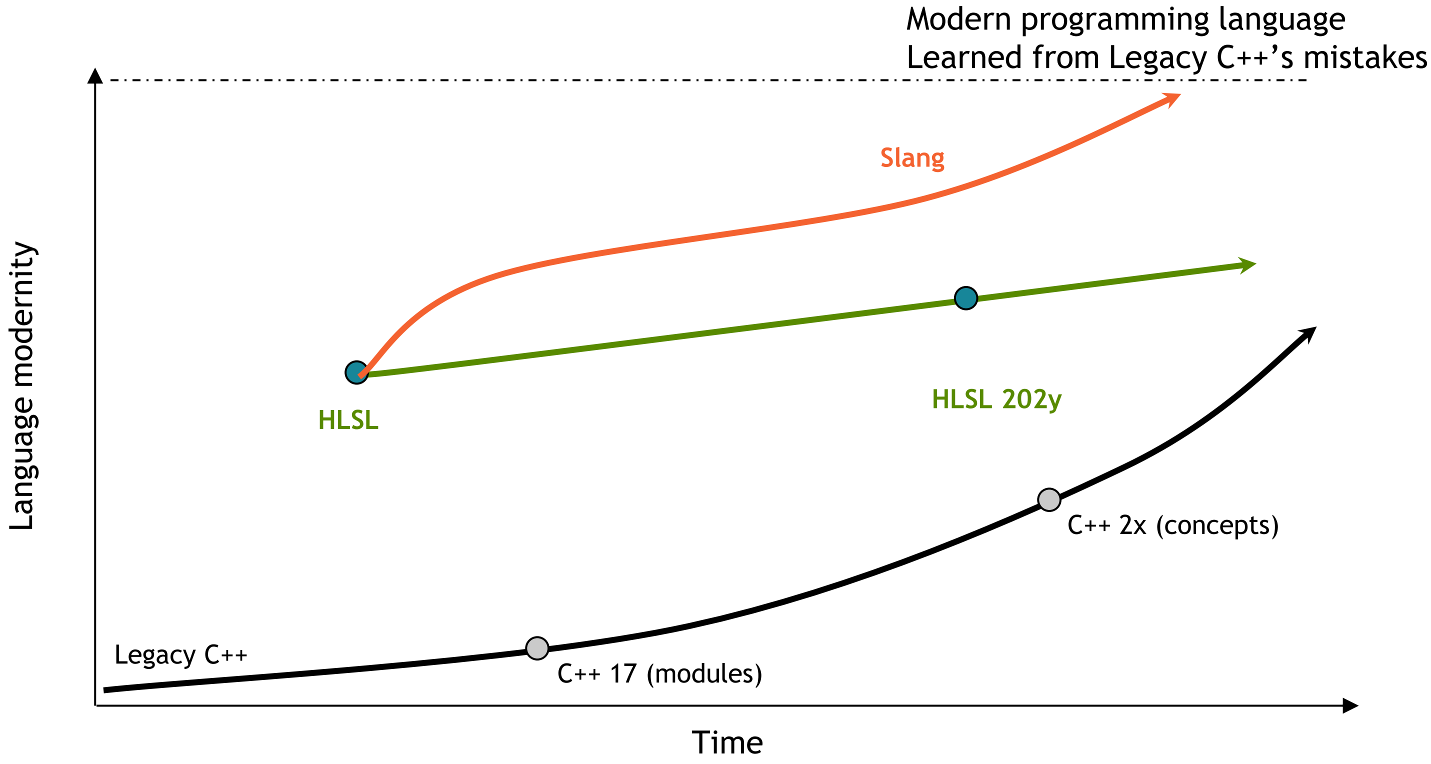

What makes Slang special?

- **Cross-compilation in Slang is easy and ergonomic**
  - a seamless way – integrated in one place
  - tooling just works
- **Automatic differentiation**
  - unique among shading languages
  - starting to show up as a necessity
  - for most AI graphics work
- **Proper solve for modularity, permutations, compile time explosion, and "string pasting"**
  - Drawing on advances from the broader language space, Slang addresses these issues with modules, generics, and interfaces

Seamless Cross-Compilation in Slang

- No need to chain together multiple libraries

- Tooling just works

- Produces human-readable code

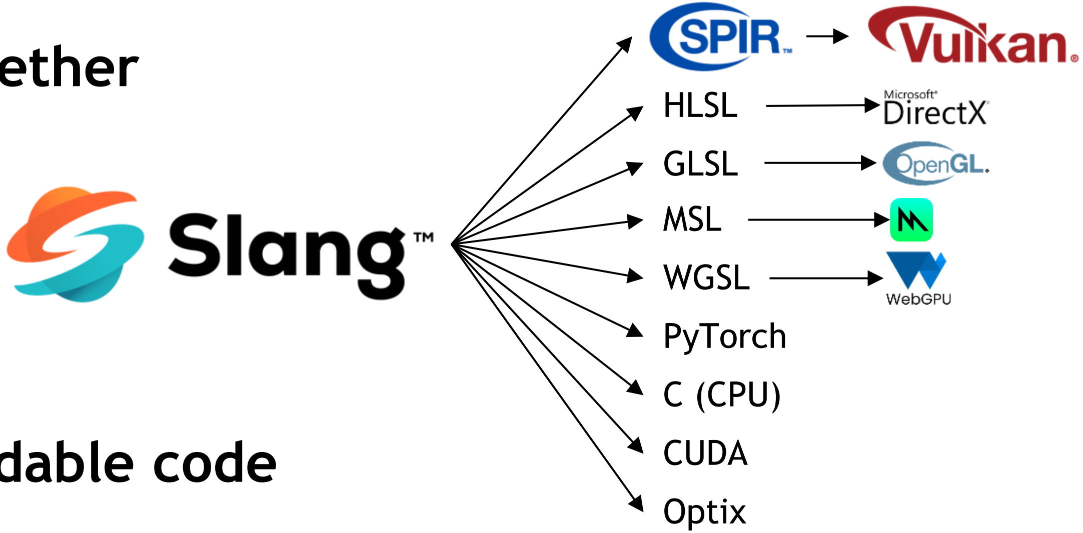

OpenXR Fast Forward

Jian Zhang

Head of XR Foundation Engineering, PICO

OpenXR Cross-Platform Portability

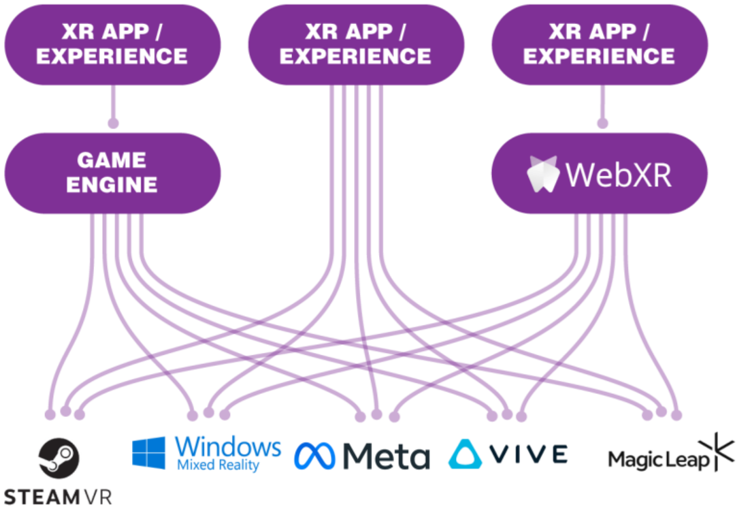

Before OpenXR: Applications and engines needed separate proprietary code for each device on the market.

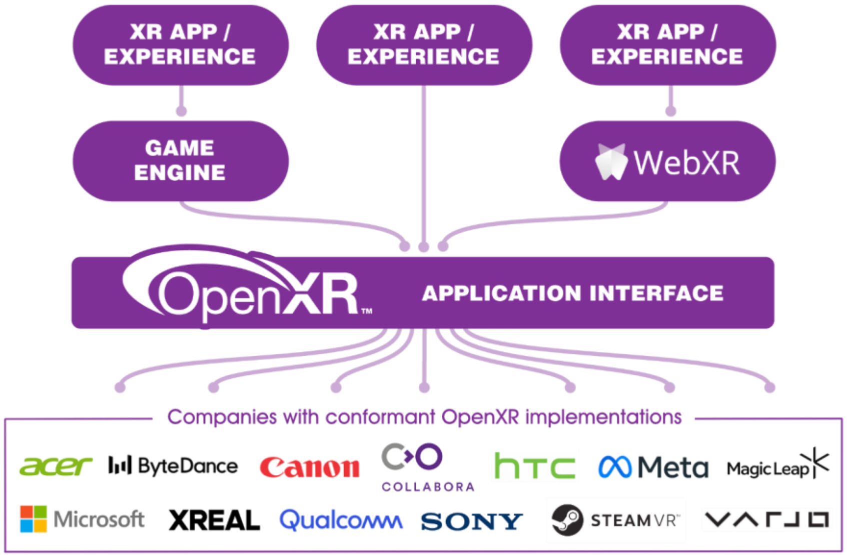

OpenXR provides a single cross-platform, high-performance API between applications and all conformant devices.

