EETech research

The State of AI in Embedded Engineering

Engineer perspectives on adoption, challenges, and real-world applications of AI within embedded systems.

eetech.com

Methodology

The AI Trends and Innovations Within Embedded Systems study was conducted by Wilson Research Group, LLC and EETech Media, LLC in partnership. The report presented here is provided by EETech Media, LLC.

The research was conducted in **June - September 2025** and includes responses from **301** qualified engineers worldwide.

The survey was distributed via direct email to members of EETech's internal engineering community databases.

Screening criteria excluded respondents who are not embedded systems hardware, software, firmware, system engineers, engineer managers or other ways involved in embedded development.

Only qualifier questions were mandatory; all other questions were optional. Respondents could choose to opt in to a random drawing for a $500 incentive.

Other, Don't Know, Unsure, or Doesn't Apply responses were removed from charts if at/under 5%, or the least answered option.

Statistical Confidence

This study is based on responses from 301 engineers worldwide. At this sample size, results are statistically accurate within approximately ±5.7% at a 95% confidence level, meaning that if the survey were repeated with similar respondents, the results would fall within five percentage points of these findings 95 out of 100 times.

Some analyses are based on smaller sub-samples (for example, around 120 respondents), where the margin of error increases to roughly ±9%. Findings from these smaller groups should be interpreted as directional insights that highlight meaningful trends rather than exact percentages.

Sample Size Margin of Error (approx.)

- n=100 95% ±9.8%
- n=200 95% ±6.9%
- n=300 95% ±5.7%
- n=385 95% ±5.0%
- n=443 95% ±4.7%

Abstract

- **Demographics of Embedded Developers -** Roles, organization size, industry, region, and purchase influence.
- **Experience & Outlook on Embedded AI -** Current involvement with AI, experience levels, and future expectations.
- **Embedded Project Characteristics -** Applications, operating systems used, open-source integration, power analysis tools, connectivity (5G/6G), and safety compliance.
- **Techniques & Performance Benefits -** Programming languages, AI methods, and realized design gains.
- **AI Tools & Frameworks -** Frameworks, model optimization methods, pre-trained model usage, cloud AI services, and hardware/development platforms.
- **Challenges, Data & Security Constraints -** Security vulnerabilities, data collection and labeling issues, explainability frameworks, connectivity limitations, and technical barriers.
- **Automotive AI Applications -** Automotive use cases, critical sensors, safety explainability, and expected areas of AI impact in vehicles.
- **Chip Type & Vendor Usage -** Processor architectures, chip families, vendor preferences, distributor relationships, and supporting software ecosystems.
- **Learning Methods & Media Consumption -** How engineers build skills and where they consume trusted technical and industry media.

Demographics

Geography

Where do you work?

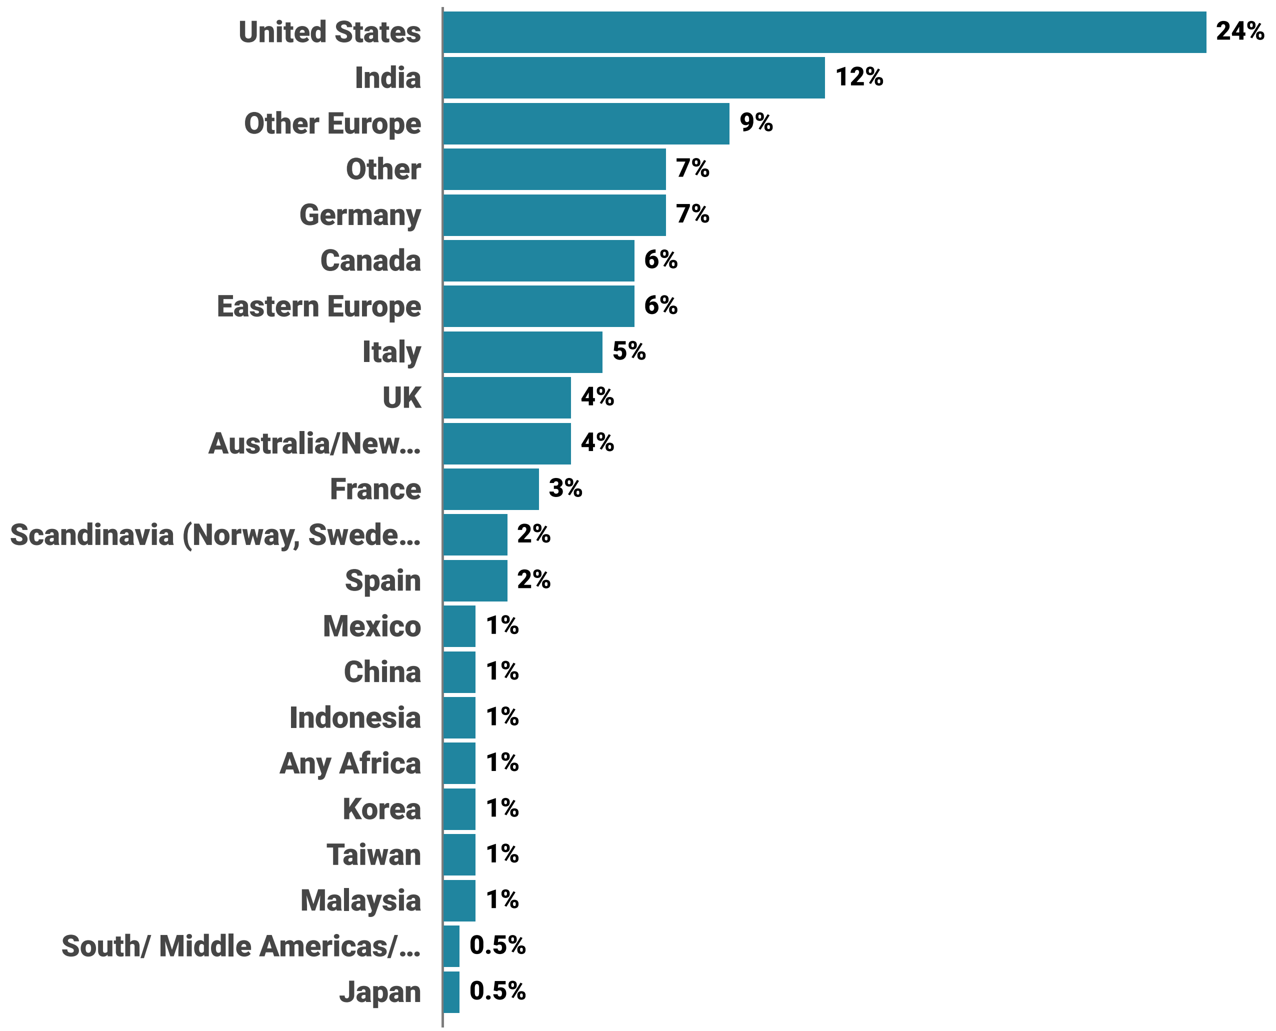

n=210

The study reached a broadly international audience, though participation was weighted toward North America and Europe.

- The United States leads as the top country represented, followed by India, confirming strong engagement from both mature and emerging engineering markets.
- Europe as a region (including Germany, Italy, the UK, France, Scandinavia, Eastern) accounts for roughly 38% of respondents, giving it slightly higher participation than North America (31% combined from the U.S., Canada, and Mexico).
- Asia-Pacific regions collectively contribute about 19%, led by India, with smaller shares from China, Indonesia, Korea, Taiwan, and others.

The distribution shows that the survey reached a truly global embedded engineering community, anchored by U.S. and European professionals but with a growing presence from India and broader Asia-Pacific markets driving future growth in embedded AI.

Age

What is your age?

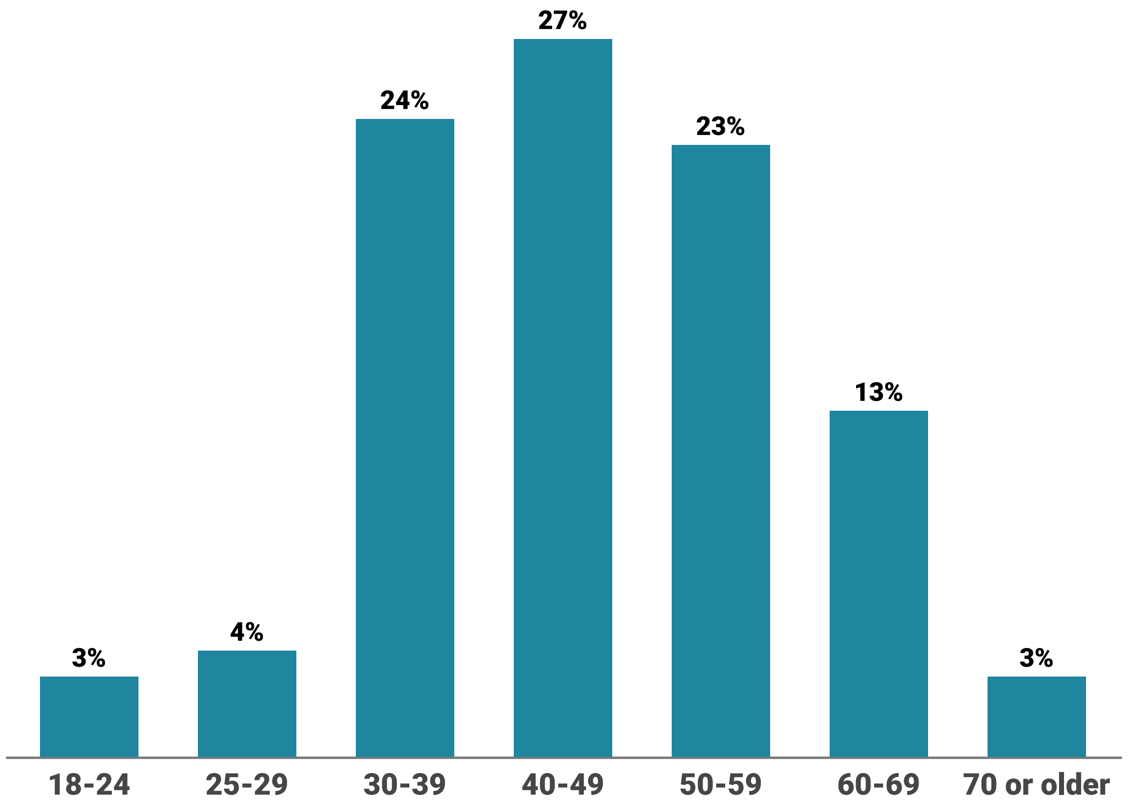

The dataset represents an experienced engineering population, concentrated in mid- to late-career stages.

- Approximately half of respondents (51%) are between 30 and 49, marking the core professional demographic for embedded systems development.
- Another 39% are aged 50 or older, highlighting the strong participation of veteran engineers with decades of industry experience.
- Only 7% are under 30, suggesting that younger professionals and recent graduates are underrepresented in this sample.

In sum, this audience primarily reflects career-established engineers, many of whom are likely to hold decision-making or leadership roles within embedded design teams.

n=202

Role

Are you an embedded systems hardware, software, firmware, systems engineer, engineering manager or other ways involved in embedded development?

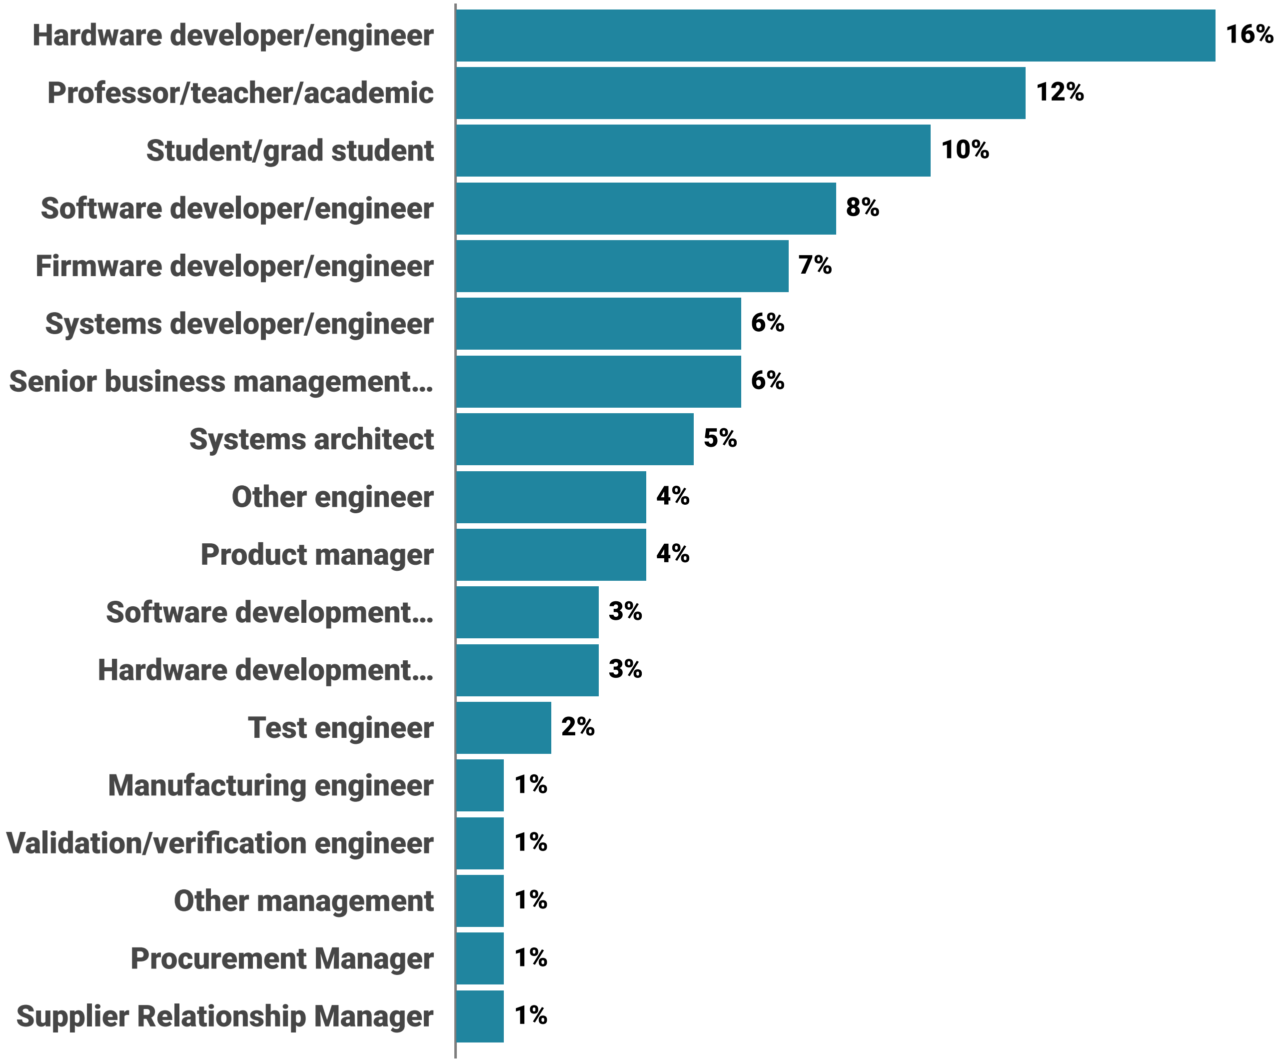

n=443

The respondent pool is primarily composed of technical practitioners, with hardware, software, and firmware engineers forming the largest segments. These three roles together represent over 30% of total respondents, confirming the survey reached a highly technical, engineering-focused audience.

- Systems developers and architects add to the technical depth, while product management and senior business leadership bring higher-level perspectives.
- Academic participation is notable: 12% educators and 10% students, suggesting strong engagement from both current and emerging engineers.
- Procurement and supply chain professionals make up only about 2% combined, reinforcing that this is a developer-driven dataset rather than a purchasing study.

Firmographics

Approximately, what is your organization's annual revenue?

