2026

the SHDgroup

# RISC-V MARKET REPORT: THE ADOPTION ACCELERATES

## AI is Driving the Growth

AUTHORED BY :
**RICHARD WAWRZYNIAK**

+1 650.400.0018
info@theshdgroup.com
www.theshdgroup.com
SHD102-26

# RISC-V Market Report: Market Adoption Accelerates

## AI is Driving the Growth

**April 2026**
**SHD102-26 – Abridged version**

# Table of Contents

## Table of Contents

# Executive Summary

This is an abridged version of our 2026 RISC-V Market Report and Ecosystem Guide, provided at no charge thanks to the support of our sponsors. An unabridged version is also available with over 200 pages and comes with a spreadsheet containing over 300 tables of detailed information. In both versions, our intention is to provide a comprehensive examination of the rapidly expanding semiconductor market, including how it is evolving alongside the concurrent emergence of RISC-V and the influence of AI. The accelerating build-out of data centers for AI inferencing and training and Large Language Models (LLMs) is having a profound impact on semiconductor revenues worldwide. This impact extends to the adoption of the RISC-V ISA in an increasing number of SoCs aimed at including some level of AI functionality in the silicon solution. These impacts also extend to the Semiconductor Intellectual Property (SIP) vendors as they look to accommodate the acceleration of the different Neural Networks being used and EDA Tool vendors as they look to infuse AI functionality into their EDA tools to aid the productivity of silicon designers.

The introduction of RISC-V has fueled extensive CPU architectural exploration, visibly impacting device revenues, unit shipments, design starts, business models and IP licensing revenues on a global basis. The pervasive integration of AI across applications is a primary catalyst in today's semiconductor market. The RISC-V architecture has notably influenced SoC designers and architects and is poised to drive a substantial share of designs, revenues, and unit shipments in the coming years.

Figure 1 shows the comparison between our new forecast (late 2025 report) and our previous forecast (early 2024 report) for unit shipments. As you can see, there has been a substantial increase in unit volume projections between the two versions. Highlights include:

- Continued growth of AI and especially adoption of RISC-V in Edge AI endpoint devices
- Penetration of RISC-V into cell phones and other high-volume consumer devices drives volume
- Adoption and deployment of RISC-V by Chinese vendors in most of their systems and devices
- Continued growth in ecosystem and support for RISC-V – imbuing confidence in SoC designers
- Adoption of RISC-V by major system companies signaling confidence in the ISA
- Heterogeneous SoC design landscape favors the selection of best CPU IP core for a given function
- Ability to customize RISC-V extensions aids in creation of domain-specific silicon solutions
- Continued introduction of higher-end RISC-V cores closes performance gap with competition

**Figure 1: RISC-V SoC Unit Shipment Forecast**

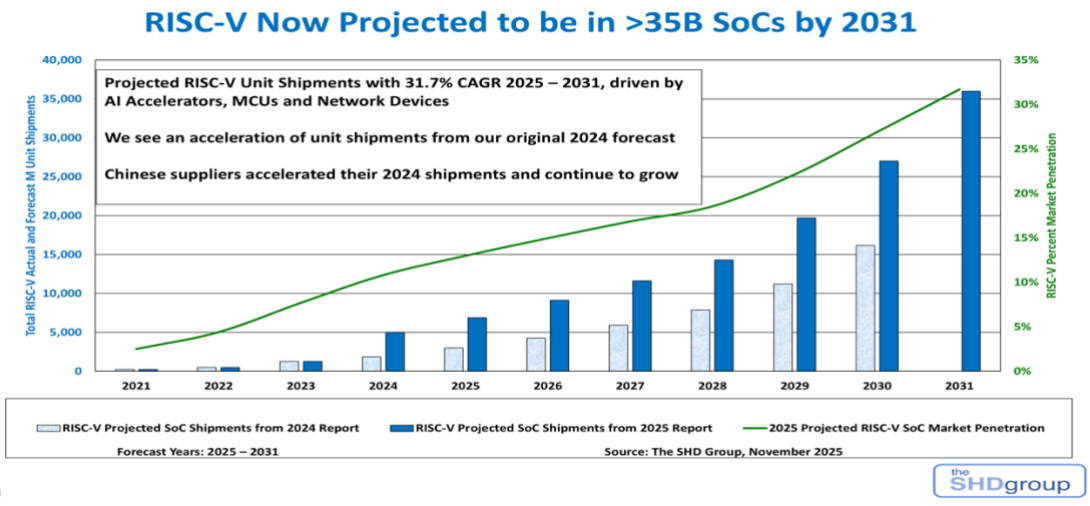

# Key Highlights

**RISC-V SoC Market Growth:**

- RISC-V-based SoC unit shipments are forecast to surge to 35.9B units, with revenues reaching $318B by 2031, boasting CAGRs of 32% and 23%, respectively.
    - The large increase in RISC-V market revenues is partly attributable to the dominance of Nvidia in the GPU market, where they are using RISC-V on every one of their SoCs.

**SoC Market Growth:**

- SoC architectures utilizing 3rd party IP exhibit substantial growth in units and revenues across Industrial, Automotive, Networking, Computer, Consumer, and other categories, notably driven by the burgeoning AI, Edge AI, and Data Center markets.
- Projections indicate SoC unit shipments reaching 113B units and revenues hitting $1.0 trillion by 2031, showcasing CAGRs of 13.7% and 6.7%, respectively.

**SoC Design Starts:**

- SoC design starts for SoCs using RISC-V CPU cores are forecast to reach 1,597 designs by 2031, a 9.7% CAGR.
- Design starts for Consumer applications are expected to show the largest number of designs by 2031, with Computer and Networking applications following closely behind.

**3rd Party IP Market:**

- In 2024, the worldwide IP market reached $9.1B, marking an 8.9% growth from 2023. Forecasts predict a 9.3% increase to $9.9B in 2025, with a potential $16.1B market by 2031, representing an 8.4% CAGR.
- The Central Processing Unit (CPU) IP market soared by 22.4% in 2024 to $3.8B and is anticipated to hit $9.7B by 2031, demonstrating a robust 13.0% CAGR.
- RISC-V IP revenues surged to $205M in 2024 and are forecasted to reach $256.9M in 2025, a growth of 25.3% over 2024. The estimated CAGR through 2031 is 39.7%
- In early 2026, SiFive announced a $400M Series G funding round, valuing the company at approximately $3.65B. We believe this underscores the increasing investor and industry confidence in the long-term market potential of RISC-V.

While this analysis excludes the impact of the emerging chiplet market, expectations lean toward increased utilization of the RISC-V architecture in future SoC designs, further bolstering IP and SoC revenues. We will continue to monitor this area.

**A special thank you to the many sponsors that contributed to the success of this report.**

# I. Introduction

This 2026 report, **RISC-V Market Analysis: Market Adoption Accelerates**, reflects how profoundly the semiconductor landscape has evolved over the past decade and a half.
In the early 2010s, the industry began to encounter the limits of traditional frequency and voltage scaling. The end of Dennard scaling meant that shrinking process geometries no longer delivered the consistent performance and power improvements that earlier generations had relied on. The primary response was a shift to multicore system-on-chip (SoC) architectures, maintaining relatively stable clock speeds while distributing workloads across an increasing number of processing cores.

During this period, most SoCs remained largely monolithic. Core counts increased, but CPU diversity remained constrained: Arm and x86 architectures together dominated the majority of major compute segments. Other CPU architectures and suppliers — including those from Andes, ARC, MIPS, PowerPC, and Tensilica — served important roles but ultimately lacked the ecosystem scale and software support required to challenge the incumbents more broadly.

That model is now evolving toward a new paradigm. The emergence of the RISC-V Instruction Set Architecture (ISA), supported by a growing global ecosystem centered around RISC-V International, has reshaped the CPU IP landscape into one that is more open, flexible, and increasingly competitive. Early commercialization was driven by SiFive—founded by the original UC Berkeley RISC-V inventors—which was instrumental in productizing the architecture, delivering the first licensable cores, development platforms, and toolchains that enabled initial ecosystem formation.

RISC-V's modular ISA and open licensing model have enabled a new class of entrants to develop CPU cores optimized for specific workloads — ranging from deeply embedded control applications to AI and data-center acceleration — while leveraging a shared and expanding ecosystem of software tools, compilers, and development frameworks.

### Heterogeneous SoCs: the new design norm

The rise of RISC-V coincides with a broader transition toward heterogeneous computing. Modern SoCs no longer rely on a single CPU architecture; instead, they integrate multiple processing elements, each optimized for a specific function. It is now common to see designs combining Arm, RISC-V, DSP, GPU, and increasingly NPU cores on the same die, operating concurrently across distinct workloads.

Rather than standardizing on a single architecture, SoC designers increasingly select cores based on workload-specific requirements. These decisions are driven by power, performance, and area (PPA), as well as safety, security, and software ecosystem considerations.

As system complexity has increased, ecosystem maturity has become a critical selection factor. Decisions once influenced primarily by licensing cost or vendor relationships now depend heavily on the availability of toolchains, verification environments, firmware stacks, and experienced engineering talent. The growing presence of EDA vendors, IP providers, and design service firms supporting RISC-V has materially reduced integration risk and accelerated time-to-market.

## A functional view of RISC-V adoption

Reflecting this heterogeneous design reality, The SHD Group structures its analysis differently from traditional CPU market studies. Rather than assigning each SoC to a single architectural category, this report focuses on the functional roles that RISC-V cores perform within a system.

This approach aligns more closely with how design teams partition compute resources and where differentiation occurs. It also provides clearer visibility into where RISC-V is gaining traction and where future growth is likely to emerge.

RISC-V deployments are segmented into four primary functional domains:

- **Deeply Embedded** — programmable replacements for fixed-function logic and state machines
- **MCU-Class** — control, sensor fusion, and low-power compute
- **Co-Processor / Accelerator** — DSP, vector, and AI-specific processing engines
- **Application Processor** — high-performance cores supporting general-purpose operating systems

This functional lens enables a more precise assessment of market dynamics than architecture-centric models.

## Industry Context: Mid-2020s Inflection Point

As the semiconductor industry moves through the middle of the 2020s, several major technology shifts are converging. These include the rapid expansion of AI and large language models (LLMs), the integration of agentic AI into edge devices, the evolution of 5G toward 6G, the commercialization of chiplet-based architectures, and the proliferation of sensor-rich, data-intensive systems across multiple industries.

Each of these trends increases both compute demand and system complexity. As a result, SoC architectures are becoming more distributed, domain-specific, and heterogeneous.

Within this environment, RISC-V's open and extensible model provides a structural advantage. Originally developed at the University of California, Berkeley in 2010 and later formalized through the RISC-V Foundation (now RISC-V International), the architecture allows companies to design and customize processors without the licensing constraints associated with proprietary ISAs.

This flexibility enables instruction-set customization and workload-specific optimization while maintaining compatibility with a growing software ecosystem—an increasingly important factor as compute becomes more specialized.

## From Academia to Commercial Deployment

RISC-V has transitioned from an academic initiative to a production-ready architecture embedded in commercial silicon programs.

- The commercial RISC-V ecosystem now includes established IP vendors such as Alibaba, Andes Technology, Codasip, Nuclei, SiFive, and others, reflecting the transition from early adoption to a competitive, multi-supplier CPU IP market.
- Leading semiconductor and system companies—including Qualcomm, Samsung, AMD, Nvidia, Amazon, Google, Microsoft, and Intel—are actively engaged with RISC-V through internal development, partnerships, or product initiatives.
- Commercial deployments are expanding across multiple domains, including wearables, edge AI, and data-center acceleration.
- Collaborations such as Google and Qualcomm's wearable platform and Meta's RISC-V-based AI inference accelerators demonstrate that RISC-V is now being deployed in high-volume production environments.
- Strategic industry activity continues to reinforce this trajectory. For example, Nvidia's investment in Intel in 2025 and their collaboration on future compute platforms suggests increasing integration of heterogeneous IP, including RISC-V subsystems, alongside established architectures.
- China represents one of the fastest-growing regions for RISC-V adoption, supported by national policy and strong participation from companies such as Alibaba and Nuclei.
- Emerging players such as Tenstorrent are contributing to ecosystem momentum through AI-focused silicon and IP offerings.

## Market Structure: Internal vs. External CPU Development

A significant portion of RISC-V adoption continues to come from internally developed cores. Many semiconductor and system companies design custom RISC-V implementations tailored to specific workloads, security requirements, or system architectures. These internal efforts represent an important source of differentiation and account for a substantial share of deployed RISC-V CPU instances.

At the same time, the commercial supply base for RISC-V IP has matured. An increasing number of vendors now offer licensable cores, development platforms, and associated services. As a result, the balance between internal development and external sourcing is shifting.

Even large semiconductor companies are increasingly incorporating third-party RISC-V IP when it improves time-to-market, reduces development cost, or complements internal design capabilities.

This normalization of external sourcing is expanding the overall market. Whether implemented through internal design, licensed IP, or off-the-shelf solutions, the outcome is the same: a rapid increase in the volume and diversity of RISC-V-enabled devices.

## Scope of This Report

This 2026 edition of *RISC-V Market Analysis: Market Adoption Accelerates* examines the current and projected impact of RISC-V across major SoC end markets, including Industrial, Automotive, Computer, Consumer, Networking, and other embedded domains.

The report provides detailed analysis of:

- RISC-V SoC market revenues across six application categories, spanning 57 end applications and 11 device types
- Unit shipments and device-level revenues for RISC-V-based SoCs
- Market penetration by device type across application segments
- Third-party semiconductor IP revenues, including CPU IP, RISC-V-specific IP, and associated licensing and royalty streams
- Competitive dynamics across CPU architectures within the IP market
- SoC design-start activity by application and device type
- Regional segmentation across the Americas, EMEA, Japan, China, and Asia-Pacific

## Market Perspective and Outlook

The RISC-V market has progressed from early exploration to structured commercialization. The interaction between internal development, commercial IP licensing, and ecosystem collaboration is creating a more diverse and competitive CPU landscape.

As heterogeneous design approaches become standard, RISC-V's flexibility positions it as an enabling architecture across a wide range of applications.

## Semiconductor Market Context (2025–2031)

Broader semiconductor industry growth provides important context for RISC-V adoption.

- Industry estimate for global semiconductor revenues is **$791.7B+ in 2025**
- Several projections indicate the market could exceed **$1T by 2026–2027**
- Longer-term estimates suggest revenues of **$1.4T–$1.6T by 2031**, assuming stable macroeconomic conditions

Within this environment, projected growth in RISC-V unit shipments and revenues aligns with broader industry expansion rather than representing an outlier.

Key growth drivers include:

- Large-scale investment in AI infrastructure
- Increasing integration of AI functionality across end systems
- Expansion of advanced wireless infrastructure (including early 6G development)
- Continued growth in sensor-driven, data-intensive applications

RISC-V is well positioned to participate in this growth due to its applicability across a wide range of compute domains.

It is also important to note that current growth rates are influenced by a relatively small installed base. As volumes increase, growth rates are expected to normalize over time, although that inflection point lies beyond the forecast horizon of this report.

## Perspective

The current transition mirrors previous inflection points in semiconductor design but operates at a broader architectural level. The move to multicore computing redefined performance scaling; the current shift toward heterogeneous, multi-ISA systems is redefining flexibility and system optimization.

RISC-V is central to this transition. It represents not only an alternative instruction set architecture, but a design model based on modularity, extensibility, and tighter alignment between hardware and software.

This report examines how that shift is unfolding—where RISC-V is being deployed, where adoption is accelerating, and how it is influencing the structure of the semiconductor ecosystem.

# II. Architectural Definition for a SoC

The SHD Group has compiled an architectural definition of the three SoC types we collect data for:

- High-end multicore SoC - silicon that functions with the most complexity, representing the highest cost range.
- Mid-range multicore SoC - silicon that represents the mid-range in device complexity and has a more moderate cost range.
- Entry-level SoC - silicon that operates with the least device complexity and occupies the lowest cost structure.

Figure 2: SoC Defined by IP Content

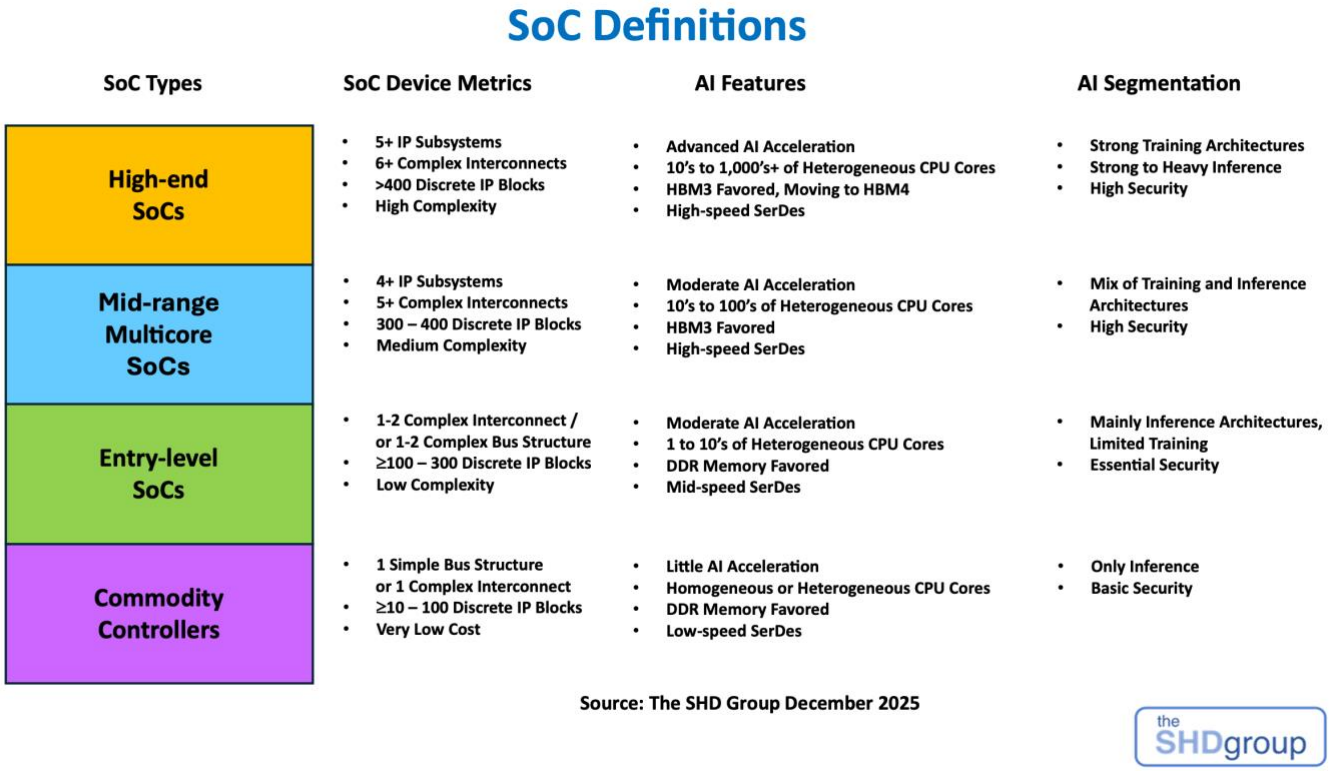


As the diagram in Figure 2 shows, we consider SoC architecture to be more about how the design is created than about the silicon itself. Contemporary SoCs today rely heavily on the use of 3rd party IP and realistically cannot be crafted any other way. With this thought in mind, we created this diagram to give some granularity to these three silicon solutions to better understand the design and market dynamics that impact each type. We have extended these definitions to include device features and segmentation of the emerging AI market solutions since this new development in the semiconductor marketplace is going to be a main driver of applications, revenues, and unit volumes for many years to come.

We expect to amend this diagram as evolving market requirements provoke responses from silicon architects and designers alike.

Figure 3: IP Support for Complex SoC Designs, Courtesy of Arteris IP

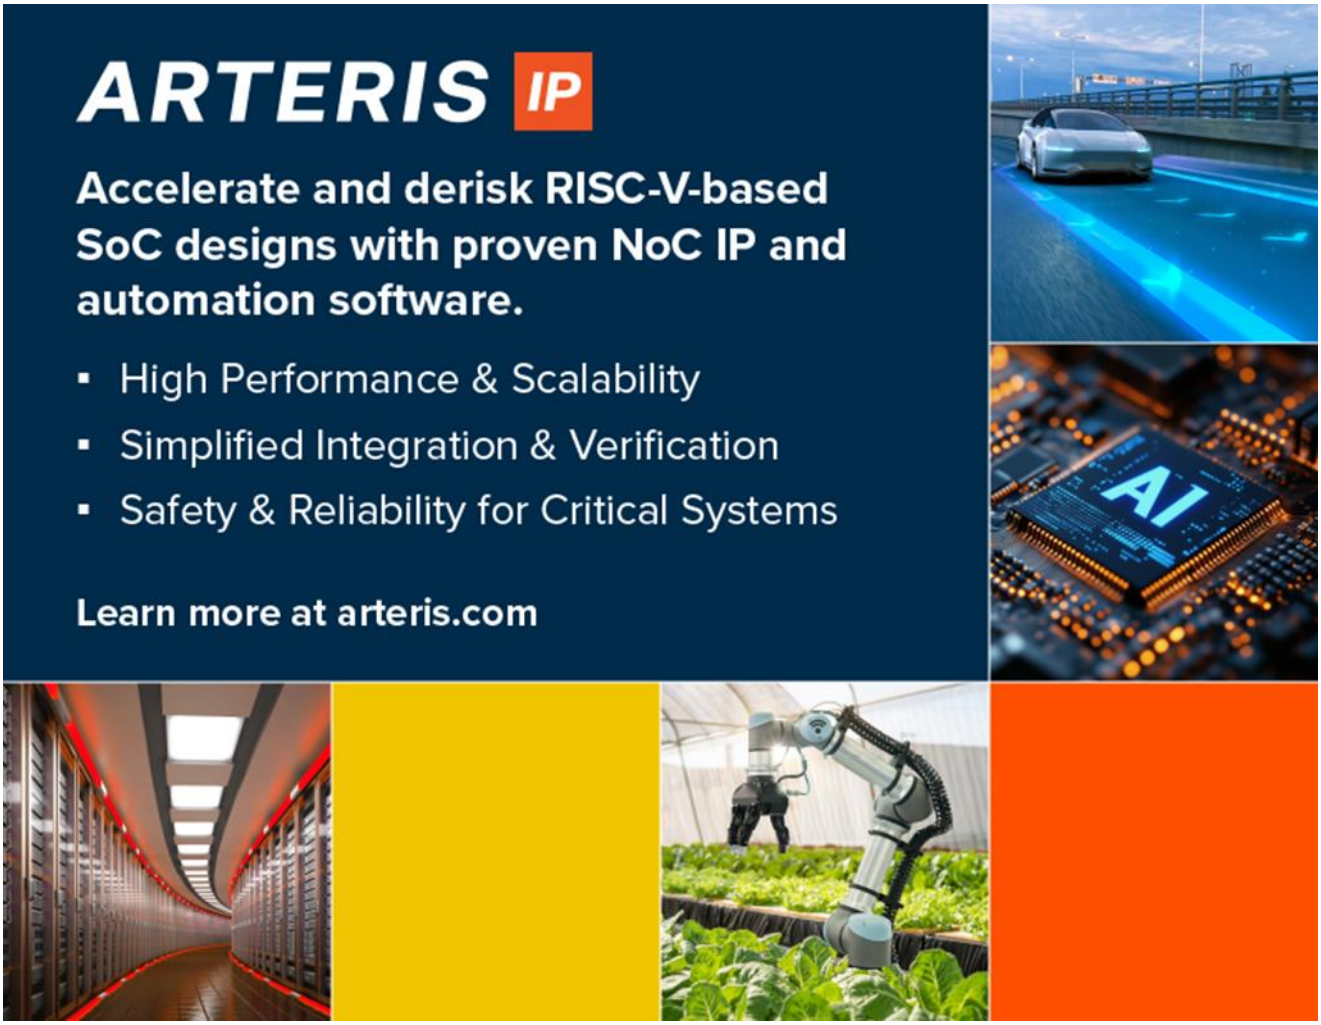


In a modern SoC, the interconnect functions as the structural backbone that links an increasing number of heterogeneous IP blocks into a cache-coherent or non-cache-coherent system. CPU clusters, AI accelerators, DSP engines, security units, memory controllers, and peripheral subsystems all generate traffic with distinct latency and bandwidth requirements. Without a well-architected interconnect, these blocks cannot meet system-level throughput, determinism, or power targets. As designs scale, dedicated interconnect IP becomes essential for predictable performance, manageable integration, and reduced verification risk.

Arteris positions its portfolio directly around this need. Its Network-on-Chip IP provides the configurable data-movement fabric required to connect all functional units within the SoC. These fabrics offer scalable topologies, high bandwidth, controlled latency, and quality-of-service management—capabilities that are central to AI processors, advanced automotive controllers, networking devices, and RISC-V-based embedded systems.

Complementing the NoC IP, Arteris supplies automation tools for topology configuration, performance modeling, and verification. These tools enable engineers to evaluate bandwidth demands, arbitration behavior, and congestion scenarios early in the design flow, shortening integration cycles and minimizing redesign.

Arteris also supports safety-critical designs with IP variants aligned to ISO 26262. These solutions incorporate error-management mechanisms, redundancy options, and diagnostic features to satisfy the reliability expectations of automotive and industrial systems.

The company's Magillem division extends this offering with system-level integration and metadata-management software. Built on IP-XACT, Magillem automates creation and validation of registers, interface descriptions, and integration files, improving consistency across large designs and streamlining handoffs between architecture, RTL, and verification teams. Together, Arteris' NoC IP and Magillem integration tools provide a unified infrastructure for assembling complex heterogeneous SoCs, combining high-performance interconnect fabrics with automation that keeps large designs organized, verifiable, and scalable.

# III. Silicon Design Trends: Rising Design Complexity

By 2025, system-on-chip (SoC) architectures have evolved into highly integrated computing platforms containing billions of transistors, hundreds of IP blocks, and diverse functional domains. A single advanced SoC may now include AI acceleration engines, dedicated NPUs, chiplet-based compute clusters, multi-standard connectivity, and increasingly heterogeneous CPU cores operating under multiple operating systems. This architectural diversity has led to exponential growth in design complexity, software content, and verification workloads.

The SHD Group's ongoing research shows that this complexity escalation is directly tied to both rising market expectations and the continuing expansion of on-chip functionality. Across all performance tiers, SoC designs incorporate ever-larger gate counts and IP portfolios to support advanced workloads spanning AI, connectivity, and real-time analytics.

Table 1 illustrates the increase in SoC complexity, measured in thousands of gates, across high-end, mid-range, and entry-level device categories. The data shows that by 2031, average high-performance devices are forecast to exceed 12 million K gates, representing a 32.8 % compound annual growth rate from 2025. Even entry-level IoT-class devices are projected to grow nearly 28 % annually as AI inference and edge connectivity requirements expand.

This growth has profound implications for design cost and methodology. Average SoC design complexity continues to outpace design-start growth by nearly 2:1, underscoring the pressure placed on Electronic Design Automation (EDA) vendors to deliver next-generation tools. By 2025, AI-driven verification, generative layout optimization, and multi-die co-design platforms have entered mainstream use, helping manage system-level complexity and accelerating tape-out cycles.

**Table 1: Device Complexity in K Gates 2025–2031**

| | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|
| **K Gates** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **High-end** | 1,652,490 | 2,199,856 | 2,974,333 | 4,032,966 | 5,338,623 | 6,974,571 | 9,198,229 | 12,061,398 | 32.8% |
| **Mid-range** | 798,039 | 988,551 | 1,132,825 | 1,314,539 | 1,555,130 | 1,832,284 | 2,197,237 | 2,639,732 | 17.8% |
| **Entry-level** | 19,401 | 24,055 | 30,337 | 38,965 | 50,023 | 64,191 | 82,334 | 105,557 | 28.0% |
| **Avg.** | 823,310 | 1,070,821 | 1,379,165 | 1,795,490 | 2,314,592 | 2,957,015 | 3,825,933 | 4,935,562 | 29.0% |
| **Growth** | 26.9% | 30.1% | 28.8% | 30.2% | 28.9% | 27.8% | 29.4% | 29.0% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 1 shows the rise in device complexity as calculated in K Gates per device over the forecast period.

Along with the EDA tool vendors, IP vendors have also stepped up with comprehensive and innovative IP solutions to help manage the increase in device complexity and functionality. Interconnect IP and Network On Chip (NOC) IP is one of these solutions.

**Figure 4: Rising Device Complexity K Gates, 2021 - 2031**

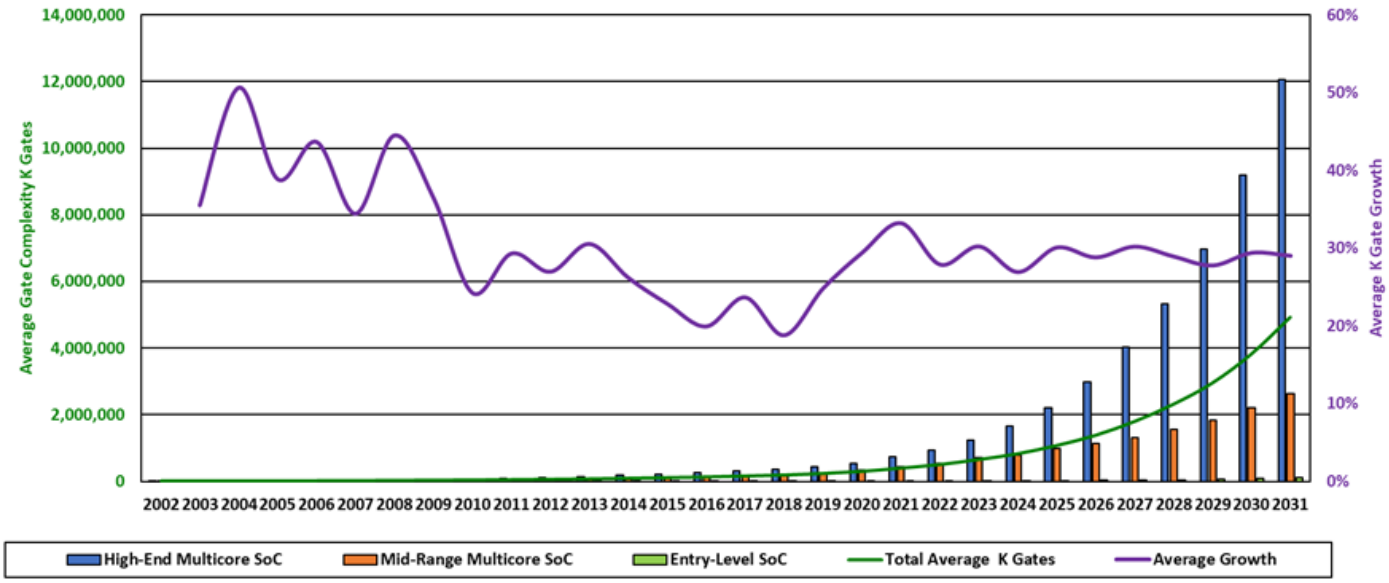


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Figure 4 looks at the three types of SoC that The SHD Group tracks by gate count and calculates how the complexity levels have risen over time. The gate counts for each device type in 2002 were used as a baseline, and then each succeeding year was divided by that value to arrive at the growth rates shown above. As this chart shows, a representative design for a high-end multicore SoC in 2025 was 462X larger than its counterpart in 2002. The mid-range multicore SoC value of 453X is larger than the high-end multicore SoCs by 2025. By 2031, advanced parts are forecast to be 1846X larger than their 2002 counterparts, and value designs are expected to be 921X larger than in 2002. Entry-level SoC designs started in 2008 and are projected to be 78X larger by 2031 than in 2008.

Looking at the SoC design landscape in this way gives some perspective on how complex silicon has become today and what those trends will look like going into the future. It is safe to say that design costs will also continue to rise, driven by the device complexity levels shown above.

As device complexity levels and silicon and software design costs have risen, the demands put upon the Electronic Design Automation (EDA) industry to continuously improve its tools have been intense. Fortunately for SoC designers, the EDA industry has delivered higher-performance tools to the semiconductor industry to aid in designing these very complex silicon solutions.

Adding to the rise in device complexity is the increasing transistor budget at each new node. Starting at 16nm in 2014 and projecting out to 2027, using TSMC's just-disclosed metrics for 2nm and 1.4nm designs, the transistor budget is increasing from 28.8M transistors/mm² in 2014 for 16nm designs to approximately 500M transistors/mm² for 1.4nm designs in 2027, which is a 17.4X increase in 10 years!

At the same time, TSMC is proposing a move from a FinFET transistor architecture to a gate-all-around FET (GAAFET) architecture starting in 2024. TSMC is also proposing a change in the reticle size from the

current 26 mm x 33 mm (858 mm²) to 2574 mm² (3X the 2020 reticle size) in 2021. TSMC plans to increase it again in 2023 to 3432 mm² (4X the 2020 reticle size).

Finally, TSMC is proposing to incorporate 2D metamaterials and carbon nanotubes (CNTs) into the manufacturing flow. The anticipated start date is late 2021 with CNTs and then branching out to include other 2D metamaterials after 2023. Material possibilities include molybdenum disulfide, hexagonal boron nitride, and titanium disulfide.

**Table 2: TSMC Metrics for Current and Advanced Process Geometries**

| | Geometry / Intro Year / Density | Transistor Type | Current Material | New Material |
|---|---|---|---|---|
| **TSMC** | 16nm: 2014 28.8M Transistors / mm² | FinFET | CMOS | |
| | 10nm: 2016 52.5M Transistors / mm² | FinFET | CMOS | |
| | 7nm: 2018 100.5M Transistors / mm² | FinFET | CMOS | |
| | 5nm: 2020 171.3M Transistors / mm² | FinFET | CMOS | |
| | 3nm: 2022 291.2M Transistors / mm² | FinFET | CMOS | Carbon Nanotube |
| | 2nm: 2024 ~379M Transistors / mm² | GAAFET | CMOS | Nano Sheets |
| | 1.4nm: 2027 ~492M Transistors / mm² | GAAFET | CMOS | Nano Wires |
| **Samsung** | 3nm: 2022 | GAAFET | CMOS | 2D Materials+ |
| | 2nm: 2025 | GAAFET | CMOS | |
| | 1.4nm: 2027 | GAAFET | CMOS | |
| **Intel** | 18Å process (~1.8nm) MP 2H2025 | RibbonFET | CMOS | |
| **TSMC** | Single Reticle Size: 26mm x 33mm = 858mm² | | | |
| | TSMC developing CoWoS (Chip-on-Wafer-on-Substrate) packaging to stitch across multiple reticles | | | |
| | 2024 generation: 2905mm² (3.5X Single Reticle Size) | | | |
| | 2025 generation: 4565mm² (5.5X Single Reticle Size) | | | |
| | 2027: Exploratory; 7470mm² (9X Single Reticle Size) | | | |

**Source: TSMC Annual Report and The SHD Group, December 2025**

Table 2 provides some metrics relating to TSMC's current process geometries and their planned 3nm and 2nm technologies. There are going to be big changes at the silicon foundries, and we can expect to see corresponding changes in the EDA tools as well. Intel has also reinvigorated its process geometry roadmap to advance four technology nodes in five years, essentially drawing even with TSMC by 2025. Samsung has also been active in introducing a 3nm process in 2022.

The transistor density race remains central to this trend. Between 2014 and 2025, transistor budgets have increased from roughly 28.8 million transistors/mm² at 16 nm to nearly 500 million transistors/mm² at 2 nm — a 17-fold improvement in one decade. TSMC's 2 nm gate-all-around (GAAFET) process entered early production in 2025, while Intel's 14A (1.4 nm) RibbonFET and Samsung's 2 nm GAAFET nodes are now in parallel development. Below 2 nm, TSMC and Samsung are exploring nanosheet and nanowire architectures with 1.4–1.6 nm research targets by 2028.

Reticle size expansion continues to enable larger stitched die and chiplet systems. TSMC's CoWoS-L packaging platform, introduced in 2024 and scaled in 2025, supports effective reticle-equivalent areas exceeding 4,500 mm² — over five times the area of 2020-era reticles. This capacity allows designers to integrate multiple logic and memory tiles on a single substrate, further contributing to system-level complexity.

Table 3 shows the transistor budget by process node from 28nm through 1.4nm and the expected years of introduction. At the time of this writing, the 2nm node is projected to be introduced in 2025. This may slip to 2026 and would push out the projected 1.4nm node to 2028 or later. The SHD Group will be closely following this in the near term.

**Table 3: Transistor Budget by Process Geometry 28nm–1.4nm**

| | 2012 | 2014 | 2015 | 2016 | 2018 | 2020 | 2022 | 2025 | 2028 |
|---|---|---|---|---|---|---|---|---|---|
| **Process Geometry** | 28nm | 20nm | 14nm | 10nm | 7nm | 5nm | 3nm | 2nm | 1.4nm |
| **Transistor Budget (M)** | 13,774 | 23,438 | 28,400 | 52,500 | 100,500 | 171,300 | 291,200 | 379,000 | 492,000 |
| **Growth Rate** | 182.6% | 70.2% | 21.2% | 84.9% | 91.4% | 70.4% | 70.0% | 30.2% | 29.8% |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

**Figure 5: Transistor Budget 1995–2029**

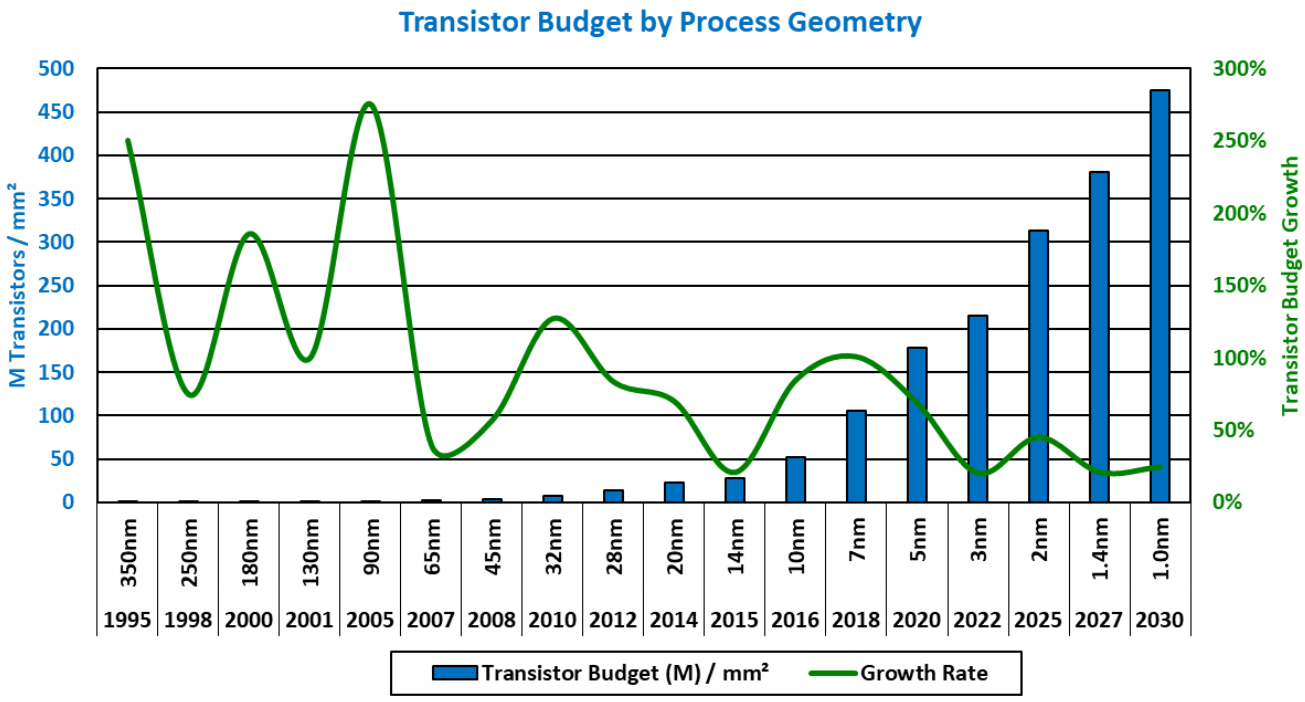


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Figure 5 shows the transistor budgets by process geometry and the projected budgets for 2nm and 1.4nm. The annual growth rates are calculated in two ways: by process node, 350nm–1.4nm, and by year, 1995–2025. Please note that projections for transistor budgets for process nodes below 1.4nm have not yet been made to any great degree. However, a hypothetical 1.0nm node would almost certainly have a transistor budget to exceed 600M-625M per mm². This is a speculative number but is well within the realm of possibility.

As these metrics show, such changes as depicted above are broad, deep and far-reaching. Given how design costs have risen, great care must be taken to ensure these expensive designs are completed as exactingly as possible and enter their markets on time. New tools to aid in the conceptualization and design of these complex silicon solutions are needed by the industry to keep the innovation and evolution of the market continuing.

To complete our analysis of the technology trends driving the increase in device complexity, we now look at the trends in the 3rd party IP market. Table 4 shows the average number of IP blocks for the three types of SoCs from 2024 to 2031.

Table 4: Average Number of IP blocks in SoC Designs 2024–2031

| | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|
| **Avg. Number of Blocks** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **High-end Multicore SoC** | 456 | 514 | 582 | 653 | 746 | 793 | 839 | 873 | 9.2% |
| **Mid-range Multicore SoC** | 308 | 355 | 388 | 419 | 455 | 486 | 515 | 547 | 7.5% |
| **Entry-level SoC** | 186 | 204 | 221 | 247 | 265 | 278 | 299 | 317 | 7.6% |
| **Total Average** | 317 | 357 | 397 | 440 | 489 | 519 | 551 | 579 | 8.4% |
| **Percent Growth** | 11.1% | 12.8% | 11.1% | 10.7% | 11.1% | 6.2% | 6.2% | 5.1% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Figure 6: Average Number of IP Blocks by SoC Category, 2021 - 2031

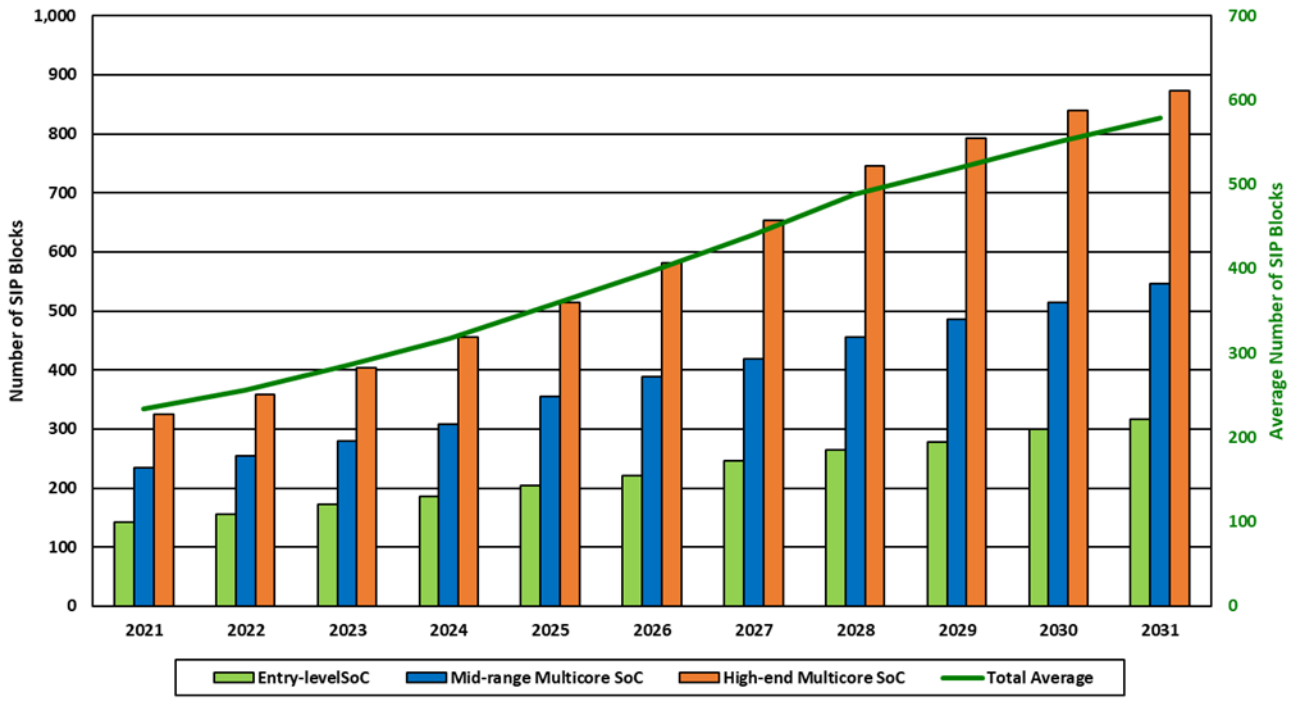


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

As Table 4 and Figure 6 show, there has been a continual increase in the number of IP blocks used in designs at all levels of complexity.

