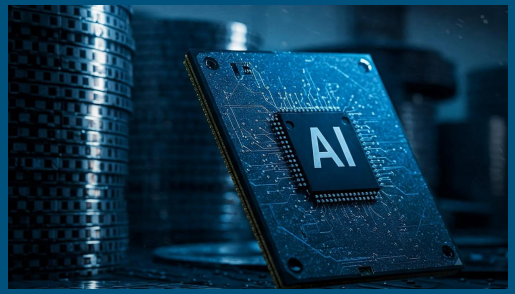

RISC-V Powering Bespoke Silicon and European Semiconductor

Pablo Valerio

Editor - Supply Chain

RISC-V SUMMIT EUROPE 2025

EE Times

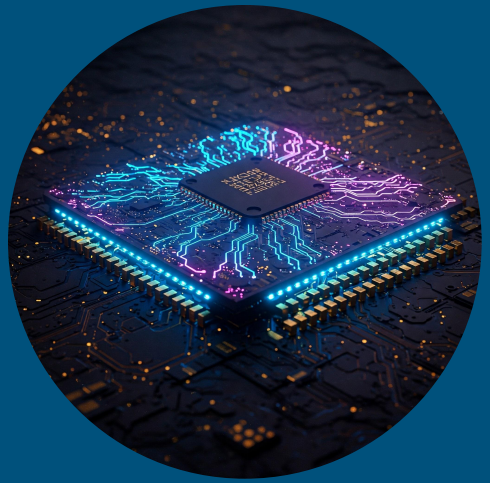

Bespoke Chips: Navigating Disruption and Supply Chain Challenges in the era of AI and HPC

- The semiconductor industry is experiencing a major shift towards custom silicon driven by the immense computational demands of AI and HPC.
- Companies are moving away from off-the-shelf processors to custom chips for optimized performance, power efficiency, and lower TCO, especially for cloud, AI/ML, and specialized devices.
- Hyperscalers and AI firms are investing heavily in bespoke silicon to gain a competitive edge through superior performance and cost reduction.

RISC-V SUMMIT EUROPE 2025

EE Times

Architectural Shift

- Bespoke silicon adoption has driven architectural diversification.
- Arm architecture now underpins custom CPU development for major companies like AWS, Google, Microsoft, Alibaba, and Apple, challenging x86 in data centers due to its power efficiency and ecosystem.
- The open-standard RISC-V ISA is a significant disruptive force, offering a compelling alternative for specialized processors in Automotive, AI, HPC, etc.
- RISC-V adoption has geopolitical implications, particularly in China, while Western companies adopt it for tactical optimizations.
- RISC-V increasingly targets AI and HPC, with major players using it in TPUs and accelerators.
- It is projected to generate substantial SoC revenue in the compute segment by 2030.

Manufacturing & Supply Chain Challenges

- A few players dominate the global foundry market, with TSMC being the primary partner for advanced logic (under 7nm), holding over 90% of high-end units and 67% of global semiconductor revenue.
- Advanced manufacturing is geographically concentrated in Taiwan (TSMC) and South Korea (Samsung), creating single points of failure for sub-10nm chips.
- Geopolitical tensions, like the US-China tech rivalry, disrupt trade and push China towards self-sufficiency.
- Innovation costs for leading-edge nodes (sub-7nm, GAAFET) are rising, limiting participation.
- Supply chains are also vulnerable to raw material and chemical disruptions. Bespoke advanced node silicon has escalating NRE costs and complexity.

