FOUNDATION
OPENHW™
PROVEN PROCESSOR IP

# How to build an AI startup on open source RISC-V Cores

Frédéric Desbiens frederic.desbiens@eclipse-foundation.org
Jeremy Bennett jeremy.bennett@embecosm.com

# Agenda

- How to Build an AI Startup
- Are there any open source, industrial-grade RISC-V cores?
- Do you want a side of software?
- Join the fun!

# How to Build an AI startup with Open Source RISC-V Designs

## Case Studies

# AI for Inference

## Meta's in-house AI Accelerator

- Training is infrequent, but inference is forever
- Enterprise Datacenters
- Significant power consumption (100's of watts)
- Large engineering teams

- Yet, a simple design
  - Make a Processing Element (PE)
  - Copy it a few times
  - Add some IO's
  - Be the best in suggesting  yet another Cat video (YACV) to watch

# AI at the Edge

## Axelera.AI edge AI

- Low-power (Could run on battery)
- Suitable for deeply embedded applications
- Available in M.2 form factor

- Again, a simple design
  - Again: Common building blocks with RISC-V Controller
  - Important: Security on-chip!

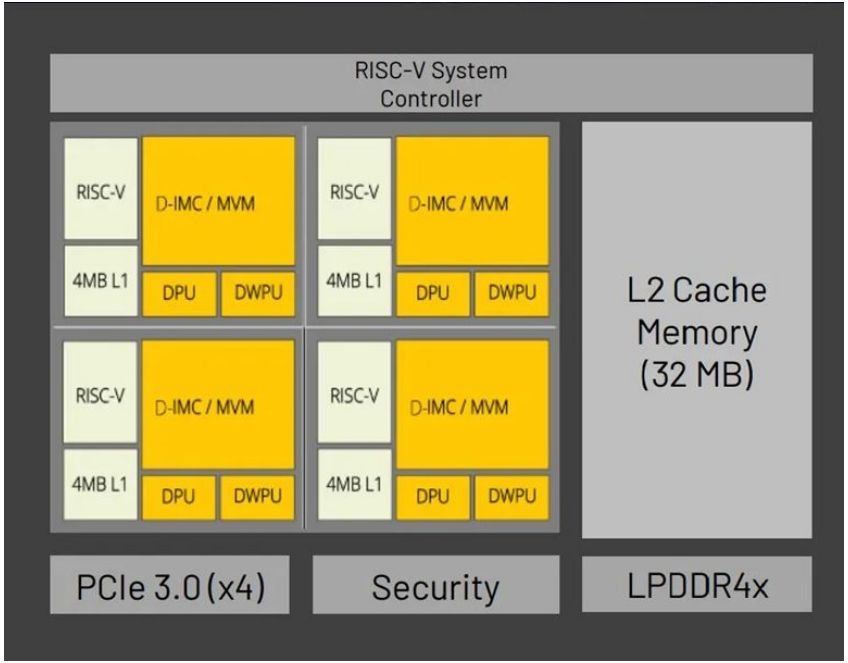

# Generic AI Building Blocks

## Andes Technology licensed AI building blocks

- RISC-V (maybe with Vector)
- DMA, General Matrix Multiply
- Extensive software stack
  - Compiler
  - RTOS
  - IDE
  - ...

