




Building the Metaverse One Standard at a Time

Neil Trevett | Khronos President
NVIDIA Vice President Developer Ecosystem
ntrevett@nvidia.com | [@neilt3d](https://twitter.com/neilt3d)

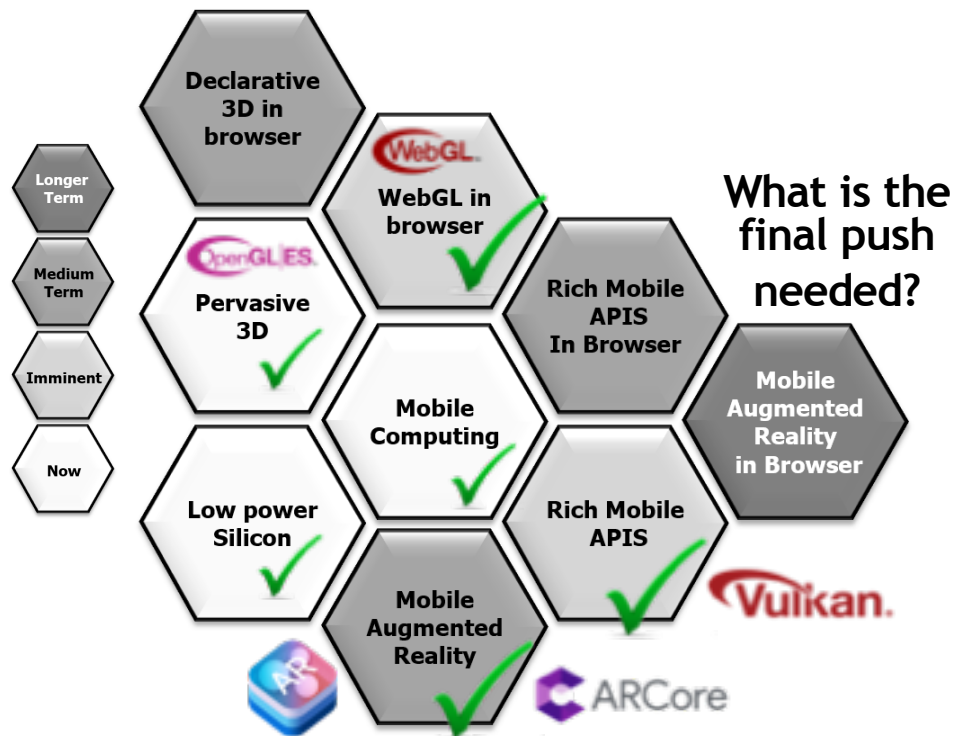
Web3D Conference
Poznań, Poland, June 22nd 2018

Khronos Mission



Khronos is an International Industry Consortium creating royalty-free, **open standards** to enable software to access hardware acceleration for **3D Graphics, Virtual and Augmented Reality, Parallel Computing, Neural Networks and Vision Processing**

Looking Back at 2011's Web3D Keynote



What is the final push needed?

The Goal

The Browser IS the metaverse
Real-world AR through Web pages not apps
Web standards interoperability overcomes platform fragmentation

What we didn't foresee in 2011

1. The rise of Machine Learning
2. Re-emergence of platform fragmentation
3. Layered ecosystems and compiler technology to combat fragmentation

Trevett Web3D Conference Keynote, Paris June 2011

Vital Elements for XR in the Browser



Efficient 3D
Asset Delivery





Cross-platform Sensor
and UI Handling



Fast, low-latency
3D Graphics









Compute, Vision and
Inferencing acceleration




All mobile and desktop have GPUs as offload processors

Ecosystem = API + File Format

	Run-time APIs	File Formats
3D Graphics		
Heterogenous Compute (Parallel Processing)		
VR and AR Vision and Inferencing		

glTF - Cross-Platform 3D Asset Transmission



All glTF spec development
on open GitHub:
<https://github.com/KhronosGroup/glTF>

- glTF™
- Compact to Transmit ✓
 - Fast to Load ✓
 - Describes Full Scenes ✓
 - Runtime Neutral ✓
 - Open and Extensible ✓

OpenGL Transmission Format
Efficient transmission of 3D
scenes and assets



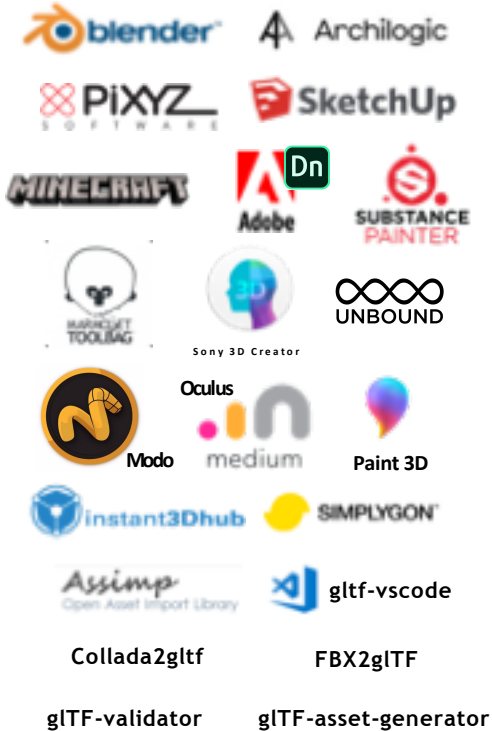
glTF 1.0 - Primarily for WebGL
Uses GLSL for materials
Released December 2015



glTF 2.0 - Physically Based Rendering
Metallic-Roughness and Specular-Glossiness Materials
Rendering API independence
Released @ Web3D 2017

glTF Momentum

Creation Tools



3D Content



Discover

Create



Experience

Drive Demand

Users




Apps and Engines



Consistent glTF Rendering on Multiple Engines




glTF considering
optional, standardized
environment and lighting
guidelines for consistent
default appearance
between viewers



<https://github.com/KhronosGroup/gltf-sample-models/pull/154>

glTF Recent Highlights

- Microsoft makes glTF files as usable as JPGs in Windows 10
 - Office , Remix 3D, Paint 3D, Simplygon, Mixed Reality Portal and Viewer
- Facebook supports drag and drop for glTF models to your feed
 - Driving exporter/tool demand: FBX2glTF, Modo, ...
- Sketchfab repository has over 150K glTF models
 - Under free, Creative Commons license
- Adobe Dimension is adopting glTF for delivery of 3D marketing assets
 - glTF publishing service for interactive marketing materials
- Mozilla integrating glTF into A-FRAME
 - Anchors for glTF objects in AR scenes
- Google Draco Mesh Compression
 - Extension is shipping and being used by tools and engines



glTF Extension and Roadmap *Discussions*



- glTF is extensible
 - Use extensions for specialist functions or testing new functionality
- Some functionality under discussion
 - Advanced lighting and shadows
 - Environmental and HDR lighting
 - Compressed Texture Transmission with efficient expansion into GPU formats
 - Advanced PBR - NVIDIA MDL?
 - Subdivision surfaces - Pixar OpenSubdiv?
 - Point clouds with Draco compression
 - Draco compression of animation data
 - Geometry streaming etc. etc.

glTF Ecosystem meeting at SIGGRAPH 2018
Let Neil know if you would like an invite!



