Scaling Open Compute: RISC-V, Chiplets, and the Future of AI and Robotics

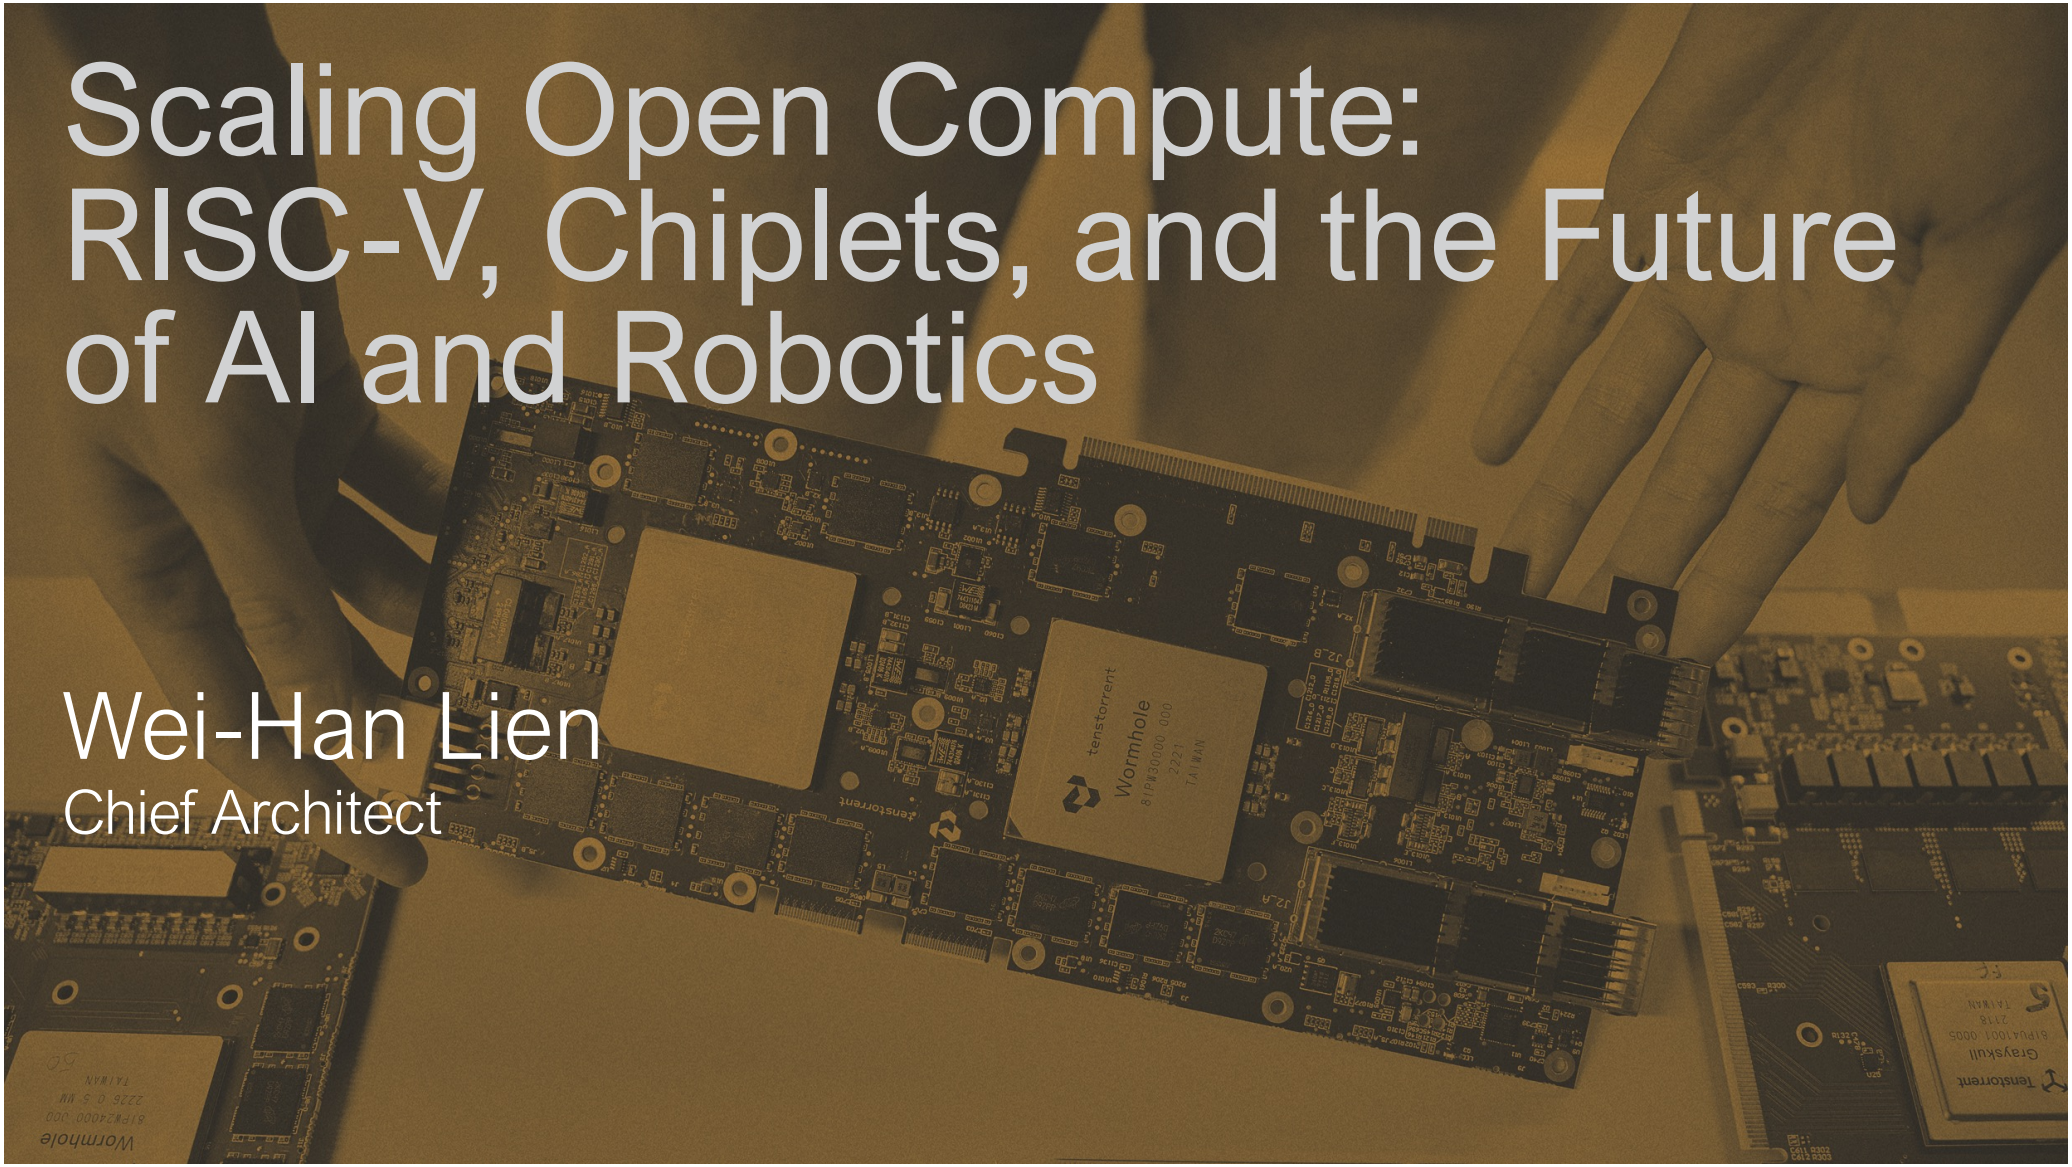

Wei-Han Lien
Chief Architect

Golden Age of Silicon Diversity

- Silicon Heterogeneity
- Architecture Variety
- Multi-vendor Silicon Landscape
- Design Compossibility