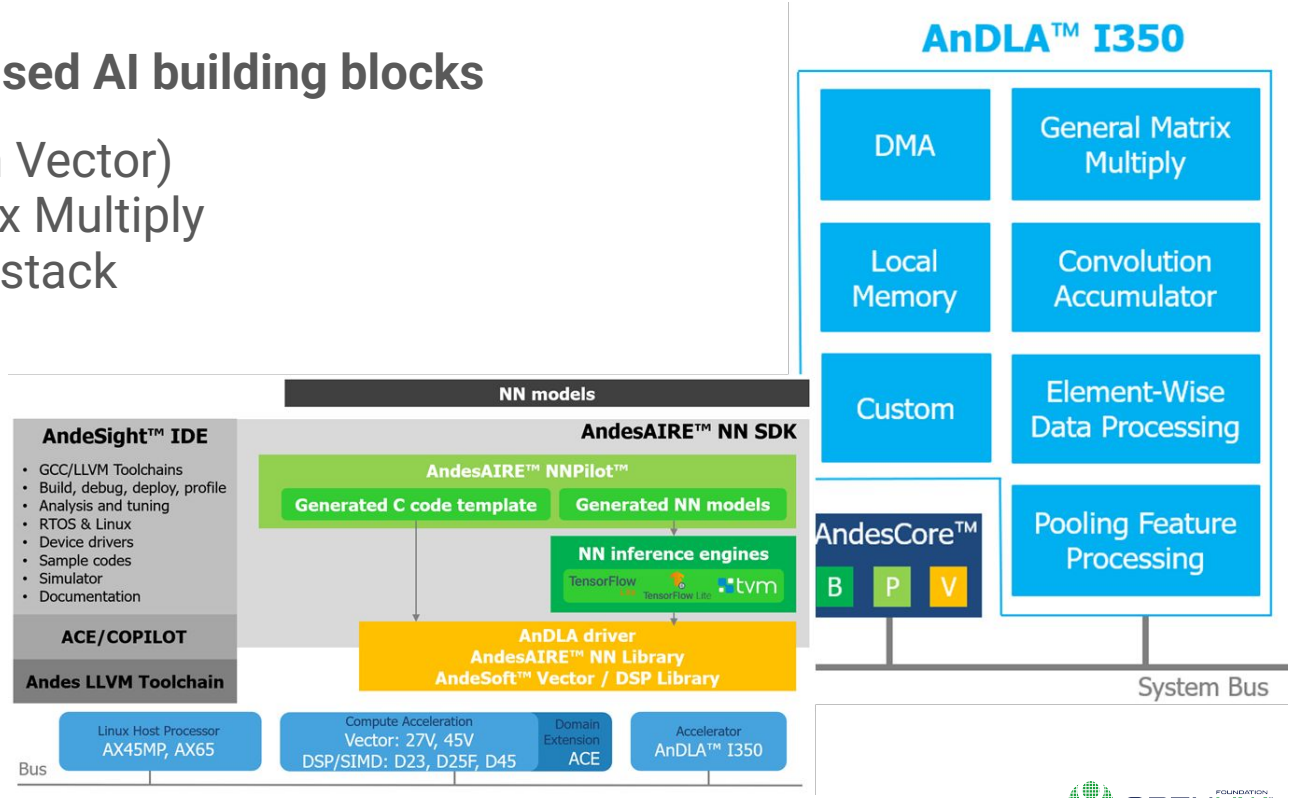

# AI in Your Pocket

## AI powered by RISC-V is already in your phone!

- Several ARM-based SOCs feature specialised cores powered by RISC-V
  - Yes, including “fruity” ones
  - Excludes the Nokia 3390