Any delivery format must balance functionality versus speed of processing and simplicity of implementing reliable importers/exporters

<https://gtr.sketchfab.com/models/damagedF1etnet.zip>

USDZ

At WWDC 2018 Apple announced their own 3D delivery format - USDZ

Pixar USD is a movie authoring interchange format - open sourced in 2016

USDZ packages USD files in a non-compressed ZIP container

Early days - not clear if Apple will subset USDZ for their platform - no spec yet



USDZ

Compact Delivery Format	Authoring Interchange Format
Multi-company Governance	Delivery subset controlled by Apple?
Open standard for use on any platform	Apple platform only?
Widespread open source tools ecosystem	?

Competition is healthy

Awaiting more details to be announced by Apple

Conversion between glTF and USDZ is certainly possible

Vital Elements for XR in the Browser



Efficient 3D
Asset Delivery



Cross-platform Sensor
and UI Handling





Fast, low-latency
3D Graphics



Compute, Vision and
Inferencing acceleration



All mobile and desktop have GPUs as offload processors



WebGL Evolution



Mobile Graphics

Programmable
Vertex and Fragment
shaders



2007

OpenGL ES 2.0

Desktop Graphics

Textures: NPOT, 3D, Depth, Arrays, Int/float
Objects: Query, Sync, Samplers
Seamless Cubemaps, Integer vertex attributes
Multiple Render Targets, Instanced rendering
Transform feedback, Uniform blocks
Vertex array objects, GLSL ES 3.0 shaders



2012

OpenGL ES 3.0

Apple does not ship
OpenGL ES 3.1 -
ending the era of
universal native 3D APIs

2014

OpenGL ES 3.1
Compute Shaders

Explicit Graphics

General Purpose Compute
Multi-threaded Rendering
Low-latency rendering



2016

Vulkan.

4 years

2011
WebGL 1.0

5 years

March 2017
WebGL 2.0

Pervasive OpenGL ES 2.0

OpenGL and OpenGL ES ships on every desktop and mobile OS
3D on the Web is enabled!

Significant ongoing work on multiple performance enhancements and getting all browsers to pass all the complete test suite

Conformance Testing is vital for Cross-Platform Reliability

WebGL 2.0 conformance tests are very thorough
10x more tests than WebGL 1.0 tests

WebGL Momentum - WebGL 2.0 is Here!



FLASH & THE FUTURE OF INTERACTIVE CONTENT

POSTED BY ADOBE CORPORATE COMMUNICATIONS ON JULY 25, 2017

Subscribe

Adobe has long played a leadership role in advancing interactivity and creative content – from video, to games and more – on the web. Where we've seen a need to push content and interactivity forward, we've innovated to meet those needs. Where a format didn't exist, we invented one – such as with Flash and Shockwave. And over time, as the web evolved, these new formats were adopted by the community, in some cases formed the basis for open standards, and became an essential part of the web.

But as open standards like WebGL have matured over the past several years, most now provide many of the capabilities that plugins pioneered and have become a viable alternative for content on the web. Over time, we've seen helper apps evolve to become plugins, and more recently, have seen many of these plugin capabilities get incorporated into open web standards. Today, most browser vendors are integrating capabilities once provided by plugins directly into browsers and deprecating plugins.

93.26% Globally

WebGL - 3D Canvas graphics - OTHER

Method of generating dynamic 3D graphics using JavaScript, accelerated through hardware

Usage % of all users
Global 93.26%

Current aligned Usage relative Date relative Show all

IE	Edge *	Firefox	Chrome	Safari	iOS Safari *	Opera Mini *	Chrome for Android	UC Browser for Android	Samsung Internet
	16	59	49		10.3				
11	17	60	66	11.1	11.2				4
	18	61	67	12	11.3	all	66	11.8	6.2
		62	68	TP					
			69						

<http://cantuse.com/#feat=webgl>

62.85% Globally

WebGL 2.0 - OTHER

Next version of WebGL. Based on OpenGL ES 3.0.

Usage % of all users
Global 62.85%

Current aligned Usage relative Date relative Show all

IE	Edge *	Firefox	Chrome	Safari	iOS Safari *	Opera Mini *	Chrome for Android	UC Browser for Android	Samsung Internet
	16	59	49		10.3				
11	17	60	66	11.1	11.2				4
	18	61	67	12	11.3	all	66	11.8	6.2
		62	68	TP					
			69						



WebGL 2.0 brings Desktop-class graphics to the Web
The time to create a new class of Web-based 3D Apps is now!

Pervasive Vulkan



Major GPU Companies supporting Vulkan for Desktop and Mobile Platforms



Platforms



Desktop



Mobile
(Android 7.0+)



Media Players



Consoles



Virtual Reality



Cloud Services



Embedded

Game Engines



Vulkan and New Generation GPU APIs

Non-proprietary, royalty-free open standard 'By the industry for the industry'

Portable across multiple platforms - desktop and mobile
Modern architecture | Low overhead | Multi-thread friendly

EXPLICIT GPU access for EFFICIENT, LOW-LATENCY,
PREDICTABLE performance



Vulkan is available on Android 7.0+

Bringing Vulkan Apps to Apple Platforms



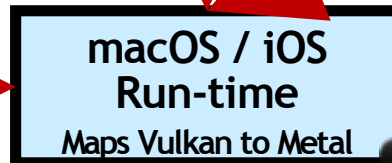

Dota 2 running on Mac up to 50% faster than native OpenGL



Efficiently map a subset of Vulkan over Metal
Very little functionality *not* supported today:
Triangle fans, Separate stencil reference masks
Vulkan Events, Allocation callbacks
Some texture-specific swizzles