AI Personalization

Ubiquitous AI Compute

Open Source Hardware

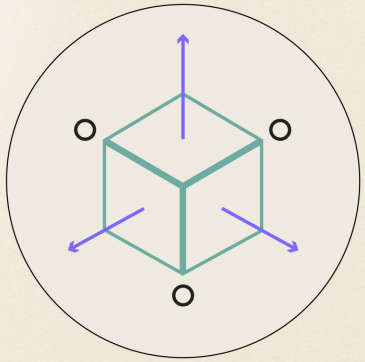

Scalable

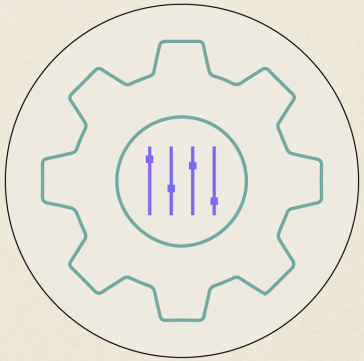

Extensible

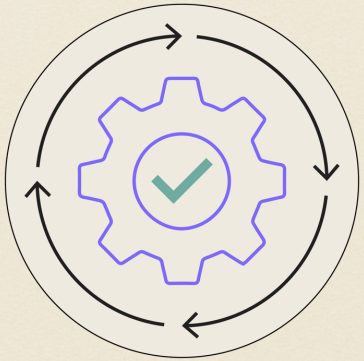

Efficient

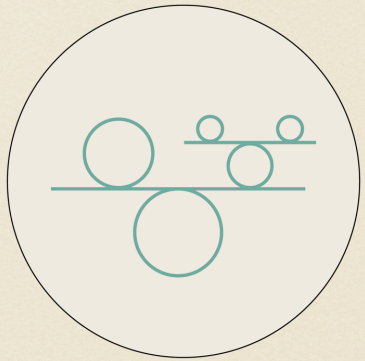

Stable

Foundational Compute IP's

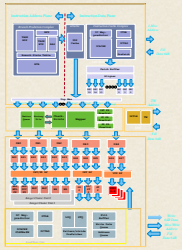

CPU

Tensix NEO

SOC Technology

PCIe

Networking

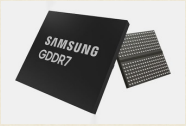

Memory

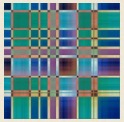

Fabric

Power/Thermal Management

RAS

Security

Debug

Chiplet Technology

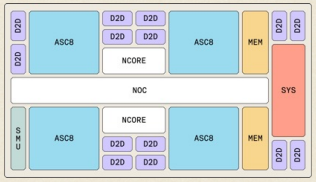

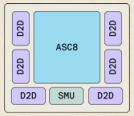

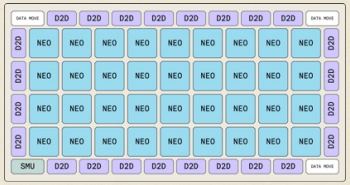

CPU, AI, IO, Memory chiplets

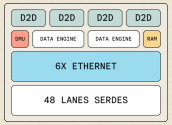

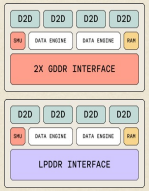

Applications