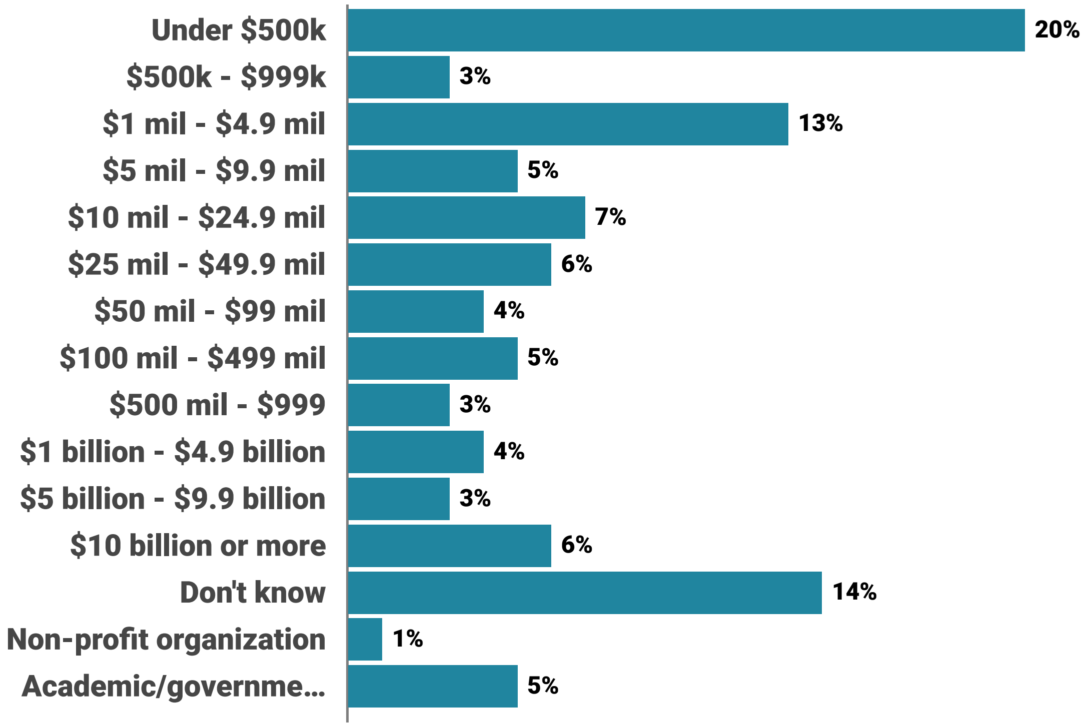

The largest single group (20%) reports revenues under $500K, reflecting many small or early-stage engineering firms, startups, and consultancies. Mid-sized firms ($1M up to $50M) make up roughly 31%, while large enterprises ($1B+) account for about 13%.

Approximately, how many employees are in your entire organization at all locations worldwide?

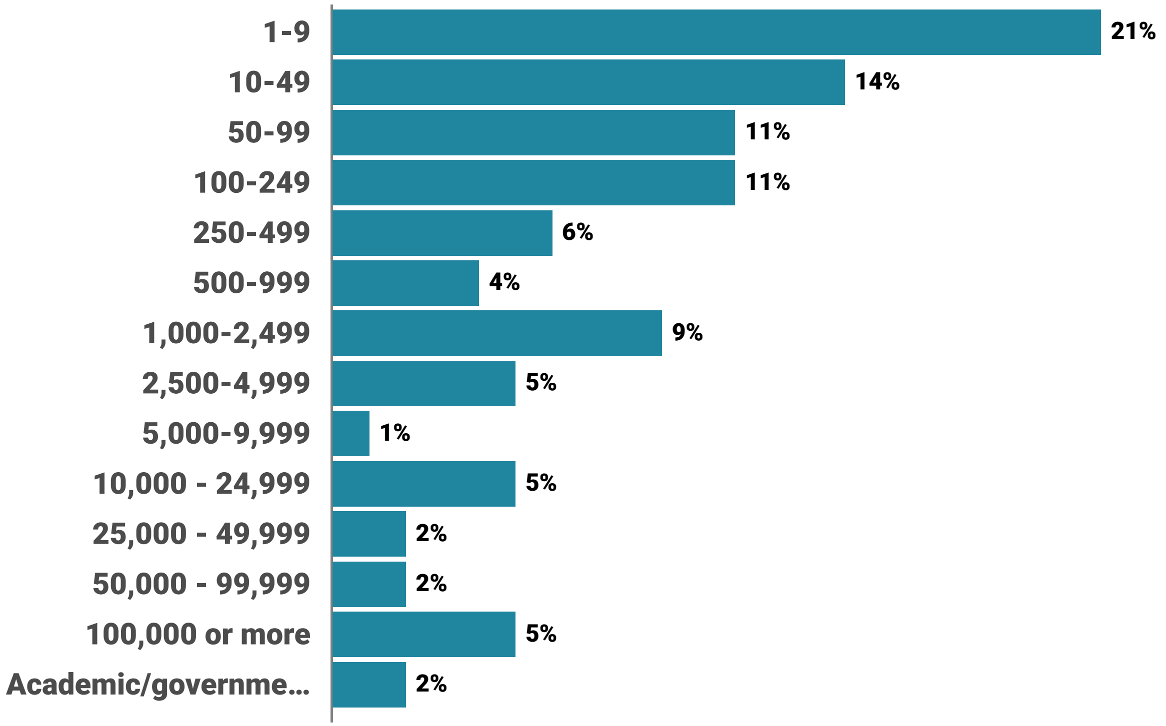

More than half of respondents work in small to mid-sized companies with fewer than 500 employees, while about one-third are part of larger enterprises. This mix suggests that the embedded community spans both agile, innovation-driven firms and established global manufacturers.

n=209

Industry

In which of the following business activities or industries is your organization involved?

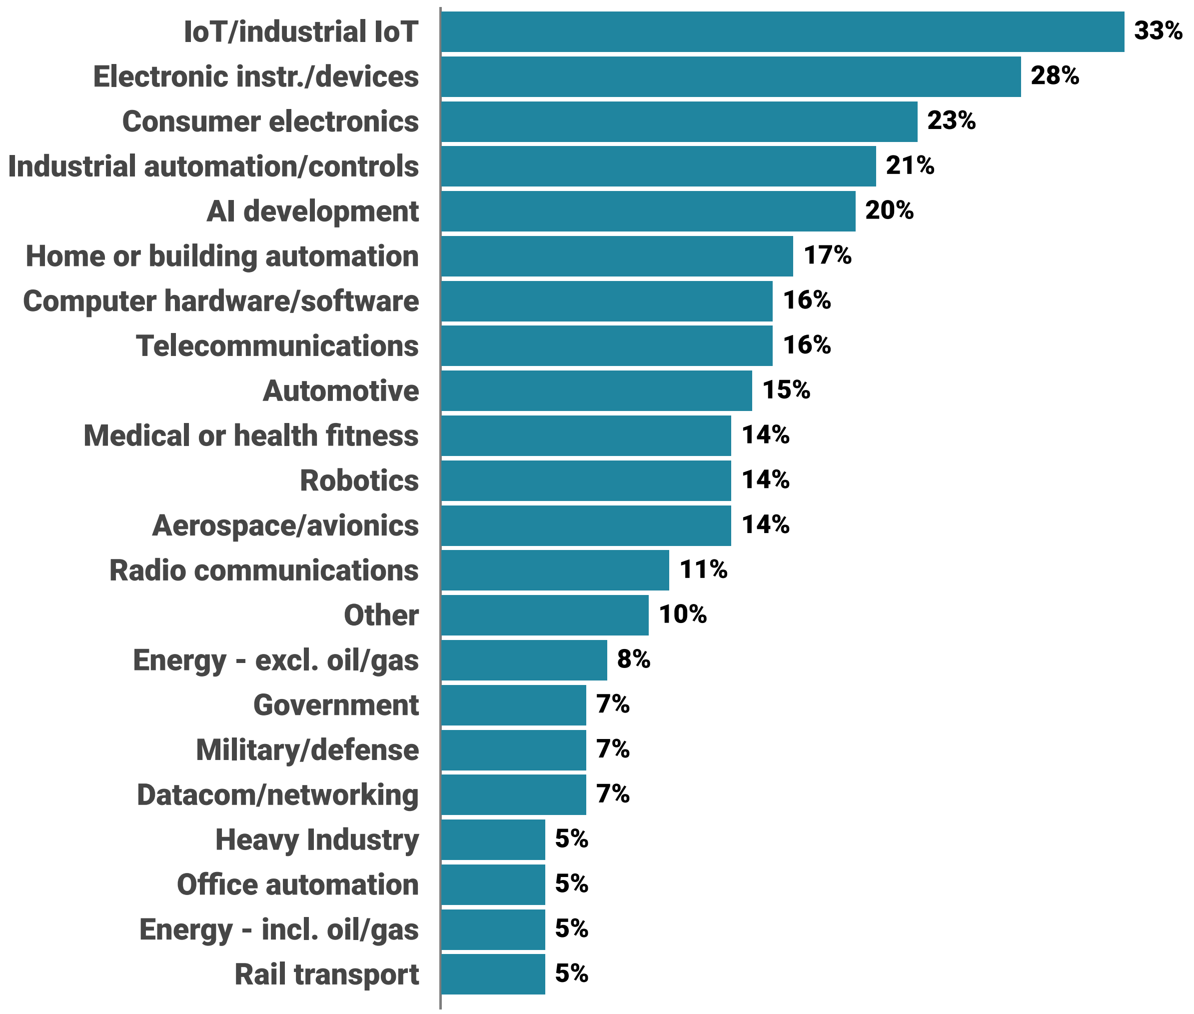

n=220

The survey population spans a broad mix of industries, anchored in connected, industrial, and automation-focused markets.

- IoT/IIoT represent the most common sector, followed by Electronic Instrumentation and Devices, with Industrial Automation and Controls in third place, highlighting that embedded development is concentrated in connected, measurement-driven, and process-control environments.
- Consumer electronics, AI development, and Home or Building Automation reflect the continued expansion of embedded systems into smart, networked devices.
- Automotive, medical, and aerospace signal strong engagement in safety-critical and performance-sensitive applications. Military/defense and government are notable segments, representing the persistent role of embedded intelligence in secure and mission-critical systems.

Overall, embedded development is deeply rooted in industrial, connected, and automation markets, with significant participation across consumer, healthcare, and transportation sectors, reflecting that embedded intelligence is foundational to nearly every product category.

Authority

In your organization, what is your scope of influence in acquiring new semiconductor hardware, software tools, or services?

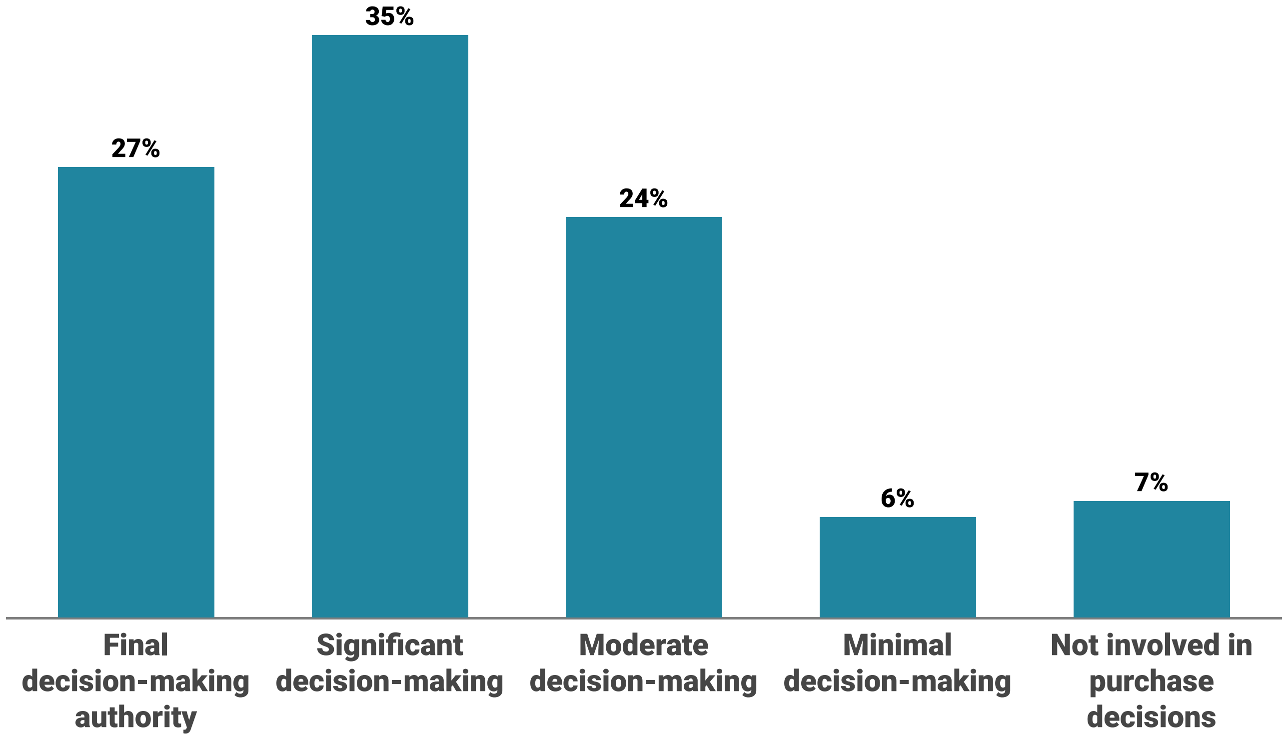

n=210

Most respondents play an active role in technology purchase or specification decisions.

- A combined 62% hold either final or significant influence, confirming that the majority are directly engaged in vendor and tool selection.
- Another 24% report moderate influence, meaning nearly nine in ten respondents have at least some input into technology acquisition.
- Only 13% have minimal or no involvement in these decisions.

These findings confirm that the surveyed individuals represent a highly influential group of practitioners and technical leaders, capable of impacting purchasing and adoption choices across semiconductor hardware, tools, and services.

Demographics - Key Takeaways

Embedded engineers represent a global, professionally mature audience working across a wide range of company sizes and industries. Most respondents are based in the United States and Europe (notably Germany, Italy, and the UK), reflecting how embedded development is distributed across established manufacturing centers and growing technology hubs.

The community is also age-diverse but leans toward experience: half of respondents are between 30 and 49 years old, suggesting a strong concentration of mid-career professionals with both design responsibility and organizational influence.

Company size data shows that embedded design is not limited to large enterprises. More than half of respondents work in organizations with fewer than 500 employees, highlighting a substantial representation of small and mid-sized firms that drive innovation through agility and specialization. At the same time, about one-third work for large companies with over 1,000 employees, indicating that embedded systems remain a core capability within global manufacturing and technology leaders.

The industries represented highlight how embedded systems span both traditional and emerging technology sectors. The largest concentrations are in IoT/IIoT, electronic instruments and devices, and consumer electronics. Other significant segments include industrial automation, AI development, and home and building automation, along with smaller but notable participation from automotive, medical and health technology, and aerospace. This variety reflects how embedded design underpins nearly every connected system, from industrial equipment to smart devices and critical infrastructure.

Demographics - Implications for Marketers

For semiconductor and technology marketers, this demographic profile signals a critical insight: embedded engineers combine technical expertise and knowledge with significant influence over purchasing decisions. Traditional marketing approaches that target only procurement or executive audiences risk missing the professionals who actively make final design decisions. Reaching engineers early in the design phase—before specifications are finalized—is the key to winning long-term loyalty.

Marketing campaigns should prioritize education over promotion. These engineers respond to content that helps them solve real problems, not sales language. Case studies showing measurable performance improvements, transparent technical comparisons, and detailed application notes build credibility far more effectively than brand messaging alone. Engineers will engage deeply with brands that demonstrate competence and respect for their technical intelligence.

Because this audience is global and cross-industry, localization matters. Regional success stories, translated technical documentation, and market-specific webinars can increase engagement significantly. The goal is to show relevance, not reach. Engineers value brands that understand their context—whether they're building an industrial control system in Germany or an IoT product in India.

Finally, marketers should view this demographic as a long-term relationship audience. Engineers remember which companies provided clarity, data, and support—not which ones pushed the hardest to sell. The brands that consistently educate, simplify, and empower engineers throughout the design process will become trusted partners, earning both mindshare and market share in the embedded community.

Experience & Outlook on Embedded AI

Experience

Do you have any experience working on projects involving embedded Artificial Intelligence (AI)?