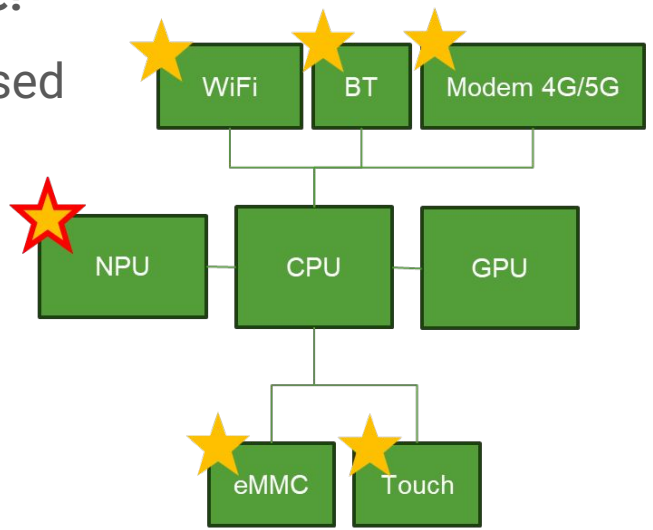


= RISC-V already today by some manufacturers, sometimes multiple Cores

# Where do I Find Industrial-Grade Open Source RISC-V Cores?

## OpenHW Foundation has them?

## For Free?

# OpenHW RISC-V Roadmap

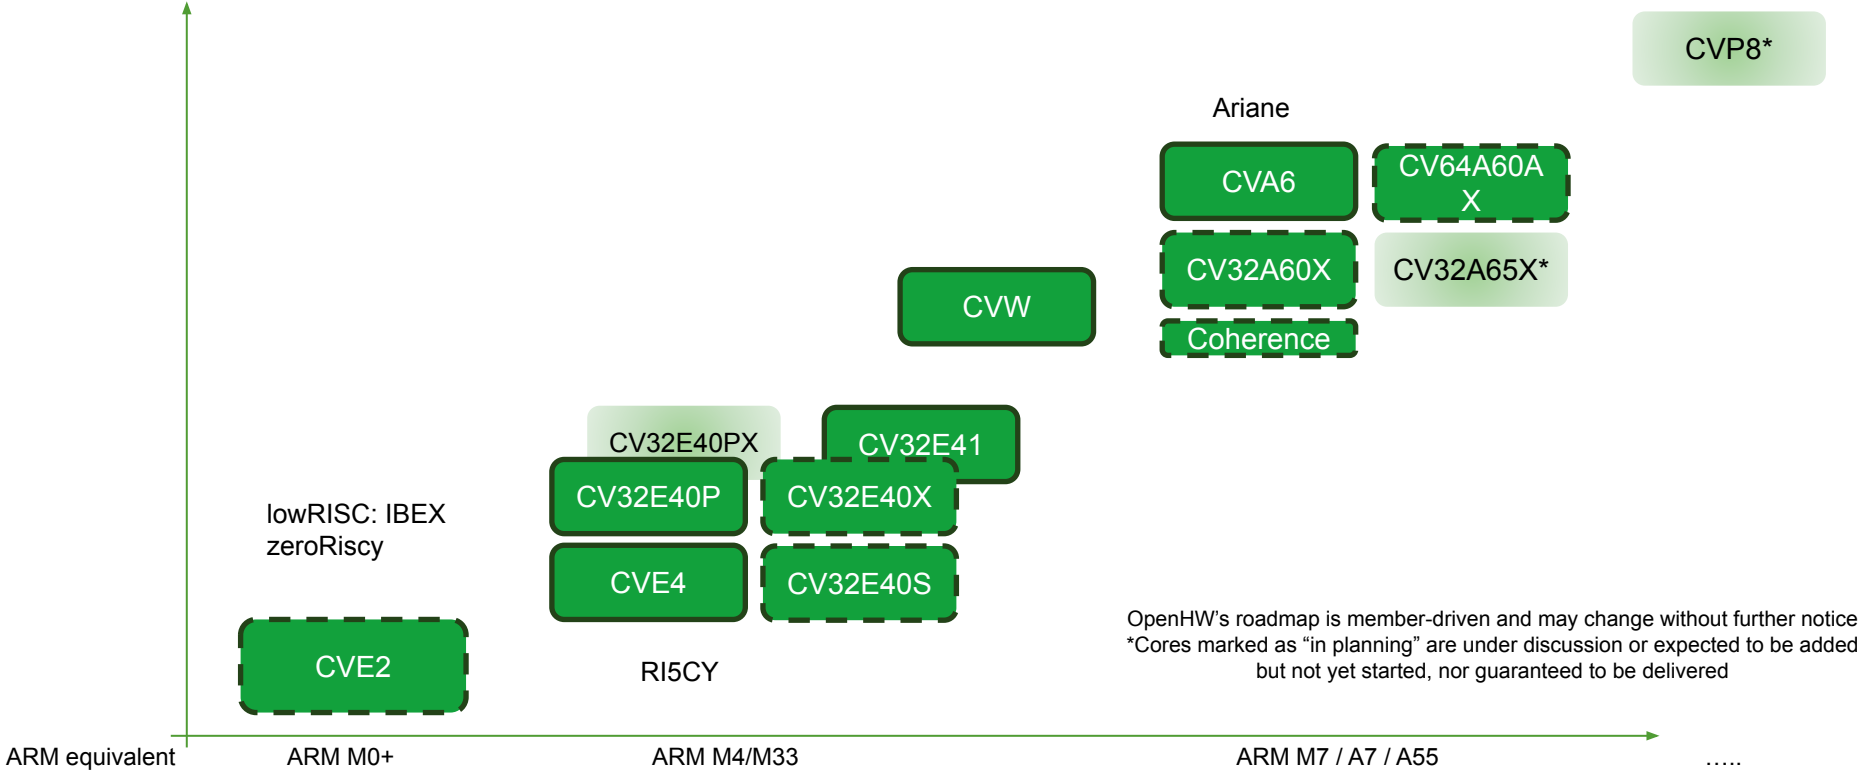

# OpenHW Foundation Deliverables

- **License**
  - Apache 2.0/Solderpad
- **Cores**
  - System Verilog

# OpenHW Foundation Deliverables

- **License**
  - Apache 2.0/Solderpad
- **Cores**
  - System Verilog
- **Test benches**
  - UVM, System Verilog, a little python and tcl
- **Tools**
  - Siemens Mentor Questa, Cadence, Synopsys, Imperas, …