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
MagicLeap	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…

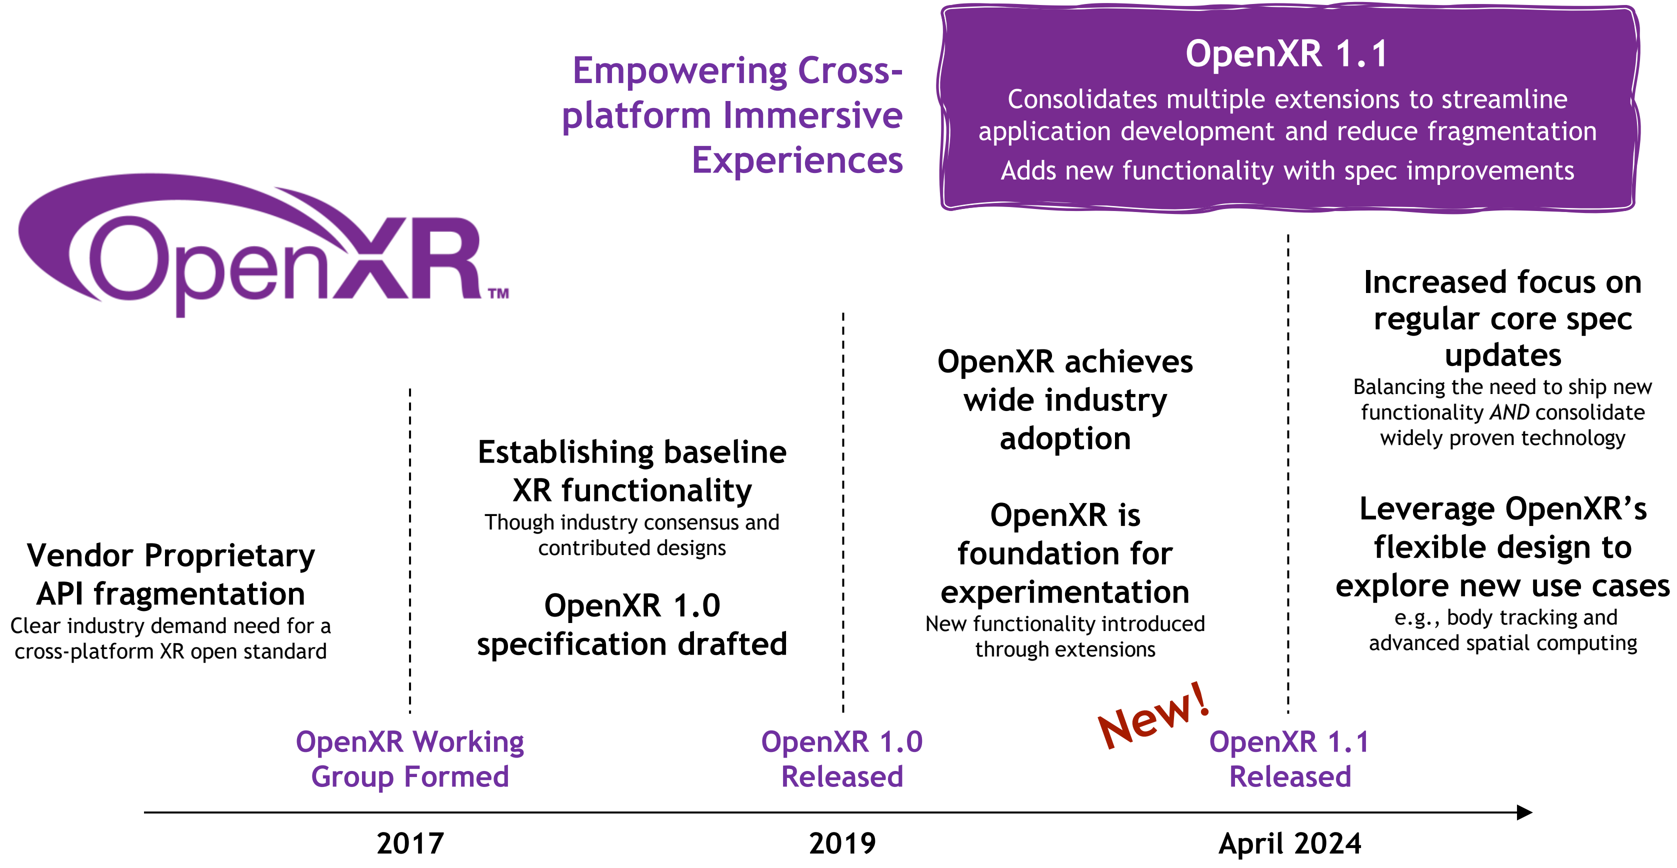

Empowering Cross-platform Immersive Experiences

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)*

The Value of Khronos Participation

Proven processes for building multi-company consensus for generation and governance of open interoperability standards

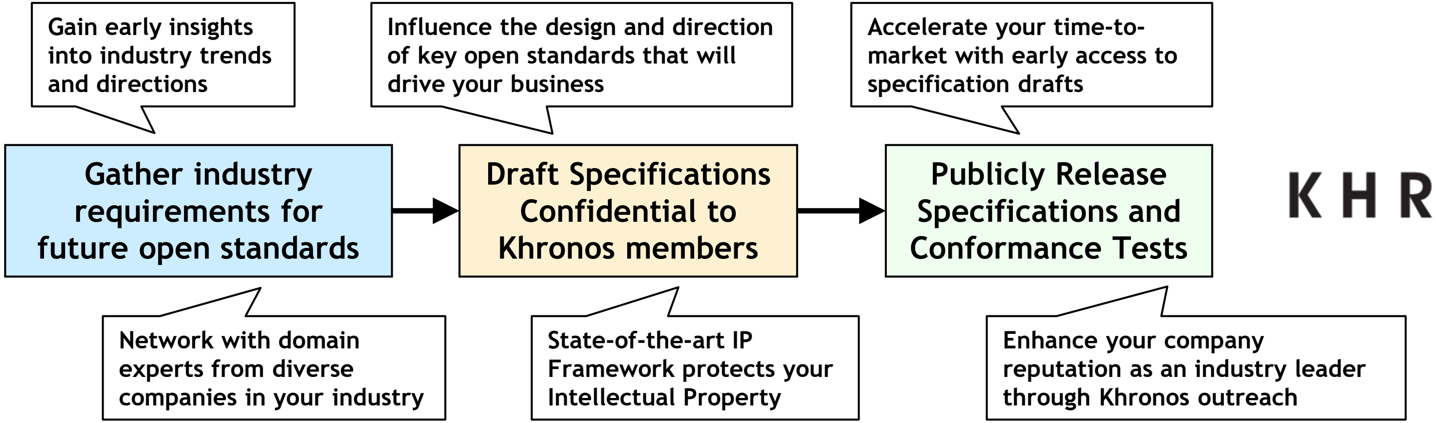

Gain early insights into industry trends and directions

Influence the design and direction of key open standards that will drive your business

Accelerate your time-to-market with early access to specification drafts

Gather industry requirements for future open standards

Draft Specifications Confidential to Khronos members

Publicly Release Specifications and Conformance Tests

Network with domain experts from diverse companies in your industry

State-of-the-art IP Framework protects your Intellectual Property

Enhance your company reputation as an industry leader through Khronos outreach

Khronos membership is open to any company for access to all standardization initiatives
Annual membership fees start at $4,000 for smaller companies
https://www.khronos.org/members/

KHRONOS® GROUP

東京

SIGGRAPH ASIA 2024 TOKYO

ENGAGE WITH KHRONOS:

Khronos IP Framework and Processes

April 2024

www.khronos.org

‘Traditional’ IP Frameworks

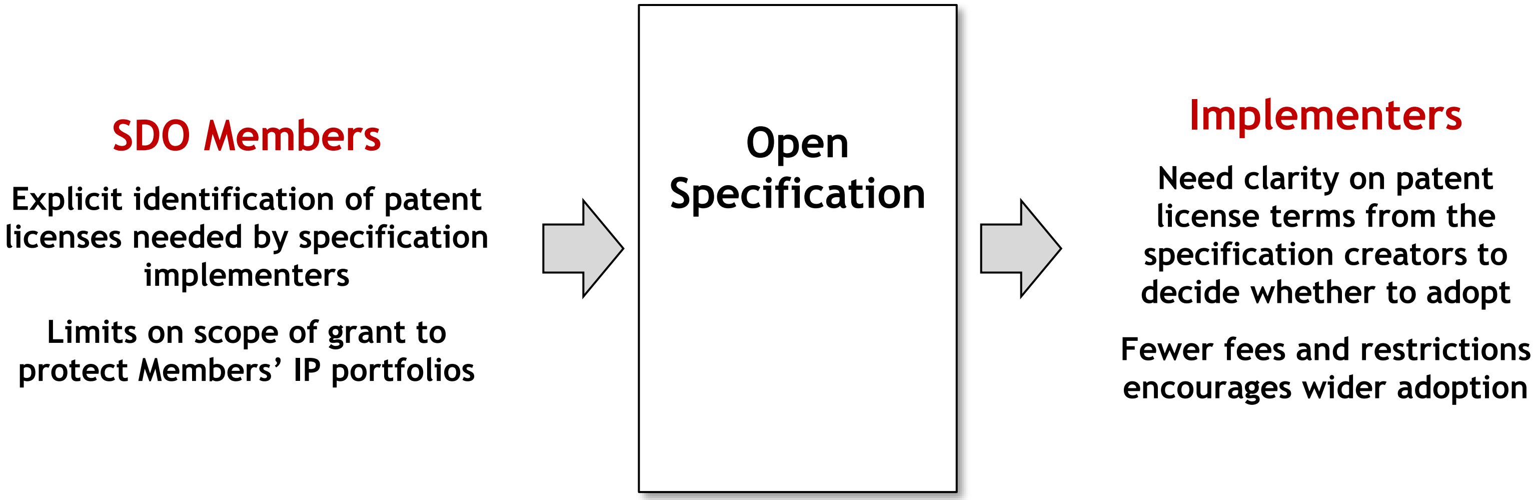

SDO Members

Explicit identification of patent licenses needed by specification implementers

Limits on scope of grant to protect Members’ IP portfolios

Open Specification

Implementers

Need clarity on patent license terms from the specification creators to decide whether to adopt