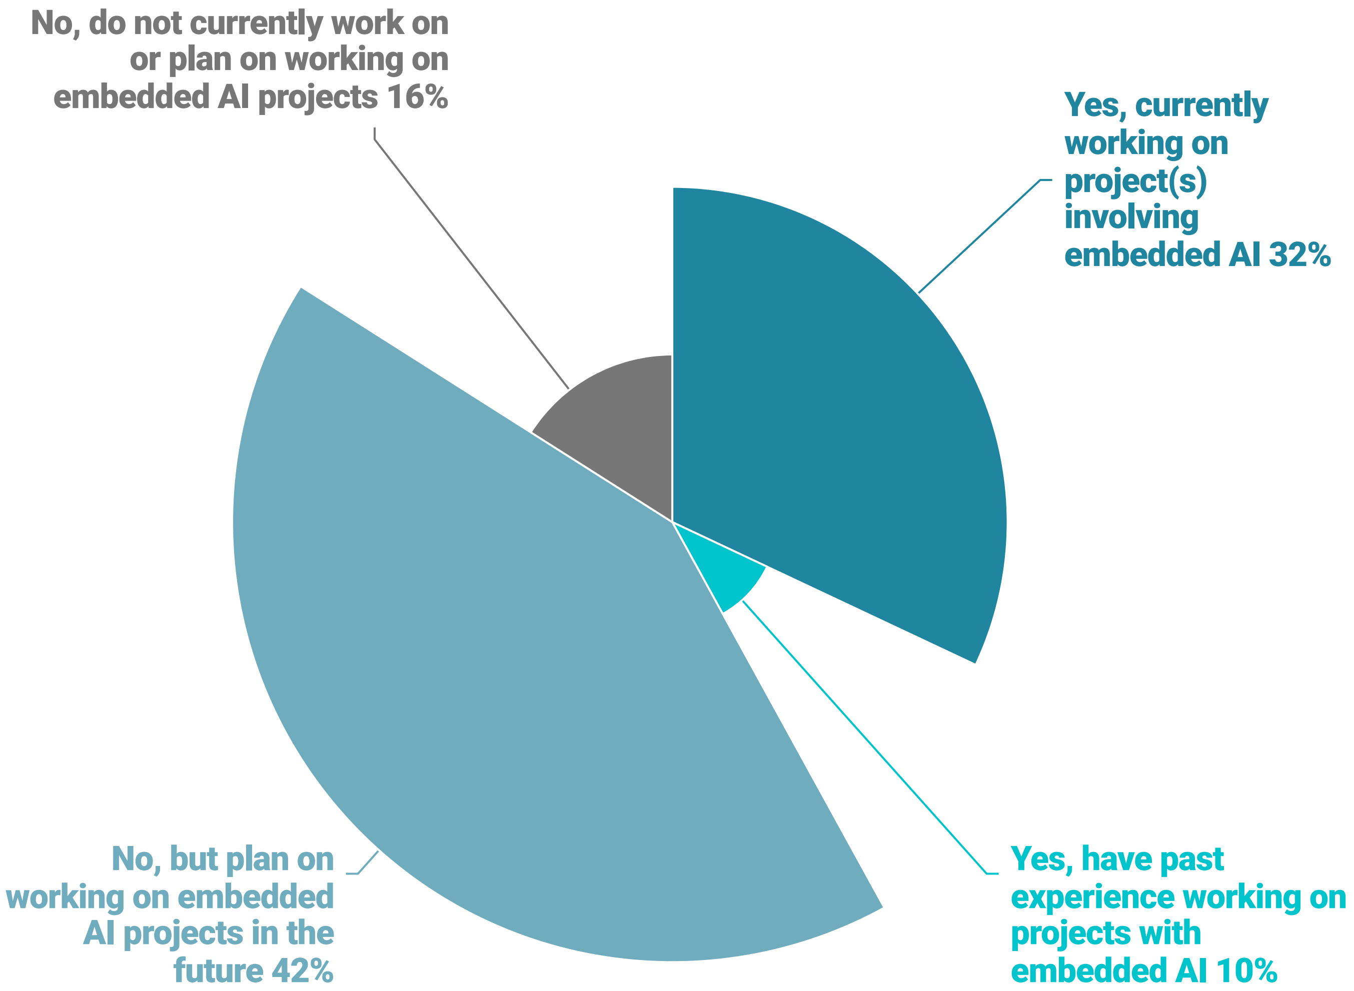

Embedded AI involvement is widespread, though active experience remains limited.

- 42% of respondents are not yet working on embedded AI but plan to in the future, indicating strong growth potential.
- 32% are currently engaged in projects involving embedded AI, and 10% report past but not current involvement, indicating that over 40% of engineers already have hands-on experience.
- Only 16% have no engagement or plans.

In total, 84% of respondents have current, past, or planned involvement with embedded AI, underscoring that AI integration has moved from niche experimentation to an expected skill area for embedded developers.

n=359

Level of Experience

How would you describe your professional experience level with AI?

The developing embedded AI workforce emphasizes the market's considerable future potential.

- The majority, 65%, identify as either Beginner or Developing, underscoring that most engineers are still building foundational AI expertise.
- Only 15% describe themselves as Advanced or Expert, revealing a notable 50-point gap between entry-level familiarity and higher-level proficiency.
- 20% fall in the middle ("Competent"), indicating some project experience but limited complexity.

Overall, the findings show that while interest and participation in AI are growing, the embedded engineering community is still early in the learning curve, suggesting strong demand for training, frameworks, and simplified tools to accelerate skill development.

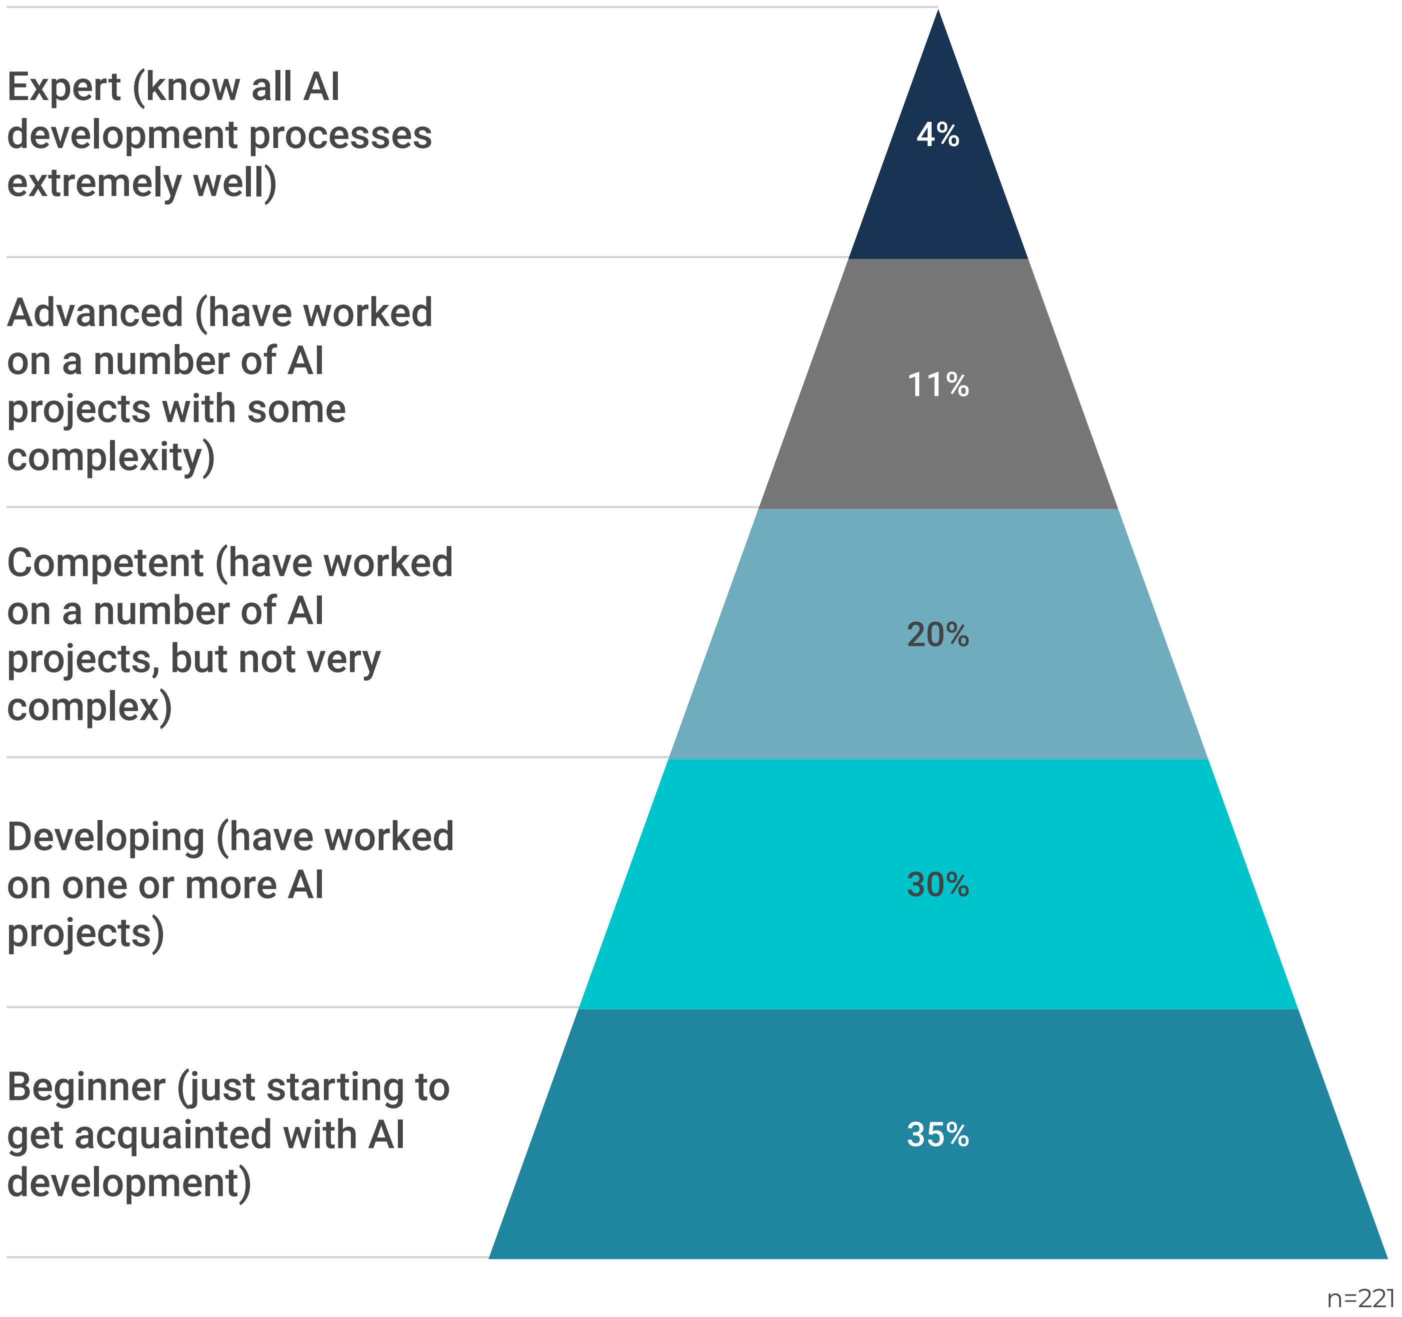

AI Implementation

Which approach best describes how you are implementing AI in your embedded projects?

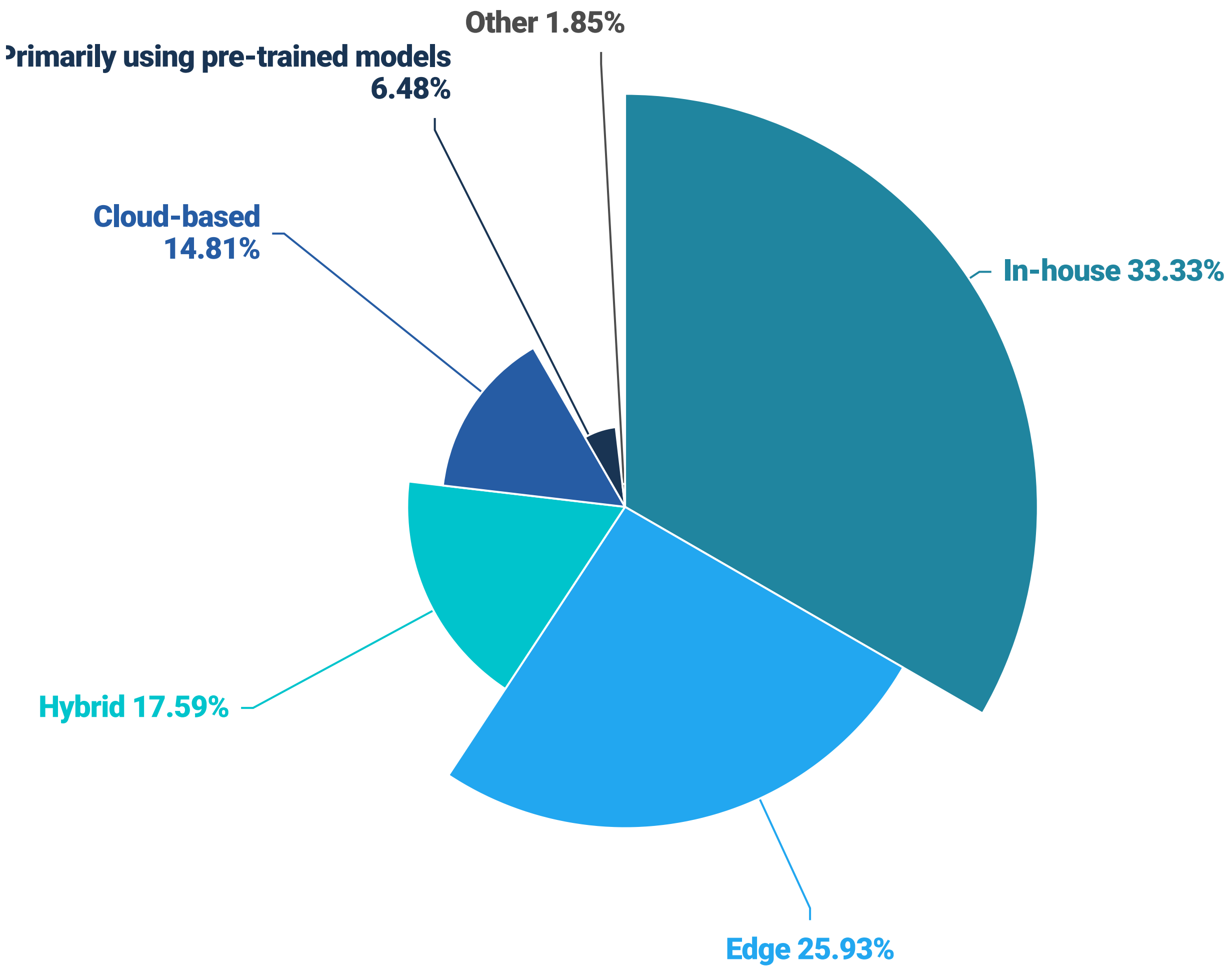

n=108

AI implementation in embedded systems is still evolving, but most engineers favor control and independence in how models are developed and deployed.

- One-third of respondents manage AI entirely in-house, reflecting a desire for customization, security, and tighter hardware integration.
- Edge processing ranks second, confirming that many teams prioritize on-device intelligence to improve responsiveness and reduce reliance on cloud connectivity.
- Hybrid and cloud-based approaches indicate growing experimentation with distributed AI models, though adoption remains lower.
- Only 6% primarily use pre-trained models, suggesting that engineers continue to prefer building or adapting their own models rather than relying solely on generic ones.

The results show that embedded AI remains highly developer-driven, with a strong emphasis on in-house and edge approaches that maximize performance, security, and design control.

AI + Cloud Integration

Does your current embedded AI design integrate with cloud-based AI platforms for model training or updates?

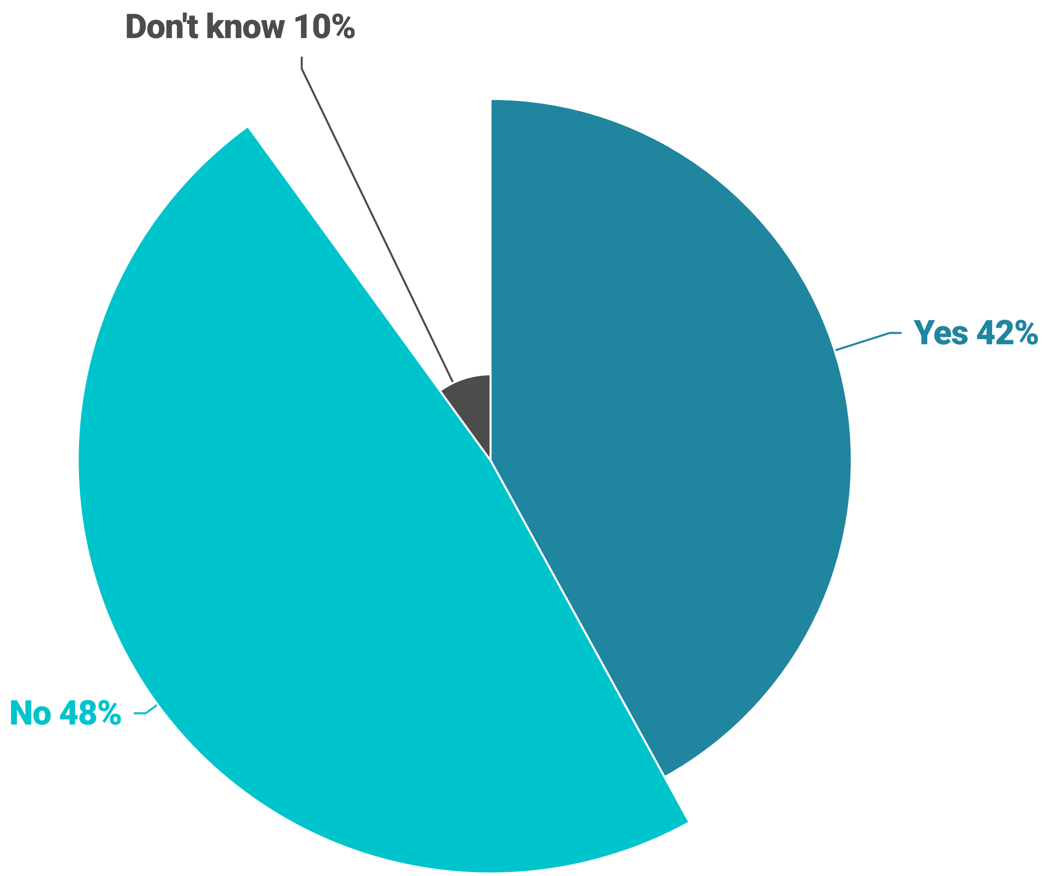

n=89

Cloud integration is emerging but not yet standard across embedded AI designs.

