LINARO CONNECT 23

HKG18

George Grey, Linaro CEO

Welcome

China Electronics Corporation (CEC)

CEC 中国电子

- Joining Linaro as a Core Member
- Major Chinese Global IT Supplier
- China GCC Board Member
- 100+ companies, 14 listed
- 130,000 employees

银河麒麟 KYLIN

飞腾 PHYTIUM

HDSC 华大半导体有限公司 HUADA SEMICONDUCTOR CO.,LTD.

Welcome

- Joining Linaro as a LITE Member
- 96Boards SC Appointment
- LITE Group Associate

Engineering Highlights - a few of many

IOTL/Tiny Linux

OpenDataPlane
.org

DeveloperCloud
www.linaro.cloud

Linaro
Downloads
linaro.org/downloads

Zephyr

SANECHIPS
中 兴 微 电 子

Rosefinch 7100
NB-IoT Chip

Linaro Value

Unique Company where competitors in the Arm Ecosystem can collaborate on Open Source projects of benefit to all

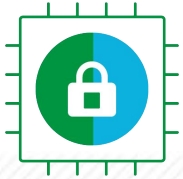

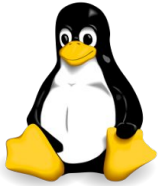

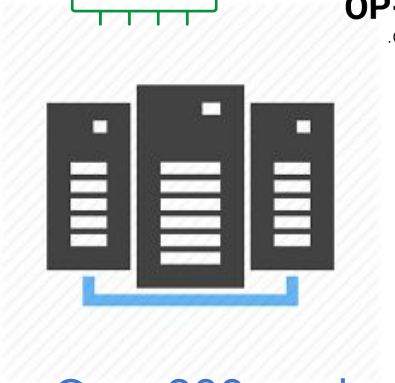

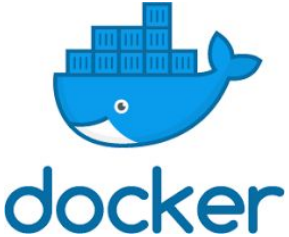

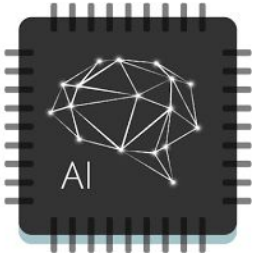

Over 300 engineers from companies across the world working together in over 70 Open Source Projects

Some of our work ...

I skate to where the puck is going to be, not where it has been

Wayne Gretzky

Our world is changing

Unprecedented Connectivity

Increasing Product Complexity

Data Everywhere - Cloud, Edge, Device

5G

Building Products with Linux

Last stable Linux kernel	4.15.10	Mar 2018
Most recent LTS kernel	4.14	Nov 2017 first release
SoC BSP #1	4.4	Jan 2016 first release
SoC BSP #2	3.18	Dec 2014 first release
Oldest current LTS kernel	3.2	Jan 2012 first release

How About Security?

Latest kernel software (upstream)

Has the latest security features

Gets security patches first

Supports the most hardware

LTS kernels get key patches backported

A Real (Embedded) Kernel

SoC BSP kernel (say 3.18)
SoC vendor code (inc. blobs)
Dev Board vendor code
End Product kernel changes
Updated LTS patches
SoC vendor updates
Final “Product Kernel” patches

Tested Product Kernel

“Based on Linux”

3.18 is about 150,000 patches behind upstream

Expect 100K-1M+ lines of “vendor code”

The result is a completely custom kernel

An industrial or automotive product lifetime is 10+ years

High lifetime maintenance costs

Delays to security updates

A Not Uncommon Example

2017 Model Car

IVI System with Bluetooth uses kernel version 3.1.10

Released January 18th 2012

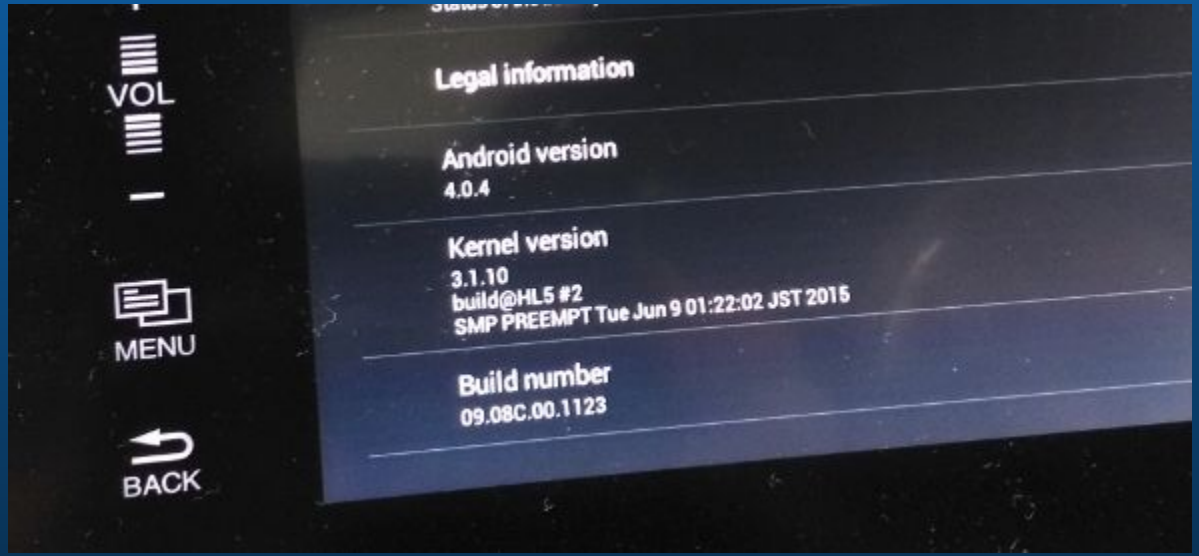

We Can Do Better

Is LTS still the best approach?

If the best software is the latest software why can't we use it?

The development cycle is too long

Certification and Compliance

Risks of making changes late in the process

And, unless it's an enterprise SoC, users don't have the option

Could we solve this?

Start with the latest (upstream) software

Design for continuous updates through the product lifetime

Use CI with use case and regression testing

Security updates, bug fixes and new features

Deliver more secure products at lower cost

What do we need to do?

Support new SoCs upstream

Improve CI testing - starting with the Kernel

Build better and better functional & regression tests

Test upstream - nothing else scales

Test in minutes, not hours or days

SoC Vendors

Every vendor should have an upstream focused team

Linaro can help

Reference Development Builds staging to upstream (like ERP)

Developer Services and Landing teams to help your engineers

Linaro

Ecosystem Development

Machine Intelligence

Autonomous Vehicles

Machine Intelligence

Early Market Adoption

AI/ML software is largely open source software

Algorithms & hardware IP are proprietary

Need to interface hardware IP to AI/ML frameworks

We need developer platforms for Datacenter, Edge & Fog

High performance, low power

Linaro Machine Intelligence Project

Provide best in class Deep Learning performance by leveraging Neural Network acceleration in IP and SoCs from the Arm ecosystem, through collaboration on the most widely adopted AI/ML software frameworks and libraries

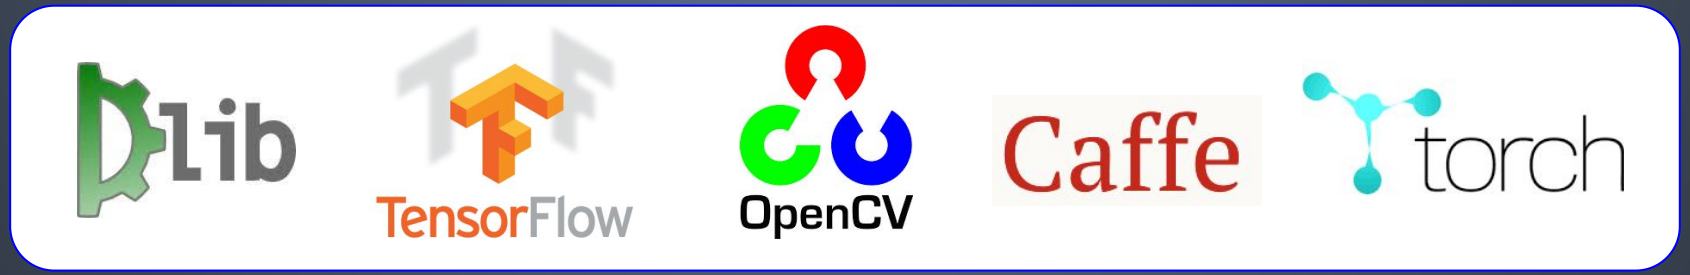

96Boards
DragonBoard 820C

DragonBoard 820C

- Uses Snapdragon 820E
  - Quad core Kryo CPU
  - 96Boards CE Extended Version
  - 3GB LPDDR4 RAM
  - 32GB UFS Flash
  - WLAN 802.11ac, Bluetooth 4.2, GPS
  - GBit Ethernet, USB 2.0/3.0
  - HDMI 4K@60fps
  - mini-PCIe
- CPU, Adreno 530 GPU and Hexagon 680 DSP
  - Heterogeneous processing
  - Linux developer support later in 2018
- Now shipping to developers
- Plan to support 96Boards.ai