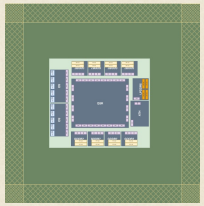

Data Center

Tenstorrent

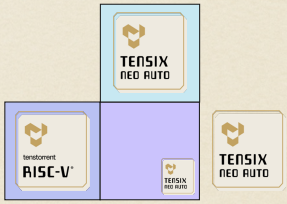

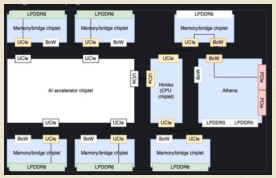

Yayui

AIDC

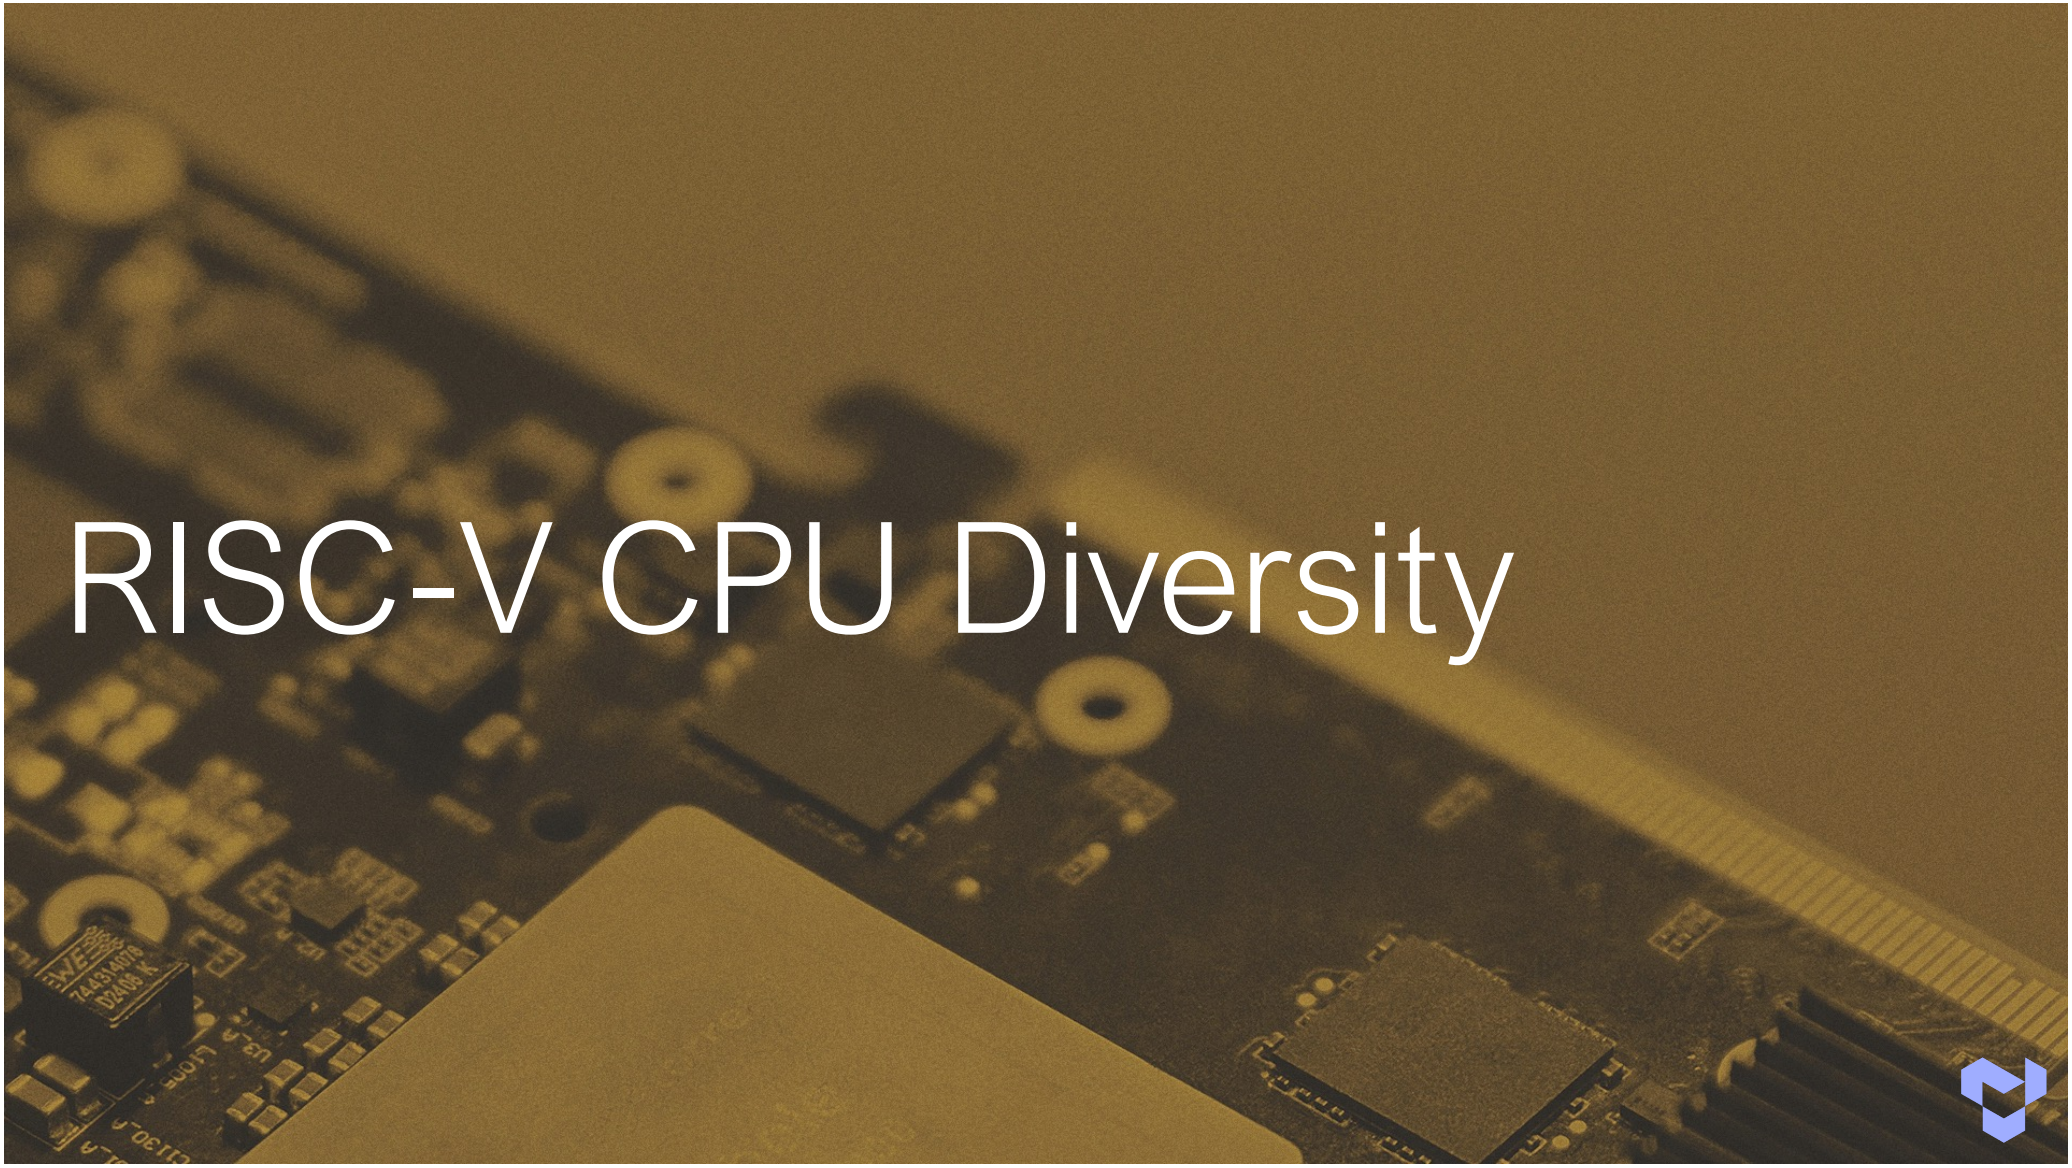
RISC-V CPU Diversity

RISC-V in the Cross-road

- RISC-V's 15-year Journey
  - An academic project in UC Berkeley
  - Mainstream adoption in embedded and industrial applications
- Sustain momentum
  - High-performance computing implementation
  - Software development
- Key Requirements
  - Advanced O-o-O CPU design with high-bandwidth memory
  - Optimized compilers, firmware, and system software
- Without high-end investment, RISC-V confined to low-cost, niche markets

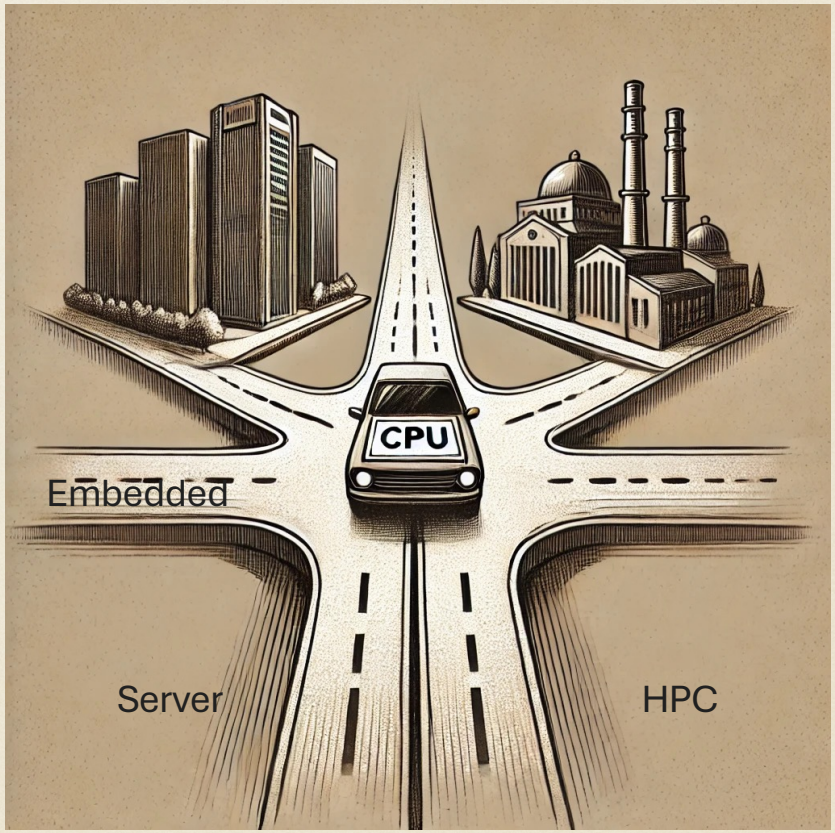

RiscV CPU - Ascalon

- Disruptive high-performance RISC-V processor for AI and server
- 18 SPECINT/Ghz