Open source SDK to build, run, and debug applications on macOS including validation layer support




MoltenVK for macOS and iOS
For macOS 10.11, iOS 9.0 and up



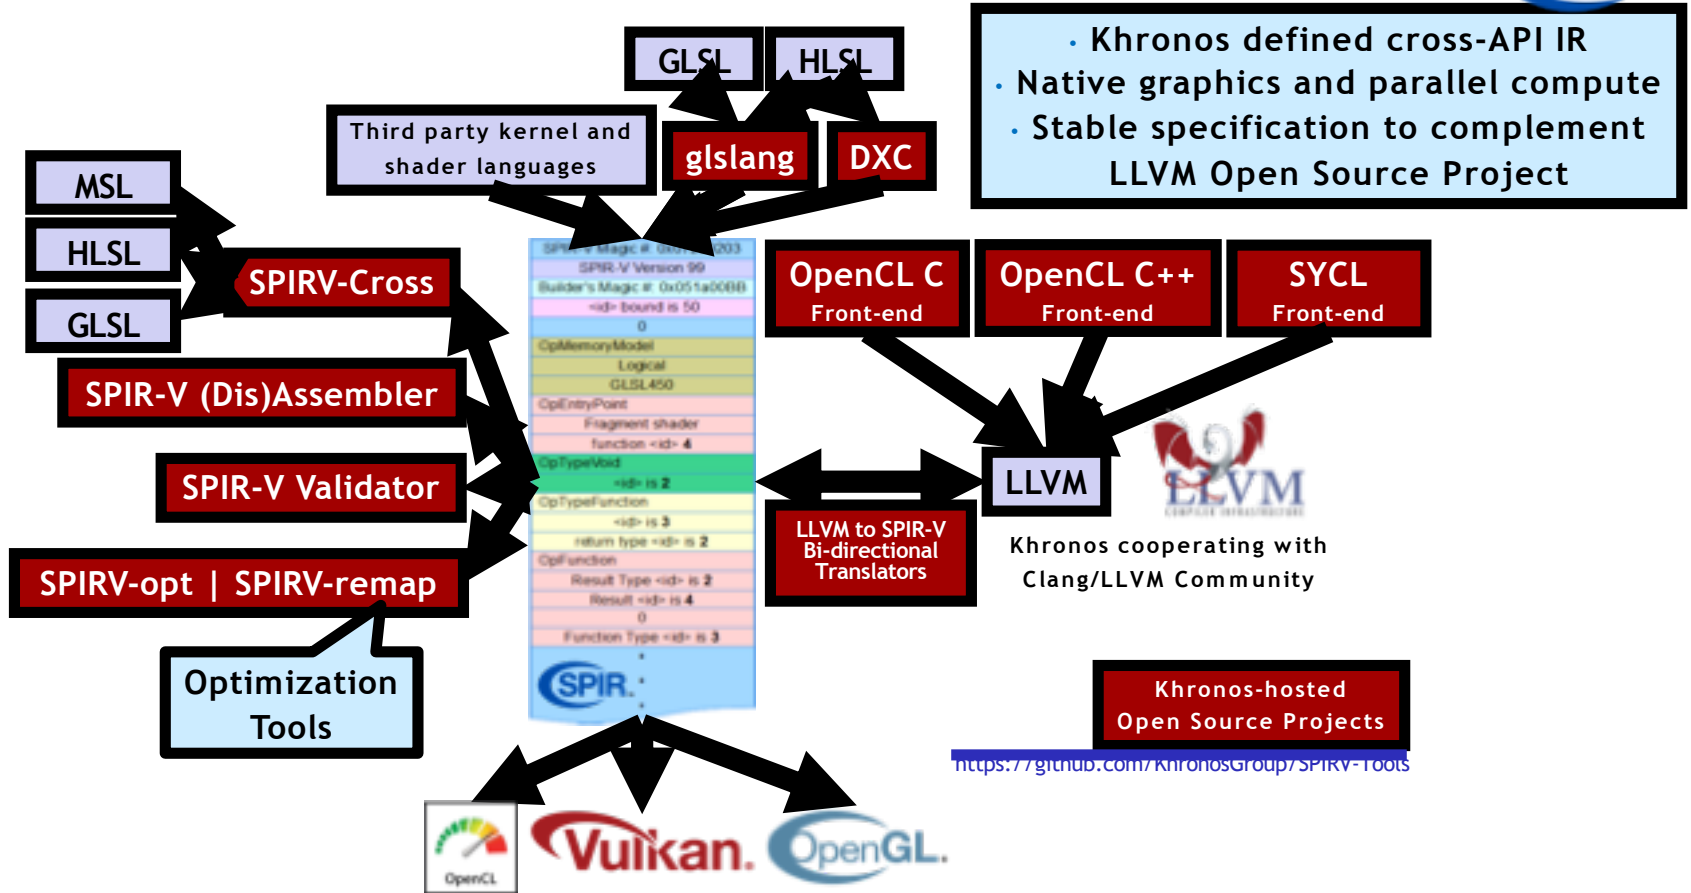
Beta release - but working to pass all applicable conformance tests

Previously a paid product
Now released into OPEN SOURCE
Completely free to use - no fees or royalties - including commercial applications

Valve - Vulkan Dota 2 on macOS



SPIR-V Ecosystem





- Khronos defined cross-API IR
- Native graphics and parallel compute
- Stable specification to complement LLVM Open Source Project

Bringing OpenCL Compute to Vulkan

- Experimental Clspv Compiler from Google, Adobe and Codeplay
 - Compiles OpenCL C to Vulkan's SPIR-V execution environment
Successfully tested on over 200K lines of Adobe OpenCL C production code
 - Open source - tracks top-of-tree LLVM and clang, not a fork



Layering and cross-compilation is making Vulkan a Universally Portable API for Graphics AND Compute



Prototype open source project

<https://github.com/google/clspv>



Possible future project - if interest

Increasing deployment options for OpenCL kernel developers e.g. Vulkan is a supported API on Android

WebGL Future

Game developers fume as Apple deprecates OpenGL

One false move and the game gets it



Fortnite has already been rewritten for Metal

IN AMONGST the announcement of [macOS 10.14 Mojave](#) on Monday, there was an alarming development for gamers and game developers. OpenGL and OpenCL are being deprecated.

No universally available 3D API since Apple did not ship OpenGL ES 3.1

That situation just got worse at Apple WWDC June 2018 - OpenGL and OpenGL ES deprecated in favor of Apple's Metal API

Don't panic!

