

KHRONOS[®]
GROUP

WEBINAR

OpenVX[™]

Vision and Inferencing Acceleration

**OpenVX Performance, Portability
and Memory Footprint Advantages**

14th October 2020

Welcome!

This slide deck will be posted at www.khronos.org today

The webinar will be recorded and made freely available

Please take a moment to answer the polls at the end of the webinar

Ask any questions via the Q&A Panel at any time

We will answer questions at the end of the webinar



Today's Speakers



**Introduction
to Khronos and
OpenVX**
Neil Trevett
Khronos President



**Cadence
and OpenVX**
Frank Brill
Cadence Systems
Design Engineering Director

**AMD
and OpenVX**
Mike Schmit
AMD

Director of Software System
Design ML Computer Vision



**Using
OpenVX**

Kiriti Nagesh Gowda
OpenVX Group Chair



Khronos Connects Software to Silicon

Open interoperability standards to enable software to effectively harness the power of 3D and multiprocessor acceleration



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

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

Open to any interested company

All Khronos standards are royalty-free

Well-defined IP Framework protects participant's intellectual property

Founded in 2000

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

Khronos Active Initiatives

3D Graphics
Desktop, Mobile, Web
Embedded and Safety Critical



3D Assets
Authoring
and Delivery



Portable XR
Augmented and
Virtual Reality



Parallel Computation
Vision, Inferencing,
Machine Learning




The Origin of OpenVX

Engines and Applications



3D Graphics API
Driver



Driver Model

An open API standard enables multiple silicon vendors to ship drivers with their silicon


Silicon vendors can aggressively optimize drivers for their own silicon architecture

OpenVX is the industry's only API standard enabling portable access to vendor-optimized vision drivers

Engines and Applications



Vision API
Driver




High-level Abstraction

3D graphics is always accelerated by a GPU - so a low-level GPU-centric API still provides cross-vendor portability

Vision processing can be accelerated by a wide variety of hardware architectures

OpenVX needs a higher-level graph abstraction to enable optimized cross-vendor drivers

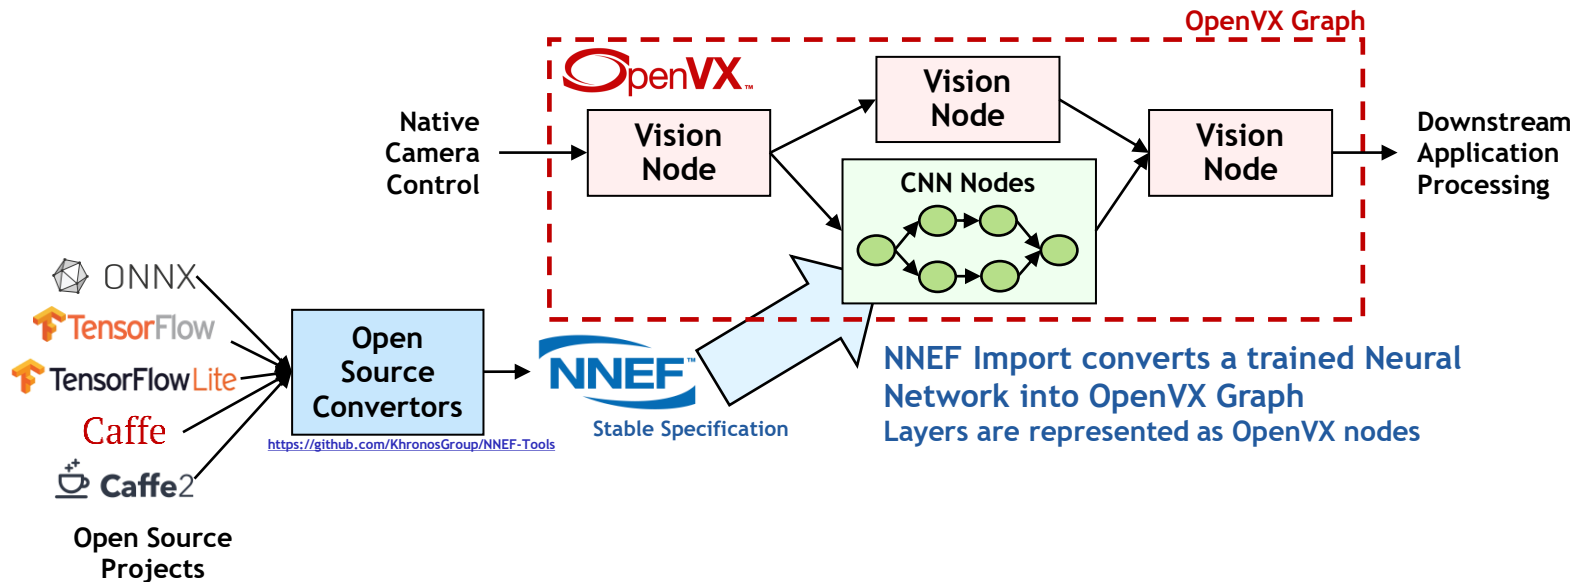


Vision Processing Graph

OpenVX Cross-Vendor Vision and Inferencing

High-level graph-based abstraction for portable, efficient vision processing

- Optimized OpenVX drivers created, optimized and shipped by processor vendors
- Implementable on almost any hardware or processor with performance portability
- Graph can contain vision processing and NN nodes for global optimization
- Run-time graph execution need very little host CPU interaction



OpenVX Efficiency through Graphs..