There are several things to consider when understanding these numbers:

1. This data is calculated by counting all IP blocks in the design regardless of whether it was acquired in the 3rd party market or internally developed for reuse. This means that even if a block was retained from a previous design or was developed internally, it was counted in these numbers.
2. This data counts the numerous instantiations of memory blocks, for example, registers, as one block and not by the number of instantiations of those blocks.

3. State machines, 4-bit micros and other design housekeeping blocks were counted as IP blocks.
4. These numbers are calculated using the number of blocks per design multiplied by the number of designs and then divided by the number of designs.

The first item that stands out is the constant rise in the number of blocks employed, even in simple implementations. An example is the entry-level SoC data that shows an 8.4% CAGR from 2021-2031. Many of these parts are used in IoT and IIoT applications and are indicative of the rise in complexity given the requirements for AI inference and connectivity functions. This trend is mirrored in the number of IP blocks used in high-end multicore SoCs with a 9.2% CAGR and mid-range multicore SoC designs with a 7.5% CAGR. In comparison, SoC design starts only show a 5.1% CAGR over the same time.

**Figure 7: Diverse IP Product Portfolio supports the Market, Courtesy of SiFive**


SiFive's portfolio is organized around a small number of leading product families, including the Essential line for embedded control, the Performance line for application and infrastructure compute, the Intelligence line for AI and machine learning processing, and the Automotive line for safety and security-oriented vehicle designs. SiFive complements this hardware portfolio with software enablement, development platforms, and integration support that help customers adopt RISC-V in more complex SoCs.

Over the last several years, SiFive has strengthened its position through steady advances in commercial processor releases and customer engagements. Leading products such as the Performance P670 and P870, the Intelligence X280 and newer X100 family, and the Automotive portfolio have broadened its reach across consumer, automotive, edge AI, and infrastructure applications. This product strategy has allowed SiFive to participate across multiple layers of the RISC-V market, from embedded control functions to Linux-capable application processors and AI acceleration.

Customer traction supports that positioning. SiFive has reported design wins with more than 100 customers, including eight of the top ten semiconductor companies, and has disclosed named engagements with Renesas in automotive, Sophgo in high-performance AI processors, Arkmicro in automotive SoCs, and Upbeat Technology in low-power MCU products. These wins indicate that SiFive's IP is being adopted across several of the most important RISC-V end markets rather than being concentrated in a single segment.

SiFive has played a central role in expanding the architecture across market segments. Its combination of processor breadth, software support, and customer adoption has contributed to the growth of RISC-V SoCs and to the broader maturation of the ecosystem. Given its portfolio depth and established market position, SiFive is likely to remain one of the major commercial drivers of RISC-V adoption through the rest of the decade.

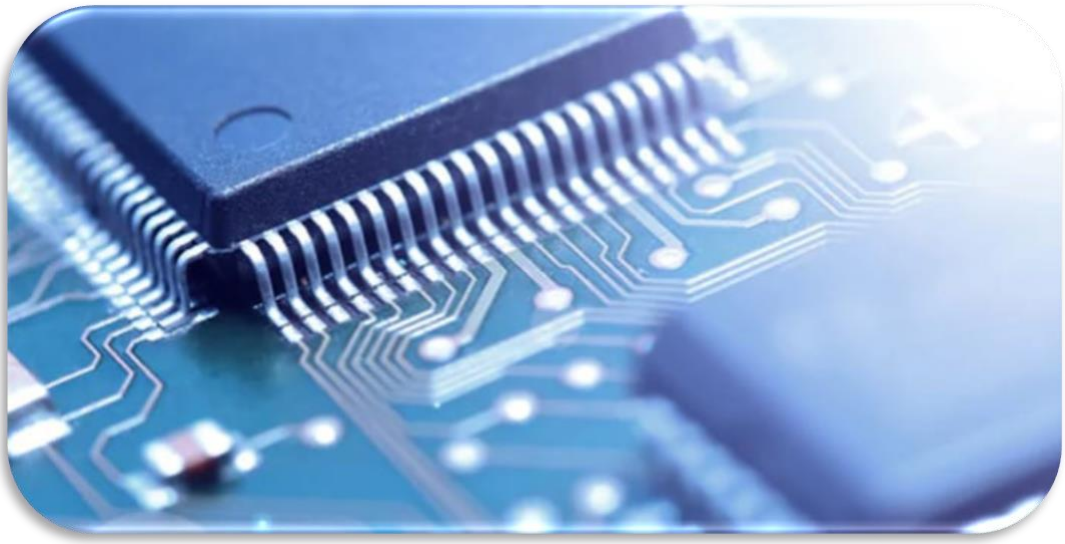

# IV. RISC-V Market Drivers and Challenges

This section looks at the trends and drivers powering the semiconductor and SoC market in the near term. Correspondingly, there are also many challenges that these markets face that could put a damper on growth going forward.

## General Market Drivers and Trends

- Pervasive deployment of sensors, notably in IoT applications.
  - Increasing usage and deployment of smart sensors that employ CPU core in the solution
- Ongoing expansion of wireless connectivity.
- Enhanced feedback and control over Industrial processes via embedded vision solutions.
- Persistent demand for increased mobile device bandwidth and requisite infrastructure.
- Growing use of FPGA solutions as accelerators in data centers alongside main CPU functions.
- Renewed focus on factory floor revitalization with connectivity and IIoT solutions.
- Electrification of automotive vehicles and the adoption of driver-assist systems.
- Extensive deployment of AI in various applications.
- Development and refinement of semiconductor architectures tailored for AI inference operations
  - Increasing incorporation of AI in endpoint devices and applications.
  - Growing significance of voice recognition across various applications.
  - Initiatives by tech giants like Microsoft, Google and Apple to integrate user-accessible AI assistance into search engines and business applications.
- Introduction and rapid adoption of OpenAI's ChatGPT application signifying a substantial impact on the economy and society, thereby influencing the IP, SoC and EDA markets.
  - Declarations by Google, Qualcomm, Samsung and Intel to enable running ChatGPT applications on their respective smartphone and desktop PC, laptop and tablet silicon, potentially rejuvenating these markets.
  - Efforts by Microsoft to infuse AI functionality into the Windows OS and Office suite of business applications.
  - Introduction of Microsoft’s own LLM models with considerably reduced parameter counts.
- Implementation of new transistor architectures at 3nm and 2nm extending Moore's Law.
- Accelerating momentum of chiplets with standardization in interconnects and efforts also underway for standardizing on packaging.
- Significant increase in reticle size by TSMC using their CoWoS packaging technology to stitch multiple reticles together. Planned increase in reticle size of 5.5X by 2025.
  - The introduction of Apple’s M1 Ultra, featuring 114B transistors in a two-die package, are likely an example of this.
  - IMEC proposed a new process technology roadmap extending beyond 2036 that reaches 0.2nm (200 picometers), reaching a potential 1 trillion transistors per square mm.
- Integration of AI functionality into EDA design tools to enhance designer efficiency.
- Ongoing rollout of 5G infrastructure worldwide.
- Renewed interest in older process geometries due to advancements in EDA toolsets accommodating higher complexity.
- Convergence of 5G, AI, IoT, and 3D Printing fostering new market demands and business models.

- Emergence of Silicon Lifecycle Management IP facilitating powerful analytics embedded into various SoCs, available from multiple vendors today.
- Multiple FPGA vendors transitioning portions of their product lines or segments to incorporate RISC-V CPU cores, boosting the momentum of RISC-V.
    - Lattice Semiconductor, Microchip, QuickLogic, Efinix.
- Substantial announcements for new fab capacity from various industry leaders to address potential capacity shortfalls in the near future.
    - New capacity by Intel, Samsung, TSMC, GlobalFoundries and Micron promises to alleviate capacity shortfalls in 2-3 years.
    - Intel has re-entered the silicon foundry market and introduced an aggressive geometry development roadmap through 2027.
- Legislative actions, such as the CHIPS and Science Act, designed to incentivize semiconductor manufacturers to build new fab capacity in the US.
    - President Biden signed the $52B CHIPS Act legislation to build new fab capacity in the US and establish new learning centers in support of the U.S. semiconductor market.
- The U.S. 3Q25 GDP came in at a revised, final read of 4.3%, signaling a growing economy compared to 2Q25. 4Q25 final for U.S. GDP is for a lower 0.5% growth.
- Very large investments made for AI infrastructure build-out. Alibaba, Alphabet (Google), Amazon, ByteDance, CoreWeave, Huawei, Meta, Microsoft, OpenAI, Oracle, and others totaling in the multiple 100s of billions of dollars.
- Large investments in Intel by the US government (US Chips Act conversion to equity), Nvidia and Softbank. Intel also sold 51% of its share in Altera to Silver Lake – all to bolster sagging revenues
    - Nvidia and Intel collaborating on joint development of next-generation CPUs.
- As part of a strategic partnership, OpenAI will receive warrants to purchase up to 10% of AMD.
- Nvidia invested $2B in Synopsys as a strategic technology alignment to create new design flows for AI and multi-die chiplet architectures.

## RISC-V-Specific Drivers

Positive developments specific to RISC-V in the market and ecosystem include:

- Nvidia announced at the 2024 RISC-V North American RISC-V International Summit they are using RISC-V in **all** of their silicon solutions, anywhere from 10 to 40+ CPU cores per part.
- Introduction of the RISC-V Software Ecosystem (RISE) aimed at accelerating the development of open-source software for RISC-V architecture.
- Embracing RISC-V ISA by well-known CPU IP vendors like Andes, MIPS, SiFive, and Synopsys/ARC (-> Global Foundries), expanding the list of IP vendors supporting RISC-V.
- Collaborative efforts by Qualcomm and Google to develop a RISC-V-based platform for wearables
- Increasing availability of EDA tool vendors offering design tools tailored to RISC-V designs.
- Ongoing enhancements in higher-performing RISC-V CPU cores offered by RISC-V CPU IP vendors.
- Qualcomm Technologies, NXP Semiconductors, Infineon Technologies AG, Nordic Semiconductor ASA, Robert Bosch GmbH and STMicroelectronics form a new company named Quintauris, based in Munich, Germany to be a single source to enable compatible RISC-V-based products, provide reference architectures, and help establish solutions to be widely used in the Automotive market and other industries as well.
- Meta introduces a training and inference accelerator SoC for data center applications based on the Andes RISC-V AX25.

- Alibaba's Damo Academy introduced a new server accelerator SoC based on a RISC-V CPU core that was internally developed.
- Red Hat has announced that it is working toward support for RISC-V, starting with a developer preview of Red Hat Enterprise Linux (RHEL) 10 on a RISC-V platform.
- NVIDIA announced that CUDA will support RISC-V as a host architecture (i.e., enabling RISC-V CPUs to run CUDA system/control logic) in addition to x86 and ARM.
- Nvidia made a $5.0B investment in Intel and the two companies will collaborate on designing new MPUs for Data Center, Computer and Consumer applications,
- Nvidia entered into a non-exclusive IP licensing agreement with Groq for a reported $20B at the end of 2025. The top management of Groq will also join Nvidia, but Groq will continue on as an independent company.
- Global Foundries has acquired the RISC-V IP and tools of MIPS and the ARC and ARC-V RISC-V IP, and tools from the CPU IP group from Synopsys. Both groups will continue licensing that IP to new customers and to support existing customers and operate as separate groups under the Global Foundries umbrella.
- The announcement by Nvidia at last year's North American RISC-V International Summit that they have been using RISC-V CPU cores, not in just some of their parts, but in ALL of their parts comes as a major validation of the RISC-V ISA. They had stated that the main reason for their support was the RVA23 extensions to the RISC-V ISA which have allowed them great flexibility in crafting solutions that exactly fit the application. This is a significant event since Nvidia currently commands a great deal of attention in the market and this announcement is sure to generate even more interest in the RISC-V ISA and removes one more potential objection to looking at or using RISC-V IP in a company's design.
- SiFive announced the expansion of its ecosystem and platform through collaboration with Red Hat to bring Red Hat Enterprise Linux to RISC-V systems and infrastructure partnerships such as NVIDIA's NVLink Fusion for next-generation AI data-center architectures.
- Tenstorrent's Blackhole AI accelerator represents a major advancement for the RISC-V market by integrating high-performance AI compute with RISC-V control cores, helping expand the open architecture into large-scale AI training and inference systems.

Figure 8: Tenstorrent – Innovating in the AI and RISC-V Landscapes, Courtesy of Tenstorrent

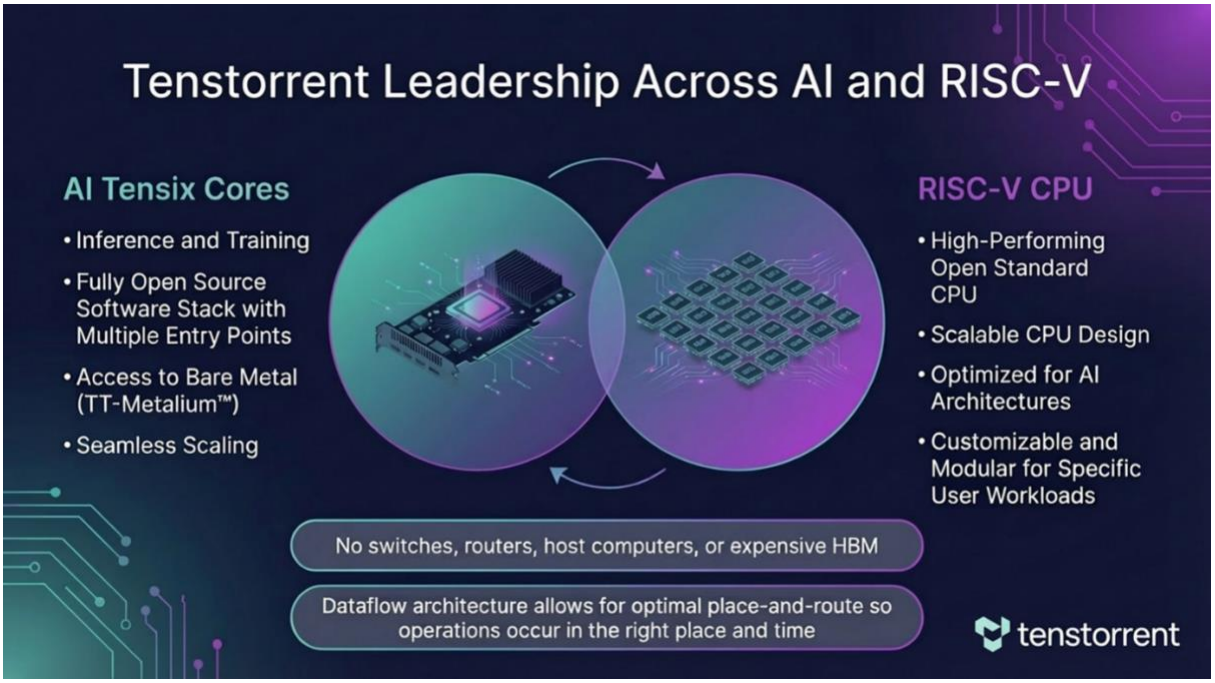


Tenstorrent is a next-generation compute company focused on high-performance RISC-V CPU and AI acceleration IP for data center, automotive, robotics, and infrastructure markets. The company offers a unified portfolio that includes Ascalon™ high-performance RISC-V CPU IP, Tensix™ AI acceleration IP, embedded and real-time RISC-V cores, and BlueLynx™ die-to-die interconnect and system IP, supported by an open software stack.

Tenstorrent's approach is centered on tightly integrated hardware–software co-design, where RISC-V serves as a foundational control and compute layer within heterogeneous architectures rather than as a standalone embedded processor. Its systems are built around dataflow-oriented compute, emphasizing high on-chip bandwidth and scalable interconnect fabrics to efficiently support AI training and inference across a range of deployment environments.

By combining licensable IP, development platforms, and shipping silicon, Tenstorrent enables customers to build differentiated, RISC-V-based SoCs and systems while retaining long-term control over their compute roadmap. This approach allows teams to integrate CPU, AI, and chiplet building blocks on a common foundation, shorten time-to-market, and reduce dependence on closed, proprietary architectures.

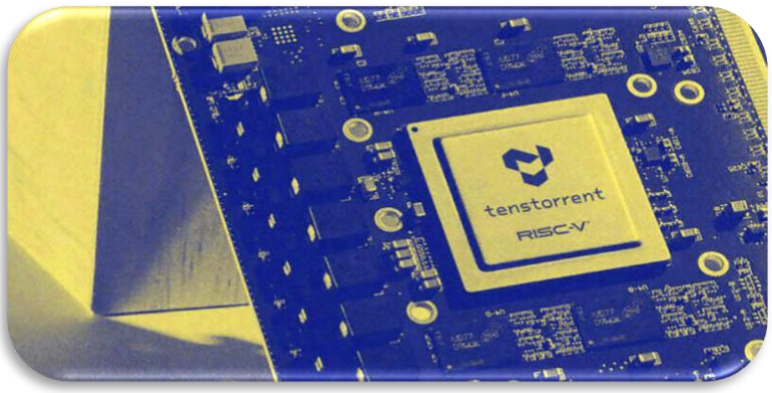

# V. Effects of AI on RISC-V SoC Designs

Artificial intelligence (AI) is no longer an emerging workload class. Rather, it is the defining force shaping semiconductor and SoC design strategies in 2025. The fusion of AI acceleration, open instruction set architectures, and advanced packaging has transformed system architecture at every level. From hyperscale datacenters to low-power edge devices.

RISC-V now plays a pivotal role in this evolution, serving as both a compute and control fabric within heterogeneous AI SoCs.

The market has clearly shifted from exploring AI integration to optimizing and scaling it across virtually all semiconductor segments.

## AI's Architectural Impact on System-on-Chip Design

The industry has entered a post-Moore's law era where specialization and workload-optimized architectures are starting to dominate. The exponential scaling of model complexity, now exceeding approximately 1 - 2 trillion parameters in foundation models, has forced innovation in efficiency rather than raw transistor count. At the same time, model compression, quantization, and retrieval-augmented inference (RAI) techniques have enabled high-value AI workloads to run entirely on-device. This removes the persistent dependency on cloud connectivity that limited earlier-generation applications. The shift has unleashed a new wave of SoC designs integrating AI capabilities directly at the silicon level, from consumer wearables to industrial controllers.

Heterogeneous computers have become the standard architecture. Modern SoCs often integrate RISC-V cores alongside Arm cores, dedicated NPUs, vector processors, and specialized DSP blocks to deliver inference capabilities with vastly improved power efficiency. RISC-V's modular instruction set extensions, particularly the vector 1.0 and p-extensions, allow vendors to tailor compute pipelines for specific AI workloads. This has resulted in SoCs that can execute diverse AI tasks at the edge while maintaining deterministic performance and low latency.

Custom instruction synergy for AI acceleration is a major strength of RISC-V. This is witnessed in AI designs utilizing the ability for hardware engineers to integrate custom instructions for domain-specific acceleration. Designers can implement specialized tensor, matrix, or sparsity operations without breaking compatibility with the base ISA. This flexibility enables co-optimization of hardware and software pipelines, shortening development cycles and delivering higher performance per watt for inference and training workloads. In practice, this has made RISC-V an ideal foundation for hybrid AI SoCs combining control, signal processing, and neural acceleration in a unified architecture.

Figure 9: Driving RISC-V In AI

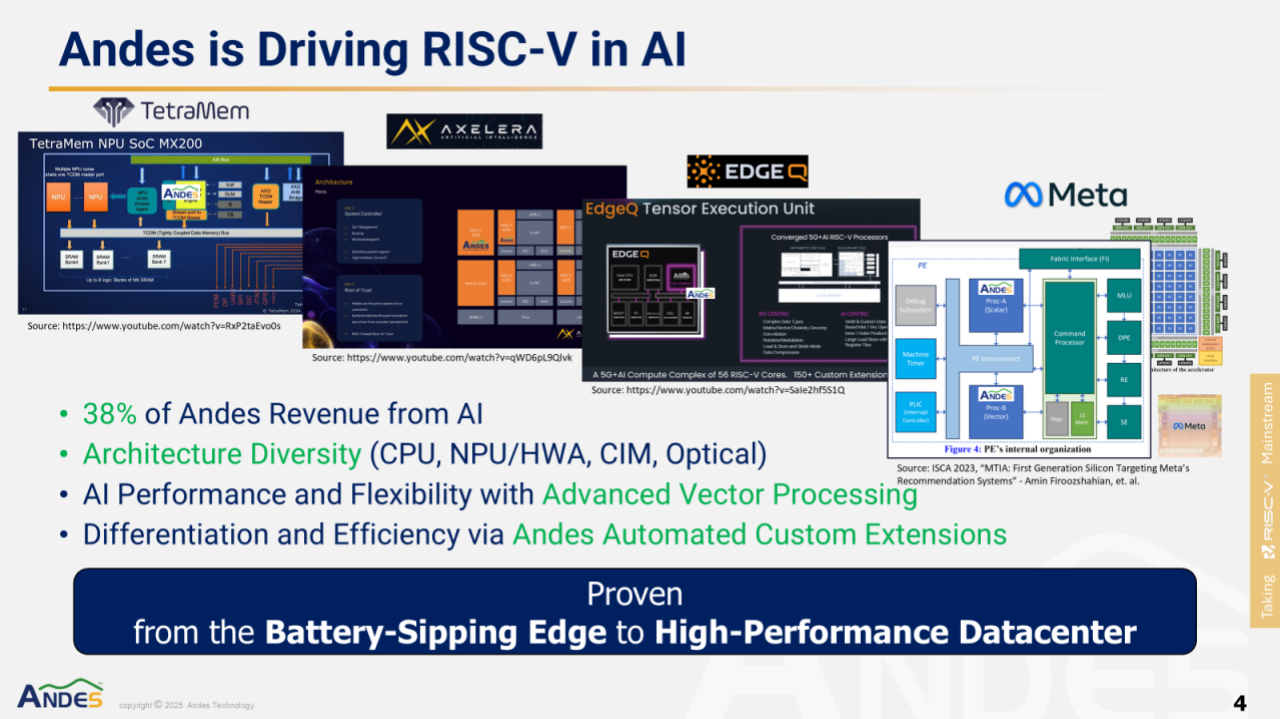


Andes holds a significant position in the RISC-V ecosystem as one of the founding members of RISC-V International. As this slide shows, they focus on supplying CPU IP, architectural extensions, and supporting technologies that enable AI-oriented, heterogeneous SoC designs. Andes provides a portfolio of RISC-V processor IP cores used across a range of compute subsystems—including NPUs, NPUs with custom hardware acceleration, compute-in-memory engines, and other AI-focused blocks. Their cores appear in products from companies such as TetraMem, Axelera, EdgeQ, and Meta, indicating broad commercial adoption.

Their IP portfolio delivers great architectural diversity spanning general-purpose CPUs, neural-network accelerators, hardware-assisted compute paths, compute-in-memory architectures, and optical-processing elements. This positions Andes as an enabler for system designers who need different compute modalities under a common RISC-V framework.

Advanced vector-processing capability is highlighted as a key differentiator. Andes offers vector extensions and performance-tuning features that target high-throughput inference and training workloads at both the edge and in datacenter environments.

The company also provides automated custom-extension technology, supporting designers who want to add domain-specific instructions and tightly integrated accelerators without breaking toolchain compatibility. This is one of their most visible contributions to RISC-V's flexibility as an open ISA.

Across these areas, Andes effectively supplies the RISC-V market with licensable CPU cores, AI-optimized extensions, customization frameworks, and toolchain support aimed at enabling designs that range from ultra-low-power edge nodes to high-performance compute clusters.

## A Framework for AI Functional Levels in 2025

To contextualize AI capability in SoC markets, The SHD Group defines five levels of AI functionality. Each is tied to hardware performance and system complexity rather than theoretical intelligence milestones.

This approach reflects the reality of how semiconductor design now segments AI performance tiers.

**Table 5: AI Functional Levels and RISC-V Hardware Tiers**

| Level | Hardware Characterization | Representative Workload |
|---|---|---|
| Level 1: Foundational AI | Fixed-function accelerators or DSP-class RISC-V cores | Keyword detection, gesture recognition, simple classifiers |
| Level 2: Narrow AI | Hybrid NPUs and RISC-V DSP extensions | Vision processing, audio intelligence, sensor fusion |
| Level 3: Advanced Narrow AI | RISC-V vector cores with integrated AI accelerators | Generative AI at the edge, multimodal small LLMs |
| Level 4: Contextual AI | Chiplet-based implementations or unified memory SoCs | Real-time decision systems, adaptive robotics, multi-agent inference |
| Level 5: Cognitive AI (Emergent) | Distributed architectures across SoC clusters | Self-optimizing systems, large-scale reasoning workloads |

Source: The SHD Group, April 2026

**Figure 10: Heat Map Diagram of AI Functional Levels and RISC-V Hardware Tiers**

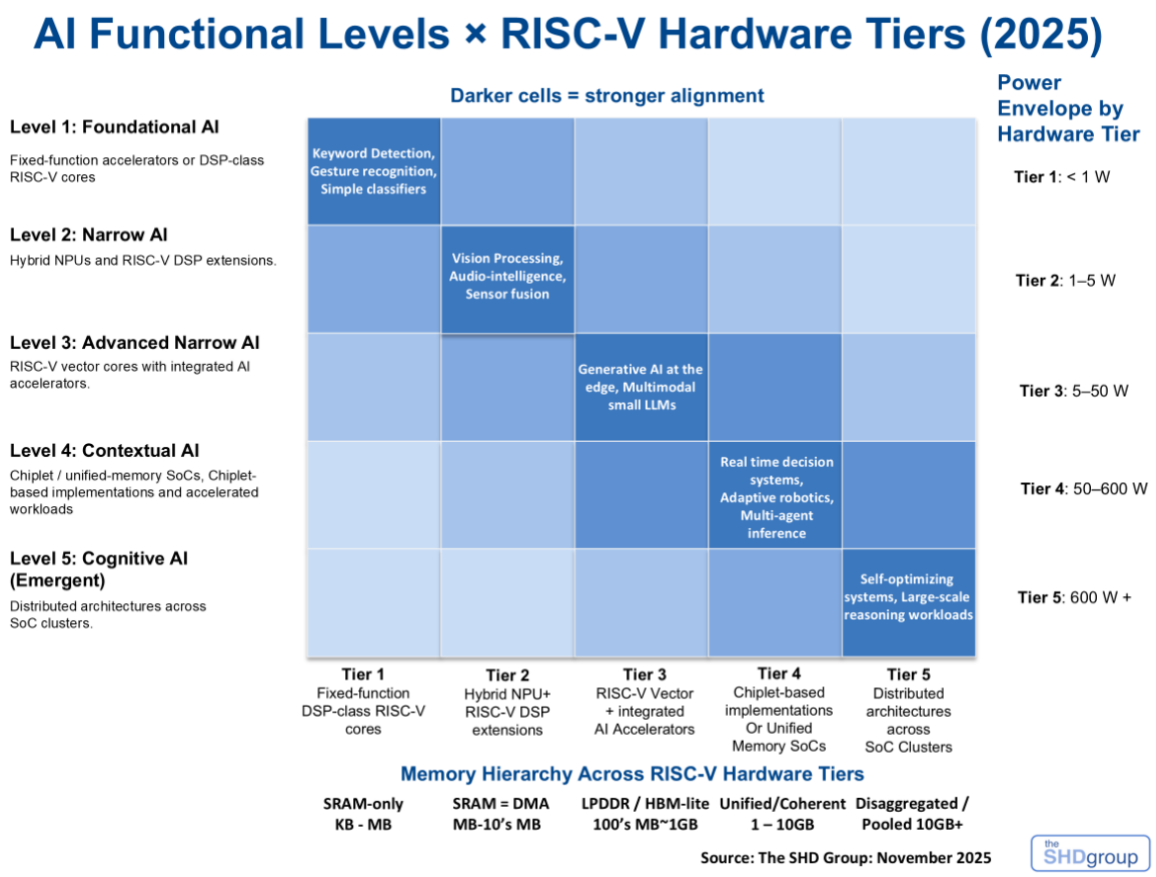

This figure illustrates the progressive mapping between artificial intelligence functional sophistication and the corresponding RISC-V hardware implementations that enable it. Beginning at the foundational level, the diagram shows how fixed-function accelerators and DSP-class RISC-V cores provide efficient compute for constrained workloads such as keyword detection, gesture recognition, and simple classification tasks. As AI workloads become more demanding, hybrid RISC-V designs incorporating DSP and NPU extensions emerge, supporting Level 2 "Narrow AI" functions like sensor fusion, audio intelligence, and early-stage vision processing within a modest power envelope.

At Level 3, labeled "Advanced Narrow AI," the chart highlights the transition to RISC-V vector cores integrated with AI accelerators, enabling edge generative AI and small multimodal language models. This represents a major inflection point where RISC-V hardware begins bridging embedded systems and higher-order inference workloads. Moving further up, Levels 4 and 5 - Contextual and Cognitive AI - reflect the evolution toward chiplet-based SoCs and distributed multi-cluster architectures. These tiers support real-time decision systems, adaptive robotics, and self-optimizing reasoning frameworks requiring unified memory and coherency across multiple dies or nodes.

The lower section of the diagram connects each tier to its corresponding memory hierarchy, ranging from simple SRAM-based implementations at the foundational level to disaggregated or pooled memory systems at the cognitive tier. This progression captures how AI functionality and memory architecture co-evolve within the RISC-V ecosystem: increasing intelligence drives more parallelism, greater coherency requirements, and higher bandwidth demands. Together, the figure demonstrates how RISC-V's modular and open ISA enables scaling AI compute efficiently - from deeply embedded intelligent sensors to distributed, cognition-class systems—while maintaining architectural consistency across the hardware spectrum.

## RISC-V's Expanding Role in AI Ecosystems

**Industry consolidation and strategic acquisitions (2025)**
This year has seen major consolidation shaping the open-ISA landscape. Meta's acquisition of RIVOS brings deep RISC-V based AI server design expertise directly into one of the world's largest AI operators. This move underscores big tech's intent to build proprietary silicon optimized for LLM inference and multimodal workloads using RISC-V's open architecture. In parallel, GlobalFoundries' acquisition of MIPS Technologies marks a turning point for traditional CPU IP. Once a competing ISA, MIPS is now fully aligned with RISC-V, enabling foundry-level design support for RISC-V customers and strengthening the end-to-end supply chain for open architectures. While consolidation occurs, other new entrants such as Ahead Computing continue to enter the RISC-V community.

Key contributors in AI-oriented RISC-V hardware include Andes Technology, and SiFive, the two dominant RISC-V IP providers globally. Their IP cores leverage RISC-V as a scalable compute substrate for inference workloads. Meanwhile, Chinese ecosystem leaders such as Nuclei, Alibaba's DAMO Academy and StarFive have accelerated development of domestic RISC-V AI processors - spurred by export restrictions and the drive for technological independence. This bifurcation between open global and nationalized AI hardware ecosystems continues to shape RISC-V's international trajectory.

Additionally, coordinated initiatives in the European Union, Japan and China are reinforcing open ISA development, supporting collaborative RISC-V research and aligning with broader goals of supply chain resilience and technological sovereignty.

**AI Compute Scaling Beyond Moore's Law**

The relationship between AI performance and transistor scaling has diverged sharply from Moore's traditional curve. Performance-per-watt for AI-centric accelerators is improving at roughly 1.4× per year - outpacing transistor scaling, driven by architectural specialization and advanced packaging (chiplets, 2.5d/3d integration). RISC-V's flexibility makes it uniquely suited to this evolution - its open ISA allows chipmakers to co-design unique and proprietary compute and memory hierarchies without legacy constraints.

As transistor scaling slows, software & hardware co-design and open ISA innovation have become the new engines of progress.

Figure 11: Performance per Watt Scaling vs. Transistor Growth 2010 - 2031

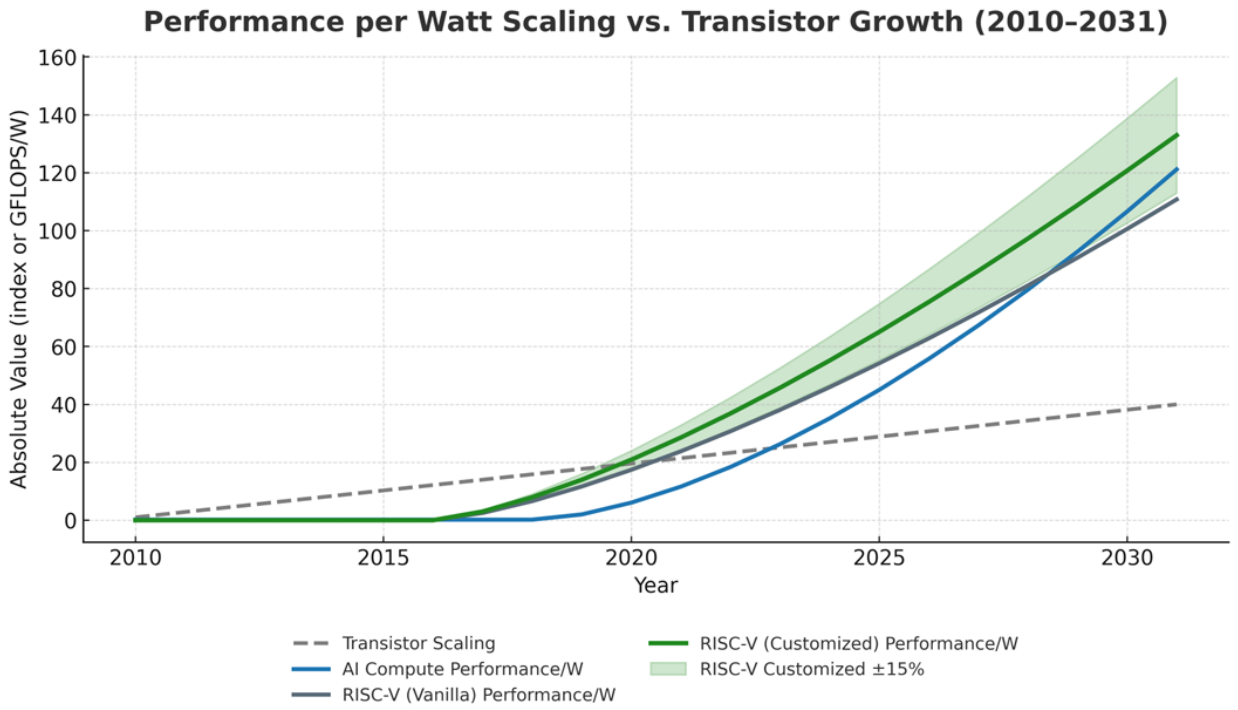


| Forecast Years: 2025 – 2031 | Source: The SHD Group, April 2026 |
|---|---|

This figure compares the long-term trajectory of transistor scaling, RISC-V architecture efficiency, and AI compute performance-per-watt improvements. Transistor scaling follows a two-year doubling trend, normalized to 2010 = 1.0. RISC-V performance data begins in 2017 with the first commercial core deployments, while AI compute efficiency is shown from 2019 onward, reflecting the emergence of dedicated accelerator architectures. The vertical blue line marks the establishment of the RISC-V Foundation in 2015, denoting the start of the open-ISA era.

Over the past fifteen years, the semiconductor industry has experienced a pronounced decoupling between transistor scaling and performance-per-watt improvements. From 2010 through 2024, transistor density has continued to increase at a near-historical rate - roughly 2.5× every five years - while energy efficiency has lagged significantly due to leakage, interconnect resistance, and memory bottlenecks.

In contrast, AI compute efficiency has accelerated sharply since 2019, driven by domain-specific architectures, advanced packaging, and workload-optimized accelerators. This performance-per-watt

growth outpaces traditional transistor scaling and demonstrates that architectural innovation now contributes more to system-level efficiency gains than lithographic advancement alone.

RISC-V architectures began showing measurable impact around 2017. Early implementations ("RISC-V Vanilla") tracked near transistor scaling trends but began diverging upward with the introduction of customized extensions and domain-optimized cores. By 2025, "RISC-V Customized" implementations - leveraging AI and vector extensions, chiplet integration, and open hardware/software co-design - are projected to deliver up to 15–20% higher efficiency than comparable fixed-ISA solutions. The performance envelope for customized RISC-V cores continues to expand, with expected gains through 2031 as SoC-level heterogeneity increases and software toolchains mature.

The chart illustrates this trend visually: transistor scaling maintains a steady but flattening trajectory, while AI compute and RISC-V performance-per-watt exhibit accelerating growth. The widening gap between process scaling and architectural efficiency underscores a key inflection point in semiconductor design strategy - the transition from process-driven performance to architecture-driven efficiency, with RISC-V positioned as a leading enabler of this shift.

**Table 6: Market Challenges and Issues – 2025 Outlook**

| Key Challenge | 2025 Market Implication |
|---|---|
| Rising AI compute demand vs. manufacturing cost | Accelerating move to chiplets and domain-specific accelerators |
| Foundry capacity concentration (TSMC, Intel, Samsung) | Supply chain resilience strategies and regional diversification |
| U.S.–China export controls | Increased adoption of RISC-V as an open alternative to restricted architectures |
| Energy consumption and thermal density | Shift toward energy-efficient RISC-V AI inference and near-memory compute |
| Software and toolchain fragmentation | Consolidation under unified compiler stacks (LLVM, MLIR, ONNX-Runtime for RISC-V) |

Source: The SHD Group, April 2026

RISC-V adoption now exceeds 1.1 billion AI-capable SoCs globally, spanning edge inference, industrial automation, and embedded vision applications.

AI has become the central organizing principle of modern semiconductor design. For RISC-V, this represents both an inflection point and a validation of the open architecture philosophy. The ISA's adaptability, competitive licensing model, and strong ecosystem support have made it a credible platform for AI across every market tier - from ultra-low-power edge devices to AI-enabled servers. As AI workloads continue to diversify and specialize, RISC-V's openness and configurability position it as a leading architecture for the next phase of intelligent computing.

Figure 12: Rise of RISC-V in Selected Markets 2021 - 2031

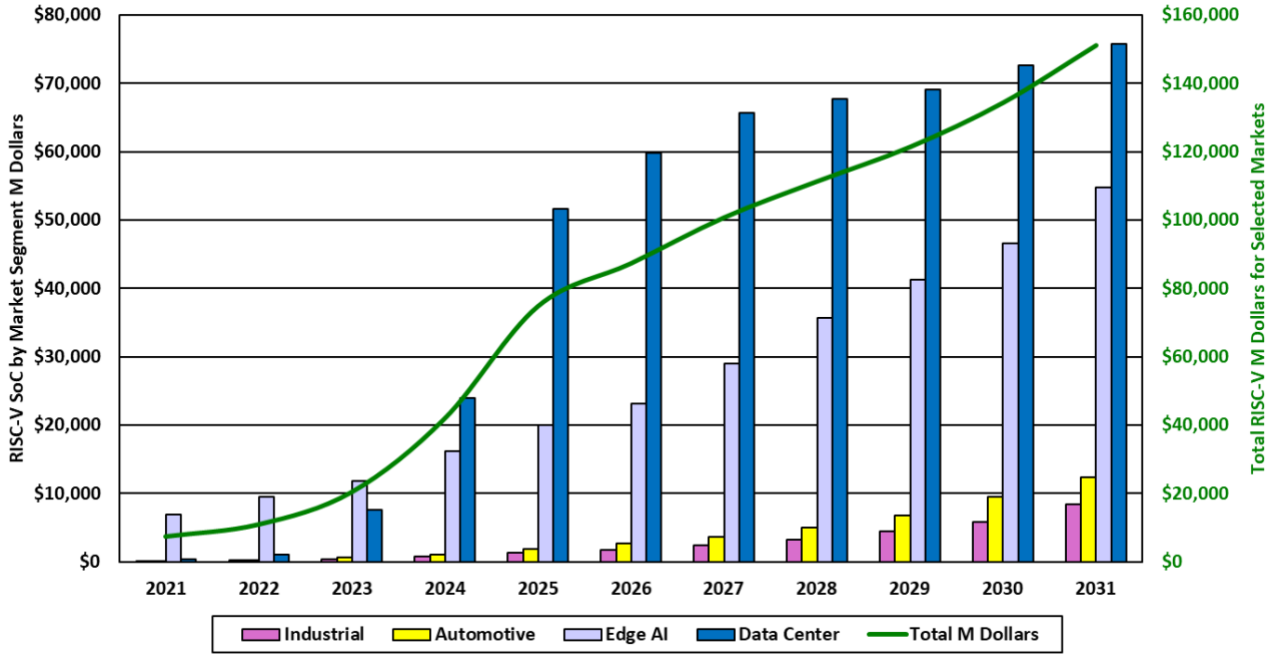


**Forecast Years 2025 – 2031** **Source: The SHD Group, April 2026**

Figure 12 shows how RISC-V-enabled SoCs are growing in the Industrial, Automotive Edge AI and Data Center markets. Starting from a very small base in 2021, RISC-V SoCs are growing at a very fast rate in these markets. This is attributable to the desire of designers to be able to customize their silicon solutions to more closely fit the end applications. This need extends from the simplest designs in Edge AI all the way highly complex designs in the Data Center. Of note in this chart is the rapid rise of RISC-V solutions in the Data Center due mostly to the dominant position of Nvidia GPUs in this market.

**Emerging Considerations: Software Portability and Regulation**

Quantization-aware training, ONNX-based model portability, emerging CUDA support, and improving cross-compilation toolchains are reducing—but not eliminating—the performance and software gap between architectures. While true parity has not yet been achieved, these advances are enabling a broader range of AI inference workloads to run across heterogeneous compute platforms, including RISC-V-based systems.

At the same time, increasing regulatory focus on AI safety, compute governance, and energy efficiency is beginning to shape semiconductor design priorities. Together, these dynamics are reinforcing interest in more flexible and energy-efficient architectures, positioning RISC-V as a viable option for certain classes of future AI SoC platforms, particularly where customization and control over the compute stack are important.

## Accelerating Neural Networks in Selected Silicon Solutions

This report looks at 57 system types and the SoC Bill of Materials used in these systems. Within these systems there are several different device types that are used to accelerate the Neural Networks being used for AI-enabled applications today. In our analysis, The SHD Group focused on several device types designers have determined are the best options to perform the acceleration. This varies by end application and functionality

We focused our analysis of AI acceleration on the following device types:

- GPU
- AI Accelerator
- DSP
- FPGA

There are other devices that can perform some forms of AI acceleration but are not listed here such as General-purpose CPU and MCUs. The following is a short explanation of our reasoning for the list of included part types and for excluding CPUs and MCUs. MCUs executing TinyML are a special case and is handled at the end of this section.