Deprecation means the APIs will be removed in the future - still available right now

WebGL 2.0 could be layered over Metal



Dean is Apple's WebGL lead



GPUWeb Community Group

Makes sense that work on nextgen Web stack takes place at W3C

But currently considering designing a fourth native GPU API - without significant input from the GPU hardware community ☹️

Could cause further industry fragmentation and be a barrier to adoption ☹️ ☹️

Watching with interest to see if GPUWeb becomes a formal working group

May use SPIR-V for shader code

Vital Elements for XR in the Browser



Efficient 3D
Asset Delivery



Cross-platform Sensor
and UI Handling



Fast, low-latency
3D Graphics

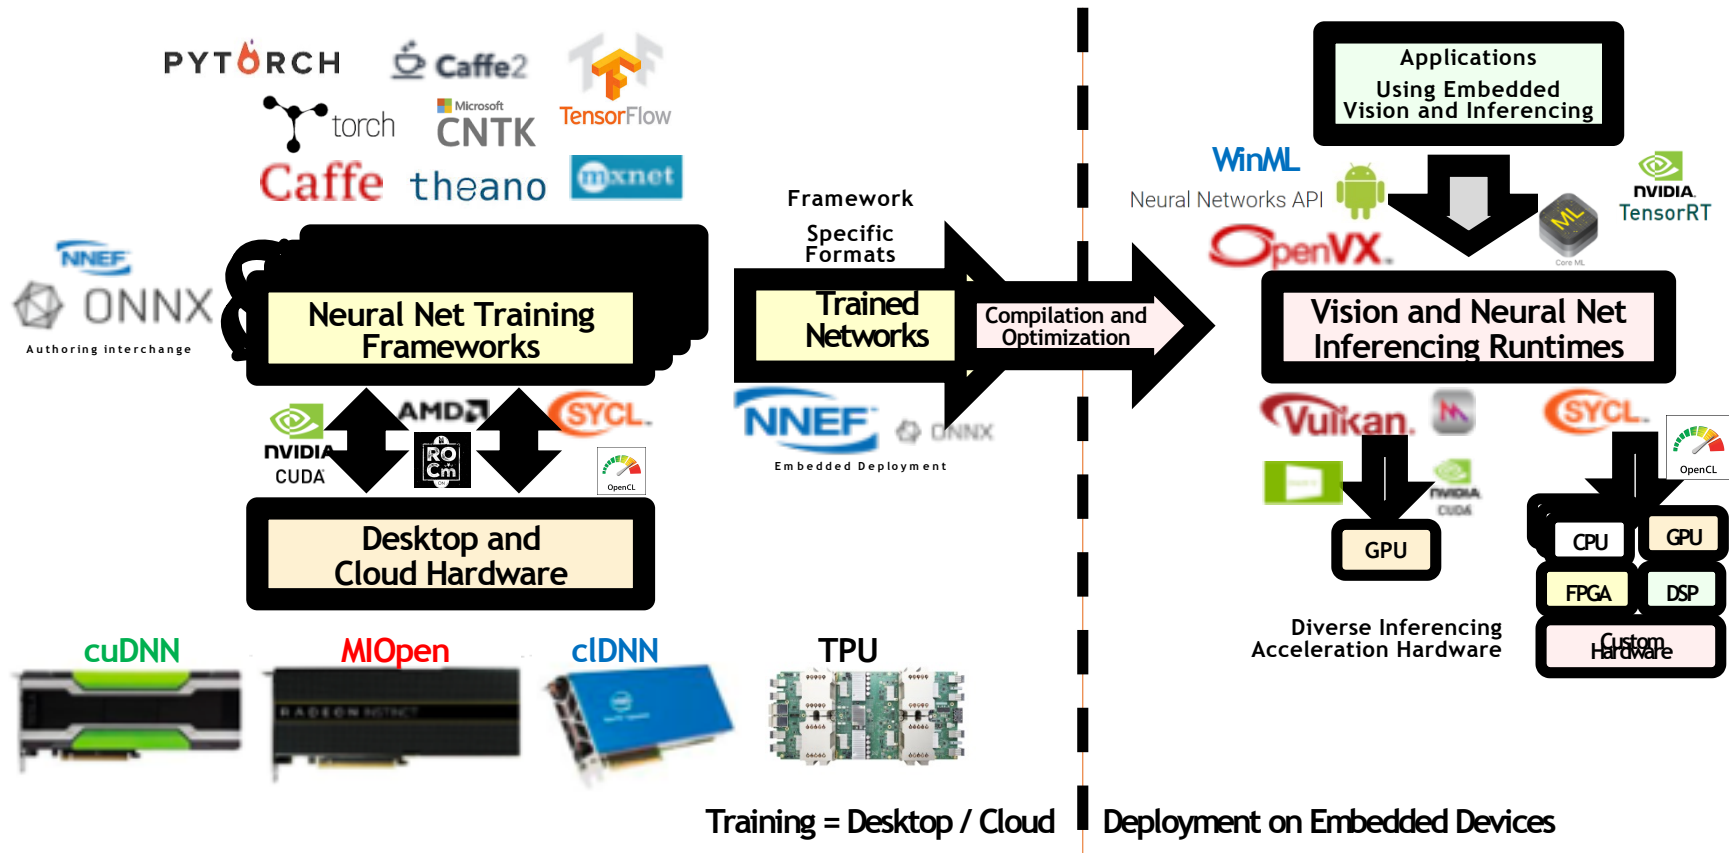


Compute, Vision and
Inferencing acceleration



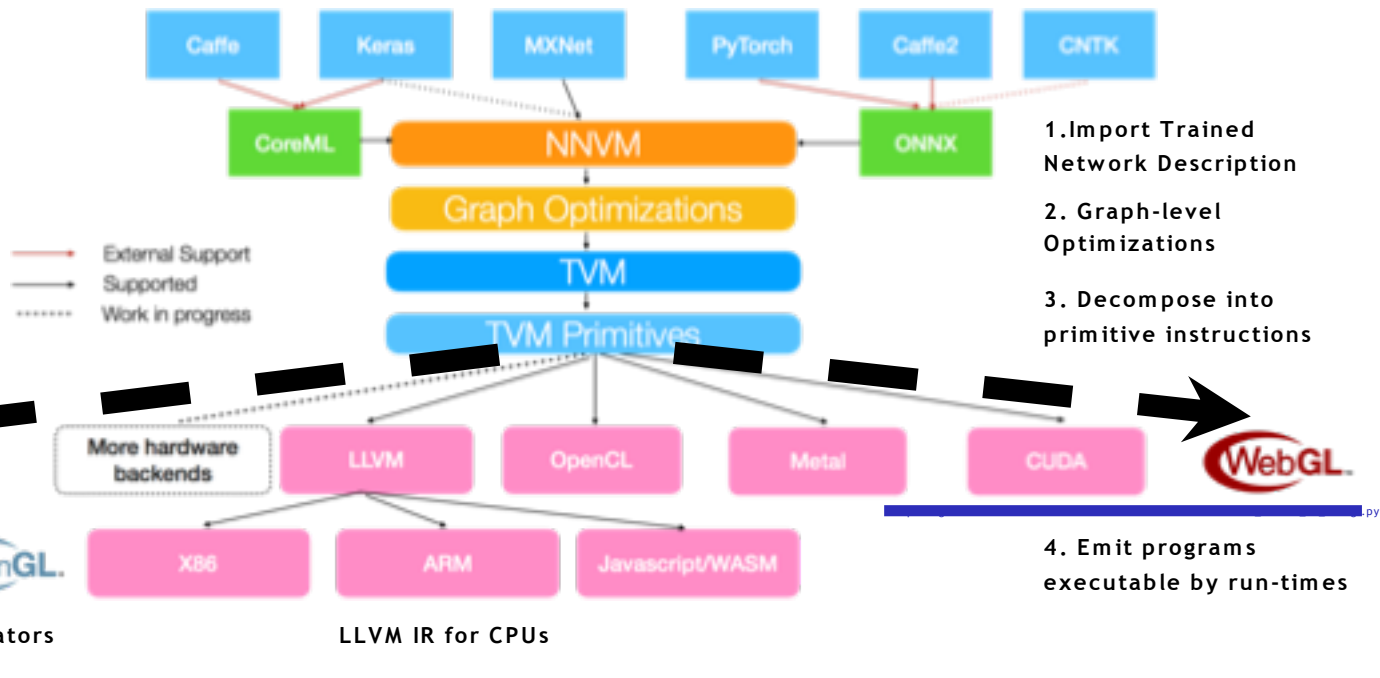