RVA-23

- Advanced branch predictions
- 8-wide decode
- 3 LD/ST with large load/store queues
- 6 ALU/2 BR
- 2 256-bit vector units
- 2 FPU units

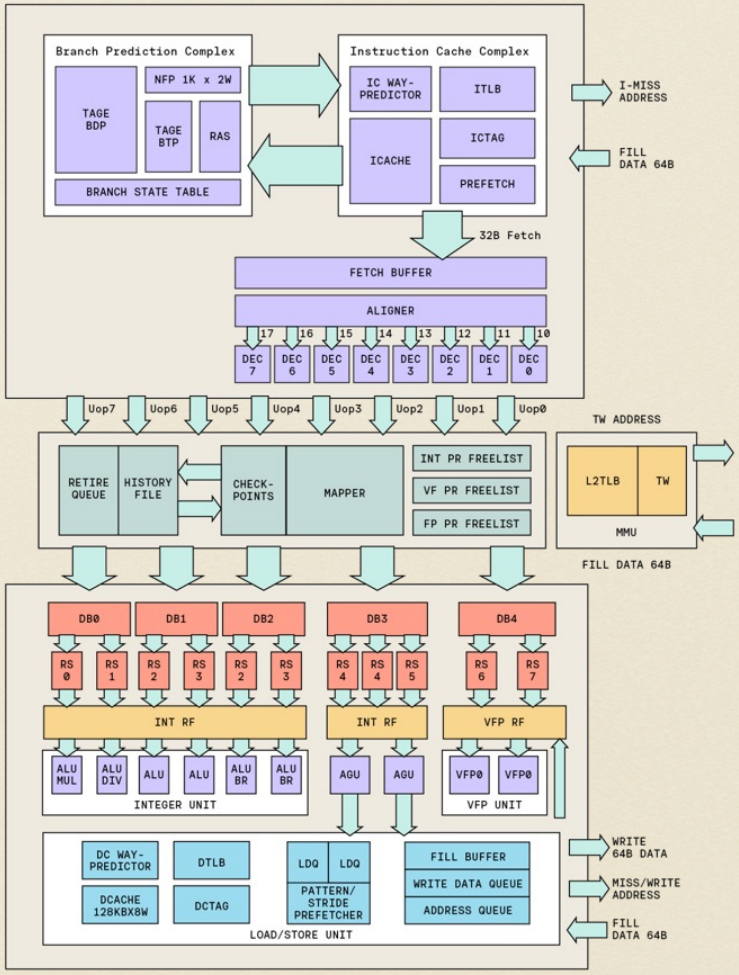

Athena Chiplet

- 8-core Ascalon CPU cluster
  - 36 SPECINT2K17 Rate
- PCIe Gen6x16
- LPDDR5X 8533