Fewer fees and restrictions encourages wider adoption

Typical Traditional IP Frameworks

SDO Members are asked to list essential patents they are aware of (with license terms for their own)
Some SDOs allow terms with royalties - but typically must at least be reasonable and non-discriminatory (RAND)
Implementers negotiate licenses with SDO Members individually or through a ‘patent pool’

Khronos IP Framework

Khronos Members

Agree to a ROYALTY-FREE reciprocal license to any essential patents they own for any CONFORMANT implementation of a ratified specification
(fail-safe and no patents need be disclosed)

License covers only the explicit contents of the specification - not other possible implementation technologies
(key to protecting member IP portfolios)

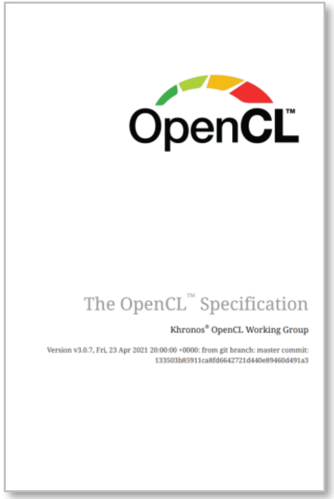

The OpenCL™ Specification

Khronos® OpenCL Working Group

Version v3.0.7, Fri, 23 Apr 2021 20:00:00 +0000: from git branch: master commit: 133503b85911ca8fd6642721d440e89460d491a3

Implementers

Any entity can use a Khronos specification with no trademark or patent licenses at any time

Formal Adopters are enabled to submit Conformance Test Results for trademark license and (optional) reciprocal patent licenses
(no negotiation with Khronos or Khronos Members is needed)

Explicit reciprocal patent license in Membership and Adopter Agreements

Enhances mutual protection and clarity

Builds network of licensing protection for the standard from Members *and* Adopters

Khronos IP Framework Flow

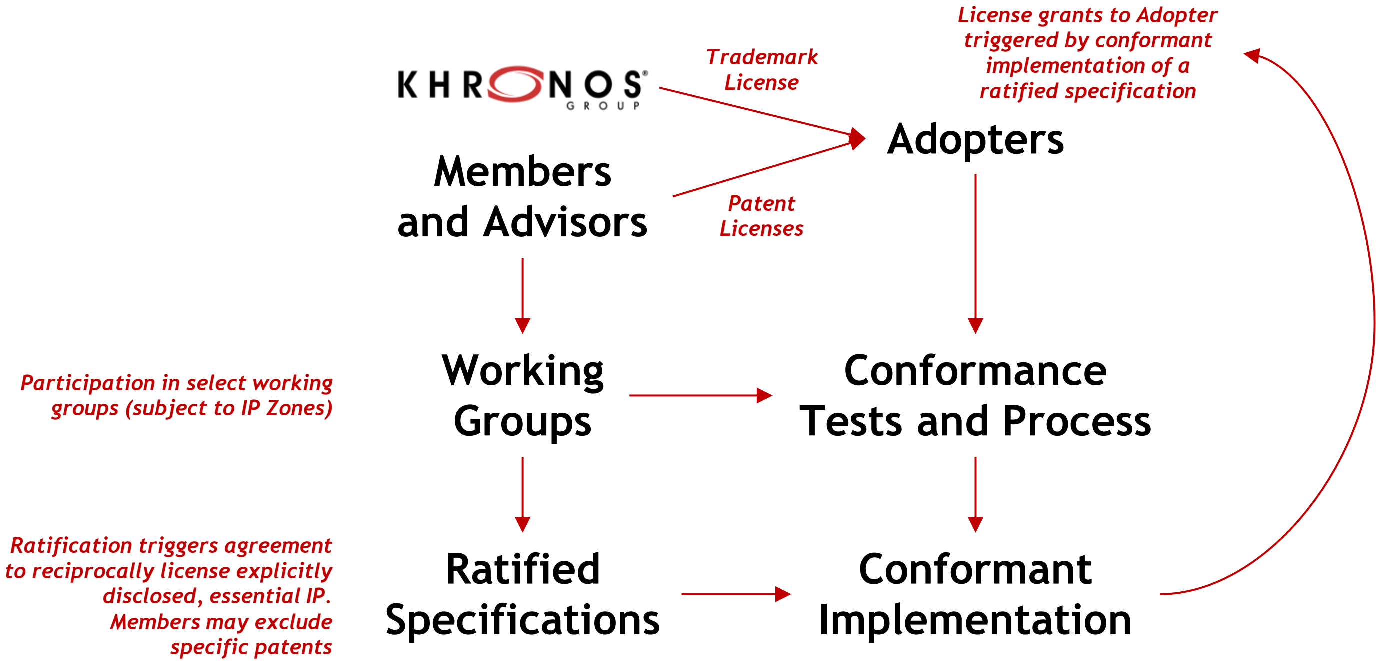

Members and Advisors

Trademark License

Patent Licenses

Adopters

License grants to Adopter triggered by conformant implementation of a ratified specification

Participation in select working groups (subject to IP Zones)

Working Groups

Conformance Tests and Process

Ratification triggers agreement to reciprocally license explicitly disclosed, essential IP. Members may exclude specific patents

Ratified Specifications

Conformant Implementation

Khronos - Balanced IP Protection

Khronos Members agree not to assert IP claims against Adopters for CONFORMANT IMPLEMENTATIONS OF RATIFIED Specifications

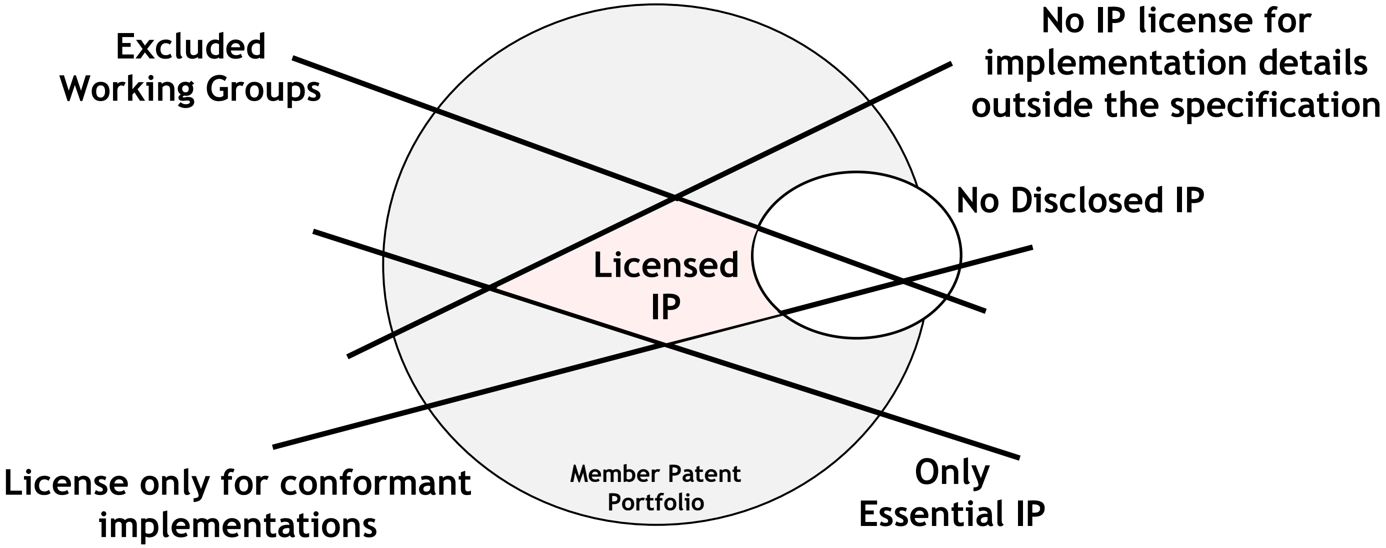
Excluded Working Groups
No IP license for implementation details outside the specification
No Disclosed IP
Licensed IP
License only for conformant implementations
Member Patent Portfolio
Only Essential IP

Scope of licensed IP is precisely and narrowly defined to protect Members IP BUT it is the IP needed to protect the specification for use in the industry

Khronos IP Zones

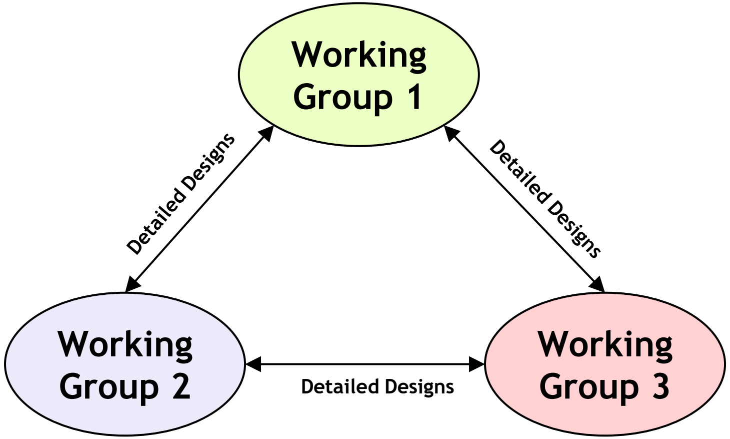

Many working groups need to share designs e.g.:

OpenGL and OpenGL ES
Vulkan and OpenGL
Vulkan, OpenCL and SPIR-V
glTF and 3D Commerce