GPUs, AI accelerators, DSPs, FPGAs, General-purpose processors differ fundamentally in how they execute neural network workloads, and this distinction determines whether a device can meaningfully accelerate inference. Neural networks rely on dense linear algebra with high degrees of parallelism and tightly structured data movement. Only architectures designed to exploit these characteristics—through wide datapaths, specialized MAC arrays, deterministic dataflows, or reconfigurable compute fabrics—deliver the throughput and efficiency required for modern AI workloads. This section outlines why CPUs fall short and why parallel compute architectures succeed in accelerating neural networks.

**GPUs** accelerate neural networks because their execution model is designed for massive data parallelism. Thousands of lightweight threads execute the same operations on large tensors, and wide SIMD/SIMT datapaths deliver high MAC throughput. High bandwidth to off-chip memory and large register files help sustain the parallel compute units, making GPUs effective for convolutional and transformer workloads.

**AI accelerators** go further by dedicating most of the silicon area to dense arrays of MAC units—systolic arrays, vector engines, or spatial compute fabrics—paired with multi-level scratchpad memories designed explicitly for ML operators. Their dataflows minimize unnecessary memory movement and exploit quantization, sparsity, and mixed precision to maximize energy efficiency.

**DSPs (digital signal processors)** improve on this by providing specialized MAC pipelines, deterministic dataflows, and tightly coupled memory architectures. Their instruction sets and execution units are optimized for vector arithmetic, which maps well to the inner loops of neural network kernels.

**FPGAs** also accelerate neural networks by allowing the compute and dataflow fabric to be customized for a specific model. The reconfigurable logic enables developers to build deeply pipelined MAC arrays, tune data reuse patterns, and match on-chip memory resources directly to the network structure. This yields strong efficiency for targeted workloads, especially where determinism, low latency, or custom operators are required.

Each of these architectures provides the parallelism, memory locality, and computational density that neural networks demand—capabilities that general-purpose CPUs are not architected to supply at competitive performance or energy efficiency.

**General-purpose CPUs** are built around control-oriented execution with deep, out-of-order pipelines, large caches, and relatively narrow SIMD units. This architecture excels at branchy, latency-sensitive workloads, but it is not efficient for the dense linear algebra at the core of neural network inference. Neural networks demand extremely high multiply–accumulate throughput and predictable data movement; CPUs simply cannot issue enough parallel MACs per cycle, nor can they feed those MACs from memory without hitting bandwidth limits. As a result, CPU-only execution quickly becomes bottlenecked by both compute density and dataflow inefficiency.

Medium-sized LLMs—anything beyond a few million parameters—are well outside the practical reach of today's MCUs. The limiting factors are absolute memory capacity, memory bandwidth, and sustained MAC throughput. Even high-end MCU-class devices rarely exceed a few megabytes of on-chip SRAM, and most lack external DRAM interfaces entirely. A "medium" LLM (tens to hundreds of millions of parameters) requires tens to hundreds of megabytes of weights even in 4-bit form, along with substantial activation storage. No shipping MCU can meet those requirements.

**MCUs and Tiny ML**

TinyML on microcontrollers and full AI accelerator SoCs represent two fundamentally different implementation points for AI inference, each optimized for distinct performance, power, and cost envelopes. While both execute neural networks, they operate under very different architectural constraints and serve different roles within the broader AI compute landscape.

MCU-based AI systems are designed for extreme power efficiency and cost sensitivity. These devices typically operate within tight energy budgets—often milliwatts or less—and rely on limited on-chip memory and modest compute resources. As a result, the neural networks deployed in this class are highly constrained, typically sub-100 kB models using aggressive quantization and simplified architectures. These implementations enable always-on sensing, anomaly detection, keyword spotting, and basic classification tasks at the edge, often directly integrated into sensor pipelines.

In contrast, AI accelerator SoCs target substantially more complex workloads, including transformer-based models, multimodal inference, and high-throughput vision processing. These systems incorporate wide vector units, systolic arrays, or other specialized compute fabrics, along with multi-level memory hierarchies and high-bandwidth external DRAM. They operate across a much broader power range—from a few watts in edge systems to hundreds of watts in data center deployments—and rely on sophisticated software stacks and runtime environments to manage heterogeneous compute resources.

This divergence reflects a structural bifurcation in the AI market rather than a single continuum. MCU-class TinyML enables distributed intelligence across billions of low-cost, power-constrained endpoints, while AI accelerator SoCs provide the compute density required for advanced inference and training workloads. Both are essential to the expansion of AI, but they address fundamentally different system requirements.

From a RISC-V perspective, this distinction is important. MCUs represent the largest unit-volume opportunity for RISC-V, forming the architectural foundation for widespread adoption across IoT, industrial, and consumer systems. These deployments may not execute large or complex AI models, but they are critical in enabling pervasive, low-level intelligence throughout the system landscape. In parallel, RISC-V is gaining traction in higher-performance roles within AI accelerator SoCs, where it is used for control, orchestration, and increasingly for domain-specific compute functions through vector and custom extensions.

Taken together, these trends position RISC-V across both ends of the AI spectrum: as the dominant architecture for low-power embedded intelligence and as an emerging component in more advanced AI compute platforms. The combination of high-volume MCU deployments and growing presence in accelerator-class systems reinforces its role as a scalable, heterogeneous compute architecture aligned with the evolving demands of AI-enabled systems.

A small but growing set of RISC-V–based MCUs can execute very small transformer-style models, though they remain in the same “ultra-small LLM” category described earlier—sub-million-parameter models with heavy quantization and carefully optimized kernels. None of these devices come close to supporting medium-scale LLMs, but several can run tiny attention blocks or low-depth sequence models.

**Figure 13: Alibaba DAMO Academy XuanTie Occupies a Prominent Position in the Chinese Market**


**Courtesy of Alibaba DAMO Academy**

Alibaba DAMO Academy is a leading RISC-V processor IP provider in the Chinese market, with a broad portfolio spanning MCU-class cores through application-class and higher-performance processors. Alibaba DAMO Academy's XuanTie RISC-V CPU IP cores are licensed extensively throughout the Chinese market, and they have engaged with over 400 ecosystem partners, largely serving cost-sensitive and high-volume-markets. XuanTie's offerings combine CPU IP across a wide performance range with associated software, tools, and configurable platforms, supporting deployments in edge intelligence and consumer electronics, high-performance and networking, servers and peripherals, and automotive electronics. Recent product introductions, including the C930, C920, C908X, and E901 series, reflect continued investment in scalable and higher-performance RISC-V platforms.

At the low end of the performance spectrum, XuanTie's MCU-class E-series processors have achieved broad adoption in IoT-oriented applications, with particular traction in high-demand segments such as power management and Bluetooth-enabled devices. This underscores the role of RISC-V IP in highly cost-optimized, high-volume semiconductor markets. As workloads evolve, XuanTie has also expanded into AI-adjacent domains, where RISC-V is increasingly used as a control or co-processor architecture alongside dedicated AI accelerators. Industry sources indicate that XuanTie holds a leading position in China's RISC-V IP market for AI accelerator control architectures.
Beyond traditional embedded markets, RISC-V designs are beginning to appear in emerging higher-performance applications, including cloud servers, video codec cards, laptops, and robotics platforms. While current shipment volumes in these segments remain modest compared with MCU and IoT deployments, market demand in these segments suggest growing commercial traction that is helping to establish new RISC-V footholds in the data center and compute-intensive domains. As a brand under Alibaba DAMO Academy, XuanTie benefits from alignment with Alibaba's capabilities in cloud computing, artificial intelligence, and data infrastructure, while maintaining a focus on building an open RISC-V ecosystem and advancing practical, deployable computing platforms across a wide range of industries.

## AI Accelerators

The following tables and figures show the metrics for AI Accelerators for M unit shipments and M Dollars.

**Table 7: AI Accelerator B Dollars by Industry Segment 2022 - 2031**

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | $0.07 | $0.15 | $0.31 | $0.53 | $0.78 | $1.05 | $1.44 | $2.00 | $2.66 | $3.64 | 38.0% |
| **Automotive** | $0.02 | $0.13 | $0.24 | $0.64 | $1.09 | $1.65 | $2.44 | $3.48 | $4.98 | $6.42 | 46.8% |
| **Networking** | $0.18 | $0.53 | $1.07 | $2.02 | $3.24 | $4.75 | $6.51 | $8.74 | $11.61 | $15.45 | 40.3% |
| **Computer** | $8.99 | $15.70 | $33.98 | $66.78 | $81.26 | $94.49 | $104.64 | $116.97 | $134.64 | $148.53 | 14.3% |
| **Consumer** | $0.54 | $1.64 | $4.47 | $7.50 | $11.06 | $14.85 | $19.45 | $24.58 | $33.20 | $42.62 | 33.6% |
| **Other** | $0.05 | $0.17 | $0.26 | $0.44 | $0.73 | $1.09 | $1.68 | $2.56 | $3.60 | $5.63 | 52.8% |
| **Total** | $9.84 | $18.32 | $40.32 | $77.91 | $98.16 | $117.88 | $136.14 | $158.33 | $190.69 | $222.29 | 19.1% |
| **Percent Growth** | 21.5% | 86.1% | 120.1% | 93.2% | 26.0% | 20.1% | 15.5% | 16.3% | 20.4% | 16.6% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

In the Total RISC-V market for Accelerators, the Computer segment is forecast to reach the highest revenues out of all the device types analyzed by The SHD Group. In 2025, the device revenues were $66,783M dollars and are forecast to grow to $148.534M dollars by 2031. This represents a CAGR of 14.3%.

In the Total RISC-V market for Accelerators, the total device revenues were $77,913M dollars in 2025 and are forecast to grow to $222,292M dollars by 2031. This represents a CAGR of 19.1%.

### Figure 14: RISC-V AI Accelerator Metrics B Dollars by Industry 2021 - 2031

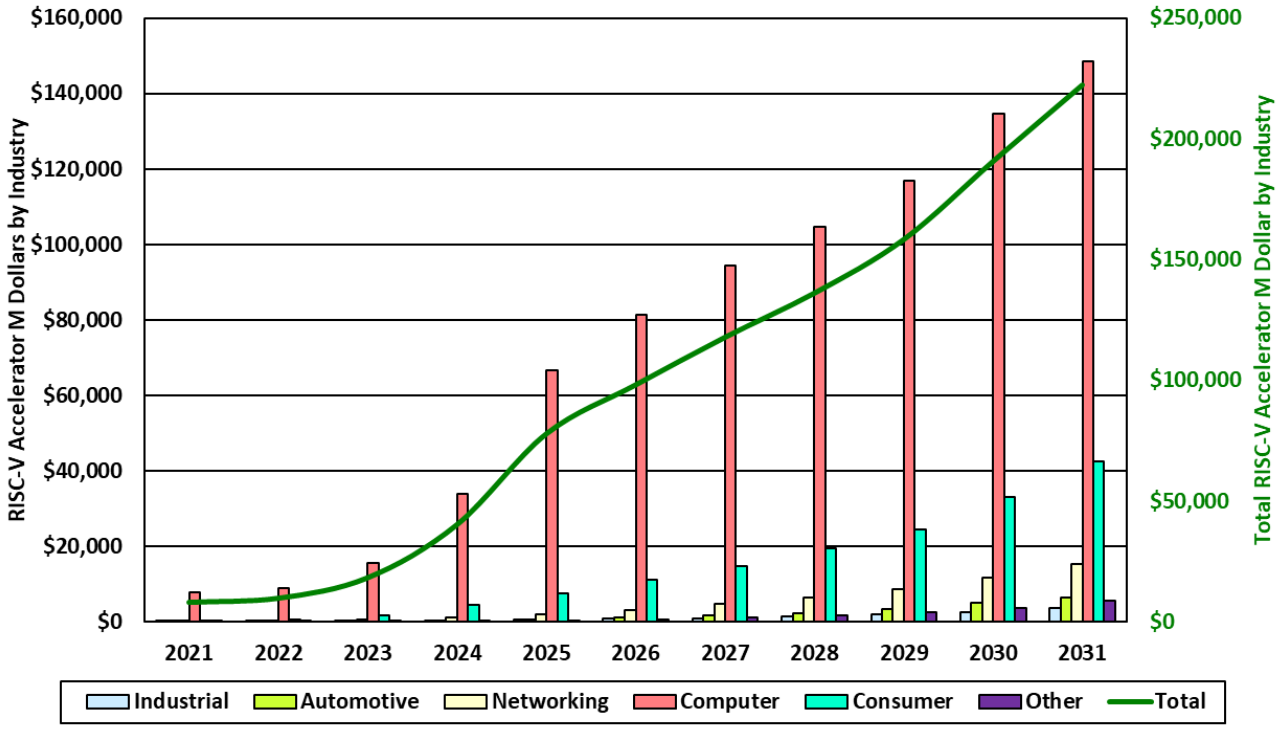


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

### Table 8: AI Accelerator M Units by Industry 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 9.7 | 22.4 | 44.6 | 81.8 | 130.8 | 198.0 | 291.5 | 438.1 | 636.2 | 873.1 | 48.4% |
| **Automotive** | 0.6 | 7.8 | 13.8 | 33.3 | 54.6 | 79.5 | 113.8 | 160.5 | 223.2 | 280.6 | 42.7% |
| **Networking** | 5.7 | 21.9 | 46.9 | 86.5 | 141.6 | 208.1 | 288.9 | 395.4 | 550.3 | 740.8 | 43.0% |
| **Computer** | 47.3 | 90.0 | 158.9 | 278.4 | 452.3 | 662.9 | 927.2 | 1,263.0 | 1,750.1 | 2,150.2 | 40.6% |
| **Consumer** | 65.5 | 161.6 | 385.0 | 653.4 | 1,009.4 | 1,418.6 | 1,914.1 | 2,511.1 | 3,376.6 | 4,587.6 | 38.4% |
| **Other** | 0.1 | 10.7 | 15.9 | 25.9 | 41.6 | 62.1 | 95.1 | 161.2 | 232.3 | 426.9 | 59.6% |
| **Total** | 128.8 | 314.3 | 665.1 | 1,159.3 | 1,830.2 | 2,629.2 | 3,630.6 | 4,929.2 | 6,768.7 | 9,059.3 | 40.9% |
| **Percent Growth** | 87.4% | 144.0% | 111.6% | 74.3% | 57.9% | 43.7% | 38.1% | 35.8% | 37.3% | 33.8% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

In the Total RISC-V market for Accelerators, the Consumer segment is forecast to reach the highest unit shipments out of all the device types analyzed by The SHD Group. In 2025, the unit shipments were 653.4M units and are forecast to grow to 4,588M units by 2031. This represents a CAGR of 38.4%.

In the Total RISC-V market for Accelerators, the total unit shipments were 1,159.3M units in 2025 and are forecast to grow to 9,059M units by 2031. This represents a CAGR of 40.9%.

**Figure 15: AI Accelerator Metrics M Units by Industry 2021 - 2031**

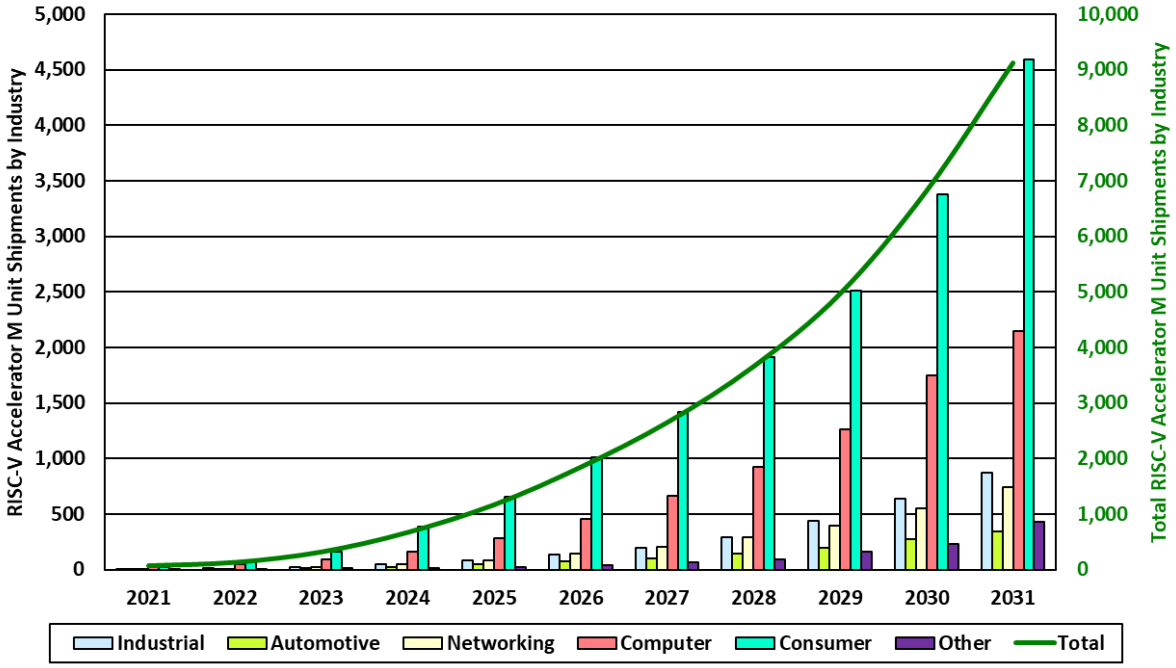


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

## Edge AI and RISC-V

The SHD Group released a study, **Edge-AI Market Analysis: Applications, Processors and Ecosystem Guide**, on the Edge AI market in April 2025. In this study we looked at which devices systems' developers were using to accelerate different types of the most popular Neural Networks, and surveyed 40+ companies to develop a profile of usage by device type and the different parameters developers use to select the right silicon for the application

The key question for device manufacturers, SIP companies and software developers today is which neural networks are going to be the most in demand for the various applications Edge-AI solutions are created for. Most of the systems where inference for Edge-AI is being performed are either low-cost or have power constraints. While an Edge-AI silicon architecture can be created to be able to run all the neural networks in the market today, such an architecture would probably be more expensive and power hungry than what most Edge-AI applications can tolerate. A more focused approach is necessary to meet the requirements of the target applications. This limits how much performance and capability can be built into the silicon for a given range of applications.

The Table 9 shows the most popular Neural Networks being used by device type for Industrial applications.

Table 9: Neural Network Usage by Device Type for the Industrial Market

| Neural Network | CPUs | GPUs | NPUs | DSPs | FPGAs | MCUs | Custom ASICs/SoCs |
|---|---|---|---|---|---|---|---|
| **Encoder-Decoder Architectures** | Moderate (small models) | High (video editing, NLP) | High (speech, video) | Moderate (lightweight) | High (media processing) | Low to Moderate | High (optimized NLP and video) |
| **Convolutional Neural Networks (CNNs)** | Moderate (small models) | High (image processing, gaming) | High (optimized models) | Moderate (lightweight) | High (customized tasks) | Low | High (optimized for tasks) |
| **Recurrent Neural Networks (RNNs)** | Moderate (small RNNs) | Moderate (sequential data) | High (text/audio tasks) | Moderate (lightweight) | Moderate (optimized) | Low | High (real-time inference) |
| **Spiking-Neural Networks (SNNs)** | Low | Low | High (neuromorphic tasks) | Moderate | High | Low | High (specialized tasks) |
| **Transformers** | Low (simplified models) | High (NLP, AI assistants) | High (large-scale NLP) | Low | Moderate (simplified) | Low | High (speech and fusion tasks) |
| **Generative Adversarial Networks (GANs)** | Low (inference only) | High (art/media generation) | Moderate (light inference) | Low | Moderate (specialized) | Low | High (image/video synthesis) |
| **Diffusion Models** | Low | High (creative tools) | Moderate (optimized) | Low | Moderate (simplified) | Very Low | High (synthetic data/image generation) |

**Source: The SHD Group, April 2025**

**High:** Neural Network optimized for use by this device type
**Moderate:** Small or moderately optimized versions of the NN for use by this device type
**Low:** low performance in this device type
**Very Low:** Poor performance in this device type

Table 10 below shows what features are important to systems developers in the silicon selection process.

System developers would typically look for the best solution to build their system around by looking at several parameters that are important for the functioning of the system, are important to their customers in their applications, are within the expertise of the system company and that are cost-competitive within the target market.

The following table lists some of the parameters that were noted by our survey respondents. It is not an all-inclusive list, but gives some perspective on how companies choose between different approaches and solutions

As this table shows, the requirements shift between different device types based on their individual strengths and weaknesses and the demands of the specific markets the solution is geared toward. Given the range of applications that have emerged in the Edge-AI market, there can be more than one device type that can be applied to the system to perform the necessary functions demanded by the market.

Table 10: Different Parameters for Choosing a Silicon Solution by Device Type

| Feature | CPU | GPU | NPU | DSP | MCU | FPGA | Custom ASIC/SoC |
|---|---|---|---|---|---|---|---|
| **Best For** | General-purpose computing, some AI tasks | AI model training, high-performance AI workloads | AI inference, low-power AI applications | Signal processing, real-time AI computations | Simple embedded AI applications | AI acceleration, reconfigurable AI workloads | Highly optimized AI applications, mass production |
| **Cost** | Low to Medium (Widely available) | High (AI-optimized GPUs are expensive) | Variable (Low for embedded NPUs, high for enterprise-class NPUs) | Low to Medium (Efficient, but not as widely used for AI) | Low (Very affordable for simple AI tasks) | High (Development cost is high, but mass production can be cost-effective) | Very High (Significant upfront cost, but low per-unit cost at scale) |
| **Power Efficiency** | Low (Consumes high power for AI tasks) | Medium (Power-hungry, but efficient for training) | High (Optimized for low-power AI inference) | High (Designed for efficient real-time signal processing) | Very High (Designed for low-power embedded systems) | Medium (Can be power-efficient depending on design) | Very High (Optimized for specific workloads) |
| **AI Inference Performance** | Slow (Software-optimized AI inference) | High (Accelerated inference but power-hungry) | Very High (Designed specifically for inference) | High (Optimized for real-time AI computations) | Low (Can run small AI models but very slow) | Variable (Depends on FPGA implementation) | Very High (Designed for AI inference at maximum efficiency) |
| **AI Training Performance** | Poor (Slow for deep learning models) | Excellent (Best for training large neural networks) | Limited (Mostly designed for inference, not training) | Low (Not suitable for AI training) | Poor (Not designed for AI training) | Moderate (Can be programmed for AI, but not ideal for large-scale training) | High (Custom-designed for AI training in some cases) |
| **Parallel Processing** | Low (Few cores, optimized for serial tasks) | High (Thousands of cores for parallel tasks) | Very High (Optimized for tensor operations) | Medium (Parallel SIMD operations for DSP) | Low (Single or few cores) | High (Custom parallelism based on design) | Very High (AI-optimized, custom-designed for workload) |
| **Memory Bandwidth** | Low (Uses DDR4/DDR5 RAM) | High (Uses HBM, GDDR6 for faster data access) | Medium to High (On-chip SRAM or LPDDR) | Medium (Optimized memory access for DSP workloads) | Low (Embedded Flash/RAM) | Variable (Depends on FPGA architecture) | Very High (Optimized memory for AI models) |
| **Ease of Programming** | High (Well-known architecture, optimized for software development) | Moderate (Requires CUDA, OpenCL, or AI SDKs) | Low (Requires specialized SDKs and AI frameworks) | Moderate (DSP programming knowledge required) | High (Simple software programming) | Low to Moderate (Requires expertise in FPGA programming) | Very Low (Fixed-function, not reprogrammable, requires hardware expertise) |
| **Software Ecosystem** | Excellent (Broad compatibility with AI frameworks) | Excellent (CUDA, ROCm, TensorRT, AI frameworks) | Limited (Vendor-specific frameworks, e.g., Edge TPU, Intel Movidius) | Moderate (DSP-optimized AI libraries) | Limited (MCU AI frameworks exist but limited) | Low to Moderate (Custom programming required) | Vendor-Specific (Highly optimized but fixed) |

**Source: The SHD Group, April 2025**

# VI. SoC Market Analysis

## Market Challenges and Issues

The global semiconductor industry continues to balance strong structural demand with ongoing regional and supply-chain challenges. Several market dynamics are shaping the near-term outlook:

- Rising design costs and chiplet standardization. The surge in SoC design complexity has accelerated adoption of chiplet-based architectures as mainstream development practice. Companies are investing heavily in advanced packaging, die-to-die interconnect IP, and verification tools to contain cost and schedule risk.
- Selective wafer-capacity constraints. While overall manufacturing capacity has improved since the supply disruptions of 2021–2023, tightness persists at the most advanced nodes—particularly 3 nm and below—limiting availability for AI accelerators and high-performance computing products.
- Smartphone market stabilization. After two years of contraction, smartphone shipments are stabilizing in 2025, supported by renewed growth in China and strong consumer interest in AI-enhanced and 5G-Advanced devices. Volume remains below pre-pandemic peaks but is trending positive.
- PC market recovery driven by AI integration. The PC sector is emerging from its post-pandemic trough, with shipments supported by next-generation AI PCs that integrate on-device neural accelerators. Desktop sales remain weak, but notebook and hybrid form factors are rebounding.
- Automotive semiconductor expansion. The sector continues to experience localized parts shortages, yet semiconductor content per vehicle is rising sharply as EV platforms and ADAS features proliferate. Average semiconductor value per car increased more than 12 % year-on-year in early 2025.
- Evolving trade and technology controls. Expanded export restrictions by the U.S. Department of Commerce and allied nations are prompting realignment of supply chains and regional design ecosystems, particularly in China. These shifts are influencing foundry investments and IP licensing models across Asia.
- Geopolitical and logistical risks. Continued regional instability and disruptions in global shipping routes remain risk factors for the broader electronics supply chain, though most companies now maintain diversified sourcing and greater inventory buffers than in prior cycles.

The SHD Group maintains a positive outlook for the semiconductor market in 2025 and beyond. Structural demand from AI infrastructure, automotive electrification, industrial automation, and the deployment of new wireless and edge-compute architectures outweigh cyclical or regional headwinds. Provided that geopolitical risks remain contained, energy costs remain stable, and inflationary pressures moderate, the industry is positioned for sustained—if uneven—growth through the latter half of the decade.

Table 11: List of Selected Applications by Segment

| | |
|---|---|
| **Industrial** | Robotics (Industrial) |
| | Smart Grid |
| | IIoT (Factory Floor) |
| | Automatic Test Equip |
| | |
| **Automotive** | ADAS |
| | Commercial Vehicles |
| | Entry-level Passenger Cars |
| | Mid-range Passenger Cars |
| | High-end Passenger Cars |
| | |
| **Communications** | 5G |
| | Central Office Switches |
| | High-end Routers |
| | Mid-range Routers |
| | Low-end Routers |
| | High-end Hubs |
| | Mid-range Hubs |
| | Low-end Hubs |
| | High-end Switches |
| | Mid-range Switches |
| | Low-end Switches |
| | ATM Switches |
| | Cable Modems |
| | DSL Modems |
| | 4G / LTE Picocell Base Stations |
| | 4G / LTE Femtocell Base Stations |
| | 6G Infrastructure |
| | |

| | |
|---|---|
| **Computing** | Computing - Datacenter / Cloud |
| | Desktop |
| | Notebook |
| | Tablet |
| | Workstations $5K - $15K |
| | Workstations > $15K |
| | Solid State Drives – Enterprise |
| | Edge Computing |
| | |
| **Consumer** | IOT |
| | Wearables |
| | AR / VR |
| | Security Cameras |
| | Cell phone - Apps Processor |
| | Cell phone - Baseband Processor |
| | Cell phone - Camera |
| | Cell phone - AI Acceleration |
| | Handheld Game Consoles |
| | HDTV |
| | Set Top Box |
| | Digital Still Camera |
| | Digital Video Camera |
| | Drones & Controllers |
| | GPS |
| | Robotic Home Appliances |
| | Sensors + RISC-V |
| | Smart Watches |
| | Smart Speakers |
| | Solid State Drives - Consumer |
| | Streaming Media Devices |
| | Video Game Consoles |
| **Other** | |

**Source: The SHD Group, April 2026**

The SoC market has grown to fill virtually every possible application niche in the semiconductor market today. This report analyzes selected applications in the Industrial, Automotive, Computer, Consumer, and Networking categories where RISC-V is thought to play a meaningful role in the very near future.

Table 12: List of Selected SoC Product Types by Segment

| SoC Device Types | |
|---|---|
| CPU - General Purpose | FPGA |
| App Processor | DSP |
| GPU | Network Device |
| RISC-V Smart Sensors | Security |
| MCU | Storage Controller |
| AI Accelerator | |

Source: The SHD Group, April 2026

# Assumptions for RISC-V Market Analysis

## General Assumptions

1. The **RISC-V ecosystem** has reached maturity, supported by full integration across major EDA vendors, compiler toolchains, and operating-system environments. Ecosystem growth continues through expanded verification IP, middleware, and domain-specific software stacks.
2. SoC designers now have demonstrated confidence that RISC-V adoption brings access to robust tools, software, and IP libraries comparable to legacy ISAs.
3. RISC-V cores have now become a standard component in many heterogeneous SoCs, routinely sharing die area with Arm, x86, or custom DSP cores.
4. RISC-V's momentum is driven by two factors: rapid adoption in MCU-class devices, establishing volume scale, and its extensible ISA, which allows vendors to implement semi-custom cores with domain-specific instructions for AI and signal-processing acceleration in higher-performance AI SoCs.
5. As higher-performance RISC-V implementations emerge, they are increasingly being assigned primary compute roles in AI SoCs, Edge-AI SoCs, and communications SoCs.
6. RISC-V continues to support a diverse range of functions—from deeply embedded state-machine control to MCU, co-processor, and full application-processor workloads—reflecting its role as a truly heterogeneous compute architecture.

## General CPU SoCs

1. RISC-V CPU IP vendors are releasing successive generations of high-performance cores, with support for vector, DSP, and AI extensions now common.
2. Each new core class expands RISC-V's addressable market toward data-centric and compute-intensive applications traditionally dominated by Arm or x86.
3. Arm remains the leading CPU IP vendor by revenue and installed base, but RISC-V is achieving measurable share gains across embedded and edge compute designs.
4. Adoption is steady to accelerating, with the technology now viewed as a credible and cost-effective alternative for new CPU SoC programs.

## Applications Processor SoCs

1. Chinese smartphone and consumer-electronics manufacturers - including Huawei, Oppo, and others—are now shipping SoCs containing RISC-V cores, replacing Western IP where feasible.
2. Western OEMs are moving more cautiously but are incorporating RISC-V cores in subsystems such as AI accelerators and security controllers within upcoming devices.

## GPU SoCs

1. Following NVIDIA's 2024 disclosure, RISC-V cores are now embedded across all NVIDIA product lines, typically serving as control, management, and telemetry processors.
2. Other GPU vendors are expected to follow, and RISC-V penetration in GPU architectures is projected to rise steadily as firmware, scheduling, and security functions migrate to open-ISA cores.

## Security SoCs

1. Dedicated security processors and controllers using RISC-V are now well-established, addressing the inadequacy of purely software-based protection.
2. Adoption is accelerating in PCs, servers, networking gear, and automotive platforms, where hardware root-of-trust and real-time cryptographic acceleration are mandatory.
3. RISC-V's open design and extensibility make it an ideal platform for national and enterprise security standards development.

## AI Accelerator

1. AI functionality has become ubiquitous, spanning data-center inference, networking, automotive, and industrial automation.
2. RISC-V cores are now routinely employed as management, control, and sometimes compute tiles within AI accelerators, leveraging vector and matrix extensions for local data processing.
3. This domain remains one of the **fastest-**growing market opportunities for RISC-V over the forecast horizon.

## FPGAs

1. Most leading FPGA vendors—including AMD/Xilinx, Altera (now independent from Intel), Lattice, Microchip, QuickLogic, Menta, Achronix, and ADI-owned Flex Logix—now provide RISC-V soft cores or support RISC-V integration into their programmable fabrics.
2. Heterogeneous architectures remain dominant, but RISC-V has become the default embedded CPU option for new FPGA platforms.
3. Full RISC-V-only product families are expected later in the decade as tool and software ecosystems converge.
4. Legacy-system inertia continues to slow migration in long-lifecycle designs that rely on historic ISAs and software stacks.

## DSP SoCs

1. Early adoption of RISC-V for DSP and signal-processing functions is now underway, driven by new vector-extension support and the need for configurable compute engines in AI-enabled systems.
2. Penetration remains moderate but is projected to grow rapidly toward the latter half of the forecast period as traditional DSP vendors modernize their architectures.

## Smart RISC-V Sensors

1. This is an expanding category in which RISC-V microcontrollers are integrated directly with sensors for **local** inference, filtering, and communication.
2. Volumes are increasing sharply across IoT, industrial automation, and automotive applications, though the category's massive scale requires keeping penetration ratios conservative in revenue models.

## Network Device SoCs

1. Networking SoCs - including Ethernet, SerDes, and wireless chipsets—are undergoing continuous redesign cycles driven by new protocol generations (PCIe 6.0, Wi-Fi 7, 800 Gb Ethernet).
2. Each redesign cycle introduces new opportunities for RISC-V adoption, primarily in control, packet-management, and embedded CPU subsystems.

## Storage Controller SoCs

1. With escalating storage densities and interface speeds, RISC-V is increasingly selected for low-power control and protocol-handling tasks in SSDs and storage controllers.
2. This segment remains one of the most promising growth areas through 2031, driven by both data-center and edge-storage expansion.

## MCUs

1. MCUs remain the largest unit-volume category for RISC-V. Their widespread adoption across embedded and industrial systems has established a solid base for the architecture.
2. As higher-performance cores are adopted into complex SoCs with higher ASPs, MCU revenue growth will moderate, but volume dominance will persist.
3. RISC-V MCUs are now foundational across IoT, industrial, and consumer segments, forming the entry point for many new adopters of the architecture.

Table 13: Market Revenues for all RISC-V SoCs by Application 2022–2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **B Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | $0.187 | $0.414 | $0.795 | $1.277 | $1.788 | $2.387 | $3.175 | $4.435 | $5.885 | $8.393 | 36.9% |
| **Automotive** | $0.190 | $0.609 | $1.096 | $1.848 | $2.698 | $3.680 | $4.974 | $6.815 | $9.511 | $12.373 | 37.3% |
| **Networking** | $0.202 | $0.998 | $2.082 | $3.731 | $5.600 | $7.835 | $10.382 | $13.892 | $18.540 | $25.265 | 37.5% |
| **Computer** | $9.083 | $16.917 | $36.441 | $70.946 | $87.297 | $102.222 | $114.360 | $130.490 | $155.225 | $175.130 | 16.3% |
| **Consumer** | $1.812 | $5.184 | $10.627 | $15.427 | $21.064 | $27.203 | $34.513 | $44.647 | $61.324 | $81.177 | 31.9% |
| **Other** | $0.122 | $0.357 | $0.584 | $1.079 | $1.860 | $2.914 | $4.443 | $6.781 | $9.455 | $15.853 | 56.5% |
| **Total** | $11.597 | $24.478 | $51.625 | $94.308 | $120.308 | $146.241 | $171.847 | $207.059 | $259.941 | $318.192 | 22.5% |
| **Market Penetration** | 3.0% | 4.8% | 8.4% | 12.6% | 14.9% | 17.3% | 19.8% | 22.9% | 27.4% | 33.1% | |
| **Percent Growth** | 30.2% | 111.1% | 110.9% | 82.7% | 27.6% | 21.6% | 17.5% | 20.5% | 25.5% | 22.4% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Total RISC-V SoC market revenues are projected to reach $94.3B in 2025, a growth of 82.7% over 2024 and is forecast to grow to $318.2B by 2031, a CAGR of 22.5%. The RISC-V SoC market, in general, is characterized by the adoption of RISC-V CPU cores into many already-existing SoCs, so the ecosystem needed to use these devices is already built out and aiding in a fast adoption rate. It is also important to remember that all of these markets are coming from very small base years, resulting in very high CAGRs.

- The Computer segment is the largest category, projected to reach $70.9B in 2025, a growth of 94.7% over 2024. It is forecast to grow to $175.1B by 2031, a CAGR of 16.3% over the forecast period of the identified applications.
    - Note that a substantial portion of this revenue is attributable to the RISC-V-enabled GPUs that Nvidia is shipping into this market.
    - Also note that given the acceleration of infrastructure buildout for AI data centers, this forecast may prove to be on the low side.

- The Consumer segment is the second largest category at $15.4B in 2025, a growth of 45.2% over 2024. It is forecast to reach $81.2B by 2031, a CAGR of 31.9%, and driven by all cell phones, 4K-8K UHDTVs, game consoles, and other high-value consumer systems.
- The Networking market represents the 3rd largest category and reached $3.7B in 2025, growing 79.2% over 2024. It is forecast to reach $25.3B by 2031, driven by networking switches and infrastructure build-out, with a CAGR of 37.5%.
- The Automotive segment is undergoing a complete transformation as ADAS, and other AI functionality is being added along with the electrification of vehicles in general. This market was $1.84B in 2025 and grew 68.6% over 2024. It is forecast to reach $12.3B by 2031, a CAGR of 37.3%
- The Other category is the 4th largest market and has the highest CAGR and is projected to reach $15.9B by 2031, a CAGR of 56.5%.
- The Industrial segment is the smallest market, reaching $1.3B in 2025, a growth of 60.5% over 2024 and is forecast to reach $8.4B by 2031, a CAGR of 36.9%.

**Figure 16: Market Revenues for All RISC-V SoCs by Application 2021–2031**

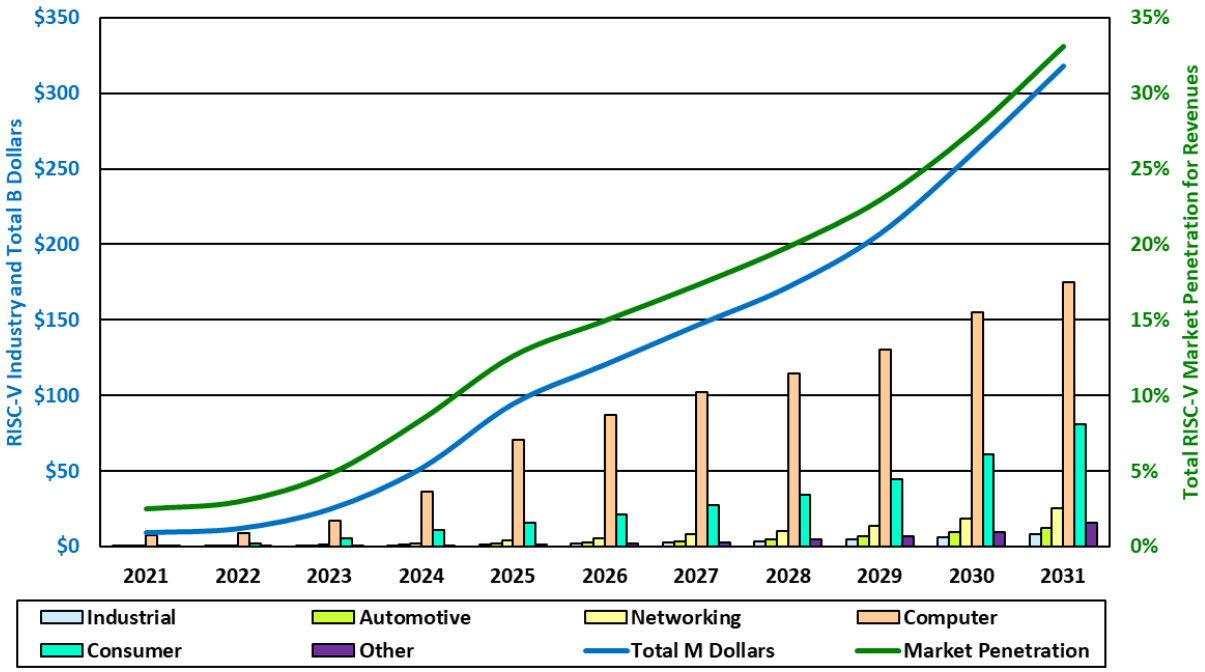


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 14: Market Unit Shipments for all RISC-V SoCs by Application 2022–2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 124.8 | 224.9 | 432.1 | 620.7 | 808.1 | 1,042.4 | 1,256.0 | 1,769.8 | 2,370.3 | 3,230.6 | 31.6% |
| **Automotive** | 51.9 | 127.6 | 240.9 | 357.6 | 463.8 | 631.8 | 778.6 | 1,034.6 | 1,370.9 | 1,814.8 | 31.1% |
| **Networking** | 27.9 | 101.0 | 188.4 | 317.2 | 469.3 | 663.8 | 882.5 | 1,234.4 | 1,743.2 | 2,473.8 | 40.8% |
| **Computer** | 119.8 | 292.2 | 505.6 | 817.6 | 1,226.3 | 1,727.0 | 2,305.3 | 3,289.7 | 4,665.7 | 5,650.3 | 38.0% |
| **Consumer** | 1,174.1 | 2,080.1 | 3,482.5 | 4,474.9 | 5,660.1 | 6,773.7 | 7,932.8 | 10,565.9 | 14,370.3 | 18,847.4 | 27.1% |
| **Other** | 68.5 | 124.2 | 176.2 | 300.8 | 487.6 | 773.2 | 1,168.1 | 1,808.3 | 2,502.3 | 3,955.3 | 53.6% |
| **Total** | 1,566.9 | 2,950.1 | 5,025.7 | 6,888.9 | 9,115.2 | 11,611.8 | 14,323.3 | 19,702.8 | 27,022.7 | 35,972.2 | 31.7% |
| **Market Penetration** | 4.4% | 7.4% | 10.8% | 13.0% | 15.0% | 16.9% | 18.6% | 22.1% | 26.9% | 31.7% | |
| **Percent Growth** | 98.7% | 88.3% | 70.4% | 37.1% | 32.3% | 27.4% | 23.4% | 37.6% | 37.2% | 33.1% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Total RISC-V SoC market unit shipments are projected to reach 6.889B units in 2025, a growth of 37.1% over 2024 and is forecast to grow to 35.9B units by 2031, a CAGR of 31.7%. The SoC market in general is characterized by rising complexity levels driven, in part, by the introduction of AI and continuing increases in connectivity and computing performance requirements. It is also important to remember that all these markets are coming from very small base years, resulting in very high CAGRs.

- The Consumer segment is the largest category at 4.5B units in 2025, a growth of 28.5% over 2024. It is forecast to reach 18.8B units by 2031, a CAGR of 27.1%, and driven by all cell phones, UHDTVs, wearables, Edge AI and other high-volume consumer systems.
- The Computer segment is the 2nd highest category, reaching 0.8B units in 2025, a growth of 61.7% over 2024. It is forecast to grow to 5.7B units by 2031, a CAGR of 38.0% over the forecast period of the identified applications.
- The Industrial segment is the 3rd largest market, reaching 0.621B units in 2025, a growth of 43.7% over 2024 and is forecast to reach 3.2B units by 2031, a CAGR of 31.6%.
- The Automotive segment is undergoing a complete transformation as ADAS, and other AI functionality is being added along with the electrification of vehicles in general. This market was 0.358B units in 2025 and grew 48.5% over 2024 as a great deal of new functionality is being added to vehicles of all types. It is forecast to reach 1.815B units by 2031, a CAGR of 31.1%
- The Networking market represents the 5th largest category and reached 0.317B units in 2025, growing 68.3% over 2024. It is forecast to reach 2.5B units by 2031, driven by 5G infrastructure and Femtocell Base Stations.
- The Other category has the largest growth rate and is projected to reach 3.955B units by 2031, a CAGR of 53.6%.