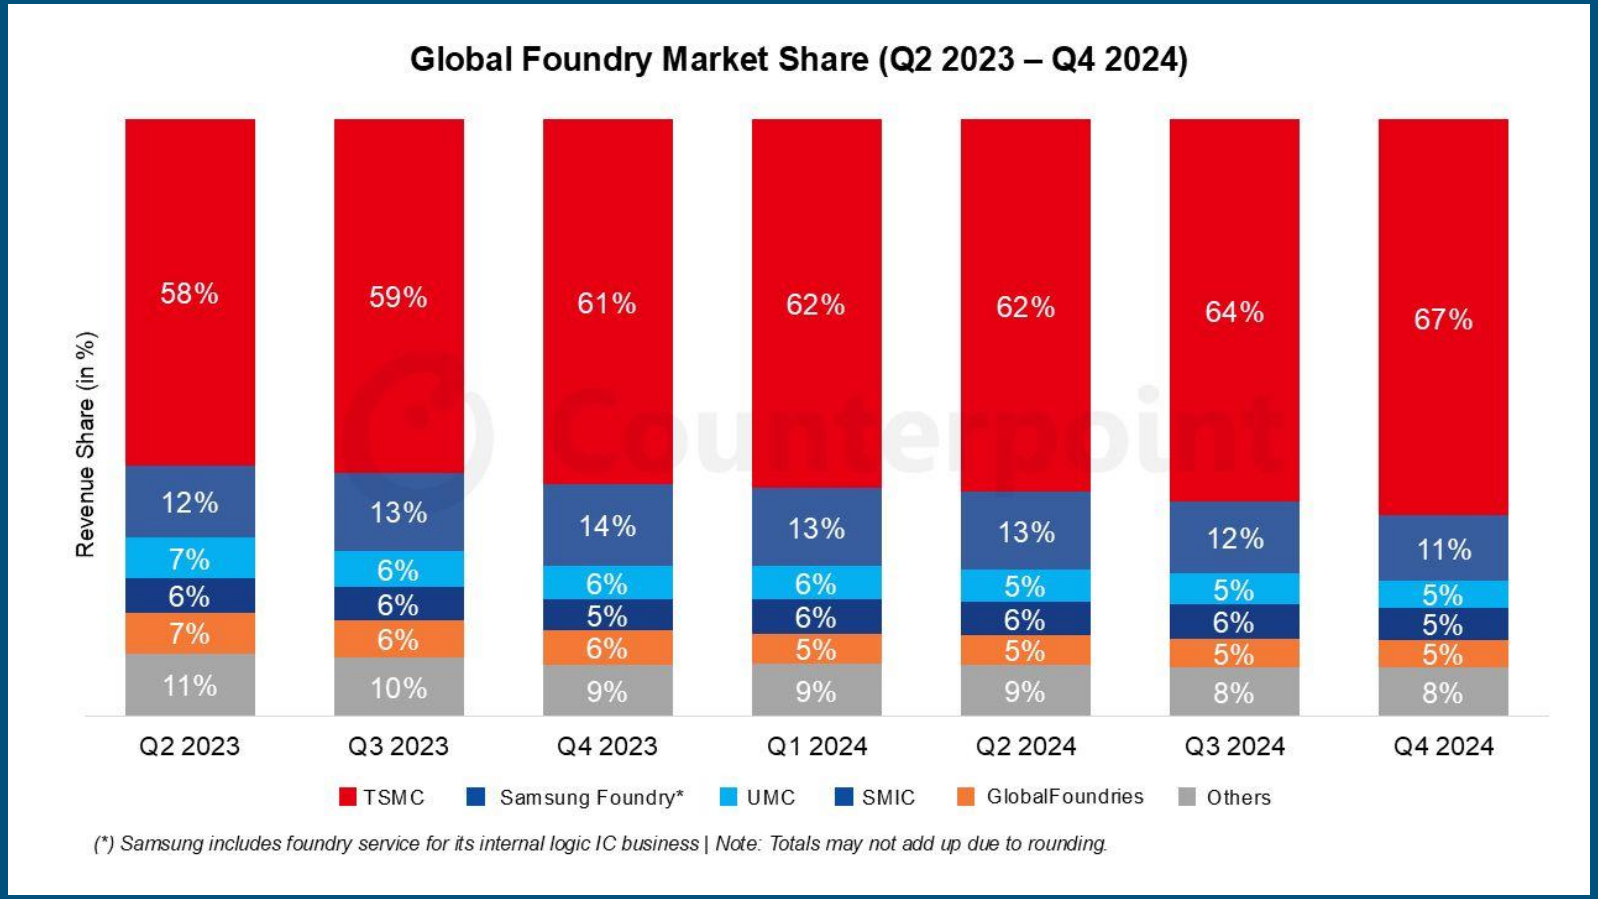

Credit: Counterpoint

RISC-V: The Open Alternative

- RISC-V offers openness, customization, and cost-effectiveness as an open-standard ISA with no licensing fees, lowering barriers to custom processor development.
- Its modularity allows for tailored instructions for domain-specific AI and HPC applications. RISC-V provides greater design and supply chain control, avoiding vendor lock-in. Its open nature fosters scrutiny for potentially more secure designs.
- Geopolitically, it offers nations a path to semiconductor sovereignty.
- Global RISC-V adoption could lead to a more resilient global supply chain.

Europe's Strategic Embrace of RISC-V

- Europe relies on non-EU sources for advanced chip manufacturing, making RISC-V a strategic priority for technological sovereignty and a resilient semiconductor ecosystem.
- The European Chips Act aims to invest over €43 billion by 2030 to strengthen the EU's semiconductor value chain, including design and manufacturing.
- The IPCEI ME/CT directly supports the development of new RISC-V processor cores.