“Submarined IP” Danger

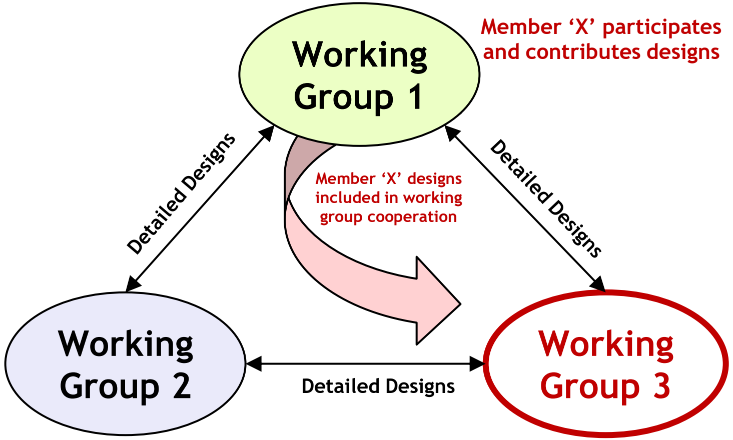

Working Group 1
Member ‘X’ participates and contributes designs
Detailed Designs
Member ‘X’ designs included in working group cooperation
Detailed Designs
Working Group 2
Detailed Designs
Working Group 3

Member ‘X’ Participates

Member ‘X’ Excludes from WG3
No IP Coverage for Member ‘X’ design contributions in specs created by WG3

Solution = IP Zones

Board maintained lists of working groups that share designs. Members should participate in ALL working groups within an IP Zone or none

IP Framework and OSS in Action: OpenVX

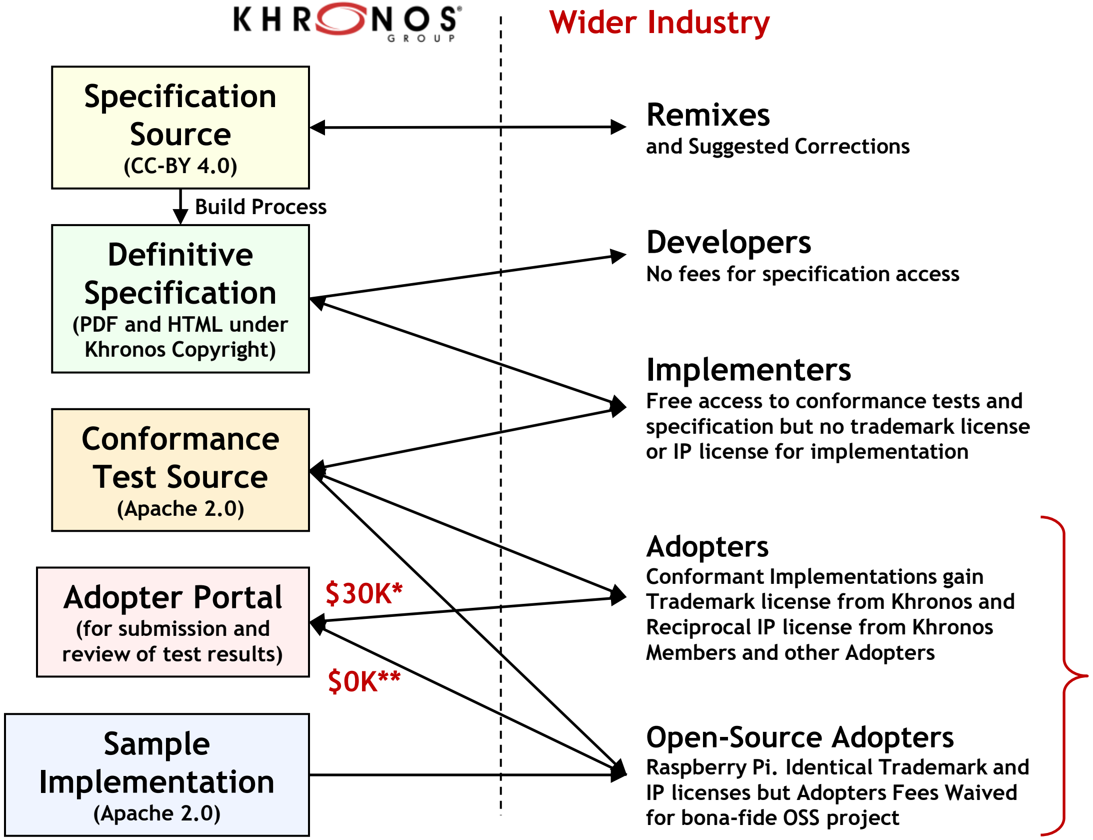

KHRONOS GROUP

Wider Industry

Specification Source (CC-BY 4.0)

Build Process

Definitive Specification (PDF and HTML under Khronos Copyright)

Conformance Test Source (Apache 2.0)

Adopter Portal (for submission and review of test results)

$30K*

$0K**

Sample Implementation (Apache 2.0)

Remixes
and Suggested Corrections

Developers
No fees for specification access

Implementers
Free access to conformance tests and specification but no trademark license or IP license for implementation

Adopters
Conformant Implementations gain Trademark license from Khronos and Reciprocal IP license from Khronos Members and other Adopters

Open-Source Adopters
Raspberry Pi. Identical Trademark and IP licenses but Adopters Fees Waived for bona-fide OSS project

Working Group also releases samples and other educational materials as Open Source

OpenVX

Essential patents from all participating Khronos Members are licensed, royalty-free, to provide protection to conformant implementations

* One-off Adopter Fee for unlimited number of products (partially covers Khronos costs)

** Adopters Fees waived for bona-fide open-source projects

KHRONOS®
GROUP

New Initiative Process

Khronos New Initiative Process Details

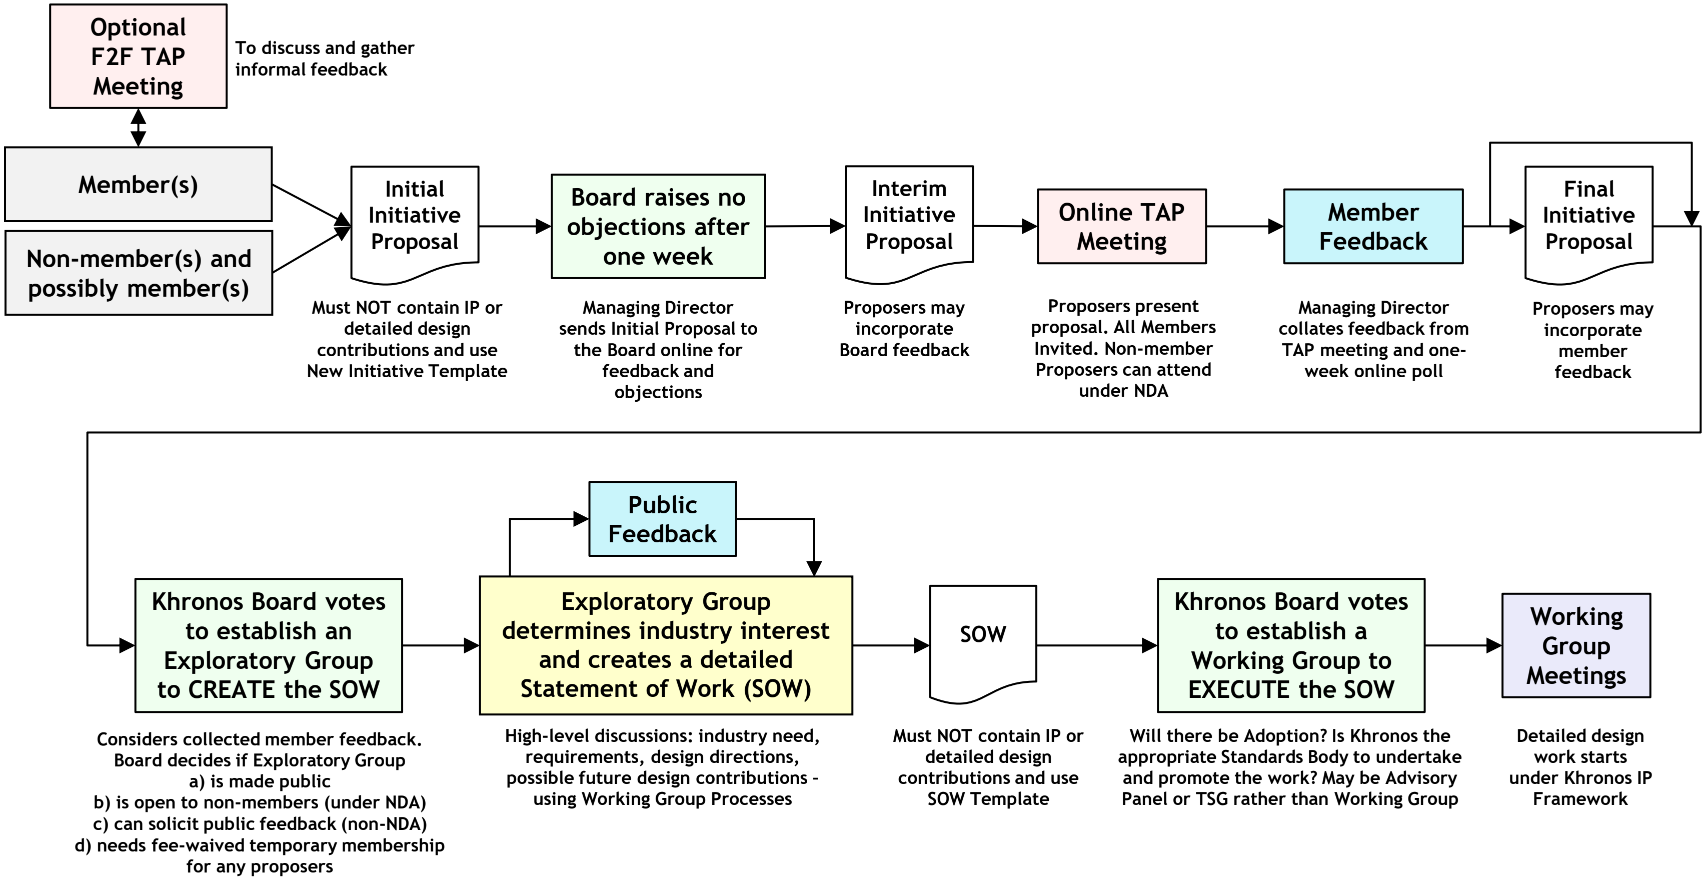

Optional F2F TAP Meeting

To discuss and gather informal feedback

Member(s)