Figure 17: Market Unit Shipments for All RISC-V SoCs by Application 2021–2031

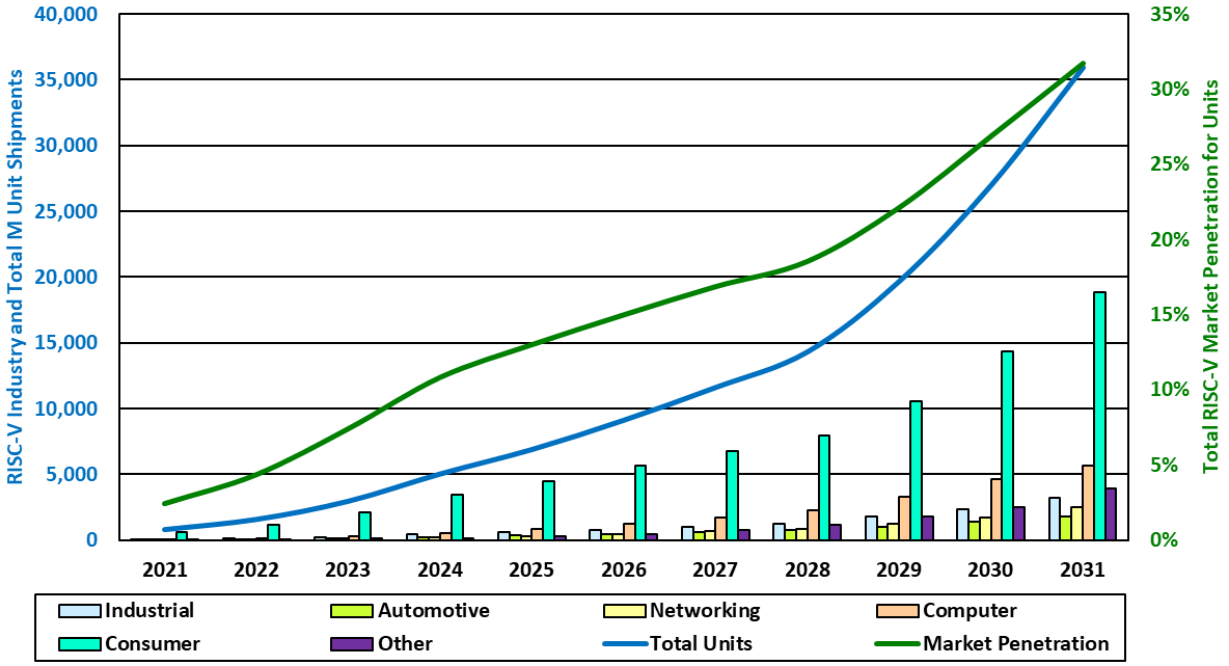


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

As the RISC-V ecosystem enters 2026, a steadily expanding base of IP vendors continues to reinforce industry confidence in the open ISA. The growing participation of established EDA and semiconductor IP suppliers signals that RISC-V has moved beyond its early adoption phase and is now regarded as a mainstream design option across multiple markets and performance tiers.

A key development over the past year has been the integration of RISC-V specifications into broader, non-RISC-V–centric design environments, such as EDA platforms, chiplet fabrics, and cross-architecture SoCs. This trend highlights RISC-V's appeal as a flexible, open standard capable of coexisting with proprietary ISAs. Its modular structure allows designers to embed RISC-V cores or instruction extensions into otherwise heterogeneous systems, improving both design efficiency and feature differentiation.

Configurability remains one of RISC-V's defining advantages. The ability to tailor instruction sets and microarchitectures for specific workloads is driving a wave of domain-optimized CPU designs spanning from deeply embedded controllers to AI and data-centric compute. This adaptability not only supports the move toward heterogeneous SoC architectures but also enables faster innovation cycles, as designers mix and match RISC-V cores with accelerators, memory subsystems, and connectivity IP to achieve application-specific performance and power targets.

Figure 18: MosChip Technologies Provides Multiple ASIC Design Services


Source: MosChip Technologies Ltd.

MosChip Technologies provides end-to-end silicon and product engineering services, spanning chip architecture, hardware design, embedded software, system-level design and verification. MosChip Technologies has the engineering resources and expertise to provide physical design, DFT (design-for-test), RTL (register-transfer-level) implementation and verification services for high-performance RISC-V cores. Their capabilities extend into mixed-signal (analog + digital) ASIC design, tape-out experience (200+ multi-million-gate ASICs) and product engineering in areas such as IoT, automotive, industrial and edge computing. For the RISC-V market specifically, MosChip acts as a trusted design services partner helping RISC-V core/IP developers and system-on-chip teams accelerate development, manage complexity and bring products to silicon. In addition, MosChip provides a range of IP porting, customization and integration services to the RISC-V community. MosChip is positioned as a service provider enabling companies leveraging RISC-V architecture to move from concept to silicon and system delivery, supplying design, verification, and product engineering support tailored for RISC-V cores, SoCs and subsystems.

## Total RISC-V SoC Market Metrics and Analysis by Market and Device Type

This section deals with the revenues and shipments for RISC-V-powered SoCs by device type and end application category.

**Please Note: Many of the CAGRs for the individual part types and the markets they are targeted for in this section can be very high, driven by the fact that they are all coming from very small base years.**

Table 15: Total RISC-V SoC Market Revenues for All Applications 2022–2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **B Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **CPU-General Purpose** | $0.042 | $0.364 | $1.288 | $2.865 | $4.266 | $5.346 | $6.876 | $9.746 | $16.412 | $23.141 | 41.6% |
| **App Processor** | $0.148 | $0.538 | $2.335 | $4.059 | $5.771 | $7.605 | $9.925 | $12.384 | $17.707 | $23.220 | 33.7% |
| **GPU** | $8.592 | $14.733 | $32.103 | $63.438 | $75.511 | $85.827 | $92.549 | $100.341 | $110.881 | $120.166 | 11.2% |
| **Security** | $0.000 | $0.046 | $0.090 | $0.259 | $0.496 | $0.843 | $1.534 | $3.450 | $6.679 | $10.059 | 84.1% |
| **Perception Engine** | $0.069 | $0.108 | $0.172 | $0.252 | $0.365 | $0.487 | $0.630 | $0.874 | $1.140 | $1.415 | 33.3% |
| **AI Accelerator** | $1.101 | $2.634 | $5.131 | $9.309 | $15.246 | $22.374 | $31.036 | $42.537 | $58.266 | $74.819 | 41.5% |
| **FPGA** | $0.004 | $0.217 | $0.452 | $0.751 | $1.138 | $1.477 | $1.960 | $2.459 | $2.990 | $3.811 | 31.1% |
| **DSP** | $0.000 | $0.206 | $0.357 | $0.529 | $0.809 | $1.056 | $1.352 | $1.652 | $2.226 | $3.115 | 34.4% |
| **Smart RISC-V Sensors** | $1.126 | $1.740 | $2.735 | $3.266 | $3.837 | $3.947 | $4.190 | $4.552 | $5.391 | $6.755 | 12.9% |
| **Network Device** | $0.000 | $2.638 | $4.970 | $6.872 | $8.827 | $11.569 | $14.339 | $18.055 | $22.777 | $31.673 | 29.0% |
| **Storage Controller** | $0.042 | $0.257 | $0.498 | $0.720 | $1.165 | $1.552 | $1.958 | $2.536 | $3.075 | $4.291 | 34.7% |
| **MCU** | $0.472 | $0.996 | $1.495 | $1.990 | $2.879 | $4.157 | $5.497 | $8.471 | $12.398 | $15.728 | 41.1% |
| **Total B Dollars** | $11.597 | $24.478 | $51.625 | $94.308 | $120.308 | $146.241 | $171.847 | $207.059 | $259.941 | $318.192 | 22.5% |
| **Percent Growth** | 30.2% | 111.1% | 110.9% | 82.7% | 27.6% | 21.6% | 17.5% | 20.5% | 25.5% | 22.4% | |
| **Market Penetration** | 3.0% | 4.8% | 8.4% | 12.6% | 14.9% | 17.3% | 19.9% | 22.9% | 27.4% | 31.1% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

The total market for RISC-V SoCs was $51.6B in 2024 and is forecast to reach $94.3B in 2025, a growth of 82.7% over 2024 and is projected to reach $318.2B by 2031, a CAGR of 22.5%.

- The total market penetration of RISC-V SoCs in 2024 was 8.4% and is forecast to reach 12.6% in 2025. By 2031, the market penetration is projected to reach 31.1%.

Within the RISC-V SoC market, the top three device types are:

- GPUs are the largest category of SoCs in 2024 reaching $32.1B, a growth of 117.9% over 2023 and are forecast to reach $63.4B in 2025. By 2031, GPUs are projected to grow to reach $120.2B, a CAGR of 11.2%
  - Of note is much of the growth of GPUs is due to the dominance of Nvidia and their use of RISC-V CPUs in this market.
- AI Accelerators were the second largest type of RISC-V SoC at $5.13B in 2024 and are forecast to reach $2.6B in 2023, a growth of 81.4%. By 2031, this part type is projected to reach $74.8B, a CAGR of 41.5%.
- The third largest SoC category is Network Devices which reached $4.97B in 2024 and are forecast to grow to $6.87B in 2025, a growth rate of 38.3%. In 2031, this part type is projected to reach $31.7B, a CAGR of 29.0%.

Figure 19: Total RISC-V SoC Market Revenues for all Applications 2021-2031

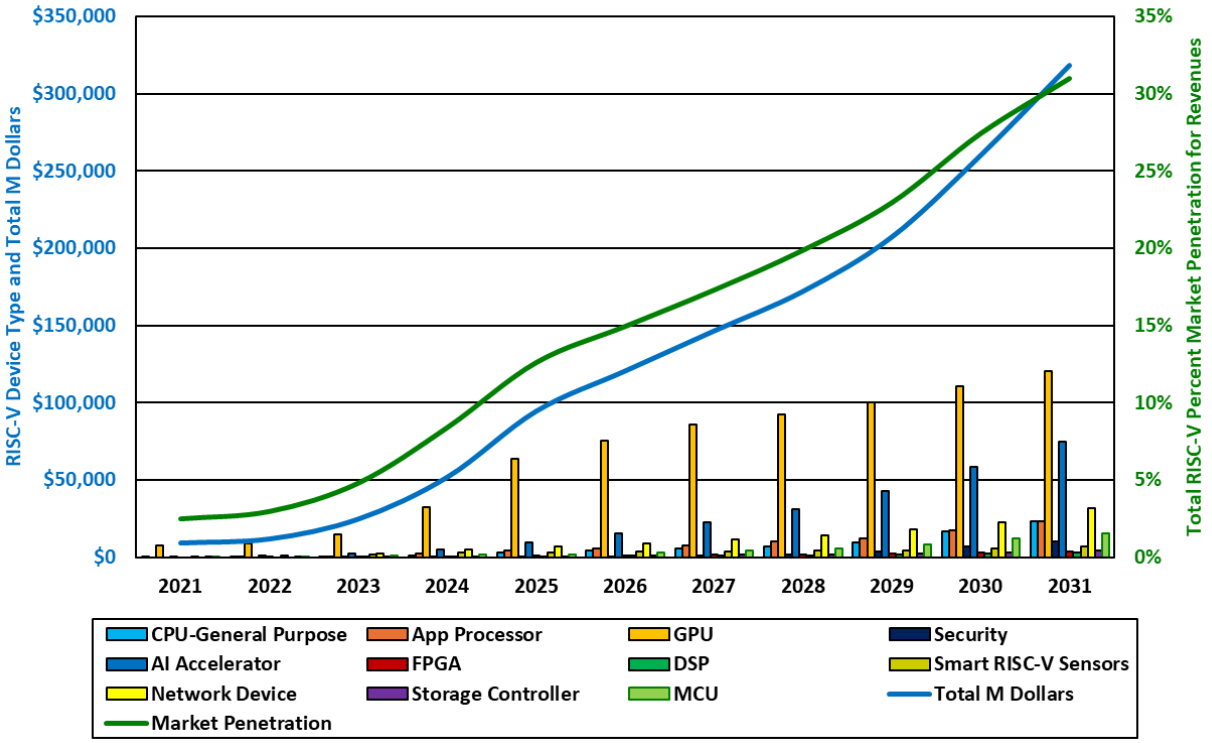


| Forecast Years: 2025 - 2031 | Source: The SHD Group, April 2026 |
|---|---|

Figure 20: RISC-V Verification Tools Provided by Breker Verification Systems

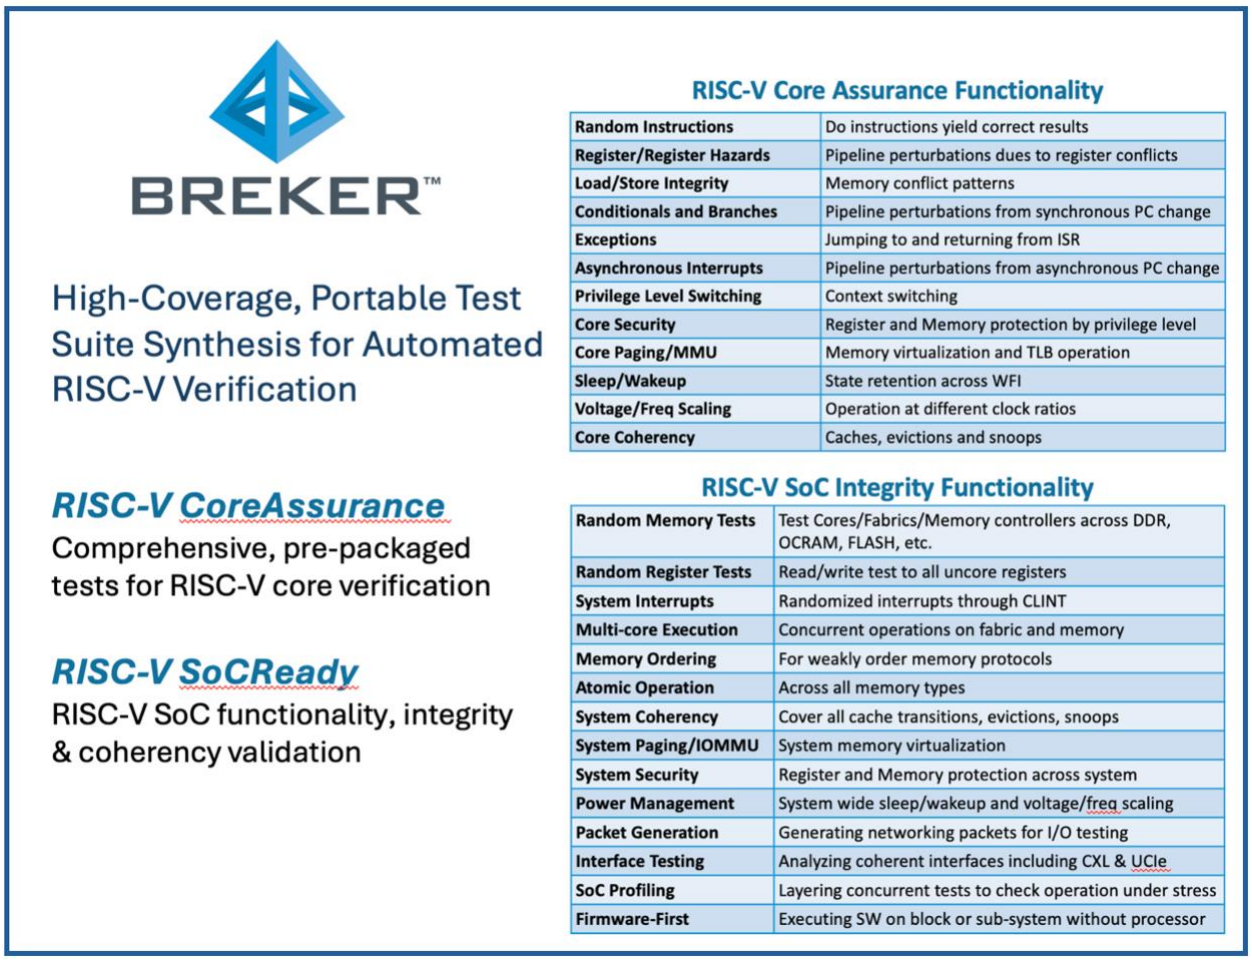


**RISC-V Core Assurance Functionality**

| | |
|---|---|
| **Random Instructions** | Do instructions yield correct results |
| **Register/Register Hazards** | Pipeline perturbations dues to register conflicts |
| **Load/Store Integrity** | Memory conflict patterns |
| **Conditionals and Branches** | Pipeline perturbations from synchronous PC change |
| **Exceptions** | Jumping to and returning from ISR |
| **Asynchronous Interrupts** | Pipeline perturbations from asynchronous PC change |
| **Privilege Level Switching** | Context switching |
| **Core Security** | Register and Memory protection by privilege level |
| **Core Paging/MMU** | Memory virtualization and TLB operation |
| **Sleep/Wakeup** | State retention across WFI |
| **Voltage/Freq Scaling** | Operation at different clock ratios |
| **Core Coherency** | Caches, evictions and snoops |

**RISC-V SoC Integrity Functionality**

| | |
|---|---|
| **Random Memory Tests** | Test Cores/Fabrics/Memory controllers across DDR, OCRAM, FLASH, etc. |
| **Random Register Tests** | Read/write test to all uncore registers |
| **System Interrupts** | Randomized interrupts through CLINT |
| **Multi-core Execution** | Concurrent operations on fabric and memory |
| **Memory Ordering** | For weakly order memory protocols |
| **Atomic Operation** | Across all memory types |
| **System Coherency** | Cover all cache transitions, evictions, snoops |
| **System Paging/IOMMU** | System memory virtualization |
| **System Security** | Register and Memory protection across system |
| **Power Management** | System wide sleep/wakeup and voltage/freq scaling |
| **Packet Generation** | Generating networking packets for I/O testing |
| **Interface Testing** | Analyzing coherent interfaces including CXL & UCIe |
| **SoC Profiling** | Layering concurrent tests to check operation under stress |
| **Firmware-First** | Executing SW on block or sub-system without processor |

**Source: Breker Verification Systems**

Breker's role in the RISC-V ecosystem centers on providing automated, high-coverage verification for both RISC-V cores and full SoCs. Its RISC-V CoreAssurance offering supplies pre-assembled test suites that stress every major aspect of a RISC-V processor, including instruction accuracy, pipeline hazards, branching behavior, exception handling, interrupt timing, privilege transitions, MMU operation, power-state changes and coherency behavior. These tests are generated through Breker's portable test-suite synthesis technology, which is designed to expose corner-case failures and ensure that RISC-V cores behave correctly under realistic and extreme operating conditions.

At the SoC level, Breker's RISC-V SoCReady package extends this methodology to full-system validation, covering memory-controller stress testing, randomized system interrupts, multi-core execution scenarios, atomic operations, security checks, coherency validation, virtualization, power-management transitions, I/O packet generation, interface-protocol behavior (including CXL and UCIe), and firmware-first operation. By delivering these system-level verification assets in a portable form that runs across simulation, emulation, prototyping and post-silicon platforms, Breker helps RISC-V adopters accelerate debug, reduce verification risk, and validate complex SoC implementations built around standard or customized RISC-V cores.

Table 16: Total RISC-V SoC Market Unit Shipments for All Applications 2022–2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **CPU-General Purpose** | 1.4 | 10.5 | 35.8 | 78.5 | 118.7 | 156.2 | 206.7 | 301.9 | 488.8 | 699.0 | 44.0% |
| **App Processor** | 5.4 | 21.4 | 100.6 | 180.7 | 267.1 | 364.4 | 496.6 | 653.0 | 960.7 | 1,369.0 | 40.1% |
| **GPU** | 26.3 | 26.1 | 31.3 | 43.8 | 54.7 | 66.2 | 78.8 | 95.4 | 113.7 | 130.1 | 19.9% |
| **Security** | 0.0 | 10.8 | 21.9 | 57.8 | 112.6 | 192.8 | 343.2 | 749.3 | 1,423.4 | 2,088.0 | 81.8% |
| **Perception Engine** | 3.0 | 4.8 | 7.6 | 11.2 | 16.3 | 21.8 | 28.3 | 39.3 | 51.4 | 63.9 | 33.6% |
| **AI Accelerator** | 96.7 | 224.7 | 451.3 | 802.4 | 1,302.9 | 1,916.4 | 2,670.5 | 3,691.9 | 5,037.5 | 6,764.0 | 42.7% |
| **FPGA** | 0.5 | 25.5 | 54.4 | 96.6 | 150.2 | 209.4 | 293.8 | 388.0 | 503.3 | 647.9 | 37.3% |
| **DSP** | 0.0 | 17.4 | 32.2 | 49.9 | 81.4 | 111.4 | 149.4 | 192.8 | 276.4 | 402.8 | 41.6% |
| **Smart RISC-V Sensors** | 1,069.0 | 1,615.1 | 2,684.5 | 3,216.7 | 3,675.2 | 3,795.6 | 3,692.0 | 4,425.9 | 5,347.5 | 6,649.1 | 12.9% |
| **Network Device** | 0.0 | 350.2 | 654.1 | 971.7 | 1,319.0 | 1,837.2 | 2,421.3 | 3,215.3 | 4,271.6 | 6,026.9 | 35.5% |
| **Storage Controller** | 10.6 | 49.4 | 101.3 | 170.5 | 252.0 | 349.2 | 452.5 | 608.0 | 761.7 | 1,062.0 | 35.6% |
| **MCU** | 354.1 | 594.1 | 850.8 | 1,209.0 | 1,785.2 | 2,591.2 | 3,490.3 | 5,341.9 | 7,786.7 | 10,069.4 | 42.4% |
| **Total M Units** | 1,566.9 | 2,950.1 | 5,025.7 | 6,888.9 | 9,135.2 | 11,611.8 | 14,323.3 | 19,702.8 | 27,022.7 | 35,972.2 | 31.7% |
| **Percent Growth** | 98.7% | 88.3% | 70.4% | 37.1% | 32.6% | 27.1% | 23.4% | 37.6% | 37.2% | 33.1% | |
| **Market Penetration** | 4.4% | 7.4% | 10.8% | 13.0% | 15.0% | 16.9% | 18.6% | 22.1% | 26.9% | 31.6% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

The total market for RISC-V SoCs is estimated at 5,025M units in 2024 and is forecast to reach 6,889M units in 2025, a growth of 37.1% over 2024 and is projected to reach 35.972B units by 2031, a CAGR of 31.7%.

- The market penetration of RISC-V SoCs in 2024 was 11.2% and is forecast to reach 13.5% in 2025. By 2031, the market penetration is projected to reach 31.6%.

Within the Other category, the top three device types are:

- Smart RISC-V Sensors were 2,685M units in 2024 and are forecast to reach 3,217M units in 2025, a growth of 19.8%. By 2031, this part type is projected to reach 6,649B units, a CAGR of 12.9%.
- The second largest SoC category is MCU SoCs which reached 850.8M units in 2024 and are forecast to grow to 1,209M units in 2025, a growth rate of 42.1%. In 2031, this part type is projected to reach 10,069M units, a CAGR of 42.4%.
- Network Devices are the third largest part type and had shipments of 654M units in 2024 and are forecast to grow to 972M units in 2025. They are forecast to reach 6,027M units by 2031, a CAGR of 35.5% over the forecast period of the identified applications.

Figure 21: Total RISC-V SoC Market Unit Shipments for all Applications 2021-2031

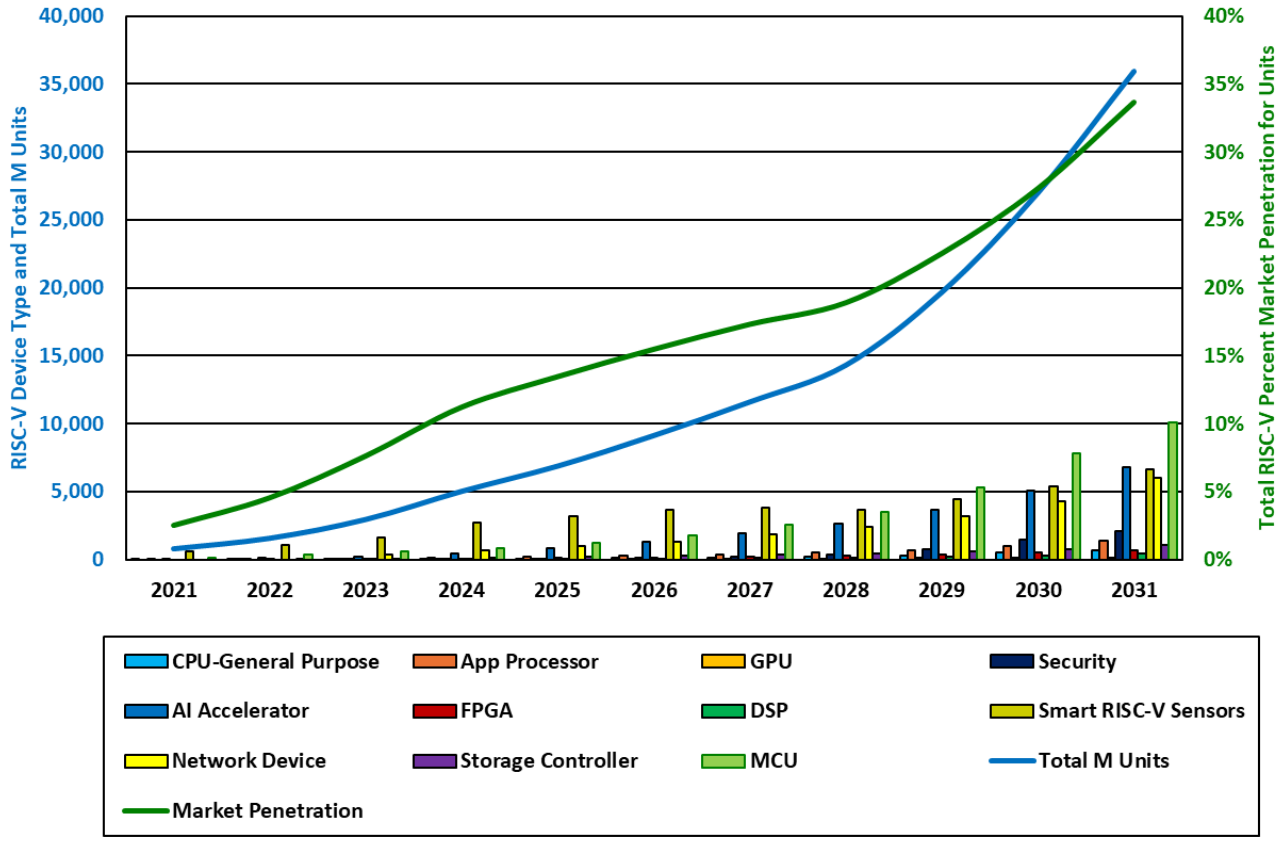


Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 17: Penetration of RISC-V SoCs by Application by Revenue 2022–2031

| Percentage | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 |
|---|---|---|---|---|---|---|---|---|---|---|
| Industrial | 2.9% | 4.9% | 8.0% | 10.1% | 12.0% | 13.8% | 16.0% | 18.3% | 21.7% | 26.0% |
| Automotive | 0.4% | 1.9% | 3.1% | 5.1% | 6.8% | 8.6% | 10.9% | 13.7% | 18.3% | 21.3% |
| Networking | 0.5% | 2.1% | 4.2% | 6.6% | 9.2% | 11.9% | 14.9% | 18.4% | 22.8% | 28.1% |
| Computer | 4.9% | 5.9% | 10.1% | 15.1% | 17.5% | 19.9% | 22.5% | 25.7% | 30.0% | 32.2% |
| Consumer | 3.8% | 7.0% | 8.8% | 9.7% | 11.5% | 14.4% | 16.6% | 21.6% | 28.2% | 32.1% |
| Other | 0.8% | 1.9% | 2.8% | 4.5% | 6.2% | 8.3% | 10.7% | 14.1% | 17.5% | 24.2% |
| Market Penetration | 3.0% | 4.8% | 8.4% | 12.6% | 14.9% | 17.3% | 19.8% | 22.9% | 27.4% | 33.1% |

Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 17 shows the penetration for RISC-V-enabled SoCs by revenue into the six application categories profiled in the report. As the table shows, RISC-V is forecast to show a 33.1% penetration into the SoC market by 2031.

The highest penetration will be in Computer and Consumer devices, followed by Networking, Industrial, Other and Automotive.

The Computer application sees the highest penetration due to the large number of SSD drives shipping into the market and also from the switch-over of general purpose SoCs in computer systems to using RISC-V-based SoCs.

Figure 22: Penetration of RISC-V SoCs by Application by Revenue 2021-2031

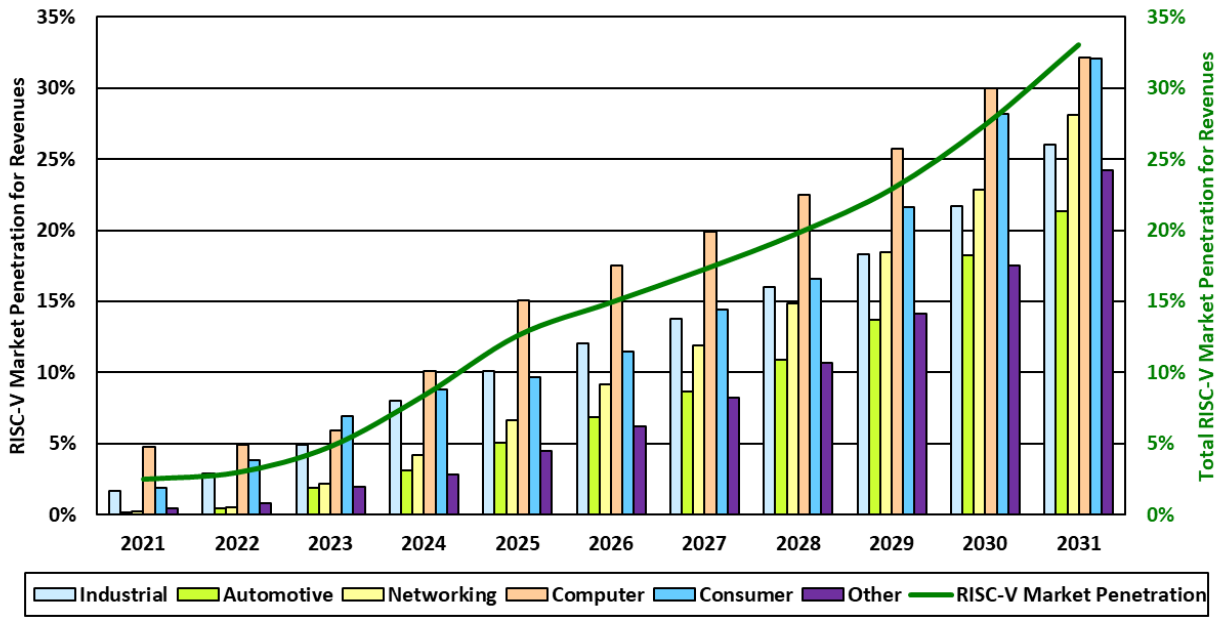


Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 18: Penetration of RISC-V SoCs by Application by Unit Shipments 2022–2031

| Percentage | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 |
|---|---|---|---|---|---|---|---|---|---|---|
| Industrial | 5.4% | 8.3% | 13.4% | 16.0% | 17.6% | 18.9% | 19.2% | 22.4% | 26.0% | 30.0% |
| Automotive | 1.5% | 4.3% | 6.9% | 10.6% | 14.3% | 18.0% | 22.1% | 27.0% | 33.9% | 38.7% |
| Networking | 0.9% | 2.9% | 4.7% | 6.9% | 9.1% | 11.5% | 13.9% | 17.1% | 21.6% | 27.4% |
| Computer | 1.6% | 3.6% | 5.9% | 8.5% | 11.5% | 14.7% | 18.0% | 22.7% | 29.1% | 33.0% |
| Consumer | 5.1% | 8.0% | 9.6% | 10.9% | 12.9% | 15.9% | 17.9% | 22.7% | 29.2% | 33.5% |
| Other | 2.2% | 3.4% | 4.4% | 6.2% | 8.1% | 10.6% | 13.2% | 16.9% | 20.3% | 26.1% |
| Market Penetration | 4.4% | 7.4% | 10.8% | 13.0% | 15.0% | 16.9% | 18.6% | 22.1% | 26.9% | 31.6% |

Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 18 shows the penetration for RISC-V-enabled SoCs by unit volume into the six application categories profiled in the report. As the table shows, RISC-V is forecast to show a 31.6% penetration into the SoC market by 2031.

The highest market penetration for RISC-V-enabled SoCs will be in Automotive and Consumer devices, followed by Computer, Industrial, Networking and Other.

RISC-V-enabled SoC in Automotive have the highest market penetration as vehicle manufacturers continually add new functionality and ADAS features to offer the most appealing products to their customers. However, it is necessary to keep in mind that Automotive has the fewest number of SoCs in the category as compared to the other market categories we analyzed. Therefore, the market penetration as a percent of the total number of parts for Automotive is the highest while this category has the fewest total number of SoCs in it.

The Consumer application category sees the second highest penetration due to the large number of consumer SSD drives shipping into the market and also from the large number of Wearables and mobile systems shipping into this category.

**Figure 23: Penetration of RISC-V SoCs by Application by Units Shipments 2021-2031**

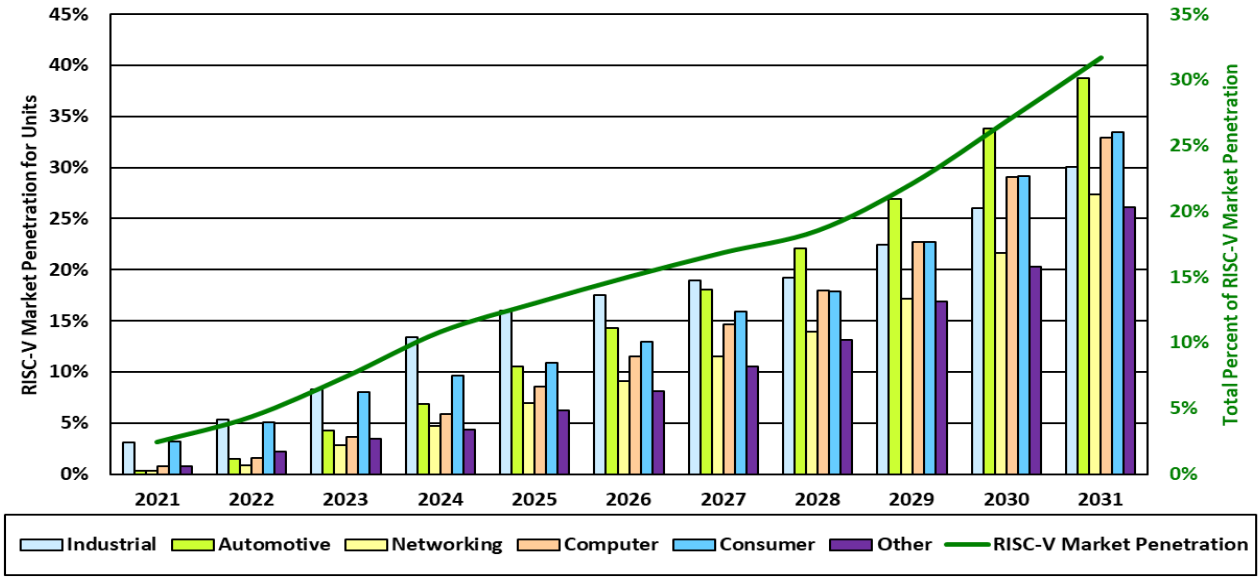


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

This next section looks at the penetration of RISC-V-enabled SoCs by the individual device types profiled in the report.

- Of note is the rise in AI accelerators using RISC-V as a design element. This is due to the interest in adopting and adding AI functionality into many existing and new applications to offer richer feature sets and usability to end users, making the products more attractive.
- Also, note that Perception Engines in Automotive using RISC-V have a high penetration as well. However, the total SoC units in Automotive are lower than for all other categories. Therefore, the penetration of RISC-V -enabled SoC is relatively higher when compared to the other categories.

Table 19: Penetration of RISC-V SoCs by Revenue by Device Type 2022–2031

| Percentage | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 |
|---|---|---|---|---|---|---|---|---|---|---|
| CPU - General Purpose | 0.1% | 0.7% | 2.3% | 4.3% | 5.7% | 6.3% | 7.4% | 9.4% | 14.4% | 18.8% |
| App Processor | 0.2% | 0.8% | 3.4% | 5.9% | 8.3% | 10.8% | 13.7% | 16.8% | 23.5% | 29.3% |
| GPU | 8.4% | 7.1% | 11.4% | 16.5% | 18.6% | 20.6% | 22.8% | 25.3% | 27.9% | 28.7% |
| Security | 0.0% | 0.6% | 0.8% | 1.7% | 2.6% | 3.7% | 5.7% | 9.5% | 14.8% | 18.9% |
| Perception Engine | 7.1% | 10.0% | 14.3% | 18.6% | 23.1% | 27.6% | 32.1% | 38.0% | 42.7% | 47.0% |
| AI Accelerator | 10.7% | 13.6% | 19.8% | 27.4% | 34.6% | 41.9% | 49.3% | 56.0% | 63.2% | 72.6% |
| FPGA | 0.1% | 3.2% | 6.3% | 9.7% | 13.1% | 16.0% | 19.9% | 23.4% | 26.6% | 30.0% |
| DSP | 0.0% | 1.0% | 1.6% | 2.2% | 3.3% | 4.1% | 5.1% | 5.9% | 7.9% | 10.0% |
| Smart RISC-V Sensors | 1.5% | 4.0% | 7.2% | 9.8% | 12.5% | 15.2% | 18.4% | 21.8% | 27.6% | 33.2% |
| Network Device | 0.0% | 3.4% | 5.0% | 6.9% | 8.5% | 10.6% | 12.5% | 14.8% | 17.7% | 20.2% |
| Storage Controller | 0.0% | 1.2% | 2.4% | 3.8% | 5.3% | 6.7% | 8.2% | 10.0% | 11.8% | 14.7% |
| MCU | 2.5% | 4.3% | 5.6% | 6.7% | 8.4% | 10.8% | 12.7% | 17.0% | 22.3% | 25.4% |
| Avg. Penetration | 3.0% | 4.8% | 8.4% | 12.6% | 14.9% | 17.3% | 19.8% | 22.9% | 27.4% | 33.1% |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 19 and Figure 24 show the penetration of RISC-V SoCs by revenue in the device types for all applications profiled in this report. These values are calculated by looking at the revenues for all SoCs and then totaling the revenues for SoCs with RISC-V CPU cores and dividing one group by the other to arrive at the penetration given in percent.

- Overall, The SHD Group projects that SoCs with RISC-V CPU cores will grow to penetrate 31.1% of all SoC revenues by 2031, powered by several different SoC types.
- AI Accelerators are forecast to reach the highest penetration out of all the device types The SHD Group analyzed. In 2024, the penetration was 19.8% and is forecast to grow to 27.4% in 2025. By 2031, we forecast the penetration of RISC-V into AI Accelerators to be 72.6%.
    - This is not hard to understand since AI functionality is being added to almost every application and SoC solution for those applications.
- The device type with the second highest penetration in 2024 was the Smart RISC-V Sensor, reaching 30.5% and forecast to grow to 31.5% in 2025. By 2031, we expect the penetration of RISC-V into the Smart Sensor market to reach 32.6%.
- There are two reasons for this. sensors are being deployed in every possible application. However, the revenues for this part type are low as compared to other RISC-V-enabled devices.
- Perception Engines with RISC-V CPU cores in Automotive have the 3rd highest penetration, with 14.3% in 2024, and are forecast to grow to 18.6% in 2025. By 2031, we expect RISC-V cores in MCUs to reach 47.0%.

Figure 24: Penetration of RISC-V SoCs by Revenue by Device Type 2021–2031

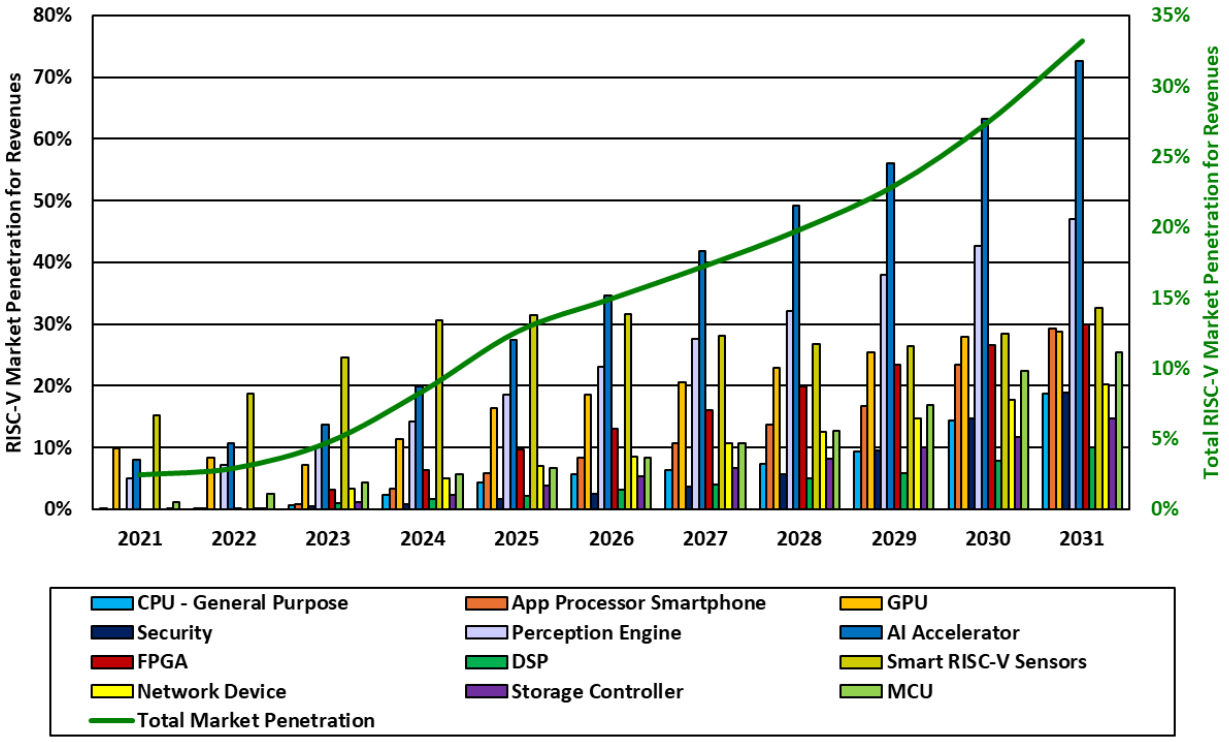


Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 20: Penetration of RISC-V SoCs by Unit Shipments by Device Type 2022–2031

| Percentage | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 |
|---|---|---|---|---|---|---|---|---|---|---|
| CPU - General Purpose | 0.1% | 0.4% | 1.3% | 2.7% | 3.8% | 4.7% | 5.8% | 7.8% | 11.9% | 15.4% |
| App Processor | 0.2% | 0.7% | 2.9% | 4.8% | 6.6% | 8.4% | 10.6% | 13.0% | 17.8% | 23.0% |
| GPU | 13.2% | 13.8% | 16.1% | 19.7% | 22.4% | 24.8% | 27.1% | 29.8% | 32.4% | 34.4% |
| Perception Engine | 0.0% | 0.7% | 1.0% | 1.8% | 2.8% | 4.0% | 6.0% | 9.5% | 14.5% | 18.2% |
| Security | 9.0% | 13.0% | 19.0% | 25.0% | 31.0% | 37.0% | 43.0% | 50.0% | 55.0% | 60.0% |
| AI Accelerator | 9.0% | 12.5% | 18.4% | 24.9% | 31.4% | 37.9% | 44.1% | 50.1% | 55.7% | 64.9% |
| FPGA | 0.1% | 2.9% | 5.6% | 8.8% | 11.9% | 14.6% | 18.3% | 21.5% | 24.7% | 27.3% |
| DSP | 0.0% | 0.7% | 1.2% | 1.7% | 2.6% | 3.3% | 4.1% | 4.8% | 6.5% | 8.5% |
| Smart RISC-V Sensors | 16.2% | 21.4% | 28.7% | 30.2% | 30.1% | 27.3% | 24.3% | 27.0% | 30.1% | 34.5% |
| Network Device | 0.0% | 3.0% | 4.6% | 6.3% | 7.8% | 9.9% | 11.9% | 14.3% | 17.2% | 20.1% |
| Storage Controller | 2.6% | 3.7% | 4.7% | 6.1% | 7.5% | 8.9% | 10.4% | 12.3% | 14.4% | 16.7% |
| MCU | 3.9% | 7.0% | 8.4% | 10.6% | 13.0% | 16.0% | 18.2% | 22.3% | 28.0% | 30.0% |
| Avg. Penetration | 4.4% | 7.4% | 10.8% | 13.0% | 15.0% | 16.9% | 18.6% | 22.1% | 26.9% | 31.6% |

Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 20 and Figure 25 show the penetration of RISC-V SoCs for units in the device types for all applications profiled in this report. These values are calculated by looking at the unit shipments for all SoCs and then totaling the unit shipments for SoCs with RISC-V CPU cores and dividing one group by the other to arrive at the penetration given in percent.