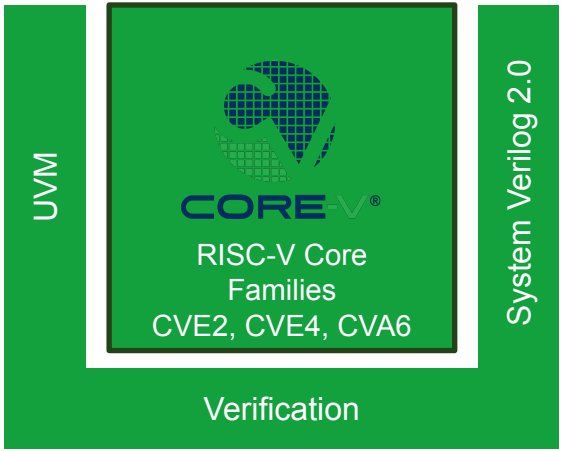

# OpenHW Foundation Deliverables

- **License**
  - Apache 2.0/Solderpad
- **Cores**
  - System Verilog
- **Test benches**
  - UVM, System Verilog, a little python and tcl
- **Tools**
  - Siemens Mentor Questa, Cadence, Synopsys, Imperas, …
- **Support software**
  - Compilers (LLVM, GCC)
  - RTOSes (FreeRTOS, Eclipse ThreadX)

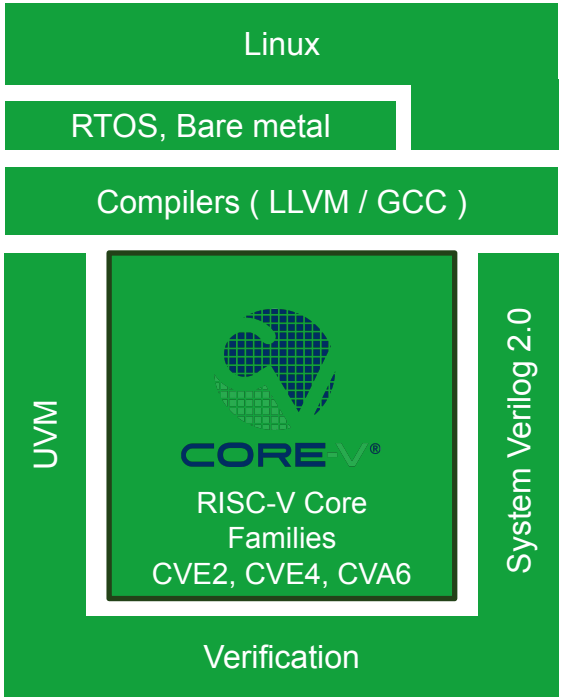

# OpenHW Foundation Deliverables

- **License**
  - Apache 2.0/Solderpad
- **Cores**
  - System Verilog
- **Test benches**
  - UVM, System Verilog, a little python and tcl
- **Tools**
  - Siemens Mentor Questa, Cadence, Synopsys, Imperas, …
- **Support software**
  - Compilers (LLVM, GCC)
  - RTOSes (FreeRTOS, Eclipse ThreadX)
- **Documentation**