Non-member(s) and possibly member(s)

Initial Initiative Proposal

Must NOT contain IP or detailed design contributions and use New Initiative Template

Board raises no objections after one week

Managing Director sends Initial Proposal to the Board online for feedback and objections

Interim Initiative Proposal

Proposers may incorporate Board feedback

Online TAP Meeting

Proposers present proposal. All Members Invited. Non-member Proposers can attend under NDA

Member Feedback

Managing Director collates feedback from TAP meeting and one-week online poll

Final Initiative Proposal

Proposers may incorporate member feedback

Khronos Board votes to establish an Exploratory Group to CREATE the SOW

Considers collected member feedback.
Board decides if Exploratory Group
a) is made public
b) is open to non-members (under NDA)
c) can solicit public feedback (non-NDA)
d) needs fee-waived temporary membership for any proposers

Public Feedback

Exploratory Group determines industry interest and creates a detailed Statement of Work (SOW)

High-level discussions: industry need, requirements, design directions, possible future design contributions - using Working Group Processes

SOW

Must NOT contain IP or detailed design contributions and use SOW Template

Khronos Board votes to establish a Working Group to EXECUTE the SOW

Will there be Adoption? Is Khronos the appropriate Standards Body to undertake and promote the work? May be Advisory Panel or TSG rather than Working Group

Working Group Meetings

Detailed design work starts under Khronos IP Framework

Conformance and Adoption

Khronos IP and Adopters Framework

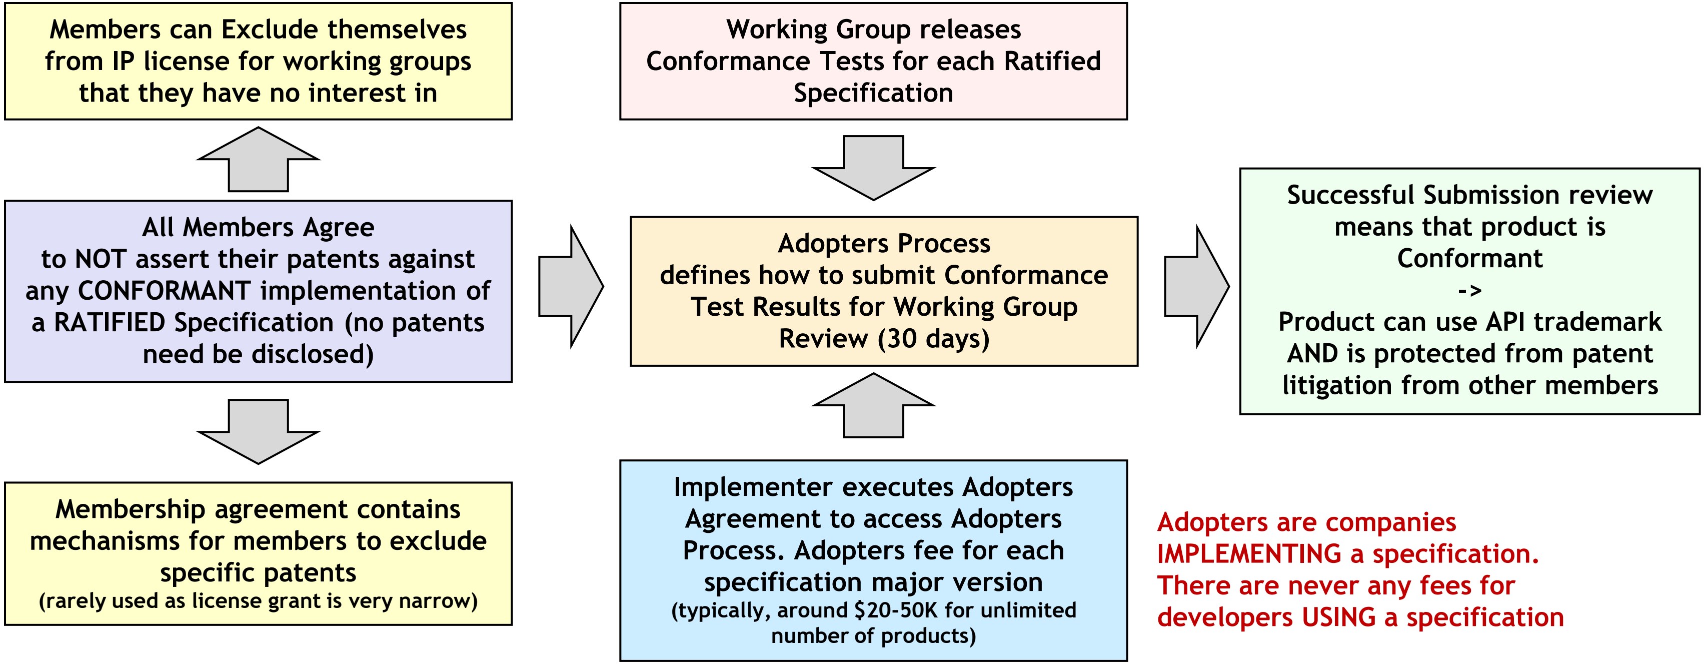

Members can Exclude themselves from IP license for working groups that they have no interest in

All Members Agree to NOT assert their patents against any CONFORMANT implementation of a RATIFIED Specification (no patents need be disclosed)

Membership agreement contains mechanisms for members to exclude specific patents (rarely used as license grant is very narrow)

Working Group releases Conformance Tests for each Ratified Specification

Adopters Process defines how to submit Conformance Test Results for Working Group Review (30 days)

Implementer executes Adopters Agreement to access Adopters Process. Adopters fee for each specification major version (typically, around $20-50K for unlimited number of products)

Successful Submission review means that product is Conformant -> Product can use API trademark AND is protected from patent litigation from other members

Adopters are companies IMPLEMENTING a specification. There are never any fees for developers USING a specification

Implementation, Conformance and Adoption

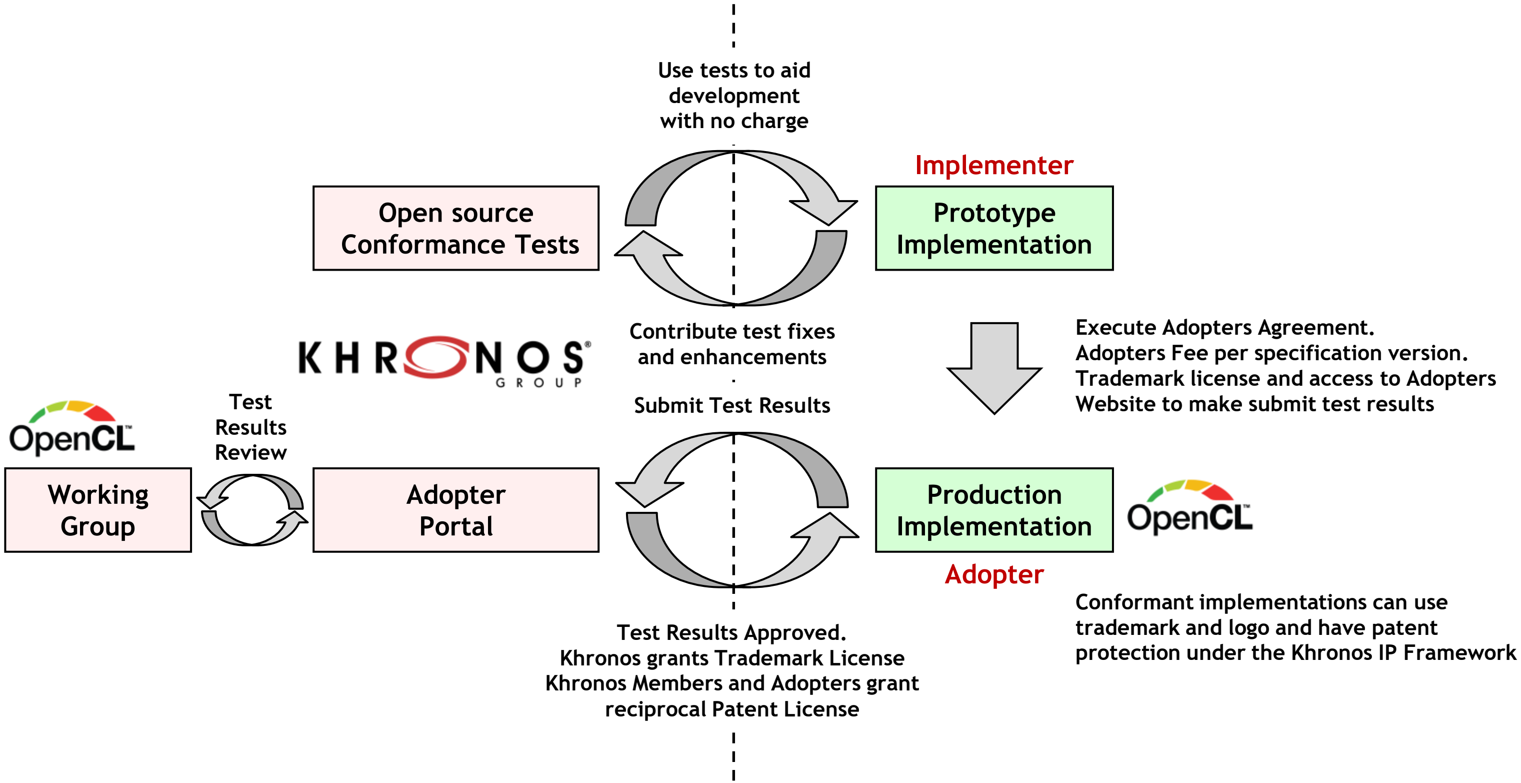

Use tests to aid development with no charge

Open source Conformance Tests

Implementer

Prototype Implementation

Contribute test fixes and enhancements

Execute Adopters Agreement. Adopters Fee per specification version. Trademark license and access to Adopters Website to make submit test results