- Overall, The SHD Group projects that SoCs with RISC-V CPU cores will grow to penetrate 31.6% of all SoC unit shipments by 2031, powered by several different SoC types.
- AI Accelerators are forecast to reach the highest penetration out of all the device types The SHD Group analyzed. In 2024, the penetration was 18.4% and is forecast to grow to 24.9% in 2025. By 2031, we forecast the penetration of RISC-V into AI Accelerators to be 64.9%.
    - This is not hard to understand since AI functionality is being added to almost every application and SoC solution for those applications.
- The device type with the highest penetration in 2024 was the Smart RISC-V Sensor, reaching 28.7% and forecast to grow to 30.2% in 2025. By 2031, we expect the penetration of RISC-V into the Smart Sensor market to reach 34.5%.
- Perception Engines with RISC-V CPU cores have the second highest penetration, with 19.0% in 2024, and are forecast to grow to 25.0% in 2025. By 2031, we expect RISC-V cores in Perception Engines to reach 60.0%.

Figure 25: Penetration of RISC-V SoCs by Unit Shipments by Device Type 2021-2031

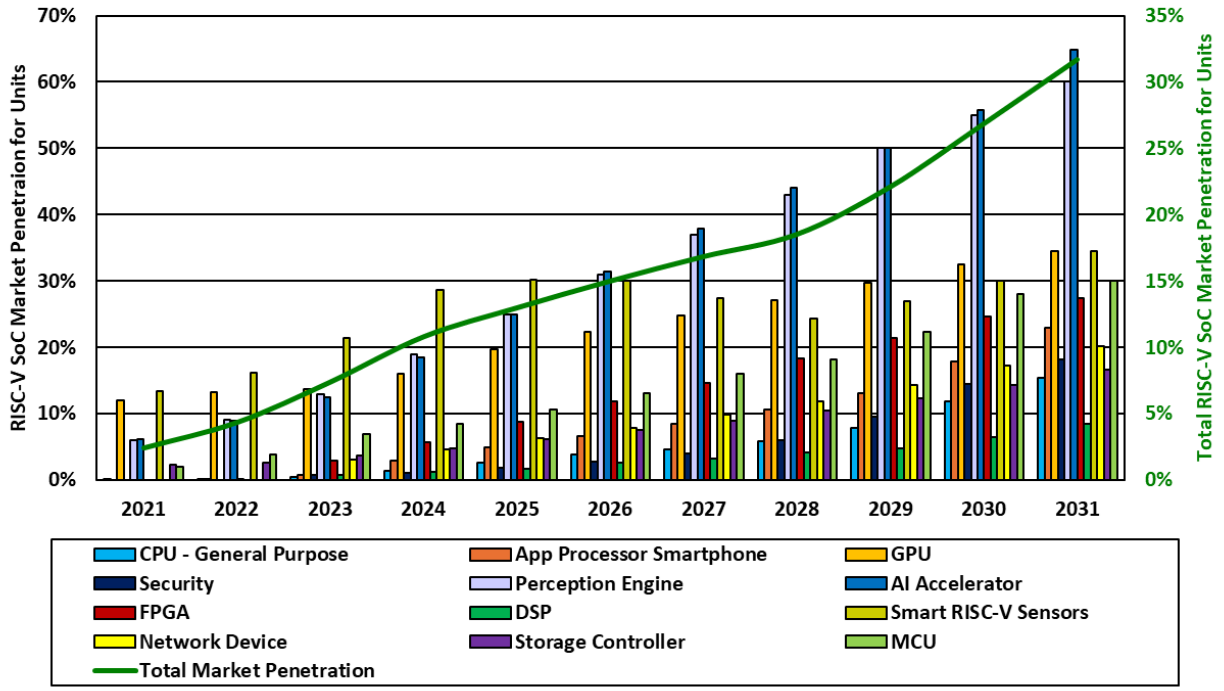


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

# VII. RISC-V SoCs by Functionality 2021 - 2031

The next section looks at what functions RISC-V CPU cores will be tasked with performing in RISC-V-powered SoCs. The functions that The SHD Group profiled for this market analysis are:

- Deeply embedded (e.g., Finite State Machine)
- MCU
- Co-processor
- Host Processor

These tables show the revenues and unit volumes for the entire SoC market by end application and the four RISC-V-enabled functions for revenues and unit volumes for the years 2021-2031.

By doing the analysis in this manner, we are no longer trying to place a device in a specific CPU-driven category, reasoning that most, if not all, SoCs going forward will have more than one type of CPU core from multiple vendors. RISC-V International already says that many high-end multicore SoCs can have up to 10-12 different ISAs, indicating the usage of multiple types of CPU cores today. These cores are likely performing the functions we have identified that are relevant to the market now and into the future.

The way to interpret these tables is to understand that the part types we have profiled for the report could have one or more of these functions' resident. So, the revenue numbers and the unit volumes are not additive but should be considered as stand-alone data points for each functional type.

## Analysis of RISC-V Usage by Functionality

Table 21: RISC-V SoC Finite State Machine Functionality Unit Volumes by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 8.5 | 14.8 | 27.6 | 45.0 | 69.0 | 100.2 | 172.0 | 218.8 | 310.3 | 448.8 | 46.7% |
| **Automotive** | 2.6 | 7.3 | 12.7 | 20.0 | 29.5 | 40.0 | 52.8 | 73.7 | 99.7 | 127.0 | 36.1% |
| **Networking** | 5.8 | 33.4 | 66.1 | 113.2 | 170.7 | 244.1 | 327.7 | 444.3 | 608.4 | 827.9 | 39.3% |
| **Computing** | 10.4 | 40.7 | 73.5 | 119.9 | 180.5 | 260.5 | 355.2 | 501.4 | 711.6 | 858.2 | 38.8% |
| **Consumer** | 48.8 | 188.8 | 352.0 | 543.4 | 798.4 | 1,117.8 | 1,576.8 | 2,006.4 | 2,695.2 | 3,666.8 | 37.5% |
| **Other** | 1.5 | 3.6 | 7.0 | 15.7 | 31.8 | 62.0 | 107.7 | 187.4 | 287.4 | 494.4 | 77.6% |
| **Total** | 77.5 | 288.6 | 538.9 | 857.3 | 1,279.9 | 1,824.5 | 2,592.1 | 3,432.0 | 4,712.6 | 6,423.1 | 39.9% |
| **Percent Growth** | 111.0% | 272.4% | 86.7% | 59.1% | 49.3% | 42.6% | 42.1% | 32.4% | 37.3% | 36.3% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Finite State Machines (FSM) are used on every SoC today. In contemporary SoCs, an FSM is a digital circuit designed to execute specific sequences of operations or control the behavior of various parts of the system based on a finite number of defined states. These states can represent different modes of operation or conditions within the chip.

FSMs in SoCs are crucial for managing the complex interactions between different hardware components, facilitating efficient control and coordination within the chip's architecture. They play a fundamental role

in controlling the system's behavior, ensuring proper functionality, and optimizing the performance of the SoC.

FSMs can range from the very simple to moderately complex depending on the SoC architecture they are designed into and the task they are performing in the silicon. In addition, there can be many instantiations of an FSM in SoCs today and can number in the hundreds on very complex parts.

In this analysis, we looked at what were likely candidates to use RISC-V-based FSMs. The numbers above show the use of RISC-V-based FSMs in the six major market categories profiled in this report.

SoCs for Consumer applications see the highest concentration of FSMs due to the very large unit volumes of this type of SoC. The Computer and Networking segments see the second and third highest usage of the FSM function.

Figure 26: RISC-V Finite State Machine Functionality Unit Volumes by Application Category, 2021-2031

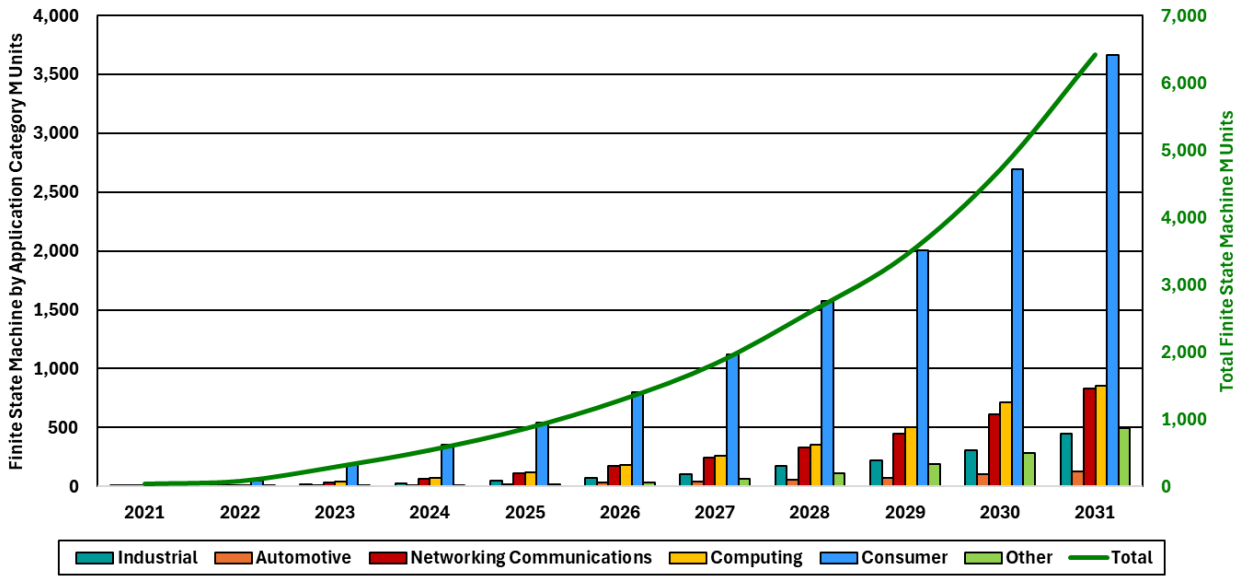


Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 22: RISC-V SoC Finite State Machine Functionality Revenues by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| M Dollars | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 | 25 - 31 |
| Industrial | $10.2 | $30.1 | $60.6 | $109.7 | $173.9 | $263.1 | $457.3 | $609.4 | $858.3 | $1,232.9 | 49.7% |
| Automotive | $54.8 | $138.2 | $238.4 | $369.4 | $535.7 | $719.3 | $938.5 | $1,279.1 | $1,682.6 | $2,193.6 | 34.6% |
| Networking | $65.0 | $405.4 | $832.9 | $1,443.0 | $2,208.8 | $3,130.8 | $4,186.1 | $5,547.6 | $7,478.4 | $10,058.2 | 38.2% |
| Computing | $179.6 | $584.1 | $1,051.8 | $1,660.2 | $2,492.7 | $3,572.0 | $4,828.8 | $6,612.9 | $9,419.6 | $11,655.1 | 38.4% |
| Consumer | $1,131.0 | $2,845.3 | $5,330.0 | $8,326.5 | $12,394.9 | $16,027.8 | $20,675.8 | $24,496.8 | $30,724.5 | $41,294.4 | 30.6% |
| Other | $2.9 | $11.4 | $25.9 | $61.1 | $128.2 | $242.2 | $418.9 | $715.1 | $1,101.7 | $2,060.2 | 79.7% |
| Total | $1,443.5 | $4,014.4 | $7,539.7 | $11,970.1 | $17,934.2 | $23,955.2 | $31,505.5 | $39,260.9 | $51,265.1 | $68,494.4 | 33.7% |
| Percent Growth | 335.9% | 178.1% | 87.8% | 58.8% | 49.8% | 33.6% | 31.5% | 24.6% | 30.6% | 33.6% | |

Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

For this analysis, we counted the device revenues of SoCs where the use of RISC-V-based FSMs occurred. We counted the entire revenue generated by shipment of the part into the markets, so we are not counting the IP revenue of the RISC-V-based FSM, but the device ASP times the units being shipped as the determining factor for counting the revenues. In this case, an SoC could have many FSMs on it and was counted the same as if it had only one FSM.

The Consumer segment generated the highest amount of revenue, followed by the Computer and Networking segments.

Figure 27: RISC-V SoC Finite State Machine Functionality Revenues by Application Category, 2021-2031

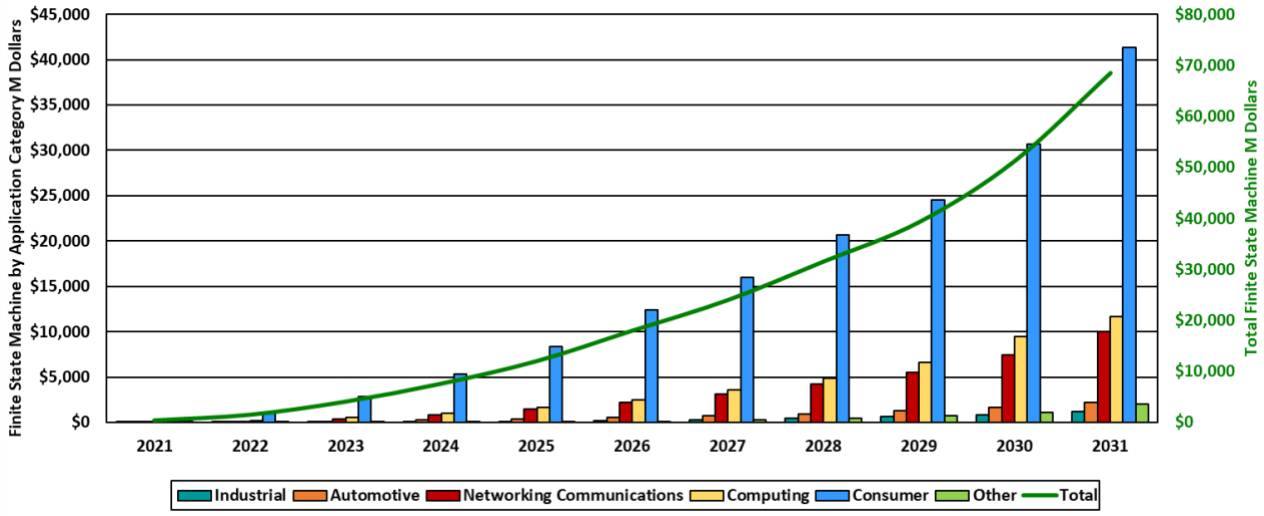


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 23: RISC-V SoC MCU Functionality Unit Volumes by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 13.2 | 23.3 | 43.1 | 71.0 | 108.8 | 159.3 | 228.3 | 350.1 | 504.8 | 709.1 | 46.7% |
| **Automotive** | 3.1 | 10.5 | 21.2 | 33.0 | 48.2 | 66.3 | 88.2 | 124.0 | 172.9 | 218.1 | 37.0% |
| **Networking** | 19.8 | 44.1 | 75.7 | 128.0 | 200.8 | 294.7 | 404.8 | 593.2 | 866.7 | 1,242.3 | 46.0% |
| **Computing** | 33.7 | 71.0 | 118.5 | 197.7 | 312.5 | 467.6 | 651.9 | 993.7 | 1,487.1 | 1,778.3 | 44.2% |
| **Consumer** | 225.6 | 450.2 | 731.1 | 1,038.6 | 1,535.8 | 2,212.7 | 2,958.0 | 4,375.1 | 6,356.5 | 8,440.7 | 41.8% |
| **Other** | 33.3 | 55.7 | 75.3 | 131.7 | 223.7 | 363.7 | 557.2 | 868.6 | 1,196.3 | 1,756.5 | 54.0% |
| **Total** | 328.8 | 654.7 | 1,064.9 | 1,600.1 | 2,429.8 | 3,564.3 | 4,888.3 | 7,304.8 | 10,584.2 | 14,144.9 | 43.8% |
| **Percent Growth** | 113.7% | 99.2% | 62.6% | 50.3% | 51.9% | 46.7% | 37.1% | 49.4% | 44.9% | 33.6% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

In the context of SoCs, an MCU function refers to the incorporation of a microcontroller unit (MCU) within the larger integrated circuit design. This involves integrating a microcontroller along with its associated components directly onto the same chip that houses other functionalities. This integration aims to

consolidate multiple features into a single chip, reducing the overall size, power consumption, and cost while enhancing efficiency and performance.

This integration is common in a wide range of devices, including IoT (Internet of Things) devices, smart appliances, automotive systems, industrial controllers, and many other embedded systems where compactness, low power consumption, and cost-effectiveness are essential.

The Consumer segment sees the largest use of RISC-V-based MCUs, followed by Computing and Networking. The other segment also has a large usage. In addition, the Automotive segment is expected to see a large and growing usage of RISC-V-based MCUs.

**Figure 28: RISC-V MCU Functionality Unit Volumes by Application Category, 2021-2031**

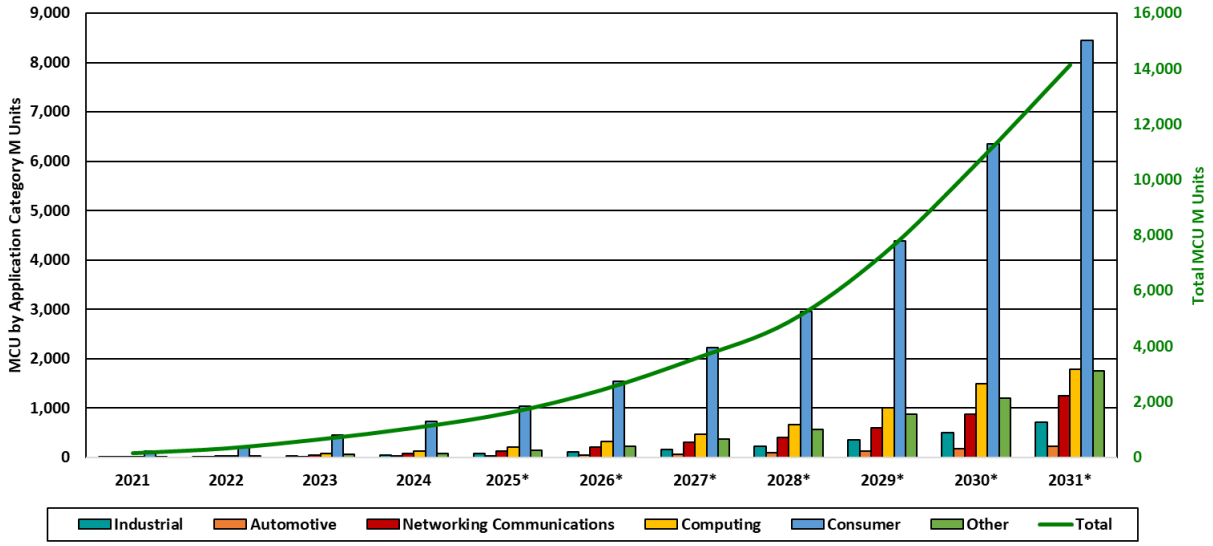


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

**Table 24: RISC-V SoC MCU Functionality Revenues by Application Category, 2022 - 2031**

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | $18.6 | $46.4 | $90.2 | $162.6 | $253.8 | $382.0 | $561.2 | $871.1 | $1,250.7 | $1,745.3 | 48.5% |
| **Automotive** | $61.3 | $215.7 | $384.6 | $608.6 | $881.0 | $1,191.3 | $1,563.4 | $2,119.4 | $2,815.9 | $3,554.0 | 34.2% |
| **Networking** | $65.0 | $274.5 | $580.3 | $1,006.7 | $1,569.9 | $2,202.6 | $2,993.1 | $4,081.6 | $5,699.9 | $8,120.7 | 41.6% |
| **Computing** | $198.9 | $497.2 | $921.6 | $1,592.2 | $2,548.4 | $3,786.8 | $5,323.8 | $7,536.7 | $11,012.6 | $13,626.8 | 43.0% |
| **Consumer** | $519.4 | $1,686.2 | $3,011.3 | $4,496.9 | $6,700.2 | $9,286.6 | $12,208.4 | $16,991.9 | $23,727.1 | $31,554.8 | 38.4% |
| **Other** | $36.6 | $88.7 | $128.8 | $217.7 | $370.9 | $598.7 | $889.7 | $1,370.0 | $1,885.2 | $3,083.3 | 55.5% |
| **Total** | $899.8 | $2,808.7 | $5,116.7 | $8,084.7 | $12,324.1 | $17,447.9 | $23,539.6 | $32,970.7 | $46,391.3 | $61,684.8 | 40.3% |
| **Percent Growth** | 126.0% | 212.2% | 82.2% | 58.0% | 52.4% | 41.6% | 34.9% | 40.1% | 40.7% | 33.0% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

For this analysis, we counted the device revenues of SoCs where the use of RISC-V-based MCUs occurred. We counted the entire revenue generated by shipment of the part into the markets, so we are not

counting the IP revenue of the RISC-V-based FSM, but the device ASP times the units being shipped as the determining factor for counting the revenues. In this case, an SoC could have many MCUs on it and was counted the same as if it had only one MCU on it.

The Consumer segment generated the highest amount of revenue, followed by the Computing and Networking segments.

**Figure 29: RISC-V SoC MCU Functionality Revenues by Application Category, 2021-2031**

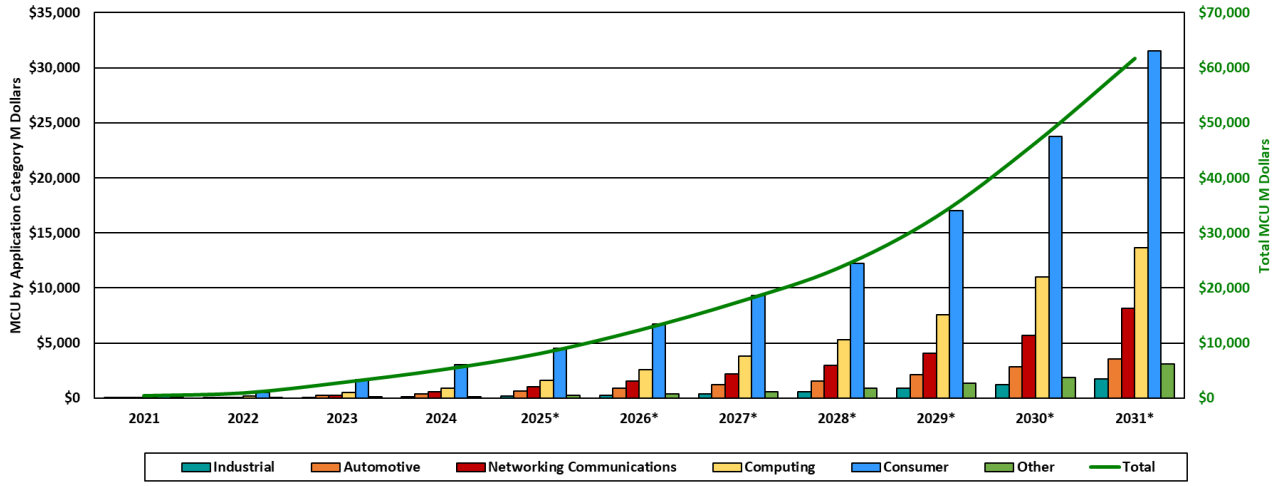


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

**Table 25: RISC-V SoC Co-processor Functionality Unit Volumes by Application Category, 2022 - 2031**

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 0.5 | 1.1 | 2.1 | 3.8 | 6.1 | 9.5 | 14.7 | 23.6 | 35.2 | 48.5 | 52.8% |
| **Automotive** | 3.6 | 25.6 | 46.6 | 74.4 | 107.3 | 148.4 | 196.2 | 267.4 | 358.2 | 458.1 | 35.4% |
| **Networking** | 3.7 | 16.2 | 36.5 | 66.6 | 108.8 | 158.2 | 224.1 | 325.5 | 477.6 | 694.0 | 47.8% |
| **Computing** | 10.4 | 28.3 | 55.2 | 93.6 | 145.8 | 207.9 | 289.4 | 401.2 | 571.3 | 712.4 | 40.3% |
| **Consumer** | 27.6 | 118.7 | 249.1 | 431.7 | 692.6 | 1,012.3 | 1,412.7 | 2,013.3 | 2,856.8 | 3,828.2 | 43.9% |
| **Other** | 0.0 | 1.1 | 2.0 | 3.5 | 6.1 | 10.1 | 17.3 | 33.7 | 53.5 | 104.3 | 76.4% |
| **Total** | 45.9 | 191.2 | 391.6 | 673.6 | 1,066.8 | 1,546.3 | 2,154.5 | 3,064.7 | 4,352.5 | 5,845.4 | 43.4% |
| **Percent Growth** | 43.4% | 316.7% | 104.8% | 72.0% | 58.4% | 45.0% | 39.3% | 42.2% | 42.0% | 34.3% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

In SoCs, a co-processor refers to a specialized processing unit that operates alongside the main processor (CPU) within the same integrated circuit. Co-processors are designed to perform specific tasks or computations more efficiently than the main CPU, thereby offloading certain types of processing to improve overall system performance and efficiency.

The integration of co-processors within an SoC architecture contributes to the overall efficiency, performance, and capabilities of the system by distributing workload and utilizing specialized hardware

for specific tasks, thus improving the user experience in various applications ranging from mobile devices to high-performance computing systems.

Consumer devices are incorporating AI functionality and are a driver for the use of co-processors. In this segment, their use in high-volume systems like smartphones, UHDTV, and other large-volume systems generates a large portion of the unit shipments. The Consumer segment is followed closely by the Computing and Networking segments.

Figure 30: RISC-V SoC Co-processor Functionality Unit Volumes by Application Category, 2021-2031

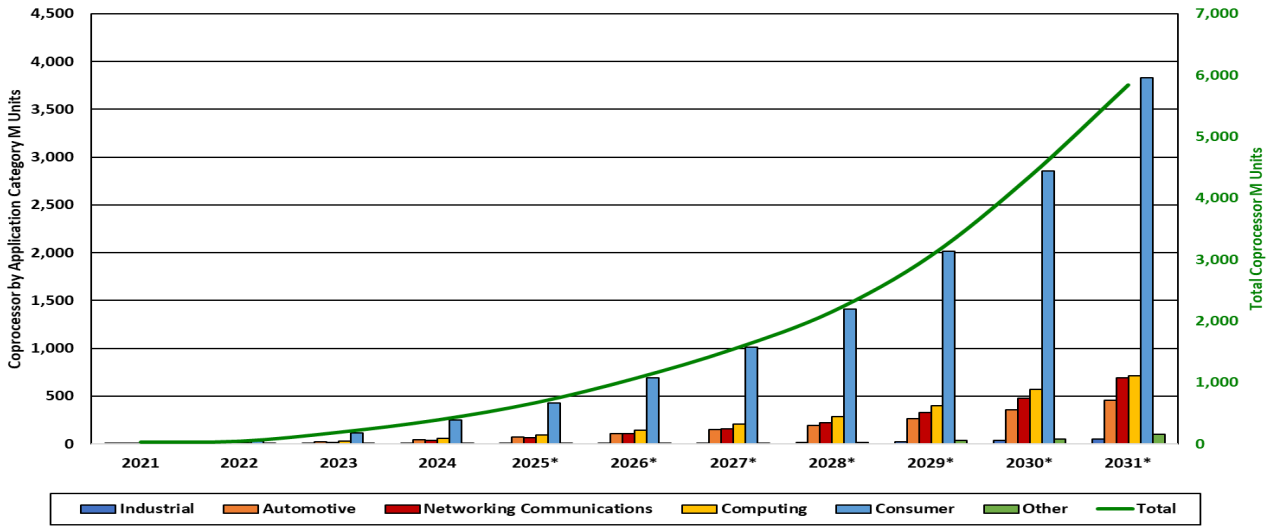


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 26: RISC-V SoC Co-processor Functionality Revenues by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | $35.1 | $64.0 | $130.4 | $220.8 | $317.8 | $405.2 | $552.3 | $755.3 | $974.9 | $1,236.7 | 33.3% |
| **Automotive** | $101.4 | $481.9 | $871.9 | $1,394.1 | $1,984.9 | $2,689.8 | $3,514.3 | $4,710.3 | $6,286.4 | $8,084.8 | 34.0% |
| **Networking** | $61.7 | $275.2 | $628.3 | $1,139.5 | $1,813.6 | $2,528.1 | $3,493.2 | $4,808.0 | $6,581.9 | $9,229.8 | 41.7% |
| **Computing** | $236.6 | $534.3 | $1,098.5 | $1,993.2 | $3,046.1 | $4,369.1 | $5,915.3 | $7,968.1 | $11,298.6 | $13,605.5 | 37.7% |
| **Consumer** | $107.2 | $622.8 | $1,243.3 | $2,160.4 | $3,433.0 | $4,938.1 | $6,924.3 | $10,076.2 | $15,017.8 | $20,663.3 | 45.7% |
| **Other** | $3.4 | $14.8 | $27.2 | $47.0 | $84.8 | $138.3 | $232.8 | $399.9 | $626.4 | $1,134.2 | 70.0% |
| **Total** | $545.4 | $1,993.0 | $3,999.6 | $6,955.0 | $10,680.1 | $15,068.6 | $20,632.2 | $28,717.8 | $40,786.0 | $53,954.4 | 40.7% |
| **Percent Growth** | 55.3% | 265.5% | 100.7% | 73.9% | 53.6% | 41.1% | 36.9% | 39.2% | 42.0% | 32.3% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Consumer devices are incorporating AI functionality and are a driver for the use of co-processors. In this segment, their use in smartphones, AR/VR, UHDTV and other higher-performing systems, and higher-ASP systems generates a large portion of the revenues.

The Consumer segment is followed closely by the Computing and Networking segments.

Figure 31: RISC-V SoC Co-processor Functionality Revenues by Application Category, 2021-2031

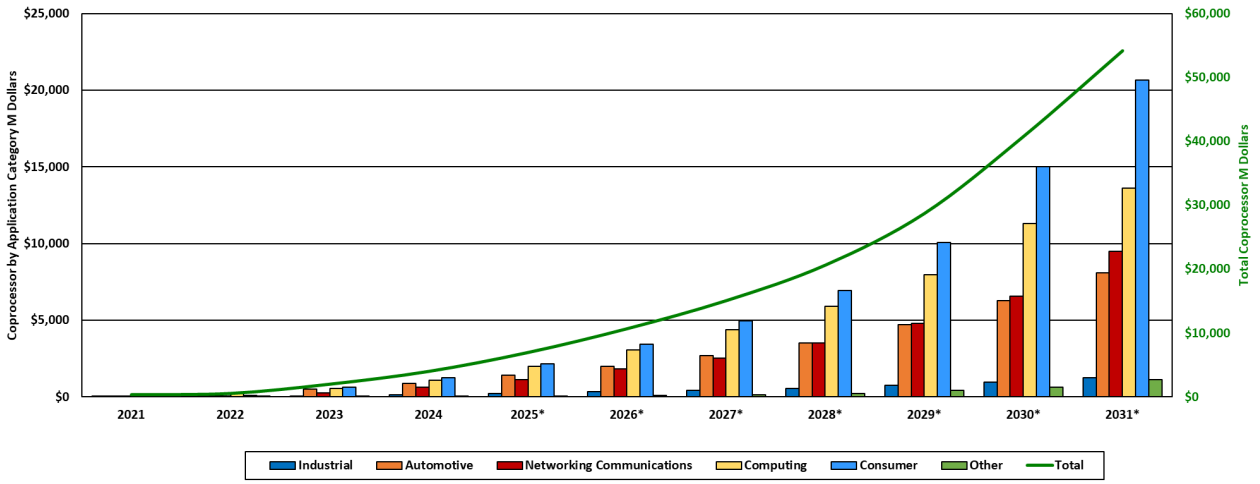


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Table 27: RISC-V SoC Host Processor Functionality Unit Volumes by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 0.1 | 0.2 | 0.4 | 0.9 | 1.8 | 3.1 | 5.5 | 10.0 | 17.1 | 23.0 | 70.3% |
| **Automotive** | 5.5 | 30.4 | 56.5 | 91.2 | 133.2 | 184.3 | 245.5 | 334.0 | 452.0 | 575.7 | 35.9% |
| **Networking** | 1.0 | 2.5 | 6.1 | 12.8 | 22.5 | 36.3 | 54.6 | 83.4 | 113.6 | 193.1 | 57.3% |
| **Computing** | 11.0 | 18.7 | 34.8 | 61.2 | 97.1 | 145.8 | 209.1 | 309.8 | 471.8 | 589.0 | 45.8% |
| **Consumer** | 48.4 | 129.3 | 328.9 | 573.1 | 897.8 | 1,290.2 | 1,798.8 | 2,557.6 | 3,296.1 | 4,598.5 | 41.5% |
| **Other** | 0.0 | 1.0 | 1.5 | 2.9 | 5.1 | 8.5 | 14.6 | 29.8 | 47.4 | 88.9 | 76.4% |
| **Total** | 65.9 | 182.0 | 428.2 | 742.1 | 1,157.5 | 1,668.3 | 2,328.0 | 3,324.5 | 4,398.0 | 6,068.1 | 41.9% |
| **Percent Growth** | 110.9% | 176.4% | 135.2% | 73.3% | 56.0% | 44.1% | 39.5% | 42.8% | 32.3% | 38.0% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

An applications processor refers to a specific type of processor core within the SoC that is primarily designed to handle general-purpose computing tasks and run various applications and software.

Applications processors are commonly found in a diverse range of devices, including smartphones, tablets, smart TVs, automotive infotainment systems, IoT devices, and other embedded systems where running software applications and handling complex tasks are necessary.

The integration of an applications processor within an SoC allows for a comprehensive solution on a single chip, enabling efficient and streamlined execution of software applications and facilitating seamless user experiences across various devices and platforms.

Many of the systems mentioned above are consumer systems, and their high unit volumes help drive the units in the analysis.

The Consumer segment is followed by usage in the Computer and Automotive segments.

Figure 32: RISC-V SoC Host Processor Functionality Unit Volumes by Application Category, 2021-2031

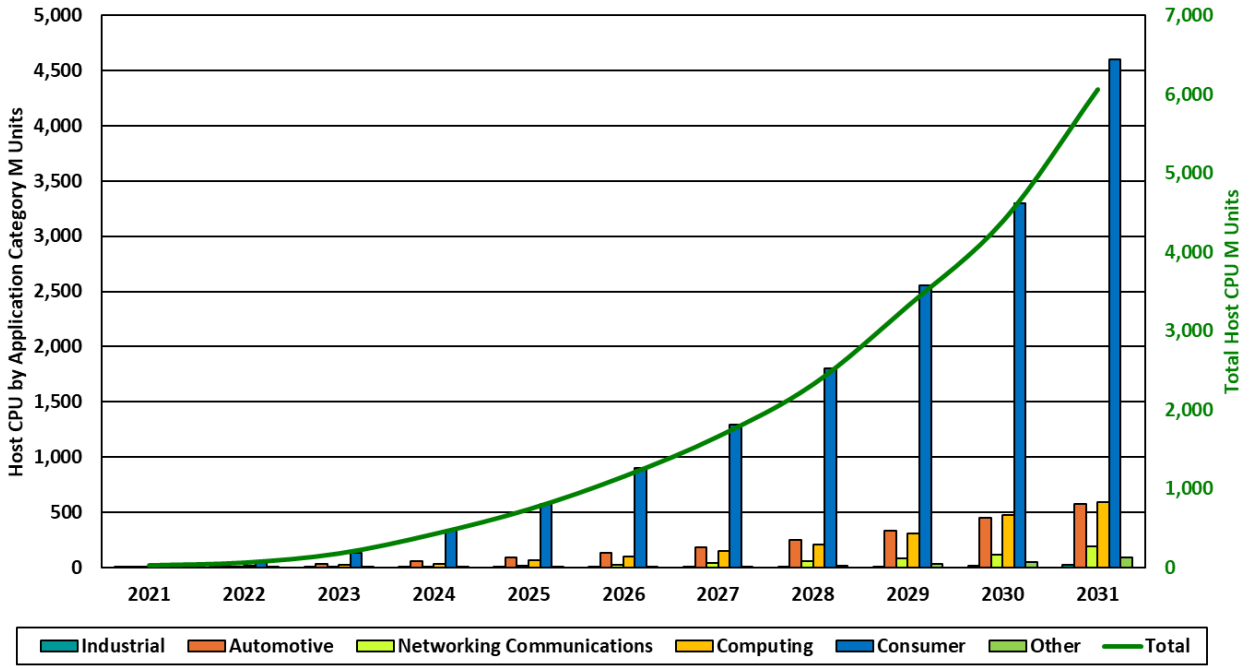


Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Table 28: RISC-V SoC Host Processor Functionality Revenues by Application Category, 2022 - 2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | $0.6 | $1.4 | $3.7 | $7.7 | $13.4 | $20.9 | $34.6 | $59.9 | $96.8 | $172.2 | 67.7% |
| **Automotive** | $101.4 | $610.1 | $1,134.8 | $1,845.2 | $2,670.5 | $3,626.9 | $4,777.1 | $6,389.3 | $7,822.6 | $10,329.5 | 33.3% |
| **Networking** | $29.2 | $82.8 | $228.6 | $510.5 | $845.3 | $1,259.9 | $1,790.2 | $2,592.1 | $4,092.1 | $5,526.1 | 48.7% |
| **Computing** | $404.5 | $835.0 | $1,955.5 | $4,042.5 | $5,796.8 | $7,901.8 | $10,166.8 | $13,185.6 | $18,121.8 | $21,723.9 | 32.3% |
| **Consumer** | $230.7 | $858.9 | $2,651.7 | $4,732.5 | $7,133.8 | $9,793.3 | $13,378.4 | $18,615.5 | $25,423.4 | $35,065.2 | 39.6% |
| **Other** | $2.9 | $13.7 | $22.6 | $43.6 | $78.7 | $128.4 | $216.9 | $377.4 | $593.2 | $1,052.8 | 70.0% |
| **Total** | $769.2 | $2,401.9 | $5,996.8 | $11,182.1 | $16,538.5 | $22,731.2 | $30,363.9 | $41,219.8 | $56,149.8 | $73,869.8 | 37.0% |
| **Percent Growth** | 74.8% | 212.3% | 149.7% | 86.5% | 47.9% | 37.4% | 33.6% | 35.8% | 36.2% | 31.6% | |

Forecast Years: 2025 - 2031 Source: The SHD Group, April 2026

Applications processors are incorporating AI functionality and are a driver for the use of functionality. In this segment, their use in smartphones, AR/VR, UHDTV, and other higher-performing systems, and higher-ASP systems, generates a large portion of the revenues.

The Consumer segment is followed closely by the Computing and Automotive segments.

Figure 33: RISC-V SoC Host Processor Functionality Revenues by Application Category, 2021-2031

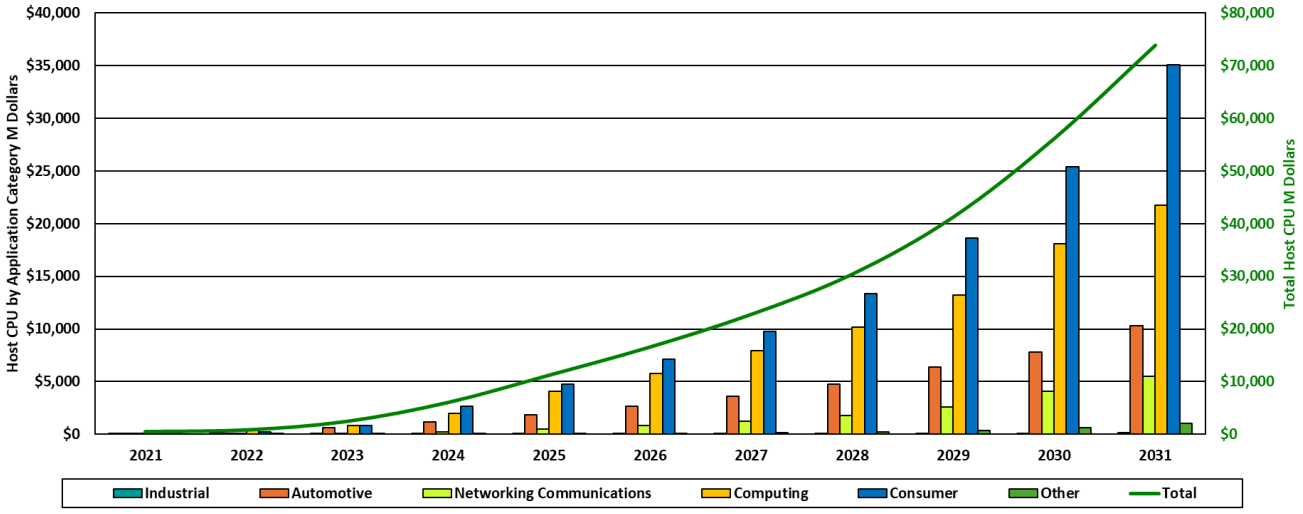


**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

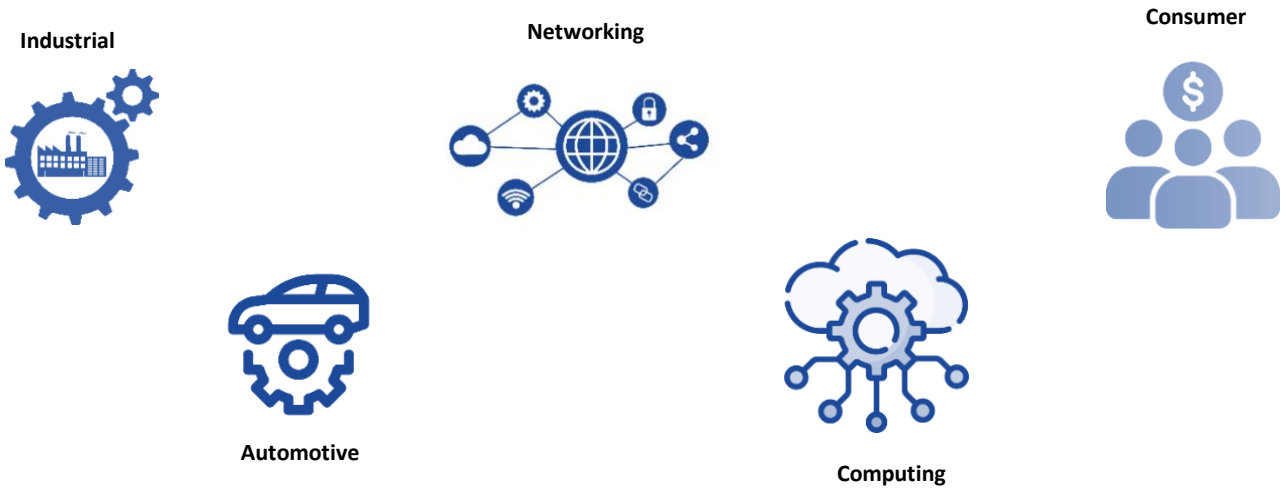

# VIII. RISC-V SoC Design Starts

The SHD Group's 2025 analysis examines RISC-V activity across a broad span of semiconductor manufacturing nodes, ranging from established mid-geometry processes to the latest advanced nodes entering volume design work. The assessment also segments SoC design starts by device category, aligned with the architectures profiled throughout this report. Application-level aggregation is provided for the major end-market groupings: Industrial, Automotive, Networking, Computing, Consumer, and a residual category for designs that fall outside these primary domains.