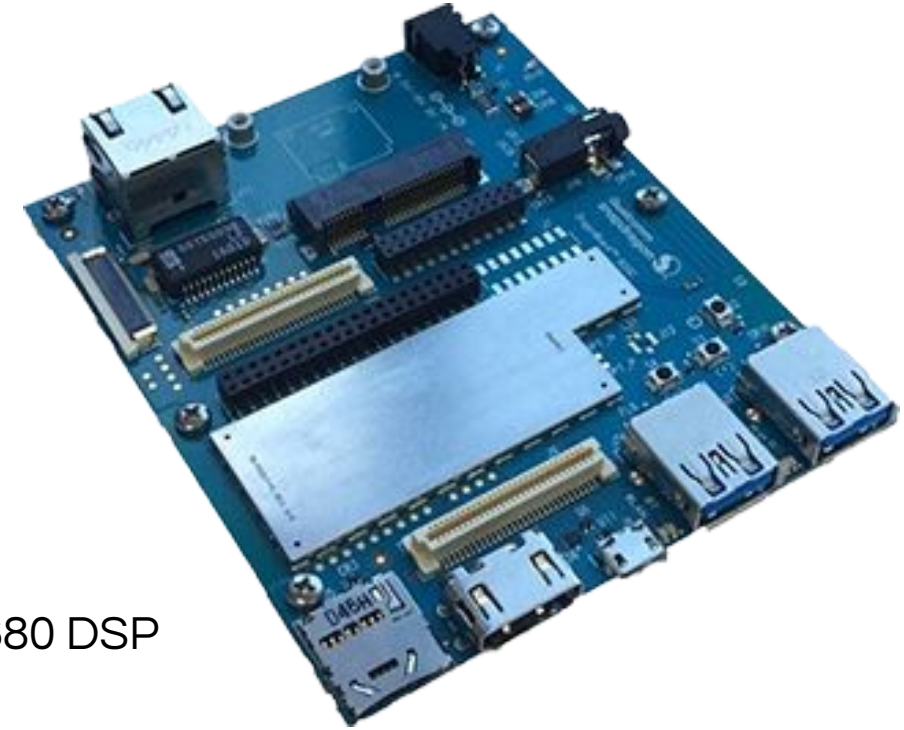

Built by Arrow

96Boards Manufacturing Partner

96Boards.ai

Announcing 96Boards.ai
Hardware for AI/ML Developers

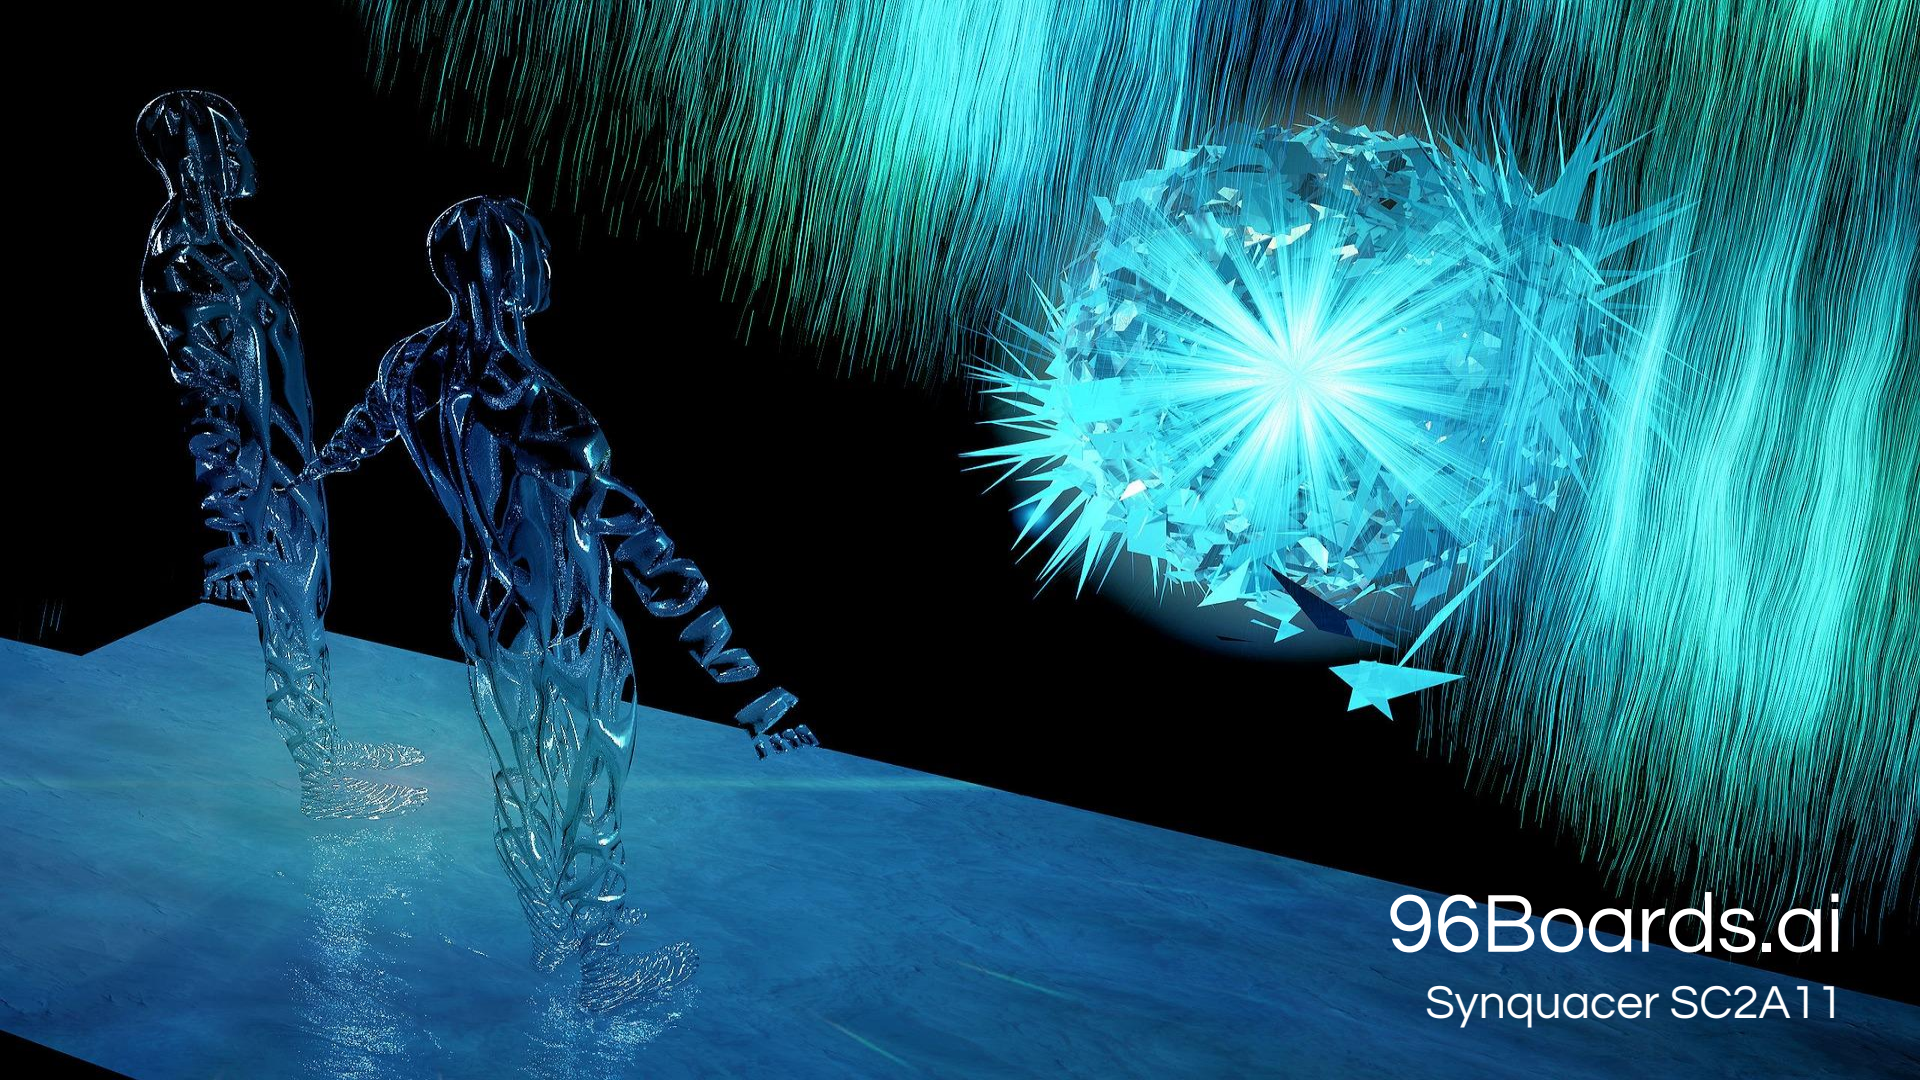
96Boards.ai
Synquacer SC2A11

Socionext Synquacer + Gyrfalcon Lightspeeur

Synquacer SC2A11

- A53@1GHz x24, 5W

96Boards EE Card

- microATX format
- up to 64GB RAM
- SATA, GBE
- DeveloperBox

Gyrfalcon Lightspeeur

- Neural Compute
- 28K cores/chip
- 9.3TOPS/W
- 8 chips on PCIe card

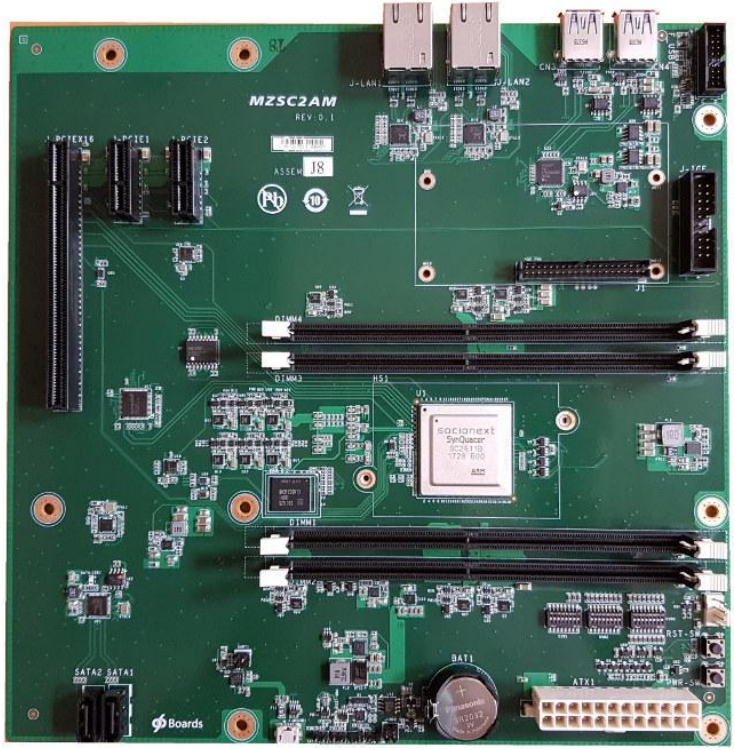

socionext™

gyrfalcon technology

Demo - Object Recognition in Multiple Simultaneous Video Streams