- Several publicly founder projects, such as DARE, intends to create a complete European HPC compute stack based on RISC-V.
-

Europe's Strategic Embrace of RISC-V (2)

- Infineon, STMicroelectronics, and NXP are launching or leveraging RISC-V in their products, mostly for the automotive industry.
- European IP providers like Codasip and Semidynamics are developing RISC-V cores.
- Foundry support is growing with GlobalFoundries Dresden and Intel's future Magdeburg fab planning to support RISC-V.
- However, Europe still lacks domestic high-volume, leading-edge (sub-7nm) logic foundry capacity crucial for advanced RISC-V processors.
- Achieving full technological sovereignty through RISC-V necessitates access to advanced manufacturing.

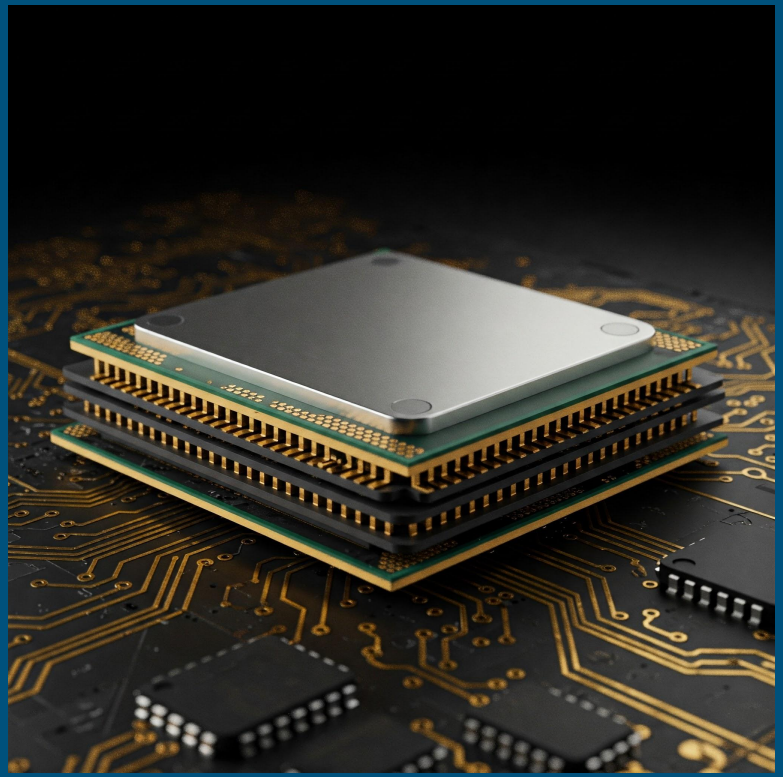

Grateful Thanks

Pablo Valerio

LinkedIn: pvalerio

RISC-V SUMMIT EUROPE 2025

EE Times

AWS Silicon Strategy

- Amazon Web Services has a comprehensive custom silicon strategy.
- AWS optimizes for performance, security, and energy efficiency.
- Nitro System offloads tasks onto dedicated hardware.
- Graviton CPUs power general-purpose cloud workloads.
- Inferentia and Trainium are used for AI acceleration.