Submit Test Results

Test Results Review

Working Group

Adopter Portal

Production Implementation

Adopter

Test Results Approved. Khronos grants Trademark License Khronos Members and Adopters grant reciprocal Patent License

Conformant implementations can use trademark and logo and have patent protection under the Khronos IP Framework

Implementation, Conformance and Adoption

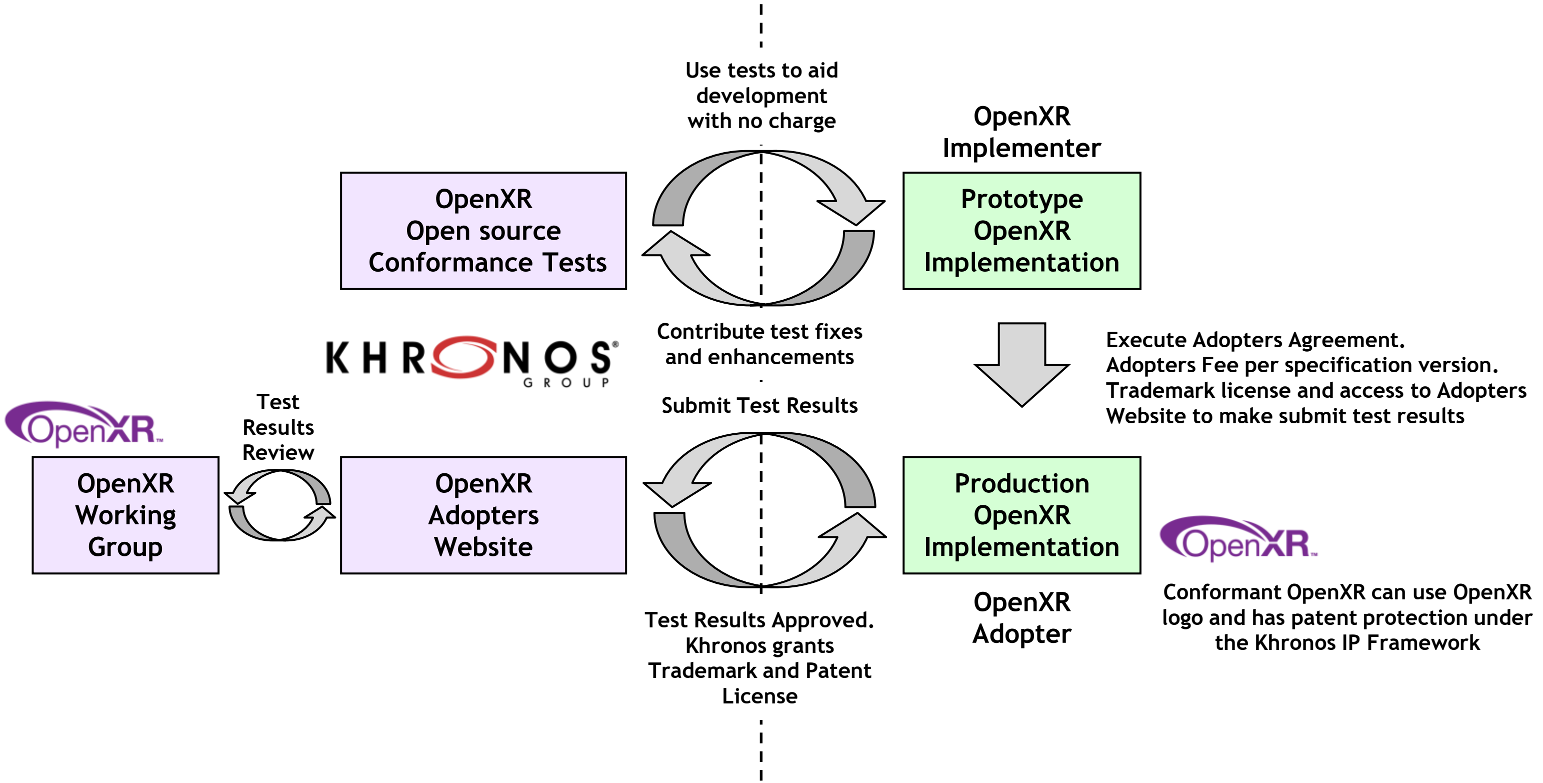

Use tests to aid development with no charge

OpenXR Implementer

OpenXR Open source Conformance Tests

Prototype OpenXR Implementation

Contribute test fixes and enhancements

Execute Adopters Agreement.
Adopters Fee per specification version.
Trademark license and access to Adopters Website to make submit test results

Test Results Review

Submit Test Results

OpenXR Working Group

OpenXR Adopters Website

Production OpenXR Implementation

OpenXR Adopter

Conformant OpenXR can use OpenXR logo and has patent protection under the Khronos IP Framework