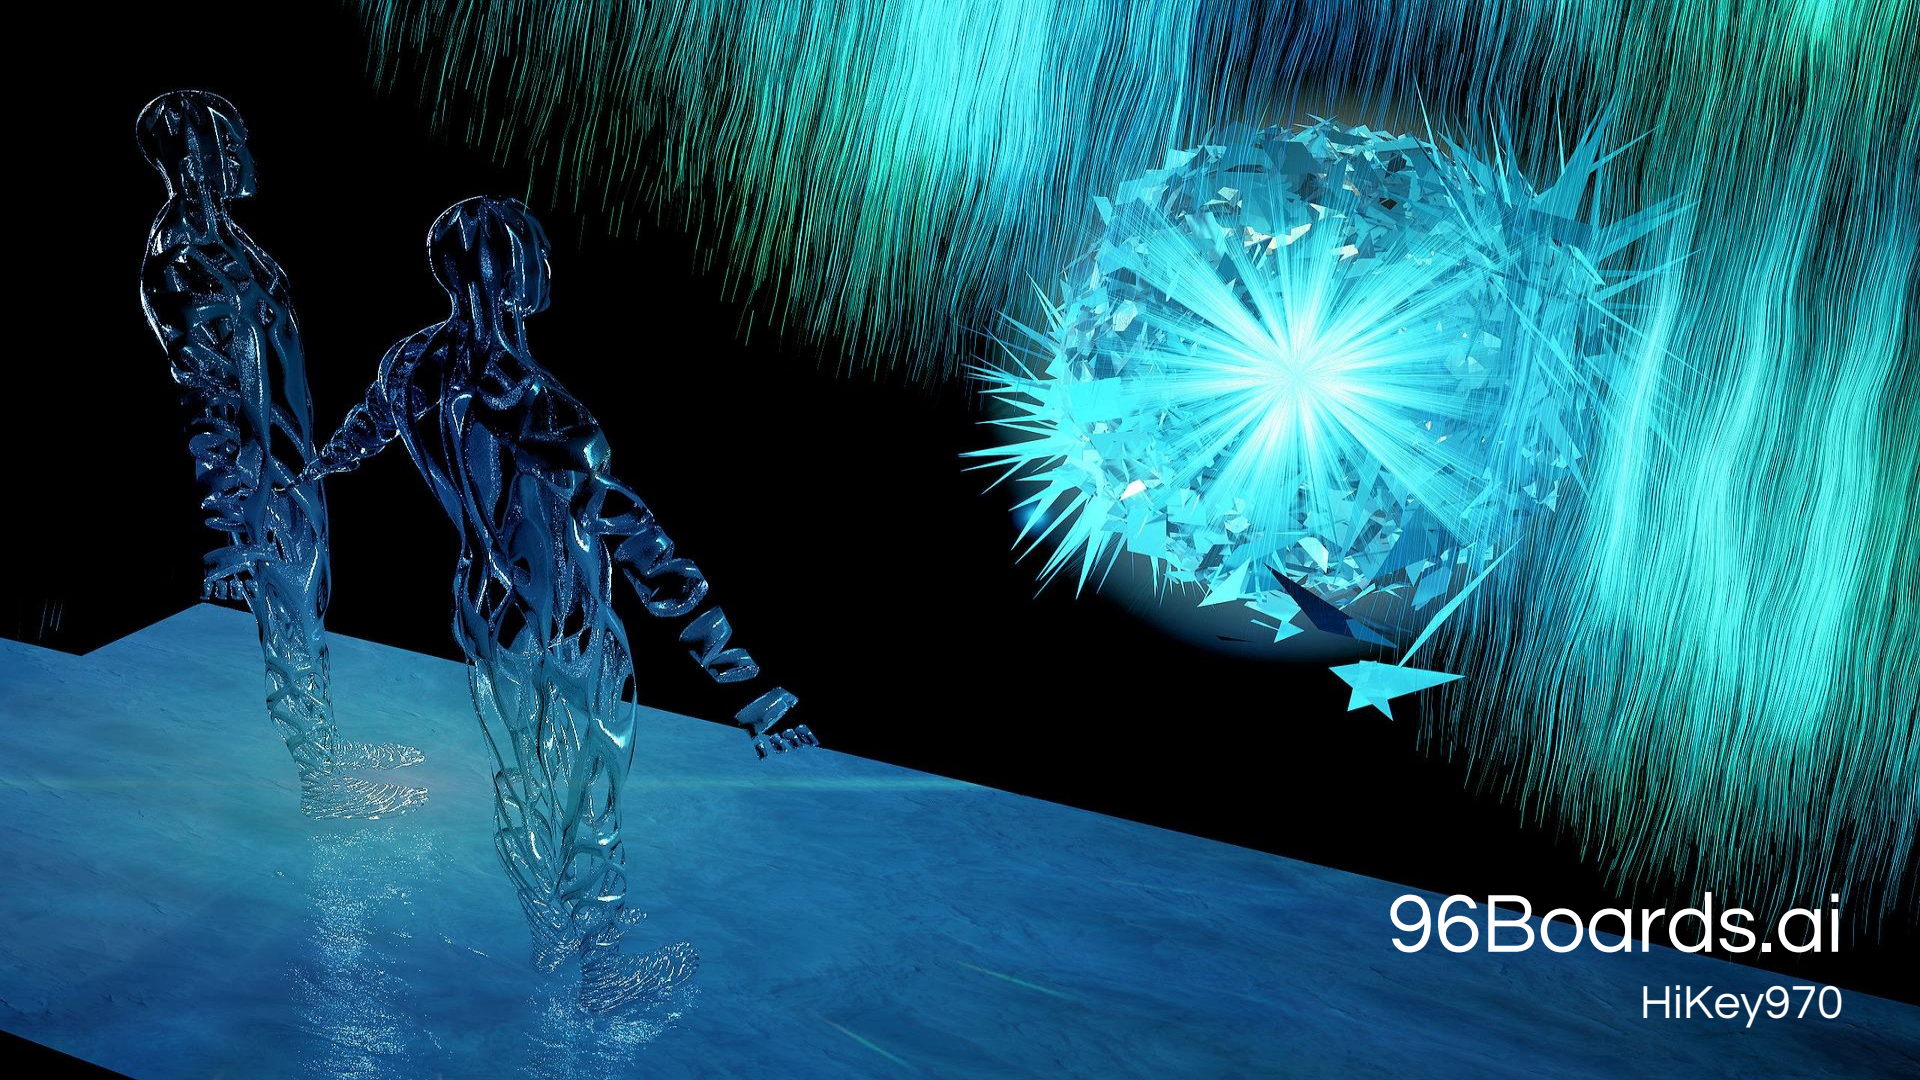
96Boards.ai
HiKey970

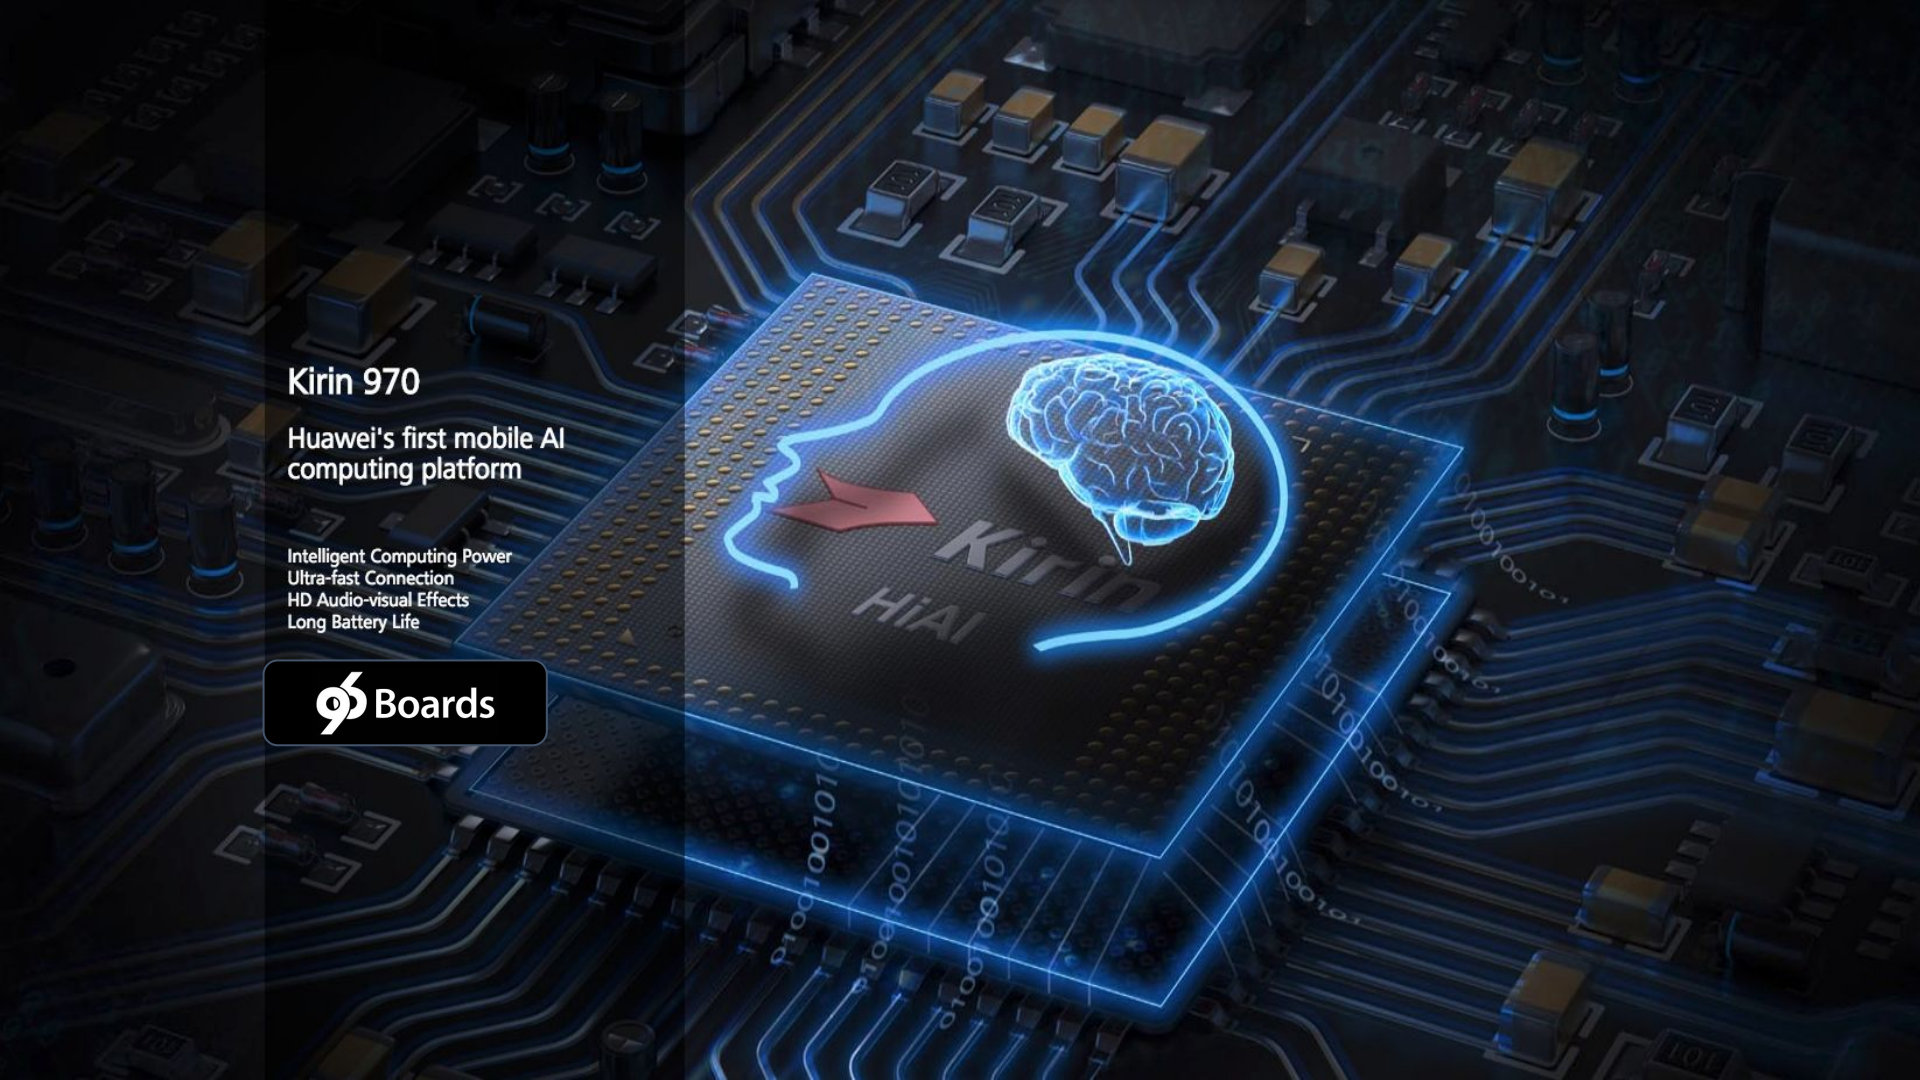
Kirin 970
Huawei's first mobile AI computing platform
Intelligent Computing Power
Ultra-fast Connection
HD Audio-visual Effects
Long Battery Life
96 Boards
Kirin
HiAI

First Super Edge AI Computing Platform with NPU HiKey 970

- Huawei HiAI SDK
  - Up to 25x Performance
  - Up to 50x Power Efficiency
- Dedicated Neural-network Processing Unit (NPU)
- Heterogeneous Resource Management

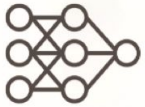

Deep Learning

Robots

Smart City

Automobile

Huawei HiAI SDK Empowers Developers