Architecture:

- CPU: Arm (Graviton series).
- AI/Other: Custom ASICs/SoCs (Nitro, Inferentia, Trainium). Proprietary architectures for AI accelerators.
- RISC-V: Supported in FreeRTOS (embedded/IoT) and FPGA platforms. No public evidence in core chips.

Foundry Partner(s):

- TSMC: Primary partner for Graviton (16nm, 7nm, 5nm, 4nm). Likely also for Inferentia/Trainium.
- Intel Foundry Services (IFS): Strategic partnership announced in 2024 for future nodes (Intel 18A, Intel 3, potentially 18AP, 14A).

Latest Known Process Nodes:

- TSMC: 4nm (Graviton4), 5nm (Graviton3).
- Intel (IFS): 18A, Intel 3, future 18AP, 14A (upcoming products).

A Strategic Imperative

- Custom silicon is essential for differentiation, cost control, and supply chain resilience in cloud, AI, and consumer tech.
- Arm dominates custom CPU designs for hyperscalers/mobile.
- AI accelerators show diverse, often proprietary, architectures.
- RISC-V is emerging, used tactically by Western firms and strategically by Chinese firms for self-sufficiency.
- TSMC remains critical for leading-edge nodes, but the foundry landscape is diversifying due to geopolitics and new players (IFS, SMIC for China).
- Chiplet architectures and advanced packaging are future trends.
- This shift impacts traditional vendors and is a central element of global tech competition.

Meta Competitive Edge

- Meta is developing custom AI accelerators named MTIA.
- MTIA uses RISC-V cores for processing elements.
- Merchant CPUs handle general processing tasks for Meta.
- TSMC is Meta's manufacturing foundry partner for MTIA.
- Meta optimizes custom AI accelerators for internal workloads.

Architecture:

- AI: Custom (MTIA). Processing Elements (PEs) within MTIA utilize RISC-V cores. MTIA v2 likely continues to use RISC-V PEs.
- CPU: Uses merchant CPUs (not custom).

Foundry Partner(s):

- TSMC: Partner for manufacturing MTIA chips (v1 and v2). Utilizes TSMC's CoWoS advanced packaging for MTIA v2.

Latest Known Process Nodes:

- TSMC: 5nm (MTIA v1 and v2).

Microsoft Azure Silicon

- Microsoft is accelerating custom silicon for Azure cloud.
- Azure Cobalt is an Arm-based CPU with 128 cores.
- Azure Maia is a custom AI accelerator for training/inference.
- Azure Boost offloads networking and storage processing.
- Azure Integrated HSM: Custom security processor.

Architecture:

- CPU: Arm (Azure Cobalt).
- AI: Custom ASIC (Azure Maia). Proprietary architecture.
- Other: Custom silicon/FPGA/DPU (Azure Boost), Custom security processor (Integrated HSM).
- RISC-V: No public indication of use in Azure Cobalt, Maia, or Boost silicon.

Foundry Partner(s):

- TSMC: Partner for manufacturing Azure Cobalt CPU and Azure Maia 100 AI accelerator.

Latest Known Process Nodes:

- TSMC: 5nm (Azure Cobalt, Azure Maia 100).

Google Silicon Mix

Key Silicon Products:

- Tensor Processing Units (TPU): Custom ASICs for accelerating neural network computations (matrix multiplication). Multiple generations for inference (v1, v7/Ironwood) and training (v2-v5, v6e/Trillium, v5p).
- Axion (CPU): Custom Arm-based CPU for general-purpose computing in GCP (C4A instances). Based on Arm Neoverse V2. Built on Titanium system using custom microcontrollers.
- Tensor G-Series (SoC): Custom SoCs for Pixel devices. Integrate Arm CPU cores, GPU, custom Edge TPU (AI), ISP, security cores. Tensor G4 (Pixel 9) uses Samsung 4nm. Tensor G5 (Pixel 10, upcoming) planned on TSMC 3nm. Tensor G6 (future) targeting TSMC 2nm.