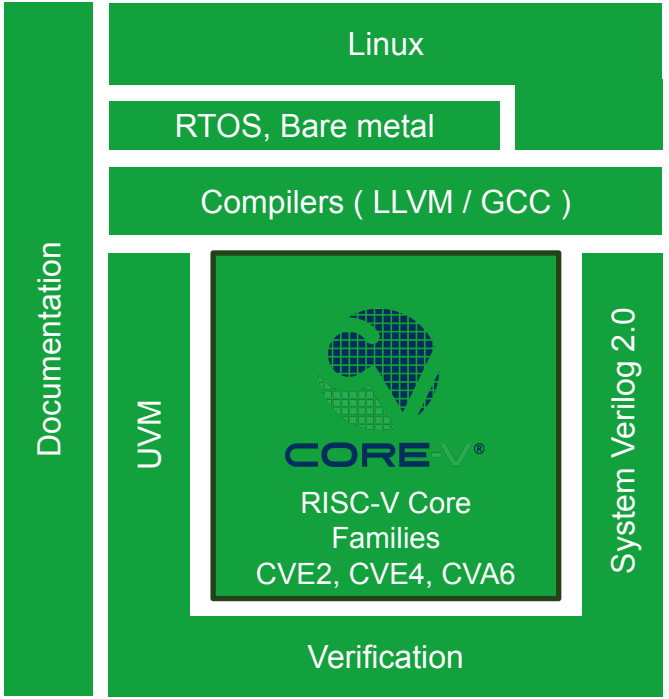

# OpenHW Foundation Deliverables

- **License**
  - Apache 2.0/Solderpad
- **Cores**
  - System Verilog
- **Test benches**
  - UVM, System Verilog, a little python and tcl
- **Tools**
  - Siemens Mentor Questa, Cadence, Synopsys, Imperas, …
- **Support software**
  - Compilers (LLVM, GCC)
  - RTOSes (FreeRTOS, Eclipse ThreadX)
- **Documentation**
- **FPGA Board Images / SOCs**
  - Digilent Nexys A7
  - Digilent Genesys 2

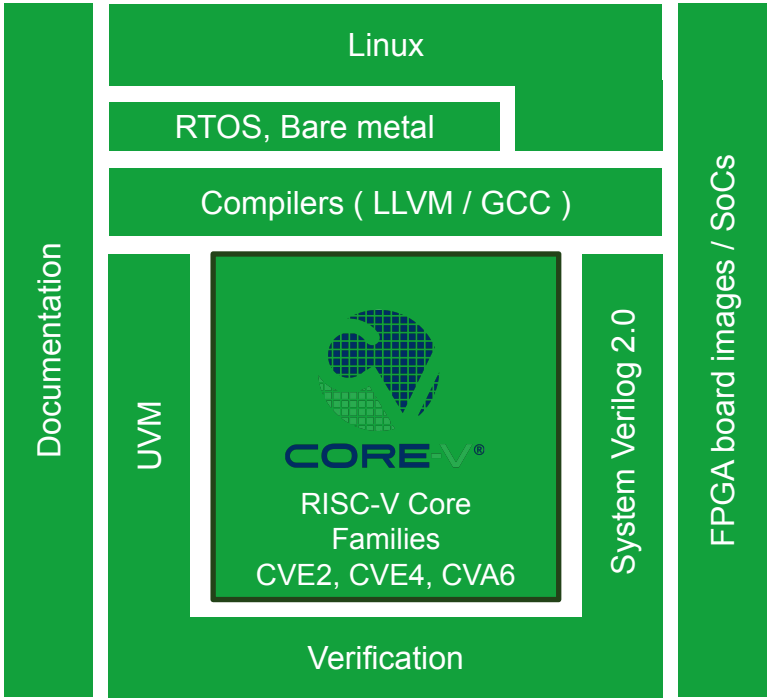

# Get Your Compiler Right

## Software Will Make or Break Your AI Chip

# OpenHW Software Task Group

- Chair: Paolo Savini, Embecosm
- Define, develop and support toolchains, operating system ports and firmware for the cores and IP developed within the OpenHW Foundation
- Active Projects:
  - GCC / LLVM
  - IDEs
  - FreeRTOS
  - CORE-V MCU SDK
- Emerging: Cooperation with the Eclipse ThreadX project

# Compilers

- The plan is for all compiler patches to go upstream
  - GCC 14.1 has 5 of the 8 CV32E40Pv2 ISA extensions
  - Clang/LLVM 18 has 4 of the 8 CV32E40Pv2 ISA extensions
  - Support for the rest available out of tree pending upstreaming
- Features:
  - Core awareness
  - Assembler and built-in support of all CORE-V ISA extensions
  - Automatic code generation for key CORE-V ISA extensions
    - including hardware loop support


# Operating Systems and Virtualisation

- **QEMU**
- **RTOS support**
  - FreeRTOS  => 10.3.0
  - Zephyr 2.4, 2.5, …
- **Upcoming: Eclipse ThreadX support**
  - Only OSS RTOS certified for safety-critical applications
  - Best-performing open source RTOS
- **Linux**
  - Linux Kernel 6.2 Support
  - uBoot and OpenSPI
  - Buildroot
  - Fedora and RedHat working on Linux Support