Huawei HiAI SDK provides AI compute libraries and APIs which enable developers to easily and effectively develop new AI applications instead of focusing on performance tuning on edge devices

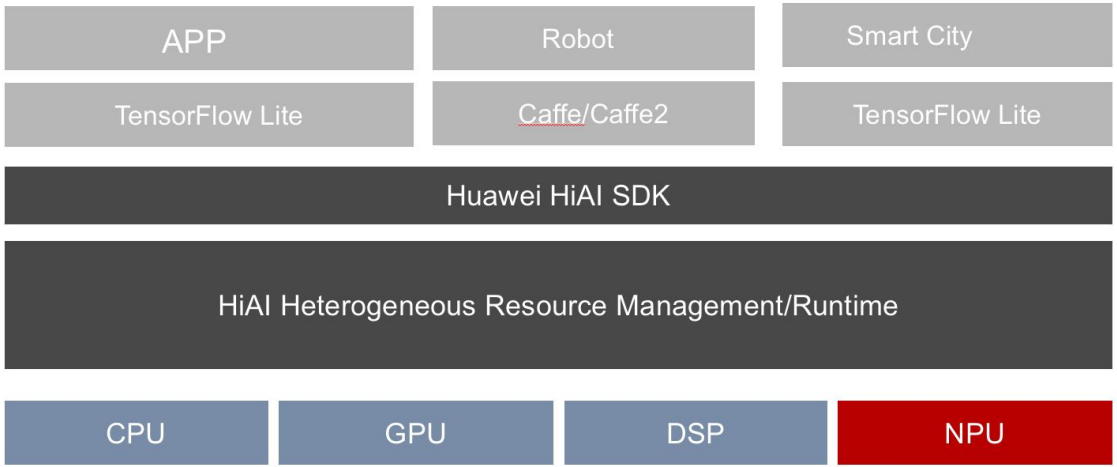

Tool Chain

Documentation

Rich set of APIs

Easy-to-use Source Code

	HiAI SDK V100	HiAI SDK V150
Framework	Caffe, TensorFlow	Caffe, TensorFlow, TensorFlow Lite, Huawei HiAI SDK, Android NN
Operators	42	99
Tools	Command Line	Android Studio plug-ins Compile Tool Sparsifying tool Code auto-generated Error code reporting
Documentation And Support	User Manual Source Code FAE	FAQ Open classes Sample code & docs More FAE