Given the increasingly diverse structure of contemporary SoCs, the dataset includes any design that integrates RISC-V in any role. In many cases, RISC-V cores serve as auxiliary controllers, secure enclave processors, or domain-specific compute elements within architectures that may otherwise rely on different primary CPU instruction sets. In other instances, RISC-V forms the central compute complex. The objective is to capture the breadth of RISC-V incorporation across the full spectrum of SoC types, independent of vendor or end-market orientation, to reflect how the architecture is being adopted and deployed in 2025's heterogeneous design environment.

**Table 29: RISC-V SoC Design Starts by End Application Category**

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Design Start Units** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Industrial** | 37 | 42 | 47 | 50 | 59 | 65 | 79 | 79 | 83 | 100 | 12.2% |
| **Automotive** | 56 | 74 | 72 | 96 | 118 | 129 | 134 | 140 | 146 | 149 | 7.6% |
| **Networking** | 52 | 92 | 126 | 208 | 234 | 269 | 290 | 303 | 340 | 357 | 9.4% |
| **Computer** | 45 | 79 | 180 | 215 | 253 | 281 | 321 | 331 | 365 | 400 | 10.9% |
| **Consumer** | 76 | 199 | 266 | 299 | 343 | 368 | 393 | 436 | 469 | 525 | 9.8% |
| **Other** | 11 | 13 | 19 | 28 | 31 | 42 | 46 | 53 | 59 | 66 | 15.4% |
| **Total** | 277 | 499 | 710 | 896 | 1,038 | 1,154 | 1,263 | 1,342 | 1,462 | 1,597 | 10.1% |
| **Percent Growth** | 149.5% | 80.1% | 42.3% | 26.2% | 15.8% | 11.2% | 9.4% | 6.3% | 8.9% | 9.2% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

- Industrial activity is expected to rise from 47 design starts in 2024 to 50 in 2025, with expansion continuing toward 100 units by 2031, corresponding to a 12.2% compound annual growth rate from 2025 through 2031.
- Automotive designs are projected to increase from 72 design starts in 2024 to 96 in 2025, advancing toward 149 units by 2031 and reflecting a 7.6% CAGR from 2025 through 2031.
- Networking designs are anticipated to grow from 126 design starts in 2024 to 208 in 2025, reaching 357 units by 2031, which equates to a 9.4% CAGR from 2025 through 2031.
- Computing applications are expected to move from 180 design starts in 2024 to 215 in 2025, progressing toward 400 units by 2031 with a 10.9% CAGR across the 2025–2031 period.
- Consumer designs are forecast to expand from 266 in 2024 to 299 in 2025, with the trajectory pointing toward 525 design starts by 2031 and a 9.8%CAGR from 2025 through 2031.
- Designs falling into the Other category are projected to increase from 19 in 2024 to 28 in 2025, eventually reaching 66 units by 2031, yielding a 15.4% CAGR between 2025 and 2031.

Table 30: Total RISC-V SoC Market Share by Application Category 2022–2031

| Percent of Market | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 |
|---|---|---|---|---|---|---|---|---|---|---|
| Industrial | 1.6% | 1.7% | 1.5% | 1.4% | 1.5% | 1.6% | 1.8% | 2.1% | 2.3% | 2.6% |
| Automotive | 1.6% | 2.5% | 2.1% | 2.0% | 2.2% | 2.5% | 2.9% | 3.3% | 3.6% | 3.9% |
| Networking | 1.7% | 4.1% | 4.0% | 4.0% | 4.7% | 5.4% | 6.0% | 6.7% | 7.1% | 7.9% |
| Computer | 78.3% | 69.1% | 70.6% | 75.2% | 72.6% | 69.9% | 66.6% | 63.0% | 59.7% | 55.0% |
| Consumer | 15.6% | 21.2% | 20.6% | 16.4% | 17.5% | 18.6% | 20.1% | 21.6% | 23.6% | 25.5% |
| Other | 1.1% | 1.4% | 1.1% | 1.1% | 1.5% | 2.0% | 2.6% | 3.3% | 3.6% | 5.0% |
| Total Market | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% | 100.0% |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

- The Industrial application market share is roughly flat over the forecast period starting at 1.6% in 2024 and forecast to end in 2031 at 2.6%.
- The Automotive application market share started in 2024 at 2.1% and is forecast to reach 3.9% by 2031.
- The Networking application market share started in 2024 at 4.0% and is forecast to reach 7.9% by 2031.
- The Computer application market share started in 2024 at 70.6% and is forecast to decline to 55.0% by 2031.
- The Consumer application market share started in 2024 at 20.6% and is forecast to increase to 25.5% by 2031.
- The Other application market share started in 2024 at 1.1% and is forecast to increase to 5.0% by 2031.

Figure 34: AION Silicon Design Services for High-Complexity, Multicore SoC Designs


**Courtesy of Aion Silicon**

AION Silicon is a design-services provider focused on advanced ASIC and SoC development, with capabilities extending from early architectural definition through tape-out and production hand-off. The company operates as a foundry-neutral partner and maintains an IP-agnostic design philosophy, which allows it to support a range of customer requirements across AI, automotive, networking, and high-performance compute markets. Their service model emphasizes risk reduction through early-stage architectural planning, structured IP selection, and tightly coordinated design execution—an approach that reflects the increasing technical complexity and cost associated with leading-edge SoC development.

The company's rebranding underscores a strategic shift toward domains where custom silicon is becoming more critical, particularly AI accelerators, automotive perception and control systems, and edge-processing devices. These segments are characterized by heterogeneous compute architectures and differentiated workload requirements, making them well suited to custom ASIC development rather than reliance on general-purpose solutions. AION Silicon's experience with multi-core designs, complex interconnects, and high-reliability architecture aligns with these trends.

Seen in the broader semiconductor market, AION Silicon reflects a structural shift: as complexity and NRE costs rise, the market increasingly depends on specialized design houses to bridge the gap between concept and manufacturable silicon. Their relevance to RISC-V growth lies not in proprietary IP but in the capability to assemble heterogeneous subsystems—RISC-V cores, AI accelerators, memory fabrics, I/O subsystems - into optimized SoCs tailored to specific vertical workloads.

# IX. RISC-V Regional Forecasts

The following section looks at the revenues for each of the six application categories on a regional basis for RISC-V-enabled SoCs.

Table 31: Total Global RISC-V SoC Regional Revenues by Application Category 2022–2031

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **B Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Americas** | $4.1 | $8.6 | $18.5 | $34.5 | $43.3 | $51.7 | $59.4 | $70.0 | $85.8 | $102.2 | 19.8% |
| **Europe & EMEA** | $1.2 | $2.5 | $5.1 | $9.3 | $11.8 | $14.3 | $16.8 | $20.2 | $25.2 | $30.7 | 22.0% |
| **Japan** | $0.9 | $1.9 | $3.8 | $6.7 | $8.5 | $10.2 | $11.8 | $14.1 | $17.6 | $21.4 | 21.3% |
| **China** | $3.0 | $6.4 | $13.3 | $24.0 | $30.9 | $37.9 | $45.0 | $54.7 | $69.6 | $86.4 | 23.8% |
| **ASIA Pacific** | $2.4 | $5.2 | $10.9 | $19.8 | $25.8 | $32.2 | $38.9 | $48.0 | $61.7 | $77.5 | 25.5% |
| **Total B Dollars** | $11.6 | $24.5 | $51.6 | $94.3 | $120.3 | $146.2 | $171.8 | $207.0 | $259.9 | $318.2 | 22.5% |
| **Percent Growth** | 30.2% | 111.0% | 110.9% | 82.7% | 27.6% | 21.6% | 17.5% | 20.5% | 25.6% | 22.4% | |
| **Market Penetration** | 3.0% | 4.8% | 8.4% | 12.6% | 14.9% | 17.3% | 19.8% | 22.9% | 27.4% | 33.1% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

- Americas: Revenue grows from $18.476B in 2024 to $34.519B in 2025, a rise of roughly 87 percent, with the region expected to reach $102.237B in 2031, corresponding to a multi-year annual pace near 20 percent.
- Europe & EMEA: Revenue moves from $5.128B in 2024 to $9.289B in 2025, an increase of about 81 percent, advancing toward $30.672B in 2031 at an annualized rate close to 22 percent.
- Japan: Revenue rises from $3.780B in 2024 to $6.706B in 2025, up around 77 percent, with the region projected to reach $21.362B in 2031, reflecting an annual trajectory near 21 percent.
- China: Revenue increases from $13.299B in 2024 to $23.971B in 2025, an expansion close to 80 percent, progressing to $86.430B in 2031 with an annual growth rate near 24 percent.
- Asia Pacific: Revenue climbs from $10.931B in 2024 to $19.807B in 2025, a gain of roughly 81 percent, trending toward $77.541B in 2031 at an annualized rate around 23 percent.
- Total Revenue expands from $51.614B in 2024 to $94.292B in 2025, an uplift close to 83 percent, with the worldwide total projected to reach $318.241B in 2031 and sustain a multi-year annual rate near 23 percent.
- Please note that the CAGRs for all these markets are very high, given that they start from very small base years.

Figure 35: Total Global RISC-V SoC Regional Revenues 2021-2031

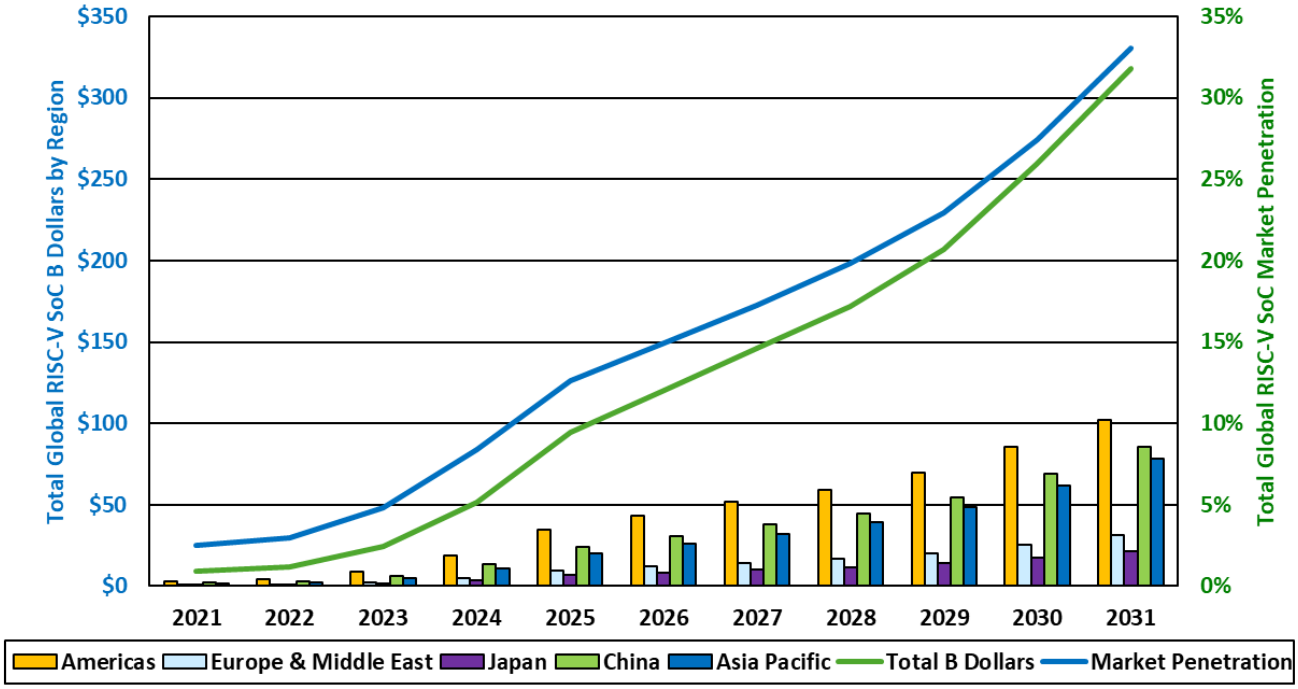


| Forecast Years: 2025 - 2031 | Source: The SHD Group, April 2026 |
|---|---|

Nuclei System Technology ("Nuclei") is an emerging semiconductor IP provider focused on the commercialization of RISC-V processor architectures. Founded in China, the company has positioned itself as a key enabler of open-architecture chip design by delivering configurable CPU IP cores spanning ultra-low-power microcontrollers to high-performance application processors. Its portfolio is designed to support a wide range of verticals, including IoT, automotive electronics, AI, and industrial control, with an emphasis on scalability and customization for system-on-chip (SoC) developers.

From a technology standpoint, Nuclei differentiates itself through a comprehensive IP stack and ecosystem approach, integrating processor cores, development tools, and software support to reduce time-to-market for semiconductor customers. The company emphasizes production readiness and flexibility, enabling customers to tailor instruction set architectures and configurations to specific application requirements. Strategic collaborations—such as its work with Siemens on trace and debugging solutions—highlight its focus on improving development efficiency and strengthening its position within the broader RISC-V ecosystem.

In terms of market positioning, Nuclei operates in a rapidly expanding segment driven by the global shift toward open-standard processor architectures as alternatives to proprietary designs. With hundreds of customers and large-scale deployment of its cores, the firm is increasingly relevant in Asia's semiconductor supply chain while also expanding its global footprint. As RISC-V adoption accelerates across edge computing and embedded systems, Nuclei's vertically integrated IP strategy and ecosystem partnerships place it in a competitive position against both established IP vendors and emerging open-source challengers.

Figure 36: Broad Product Portfolio for Nuclei System Technologies

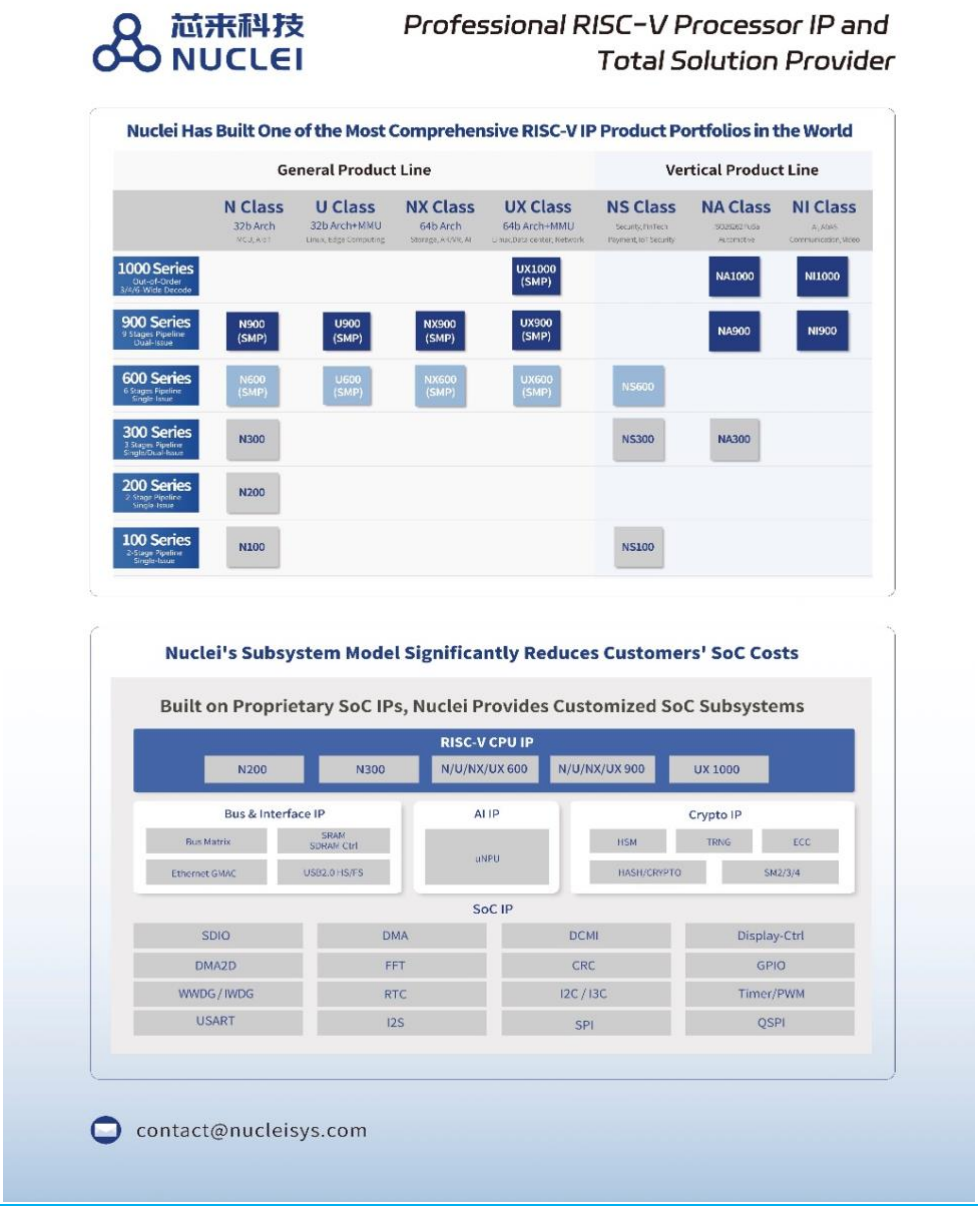


Source: Courtesy of Nuclei Systems Technology

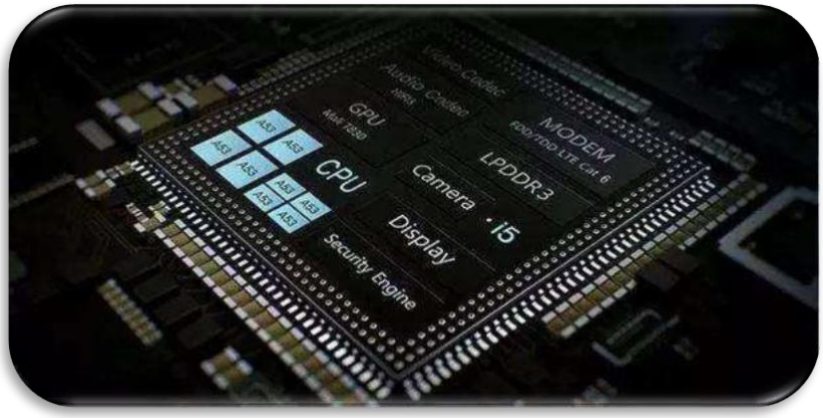

# X. RISC-V CPU IP Market Analysis

Figure 37: RISC-V IP Market Forecast 2021-2031

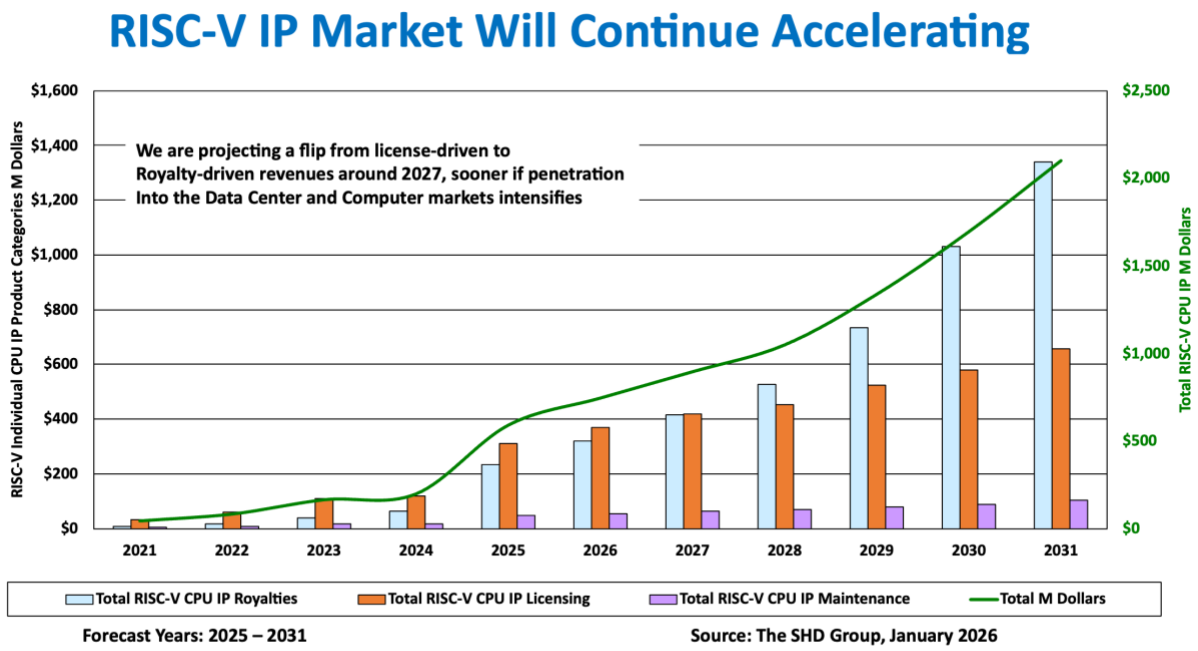


## Projecting RISC-V IP Into the Future

Figure 37 shows The SHD Group's forecast for the RISC-V CPU IP market from 2021 to 2031. In this forecast, we are projecting that around 2027 or 2028, the royalty revenues will start to exceed the licensing revenues. By the end of this forecast period, we expect the royalties to be about 1.2X the licensing revenues.

This change in the mix of revenues is driven by the ramping up of SoCs with higher performing RISC-V CPU cores shipping in greater volumes in the later years of the forecast period. A similar occurrence happened in the early years of the original CPU IP market and in the broader IP market as well. Eventually, the gap between royalty and licensing revenues tends to be reduced as IP types enter the market over time and licensing for the new types of IP blocks occurs. So, there is a tendency for these trend lines to tighten and loosen as these different IP types enter the market and are adopted at different rates by SoC designers.

Figure 38 presents The SHD Group's revised forecast for the RISC-V CPU IP market spanning 2021 through 2031. Our latest analysis projects that by roughly 2028, royalty revenues will begin to outpace licensing revenues. By the end of the forecast window, royalties are expected to reach approximately 1.3 times the level of licensing income—reflecting the maturation of RISC-V-based SoCs entering high-volume production.

This crossover is driven by the accelerating shipment of higher-performance RISC-V cores in complex SoCs now ramping into mainstream applications such as automotive, industrial, and AI-enabled computing. The same pattern has repeated historically in other IP markets: initial licensing peaks as early adopters enter,

then royalties dominate as production volume scales. Comparable transitions occurred in the early HDMI and USB IP markets, where competition compressed licensing costs even as cumulative royalties increased with shipment growth.

Unlike those earlier IP domains, however, the current RISC-V CPU IP market remains far from maturity. It continues to exhibit hallmarks of a high-growth, rapidly evolving sector—continuous architectural innovation, new entrants, and deepening ecosystem support across EDA, software, and IP subsystems. Market revenues are expected to surpass $1.8 billion by 2031, with compound annual growth remaining above 25 percent through the latter half of the decade.

The CPU IP market's intrinsic dynamism also drives a structurally higher royalty-to-license ratio than other IP categories. The constant introduction of new cores and the industry's ongoing shift toward heterogeneous computing architectures ensure continued differentiation and royalty uplift. The SHD Group's updated interviews—covering over 60 ecosystem companies under NDA and public sources—show that globally, most RISC-V chip shipments are enabled by (in alphabetical order): Alibaba's DAMO Academy, Andes Technology, Nuclei and SiFive. With Cadence expanding their Tensilica offerings to include RISC-V, Synopsys's (going to Global Foundries) transitioning their ARC-V line to RISC-V and MIPS (now a Global Foundries company) expanding to support the RISC-V architecture, all further strengthen the ecosystem.

Momentum among system OEMs also continues to build. Nvidia, Qualcomm, Samsung, Western Digital and several automotive Tier-1 suppliers have disclosed internal deployments of RISC-V IP for control, security, and co-processing functions. The next inflection point will occur as leading server, PC, and industrial OEMs bring RISC-V-based SoCs to market in volume starting 2026–2027.

Importantly, The SHD Group maintains that the market will not transition to a single-ISA world. Designers will continue to adopt heterogeneous SoCs combining RISC-V with cores based on Arm, x86, or others — each optimized for specific functions such as real-time control, vector processing, AI acceleration, or host-level orchestration. RISC-V's openness and configurability position it as the natural complement in these multi-ISA designs rather than a full replacement for incumbent cores.

In short, the RISC-V CPU IP market in 2025 is characterized by expanding adoption, growing royalty leverage, and continuous innovation across the ecosystem. The market is entering a phase of broad commercial deployment—still dynamic, not mature—and positioned for sustained outperformance relative to traditional CPU IP segments through 2031.

**Table 32: Total RISC-V CPU IP Market Revenues 2022–2031**

| | | | | | | | | | | | CAGR % |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **M Dollars** | **2022** | **2023** | **2024** | **2025** | **2026** | **2027** | **2028** | **2029** | **2030** | **2031** | **25 - 31** |
| **Total RISC-V CPU IP Royalties** | $16.7 | $35.8 | $49.2 | $79.6 | $120.8 | $194.6 | $281.0 | $415.3 | $630.3 | $944.0 | 51.0% |
| **Total RISC-V CPU IP Licensing** | $59.1 | $103.2 | $129.4 | $142.7 | $165.5 | $195.2 | $258.9 | $365.7 | $531.5 | $799.6 | 33.3% |
| **Total RISC-V CPU IP Maintenance** | $8.9 | $16.7 | $26.4 | $34.6 | $44.5 | $48.9 | $62.0 | $83.9 | $117.4 | $168.8 | 30.2% |
| **Total Revenues** | $84.6 | $155.7 | $205.0 | $256.9 | $330.8 | $438.7 | $601.9 | $864.9 | $1,279.2 | $1,912.4 | 39.7% |
| **Percent Growth** | 84.7% | 84.0% | 31.6% | 25.3% | 28.8% | 32.6% | 37.2% | 43.7% | 47.9% | 49.5% | |

**Forecast Years: 2025 - 2031** **Source: The SHD Group, April 2026**

Figure 38: Total RISC-V CPU IP Market Revenues 2021–2031

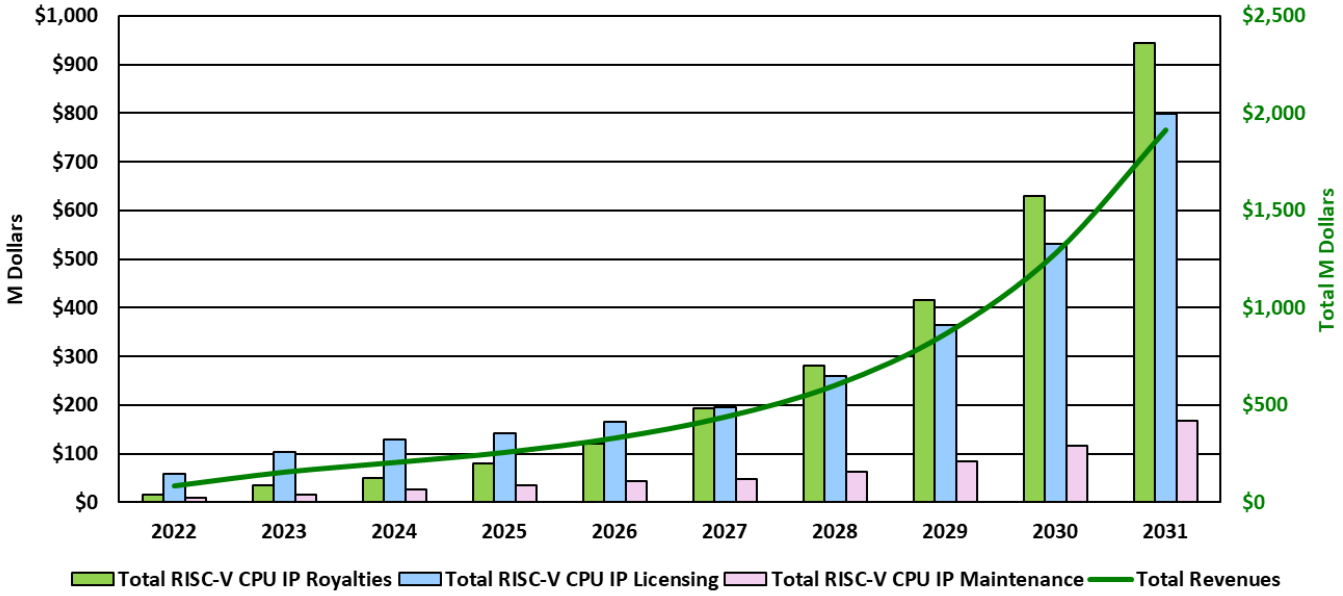


| Forecast Years: 2025 - 2031 | Source: The SHD Group, April 2026 |
|---|---|

- Total RISC-V CPU IP royalties reached $49.2M in 2024 and are forecast to grow to $944.0M by 2031, a CAGR of 51.0%.
- Total RISC-V CPU IP licensing revenues were $129.4M in 2024 and are forecast to reach $799.6M by 2031, a CAGR of 33.3%.
- Total RISC-V CPU IP maintenance revenues were $26.4M in 2024 and are forecast to reach $168.8M by 2031, a CAGR of 30.2%.
- Total RISC-V CPU IP revenues were $205.0M in 2024 and are forecast to reach $1,912.4M by 2031, a CAGR of 39.7%.

The SHD Group believes this forecast is reasonable given the high degree of interest in the RISC-V ISA and the number of systems it can be deployed into. The current massive buildout for AI Data Centers indicates increased numbers of silicon solutions will be required. This also indicates that as device complexity continues to increase, higher levels of IP will also be needed to meet the higher market requirements for more functionality. RISC-V will benefit from this trend going forward.

It is also necessary to state that the very high CAGR growth rates are due to coming from very low starting base numbers and then growing from there. These rates will moderate over time

## RISC-V IP Vendor List

During the course of 2025, there have been several changes to the RISC-V CPU IP market with acquisitions of companies, new entrants to the market and several exits of companies that left the RISC-V market for various reasons.

The following table shows these changes to the RISC-V Market.

Table 33: Changes to the RISC-V CPU IP Market

| Company | RISC-V IP Providers | New Entrants | Acquisition or Exit |
|---|---|---|---|
| **Ahead Computing** | Yes | Yes | |
| **Alibaba Damo Academy** | Yes | | |
| **Ainekko/Veevx** | Yes | Yes | Ainekko acquired Esperanto RISC-V CPU IP |
| **Akeana** | Yes | Yes | |
| **Andes Technology** | Yes | | |
| **Bluespec** | Yes | | |
| **Codasip** | Yes | | |
| **Cortus** | Yes | | |
| **Global Foundries** | Yes | Yes | Synopsys CPU IP group acquired by Global Foundries |
| **InCore Semiconductor** | Yes | | |
| **MIPS** | Yes | | Acquired by Global Foundries |
| **Nuclei System Technology** | Yes | | |
| **OpenHW Foundation** | Yes | | |
| **Oxmiq Labs** | Yes | Yes | |
| **Quintauris** | Yes | Yes | |
| **Semidynamics** | Yes | | |
| **SiFive** | Yes | | |
| **StarFive Technology** | Yes | | |
| **Tenstorrent** | Yes | Yes | |
| **Ventana Micro Systems** | No | | Acquired by Qualcomm |
| **VyperCore** | Yes | Yes | |

**Source: The SHD Group, April 2026**

Vibrant markets in their early years often exhibit growth spurts and sudden declines. The RISC-V CPU IP market is no exception with companies entering, exiting and being acquired by other, larger companies in this market over time. Some of the recent acquisitions might be put down to being 'acqui-hires' where the larger acquiring company is only interested in the engineering talent that come with the acquisition. However, in the case of RISC-V CPU IP, this is not really happening since, in most cases the acquiring company has signaled they intend to use the IP for their own, internal uses, or intend to continue to license the IP to external customers.

There have been more companies entering the market than exiting it and The SHD Group thought it important to make this point: the RISC-V Market is growing very quickly and continues to attract people with good ideas who want to put them into the marketplace.

Until we see more exits compared to entrants, the RISC-V market is doing just fine!

# XI. RISC-V Landscape

By 2025, the RISC-V architecture has moved decisively from early adoption to broad commercial engagement. The ecosystem now spans dozens of suppliers, encompassing CPU IP vendors, toolchain providers, foundries, and system companies building silicon across nearly every major market segment. The transition from curiosity to commitment is evident in expanding design activity for AI acceleration, automotive control, and edge computing, where customization and openness provide tangible design advantages. The entry of major EDA firms, new processor startups, and collaborative ventures such as Quintauris has reinforced RISC-V's legitimacy as a foundational compute architecture. As licensing models mature and royalty revenues begin to scale, the RISC-V IP market is entering a sustained growth phase characterized by architectural diversity, competitive innovation, and deepening ecosystem maturity.

## RISC-V IP Companies

- **Ahead Computing** is a RISC-V CPU IP vendor developing high-performance, out-of-order processor cores for data center, AI, and advanced compute systems. The company is positioning RISC-V beyond embedded roles and into performance-driven SoCs that integrate CPUs with accelerators and modern memory and interconnect architectures.
- **Ainekko** has acquired the IP, existing silicon and development tools and frameworks from Esperanto, who has exited the RISC-V IP market. their mission is to open-source a fully composable AI infrastructure stack — from silicon (RTL) through runtime, tooling, and APIs - under a community-governed model.
- **Akeana**, launched from stealth in 2024, has entered the RISC-V market with a family of scalable CPU cores and supporting system IP. Its portfolio spans microcontroller-class, application-class, and data-center-grade cores, each designed for configurability and rapid SoC integration. The company positions itself as a next-generation alternative to incumbent CPU IP suppliers, with a focus on performance-per-watt efficiency and licensing flexibility.
- **Alibaba Damo Academy,** Alibaba's semiconductor and IP division, is one of China's largest RISC-V CPU IP suppliers and a major contributor to the global RISC-V ecosystem. Its Xuantie family spans high-performance 64-bit Linux-class processors (C910/C920/C930) to ultra-low-power 32-bit embedded cores (E902/E906/E907). The company licenses its IP through the OpenXuantie platform and integrates these cores across Alibaba's cloud, IoT, and AI chips. By 2025, the XuanTie RISC-V designs serve as a national reference for open-ISA computing within China's semiconductor strategy.
- **Andes Technology,** a founding member of RISC-V International. With more than 17 years of experience, the company provides a comprehensive lineup of efficient, highly configurable 32-/64-bit cores. Its products address domains from deeply embedded controllers to data center–class multicore designs, including support for AI, DSP, vector, and floating-point operations. Andes maintains strong traction in automotive, where its ISO 26262-certified cores meet ASIL B–D requirements. Its integrated toolchain and development ecosystem enable customers to tailor hardware/software partitioning for specific workloads, making it one of the most complete commercial RISC-V offerings worldwide.
- **Bluespec** provides RISC-V processor IP, virtual development environments, and hardware acceleration tools. Its cloud-based emulation and high-speed prototyping platforms enable faster software bring-up and design iteration. Bluespec's model-driven design tools continue to reduce

development risk and time to market, especially for teams integrating RISC-V and heterogeneous compute architectures.

- **Codasip** continues to expand its customizable RISC-V processor lineup built using Codasip Studio and the CodAL architectural description language. This methodology allows customers to introduce proprietary extensions and performance optimizations while maintaining software compatibility. Codasip's new high-performance L31 and L70 cores reflect growing demand for design-specific compute efficiency.

- **Cortus** offers a mixed portfolio of 32-/64-bit RISC-V and proprietary cores targeting industrial control, IoT, edge compute, and automotive applications. In 2025 the company enhanced its mixed-signal and security IP suite to complement its processor offerings, emphasizing integration flexibility across diverse market verticals.

- **InCore Semiconductor** focuses on automation-driven design methodologies for RISC-V IP and SoC development. Its platform simplifies CPU customization and integrates accelerator design into heterogeneous multicore architectures. InCore's emphasis on modularity aligns with the industry's transition toward domain-specific compute and chiplet-based architectures, particularly in industrial and automotive segments.

- **MIPS** (acquired by GlobalFoundries), transition to RISC-V architecture is now complete. Its latest generation of CPU IP targets high-performance compute, automotive, and networking markets, leveraging the company's long experience in virtualization, real-time processing, and functional-safety design.

- **Nuclei System Technology,** one of China's leading RISC-V CPU IP vendors and its embedded and application-grade 32-/64-bit cores power a wide range of consumer, industrial, and IoT SoCs across Asia. Nuclei has signaled their intent to expand their presence in the U.S, Europe, Latin America and Africa to accelerate the adoption of RISC-V in 2026.

- **Oxmiq Labs**, launched in 2025 by GPU veteran Raja Koduri, Oxmiq Labs is building RISC-V-based GPU and tensor-processing IP compatible with existing CUDA software environments. Its goal is to enable seamless AI compute portability across non-proprietary hardware, representing an important frontier for RISC-V in heterogeneous compute

- **Quintauris**, the European RISC-V joint venture formed by Bosch, Infineon, Nordic Semiconductor, NXP, Qualcomm, and later STMicroelectronics, became operational in 2024. The consortium's mandate is to accelerate industrial adoption of RISC-V through shared IP development, standardization, and ecosystem enablement within Europe.

- **Semidynamics** continues to focus on fully customizable 64-bit Out-of-Order and In-Order RISC-V cores. Its Tensor Unit and Vector Unit architectures now support wider AI/ML workloads, offering scalable options for both datacenter and edge inference designs.

- SiFive remains a cornerstone of the RISC-V movement. The company partners with leading semiconductor manufacturers, hyperscale operators, and consumer OEMs to deliver scalable high-performance, low-power compute. Its product portfolio spans wearables, mobile, data center, and aerospace, all built on standardized yet customizable RISC-V cores. SiFive's collaborations with TSMC and Intel Foundry have strengthened its position as a reference supplier for commercial RISC-V IP.

- **StarFive Technology** has expanded its RISC-V portfolio to include the Dubhe-90 and JH-series SoCs, supporting applications from consumer electronics to edge servers. The company is among the first to ship RISC-V-based PC-class SoCs and continues to strengthen the Chinese ecosystem around open ISA computing.

- **Synopsys** (ARC and ARC-V CPU IP group acquired by Global Foundries) is a premier member of RISC-V International, the leader in the EDA market and the #2 IP vendor worldwide. Following its 2024 introduction of the ARC-V family, Synopsys has now commercialized this RISC-V-based CPU IP line, combining its deep EDA expertise with processor design experience. The ARC-V family targets embedded control, signal processing, and AI workloads, with full toolchain and software-stack support. Synopsys' participation underscores the mainstreaming of RISC-V in production-grade design flows, complementing its extensive verification, interface, and security IP portfolios. The CPU IP group at Synopsys will continue to license their CPU IP and service their customers until the acquisition is finalized.

- **Tenstorrent's** family of RISC-V CPU IP now extends from two-wide to eight-wide decode implementations, integrated into its own data-center processors and licensed to select partners. The company's approach—combining RISC-V cores with its AI-focused interconnect and accelerator IP—positions it as a bridge between CPU and neural-compute domains.

- **Ventana Micro Systems** (acquired by Qualcomm) continues to advance the use of chiplets for scalable RISC-V computing. Its 64-bit Veyron V2 platform targets data center and HPC environments, offering customers both IP licensing and chiplet-based integration models. The company's approach enables rapid development of high-performance custom RISC-V processors and highlights the architectural flexibility of chiplet ecosystems in 2025.

- **VyperCore**, founded in 2024 in the UK, VyperCore develops high-performance RISC-V CPU IP optimized for cloud and data-center workloads. Its evaluation and co-design environment, delivered as a cloud service, enables customers to explore core configurations and performance trade-offs rapidly.

## EDA Tool Providers

Discussing the 3rd party IP market only tells part of the story. The other essential element is the ecosystem of Electronic Design Automation (EDA) tools available from a wide range of suppliers. Without these tools, neither complex SoC development nor the creation of the IP blocks themselves would be possible. EDA remains the critical enabler that transforms architectural concepts into manufacturable silicon.

Most EDA vendors continue to operate with a CPU-agnostic philosophy, ensuring that their design, verification, simulation, and emulation platforms can support any processor architecture—including RISC-V. Over time, however, these companies have refined and extended their tools to address the increasingly intricate challenges of SoC design. This evolution reflects a broader industry shift toward tighter integration between software and hardware development flows. The growing incorporation of AI-driven automation and machine learning capabilities within EDA platforms further highlights the industry's push to accelerate design closure, improve yield prediction, and manage the rising complexity of advanced process nodes.

In 2025, major EDA vendors have deepened their alignment with the RISC-V ecosystem, integrating dedicated support into their toolchains and design automation environments. Synopsys, now incorporating technology from both Imperas and Ansys, has enhanced its verification and power analysis

workflows to accelerate RISC-V-based SoC development. Cadence has expanded its AI-assisted design and simulation capabilities across its Innovus and Xcelium platforms, improving efficiency for multi-core RISC-V integration and validation. Siemens EDA continues to strengthen its verification and emulation portfolio, offering tools optimized for RISC-V hardware/software co-design. Collectively, these advances demonstrate a maturing EDA landscape where RISC-V is no longer treated as an emerging alternative but as a mainstream architecture supported across the full design-to-silicon workflow.

- **Ansys (acquired by Synopsys),** remains a leading provider of analysis tools essential to CPU integration, including RedHawk-SC and PowerArtist for power, IR drop, and thermal reliability. Its tools now incorporate machine-learning algorithms that predict thermal gradients and electromigration in RISC-V SoCs more accurately. As part of its ongoing integration with Synopsys announced in 2024, Ansys continues to deliver independent multiphysics solutions widely used in RISC-V processor and chiplet designs.

- **Arteris** As the pioneer of NoC technology—with 71+ patents, 4B SoCs shipped, 200+ customers, and 725+ design starts—**Arteris** enables RISC-V teams to seamlessly connect compute and AI/ML subsystems using silicon-proven NoC IP and SoC Integration Automation. Our technology unifies the "protocol salad" of hundreds of reused IPs (ACE-Lite, AXI, AHB, APB, CHI, and more) and helps de-risk SoC integration with low-latency, low-power, high-bandwidth, flexible, and proven system IP backed by expert support.

- **Breker Verification Systems** focuses on scenario-based verification with its TrekSoC™ and TrekSoC-XE products, which automatically generate self-checking C test cases for RISC-V cores and subsystems. In 2025, the company expanded its toolset to support complex cache-coherency and memory-ordering verification, addressing multicore RISC-V SoCs designed for AI inference and networking.

- **Cadence Design Systems** continues to expand its presence in RISC-V design enablement through its Genus and Innovus digital implementation tools and the Xcelium verification suite. Its Helium Virtual and Palladium Z2 emulation platforms now feature optimized RISC-V models, allowing software bring-up before silicon availability. Cadence's Tensilica DSP team has also partnered with RISC-V ecosystem developers to ensure cross-ISA co-simulation and unified tool flows.

- **OneSpin Solutions**, now operating as Siemens EDA Formal, continues to provide industry-proven formal verification solutions used extensively for RISC-V functional-safety and security certification. Its pre-packaged assertion libraries and equivalence-checking methodologies have become central to automotive and aerospace verification workflows.

- **Synopsys** remains the largest EDA supplier and a major catalyst for RISC-V enablement. Its Fusion Compiler and Design Compiler platforms form the core of physical synthesis and optimization flows used throughout the industry. As of 2025, Synopsys has fully migrated its MetaWare tools to support the ARC-V RISC-V CPU family, and its verification suite includes dedicated flows for RISC-V compliance. Following the 2024 acquisition of Imperas, Synopsys now integrates virtual prototyping, instruction-set simulation, and verification IP for RISC-V under a single portfolio—solidifying its role as both EDA and processor-IP provider.