- 42% of respondents currently integrate with cloud-based AI platforms for model training or updates, indicating that many teams are leveraging cloud resources for scalability and faster iteration.
- Nearly half report no integration, reflecting continued preference for local or offline workflows, often due to security, bandwidth, or latency constraints.
- 10% are unsure, likely representing teams that rely on vendor tools without direct oversight of where model training occurs.

We conclude that while cloud connectivity is becoming more common, embedded developers still prioritize control and autonomy in how models are managed and deployed.

System Performance Benefits

What type of system performance benefits are you realizing from your current utilization of AI techniques?

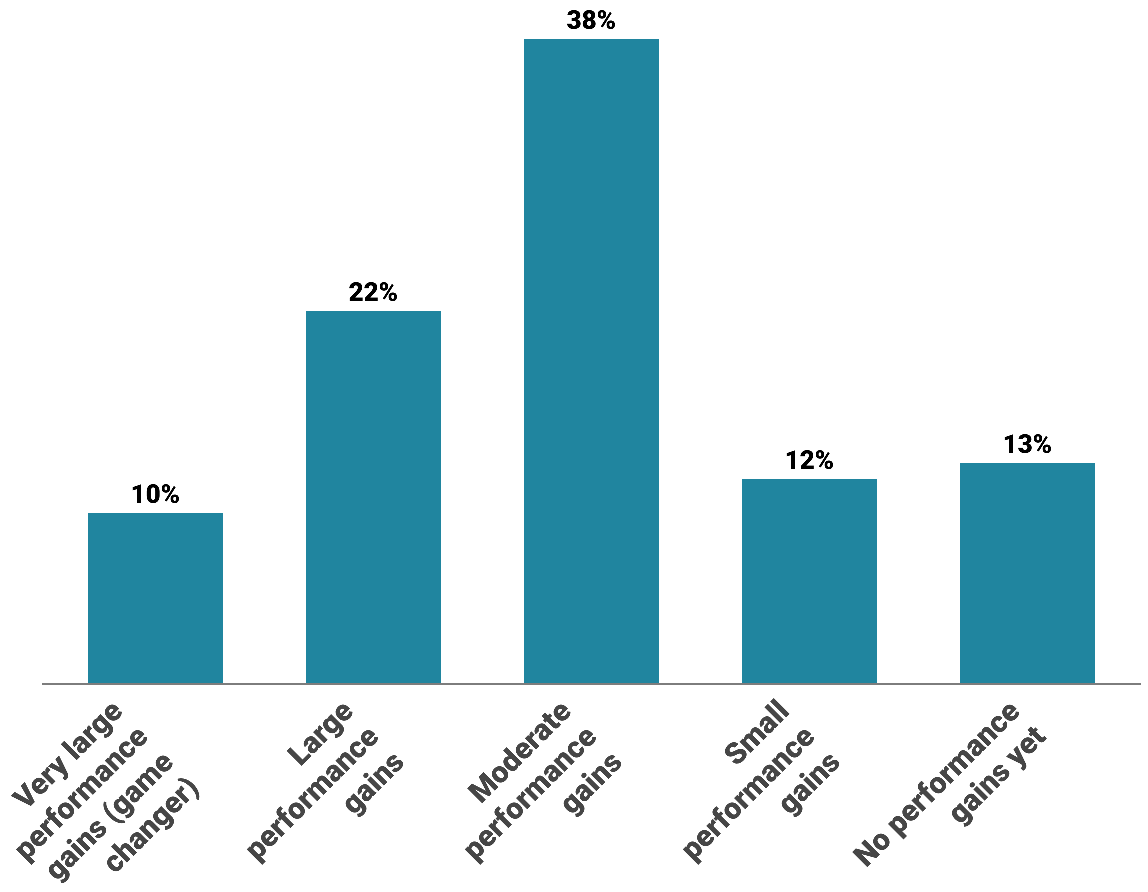

n=100

Engineers report measurable system-level performance improvements from AI integration, though most describe the gains as incremental rather than transformational.

- A combined 60% of respondents experience moderate to large improvements, confirming that AI techniques are already contributing to better system efficiency, speed, or responsiveness.
- 10% characterize the gains as a "game changer," suggesting that breakthrough results are currently limited to select use cases or mature implementations.
- About 13% have not yet realized performance benefits, indicating potential for optimization as AI becomes more tightly integrated with embedded hardware and software.

Overall, the findings show that AI delivers consistent system-level advantages, with the majority of engineers seeing moderate but meaningful improvements in real-world performance.

Important to AI

Which of the following do you anticipate being most important to new AI embedded hardware designs in the next 1-3 years?

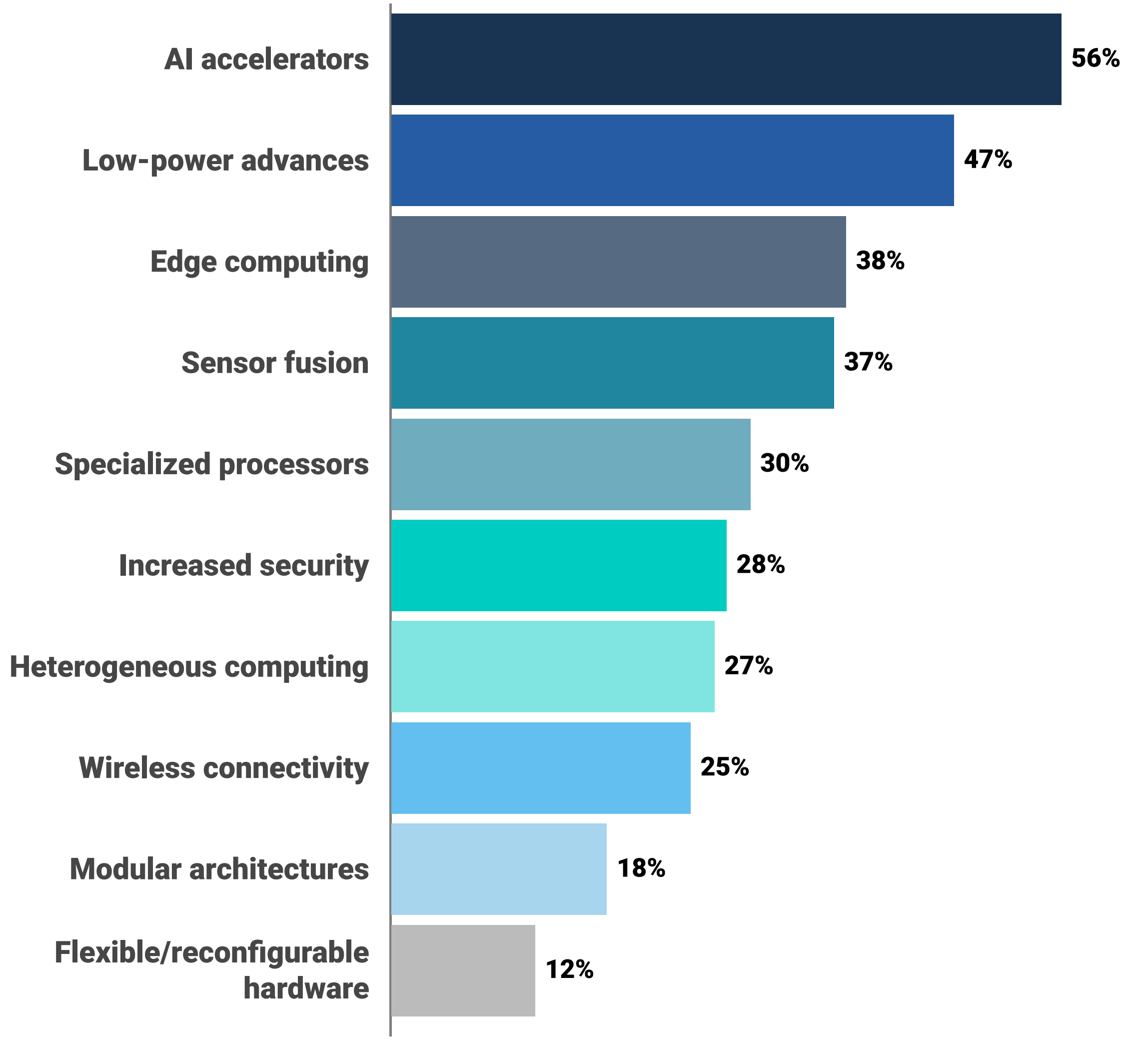

n=89

Engineers expect the next wave of embedded hardware innovation to focus on performance efficiency and AI specialization.

- AI accelerators are seen as the top priority, emphasizing the demand for hardware that can handle intensive AI workloads locally.
- Low-power advances rank second, highlighting ongoing pressure to balance performance with energy efficiency in edge and IoT environments.
- Edge computing and sensor fusion further confirm the industry's shift toward real-time, distributed processing across connected devices.
- Emerging priorities such as specialized processors and increased security indicate that new designs are expected to combine AI performance, safety, and scalability rather than focusing on a single metric.

The findings point to a maturing hardware landscape where AI acceleration and power efficiency drive design decisions for the next generation of embedded systems.

Experience with Embedded AI - Key Takeaways

Embedded developers have moved beyond experimentation—AI is now an integral part of modern design workflows. Nearly all respondents report some level of AI involvement, with many already implementing inference at the edge or integrating ML models within embedded systems.

Experience levels vary, but the trend line is clear: more engineers are shifting from evaluating AI concepts to deploying production-ready designs. When asked about their professional experience with AI, most respondents (65%) stated that they have worked on one or more AI projects, reflecting hands-on application rather than theoretical research.

Developers describe clear technical advantages from AI integration. Many cite improved workflow efficiency, faster optimization cycles, and more adaptive system behavior as measurable benefits—indicating that AI's impact is being realized at the engineering level, even as commercial outcomes continue to mature.

Looking forward, developers expect continued evolution in both deployment and hardware priorities. Roughly one-third anticipate AI processing to remain primarily in-house, while about four in ten expect a balance between edge and cloud—underscoring a future defined by hybrid architectures.

When asked about what will shape next-generation hardware, engineers highlight AI accelerators, low-power advances, and edge computing as top priorities, signaling that performance, efficiency, and localized intelligence will guide near-term innovation.

Collectively, these findings portray an engineering community that's no longer learning about AI—they're engineering with it.

Experience with Embedded AI - Implications for Marketers

For technology marketers, these results point to a clear opportunity: the AI conversation in embedded design has shifted from “if” to “how.” Engineers are no longer deciding whether to use AI—they’re deciding which tools, models, and hardware make integration easier and more reliable. Marketers who can clearly articulate the path to implementation—not just the promise of AI—will stand out.

Effective marketing should focus on practical enablement. Content that walks engineers through optimization examples, training workflows, or on-device inference performance will surpass generic AI narratives. Case studies that demonstrate how real teams reduced latency, improved power efficiency, or achieved real-time performance resonate more strongly than visionary claims about “intelligent systems.” Engineers trust brands that help them apply AI, not just talk about it.

Another key insight is the emergence of hybrid AI strategies—balancing edge performance with cloud flexibility. This creates messaging space around interoperability, security, and scalability. Marketers who show how their solutions support both architectures and enable seamless transitions will earn the trust of design teams focused on long-term adaptability.

Finally, marketers should emphasize evidence over aspiration. Engineers respond to data that verifies real-world performance and efficiency gains. Highlighting measurable outcomes—such as inference speed, energy consumption, or deployment reliability—translates brand promises into technical credibility. The brands that simplify AI integration through clear documentation, ready-to-use frameworks, and transparent performance data will become the trusted partners engineers rely on as AI becomes embedded design’s default feature set.

Embedded Project Characteristics

Applications

Which one of the following is the primary application for your current embedded design?

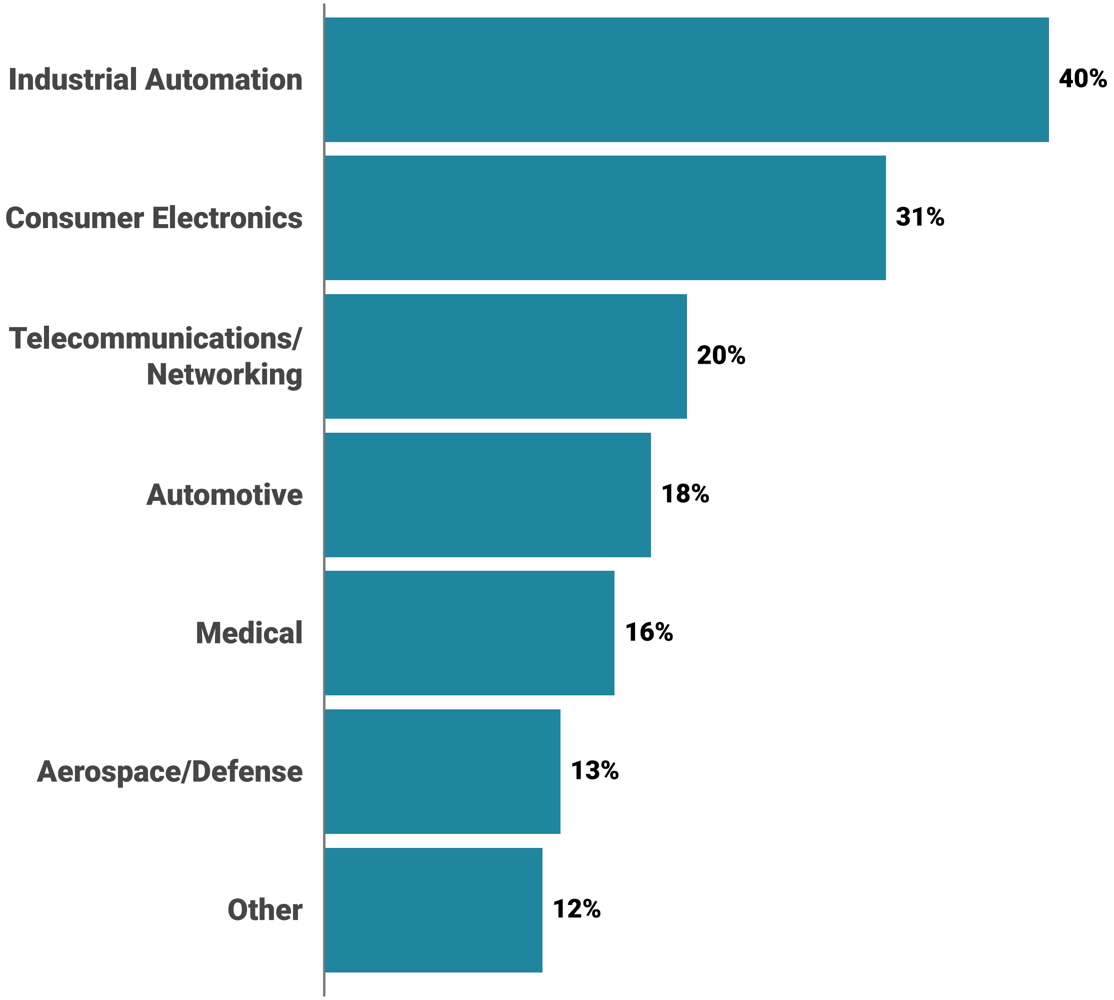

n=179

Industrial and consumer applications dominate current embedded design activity.

- Industrial automation remains the most common focus area, with consumer electronics also commanding strong representation, showing that embedded innovation drives progress across factory floors and consumer devices alike.
- Telecom/networking and automotive form the next tier, underscoring AI's role in connected and mobility systems.
- Medical and aerospace/defense remain important but more specialized niches.

The pattern reinforces that embedded AI development is concentrated in automation, consumer interaction, and connectivity, the three markets seeing the fastest convergence of hardware, firmware, and intelligent systems.