HiKey970 Empowers More Product Lines

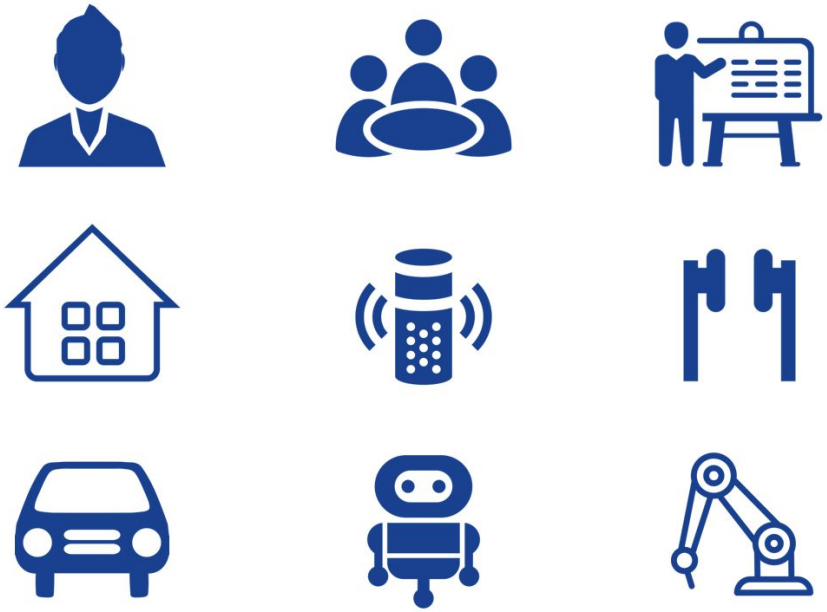

Popular AI Stacks

Mainstream OS

More Hardware Interfaces

High Performance Compute

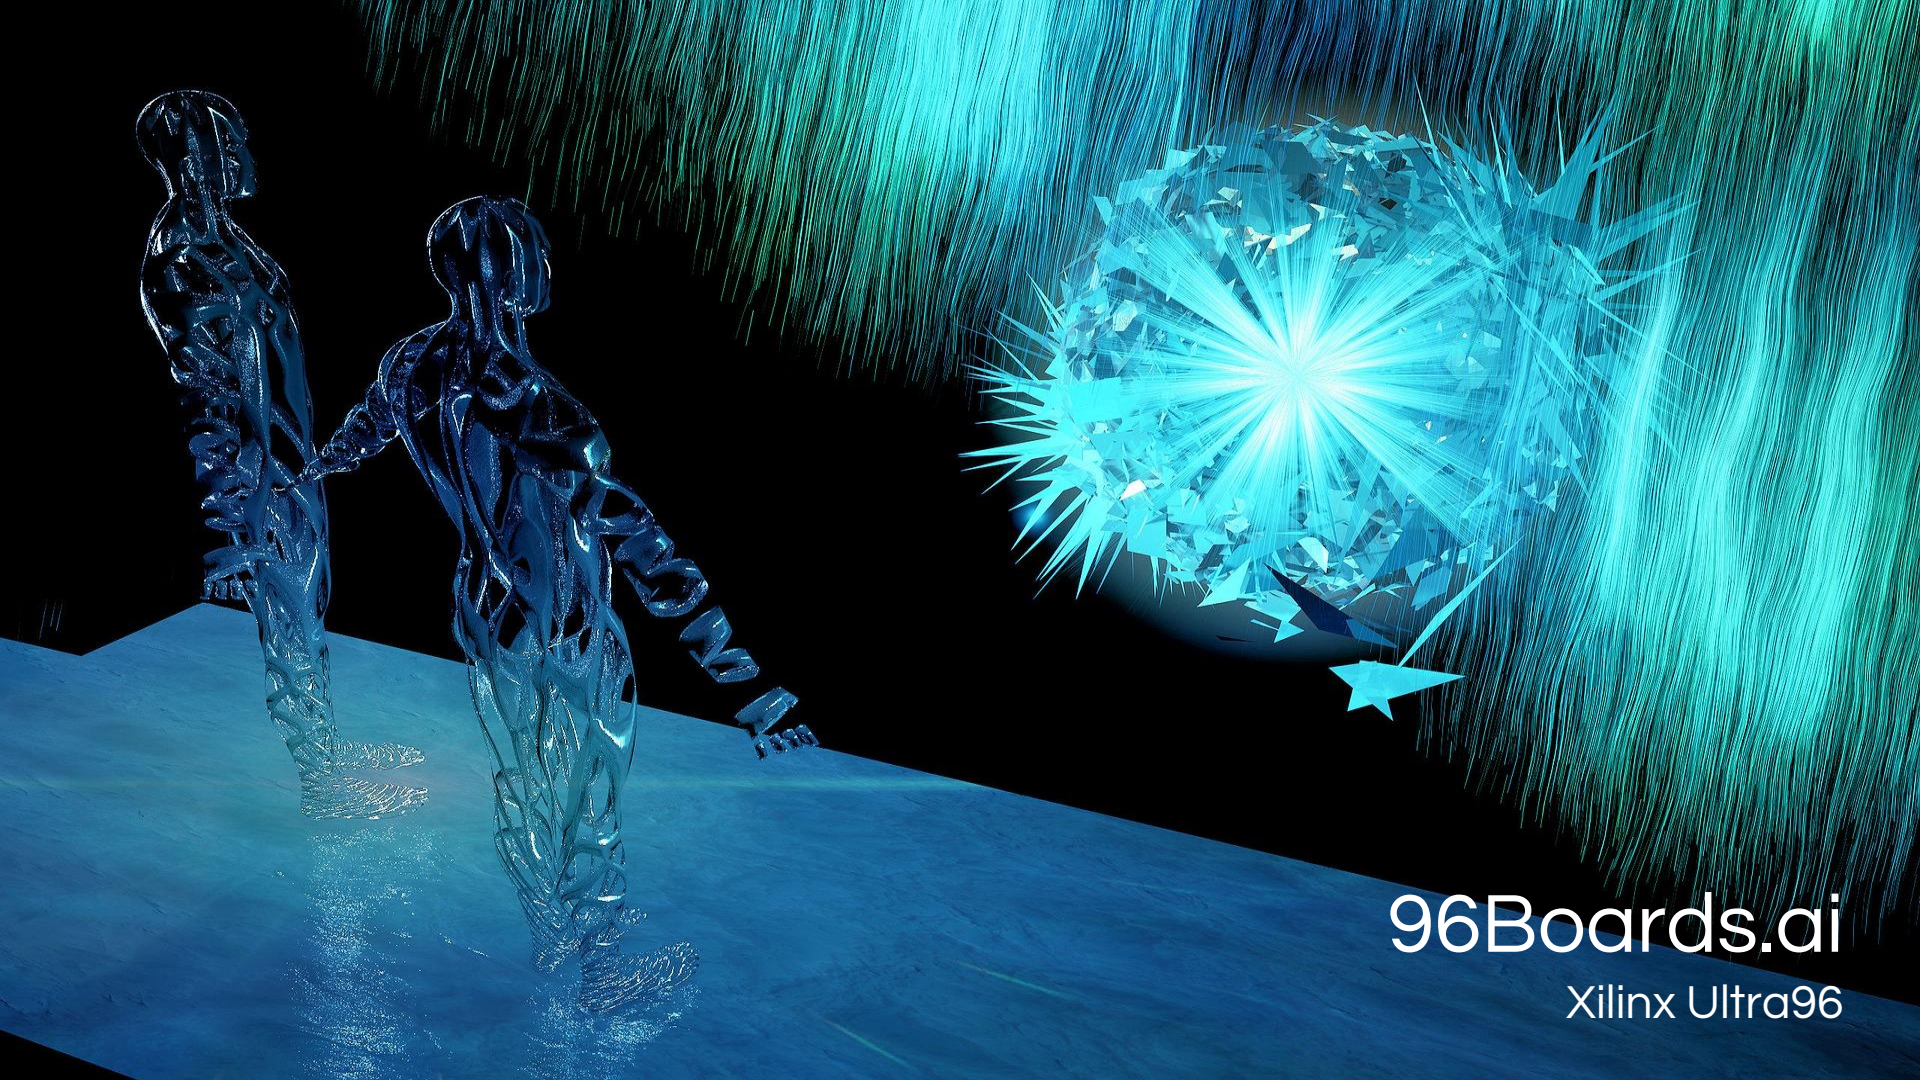
96Boards.ai
Xilinx Ultra96

Ultra96 Features the Zynq® UltraScale+ MPSoC

Memory Subsystem
High Bandwidth
Low Latency

Application Processor
64-bit Quad-Core

Graphics Processor
ARM Mali-400MP2

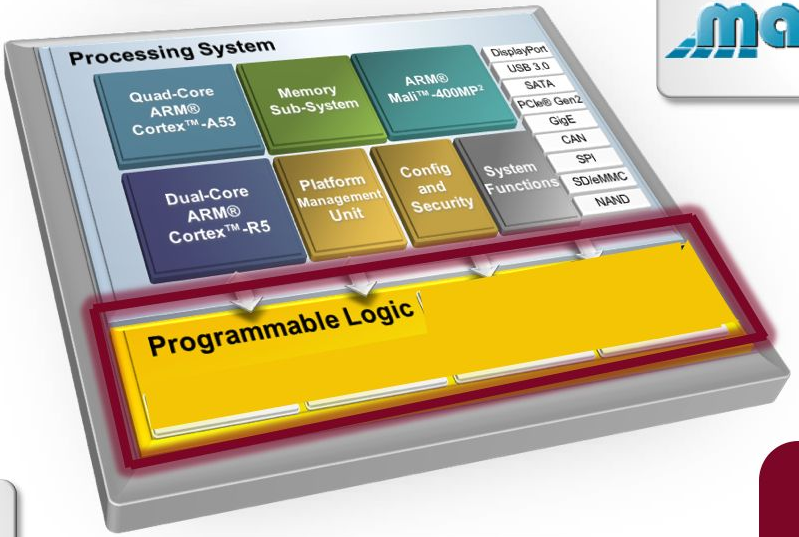

Real-Time Processors
32-bit Dual-Core

Fabric Acceleration
Customizable Engines
High Speed Connectivity

Platform & Power Management
Granular Power Control
Functional Safety

Configuration & Security Unit
Anti-Tamper & Trust
Industry Standards

Image processing
Machine learning
Computer vision
…virtually anything else

© Copyright 2018 Xilinx

FPGA Logic Example Application: Tiny YOLO Object Detection

Dotted line encapsulates FPGA logic

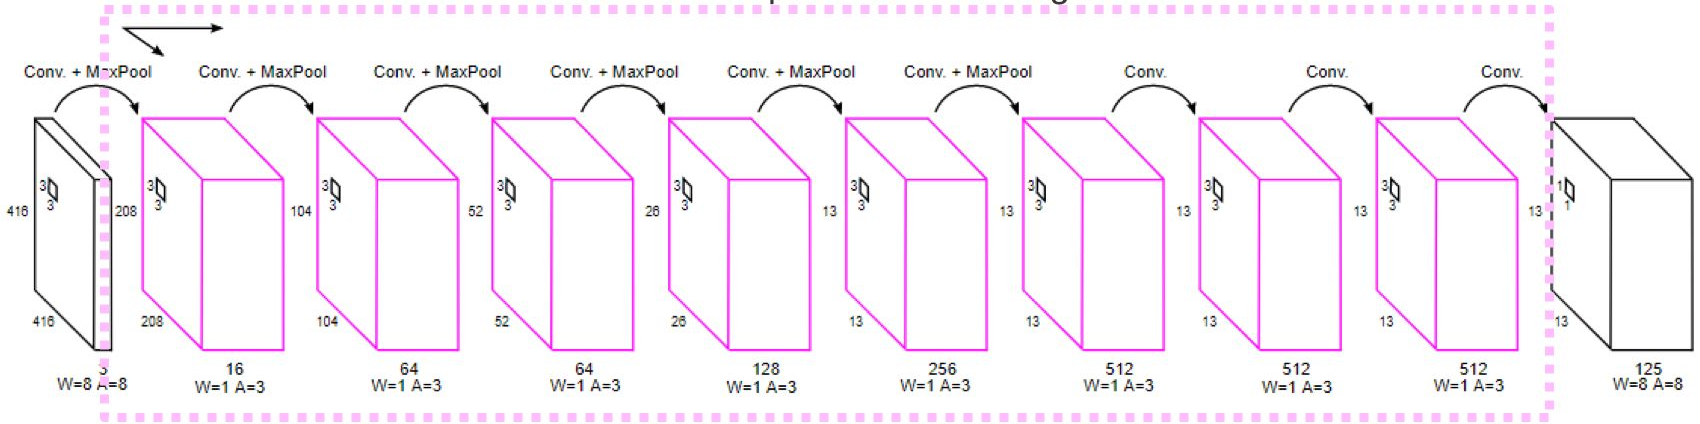

Trained using PASCAL Visual Object Classes

- Person: person
- Animal: bird, cat, cow, dog, horse, sheep
- Vehicle: airplane, bicycle, boat, bus, car, motorbike, train
- Indoor: bottle, chair, dining table, potted plant, sofa, tv/monitor