# CORE-V IDE

- CORE-V IDE is an open-source project under the SW TG at the OpenHW Group
- Based on Eclipse IDE, with native support for CORE-V development
- Includes the GCC Toolchain for CORE-V
- OpenOCD Debug Support
- “Ready-to-run” examples for Digilent FPGA boards
- Getting started guides

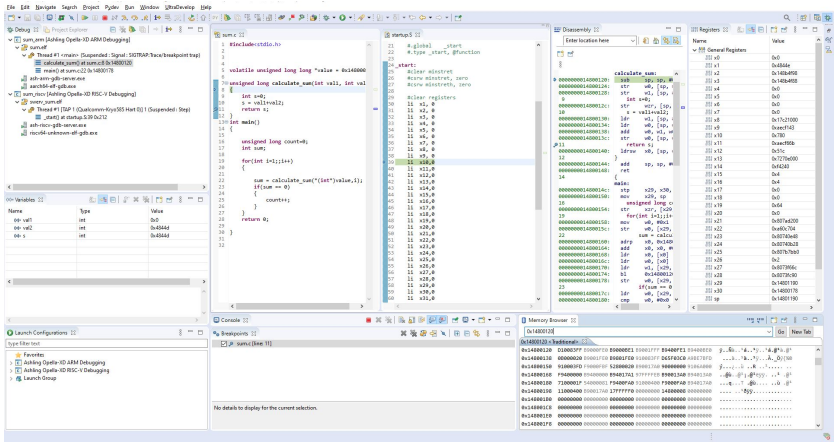

ASHLING

EMBECOSM®

eclipse®

CORE-V®

OPENHW FOUNDATION

# CORE-V MCU

- Real Time Operating System (e.g. FreeRTOS) capable ~400+MHz CV32E4 MCU
- Embedded FPGA fabric with hardware accelerators from QuickLogic
- Multiple low power peripheral interfaces (SPI, GPIO, I2C, HyperRAM, CAMIF, etc) for interfacing with sensors, displays, and connectivity modules
- Built in 22FDX with GlobalFoundries