Operating Systems

For which of the following operating systems are you developing?

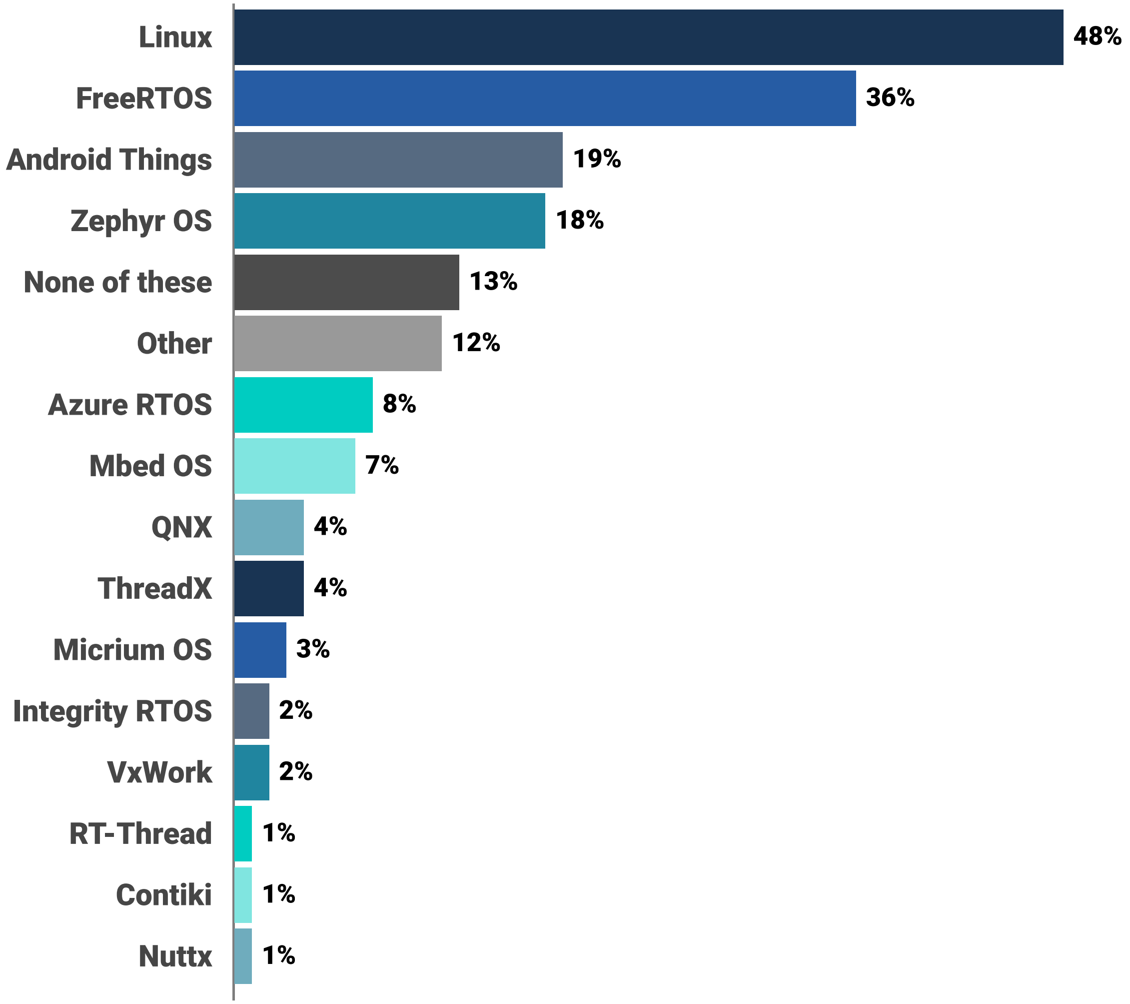

n=179

Linux leads as the dominant development platform for embedded projects, reflecting its flexibility, maturity, and strong community support.

- FreeRTOS remains a key choice for embedded developers, highlighting the continued importance of lightweight, open-source systems in constrained environments.
- A second tier of emerging and specialized platforms includes Android Things and Zephyr OS, both showing meaningful adoption as developers explore newer frameworks for connected and IoT devices.

The embedded ecosystem remains heavily rooted in open-source operating systems, with Linux and FreeRTOS setting the standard while Zephyr and Android Things gain traction in next-generation designs.

Open Source

Approximately what percentage of your current embedded project includes open-source elements?

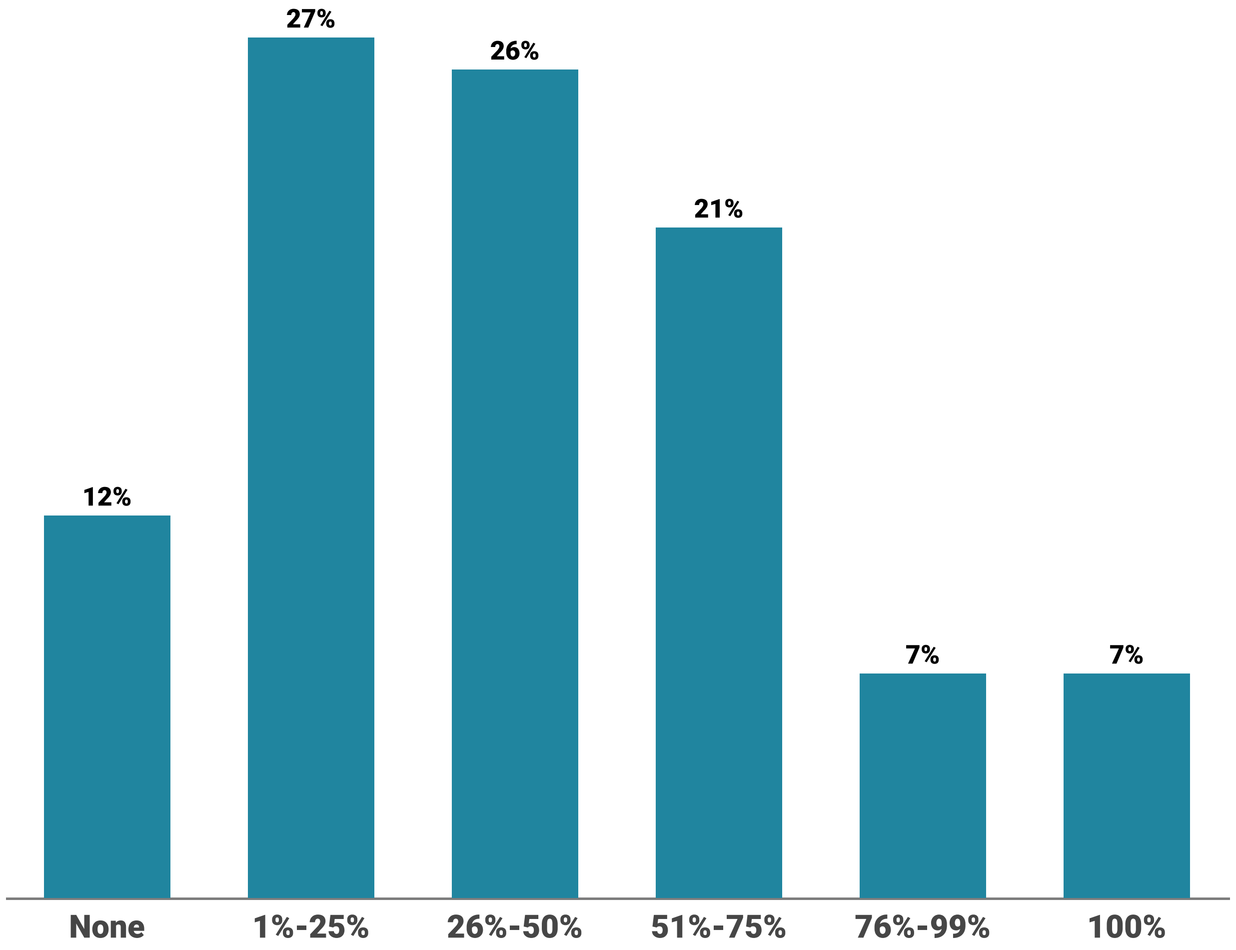

n=177

Nearly all engineers incorporate open-source components into their embedded designs.

- Only 12% report using no open-source elements, while the remaining 88% include at least some level of open-source code or tools.
- The largest segments fall between 1%–50% inclusion, suggesting that open-source software is often used selectively within broader proprietary systems.
- A smaller but meaningful 14% use open-source for more than three-quarters of their project, representing a subset of developers deeply committed to open ecosystems.

In total, the data highlights that open-source software is now a standard part of embedded development, not an exception—integrated at varying degrees across nearly every project.

Top 3 Open Sources

What are the top 3 open-source components currently used in your embedded designs?

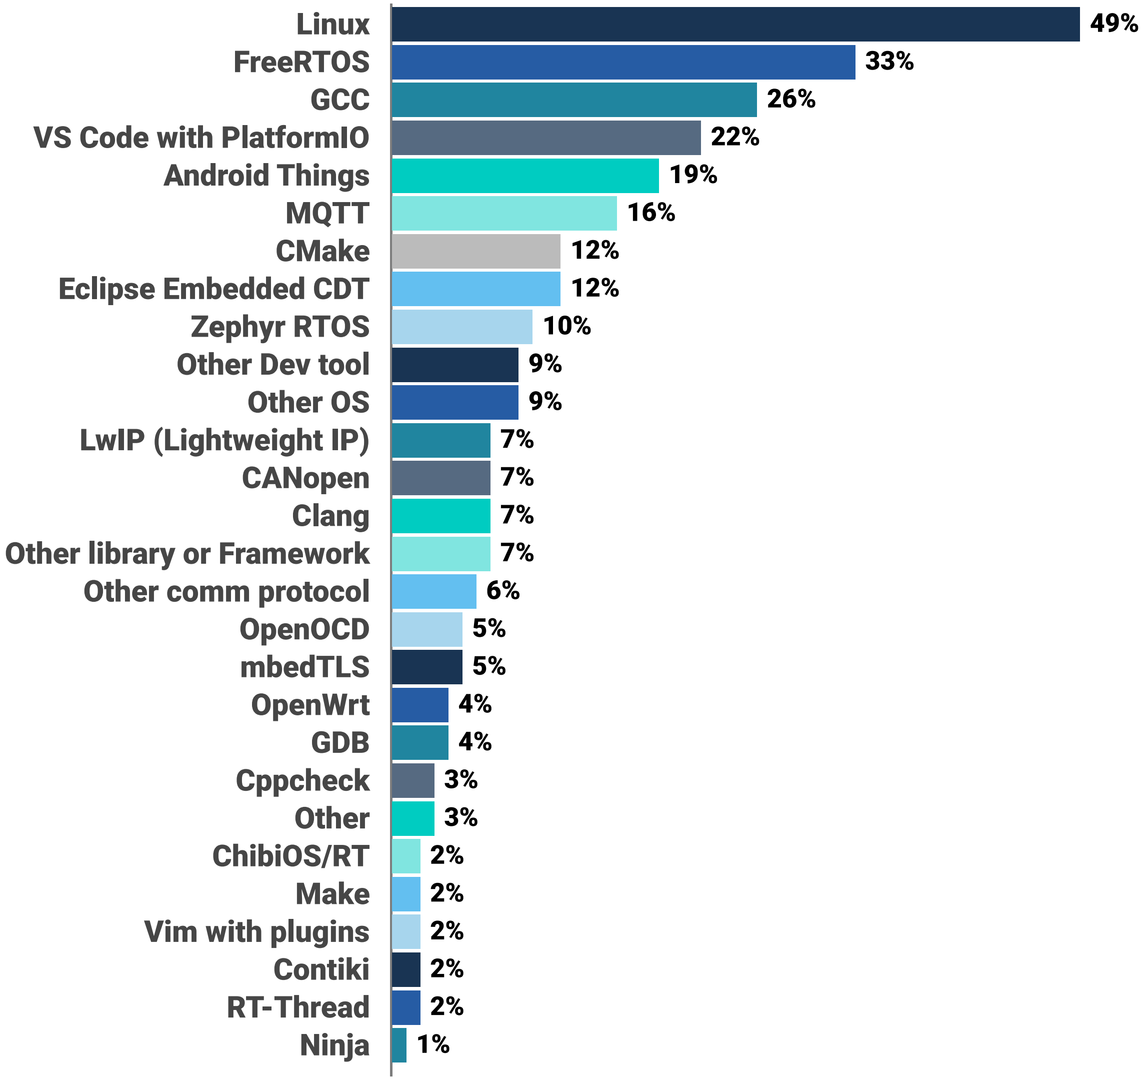

n=174

The open-source landscape in embedded development is dominated by a small group of widely adopted tools and operating systems.

- Linux is the clear leader, reflecting its established role as the foundation for a wide range of embedded applications.
- FreeRTOS and GCC follow as key building blocks for developers seeking flexibility and portability.
- Open-source tool adoption extends beyond operating systems, as seen in the widespread use of development environments like VS Code with PlatformIO and Eclipse Embedded CDT.
- Other elements such as MQTT, CMake, and Zephyr RTOS reinforce the ecosystem's depth and diversity.

The findings show that open-source tools, compilers, and operating systems form the backbone of modern embedded development, giving engineers broad flexibility across architectures and applications.

5G Connectivity

Will your next embedded project incorporate 5G?

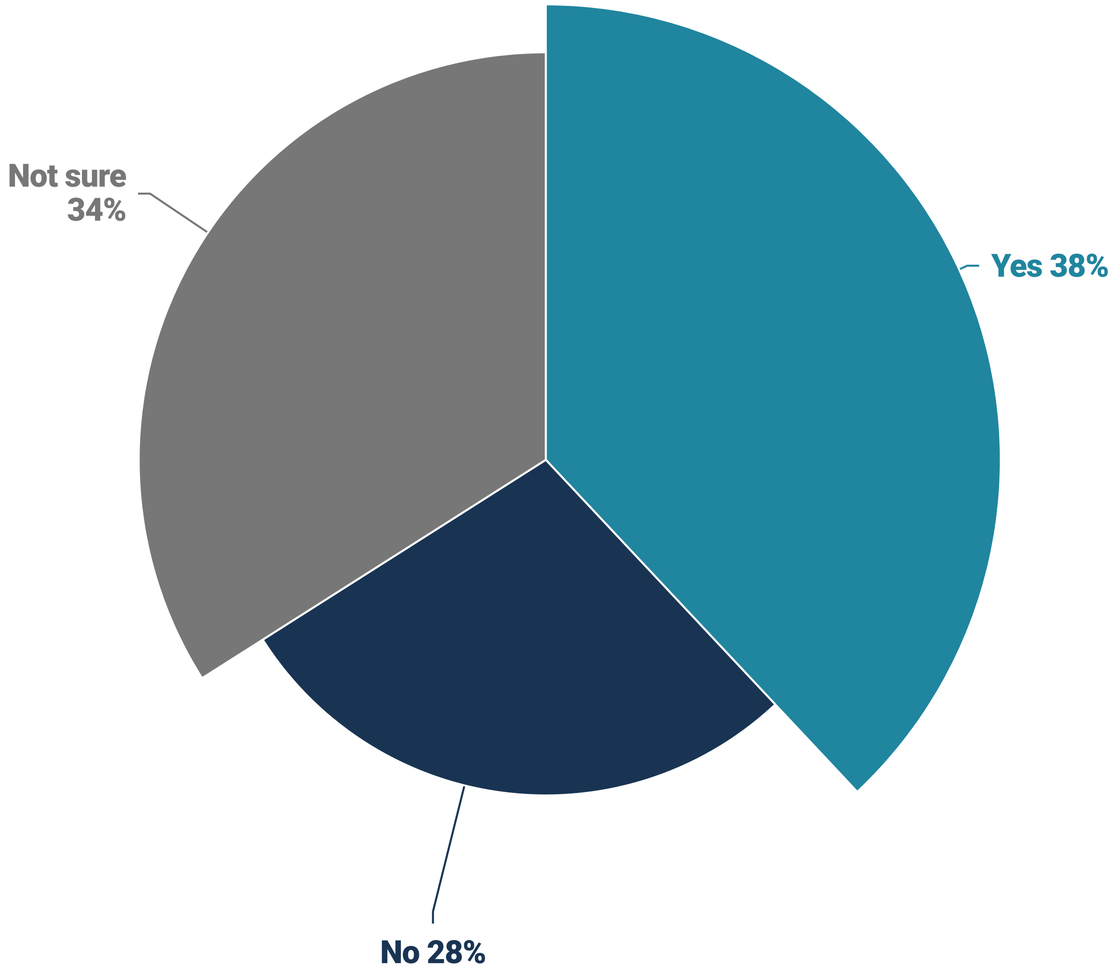

n=159

Interest in 5G integration is strong but not yet universal among embedded developers.