Architecture:

- AI: Custom ASIC (TPU - systolic array). Custom (Tensor G Edge TPU).
- CPU/SoC: Arm (Axion, Tensor G). Google-designed Arm-compatible cores/systems for Tensor G.
- RISC-V: Used for flexible compute hosts alongside TPUs (SiFive Intelligence X280 cores), and supported in Android OS starting with Android 15. No public confirmation for main Axion or Tensor G CPU cores.

Foundry Partner(s):

- TSMC: Likely foundry for latest TPUs and Axion. Partner for upcoming Tensor G5 (3nm) and G6 (2nm).
- Samsung Foundry: Manufactured early TPUs and Tensor G-series chips (G1-G4).

Latest Known Process Nodes:

- TSMC: 3nm (Tensor G5 planned), 2nm (Tensor G6 planned), 7nm (TPUv4). Axion rumored on 3nm N3E.
- Samsung: 4nm (Tensor G4).

Huawei's AI Silicon Push

- Driven by sanctions, Huawei pursues self-sufficiency.
- Kirin SoCs use custom Arm-based cores.
- Ascend AI processors use Da Vinci architecture.
- SMIC is the primary foundry partner for Huawei.
- Huawei targets 5nm processing node with SMIC.

Key Silicon Products:

- Kirin (SoC): Arm-based SoCs for consumer devices (smartphones, tablets). Feature custom CPU/GPU cores, NPUs. Kirin 9000s uses homegrown cores, based on Arm ISA.
- Kunpeng (CPU): Arm-based server CPUs for TaiShan servers and cloud services. Kunpeng 920 features 64 custom Armv8 cores.
- Ascend (AI Accelerator): AI processors based on proprietary Da Vinci architecture. Series includes Ascend 310 (inference), Ascend 910 series (training/inference), and upcoming Ascend 920. Positions as alternative to Nvidia GPUs in China.

Architecture:

- CPU: Arm (Kirin, Kunpeng - custom Arm-compatible cores).
- AI: Proprietary Da Vinci architecture (Ascend).
- RISC-V: Part of broader Chinese exploration. No confirmation of use in core Kirin, Kunpeng, or Ascend chips.

Foundry Partner(s):

- SMIC: Primary foundry partner due to sanctions. Manufactures Kirin 9000s, Ascend 910B/C on SMIC 7nm (N+2). Targeting 5nm with SMIC.
- TSMC: Likely manufactured Kunpeng 920 initially before sanctions tightened.

Latest Known Process Nodes:

- SMIC: 7nm (N+2) (Kirin 9000s, Ascend 910B/C). 6nm (claimed for Ascend 920). 5nm (target).

Alibaba RISC-V Champion

Dual-architecture approach:

- High-performance Arm-based CPUs/custom AI accelerators for internal Alibaba Cloud.
- Major strategic investment in RISC-V through T-Head to build a domestic ecosystem and promote self-sufficiency.

Key Silicon Products:

- Yitian (CPU): Arm-based server CPUs for Alibaba Cloud instances. Yitian 710 features 128 Armv9-compatible cores.
- Hanguang (AI Accelerator): AI inference chip (Hanguang 800) for internal use (e-commerce, search). Optimized for computer vision.
- T-Head / XuanTie (RISC-V Cores): Alibaba's RISC-V development arm. Designs and licenses XuanTie RISC-V CPU cores (IoT to high-performance). XuanTie C930 is a high-performance, 64-bit, multi-core RISC-V design (RVA23 profile) for servers/PCs. Major RISC-V IP supplier.

Architecture:

- CPU: Arm (Yitian).
- AI: Custom (Hanguang).
- RISC-V: Core architecture for XuanTie IP cores (XuanTie C930 and others). Strategic pillar for Alibaba and China.

Foundry Partner(s):

- TSMC: Manufacturer for Yitian 710 (5nm) and likely Hanguang 800 (12nm).
- Other: Foundries for XuanTie cores vary by licensee.

Latest Known Process Nodes:

- TSMC: 5nm (Yitian 710), 12nm (Hanguang 800).