Graph Scheduling

Split graph execution across the whole system:
CPU / GPU / dedicated HW

Faster execution or lower power consumption

Memory Management

Reuse pre-allocated memory for multiple intermediate data

Less allocation overhead, more memory for other applications

Kernel Fusion

Replace a sub-graph with a single faster node

Better memory locality, less kernel launch overhead





Data Tiling

Execute a sub-graph at tile granularity instead of image granularity

Better use of data cache and local memory

Performance comparable to hand-optimized, non-portable code
Real, complex applications on real-world hardware
Much lower development effort and higher portability than hand-optimized code

OpenVX and OpenCV are Complementary

		
Implementation	Community driven open source library	API drivers implemented, optimized and shipped by hardware vendors
Conformance	Extensive OpenCV Test Suite but no formal Conformance program	Implementations pass Khronos Conformance Testing to use trademark and ensure cross vendor consistency
Scope	100s of imaging and vision functions Multiple camera APIs/interfaces	Tight focus on dozens of core hardware accelerated functions plus extensions and accelerated custom nodes. Uses external camera drivers
Inferencing	Deep Neural Network module to construct networks from layers for forward pass computations only. Import from ONNX, TensorFlow, Torch, Caffe	Neural Network layers and operations represented directly in the OpenVX Graph. NNEF direct import with multiple convertors
Acceleration	OpenCV 3.0 Transparent API (or T-API) enables function offload to OpenCL devices 	Implementation free to use any underlying API such as OpenCL. OpenCL for accelerated Custom Nodes 
Efficiency	OpenCV 4.0 G-API graph model for some filters, arithmetic/binary operations, and well-defined geometrical transformations	Graph-based execution of all Nodes. Optimizable computation and data transfer
IP Protection	None. Source code licensed under BSD. Some modules require royalties/licensing	Protected under Khronos IP Framework - Khronos members agree not to assert patents against API when used in Conformant implementations

OpenVX Extensibility

OpenVX core specification defines market-targeted feature sets

Baseline Graph Infrastructure (enables other Feature Sets)

Default Vision Functions

Enhanced Vision Functions

Neural Network Inferencing (including tensor objects)

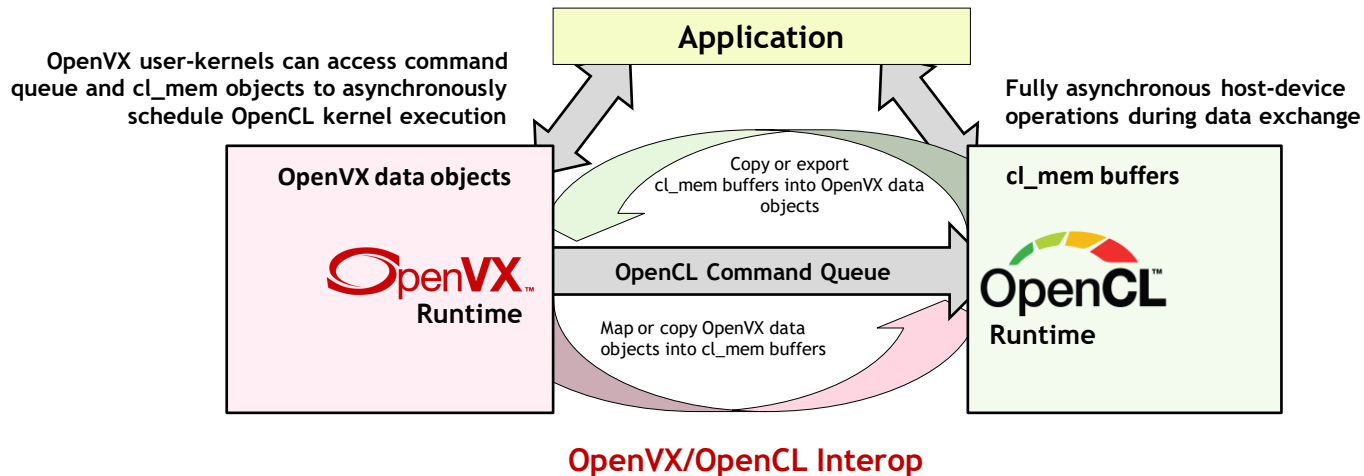
NNEF Kernel import (including tensor objects)