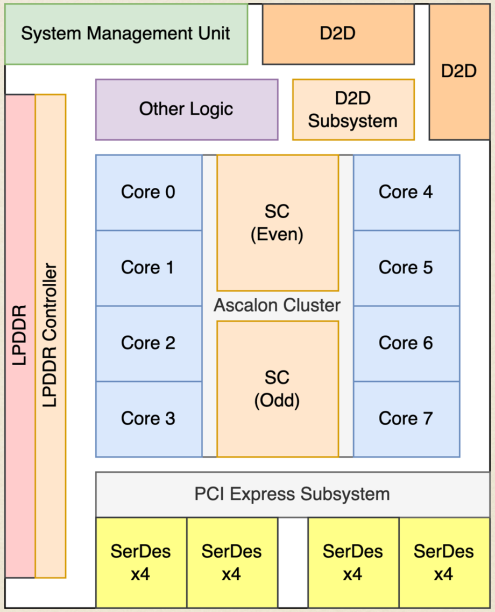

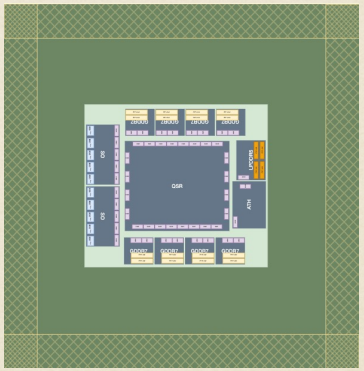

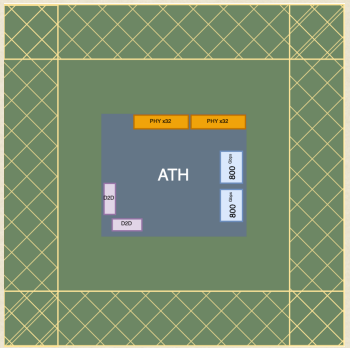

Ascalon Cluster

- Samsung SF4
- Performance correlated 18 SPEC2k6INT/GHz

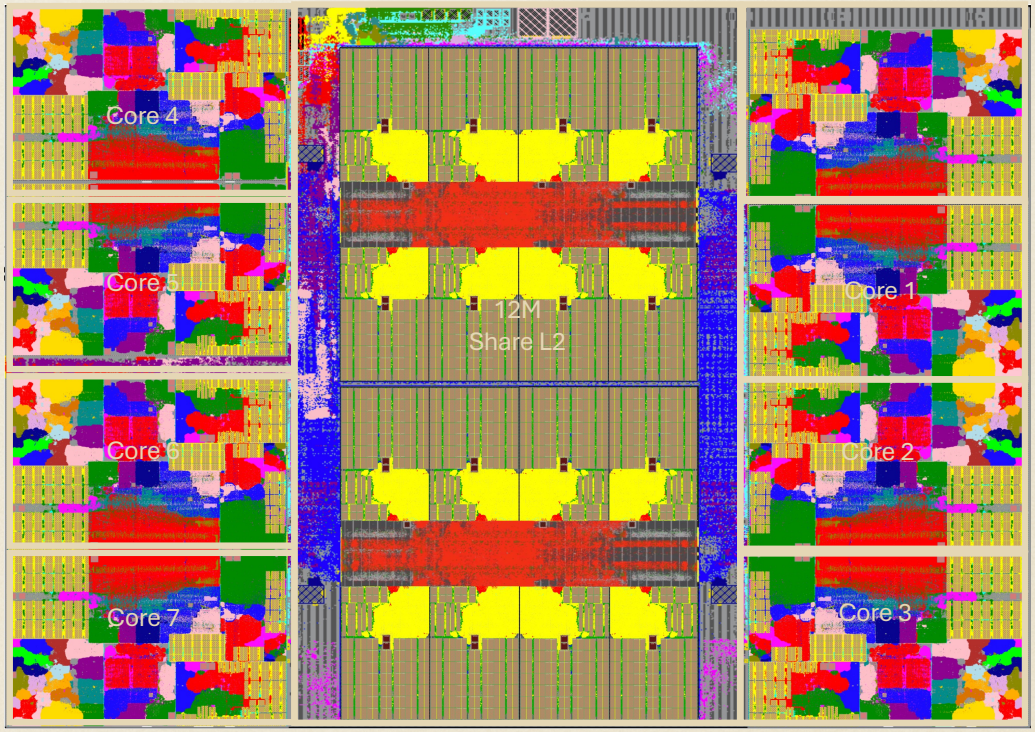

Ascalon Core

8 Ascalon Cluster
12M Shared L2

RISC-V CPU Roadmap

- Ascalon yearly ~10% upgrade
  - Ascalon+ 21 SPEC2k6INT/GHz*
  - Ascalon++ 23 SPEC2K6INT/GHz*
- **Callandor 100%**
  - **35 SPECINT2K6/GHz***

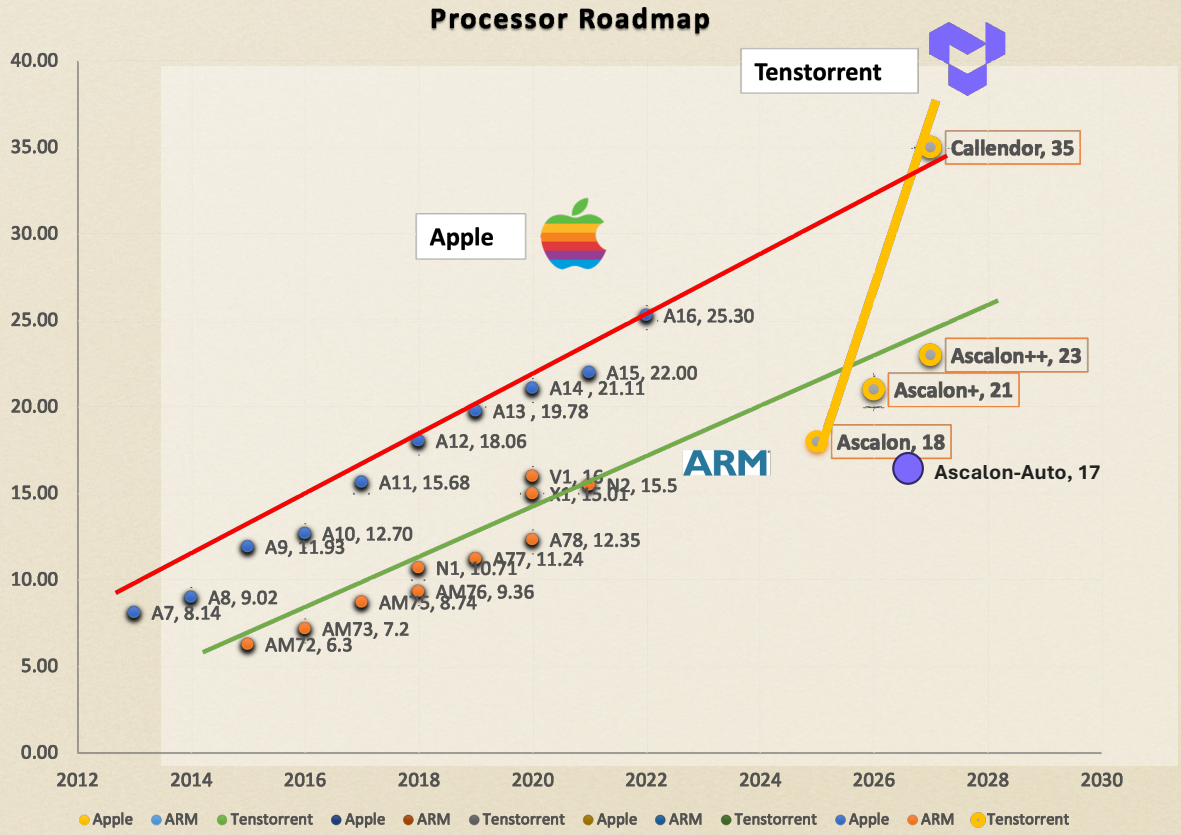

*Future roadmap subject to changes

Callandor O-o-O Superscalar Processor

- At the forefront of CPU performance innovation in Q1/2027
- 35 SPEC2K6INT/GHz
- 3.5 SPEC2K17INT/GHz
- RVA 23+
  - Front-end
    - Advanced branch predictors: Two taken branches
    - Decouple Front-end
    - uOp cache
  - Mid-core
    - 16-wide decode
    - Register file sharing
    - 1K ROB
  - Execution uits
    - 6 wide LD/ST with large load/store queues
    - 8 ALU/4 BR
    - 4 256-bit vector units
    - 4 FPU units
    - Matrix engine

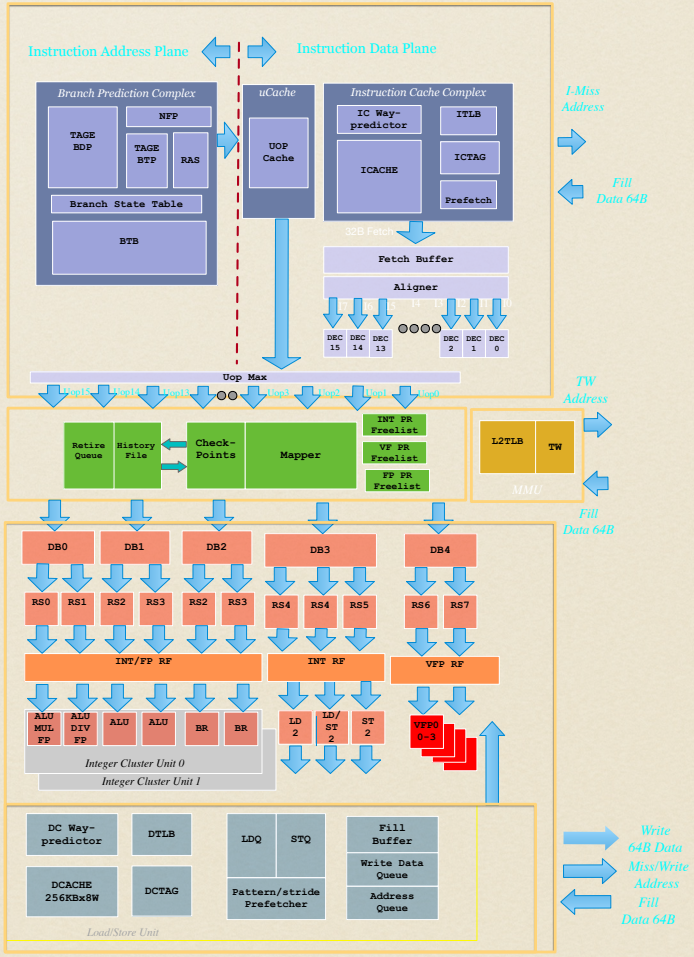

Ascalon-Auto IP

- ISO26262 Functional Safety Features
  - Dual Core Lock Stepped with Time Disparity
  - Coherent & Non-Coherent Bus Protection with CRC or ECC
  - EEC Protection for L1 & L2 Cache
  - AXI Interface Parity Protection
  - RAS & Fault Controller
  - Safety Bus for Debug
  - Software Test Library Support