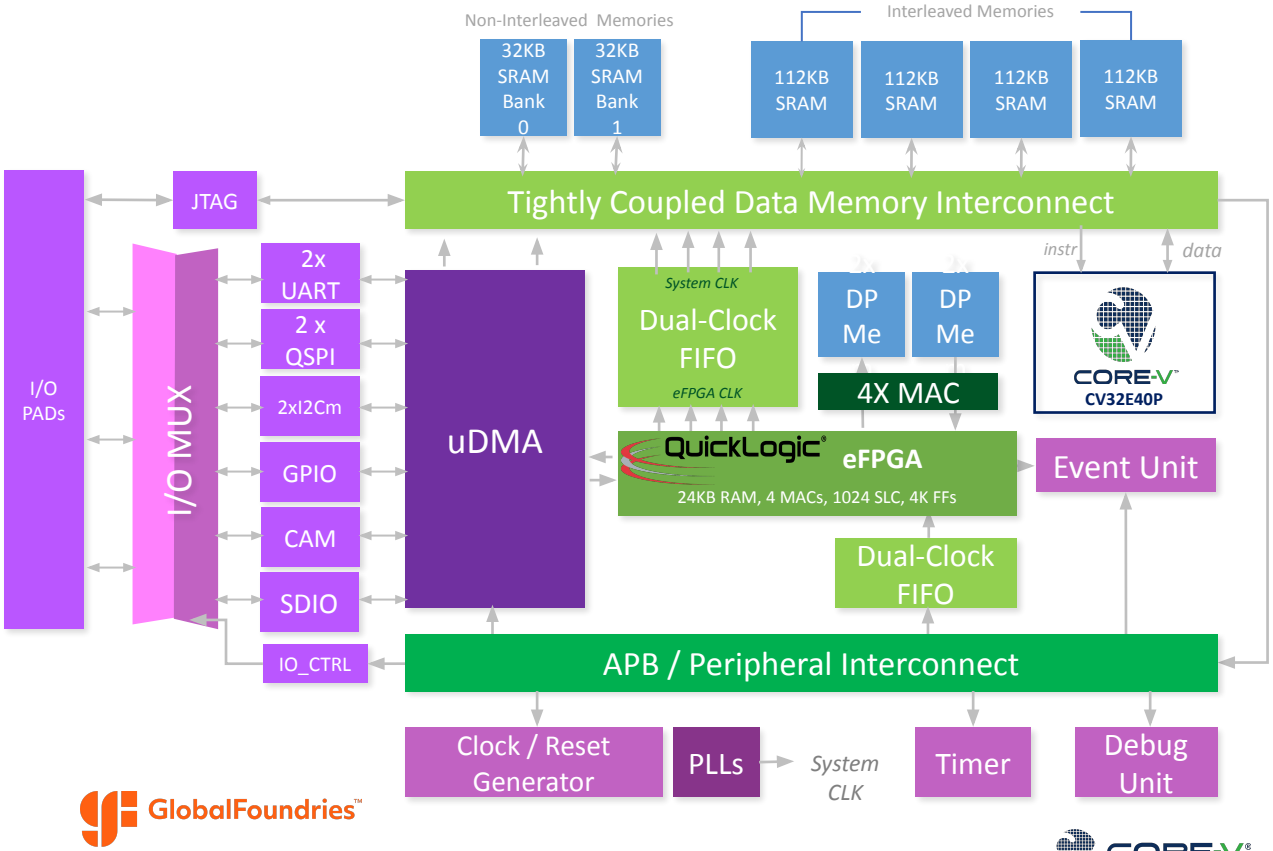

# CORE-V DevKit

- CORE-V MCU SoC
  - CV32E40P processor core
  - Quicklogic ArticPro eFPGA
  - Global Foundries 22FDX
- Ashling Opella-LD onboard JTAG debug module
- USB-C for terminal and onboard debug access
- JTAG connector for external debug access
- Espressif AWS IoT ExpressLink Module for AWS IoT cloud interconnect
- mikroBUS onboard socket, allowing access to a vast range of mikroBUS modules
- 40 pin expansion header
- I2C temperature sensor

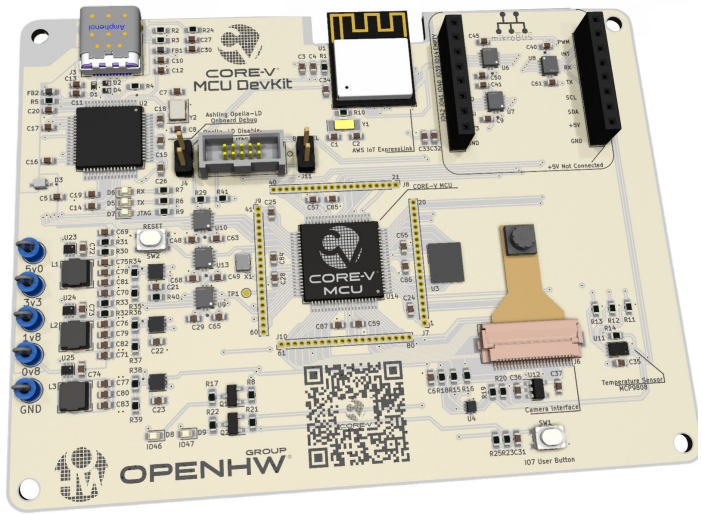

# Low Power AI Edge Inference Demonstrator

- Developed in 2021 at the University of Southampton
- How few instructions are needed to accelerate AI inference?
- CV32E40P extended with 8 RVV instructions
- Result: 5-7x speedup in TinyMLPerf benchmark
- Video: youtu.be/t0Fpy4TzLUE
- GitHub: github.com/AI-Vector-Accelerator