Binary Images

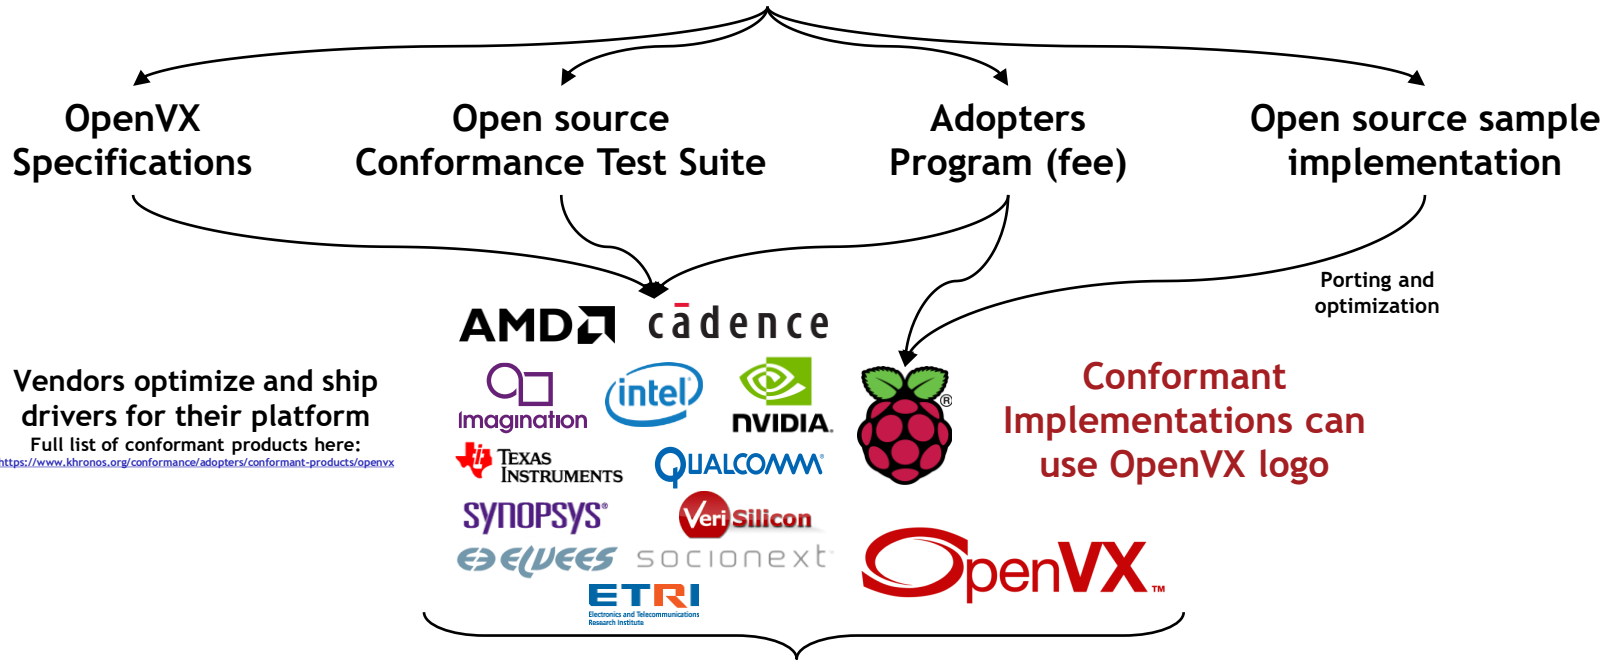
Safety Critical (reduced features and graph import for easier safety certification)

OpenVX is Extensible

Fully accelerated custom nodes can be integrated into the OpenVX graph with OpenCL interop



OpenVX Industry Deployment



Developers can freely use any OpenVX Implementation

Khronos for Global Industry Collaboration

Promoter Members

Participate and vote in Working Groups, Board seat for setting strategy and budget

Conformance is Key

Comprehensive testing frameworks available

Adopters

Build conformant implementations

Developers

Freely develop software using Khronos standards

KHRONOS
GROUP

Promoter, Contributor, Non-Profit, Associate, and Academic Members



Conformance Tests, Adopters Program

Ratified Specs, SDKs, Samples, Reference Cards

Adopters

Developers

Education

Contributor Members

Participate & vote in Working Groups

Non-Profit, Associate, and Academic Members

Participate in Working Groups

Working Groups

For each Standard, open to all members

Specifications & Learning Materials

Public & free of charge

Ecosystem

Samples, tools, webinars, tutorials, meetups

Khronos membership is open to any company

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

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

Provide industry thought leadership and gain insights into industry trends and directions

Benefit from Adopter discounts

www.khronos.org/members/