- **Siemens EDA (formerly Mentor Graphics),** delivers comprehensive verification and system design solutions through Questa, ModelSim, and Tessent. In 2025, Siemens added RISC-V-specific formal verification libraries and integrated support for AI-assisted coverage closure. Its Calibre platform continues to dominate physical verification, ensuring sign-off accuracy for SoCs using RISC-V

cores. The company's shift toward digital-twin methodologies enables earlier validation of system-level performance for RISC-V-based architectures.

- **Silvaco** provides design software for analog, mixed-signal, and process modeling that complements CPU IP design. Its TCAD and circuit-simulation tools are now used by RISC-V implementers for early-stage device characterization and PDK optimization. In 2025, Silvaco introduced AI-assisted layout tools aimed at reducing turnaround time for RISC-V-based SoCs targeting emerging-node geometries.
- **S2C** has become a visible participant in the RISC-V ecosystem through its FPGA-based prototyping and hardware-acceleration platforms. Its Prodigy Logic System supports real-time verification of RISC-V SoCs, and the company's cloud-enabled emulation solutions are increasingly used in Asia for rapid hardware validation. S2C's partnerships with leading CPU-IP vendors illustrate the growing demand for fast, scalable prototyping in the RISC-V domain.

The broader EDA landscape surrounding RISC-V has matured significantly over the past year. Vendors are deploying AI-driven optimization across place-and-route, formal verification, and test generation, while cloud collaboration platforms now support multi-team co-design. The emergence of chiplet-based design strategies has also pushed tool vendors to enhance hierarchical partitioning, 3D-IC analysis, and package-level thermal modeling. Together, these developments mark the transition of RISC-V design flows from adaptation to full industrialization.

## RISC-V Software Development Tools & OS:

- **Red Hat**: Red Hat continues to lead in enterprise open-source software, and by 2025 its RHEL (RISC-V) variant is certified or in customer engagements on RISC-V hardware. The company has adapted core libraries, kernel optimization and virtualization stacks to the open-ISA ecosystem, strengthening software portability for RISC-V-based hosts.
- **Synopsys (Metaware / ARC-V toolchain)**: Synopsys offers a full suite of development tools for embedded and host RISC-V processors via its MetaWare and ARC-V lines. Support spans 32-bit RMX real-time cores, RHX signal-processing engines and 64-bit RPX application-class processors. As of 2025, the toolchain includes compiler optimizations, profile-guided tuning, full debug support and turnkey enablement for safety- and security-critical applications in automotive, industrial and aerospace domains. It has been announced that the ARC-V product line will moving to Global Foundries.
- **Emproof**: Emproof remains a specialist in embedded security, delivering integrity, boot-chain and algorithm-protection IP for RISC-V-based systems, especially in constrained-edge, automotive and IoT domains.
- **Canonical (Ubuntu)**: Ubuntu's RISC-V support continues to mature. Canonical supplies a full Ubuntu distribution ported to RISC-V64, optimized for cloud-edge devices, and contributes to upstream kernel and hardware-abstraction support for RISC-V platforms.
- **Wind River**: Wind River's VxWorks® and the Wind River Linux product lines fully support RISC-V-based SoCs. Their multi-OS hypervisor and safety-certified stacks (ASIL B–D) are now deployed in edge, aerospace and automotive contexts using RISC-V controllers and accelerators.
- **SUSE**: SUSE offers enterprise-grade Linux distributions with RISC-V support, delivering infrastructure for servers, cloud and edge nodes built around open-ISA compute.
- **Codeplay (Intel)**: Codeplay supports heterogeneous computing on RISC-V platforms, offering compilers, language runtimes and accelerator-programming frameworks tailored to AI/ML and automotive workloads on RISC-V hosts and accelerators.
- **RISE Project (Linux Foundation)**: The RISC-V Software Ecosystem (RISE) project coordinates foundational software support—compilers, toolchains, libraries, kernel subsystems, virtualization and debug/profiling infrastructure—across mobile, IoT, data-center and automotive segments.
- **Google (Android)** continues to support RISC-V within the Android ecosystem, primarily through its participation in the RISC-V International Android SIG and ongoing AOSP development efforts. However, in 2024 Google removed RISC-V from the Android Common Kernel and Generic Kernel Image flows, citing the need for greater ecosystem stability before unified support can be maintained. As of 2025, Android can still be built for RISC-V platforms, but commercial deployments require additional vendor engineering and integration work. Formal certification and full kernel-level support are expected to resume once RISC-V platform maturity and standardization progress further.
- **Syntacore** has added a full development-toolkit release (2025.09) that bundles IDEs, debugging tools, simulators and board-support packages for its RISC-V core IP family, addressing software-partner readiness from day one.

- **NVIDIA (CUDA on RISC-V)** In 2025, NVIDIA announced that the CUDA programming model now supports RISC-V CPUs in heterogeneous systems, enabling RISC-V hosts to run high-level CUDA-based AI workloads while GPUs handle parallel processing.

## FPGA (with RISC-V)

By 2025, the convergence of open-architecture compute and programmable logic has matured into a core design strategy across multiple markets. FPGA and eFPGA vendors now provide verified flows, RISC-V-ready soft cores, and even Linux-capable processor subsystems that accelerate time-to-prototype and enable differentiated silicon in production. Programmable logic has become a key vehicle for innovation in RISC-V SoCs, spanning edge AI, data-center acceleration, industrial automation, and defense systems.

- **Achronix** Speedster 7t family combines a 2D Network-on-Chip architecture with embedded AI acceleration blocks and supports Linux-capable RISC-V soft processors through its partnership with Bluespec. These cores can be deployed directly in the FPGA fabric, enabling scalable multicore compute for networking and data-center applications.
- **AMD (Xilinx),** AMD's Xilinx division maintains its leadership in adaptive compute through the Versal™ ACAP family and Vivado® Design Suite. These platforms support RISC-V core integration alongside Arm® and AI engines, enabling heterogeneous compute architectures that pair open-ISA flexibility with established tool flows.
- **Altera (a standalone company spun out of Intel in 2024),** following its separation from Intel, Altera re-emerged as an independent FPGA supplier while continuing to leverage Intel Foundry technology. Its Agilex® FPGA family and Quartus® Prime software suite now provide full support for integrating RISC-V cores as host or co-processor elements. Altera's 2025 roadmap emphasizes hybrid FPGA–SoC designs that combine proprietary blocks with open-ISA subsystems for communications, networking, and automotive markets.
- **Efinix** continues to advance its Trion® and Titanium® FPGAs, which integrate turnkey RISC-V soft-core options and toolchain support. In 2025, these devices are widely adopted for low-power, high-efficiency compute at the edge, combining RISC-V control processing with configurable AI and DSP blocks.
- **Flex Logix (now part of Analog Devices),** acquired by Analog Devices in late 2024, continues to supply its EFLX® eFPGA platform optimized for RISC-V integration. Under ADI's ownership, Flex Logix is extending its AI inference IP and software toolchain for mixed-signal and sensor-fusion systems, broadening RISC-V's reach into analog-digital co-processing and embedded AI.
- **Gowin Semiconductor,** Gowin's GW-series FPGAs advertise RISC-V compatibility across tool flows and soft-core templates, targeting consumer, industrial, and IoT markets. The company's rapid expansion in Asia makes it an increasingly visible mid-tier competitor in programmable logic.
- **Lattice Semiconductor's** Nexus™ platform and Radiant® Design Suite focus on compact, power-efficient applications. As of 2025, Lattice offers validated RISC-V soft-core templates and reference designs for IoT, industrial control, and embedded security systems.
- **Microchip (PolarFire® Family),** Microchip's PolarFire®FPGA family—together with the Libero® SoC Design Suite—now supports extensive RISC-V integration. In 2025 the company expanded its spaceflight offering by announcing a new RT PolarFire® SoC FPGA series, alongside its role in NASA's HPSC RISC-V computing initiative for spacecraft systems.

- **QuickLogic's** EOS S3 and ArcticPro™ platforms feature open-source RISC-V integration with the SensiML Analytics Toolkit. Designed for always-on sensing and edge-AI applications, they demonstrate how open architectures and FPGA flexibility can coexist in ultra-low-power designs.
- **S2C** Introduced in 2025, S2C's Prodigy™ S8-100 FPGA prototyping platform offers 100 million-gate capacity and optimized verification flows for advanced RISC-V SoCs. Collaborations with Andes Technology and other IP vendors position it as a leading prototyping environment for high-complexity designs.
- **Menta** provides embedded-FPGA (eFPGA) IP blocks that integrate seamlessly with RISC-V processors. This approach allows SoC designers to embed reconfigurable logic alongside RISC-V cores, enabling post-deployment algorithm updates and instruction-set customization—ideal for industrial and defense applications.

## Related IP and Service Providers

- **AION Silicon** is an ASIC and SoC design-services provider supporting advanced custom silicon development from requirements to architectural definition through tape-out and production, operating with a foundry-neutral and IP-agnostic model. The company focuses on complex, heterogeneous designs for AI, automotive, networking, and high-performance compute, where custom architectures are increasingly favored over general-purpose solutions. Within the RISC-V ecosystem, AION Silicon's role centers on SoC integration, assembling RISC-V cores, accelerators, interconnects, and I/O into workload-optimized devices rather than supplying proprietary processor IP.
- **Alphawave Semi,** now part of Qualcomm, provides configurable IP subsystems, high-speed connectivity IP, and chiplet-ready design frameworks for RISC-V SoCs. The company's portfolio includes die-to-die interconnects, PCIe/CXL interface IP, and multi-chip integration solutions, positioning it strongly in the chiplet-based computing ecosystem.
- **Arteris IP** continues to lead the interconnect IP market, supplying both coherent and non-coherent Network-on-Chip (NoC) technologies that link the diverse IP blocks within advanced SoCs. Its Ncore cache-coherent and FlexNoC interconnects have evolved to improve scalability, reduce latency, and optimize power efficiency. As of 2025, Arteris is marking its third consecutive year of ISO 26262 Tool Confidence Level 1 (TCL1) certification for its SoC Integration Automation platform—further solidifying its leadership in automotive and functional-safety-qualified design automation.
- **Baya Systems** develops automation frameworks and IP for chiplet-based design, focusing on rapid partitioning and integration across heterogeneous SoCs. Its tools complement RISC-V-based chiplet ecosystems by enabling fast architectural exploration, design reuse, and optimized physical integration for multi-die systems.
- **BrainChip** develops neuromorphic AI processors based on event-driven, spiking neural network architectures, with RISC-V used for system control and SoC integration. Its Akida technology is offered both as packaged devices and as licensable IP, enabling integration into 3rd party SoCs for ultra-low-power edge AI inference in applications such as vision, audio, and sensor processing across industrial, automotive, and consumer markets.
- **Dover Microsystems** advances its CoreGuard® security IP, embedding dedicated sentry logic beside the CPU to verify instruction execution against hardware-enforced safety and privacy rules. This architecture blocks malicious or unsafe operations before they can execute, making CoreGuard® a compelling choice for RISC-V-based defense, industrial, and critical-infrastructure

systems.

- **Eliyan Labs** has gained prominence for its innovative interconnect technology, notably its high-efficiency physical layer IP supporting ultra-low-power, high-bandwidth chiplet communication. Its NuLink™ die-to-die PHY enables near-monolithic performance across advanced packaging architectures, including those combining RISC-V and non-RISC-V cores within multi-chip systems.
- **Metrics Technology** develops RISC-V-specific verification and compliance automation tools that accelerate the validation of both standard and custom extensions. Its cloud-based verification environment supports large-scale simulation and regression testing, helping design teams close verification gaps earlier in the design cycle.
- **MosChip Technologies** is an engineering services firm supporting semiconductor development from architecture definition through silicon implementation and system delivery. The company provides RISC-V–related services including RTL development, verification, physical design, DFT, and IP integration, with experience spanning high-performance cores, mixed-signal ASICs, and large-scale tape-outs. In the RISC-V ecosystem, MosChip functions as a technical enablement partner, helping IP providers and SoC teams reduce development risk and accelerate time to silicon across automotive, industrial, IoT, and edge platforms.
- **proteanTecs** has become a key enabler of reliability and lifecycle analytics. Its embedded telemetry IP embeds deep sensing agents throughout the SoC, collecting performance and aging data during production and field operation. The company's integration with RISC-V designs enables predictive maintenance, yield optimization, and in-field health monitoring for automotive, aerospace, and data-center applications.
- **Rambus** continues to enhance its high-speed interface, memory, and security IP portfolios with expanded support for RISC-V system integration. Its IP solutions for encryption, key management, and interface verification play a central role in protecting heterogeneous compute architectures and high-bandwidth SoCs.
- **Rapid Silicon** has developed a specialized EDA tool suite tied to its FPGA and programmable logic ecosystem. The company's Raptor Design Suite integrates synthesis (e.g., Yosys), place-and-route (e.g., VPR), bitstream generation, debugging, and other common design flows under a unified framework, with tooling that supports its devices and IP infrastructure.
- **UltraSoC, fully integrated into Siemens EDA**, continues to expand its embedded analytics IP for RISC-V systems. Its real-time on-chip monitoring solutions enable deep visibility into system performance, power behavior, and reliability—supporting predictive diagnostics and faster design validation cycles.

Together, these companies provide some of the surrounding the IP foundation that enables the RISC-V ecosystem to mature toward parity with proprietary CPU architectures. Their combined contributions—in interconnect, verification, analytics, chiplet design, and reliability - underscore the accelerating industrialization of RISC-V as the architecture of choice for next-generation heterogeneous SoCs.

Figure 39: proteanTecs Provides Analytic Tools for In-Device Monitoring and Predictive Maintenance

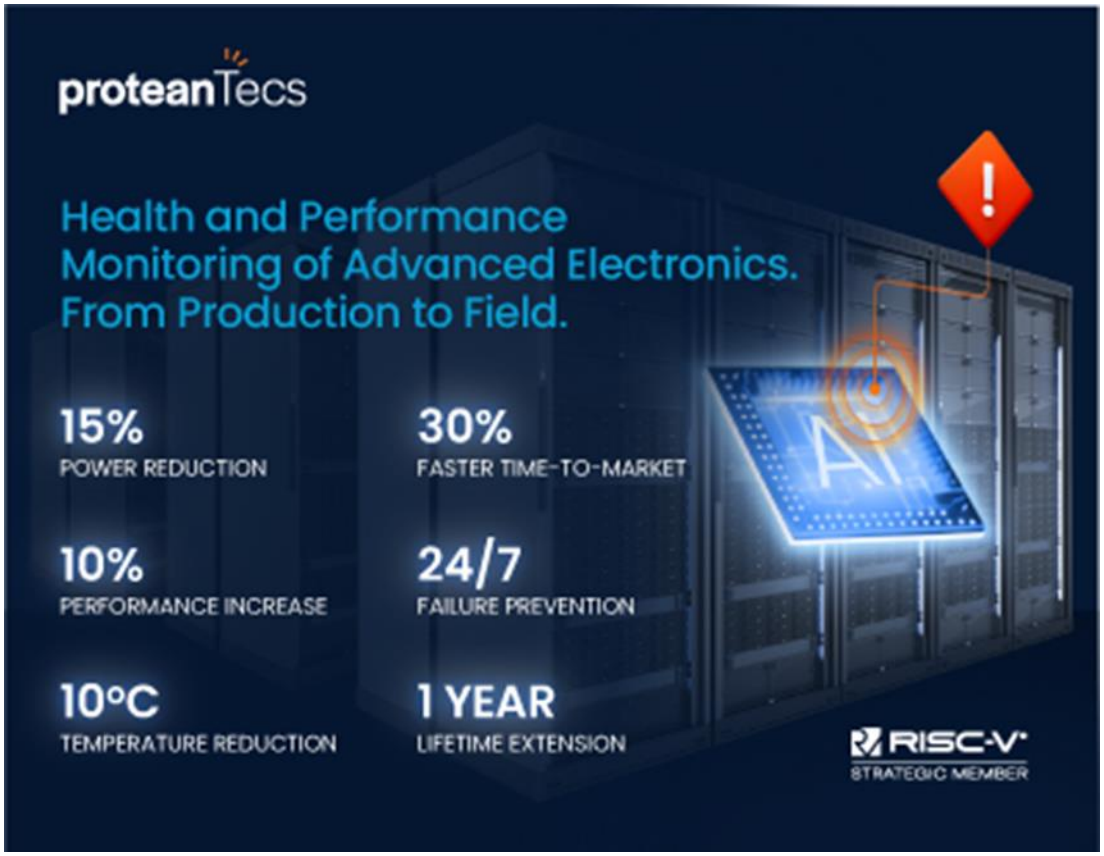


Courtesy of proteanTecs

proteanTecs provides embedded monitoring and analytics technology that enables continuous visibility into the health, performance, power, and thermal behavior of advanced semiconductor devices from production through in-field operation. Its approach is based on integrating on-chip monitoring Agents that generate high-resolution telemetry, allowing systems to report their own operating conditions in real time. This data is analyzed using machine-learning-based analytics to support design validation, manufacturing test and binning, system bring-up, and ongoing reliability monitoring, forming a closed feedback loop across the silicon lifecycle.

The platform supports measurable operational outcomes, including power and temperature reduction, performance optimization, accelerated time-to-market, and predictive failure prevention that can extend system lifetime. proteanTecs is deployed in complex, high-availability markets such as data centers, AI and accelerator platforms, networking infrastructure, and automotive electronics, where reliability and efficiency are critical. Within the RISC-V ecosystem, its architecture-agnostic telemetry and analytics capabilities are increasingly relevant as RISC-V designs move into larger, more heterogeneous, and performance-constrained systems. proteanTecs targets complex semiconductor and electronics markets where reliability, uptime, and performance consistency are critical. Key segments include datacenter infrastructure, telecommunications, automotive (including ADAS and safety-critical electronics), mobile and consumer devices, and aerospace/defense systems.

# XII. Ecosystem Guide

The RISC-V ecosystem is dynamic and fast-growing. This report openly invited all RISC-V International member companies, and others, to participate. All content in this Directory was provided by the companies listed. ***Any company wishing to be included can be added at no charge by completing the entry information at:*** https://theshdgroup.com/company-profile-form/. (Note: Sponsors will have the opportunity to include additional product and company details in the Ecosystem Guide.).

## What is the True Value of the RISC-V ISA and its Ecosystem?

When assessing the evolution of any technology market, it's critical to identify its enduring value drivers - the elements that sustain growth regardless of external factors or competitive pressures. The RISC-V market is no exception. As we move through 2025, adoption continues to accelerate, and the global ecosystem has matured into a genuine alternative to proprietary CPU architectures.

One of the most common claims heard early in RISC-V's rise was that its popularity was simply a reaction to dissatisfaction with traditional ISAs like Arm - particularly the rising cost of CPU IP licenses and royalties.

A fair question therefore remains: If Arm were to drastically reduce its licensing fees, would RISC-V's momentum fade?

At The SHD Group, we believe the answer remains a firm no. The value of RISC-V extends far beyond cost considerations. Its appeal lies in the structural advantages of an open, extensible, and collaborative model that changes how the semiconductor industry innovates.

**1. Supplier Independence**
RISC-V enables true supplier independence. Software written for the RISC-V ISA is largely portable across hardware from different vendors. This allows companies to select or change suppliers based on performance, roadmap, or geopolitical alignment - without rewriting software stacks from the ground up. Before RISC-V, switching between proprietary architectures often required a full revalidation of toolchains and applications, a process so costly that most companies simply avoided it. In an era of supply-chain diversification and regional design strategies, this flexibility has become a decisive competitive factor.

**2. Broad Market Access and Choice**
Proprietary ISAs centralize control under a single company, where every new feature, extension, or license flows through that organization.

RISC-V operates differently: RISC-V International coordinates the ISA standard, but implementers remain free to design custom extensions or microarchitectures without central approval. This has encouraged dozens of companies - from startups to global semiconductor leaders - to enter the CPU IP space. The result is a marketplace of interchangeable, specialized RISC-V cores rather than a single-source dependency.

**3. Freedom to Customize and Differentiate**
Closed architectures require an architectural license for any modification, typically costing millions of dollars and imposing legal and design constraints.

RISC-V, by contrast, is an open standard. Any company can download the specification, design a compatible core, and tailor extensions to specific workloads - from AI inference to sensor fusion - at minimal cost. This freedom has lowered the barrier to architectural innovation and enabled regional initiatives, such as Europe's Quintauris consortium and China's domestic RISC-V programs, to build unique SoC solutions optimized for their markets.

**4. Accelerated Innovation Through Collaboration**
Proprietary architectures innovate within a single company's R&D bandwidth. RISC-V, by contrast, leverages the creativity of an entire global community. Each vendor focuses on its niche - microcontrollers, HPC, automotive, or edge AI - collectively generating innovation at a scale no single company can match.

The pace of progress has become self-reinforcing: tool vendors, IP suppliers, and software developers are iterating in parallel, feeding improvements back into open specifications. As the ecosystem matures, new extensions for vector, DSP, and security workloads are being standardized faster than ever before, shortening the time from concept to silicon deployment.

## The Core Value Proposition for 2025

The essence of RISC-V's value today lies in independence, flexibility, and acceleration. It enables hardware developers to control their roadmaps, expand market choice, and integrate custom IP without licensing friction. It also invites collaboration across nations and companies - an especially important attribute in a geopolitically fragmented semiconductor landscape.

In short, RISC-V has evolved from a cost-driven alternative to a strategic enabler of innovation. It's open-architecture model is not simply a different licensing framework; it represents a different way of building technology - one that prioritizes openness, competition, and shared advancement over control and exclusivity.

## Bottom Line

By 2025, RISC-V stands not just as another ISA, but as a structural shift in how the semiconductor industry approaches compute design.

It delivers:

- **Vendor independence**, allowing market-driven hardware sourcing.
- **Architectural flexibility**, enabling both standardization and differentiation.
- **Collaborative innovation**, accelerating technical progress through diversity.

These principles — more than only licensing terms or cost savings — represent a true and lasting value of the RISC-V ecosystem.

| | |
|---|---|
| **ACCELR** | **Headquarters**:<br>Colombo, Sri Lanka<br><br>**Website:** https://accelr.lk |

**About the Company:**
ACCELR is a technology partner specializing in software and hardware system design and development. The company focuses on building maintainable, high-quality solutions that enhance client value, with strong experience delivering complex, mission-critical systems and a commitment to engineering excellence.

| | | |
|---|---|---|
| AION SILICON<br>**AION Silicon** | **Headquarters**:<br>London, UK | **Founded**:<br>2002 |
| | **Website:** https://aionsilicon.com | |

**About the Company:**
Aion Silicon is a UK-based semiconductor design company specializing in end-to-end ASIC and advanced SoC development. Founded in 2002, the company supports global markets including AI, automotive, high-performance computing, networking, and consumer electronics.

Aion delivers architecture-to-production services through a global engineering organization spanning the UK, US, Spain, Morocco, India, and other regions, with expertise across RTL design, verification, physical design, and DFT. With a strong track record of hundreds of successful silicon tapeouts and deep partnerships across the IP and EDA ecosystem, Aion is recognized as a trusted partner for developing complex, custom SoC solutions that reduce risk and accelerate time to market.

**Key Products Attribute:**
Aion Silicon provides end-to-end custom ASIC and advanced SoC solutions through a high-touch, consultative engineering model covering architecture, IP selection, design, tapeout, and volume production.

Its approach is both foundry-neutral and IP-agnostic, enabling flexibility across process technologies while leveraging deep expertise in AI, automotive, HPC, networking, and 5G applications. The company emphasizes risk-optimized design methodologies, scalability, and rapid time-to-market, supported by a strong execution track record across hundreds of tapeouts.

In addition, Aion offers turnkey silicon lifecycle management—from initial architecture through packaged silicon—as a design partner for products manufactured at Intel, Global Foundries and TSMC foundries, Aion helps customers efficiently bring complex, next-generation semiconductor products to market.

| XUANTIE<br>Alibaba | **Headquarters**:<br>Hangzhou, Zhejiang, China | **Founded**:<br>2001 |
|---|---|---|
| | **Website:** https://www.xrvm.com/ | |

**About the Company:**

**XuanTie** is committed to promoting the cutting-edge research of RISC-V architecture and building the RISC-V ecosystem. It is an international RISC-V ecological leader and provides powerful, intelligent, secure and open new computing architecture services for the digital era.

Since its establishment, XuanTie has continued to delve into RISC-V technology R&D and ecosystem building. It has a series of XuanTie CPUs, which can meet a full range of performance needs at high, medium, and low levels. XuanTie actively build the RISC-V full-stack software and hardware technology in multiple fields and accelerate the implementation of RISC-V. XuanTie processors are widely used in industries such as computer vision, data storage, IoT, networking and communication, smart home, biometrics, and information security protection.

**Key Products Attribute:**

XuanTie (玄铁), the RISC-V processor brand under Alibaba DAMO Academy, delivers a full spectrum of RISC-V CPU IP cores designed to meet high-, mid-, and low-performance requirements across applications ranging from IoT to high-performance computing. Its processors emphasize powerful, intelligent, secure, and open computing architectures, leveraging Alibaba's strengths in cloud computing, AI, and big data to provide reliable IP for the digital era. XuanTie persistently advances self-developed RISC-V technologies, releasing multiple CPU series that incorporate innovations such as vector processing, matrix extensions, virtualization support, and confidential computing to address AI acceleration, real-time control, and high-reliability needs in edge-to-cloud systems. The brand also maintains a robust open-source ecosystem, offering software tools, development platforms, and OS support to enable holistic soft-hardware co-development across diverse industries.

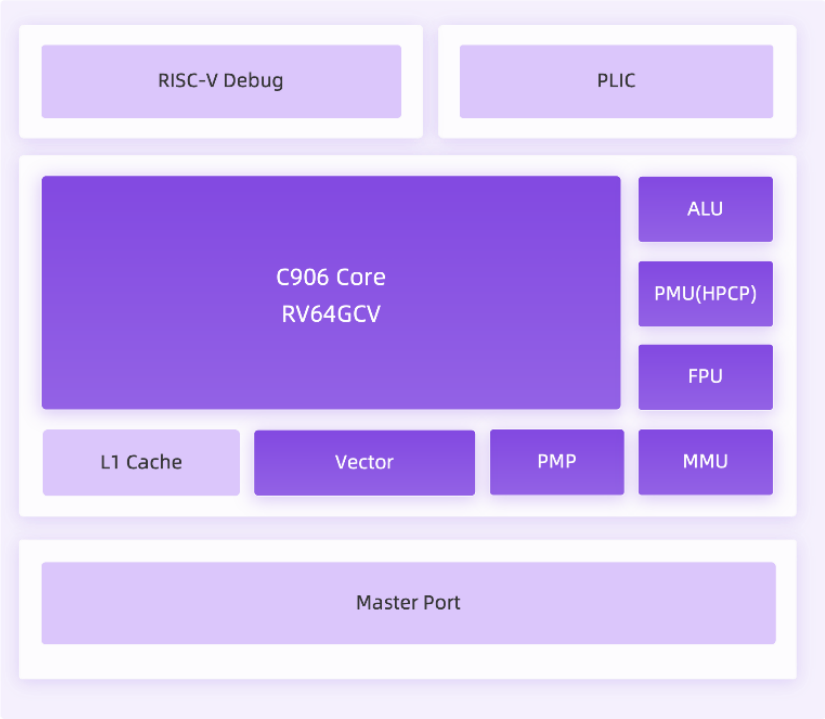

| Allwinner Technology | **Headquarters**:<br>Zhuhai, China |
|---|---|
| | **Website:** http://www.allwinnertech.com/ |

**About the Company:**

Allwinner Technology is a semiconductor company specializing in system-on-chip solutions for smart devices and single-board computers. The company develops processors and platforms for applications including consumer electronics, IoT, and multimedia systems.

| Andes Technology Corp | **Headquarters**:<br>Hsinchu City, Taiwan<br>Sunnyvale CA USA | **Founded**:<br>2005 |
|---|---|---|
| | **Website:** https://www.andestech.com | |

**About the Company:**

A Founding Premier member of RISC-V International, Andes is a publicly listed company and a leading supplier of high-performance/low-power 32/64-bit embedded processor IP solutions, and the driving force in taking RISC-V mainstream. Its V5 RISC-V CPU families range from tiny 32-bit cores to advanced 64-bit Out-of-Order processors with DSP, FPU, Vector, Linux, superscalar, and/or multi/many-core capabilities. By the end of 2022, the cumulative volume of Andes-Embedded™ SoCs has surpassed 12 billion.

**Key Products Attribute:**

Andes Custom Extension (ACE) enables designers to create unique CPU instructions on the performance optimized AndesCore processors. ACE instructions designed specifically for the target applications eliminate the software bottlenecks and significantly improve runtime performance.

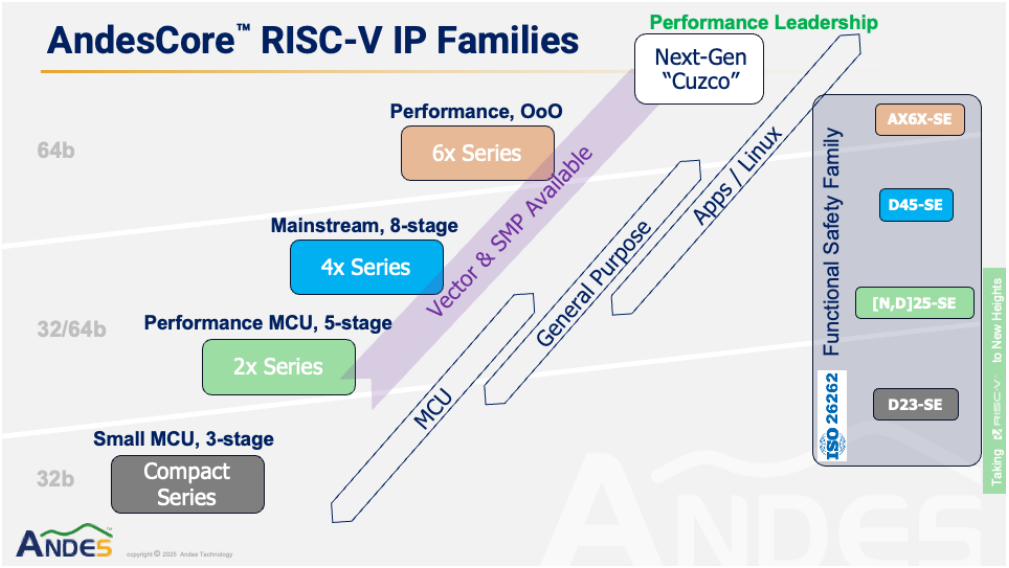

| Antmicro | **Headquarters**:<br>Sweden |
|---|---|
| | **Website:** https://antmicro.com/ |

**About the Company:**

Antmicro is a software-driven technology company providing development services, platforms, and expertise for advanced hardware and software systems. The company focuses on applied R&D across domains including robotics, aerospace, automotive, industrial systems, and edge AI. Antmicro supports innovation in FPGA, ASIC, and edge-to-cloud architectures, helping organizations adopt emerging technologies effectively.

| Arowanie Group | **Headquarters**:<br>Czech Republic |
|---|---|
| | **Website:** https://arowanie.net/ |

**About the Company:**

Arowanie Group focuses on system-on-chip (SoC) research and development, supporting advanced semiconductor design initiatives. The company works on enabling efficient and scalable SoC architectures tailored to emerging computing demands.

| Arteris IP | **Headquarters**: Campbell, CA, USA | **Founded**: 2004 |
|---|---|---|
| | **Website:** https://arteris.com | |

## About the Company:

Arteris specializes in semiconductor technology that accelerates the creation of high-performance, power-efficient silicon with built-in safety, reliability, and security. Arteris' leading solutions include network-on-chip (NoC) interconnect IP, SoC integration automation software, and hardware security assurance, enabling power-efficient, scalable, cost-effective, and secure chip or chiplet designs. The Arteris silicon-proven network-on-chip interconnect IP streamlines communication and optimizes data movement between on-chip components, accelerating the development of sophisticated chips while reducing costs. The Arteris SoC integration automation software delivers cutting-edge tools for packaging, reuse, and integration of diverse IP blocks using the IP-XACT (IEEE 1685) standard. Designers experience a streamlined design environment, from early architecture stages to a fully documented and traceable chip design. Cycuity Radix software from Arteris uncovers vulnerabilities throughout the design process, enabling designers to address security risks, and helping to ensure robust protection for chips in advanced electronic systems.

## Key Products Attribute:

Arteris is a leading provider of semiconductor technology that accelerates the creation of high-performance, power-efficient silicon with built-in safety, reliability, and security. Innovative Arteris products are designed to optimize data movement and help ease complexity in the modern AI era with network-on-chip (NoC) interconnect intellectual property (IP), system-on-chip (SoC) software for integration automation and hardware security assurance. All are used by the world's top technology companies to improve overall performance and engineering productivity, reduce risk, lower costs, and bring cutting-edge designs to market faster.

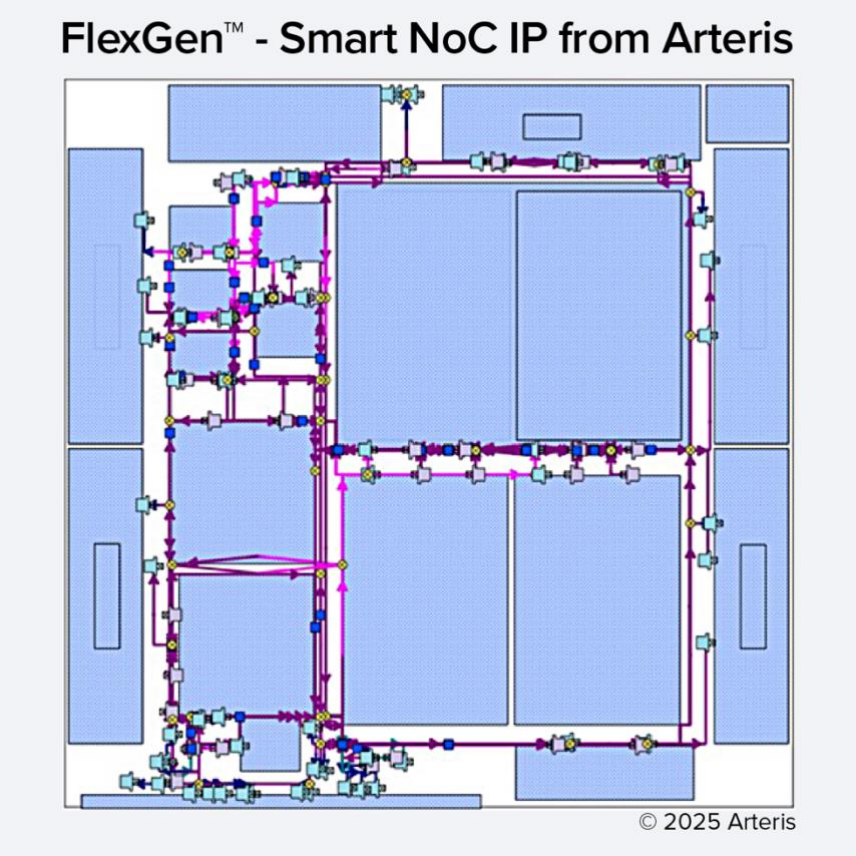

| Ashling | **Headquarters**:<br>Limerick, Ireland |
|---|---|
| | **Website:** https://www.ashling.com |

**About the Company:**
Ashling is a long-established provider of embedded development tools and services, with over 30 years of experience supporting embedded systems engineers. The company offers a comprehensive suite of tools including debug probes, trace solutions, SDKs, IDEs, compilers, and simulators, supporting a wide range of architectures including RISC-V, Arm, and others. Ashling is recognized for its early leadership in heterogeneous debugging for RISC-V-based systems.

| Beijing ESWIN Computing Technology Co., Ltd. | **Headquarters**:<br>Beijing, China |
|---|---|
| | **Website:** https://www.eswincomputing.com |

**About the Company:**
ESWIN Computing provides intelligent computing solutions focused on next-generation RISC-V architectures. The company develops system-level platforms combining domain-specific algorithms, IP modules, and integrated hardware-software ecosystems to support applications in smart devices and embodied intelligence. ESWIN aims to deliver scalable and competitive solutions for global markets.

| Beijing Institute of Open Source Chip | **Headquarters**:<br>Limerick, Ireland | **Founded**:<br>2021 |
|---|---|---|
| | **Website:** https://www.ashling.com | |

**About the Company:**
The Beijing Institute of Open Source Chip (BOSC) is a nonprofit organization established to advance open-source chip innovation, particularly around the XiangShan RISC-V processor project. Originating from research led by the Chinese Academy of Sciences, BOSC collaborates with industry partners to further develop high-performance open-source processor technologies and accelerate ecosystem adoption.

| Beijing Sophgo | **Headquarters**:<br>Beijing, China |
|---|---|
| | **Website:** https://www.sophgo.com |

**About the Company:**
Beijing Sophgo develops advanced computing solutions based on RISC-V and AI acceleration technologies. The company focuses on delivering high-performance processors and system platforms for intelligent computing applications.

| | | |
|---|---|---|
| **Breker Verification Systems** | **Headquarters**: San Jose, USA | **Founded**: 2004 |
| | **Website:** https://brekersystems.com | |

**About the Company:**

Breker Verification Systems solves complex semiconductor challenges across the functional verification process from streamlining UVM-based testbench composition to execution for IP block verification, significantly enhancing SoC integration and firmware verification with automated solutions that provide test content portability and reuse. Breker solutions easily layer into existing environments and operate across simulation, emulation and prototyping, and post-silicon execution platforms. Its Trek family is production-proven at leading semiconductor companies worldwide and enables design managers and verification engineers to realize measurable productivity gains, speed coverage closure and easy verification knowledge reuse. As a leader in the development of the Accellera Portable Stimulus Standard (PSS), privately held Breker has a reputation for dramatically reducing verification schedules in advanced development environments. Case studies that feature Altera (now Intel), Analog Devices, Broadcom, IBM and other companies leveraging Breker's solutions are available on the Breker website.

**Key Products Attribute:**

EDA company selling SystemVIP driving Test Suite Synthesis for IP and SoC verification. Particularly focused on RISC-V Core and SoC verification. Breker has sold products to most of the companies developing RISC-V cores, as well as many of the major semiconductor companies.