- SPECInt2K17@2.25GHZ up to 30 per cluster
- Dhrystone MIPS : 12.38 DMIPS Rate/ MHZ

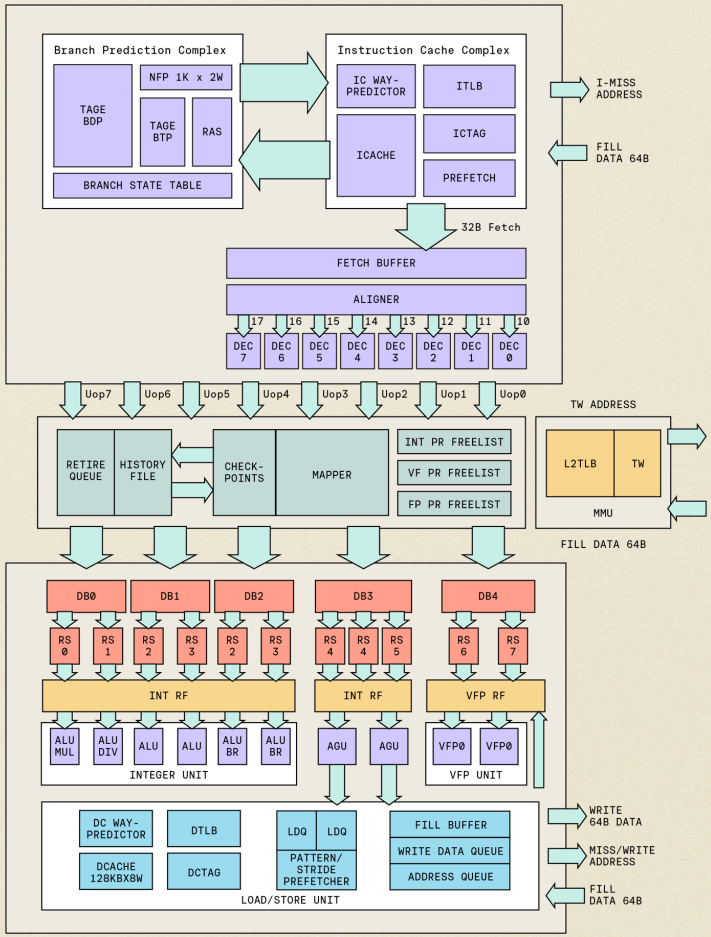


Chiplet Diversity

Chiplet

- Silicon SiP diversity
  - Design reuse
  - Low-cost development
  - Composability
  - Heterogeneity
- Diversity through composing chiplets from different organizational entities

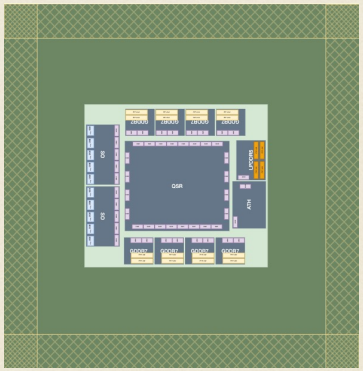

Open Chiplet Architecture (OCA)

- Open Chiplet Architecture (OCA) defines chiplet-interoperability in 5 layers:
  - Physical
  - Transport
  - Protocol
  - System
  - Software
- **Plug-Talk-Play**

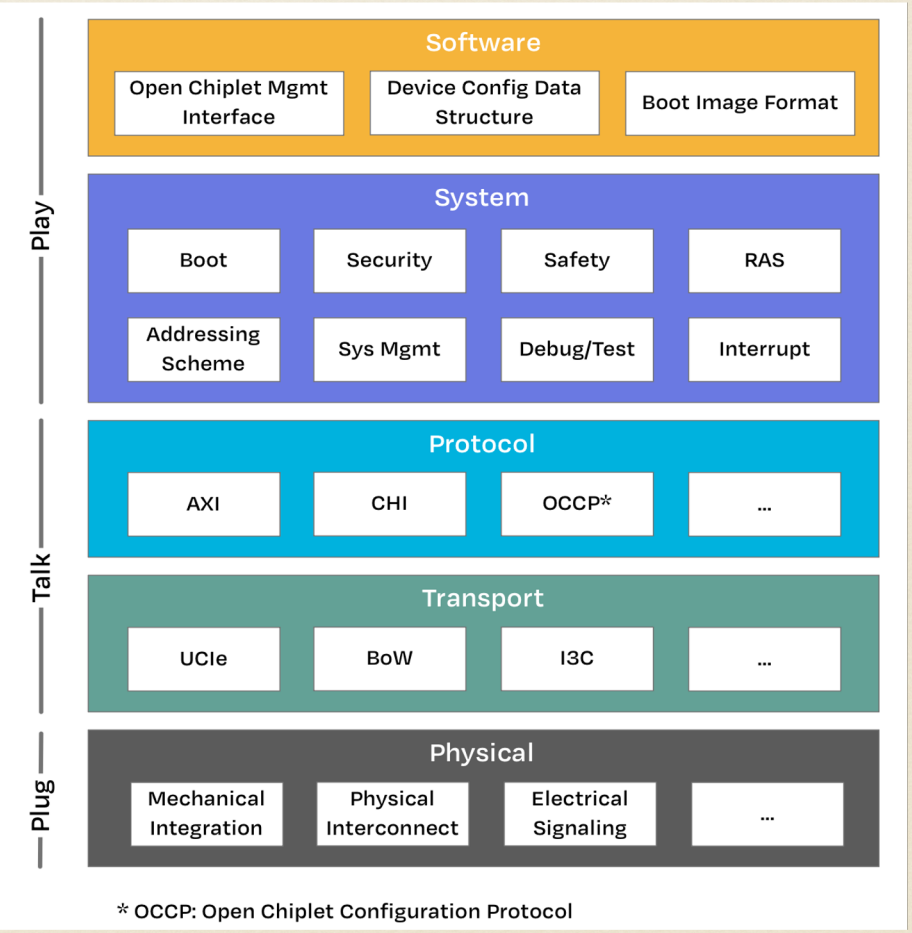

Open Chiplet Architecture

- OCA takes care of chiplet compatbility
- Each company focus on its core-competency design