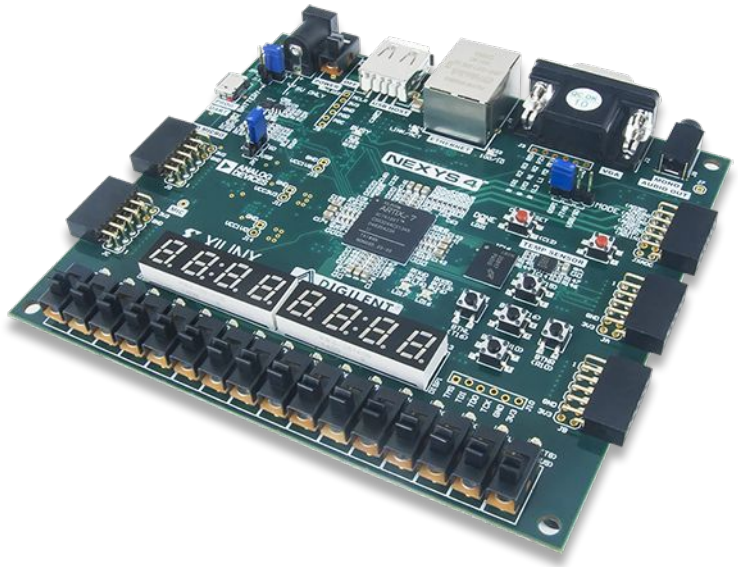

# Why Work With Us?

# Why Build Your Own AI Chip With our IP?

- RISC-V is already everywhere in AI
- OpenHW as a platform to host, maintain and verify high quality, industrial-grade cores
  - Open
  - Transparent
  - Meritocratic
- Academia and industry working together
- Fully open cores (RTL, Verification)
  - Challenge our RTL!
- The Eclipse Foundation is the largest open source organisation in Europe

# Call to Action

- Learn
  - Visit openhwfoundation.org
  - Watch OpenHW TV episodes on YouTube
- Try
  - Get the code on GitHub
- Engage

# Towards a Comprehensive Open Source Embedded RISC-V stack

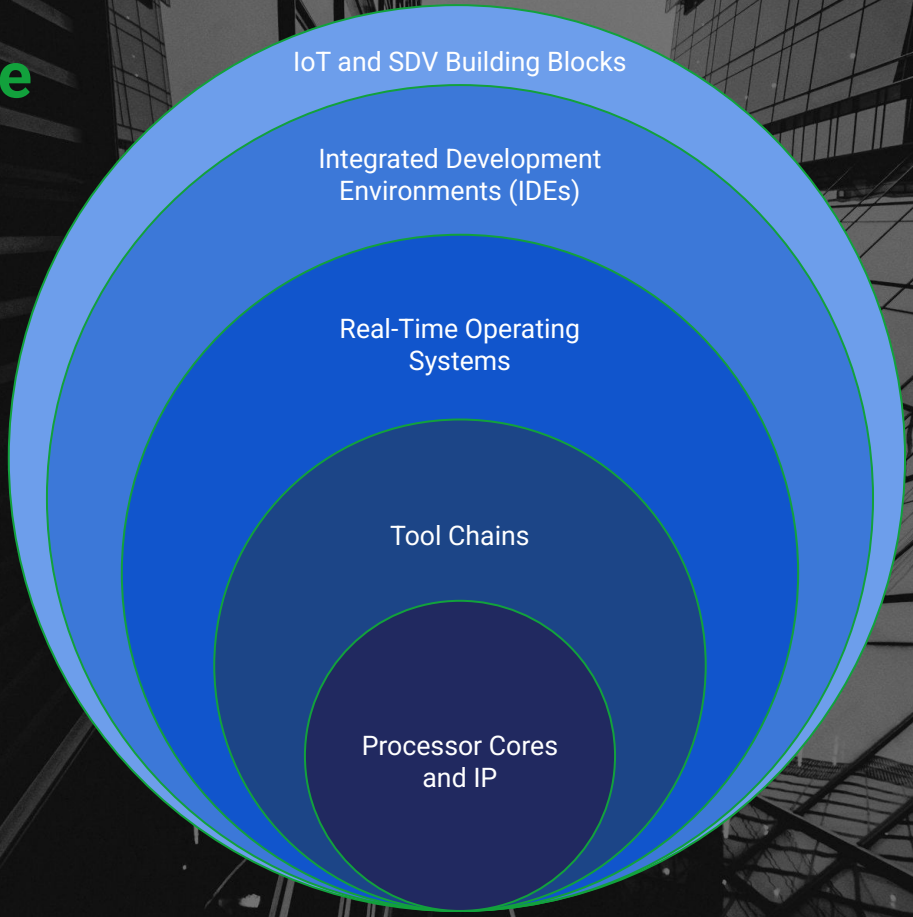


eclipse

eclipse iot

THEIA

ECLIPSE THREADX

FOUNDATION OPENHW PROVEN PROCESSOR IP

FOUNDATION
OPENHW™
PROVEN PROCESSOR IP

# Thank you!

## openhwfoundation.org