Test Results Approved. Khronos grants Trademark and Patent License

Khronos Specification and Adoption Flow

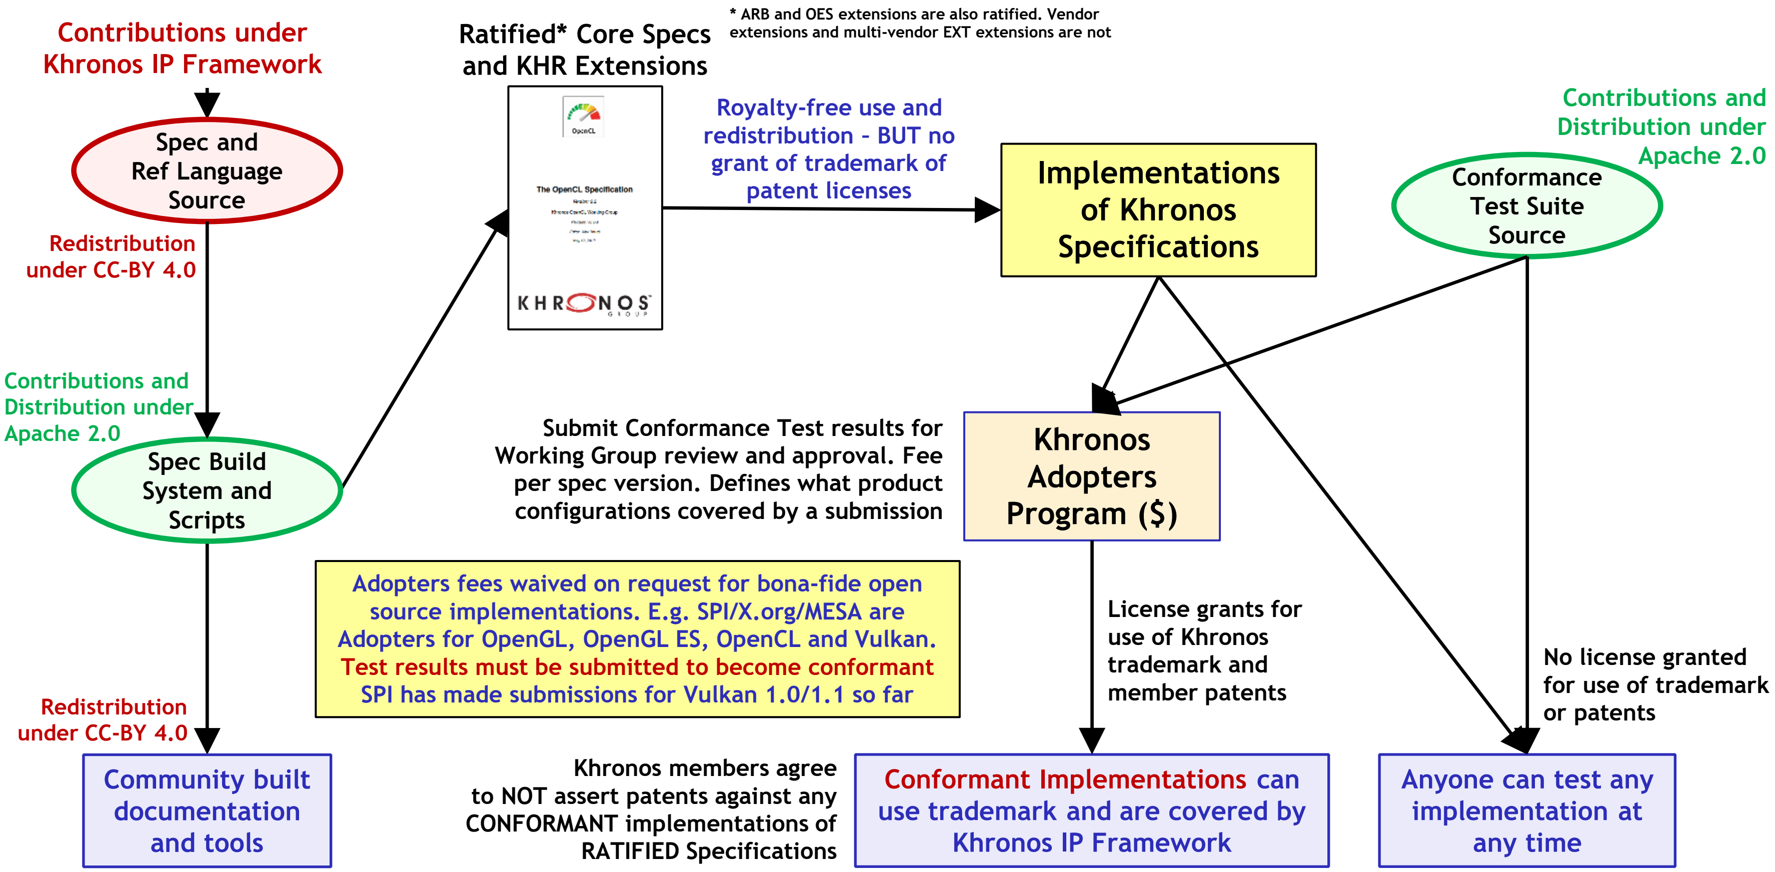

Contributions under Khronos IP Framework

Spec and Ref Language Source

Redistribution under CC-BY 4.0

Contributions and Distribution under Apache 2.0

Spec Build System and Scripts

Redistribution under CC-BY 4.0

Community built documentation and tools

Ratified* Core Specs and KHR Extensions

* ARB and OES extensions are also ratified. Vendor extensions and multi-vendor EXT extensions are not

Royalty-free use and redistribution - BUT no grant of trademark of patent licenses

Implementations of Khronos Specifications

Contributions and Distribution under Apache 2.0

Conformance Test Suite Source

Submit Conformance Test results for Working Group review and approval. Fee per spec version. Defines what product configurations covered by a submission

Khronos Adopters Program ($)

Adopters fees waived on request for bona-fide open source implementations. E.g. SPI/X.org/MESA are Adopters for OpenGL, OpenGL ES, OpenCL and Vulkan.
Test results must be submitted to become conformant
SPI has made submissions for Vulkan 1.0/1.1 so far

License grants for use of Khronos trademark and member patents

No license granted for use of trademark or patents

Khronos members agree to NOT assert patents against any CONFORMANT implementations of RATIFIED Specifications

Conformant Implementations can use trademark and are covered by Khronos IP Framework

Anyone can test any implementation at any time

Participation and Processes

Khronos Spec Standardization Process

Khronos processes is designed to encourage and create consensus

Working Groups are goal-focused
with very few 'politics'
Voting rarely needed to break deadlock

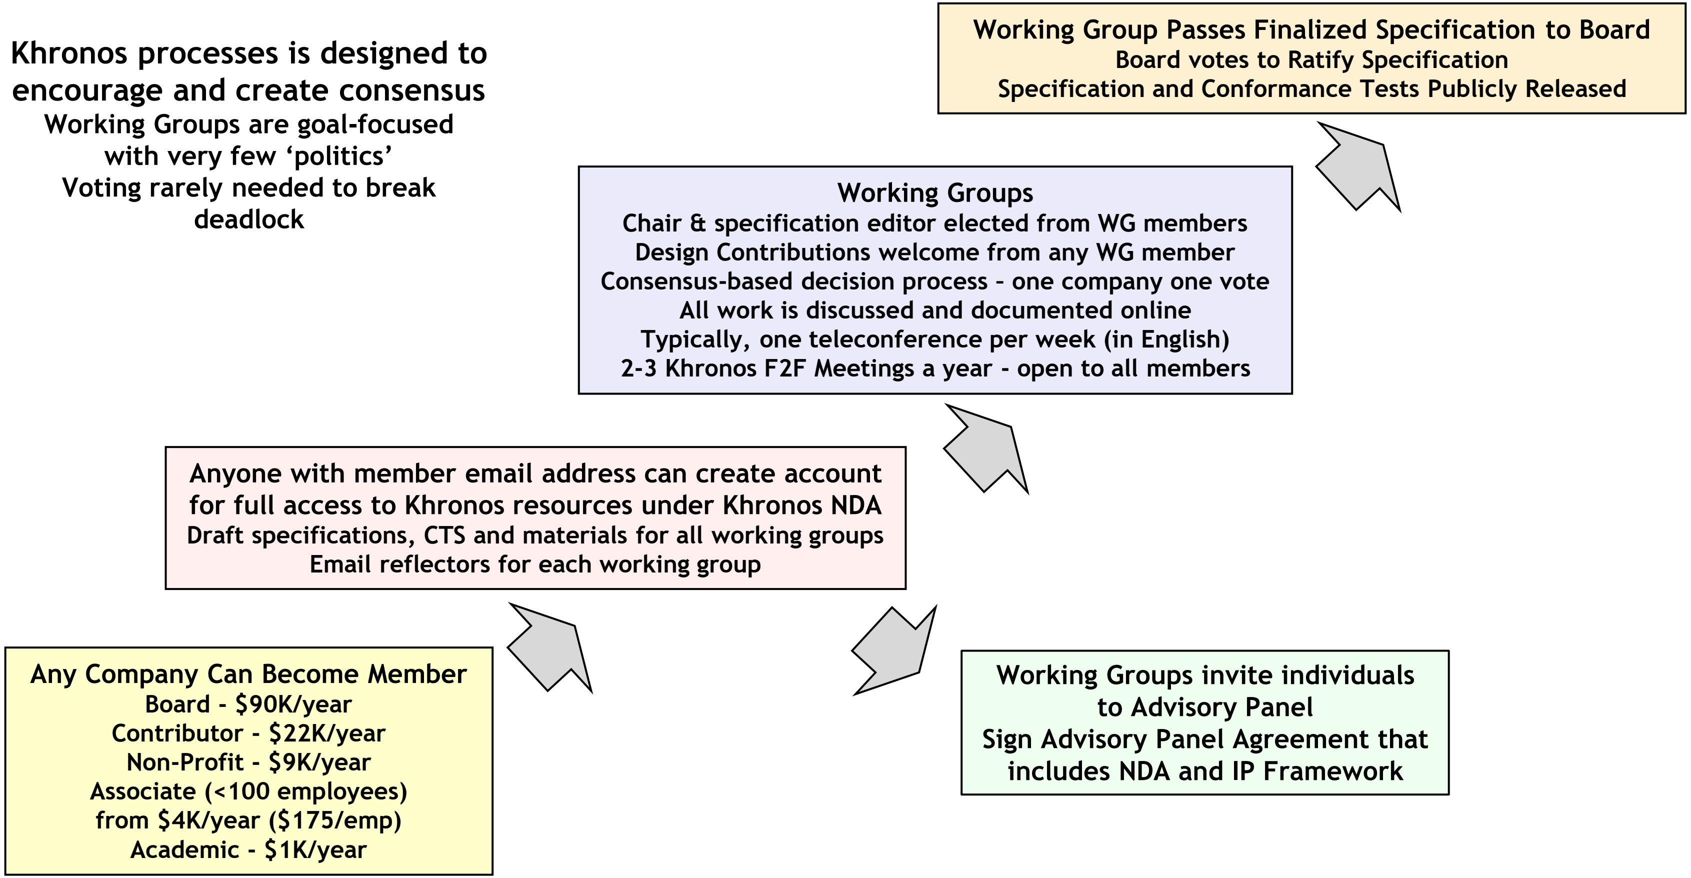

Any Company Can Become Member
Board - $90K/year
Contributor - $22K/year
Non-Profit - $9K/year
Associate (<100 employees)
from $4K/year ($175/emp)
Academic - $1K/year

Anyone with member email address can create account for full access to Khronos resources under Khronos NDA
Draft specifications, CTS and materials for all working groups
Email reflectors for each working group

Working Groups invite individuals to Advisory Panel
Sign Advisory Panel Agreement that includes NDA and IP Framework

Working Groups
Chair & specification editor elected from WG members
Design Contributions welcome from any WG member
Consensus-based decision process - one company one vote
All work is discussed and documented online
Typically, one teleconference per week (in English)
2-3 Khronos F2F Meetings a year - open to all members

Working Group Passes Finalized Specification to Board
Board votes to Ratify Specification
Specification and Conformance Tests Publicly Released