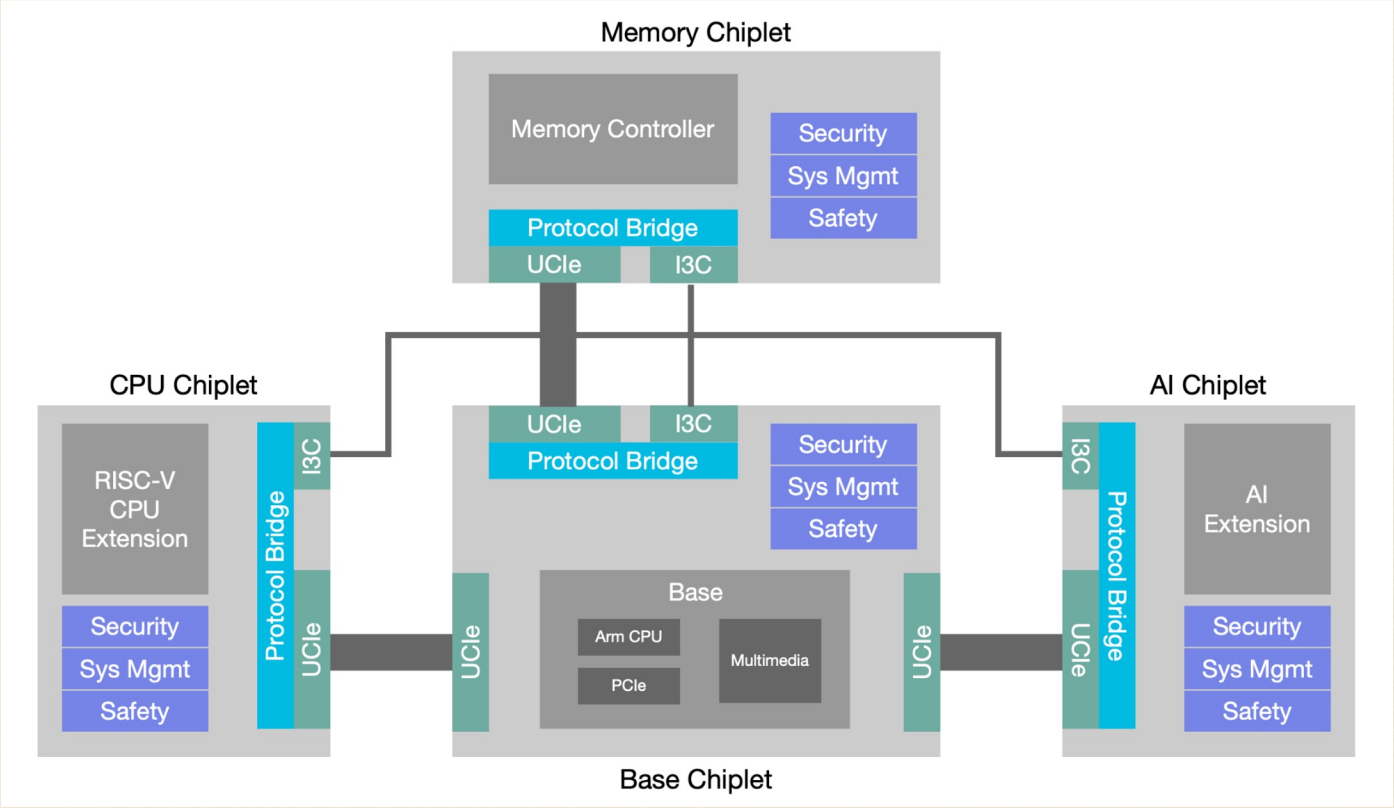

Robotic – Auto

Level 2+/3 → Level 3/4 → Level 4/5 → Data Center

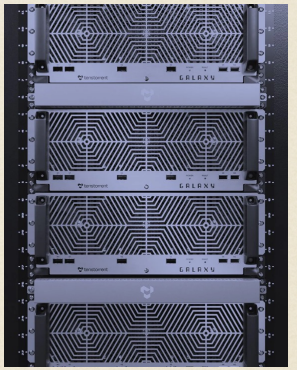

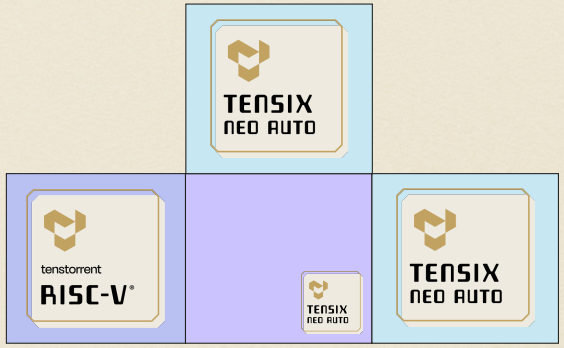

Dual SoCs

Dual SoCs + AI Co Processors

Central Compute Unit - Heterogenous

Neuro-symbolic AI Chiplet

- Heterogenous Compute with high-performance Ascalon
  - Neuro-symbolic AI
  - Optimize CPU/GPU data transfer latency