- 38% expect to include 5G connectivity in their next project, confirming that adoption is well underway across advanced and connected applications.
- Another 34% are unsure, indicating that many teams are still evaluating feasibility, cost, or relevance to their use case.
- 28% do not plan to use 5G, likely reflecting projects where alternative wireless standards (Wi-Fi, LTE, or proprietary networks) remain sufficient.

We see that 5G adoption in embedded systems is moving from early experimentation toward broader implementation, but many organizations are still determining where the technology fits best within their product roadmaps.

6G Technology

Considering your future wireless connectivity needs, how important is 6G technology?

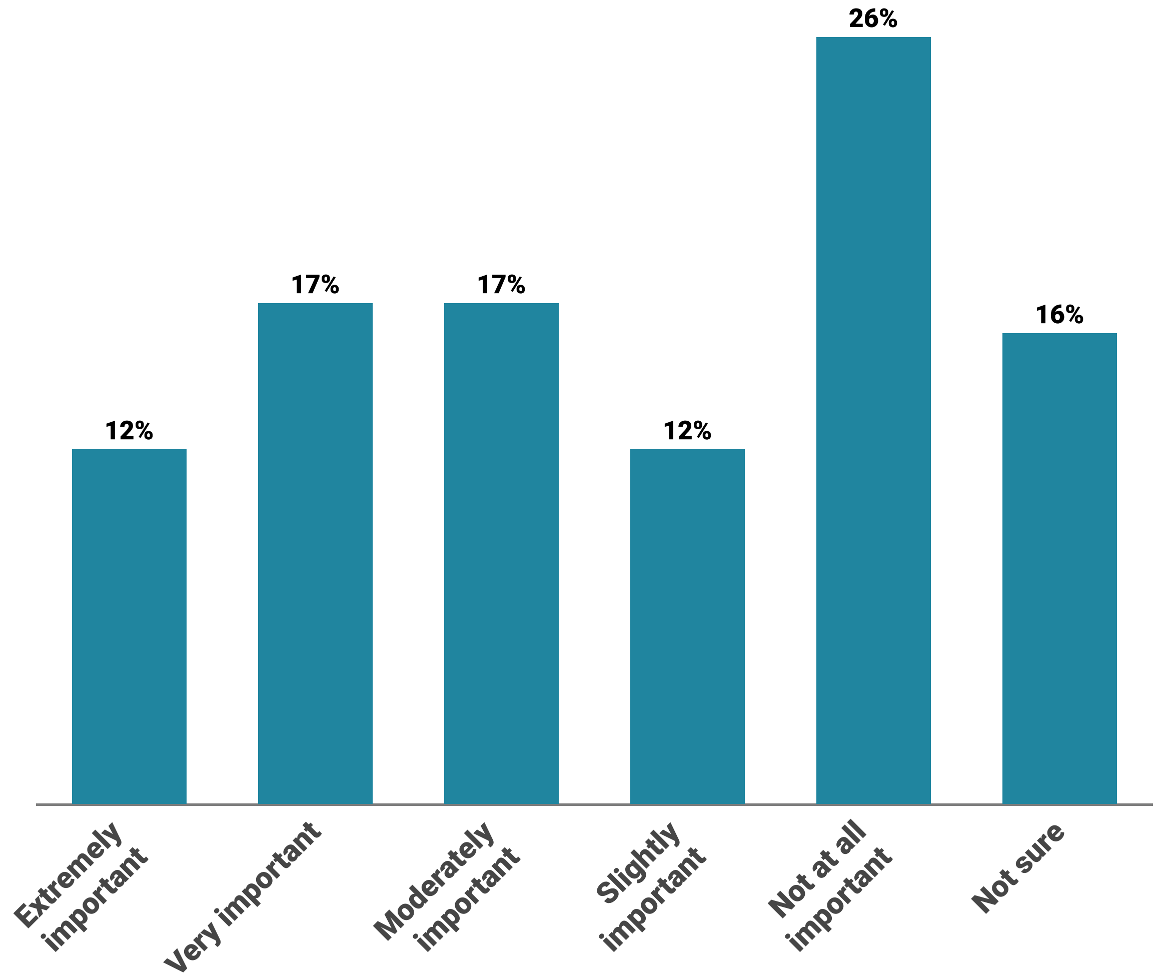

n=155

Most engineers view 6G as a future consideration rather than an immediate priority.

- Only 29% rate it as extremely or very important, while another 17% see it as moderately important—indicating that just under half of respondents are beginning to plan for 6G in some capacity.
- 26% say it is not important, and 16% are unsure, suggesting that the technology remains largely in the exploratory phase for embedded teams.
- The combined 28% who consider it slightly or not at all important show that many organizations are still focused on maximizing current 4G and 5G deployments before considering next-generation infrastructure.

Awareness of 6G is emerging, but widespread practical adoption is still several years away, with most developers monitoring rather than actively preparing for its integration.

Safety Standards

Does your current embedded project follow a development safety standard?

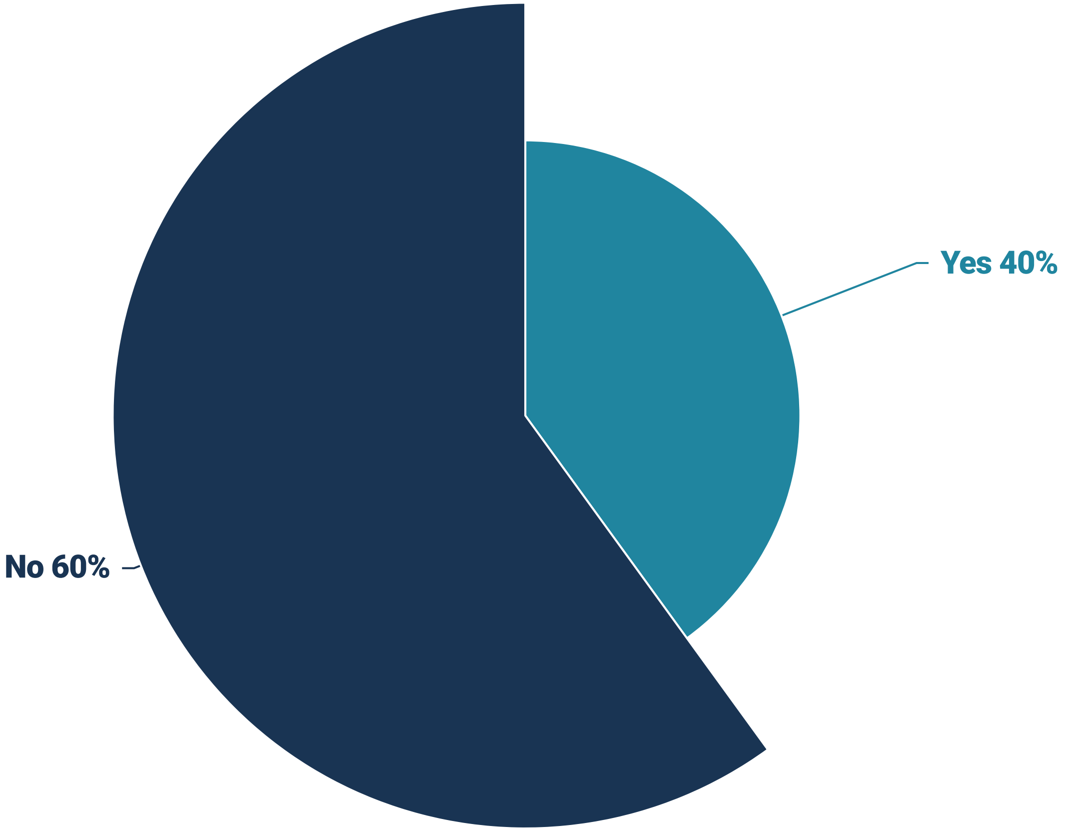

n=173

Four in ten engineers report that their current embedded projects adhere to recognized safety or regulatory development standards, such as DO-254, ISO 26262, or IEC 60601.

The remaining 60% do not operate under these frameworks, indicating that while safety compliance is essential in sectors like automotive, aerospace, and medical, it is not yet a universal requirement across all embedded applications.

This mix reflects the diverse range of industries represented in the study—where safety standards are mission-critical for some markets but not directly applicable to others.

Safety Standards

For which of the following functional safety standards do you or your company develop?

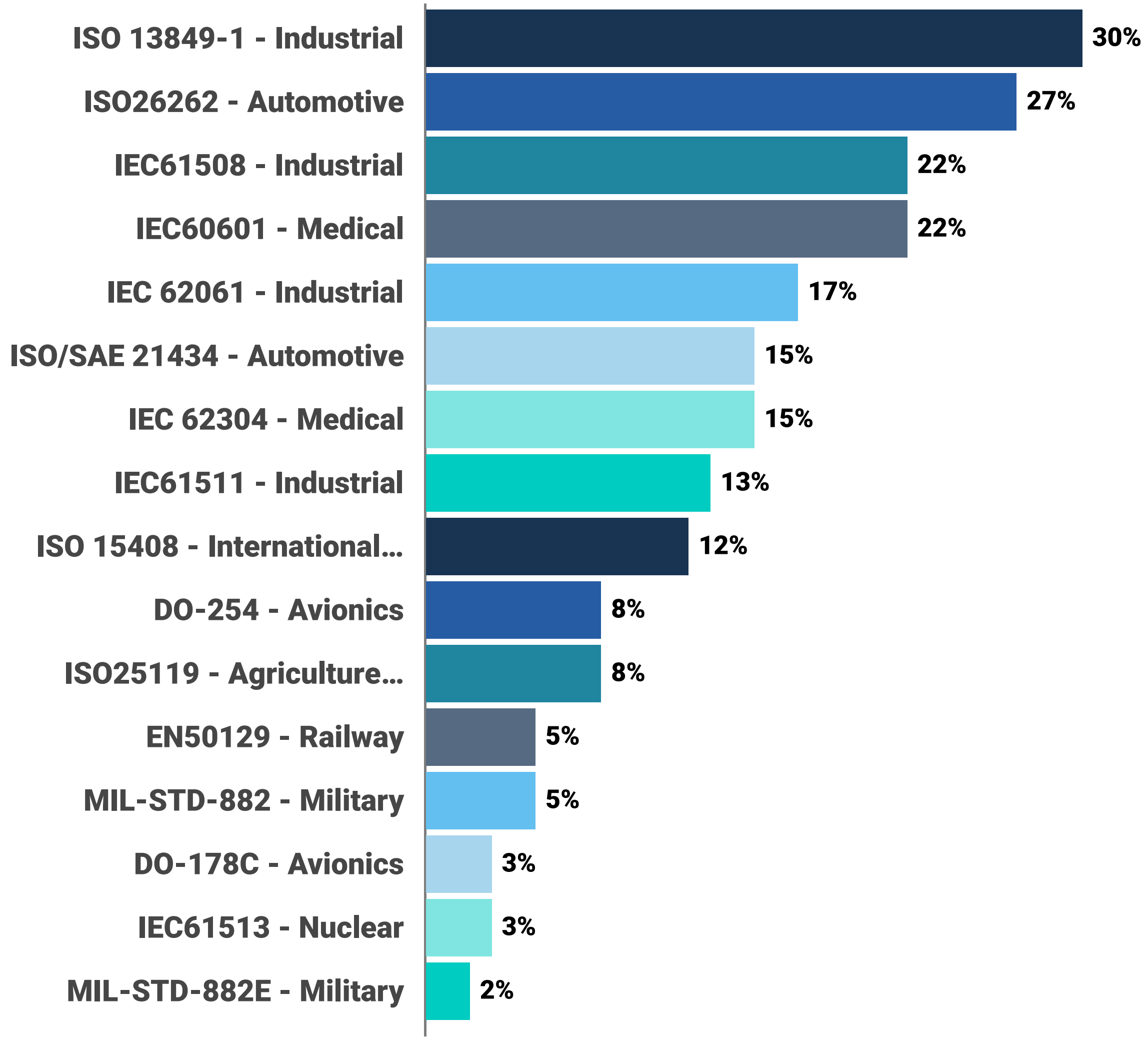

n=60

The safety landscape in embedded design spans multiple industries, led by industrial and automotive applications.

- ISO 13849-1 and ISO 26262 are the most frequently cited standards, showing strong activity in both factory automation and vehicle safety.
- IEC 61508 and IEC 60601 are also widely used, reflecting demand for compliance in industrial controls and medical devices.
- Beyond those, smaller groups are engaged in IEC 62061, ISO/SAE 21434, and IEC 62304, underscoring the spread of functional safety across domains.

The data confirms that functional safety development is concentrated in industrial, automotive, and medical sectors, with more limited adoption in aerospace, defense, and niche applications like nuclear or agriculture.

Embedded Project Characteristics - Key Takeaways

Embedded developers work across a diverse landscape of applications, with industrial automation, consumer electronics, telecommunications, and automotive emerging as the top focus areas. This mix reflects how embedded technology now sits at the intersection of manufacturing, consumer innovation, and mobility—each with unique demands for reliability, real-time performance, and connectivity. Developers are designing for increasingly connected ecosystems, not isolated products.

From a technical standpoint, Linux remains the dominant operating system, followed by FreeRTOS, illustrating the dual demand for open-source flexibility and real-time responsiveness. Most projects include some level of open-source integration: nearly nine in ten developers use open components, and over half say open-source elements make up more than one-quarter of their design. The most common tools—Linux, FreeRTOS, GCC, and PlatformIO—highlight a community grounded in collaboration and customization rather than proprietary lock-in.

Connectivity and safety are also key attributes of modern embedded projects. Over one-third plan to incorporate 5G in their next design, and an additional 12% already view 6G as extremely important to future applications. Meanwhile, 40% of respondents follow recognized safety standards such as ISO 26262, IEC 61508, or IEC 60601, highlighting the central role of compliance in modern embedded development. Engineers are designing for safety-critical systems as AI and connectivity push more functionality closer to the edge.

Embedded Project Characteristics - Implications for Marketers

For semiconductor and technology marketers, these results point to a major shift: engineers are designing for integration and longevity, not one-off product cycles. Projects operate within a broader network of connected devices and regulatory expectations. Marketers who position their products as reliable building blocks in a larger system—highlighting interoperability, lifecycle support, and standards compliance—will earn credibility and preference in design decisions.

The dominance of open-source tools offers another strategic takeaway. Engineers prioritize ecosystems that are transparent, flexible, and well-documented. Marketers can strengthen brand trust by participating in open communities, publishing reference implementations, or offering open APIs and SDKs. When developers see a brand actively contributing to the same ecosystems they depend on, that brand becomes part of their workflow rather than a vendor on the outside.

The growing emphasis on 5G and 6G connectivity suggests a marketing opportunity around real-time performance and reliability. Engineers are eager for data showing how components, modules, and software stacks behave under latency-sensitive conditions. Brands that supply benchmark data, integration guides, and case studies tied to real-world connectivity scenarios can position themselves as enablers of next-generation design.

As safety and compliance continue to rise in importance, marketers should treat certification not as a footnote but as a proof point of reliability and professionalism. Highlighting alignment with recognized standards like ISO 26262 or IEC 61508 can build immediate trust among design engineers. In the embedded market, confidence in quality and compliance often determines which brands become long-term partners in critical systems.

Techniques & Performance Benefits

Programming Languages

Which of the following programming languages are you using in your current embedded project?