© Copyright 2018 Xilinx

Tiny Yolo On An Image: Performance

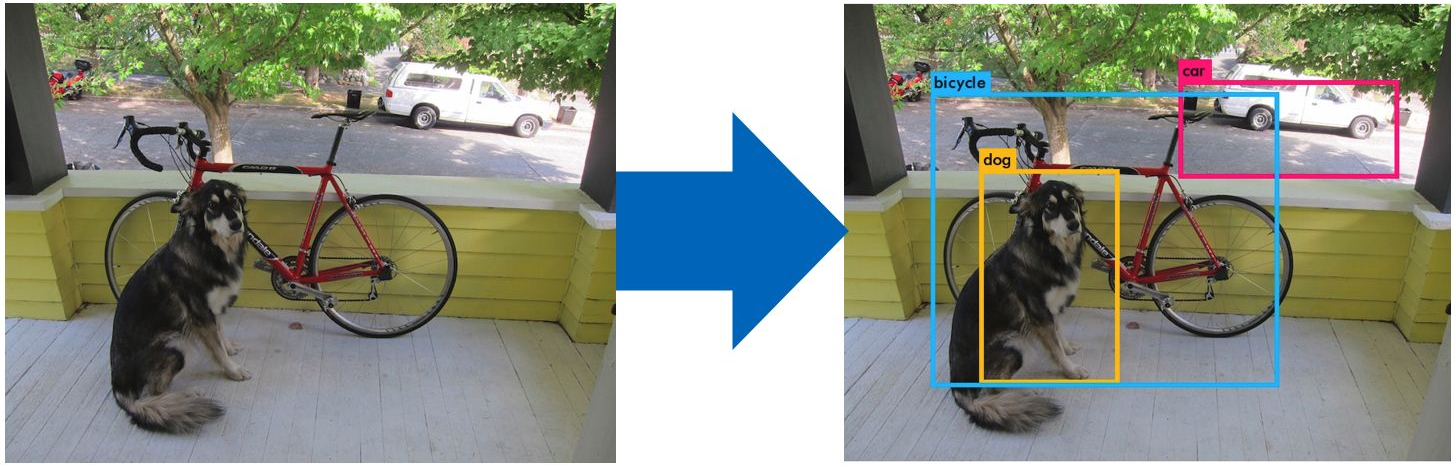

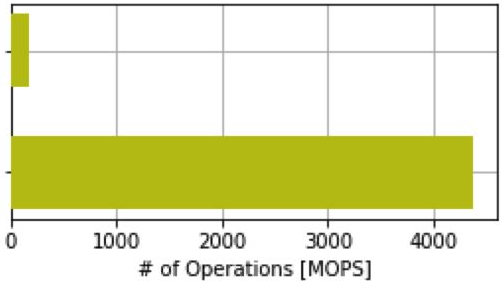

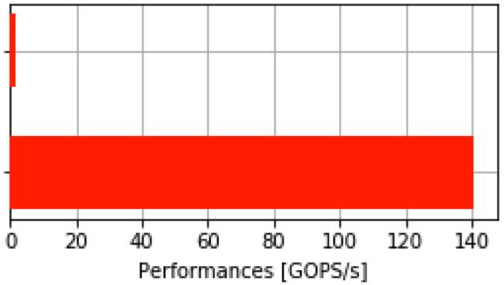

© Copyright 2018 Xilinx

Demo - Ultra96 Live HDMI Feed

ultra96.org

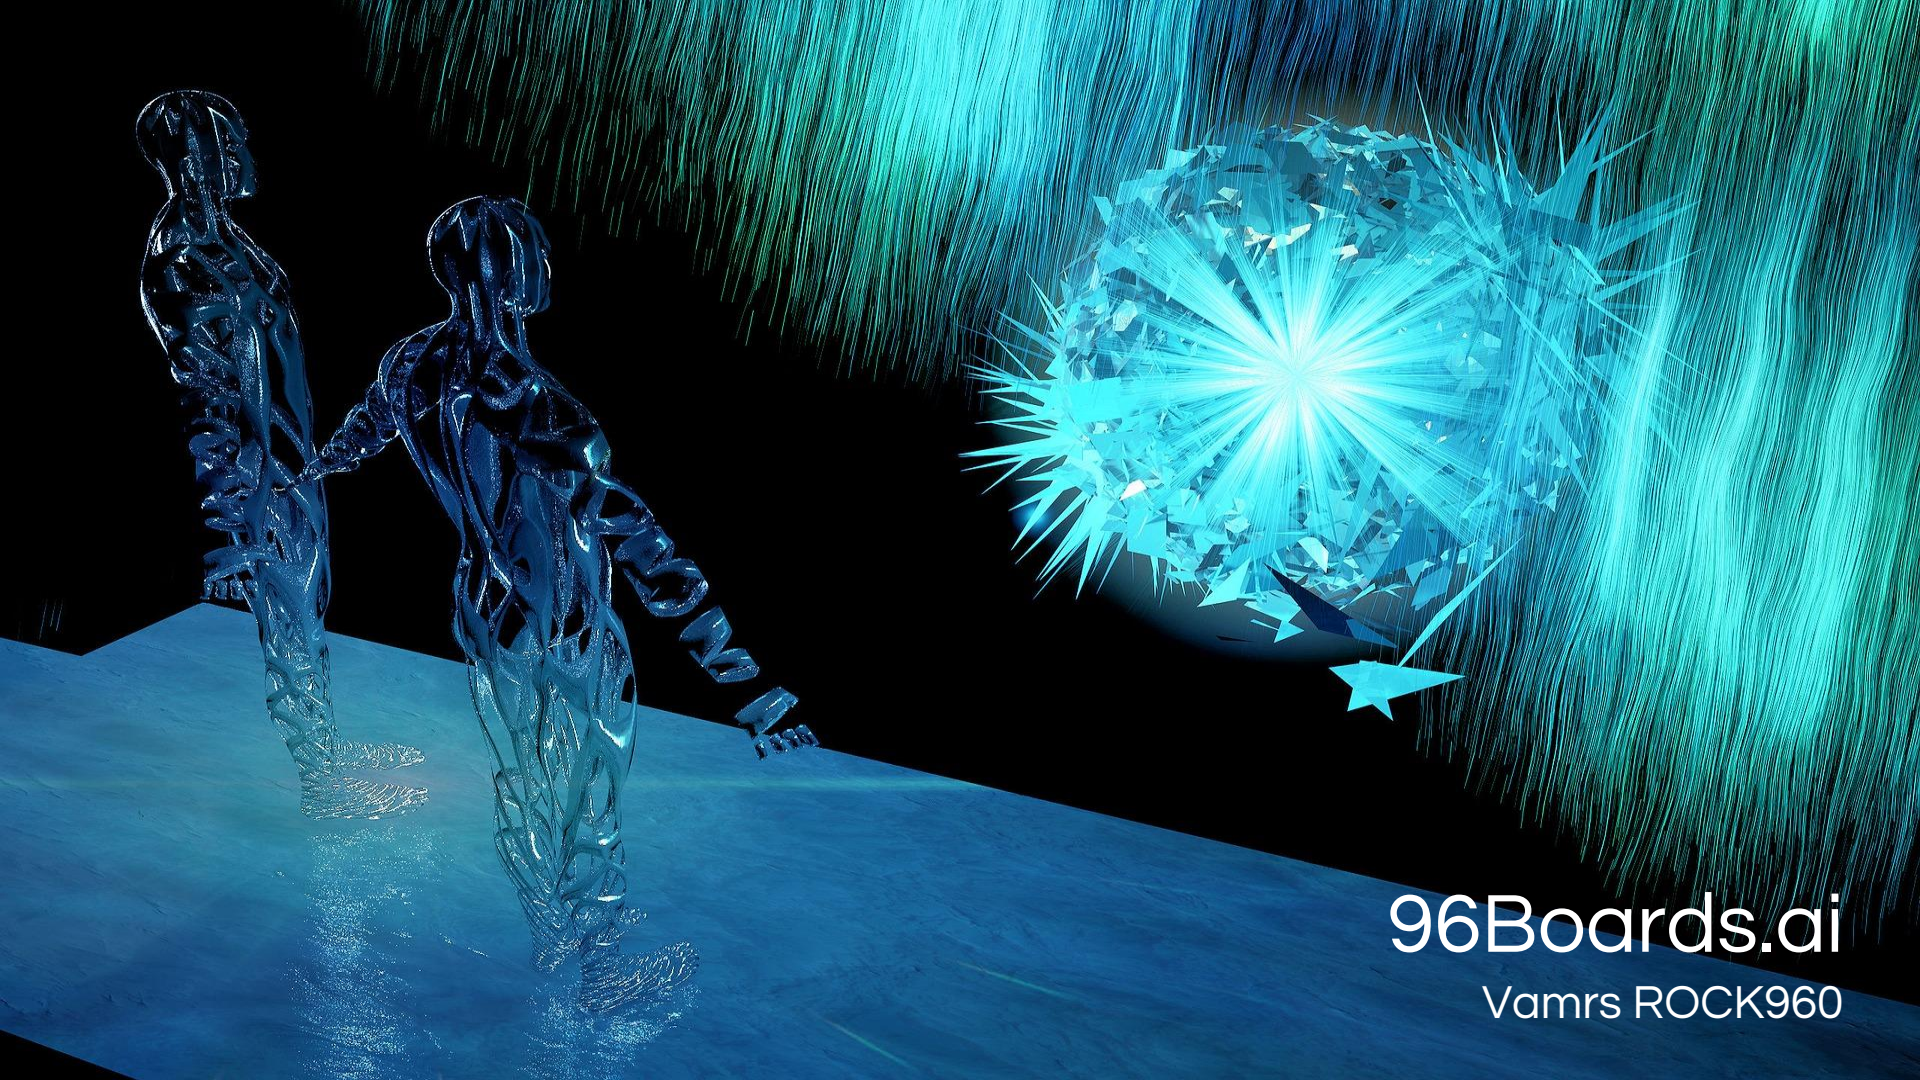
96Boards.ai
Vamrs ROCK960

Empower your intelligence

Introducing ROCK960 - AI Ready Platform

www.96rocks.com

Designed and
Built by VAMRS

96Boards Manufacturing Partner

Rockchip

VAMRS

ROCK960

- RK3399PRO CPU Complex*
  - 2x Cortex-A72 2GHz
  - 4x Cortex-A53 1.5GHz
  - Mali-T860-4
  - Up to 4GB RAM
- Dedicated NPU
  - Up to 2.4 TOPs
- High Speed Interfaces
  - USB 3.0
  - mini-PCIe
  - 802.11ac WIFI