Khronos Conformance Submission Process

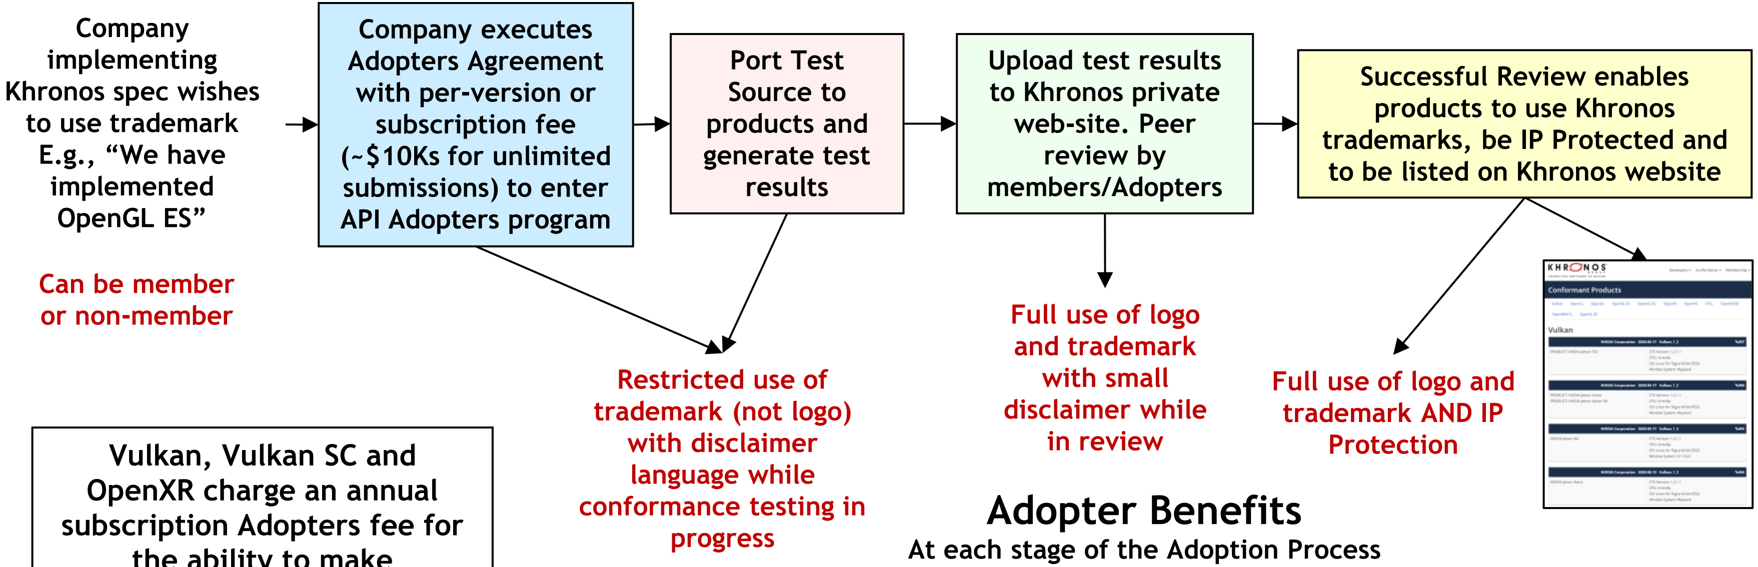

Company implementing Khronos spec wishes to use trademark E.g., “We have implemented OpenGL ES”

Can be member or non-member

Company executes Adopters Agreement with per-version or subscription fee (~$10Ks for unlimited submissions) to enter API Adopters program

Port Test Source to products and generate test results

Upload test results to Khronos private web-site. Peer review by members/Adopters

Successful Review enables products to use Khronos trademarks, be IP Protected and to be listed on Khronos website

Restricted use of trademark (not logo) with disclaimer language while conformance testing in progress

Full use of logo and trademark with small disclaimer while in review

Full use of logo and trademark AND IP Protection

Adopter Benefits

At each stage of the Adoption Process

Vulkan, Vulkan SC and OpenXR charge an annual subscription Adopters fee for the ability to make conformance submissions

Khronos Advisory Panels

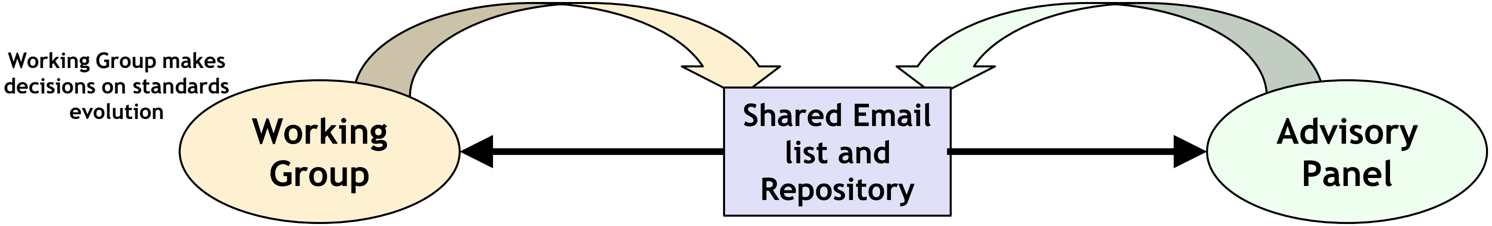

Working Group makes decisions on standards evolution

Working Group

Shared Email list and Repository

Advisory Panel

Hosted by Khronos.
Under Khronos NDA

Khronos Members

Any company can join.
Membership Fee.
Covered by NDA *and*
IP Framework

Working groups can share draft specifications and accept detailed design contributions as Panel Members are covered by IP Framework

Panel Members

Invited industry experts.
$0 Cost.
Covered by NDA *and*
IP Framework

Khronos Advisory Forums

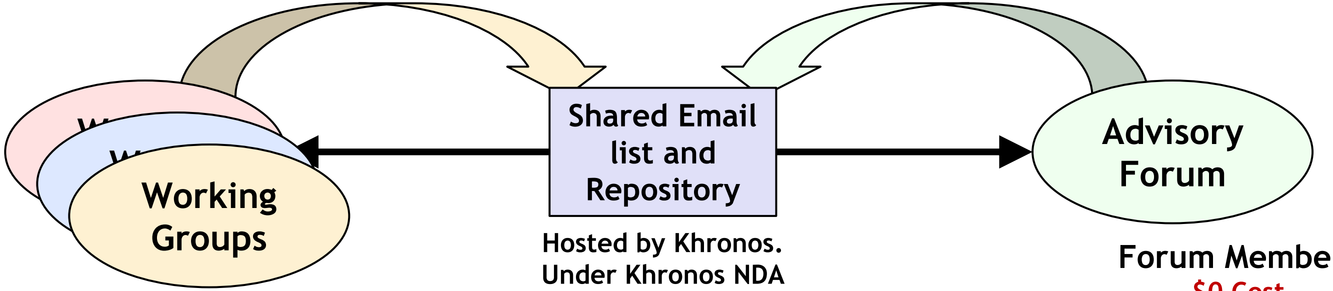

Hosted by Khronos.
Under Khronos NDA

Khronos Members
Can be from multiple working groups

Working groups can share draft specifications under NDA but NOT accept detailed design contributions as Panel Members are NOT covered by IP Framework

Forum Members
$0 Cost
Open to any interested company
MAY be covered by NDA but
NOT IP Framework

Khronos Ecosystem Engagement

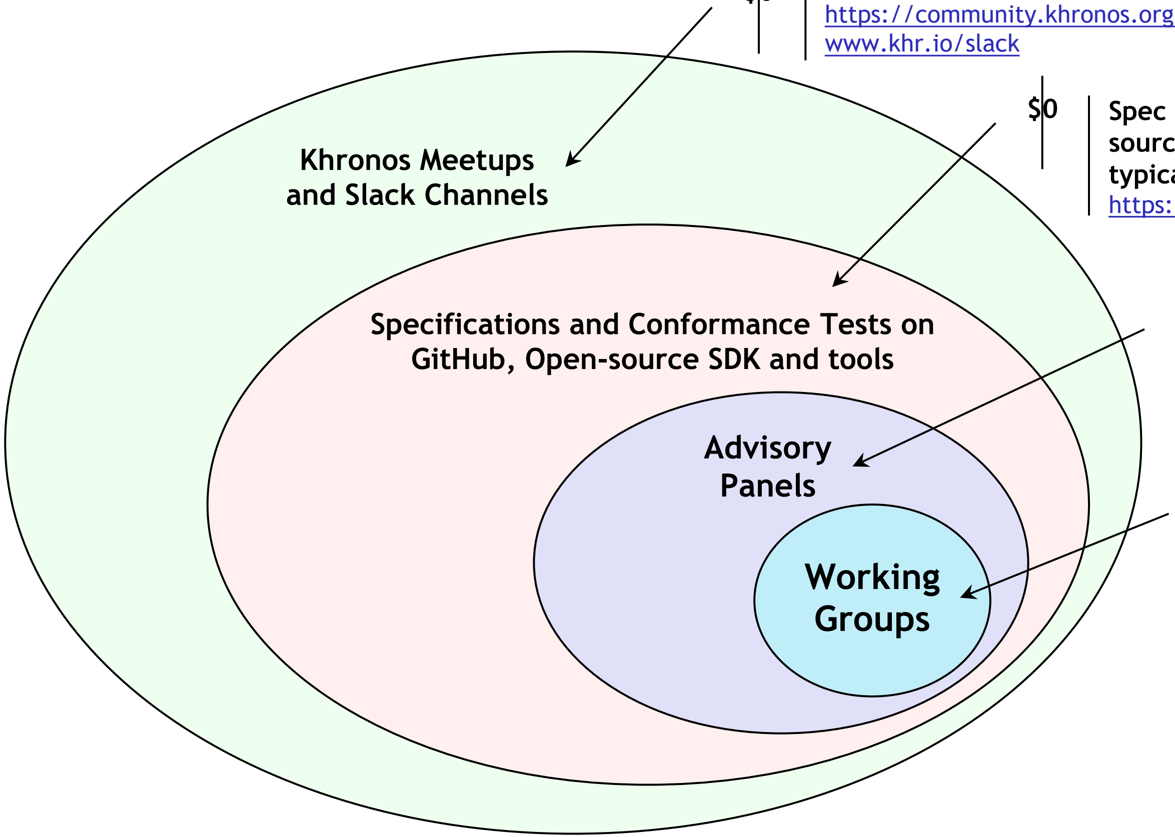

Khronos Meetups and Slack Channels

Specifications and Conformance Tests on GitHub, Open-source SDK and tools

Advisory Panels

Working Groups

$0 | Open to all!
https://community.khronos.org/
www.khr.io/slack

$0 | Spec fixes and suggestions. Open source contributions under repo’s CLA, typically Apache 2.0
https://github.com/KhronosGroup

$0 | Advisors under the Khronos NDA and IP Framework can comment and contribute to requirements and draft specifications
https://www.khronos.org/advisors/

$ | Khronos members under Khronos NDA and IP Framework participate and vote in working group meetings
https://www.khronos.org/members/