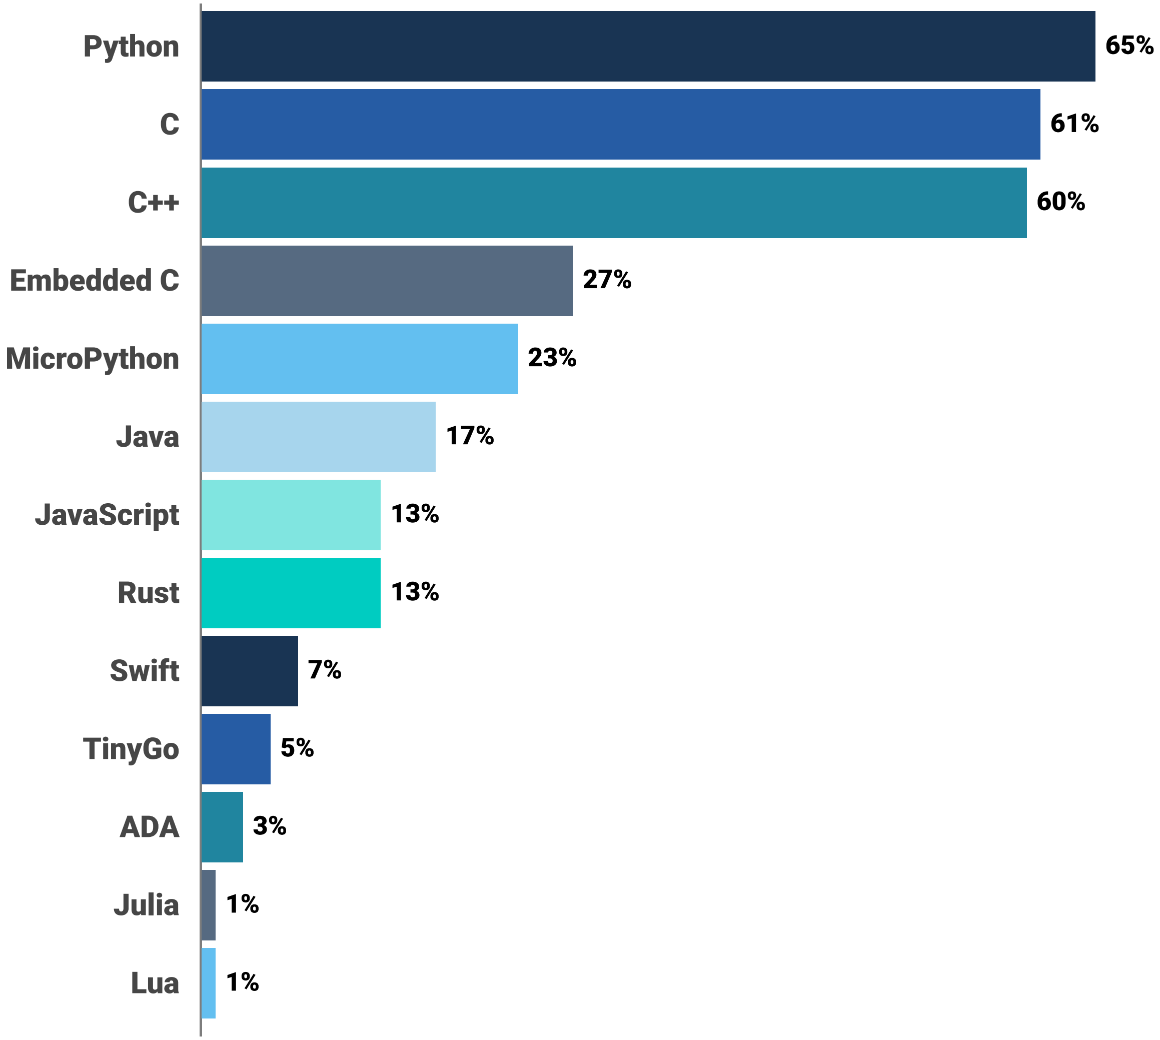

n=103

Embedded developers continue to rely on established languages while integrating newer, higher-level options for AI and automation.

- Python, C, and C++ form the dominant trio, showing that engineers blend traditional low-level control with Python's established role in data processing, automation, and AI workflows.
- Embedded C and MicroPython highlight how developers are adapting existing toolchains for modern embedded environments.
- Emerging use of Rust and JavaScript suggests growing experimentation with programming frameworks traditionally associated with other industries.

Overall, the results confirm that while C and C++ remain foundational, Python has become an essential programming language for embedded engineers, bridging low-level hardware control and AI development tasks.

AI Techniques

Which of the following AI techniques are you currently using within your embedded systems designs?

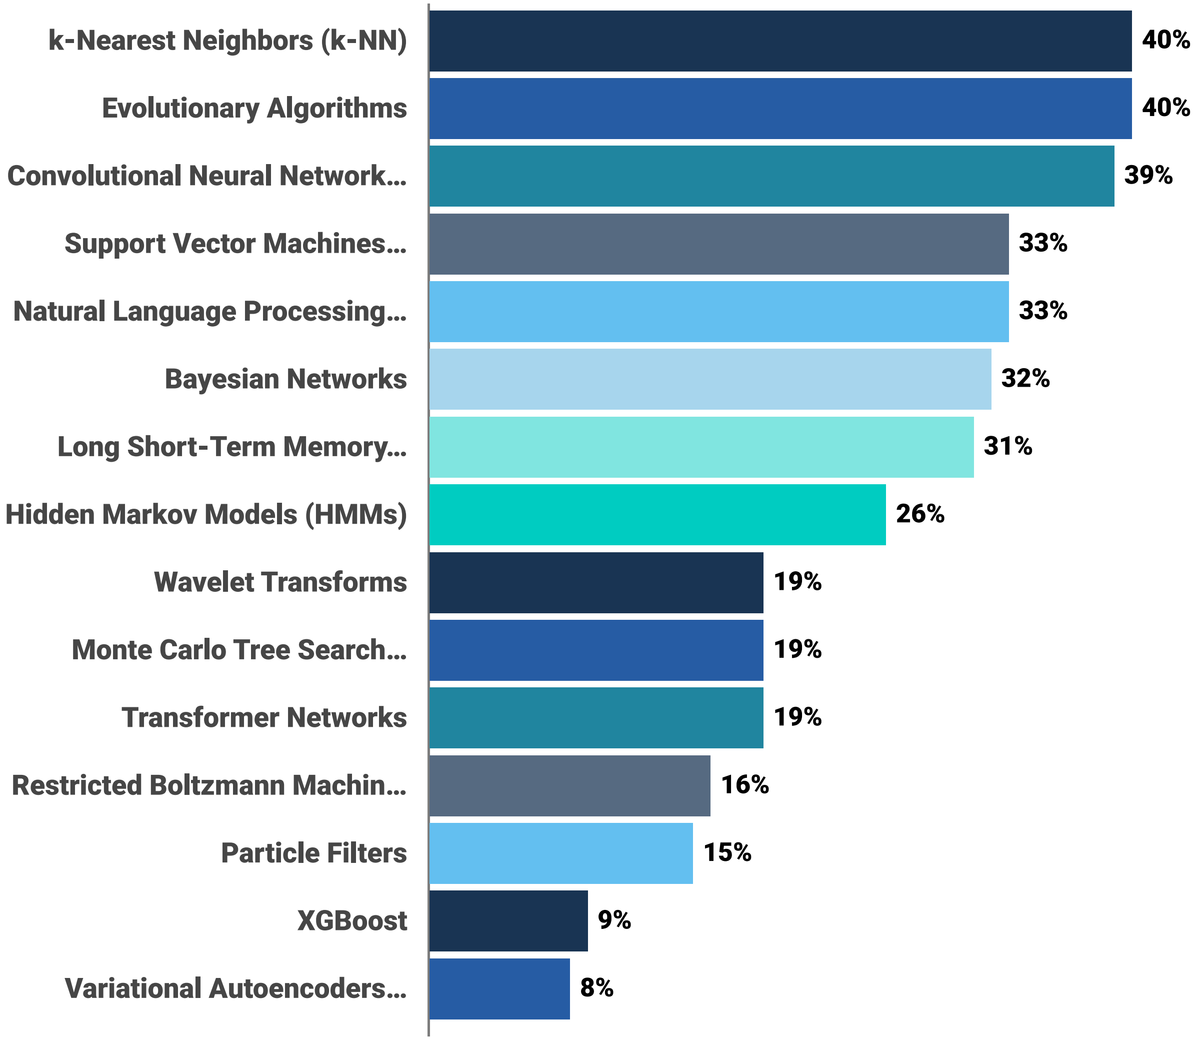

n=102

Answer options at/below 5% were removed

Embedded engineers are applying a diverse mix of AI techniques, with both classical and deep-learning methods in active use.

- The most common approaches include k-Nearest Neighbors, Evolutionary Algorithms, and Convolutional Neural Networks, highlighting the balance between optimization-focused and perception-oriented AI applications.
- A second tier of adoption—SVMs, NLP, and Bayesian Networks—reflects broader experimentation beyond vision systems, extending AI into pattern recognition, decision-making, and predictive modeling.
- Advanced methods such as LSTM networks demonstrate growing engagement with sequential and temporal data.

The findings highlight that AI in embedded design is broad and maturing, with developers employing a wide toolkit that spans both traditional algorithms and modern techniques.

AI Technique Application

Within your current embedded system design, how are you applying AI techniques?

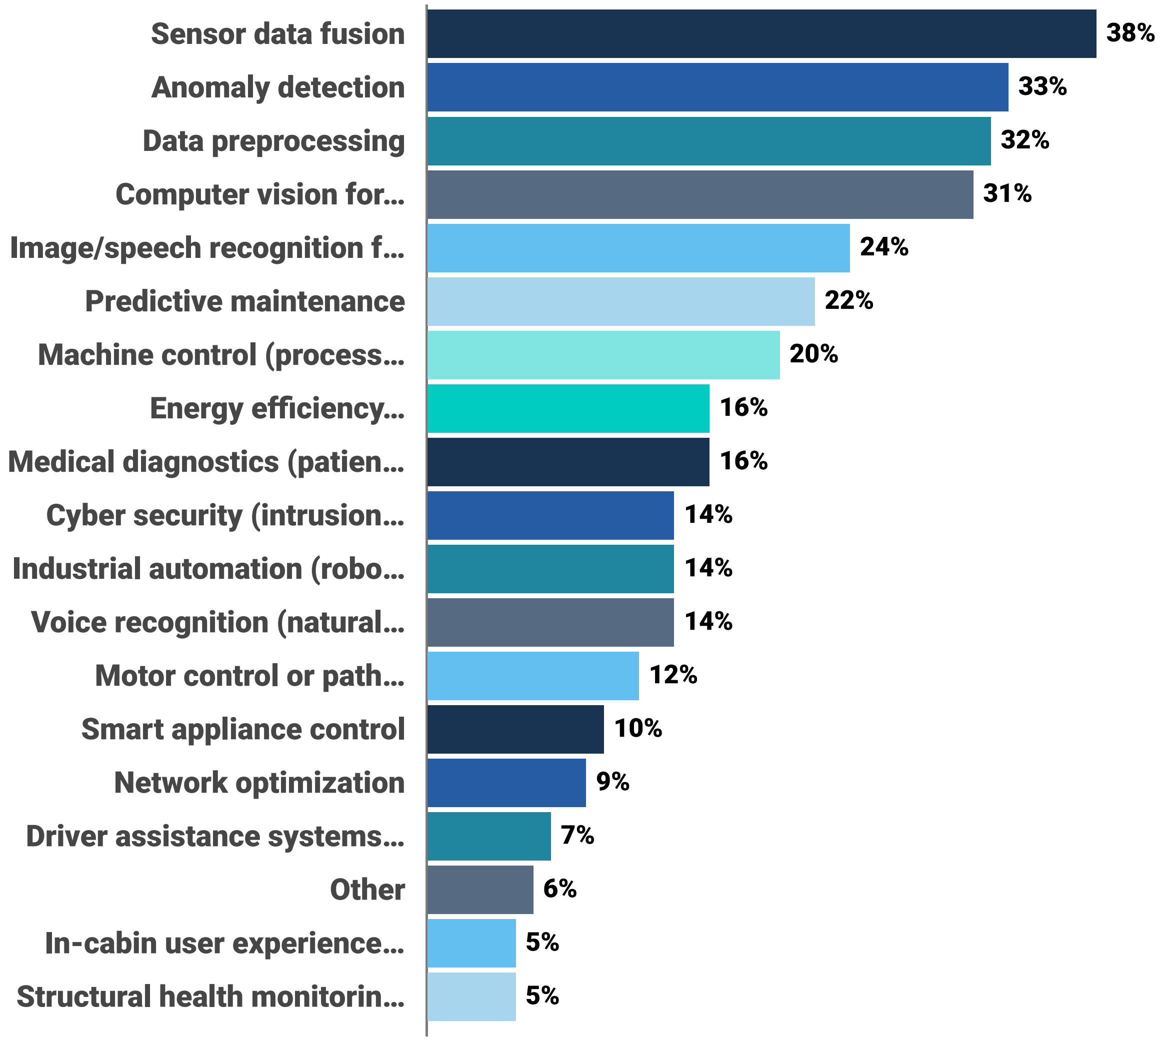

n=102

AI is being applied across a wide range of embedded use cases, with the highest adoption centered on data interpretation and quality control.

- The top three applications—sensor data fusion, anomaly detection, and data preprocessing—show that engineers are primarily using AI to enhance signal accuracy, reliability, and system intelligence.
- Computer vision and predictive maintenance represent major practical implementations, especially in industrial and manufacturing environments.
- Emerging applications such as energy management, medical diagnostics and cybersecurity highlight the expanding reach of AI beyond automation into health, sustainability, and safety.

This data indicates that AI is most widely used to make embedded systems smarter, more efficient, and more adaptive, with strong representation in data processing and machine vision tasks.

Tools for Power Consumption

Which of the following tools do you use to profile and analyze the power consumption of your embedded systems?

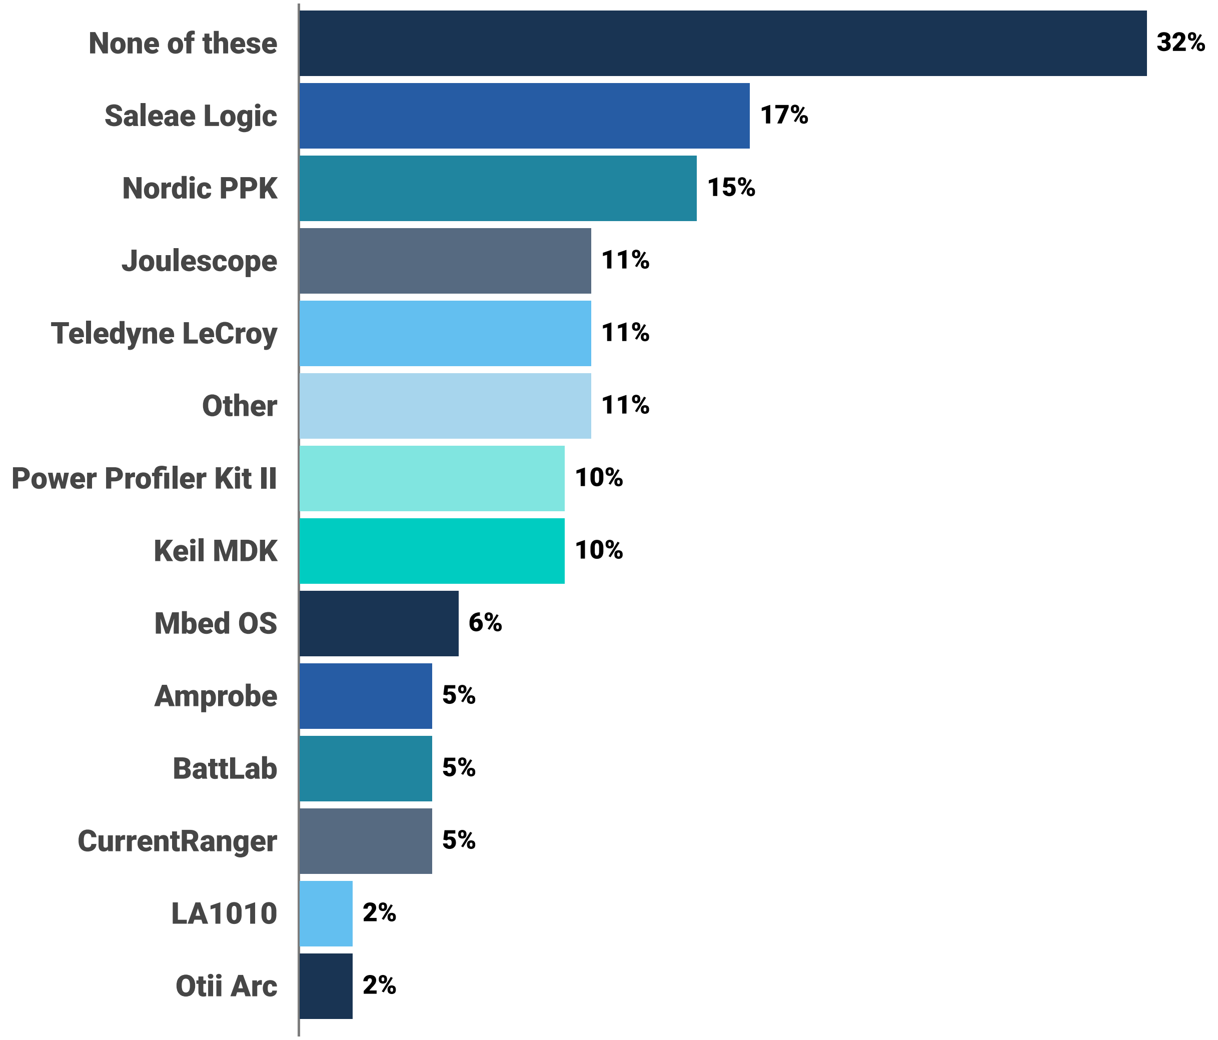

n=157

A significant portion of developers do not currently use the listed power analysis tools, with 32% selecting "none of these." This figure likely reflects a combination of engineers who aren't measuring power consumption at all and those who rely on proprietary or unlisted solutions, since another 11% chose "other."