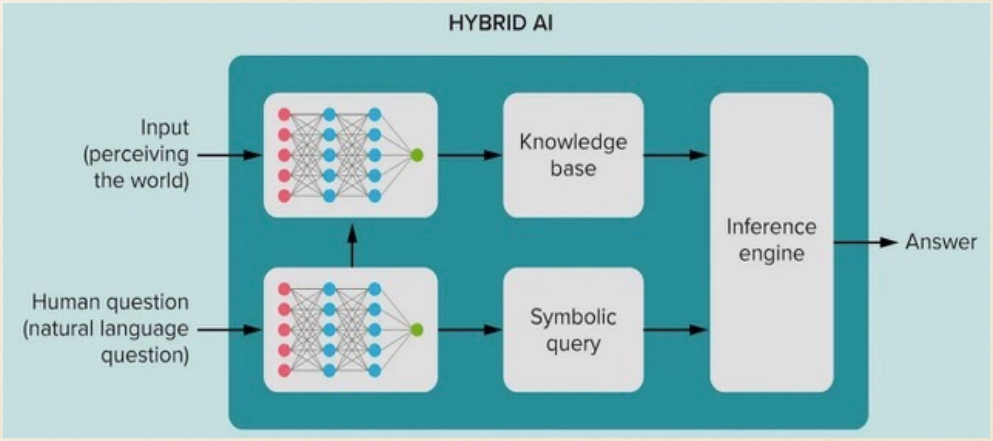

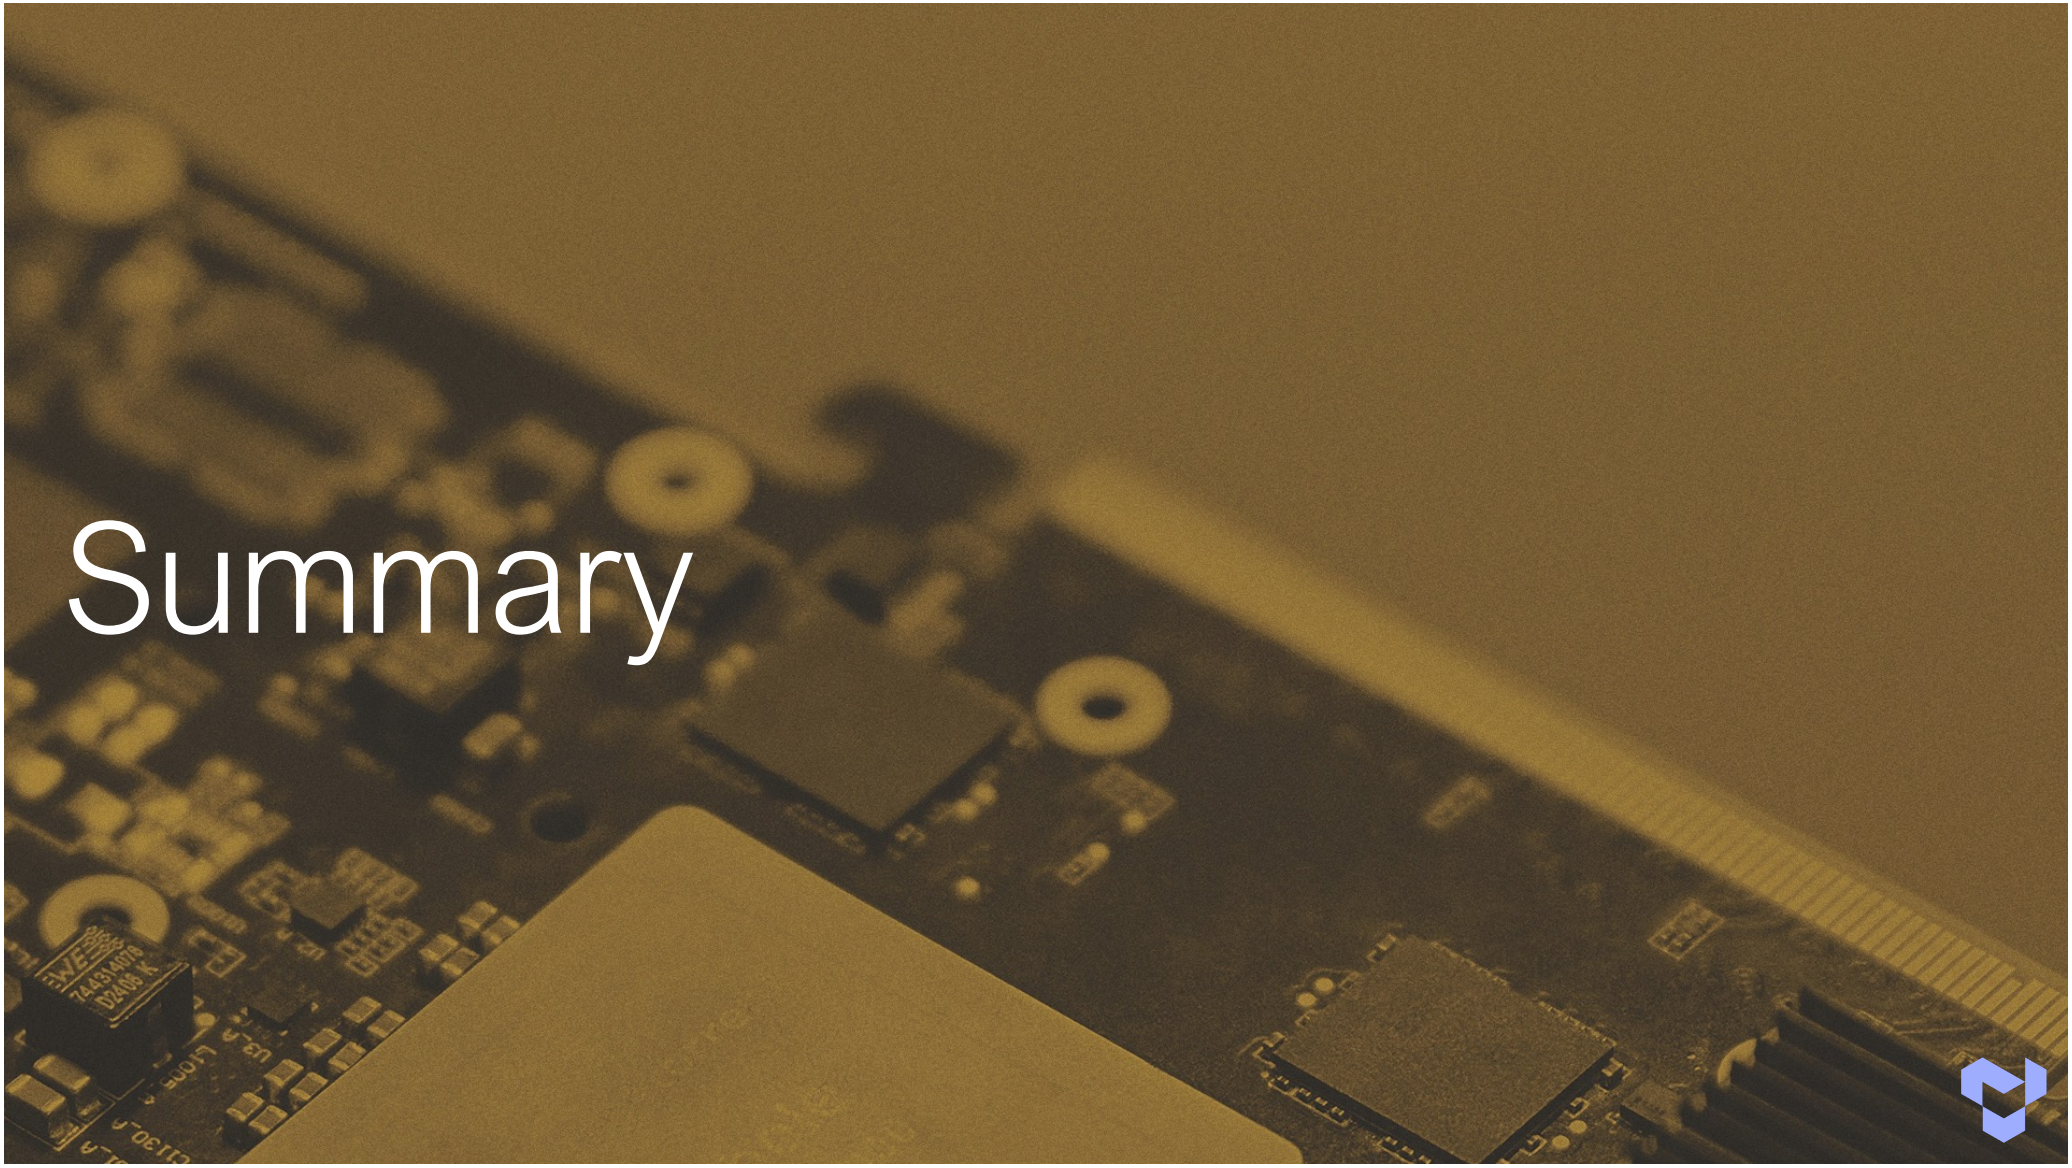
Summary

Summary

- AI personalization
- Ubiquitous AI computing
- Golden Age of Silicon Diversity
  - High performance CPU roadmap
  - Open Chiplet Architecture (OCA)