All mobile and desktop have GPUs as offload processors

Neural Network Workflow



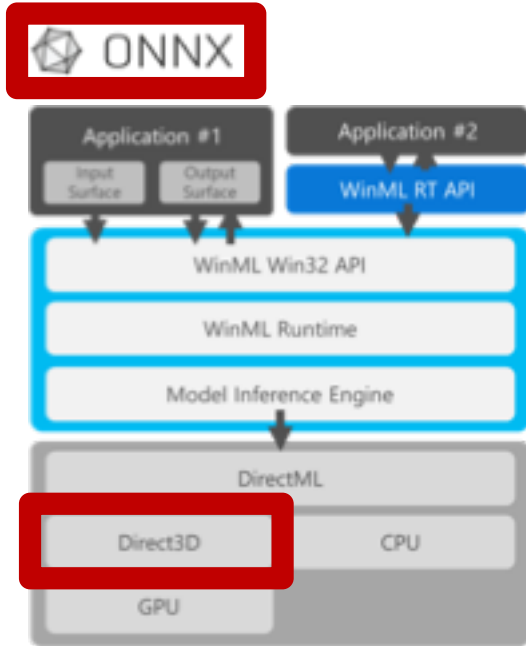
NNVM - Open Compiler for AI Inferencing

Paul G. Allen School of Computer Science & Engineering, University of Washington

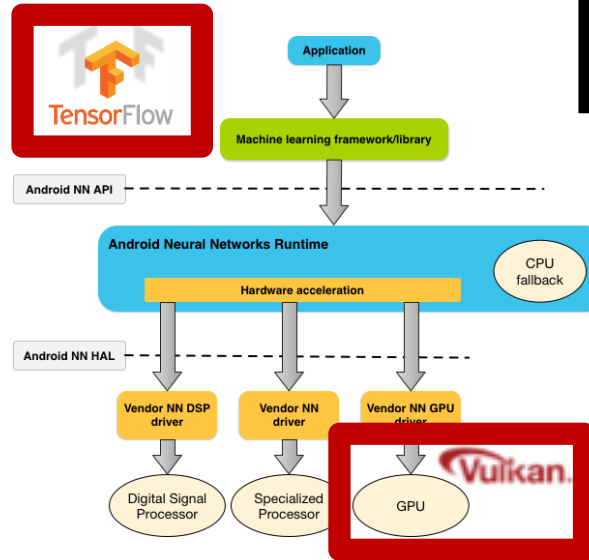


TensorFlow over WebGL demo:
 TensorFlow.js powered implementation of the tSNE algorithm for high-dimensional data analysis
<https://nicola17.github.io/tfis-tsne-demo/>

Platform Neural Network Stacks



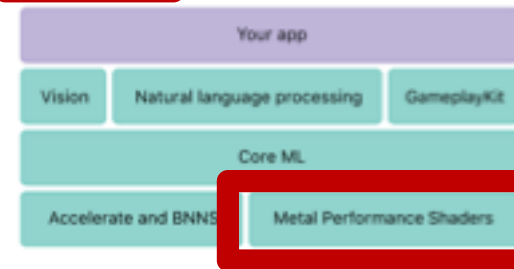
Microsoft Windows
Windows Machine Learning (WinML)



Google Android
Neural Network API (NNAPI)

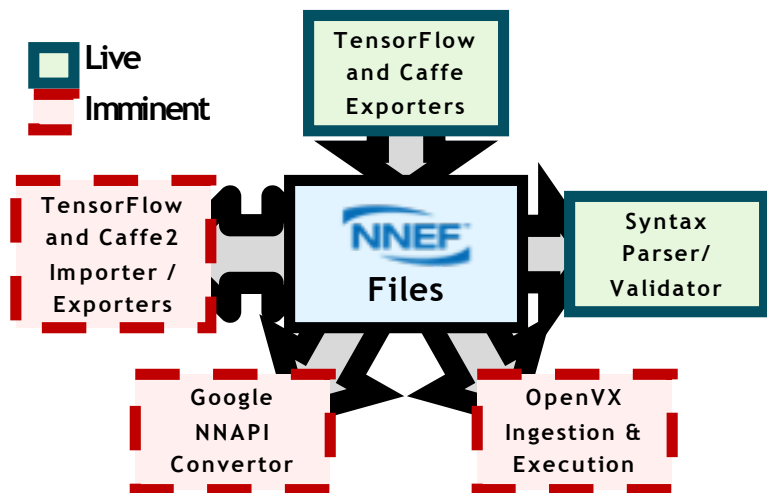
Consistent Fundamental Steps

1. Import trained NN model file
2. Build optimized version of graph
3. Accelerate on GPU or other processor using available low-level API