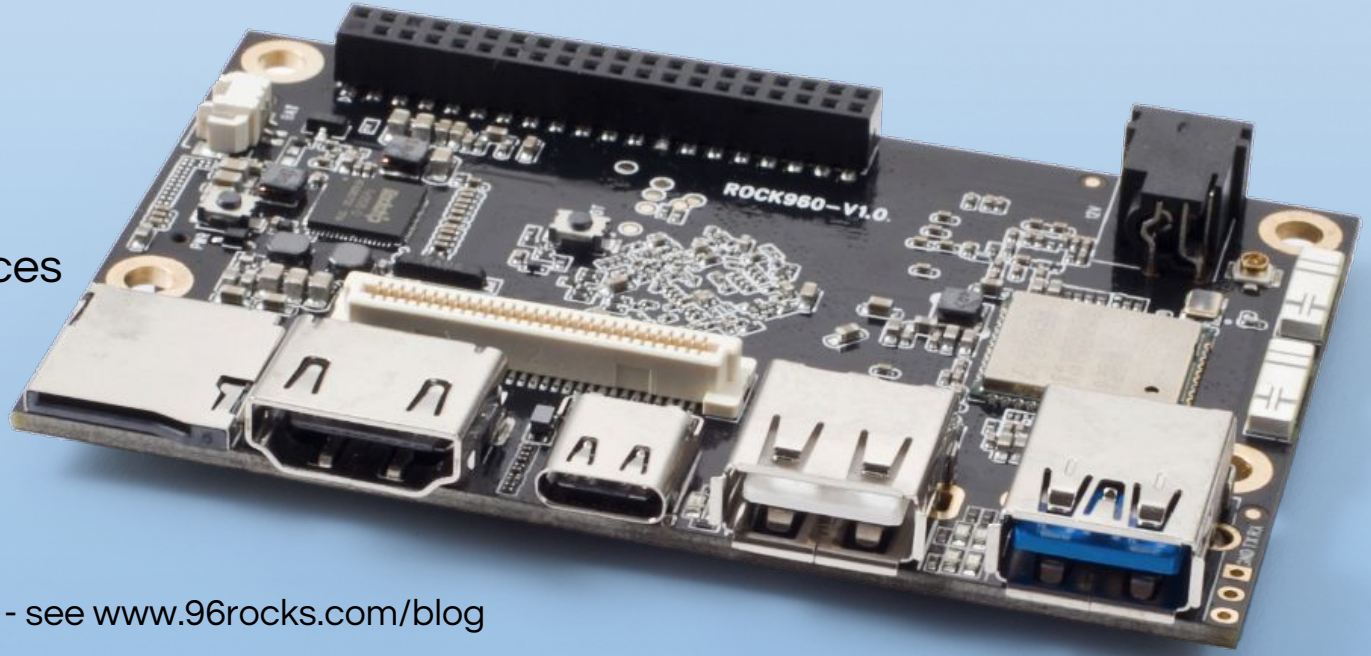

*PRO version Release Q2 - see www.96rocks.com/blog

ROCK960 - Enterprise Edition (Release Q2)

- RK3399PRO CPU Complex
  - Up to 8GB RAM
- Dedicated NPU
  - Up to 2.4 TOPs
- High Speed Interfaces
  - USB 3.0 x3, USB 2.0 x5
  - PCIe 2.1 x16 slot
  - Dual SATA 3.0, RAID 0/1
  - GBE, 802.11ac WIFI
  - HDMI 2.0/eDP 4K @ 60Hz
  - Dual MIPI CSI
- Remote Management
  - 4G LTE
  - Management Port

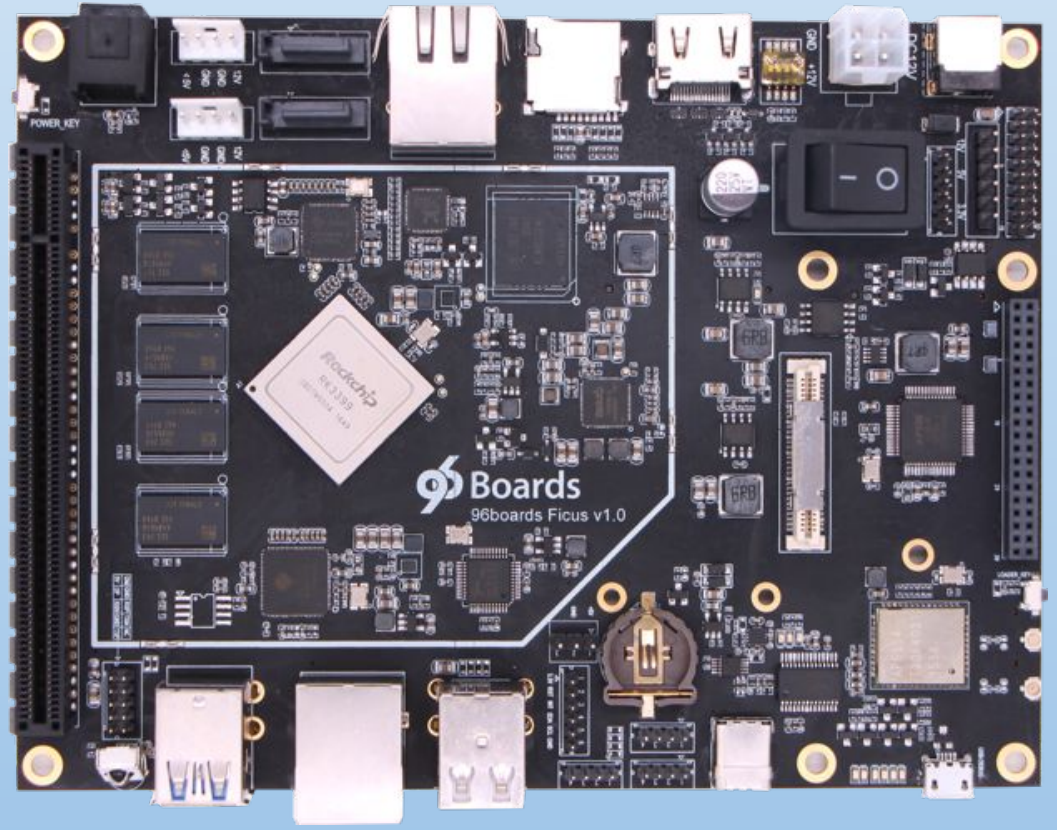

ROCK960 & Open AI Lab

Open AID on ROCK960

Open AID simplifies the application development of Vision and Speech on the edge. It brings Domain Libraries for Vision and Speech inferences along with a unified API for developers, Tengine with improved DL frameworks on Caffe, MXNet and TensorFlow for inference, and Heterogeneous Computing Library (**HCL**) for optimized Arm CPU and Mali GPU utilization