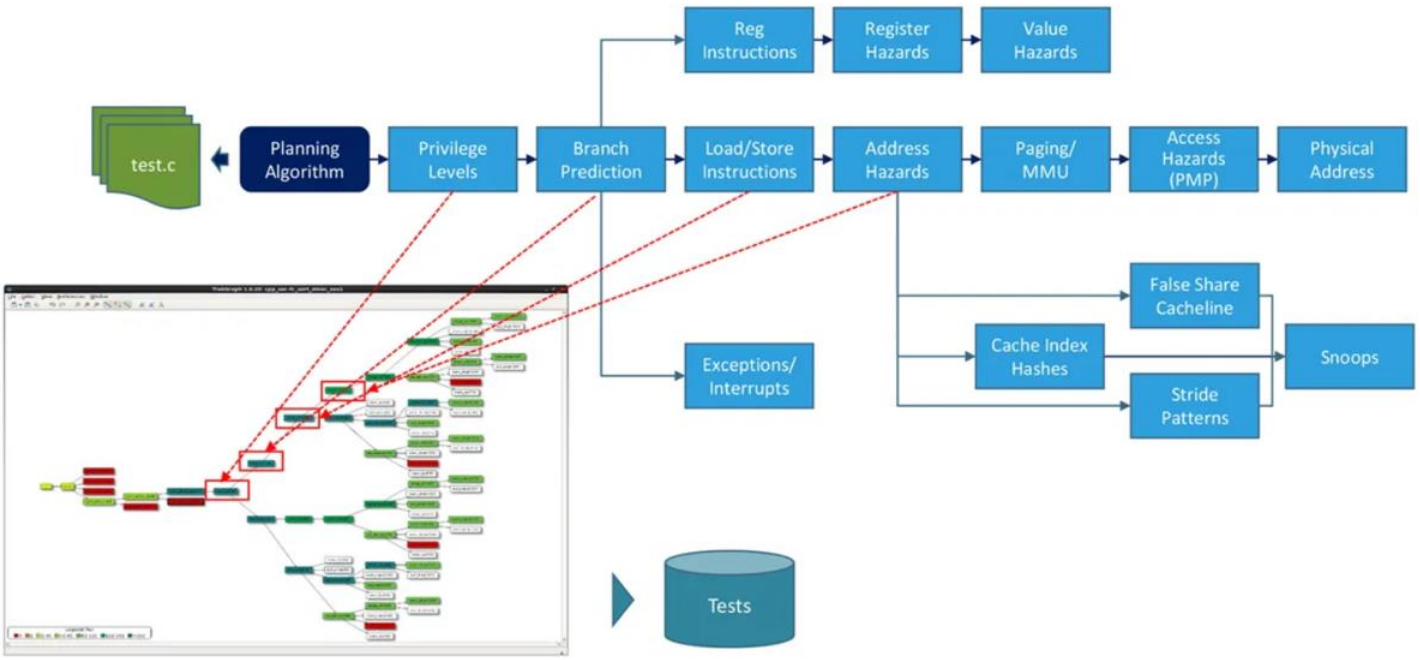

| Calsoft Labs Inc. dba ACL Digital, an Alten Group Company | **Headquarters**: San Jose, CA |
|---|---|
| | **Website:** https://www.acldigital.com |

**About the Company:**
ACL Digital is a global engineering services company providing expertise in semiconductor design, embedded systems, and digital transformation. The company supports industries including automotive, high-performance computing, AI/ML, and networking with end-to-end product engineering solutions.

| Cloud Native Computing Foundation (CNCF) | **Headquarters**: San Francisco, CA |
|---|---|
| | **Website:** https://wasmedge.org/ |

**About the Company:**
The Cloud Native Computing Foundation (CNCF) supports open-source technologies that enable cloud-native architectures. Through projects such as WasmEdge, CNCF contributes to runtime environments and platforms that support modern, scalable application development, including emerging RISC-V ecosystems.

| Codasip | **Headquarters**: Munich, Germany | **Founded**: 2014 |
|---|---|---|
| | **Website:** https://www.codasip.com | |

**About the Company:**
Codasip is a processor solutions company focused on enabling differentiation through customizable RISC-V processors. As a leading European RISC-V provider, Codasip delivers design tools and processor IP used in billions of chips, supporting applications across embedded, AI/ML, automotive, and IoT markets.

| Codeplay Software | **Headquarters**: Edinburgh, Scotland, UK | **Founded**: 2002 |
|---|---|---|
| | **Website:** https://www.codeplay.com | |

**About the Company:**
Codeplay develops software platforms that connect advanced processors with modern application development frameworks. Its ComputeSuite™ products enable developers to target AI accelerators using open standards, supporting industries such as automotive, industrial, and high-performance computing.

| Cortus | **Headquarters**:<br>France | **Founded**:<br>2005 |
|---|---|---|
| | **Website:** https://www.cortus.com | |

**About the Company:**
Cortus is a fabless semiconductor company designing microcontrollers and system-on-chip solutions using both proprietary and RISC-V architectures. Its portfolio includes digital, analog, RF, and security IP, supporting applications in automotive, avionics, and server markets.

| Daol Investment and Securities | **Headquarters**:<br>Seoul, South Korea |
|---|---|
| | **Website:** https://www.daolsecurities.com/top.jsp |

**About the Company:**
Daol Investment and Securities operates in capital markets, providing financial services and investment solutions. The company supports technology sectors including semiconductor and emerging computing ecosystems through strategic investment activities.

| DeepComputing | **Headquarters**:<br>Hong Kong, China | **Founded**:<br>2022 |
|---|---|---|
| | **Website:** https://deepcomputing.io | |

**About the Company:**
DeepComputing is focused on advancing the adoption and implementation of RISC-V technologies through innovative hardware platforms and ecosystem development. The company develops products such as laptops, tablets, and edge devices while collaborating with partners to expand RISC-V use cases across consumer, industrial, and research applications.

| Digital Core Design | **Headquarters**:<br>Bytom, Poland |
|---|---|
| | **Website:** https://www.dcd.pl/product-category/cpus/#riscv-2 |

**About the Company:**
Digital Core Design specializes in processor IP development, offering a portfolio of CPU cores including RISC-V-based designs. The company provides customizable IP solutions for embedded and industrial applications, supporting efficient and flexible system integration.

| E4 Computer Engineering | **Headquarters**: Scandiano, Italy | **Founded**: 2002 |
|---|---|---|
| | **Website:** https://www.e4company.com | |

**About the Company:**
E4 Computer Engineering is a provider of high-performance computing (HPC), AI, and data analytics solutions. The company collaborates with leading research institutions and enterprises, delivering advanced computing infrastructure and expertise for scientific and industrial applications.

| Emproof B.V. | **Headquarters**: Eindhoven, Netherlands | **Founded**: 2021 |
|---|---|---|
| | **Website:** https://www.emproof.com | |

**About the Company:**
Emproof provides advanced embedded software security solutions designed to protect intellectual property and ensure system integrity. Its platform enables protection against reverse engineering and exploitation attacks, supporting secure deployment across embedded and resource-constrained systems.

| E-Soft Technologies | **Headquarters**: Taipei City, Taiwan |
|---|---|
| | **Website:** https://www.esoft.com.tw/ |

**About the Company:**
E-Soft Technologies develops embedded software solutions supporting a range of hardware platforms and applications. The company provides tools and services that enable efficient development and deployment of embedded systems across industries.

| ETH Zurich | **Headquarters**: Zürich, Switzerland |
|---|---|
| | **Website:** https://pulp-platform.org/ |

**About the Company:**
ETH Zurich is a leading research institution contributing to open-source hardware and RISC-V innovation through initiatives such as the PULP platform. The institution develops advanced processor architectures and system designs, supporting academic and industrial collaboration.

| Google | **Headquarters**: Mountain View, CA | **Founded**: 2021 |
|---|---|---|
| | **Website:** https://go.dev/ | |

**About the Company:**
Google develops infrastructure, software platforms, and tools that support modern computing ecosystems. Through its work on programming languages, cloud infrastructure, and open technologies, Google contributes to scalable and high-performance computing environments, including those leveraging RISC-V architectures.

| Hexin Technology | **Headquarters**: Shanghai, China |
|---|---|
| | **Website:** https://shingroup.cn |

**About the Company:**
Hexin Technology focuses on the development of server-class processors and computing solutions. The company works on advancing high-performance CPU architectures to support data center and enterprise computing needs.

| Imagination Technologies | **Headquarters**: Kings Langley, England, UK | **Founded**: 2005 |
|---|---|---|
| | **Website:** https://imgtec.com/ | |

**About the Company:**
Imagination Technologies develops semiconductor IP solutions that enable advanced graphics, AI, and processing capabilities. The company partners with global technology leaders to deliver innovative solutions that power next-generation electronic devices.

| InCore Semiconductors | **Headquarters**: Chennai, India | **Founded**: 2018 |
|---|---|---|
| | **Website:** https://incoresemi.com/ | |

**About the Company:**
InCore Semiconductors is a fabless semiconductor company developing customizable RISC-V processor IP and SoC solutions. Built on expertise from the Shakti processor initiative, the company enables high-performance, low-power designs for embedded and edge computing applications.

| Lattice Semiconductor | **Headquarters**:<br>Portland, OR | **Founded**:<br>1983 |
|---|---|---|
| | **Website:** https://www.latticesemi.com | |

**About the Company:**
Lattice Semiconductor is a leading provider of low-power programmable solutions, enabling innovation across communications, computing, industrial, automotive, and consumer markets. The company's technologies support edge-to-cloud applications with a focus on efficiency, security, and rapid deployment.

| Lauterbach | **Headquarters**:<br>Hoehenkirchen, Germany | **Founded**:<br>1979 |
|---|---|---|
| | **Website:** https://www.lauterbach.com/ | |

**About the Company:**
Lauterbach is a global leader in development tools for embedded systems, offering advanced debugging and trace solutions through its TRACE32® platform. With over four decades of experience, the company supports semiconductor manufacturers and system developers worldwide across a wide range of architectures, including RISC-V.

| Mega Securities | **Headquarters**:<br>Taipei, Taiwan |
|---|---|
| | **Website:** https://www.emega.com.tw/index.do |

**About the Company:**
Mega Securities provides financial services and investment solutions within capital markets. The company supports technology and innovation sectors through strategic financial services and market expertise.

| Microchip Technology | **Headquarters**:<br>Chandler, AZ | **Founded**:<br>1989 |
|---|---|---|
| | **Website:** https://microchip.com | |

**About the Company:**
Microchip Technology is a leading provider of embedded control solutions, offering a comprehensive portfolio of microcontrollers, analog devices, and connectivity solutions. The company supports a wide range of industries including automotive, industrial, aerospace, and IoT, enabling efficient and secure system design.

| MINRES Technologies GmbH | **Headquarters**: Neubiberg, Germany | **Founded**: 2019 |
|---|---|---|
| | **Website:** https://www.minres.com | |

**About the Company:**
MINRES Technologies provides consulting and engineering services focused on embedded software development and virtual platform technologies. The company helps organizations optimize development processes and leverage virtual platforms to improve system design and validation.

| MIPS | **Headquarters**: San Jose, CA | **Founded**: 1989 |
|---|---|---|
| | **Website:** https://mips.com | |

**About the Company:**
MIPS develops scalable processor IP for automotive, data center, and communications applications. Building on decades of architectural expertise, the company is advancing RISC-V-based solutions for high-performance and heterogeneous computing environments.

| Morphing Machines | **Headquarters**: Bangalore, India |
|---|---|
| | **Website:** https://www.morphing.in/ |

**About the Company:**
Morphing a closely held fabless semiconductor IP products and solutions company launched through the Technology Entrepreneurship initiative of the Indian Institute of Science at Bangalore. Our focus is on innovations centered around **REDEFINE™**, a runtime reconfigurable many-core processor. We can enable customers and businesses around the world to build energy efficient high-performance systems and solutions. Our semiconductor IP, systems, and solutions can be licensed to key customers around the world.

| MosChip Technologies | **Headquarters**: India / USA | **Founded**: 1999 |
|---|---|---|
| | **Website:** https://www.moschip.com | |

**About the Company:**

MosChip Technologies is a global semiconductor and product engineering company offering custom silicon solutions, embedded systems, and AI accelerators. They operate in two main segments: products and semiconductor engineering which include ASIC platforms, mixed-signal IP, OSAT solutions and design verification. They cover hardware and embedded systems, multimedia and vision, mobility, connectivity and automation engineering. Also, they provide AI/ML engineering and solutions in computer vision, generative AI, cognitive computing, and AI/FPGA acceleration.

**Key Products Attribute:**

Their products are characterized by high scalability, modular design, and industry-specific optimization, supporting applications in automotive, aerospace, healthcare, and consumer electronics. With strong ecosystem integration and advanced AI frameworks, MosChip enables intelligent, connected, and adaptive devices across diverse markets. They serve industries like automotive, aerospace, healthcare, and consumer electronics, and have global partnerships to support advanced chip development and intelligent product ecosystems.

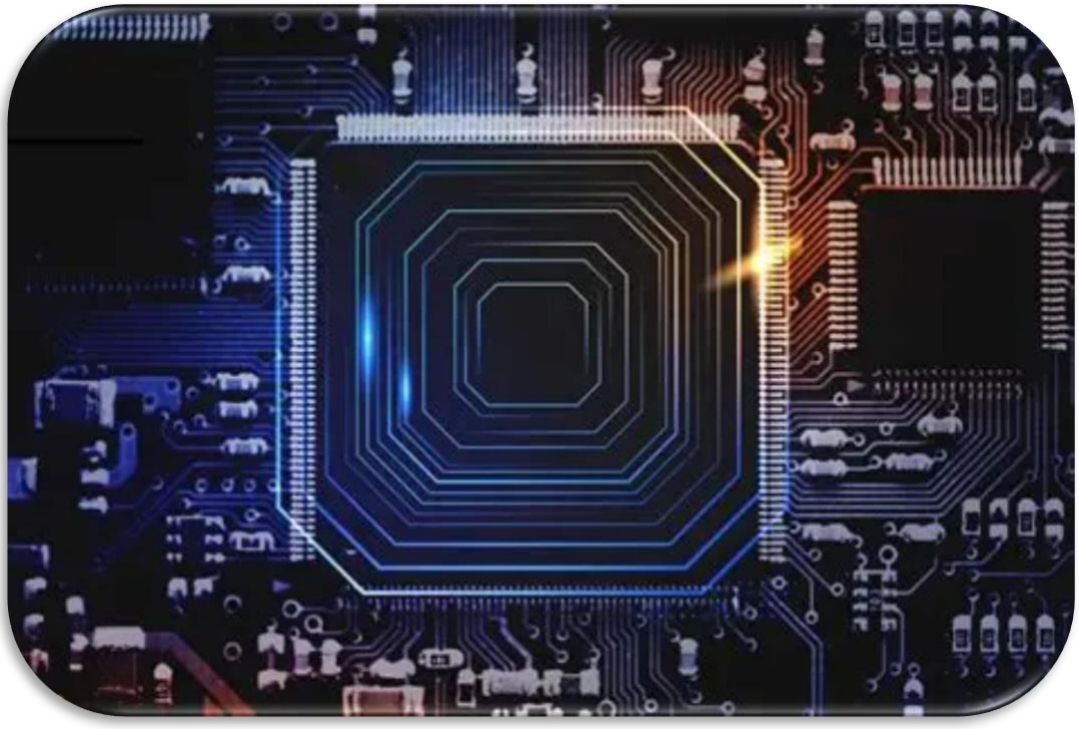

| Nuclei System Technology Co., Ltd | **Headquarters**: Shanghai, China | **Founded**: 2018 |
|---|---|---|
| | **Website:** https://www.nucleisys.com/product.php | |

**About the Company:**

Nuclei System Technology a Chinese company focused on RISC-V processor IP and chip solutions. It provides a comprehensive portfolio of RISC-V CPU IP cores and SoC subsystem solutions, enabling cost-effective and customizable designs for various applications. Nuclei's offerings include flexible, high-performance processors supporting security, functional safety (ISO26262), and AI acceleration features. The company also delivers a complete ecosystem with toolchains, SDKs, RTOS/Linux support, and localized R&D to meet diverse market needs. Its technology powers products in AIoT, automotive, industrial, and consumer electronics sectors, and Nuclei has partnered with leading chipmakers to bring RISC-V-based MCUs and AI solutions to mass production.

**Key Products Attribute:**

RVA23 with Hypervisor, IOMMU and AIA; ASIL-D; Vector; SMP; DSP; Low-power, etc.

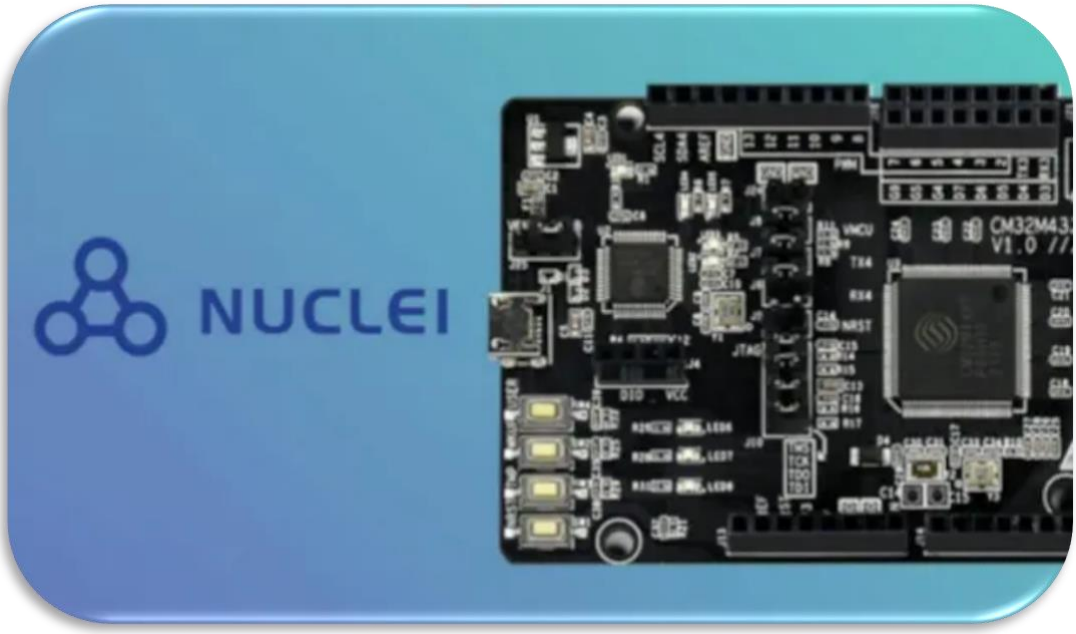

| OpenHW Group | **Headquarters**:<br>Ottawa, Canada |
|---|---|
| | **Website:** https://www.openhwgroup.org/ |

**About the Company:**

OpenHW Group is a global, not-for-profit organization that collaborates on the development of open-source hardware based on the RISC-V architecture. The group brings together industry and academia to create high-quality processor cores, tools, and verification environments that accelerate adoption of open hardware.

| proteanTecs | **Headquarters**:<br>Haifa, Israel | **Founded**:<br>2017 |
|---|---|---|
| | **Website:** https://www.proteantecs.com | |

**About the Company:**

proteanTecs is a deep-tech semiconductor analytics company founded in 2017 in Haifa, Israel. The company develops cloud-based lifecycle analytics software that uses on-chip telemetry ("Agents") and machine learning to provide unprecedented visibility into chip health, reliability, and performance throughout manufacturing, deployment, and operation. Its solutions are widely adopted in datacenter, automotive, communications, and mobile sectors, enabling predictive maintenance, early fault detection, and long-term reliability improvements for advanced electronic systems. proteanTecs is headquartered in Haifa, Israel, with additional offices in the United States, India, South Korea, and Taiwan, serving a global customer base across critical electronics markets.

**Key Products Attribute:**


proteanTecs' core product offering is a cloud-based lifecycle analytics platform for chips, designed to continuously monitor semiconductor health and performance throughout production, deployment, and field operation. Its technology relies on on-chip monitors ("Agents") that generate novel internal telemetry data, which is then analyzed using advanced machine learning to deliver deep visibility into system behavior, predictive failure alerts, and reliability insights far beyond traditional test methods. Key attributes of the platform include lifetime product tracking, static timing analysis, customer usage tracking, predictive maintenance, and early fault detection, enabling proactive quality assurance and improved operational efficiency in data centers, automotive, communications, and mobile markets. This combination of embedded chip-level sensing and cloud-scale analytics—referred to as deep data analytics—is the company's differentiating capability, offering unparalleled visibility and enabling electronics to self-report actionable health metrics for enhanced reliability and long-term performance management.

| Qamcom | **Headquarters**: Gothenburg, Sweden | **Founded**: 2001 |
|---|---|---|
| | **Website:** https://www.qamcom.com/ | |

**About the Company:**
Qamcom is a research and technology company specializing in system development across hardware, software, and signal processing domains. The company delivers solutions in areas such as wireless communication, industrial AI, and IoT, helping bridge the gap between advanced technology and real-world applications.

| Renesas Electronics Corporation | **Headquarters**: Tokyo, Japan |
|---|---|
| | **Website:** https://www.renesas.com |

**About the Company:**
Renesas Electronics is a global leader in microcontrollers, analog, and power devices, serving industries including automotive, industrial, IoT, and data center infrastructure. The company delivers scalable and reliable semiconductor solutions that support billions of devices worldwide.

| RISC-V Alliance Japan | **Headquarters**: Tokyo, Japan | **Founded**: 2017 |
|---|---|---|
| | **Website:** https://riscv.or.jp | |

**About the Company:**
RISC-V Alliance Japan is a nonprofit organization dedicated to promoting RISC-V technologies and open silicon innovation across Japan and neighboring regions. The alliance collaborates with global partners to advance ecosystem development and adoption.

| RiVAI Technologies (Shenzhen) Co. | **Headquarters**: Shenzhen, China | **Founded**: 2018 |
|---|---|---|
| | **Website:** https://rivai-ic.com.cn/ | |

**About the Company:**
RiVAI Technologies develops high-performance RISC-V processor IP and custom silicon solutions. Founded by members of the original UC Berkeley RISC-V team, the company focuses on scalable CPU architectures for applications ranging from edge devices to data centers.

| Semidynamics | **Headquarters**:<br>Barcelona, Spain | **Founded**:<br>2018 |
|---|---|---|
| | **Website:** https://www.semidynamics.com/ | |

**About the Company:**

Semidynamics develops high-performance processor IP solutions with a focus on customizable RISC-V architectures. The company enables advanced computing designs through flexible and scalable processor technologies tailored to modern workloads.

| Siemens | **Headquarters**:<br>Wilsonville, OR |
|---|---|
| | **Website:** https://eda.sw.siemens.com/en-US/ic/tessent/embedded-analytics/ |

**About the Company:**

Siemens provides electronic design automation (EDA) solutions that support advanced semiconductor development. Its Tessent Embedded Analytics platform enables real-time monitoring, validation, and optimization of SoC performance, including support for RISC-V trace standards.

| SiFive | **Headquarters**: Santa Clara, USA | **Founded**: 2015 |
|---|---|---|
| | **Website:** https://www.sifive.com | |

## About the Company:

SiFive is a leading pioneer in the RISC-V revolution, providing customizable, high-performance processor IP solutions based on the open-standard RISC-V architecture. The company enables customers to design tailored chips for applications across AI/ML, automotive, IoT, data centers, storage, and consumer electronics. SiFive's portfolio includes scalar, vector, and matrix compute cores optimized for scalability, power efficiency, and rapid time-to-market. With over 400 design wins and 2+ billion devices powered by SiFive IP, the company drives innovation through open architecture, flexible customization, and global partnerships, empowering engineers to break free from proprietary architectures and accelerate technological advancement.

## Key Products Attribute:

Low power, flexible, high performance, superior compute density.

| StarFive Technology | **Headquarters**:<br>Shanghai, China | **Founded**:<br>2018 |
|---|---|---|
| | **Website:** https://www.starfivetech.com/en/site/boards | |

**About the Company:**

StarFive Technology develops RISC-V-based processors, chips, and development platforms, including single-board computers. The company focuses on enabling accessible and scalable hardware solutions for developers and commercial applications.

| Syntacore | **Headquarters**:<br>Moscow, Russian Federation | **Founded**:<br>2018 |
|---|---|---|
| | **Website:** http://www.syntacore.com/ | |

**About the Company:**

Syntacore provides RISC-V processor IP and development tools designed for embedded and industrial applications. The company supports customizable processor solutions and software ecosystems for efficient system development.

| Tenstorrent | **Headquarters**: Toronto, Canada | **Founded**: 2016 |
|---|---|---|
| | **Website:** https://tenstorrent.com | |

**About the Company:**

Tenstorrent is a next-generation computing company that builds computers for AI, bringing together experts in computer architecture, ASIC design, advanced systems, RISC-V technology, and neural-network compilers to create the next wave of high-performance AI hardware and software solutions

Headquartered in the United States, with offices across Austin, Silicon Valley, Toronto, Belgrade, Seoul, Tokyo, and Bangalore, the company brings together deep expertise in computer architecture, ASIC design, RISC-V CPU technology, and neural-network compilers.

**Key Products Attribute:**

Tenstorrent's products are built specifically for advanced AI computing, centered on its AI graph processors, which are designed to run complex neural-network graphs efficiently at scale. The company also develops high-performance RISC-V CPUs and configurable chiplets, giving customers modular, flexible building blocks for creating custom AI and compute solutions. Their hardware is supported by a fully open-source software stack, including the MLIR-based TT-Forge compiler, enabling developers to customize, adapt, and optimize models for Tenstorrent systems with full transparency and control. Overall, Tenstorrent's product philosophy emphasizes scalability, openness, and performance, enabling deployments from single accelerator cards to full AI clusters.

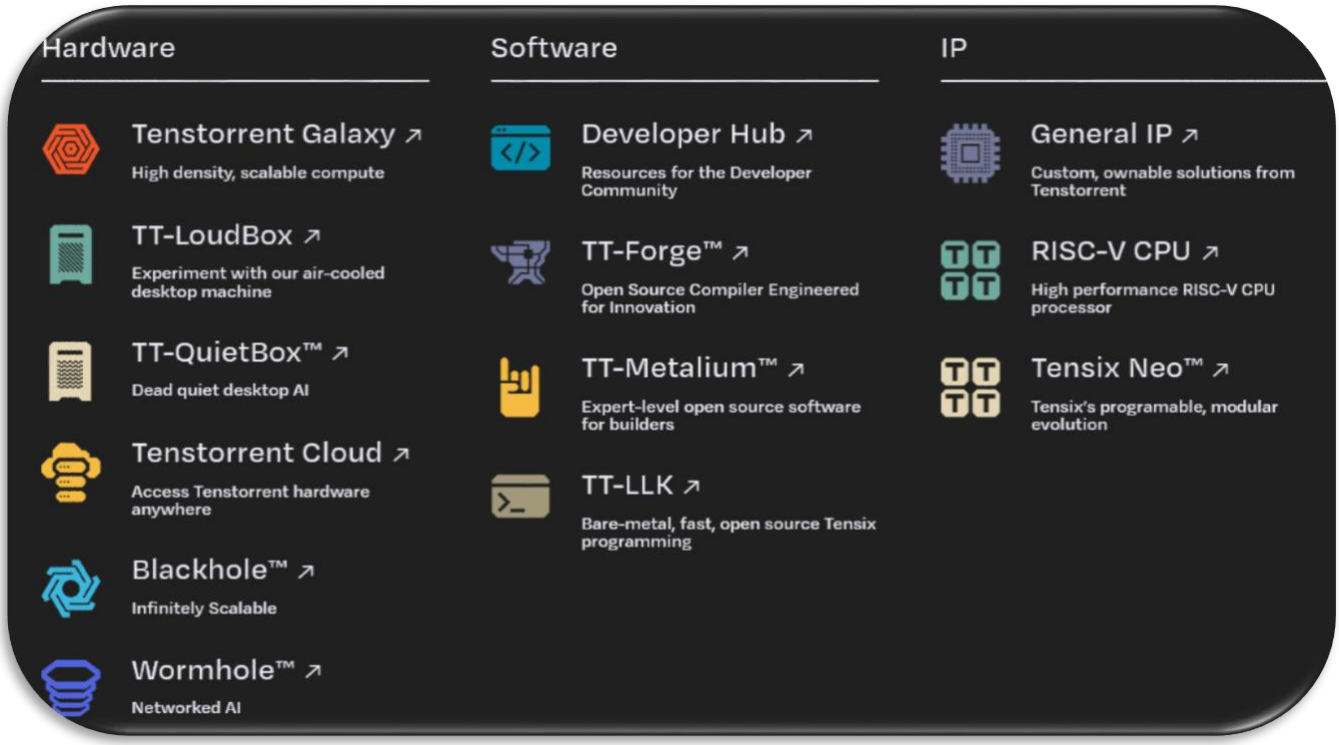

| The Edge Marketing | **Headquarters**:<br>San Jose, CA |
|---|---|

**About the Company:**
The Edge Marketing provides strategic marketing communications, media relations, and content creation services. The company supports technology organizations in building visibility and engagement across digital and industry platforms.

| University of Bremen | **Headquarters**:<br>Bremen, Germany |
|---|---|
| | **Website:** https://agra.informatik.uni-bremen.de/projects/risc-v/ |

**About the Company:**
The University of Bremen contributes to RISC-V research and development through academic initiatives focused on processor architecture and system design. The institution supports innovation through collaborative research and open-source projects.

| ViShare Technology Limited | **Headquarters**:<br>Hong Kong, China |
|---|---|
| | **Website:** https://www.visharetech.com |

**About the Company:**
ViShare Technology develops solutions for low-latency video streaming and high-performance computing architectures. The company focuses on heterogeneous multi-core systems and advanced accelerator technologies for real-time applications.

| Vyoma Systems Private Limited | **Headquarters**:<br>Chennai, India |
|---|---|
| | **Website:** https://vyomasystems.com |

**About the Company**
Vyoma Systems develops solutions focused on semiconductor verification and system-level validation. The company supports design and testing processes for advanced chip development.

| Win Source Electronic Technology Ltd. | **Headquarters**: Shenzhen, Guangdong, China | **Founded**: 1999 |
|---|---|---|
| | **Website:** https://www.win-source.net/ | |

**About the Company:**
Win Source is a global electronic components distributor offering a wide range of products and procurement services. The company provides efficient sourcing solutions supported by real-time inventory access and value-added services.

| Wind River | **Headquarters**: Alameda, CA, USA |
|---|---|
| | **Website:** https://www.windriver.com/products/vxworks |

**About the Company:**
Wind River develops software platforms for intelligent edge systems, including real-time operating systems and cloud-native infrastructure. The company supports mission-critical applications across industries such as aerospace, automotive, and industrial systems.

| XMOS | **Headquarters**: Bristol, England, UK |
|---|---|
| | **Website:** https://xmos.ai/ |

**About the Company:**
XMOS develops semiconductor solutions focused on audio, control, and edge AI processing. The company provides processors and software tools that enable real-time, deterministic performance in embedded systems.

| ZTE | **Headquarters**: China | **Founded**: 2018 |
|---|---|---|
| | **Website:** https://www.zte.com.cn | |

**About the company:**
ZTE is a global provider of telecommunications and information technology solutions, offering products and services across wireless networks, devices, and digital infrastructure. The company supports large-scale connectivity and computing systems worldwide.

# XIII. Conclusions and Recommendations

As system complexity continues to escalate across all semiconductor sectors, the EDA and IP industries have responded with another wave of innovation - developing increasingly intelligent, automated, and integrated design solutions. This ongoing evolution has given SoC architects an unprecedented level of creative control, allowing them to design silicon that is more specialized, efficient, and powerful than at any previous point in semiconductor history.

The third - party IP ecosystem remains a central driver of this innovation, continuously advancing interconnects, processors, memory subsystems, high - speed communications channels and verification technologies in step with the industry's accelerating demands. The emergence of the RISC - V instruction set architecture represents both a continuation and a redefinition of that process. RISC - V has matured from an academic initiative into a global design movement - open, extensible, and now fully supported by the mainstream EDA and software toolchain providers.

Its timing could hardly be better. The rapid rise of AI, edge intelligence, and chiplet - based architectures has created a once - in - a - generation opportunity for open - standard compute. RISC - V designs are now scaling alongside AI workloads, offering flexible cores that can be tightly coupled to accelerators, DSPs, and domain - specific engines. This combination of openness and adaptability positions RISC - V to capture the same type of market momentum that Arm achieved during the mobile era, though now in a more heterogeneous and diversified landscape.

In addition, the public announcement by Nvidia at the 2024 North American RISC - V Summit that they are using multiple RISC - V CPU cores in every one of their GPUs and CPUs reinforces the validity and utility of the RISC - V ISA in the marketplace and potentially removes a major hurdle for adoption moving forward: that RISC - V is only acceptable for low complexity silicon solutions. No one can make that argument now since GPUs are one of the most high - complexity parts available and demand high - performance CPU cores of every type. This has given a great boost to the RISC-V ISA and has built confidence in the minds of designers that they are making the right choice in selecting the RISC-V ISA for their design.

The SHD Group views RISC - V not only as an enabling CPU architecture but as a foundation for broader ecosystem expansion - encompassing IP subsystems, chiplets, and vertically integrated design frameworks. The result is an environment primed for continuous innovation, competitive diversity, and global participation. In 2025 and beyond, RISC - V stands as one of the most significant catalysts for growth in the semiconductor industry's next phase of architectural evolution.

# XIV. Appendix

## Report Methodology

Data sources were a combination of public and private data gathered by company briefings and interviews.

The SHD Group has conducted both primary and secondary research to compile this report. We have also completed 50+ in-person interviews with companies prominent in the RISC-V ecosystem to get their perspectives on how they view the market today and into the future. We have kept responses by interviewees anonymous, if so requested.

## SoC Device Definitions

This report looks at 10 separate SoC devices in each application we are profiling. The following table gives this listing, and the following text defines these devices as they apply to the report.

Table 34: SoC Definitions for Devices Profiled in Each Application

| | |
|---|---|
| **CPU - General Purpose** | **FPGA** |
| **Host Processor** | **DSP** |
| **GPU** | **Network Device** |
| **RISC-V Smart Sensor** | **Storage Controller** |
| **Security** | **AI Accelerator** |
| **Microcontroller** | |

**Source: The SHD Group, April 2026**

**CPU - General Purpose:**

A General Purpose CPU is a central processing unit designed to handle a wide variety of computing tasks. It is the core component in traditional computing devices such as desktops, laptops, and servers. This type of CPU is characterized by its versatility, ability to efficiently run operating systems, perform complex computations, and execute a broad range of software applications. It is not specialized for any specific type of application but is built to be flexible and adaptable to various computing needs. Typical operating systems would be versions of Linux or similar.

**App Processor:**

In this report, we are defining application processor as the host processor developed for primary use in cell phones and tablet devices. As such, it is a type of CPU specifically tailored to manage the unique demands of mobile/portable applications. While it shares fundamental characteristics with general purpose CPUs, including the ability to run a variety of software, it is optimized for the performance, power efficiency, and space constraints inherent in mobile devices. Application processors handle the computing tasks related to user interface, multimedia processing, connectivity, and other application-specific functions, making them distinct in their design and

use from more generalized CPUs. App processors would typically host an operating system such as IOS or Android.

**GPU:**

A Graphics Processing Unit is specialized for rendering graphics and parallel processing. Used in graphics-intensive applications, including gaming and visual simulations.

**Security:**

Semiconductor device dedicated to ensuring the security of data and systems. Includes components like encryption modules and secure key storage. Use case examples would include hardware-based security, secure boot processes, or cryptographic operations.

**AI Accelerator:**

Accelerates artificial intelligence computations, enhancing performance for AI-related tasks. Commonly used in machine learning and deep learning applications. Examples include Google's TPU, NVIDIA's Tensor Cores, Meta's RISC-V-based MTIA, and many new products under development.

**FPGA:**

Field-Programmable Gate Array is a reconfigurable semiconductor device. It can be customized for specific applications after manufacturing, offering flexibility in hardware design.

**DSP:**

Digital Signal Processor is optimized for processing digital signals, such as audio or video. Used in applications requiring real-time signal processing, such as audio signal processing, radar systems, and even mobile phones.

**Network Device:**

Semiconductor device facilitating communication within a network. This could be either for wired or wireless networks. Examples would include network interface cards or wireless communication chips.

**RISC-V Smart Sensor**

A RISC-V smart sensor combines a sensor element with an embedded RISC-V processor to perform local signal processing, control, and basic analytics directly at the point of data capture. By handling tasks such as filtering, feature extraction, or event detection on-device, it reduces power consumption, latency, and data bandwidth compared with streaming raw sensor data to a central processor.

**Storage Controller:**

Manages data storage devices, controlling data flow between a computer and storage media like hard drives or SSDs with optimization for read/write performance and data integrity.

**Microcontroller:**

Compact integrated circuit containing a processor core, memory and programmable input/output peripherals. Used in embedded systems for specific control functions. Microcontrollers might run a real-time operating system or could run software in a "base metal" mode.

These are generally accepted definitions for these devices within the semiconductor industry today. There can be derivations of these devices based upon special use cases within different applications. In this report, we are adhering to these definitions wherever possible, with exceptions noted as they occur.

## Glossary of Terms

**AI Accelerator:**
Dedicated compute hardware optimized for matrix, tensor, and inferencing operations used in modern AI workflows. These engines increasingly anchor data-center and edge inference platforms by offloading high-intensity neural network operations from the host CPU.

**Application Processor:**
The primary compute engine in mobile systems and other consumer platforms, responsible for executing the operating system, user applications, and high-level workloads such as multimedia, AI inference, and secure processing.

**ASSP (Application-Specific Standard Product):**
A semiconductor device built for a well-defined application segment and sold broadly, rather than being custom-built for a single customer. ASSPs often incorporate RISC-V cores for domain-specific control functions.

**CAGR (Compound Annual Growth Rate):**
A normalized annualized growth metric that smooths out year-to-year variability by assuming compounding across the measured period.

**Chiplets:**
Modular silicon blocks that can be interconnected to build heterogeneous multi-die SoCs. The chiplet model is emerging as a strong insertion point for RISC-V due to its configurability, small footprint, and suitability for tightly scoped compute and control roles.

**Co-processor:**
A supplementary compute unit that executes specialized or accelerated operations alongside a primary CPU, often improving determinism, throughput, or power efficiency.

**Custom Extensions:**
Instruction-set or microarchitectural additions layered onto standard RISC-V cores to support domain-optimized interfaces, accelerators, or datapaths. These extensions are increasingly used to tailor RISC-V implementations for AI, storage, and networking workloads.

**Data Center – AI Accelerator:**
Large-scale, high-throughput compute engines designed specifically to enhance machine-learning performance in hyperscale environments.

**Deeply Embedded:**
Logic elements or microcontrollers that reside within the functional core of a system, typically implementing state machines, low-latency control loops, or protocol management functions with fixed or narrowly scoped behaviors.

**DIY (Home-Grown RISC-V):**
Internally developed RISC-V CPU cores created by companies seeking to optimize performance, security, or power for proprietary workloads rather than relying solely on commercial IP vendors.

**Edge Devices:**
Compute endpoints positioned near the data source—such as industrial nodes, cameras, or consumer electronics—where low-latency processing and power efficiency are critical. RISC-V is broadly adopted in these systems due to its configurability and small footprint.

**Finite State Machine (FSM):**
A design construct used to implement structured, predictable transitions between operational states—fundamental to deeply embedded control logic.

**HPC (High-Performance Computing):**
A computation modality leveraging parallelism and advanced architectures to support simulation, modeling, and scientific workloads. Several emerging RISC-V initiatives are targeted at HPC-class vector and accelerator designs.

**Hyperscale Data Centers:**
Massive compute infrastructures engineered for horizontal scaling, capable of serving large AI, cloud, and storage workloads. RISC-V is increasingly being evaluated as a control-plane compute element in these deployments.

**IP (Intellectual Property) Market:**
The commercial ecosystem providing reusable blocks of silicon design—CPU cores, interfaces, accelerators, memory subsystems—used as foundational elements for SoC development.

**License-Driven Revenue:**
Upfront and recurring income associated with the right to integrate IP blocks into new SoC designs. This revenue dominates early-stage markets before volume production drives substantial royalties.

**Linux Debian and Yocto:**
Linux distributions widely used on RISC-V platforms. Debian provides a general-purpose environment, while Yocto enables customizable, production-grade embedded Linux builds.

**LLMs (Large Language Models):**
Neural network architectures trained on large corpora, capable of text synthesis, inference, and reasoning. These models are shaping requirements for high-performance, energy-efficient SoC designs, including RISC-V vector and AI-accelerated architectures.

**Microcontroller (MCU):**
A compact processing subsystem that integrates CPU, memory, and I/O into a single low-power device. MCUs have been a leading early-adoption category for RISC-V.

**MTIA (Meta Training and Inference Accelerator):**
A compute platform from Meta designed for machine-learning workloads, now incorporating RISC-V for control-path and subsystem management.

**Phison X1 Enterprise Storage:**
A storage controller architecture integrating RISC-V compute elements to support next-generation SSD performance for AI, HPC, and hyperscale deployments.

**Renesas Voice-Control ASSP Solution:**
A consumer-facing semiconductor product using RISC-V for speech-recognition and voice command processing.

**RISC-V SoC Revenue Growth:**
The expanding contribution of RISC-V-based SoCs to regional semiconductor markets—including China, the broader Asia-Pacific region, and the Americas—where adoption is accelerating across consumer, industrial, and communications sectors.

**RISC-V:**
An open and extensible instruction-set architecture enabling architectural transparency, multi-vendor competition, and domain customization across the semiconductor industry.

**Royalty-Driven Revenue:**
Income derived from per-unit shipments of SoCs integrating licensed IP. This revenue becomes dominant once RISC-V-based silicon reaches sustained volume production.

**System-on-Chip (SoC):**
A complete electronics system integrated onto a single die or package, incorporating CPUs, accelerators, memory controllers, and interface logic.

**SoC Revenues on a Regional Basis:**
Geographic segmentation of SoC revenue streams, highlighting the accelerating penetration of RISC-V in regions with high system design activity.

**Share of Market (SOM):**
A measure of the relative market capture of a given company or architecture, calculated as its portion of total market revenue or unit shipments.

**Smart RISC-V Sensors (New Definition):**
Sensor devices that integrate a small RISC-V core with on-sensor compute to enable local preprocessing, noise filtering, event detection, or lightweight AI inference before data is sent upstream. These devices reduce system power, bandwidth demand, and latency, and represent one of the fastest-growing early-adoption domains for RISC-V.

**SSD/Flash:**
High-speed storage technologies used in computing, cloud, and consumer electronics, where RISC-V increasingly appears in controller functions.

# XV. Acknowledgements

The SHD Group wishes to thank the following companies who have either contributed direction to the insights and data points in this report, agreed to be interviewed, or advised us in the overall report construction.

| | |
|---|---|
| Ahead Computing | MediaTek |
| Ainekko | Mega Securities |
| AION Silicon | Menta |
| Akeana, Inc. | Microchip |
| Alibaba DAMO Academy | MIPS (Global Foundries) |
| Andes Technology Corporation | MosChip Technologies Ltd. |
| Apple | Nanoveu |
| Arteris IP | Nuclei System Technology |
| Ashling Systems Corporation | Nvidia Corporation |
| Axelera AI | Open Hardware Foundation |
| Baya Systems, Inc. | proteanTecs |
| Beijing Institute of Open-Source Chip | Qualcomm |
| BrainChip Holdings Ltd. | Red Hat |
| Breker Verification Systems | Red Semiconductor |
| Cadence Design Systems | Renesas |
| Canonical | Samsung Electronics Co., Ltd. |
| Codasip | Semidynamics |
| Centre for Development of Advanced Computing | Siemens EDA |
| Deep Computing | SiFive |
| Electronics and Telecommunications Research Institute | Sophgo |
| Eliyan Labs | SUSE |
| Eswin Computing | Synopsys (Global Foundries) |
| Imagination Technologies | Tenstorrent Holdings Inc. |
| Innatera | Ventana Micro Systems (Qualcomm) |
| Inspire Semiconductors | VeriSilicon |
| Intel Corporation | VRULL GmbH |

# XVI. Addendum Spreadsheet TOC

The following Table of Contents shows tables and figures that are available in an accompanying spreadsheet, included with the unabridged version of this report. These tables and figures provide a finer granularity to the data contained in the abridged and the full versions of this report. In total, there are 338 tables and 298 figures in this spreadsheet. Please contact us at info@theshdgroup.com to get a copy.

## 2026 RISC-V Market Analysis Tables and Figures

| Tables | Figures |
|---|---|
| RISC-V SoC Unit Shipment Actuals and Forecast | RISC-V SoC Unit Shipment Actuals and Forecast |
| Rising Device Complexity | SoC Definitions by IP Content |
| TSMC Metrics for Current and Advanced Process Geometries | Average Gate Complexity |
| Transistor Budget by Process Geometry 28nm–1.4nm | Transistor Budget 1995–2030 |
| Average Number of IP blocks in SoC Designs 2024–2031 | Average Number of IP Blocks by SoC Category, 2021 - 2031 |
| AI Functional Levels and RISC-V Hardware Tiers | Heat Map Diagram of AI Functional Levels and RISC-V Hardware Tiers |
| Market Challenges and Issues – 2025 Outlook | Performance per Watt Scaling vs. Transistor Growth 2010 - 2031 |
| | Rise of RISC-V in Selected Markets 2021 - 2031 |

## Application Table

List of Market Segments Covered and Analyzed

## Industrial Segment

Industrial RISC-V SoC Unit Shipment Tables
Industrial RISC-V SoC Unit Revenue Tables
Industrial RISC-V SoC Unit Shipment Figures
Industrial RISC-V SoC Revenue Tables
Industrial RISC-V SoC Revenue Figures

## Automotive Segment

Automotive RISC-V SoC Unit Shipment Tables
Automotive RISC-V SoC Revenue Tables
Automotive RISC-V SoC Unit Shipment Figures
Automotive RISC-V SoC Revenue Figures

## Networking Segment

Networking RISC-V SoC Unit Shipment Tables
Networking RISC-V SoC Revenue Tables
Networking RISC-V SoC Unit Shipment Figures
Networking RISC-V SoC Revenue Figures

## Computer Segment

Computer RISC-V SoC Unit Shipment Tables
Computer RISC-V SoC Revenue Tables
Computer RISC-V SoC Unit Shipment Figures
Computer RISC-V SoC Revenue Figures
Consumer RISC-V SoC Unit Shipment Tables
Consumer RISC-V SoC Revenue Tables
Consumer RISC-V SoC Unit Shipment Figures
Consumer RISC-V SoC Revenue Figures

## Other Segment

Other RISC-V SoC Unit Shipment Table
Other RISC-V SoC Revenue Table
Other RISC-V SoC Unit Shipment Figure
Other RISC-V SoC Revenue Figure

## Total RISC-V SoC Market

Total RISC-V SoC Unit Shipment Table
Total RISC-V Market SoC Revenue Table
Total RISC-V Market SoC Unit Shipment Figure
Total RISC-V Market SoC Revenue Figure

## RISC-V Use Cases: FSM, MCU, Coprocessor, Host CPU

Industrial & Automotive Applications M Units and Revenues 2021 - 2031
Networking & Computing Applications M Units and Revenues 2021 - 2031
Consumer & Other Applications M Units and Revenues 2021 – 2031

## RISC-V Function Totals

RISC-V Function Shipment Tables
RISC-V Function Revenue Tables
RISC-V Function Unit Shipment Figures
RISC-V Function Revenue Figures

## RISC-V SoC Accelerators

RISC-V SoC Accelerator Unit Shipment Tables
RISC-V SoC Accelerator Revenue Tables
RISC-V SoC Accelerator Unit Shipment Figures
RISC-V SoC Accelerator Revenue Figures

## RISC-V SoC Market Penetration

RISC-V SoC Market Penetration by Device Type for Unit Shipment Tables
RISC-V SoC Market Penetration by Device Type for Revenue Tables
RISC-V SoC Market Penetration by Device Type for Unit Shipment Figures
RISC-V SoC Market Penetration by Device Type for Revenue Figures

## RISC-V SoC Totals by Device Type by Industry

RISC-V SoC Unit Shipments by Device type by Industry Tables 2021 - 2031
RISC-V SoC Revenues by Device type by Industry Tables 2021 - 2031
RISC-V SoC Unit Shipments by Device type by Industry Figures 2021 - 2031
RISC-V SoC Revenues by Device type by Industry Figures 2021 - 2031

## RISC-V CPU IP Market 2021 - 2031

IP Market and RISC-V CPU IP Market Revenue Tables 2021 – 2031
IP Market and RISC-V CPU IP Market Revenue Figures 2021 – 2031

## RISC-V SoC Revenues by Region

RISC-V SoC Revenues by Region Tables 2021 - 2031
RISC-V SoC Revenues by Region Figures 2021 - 2031

## RISC-V SoC Design Starts 2021 – 2031

RISC-V SoC Design Starts by Process Geometry Tables 2021 - 2031
RISC-V SoC Design Starts by Device Type Tables 2021 - 2031
RISC-V SoC Design Starts by Process Geometry Figures 2021 - 2031
RISC-V SoC Design Starts by Device Type Figures 2021 - 2031