Apple MacOS and iOS
CoreML

NEF Ecosystem



NEF open source projects hosted on
 Khronos NEF GitHub repository
 Apache 2.0 license
<https://github.com/KhronosGroup/NEF-Tools>

NEF = Neural Network Exchange Format
 V1.0 Provisional Released for industry feedback before finalization

Comparing Neural Network Exchange Industry Initiatives



NEF	ONNX
Defined Specification	Open Source Project
Stability for hardware deployment	Software stack flexibility
Multi-company Governance	Initiated by Facebook
Flexible Precision	32-bit Floating Point only
Flat and Compound Ops	Flat Ops Only

Vital Elements for XR in the Browser



Efficient 3D
Asset Delivery




Cross-platform Sensor
and UI Handling



Fast, low-latency
3D Graphics

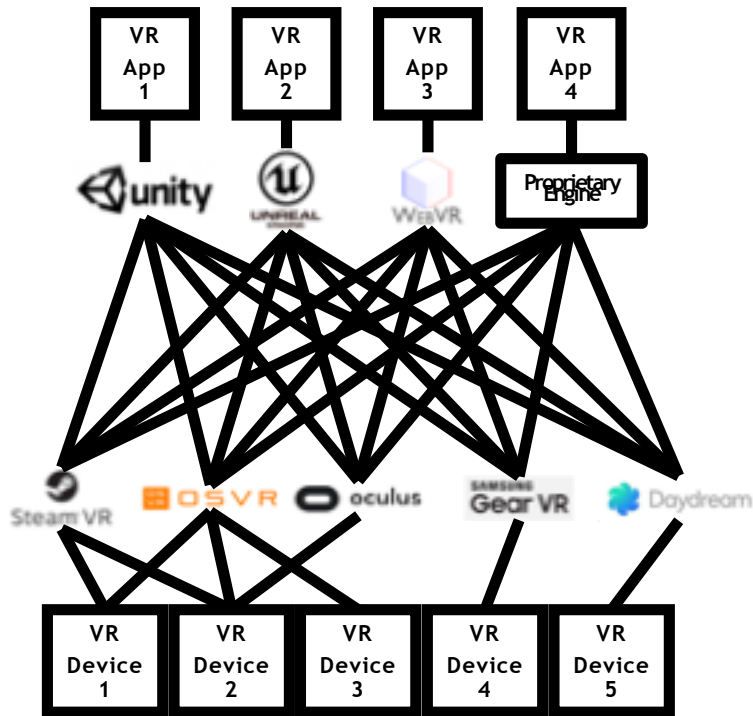


Compute, Vision and
Inferencing acceleration

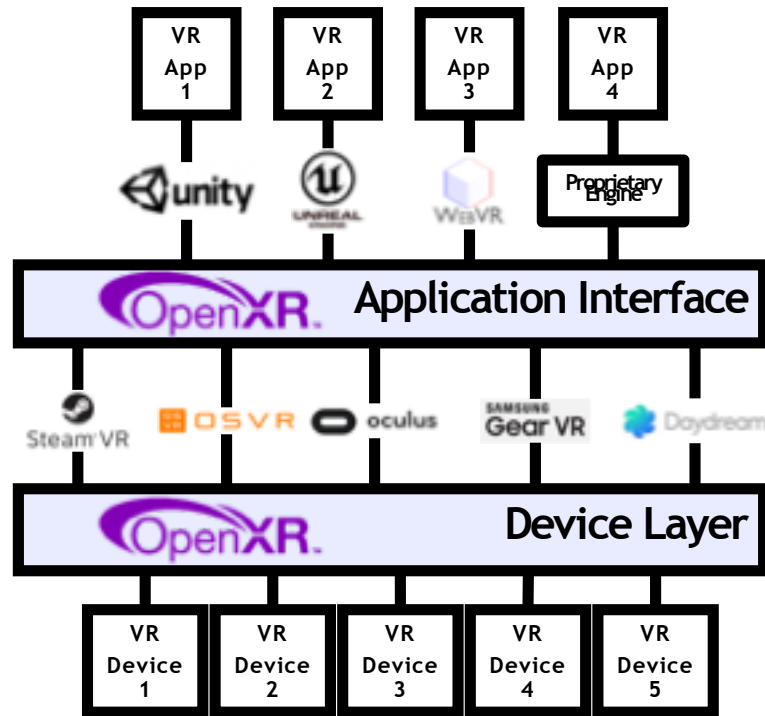


All mobile and desktop have GPUs as offload processors

OpenXR - Solving VR Fragmentation



Before OpenXR
VR Market
Fragmentation



After OpenXR
Wide interoperability of
VR apps and devices

Companies Publicly Supporting OpenXR

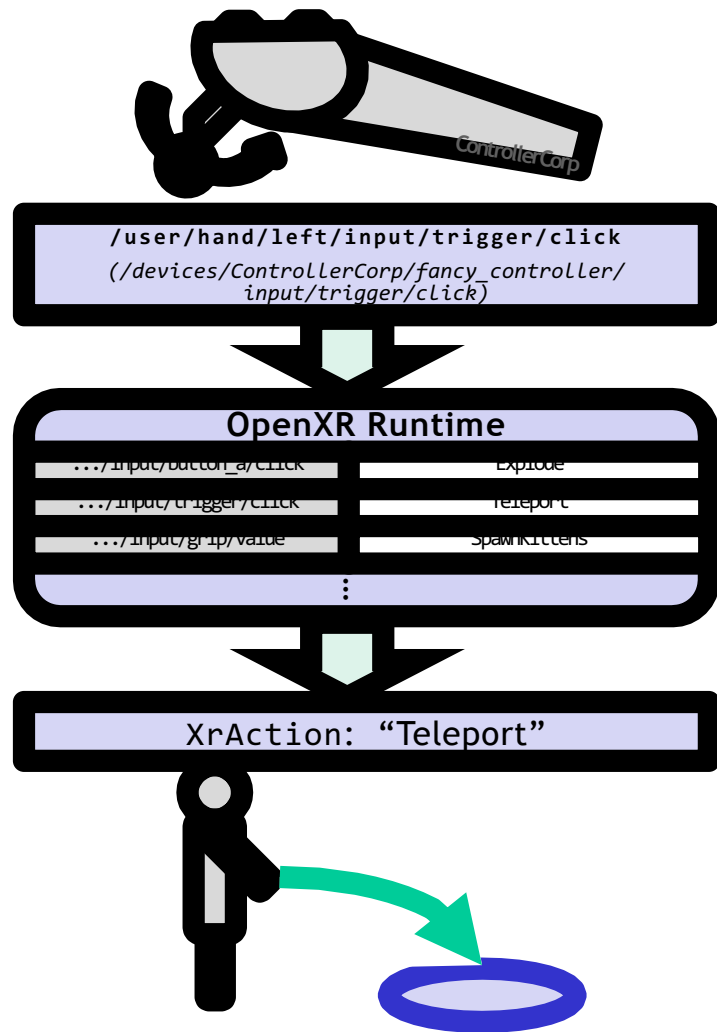


OpenXR is a collaborative design

- 1) For cross-platform XR portability - VR in V1.0, then add AR