WHITE PAPER

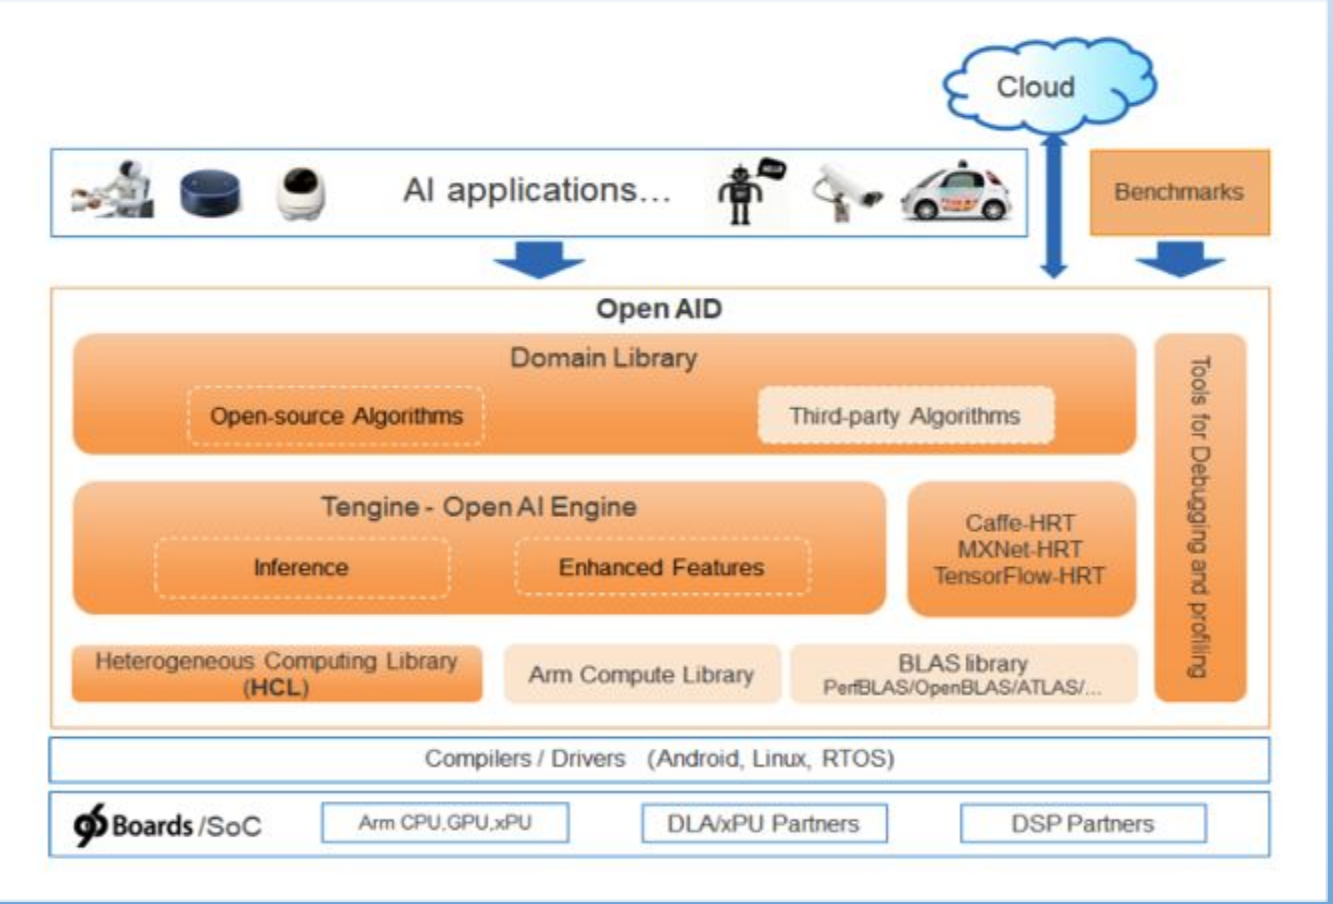

Jelly, an industry-oriented business robot with remarkable openness and customizability, is widely used in banks, hospitals, restaurants, shopping malls, etc.

招商銀行
CHINA MERCHANTS BANK

The China Merchants Bank customization project has successfully been completed. It would provide customized services in greeting clients and presenting business and other scenarios.

96Boards.ai

Socionext Synquacer EE with Gyrfalcon PCIe AI accelerator

HiKey970 CE Extended

Xilinx Ultra96 CE

Vamrs ROCK960 CE and EE

Arrow Qualcomm DragonBoard820C CE Extended*

* Support planned 2018 H2

Automotive

Linaro Autonomous Vehicle Project

Accelerate the Arm ecosystem in autonomous vehicle control by defining, adopting, improving and promoting open software platforms and scalable software architecture suitable for cars, other vehicles, robotics & drones

Where is the puck going to … 2021-2025

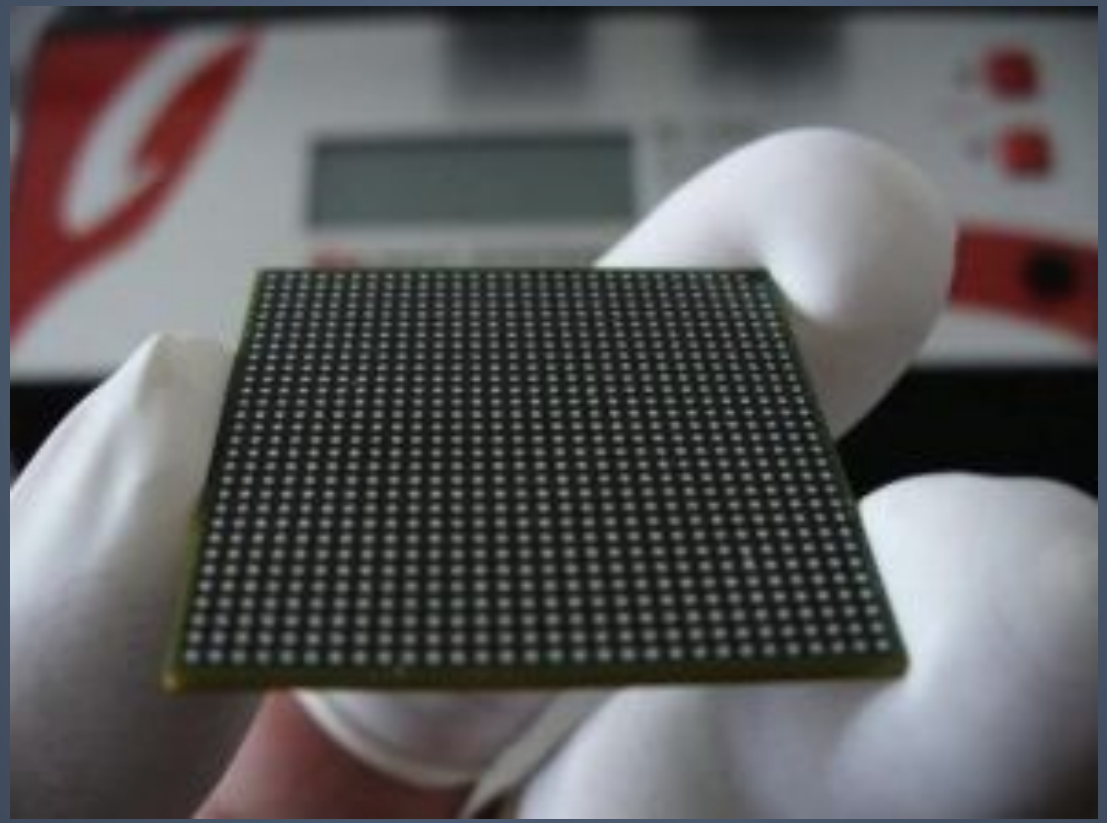

Where is the puck going to … 2021-2025

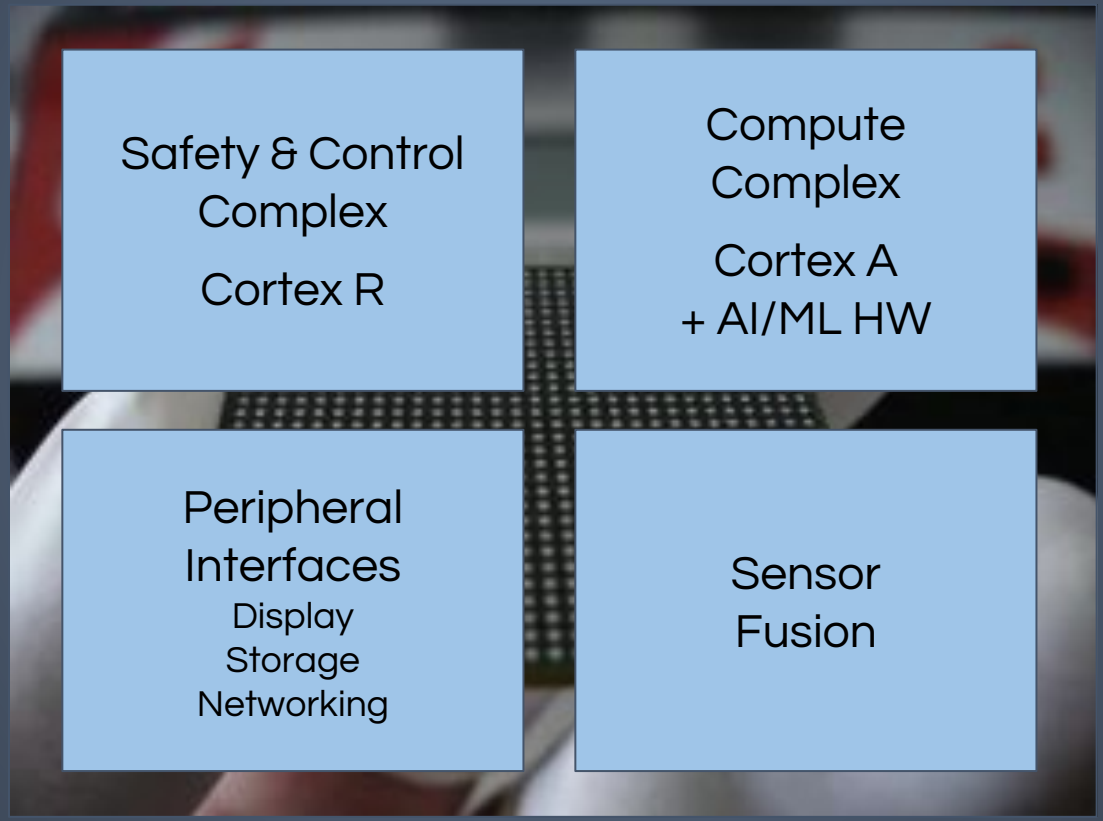

Architecting an Open Platform

Identify key open source technologies

Involve the OEMs/Tier I & II vendors

Build working Proof of Concept on commodity hardware

Boards, Software and Simulators/Cars

Iterate architecture and components

Enable substitution of proprietary technologies

Architecting a Platform

OTA Managed Containers

IVI • L5 Autonomy

Sensor Fusion • AI/ML Inference

Certified Safety RTOS

Secure, OTA Updatable minimal Linux

Certified uKernel/L1 Hypervisor

Secure, Trusted Firmware

Evaluating PoC Components

Kernkonzept L4RE microkernel • ACRN

Certifiable Zephyr

Open Source Foundries Linux microPlatform

Uptane • TUF

AGL

OSRF ROS 2.0 • Gazebo

PX4 Autopilot

Tier IV Autoware • Baidu Apollo

Streetdrone Autonomous EV

Please Collaborate, and Enjoy HKG18