- 2) Integrating many lessons from proprietary 'first-generation' API designs

Input and Haptics

- Input uses abstracted Input Actions
 - E.g. “Move,” “Jump,” “Teleport”
- Many advantages
 - Existing content can easily use new devices
 - Mix-and-match multiple input sources to create a unified UI
 - Easy optional feature support (e.g. eye and body tracking)
 - Future-proofing for innovation in input devices and form factors



OpenXR Viewport Configurations

- Applications can:
 - Query for runtime supported Viewport Configurations
 - Applications can then set the Viewport Configurations that they plan to use
 - Select and change their active configuration over the lifetime of the session

Camera Passthrough AR



Stereoscopic VR / AR



Projection CAVE



Photo Credit: Dave Papp

One Viewport

`/viewport_configuration/ar_mono/magic_window`

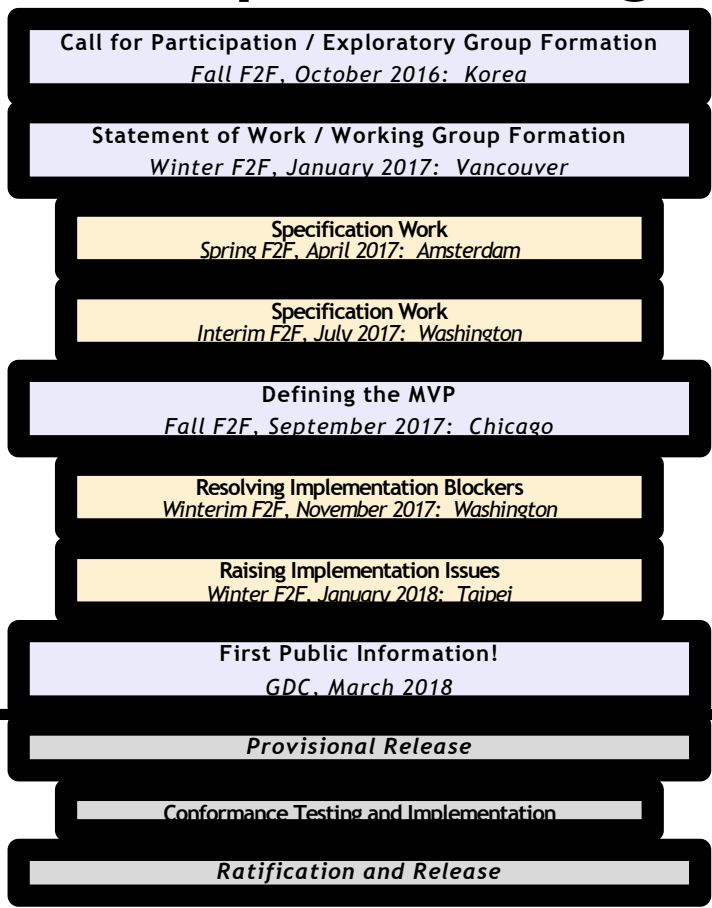
Two Viewports (one per eye)

`/viewport_configuration/vr/hmd`

Twelve Viewports (six per eye)

`/viewport_configuration/vr_cube/cave_vr`

OpenXR Development Progress



Much more detailed specification overview
and GDC session videos:
<https://www.khronos.org/developers/library/2018-gdc>

Present Day
Coming Soon

Implementations Underway!
Spec Finalization - incorporating
implementation experience

Mobile Augmented Reality Libraries

Encapsulated Vision-based
Functionality
Also leveraging motion sensors



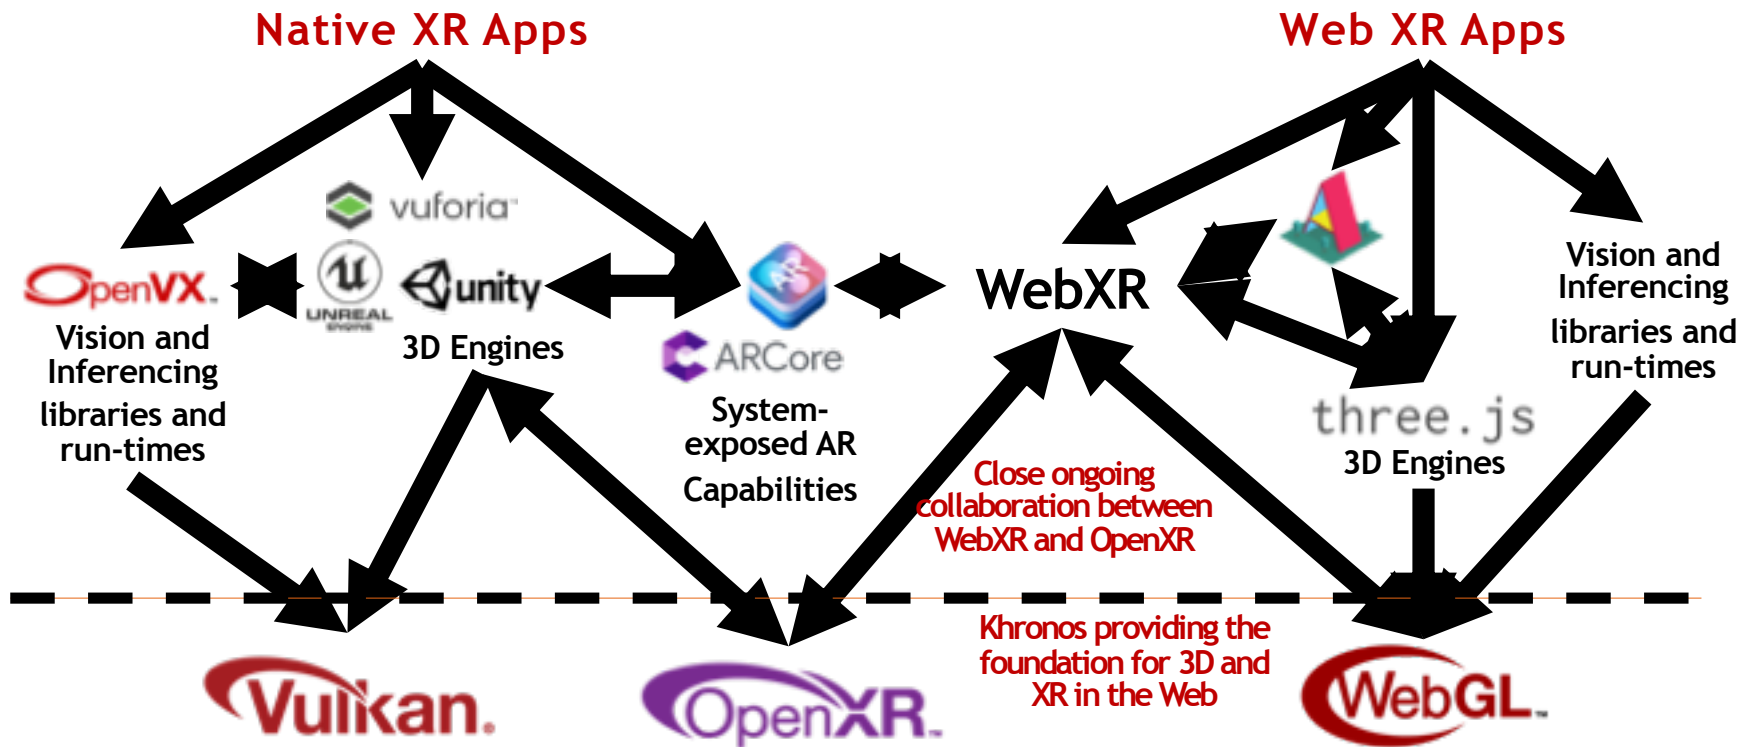
	ARKit	ARCore
Pose Tracking	Yes	Yes
Plane Detection and Tracking	Yes	Yes
Image Recognition and Tracking	Yes	Yes
Ambient light level and temperature	Yes	Yes
Link to Neural Net-based Object Detection	Yes	Yes
Access to Point Cloud	Yes	Yes
Camera Intrinsic	Yes	Yes
Multi-user and persistent cloud-based anchors	Yes	Yes
Face tracking (iPhone X)	Yes	No
Light probe (using a tracked face)	Yes	No
OS Availability	iOS	Android and iOS

3rd party AR Libraries typically use ARKit/ARCore if available or implement own tracking if not

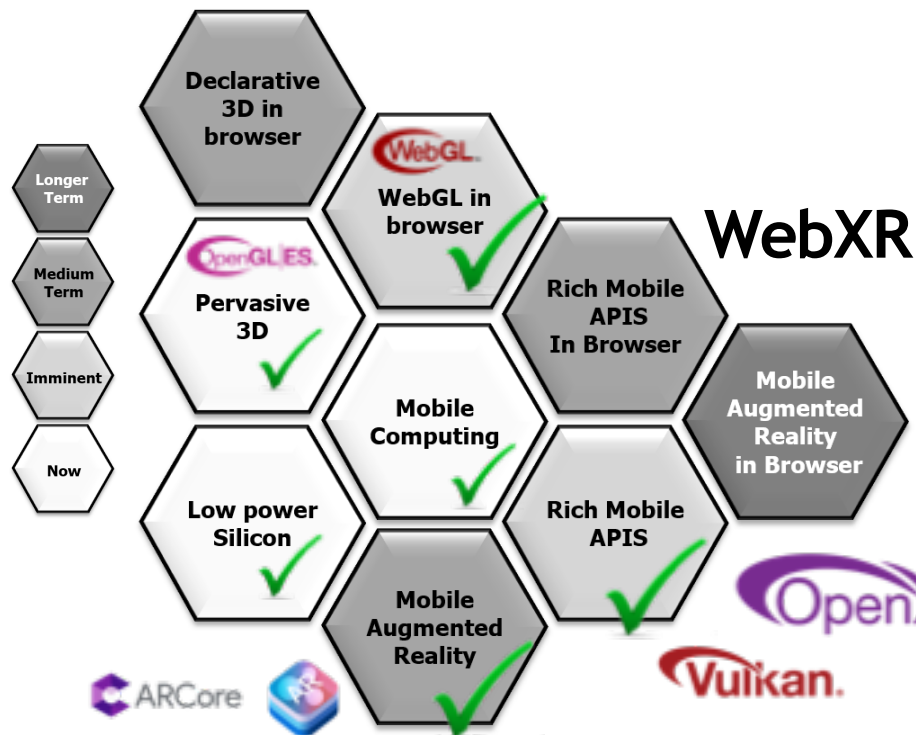


V1.0 is VR focused. Subsequent versions will extend to cross-platform AR

Layered XR Ecosystems with GPU Acceleration



Looking Back at 2011's Web3D Keynote



With WebXR and OpenXR, portable, accelerated XR in the Web is NOT going to take another 7 years 😊 It will be here *much* sooner!

Trevett Web3D Conference Keynote, Paris June 2011

Please Get Involved!

- These slides and information on all Khronos standards at:
 - www.khronos.org
- Khronos is driving to new levels of community engagement
 - Specifications and resources increasingly open for industry input and feedback
- Any company or organization is welcome to join Khronos
 - For a voice and a vote in any of these standards
- If joining is not possible - ask about a Khronos Advisory Panel
 - Free of charge - enables design reviews, requirements and contributions
- Khronos is always open to liaison projects with the 3D Web community
 - The industry benefits from everyone working more closely together
- Neil Trevett
 - ntrevett@nvidia.com
 - [@neilt3d](https://twitter.com/neilt3d)