- Among active users, Saleae Logic and Nordic PPK are the most common choices, showing strong adoption of accessible, low-cost measurement platforms.
- Smaller groups rely on Joulescope, Teledyne LeCroy, and Power Profiler Kit II, indicating use of more specialized or precision instruments.

The findings suggest that while power profiling is gaining traction, a large share of engineers either lack dedicated tools or use proprietary setups, representing a clear opportunity for tool vendors to drive greater awareness and integration of power-optimization solutions.

Power Saving Techniques

Which of the following power-saving techniques do you currently use in your embedded systems design?

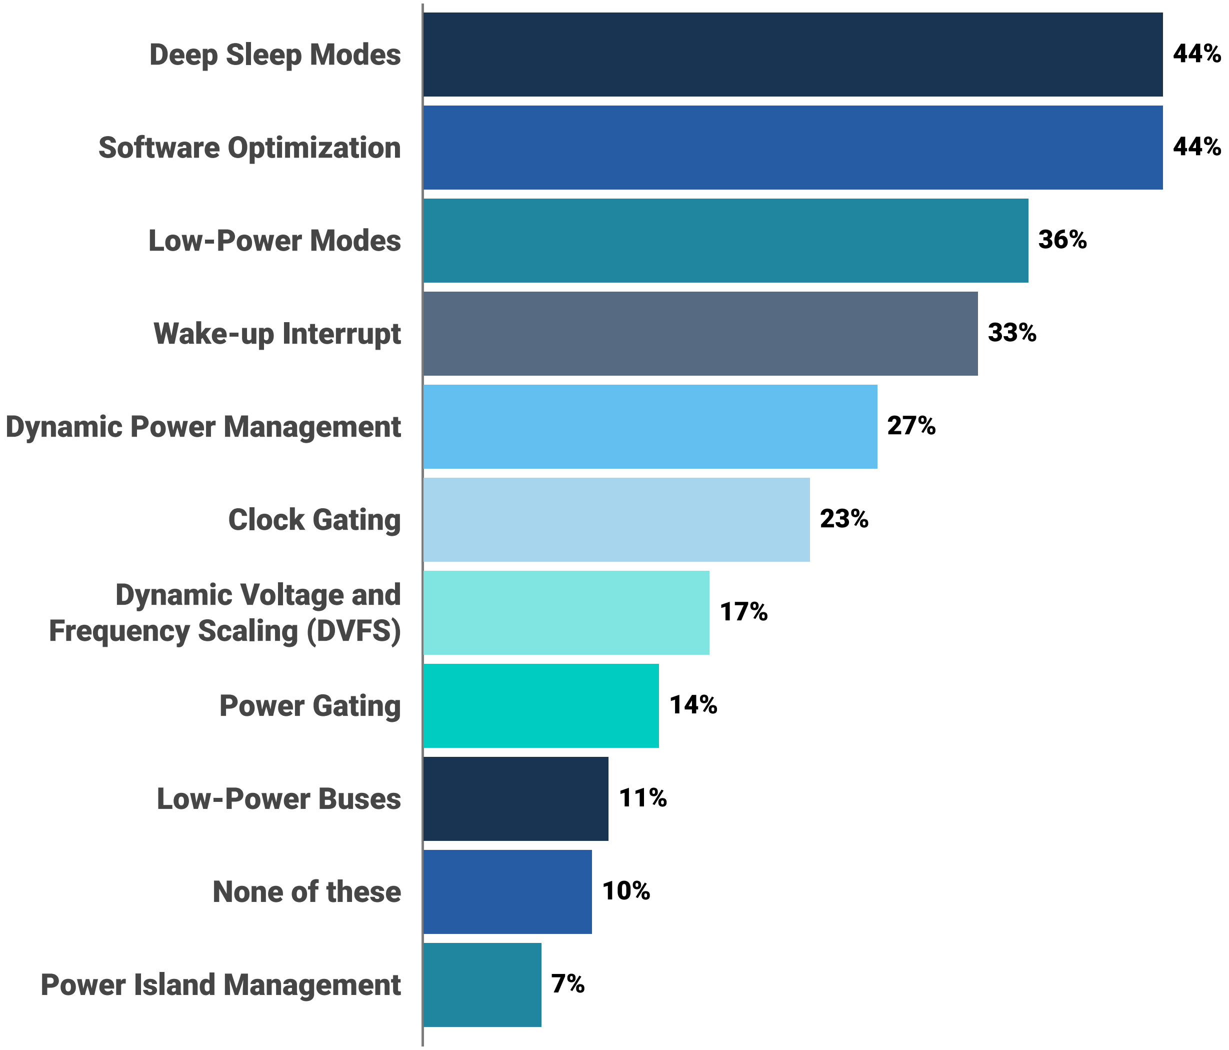

n=166

Power efficiency remains a central focus in embedded design, with most engineers employing multiple strategies to extend battery life and reduce energy use.

- Deep sleep modes and software optimization are the most common techniques, reflecting an emphasis on both hardware-level control and code efficiency.
- Low-power modes and wake-up interrupts follow, showing that many developers balance sleep states with responsiveness to external events.
- More advanced strategies such as dynamic power management and clock gating are less common but indicate a demand for fine-grained power control.

Only 10% report using none of these methods, confirming that power optimization is now a standard design consideration across almost all embedded projects.

Programming Languages

What type of design benefits are you realizing from your current utilization of AI techniques?

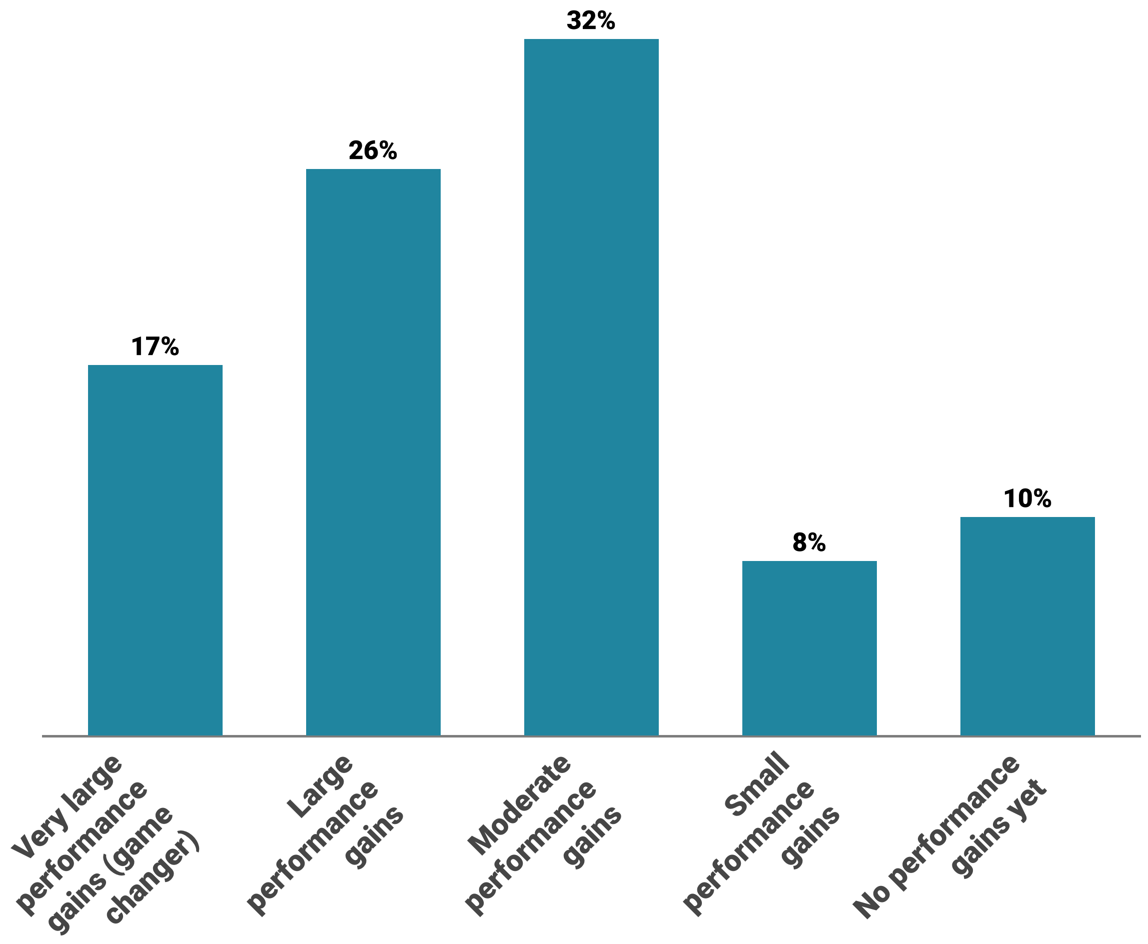

n=99

Most engineers are seeing tangible performance improvements from AI integration, though the scale of benefit varies.

- 75% of respondents report moderate to large gains, with 17% describing the impact as a “game changer.”
- Only 10% say they have not yet realized measurable performance improvements, suggesting that even early AI adoption delivers meaningful results in embedded applications.
- Smaller gains are likely tied to teams still in the optimization or proof-of-concept stage.

Overall, the data shows that AI is already delivering tangible performance value in embedded design, moving beyond experimentation toward measurable engineering outcomes.

Techniques & Performance Benefits - Key Takeaways

Embedded engineers are employing a blend of traditional and advanced design techniques to optimize for both performance and efficiency. C and C++ remain foundational languages in embedded development, underscoring their continued dominance in performance-critical environments. However, the ecosystem is evolving: Python and MicroPython are increasingly used to prototype, test, and extend functionality, signaling a stronger overlap between embedded and data-centric workflows.

On the AI side, developers are experimenting broadly across methodologies. k-Nearest Neighbors, evolutionary algorithms, and convolutional neural networks top the list, showing that engineers are focusing on proven techniques rather than bleeding-edge experimentation. These algorithms are being applied to sensor data fusion, anomaly detection, and data preprocessing, with additional applications in computer vision and predictive maintenance. Engineers are using AI to improve system responsiveness and quality control rather than automate human decision-making.

When it comes to optimization, power management and efficiency benchmarking remain central to design success. Roughly one-third of engineers use power analysis tools such as Saleae Logic or Nordic PPK, and another 44% report using deep sleep or software-based power reduction modes. Most describe their AI-related performance improvements as moderate to large, demonstrating that embedded AI is delivering measurable results in both speed and efficiency.

Techniques & Performance Benefits - Implications for Marketers

These findings reinforce that embedded developers respond best to data-backed, application-focused documentation. Engineers want to see real-world examples of how a chip, framework, or algorithm performs under specific workloads, particularly regarding latency, power efficiency, and model accuracy. Marketing that focuses on outcomes rather than buzzwords resonates most. For instance, content showing how a specific microcontroller achieves sub-millisecond inference or extends battery life by a measurable amount will outperform generic claims of “AI readiness.”

The combination of traditional coding languages and AI techniques also presents a strong opportunity for education-driven marketing. Marketers can bridge the gap between what engineers already know (C/C++) and what they are learning (Python and machine learning frameworks) by offering practical guides, sample code, and application notes. Brands that simplify integration across these toolchains can position themselves as trusted partners in engineers’ ongoing skill development.

Power and performance optimization represent another high-value messaging area. Engineers are motivated by efficiency, and every watt, millisecond, and megabyte matters. Highlighting measurable results through real benchmark data, power-usage metrics, and reproducible testing scenarios provides credibility and differentiation. Marketers can strengthen this further by framing their technology as part of the engineer’s toolkit—a reliable enabler that simplifies optimization without sacrificing performance.

Finally, these findings show that AI adoption within embedded design is incremental and evidence-driven. Marketing strategies should emphasize performance proof over innovation hype. Engineers do not want to be sold on the concept of AI; they want to see how it works, what it saves, and how it scales. The brands that meet that need with clarity, transparency, credibility, and technical depth will be the ones engineers turn to when designing next-generation systems.

AI Tools & Frameworks

Cloud Services

Which (if any) of the following cloud-based services are you currently using?

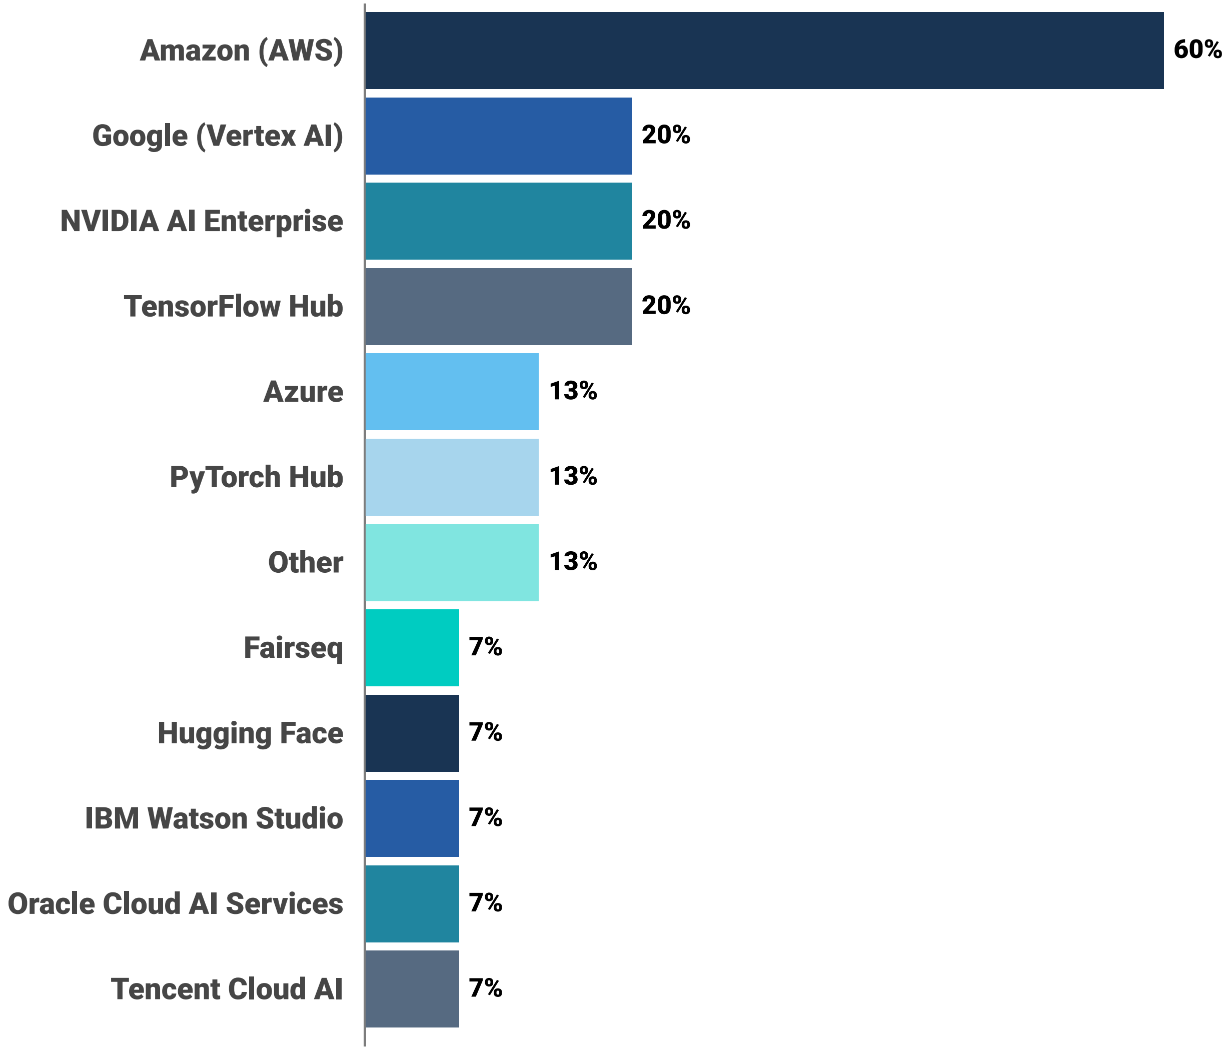

n=15

Among the small group of developers currently using cloud-based AI services, Amazon Web Services dominates the landscape.

- Secondary use is spread across multiple providers—Google Vertex AI, NVIDIA AI Enterprise, and TensorFlow Hub—suggesting no clear challenger to AWS's leadership.
- Adoption across smaller platforms such as Azure, PyTorch Hub , and others remains limited, reflecting both the specialized nature of embedded AI projects and the small sample size for this question.

While the base is small, the data reinforces that AWS is the primary entry point for cloud-based AI development in embedded systems, with engineers selectively leveraging other platforms depending on their toolchain or ecosystem.

AI Frameworks

Which AI frameworks or development environments do you primarily use for creating and training your embedded AI models?

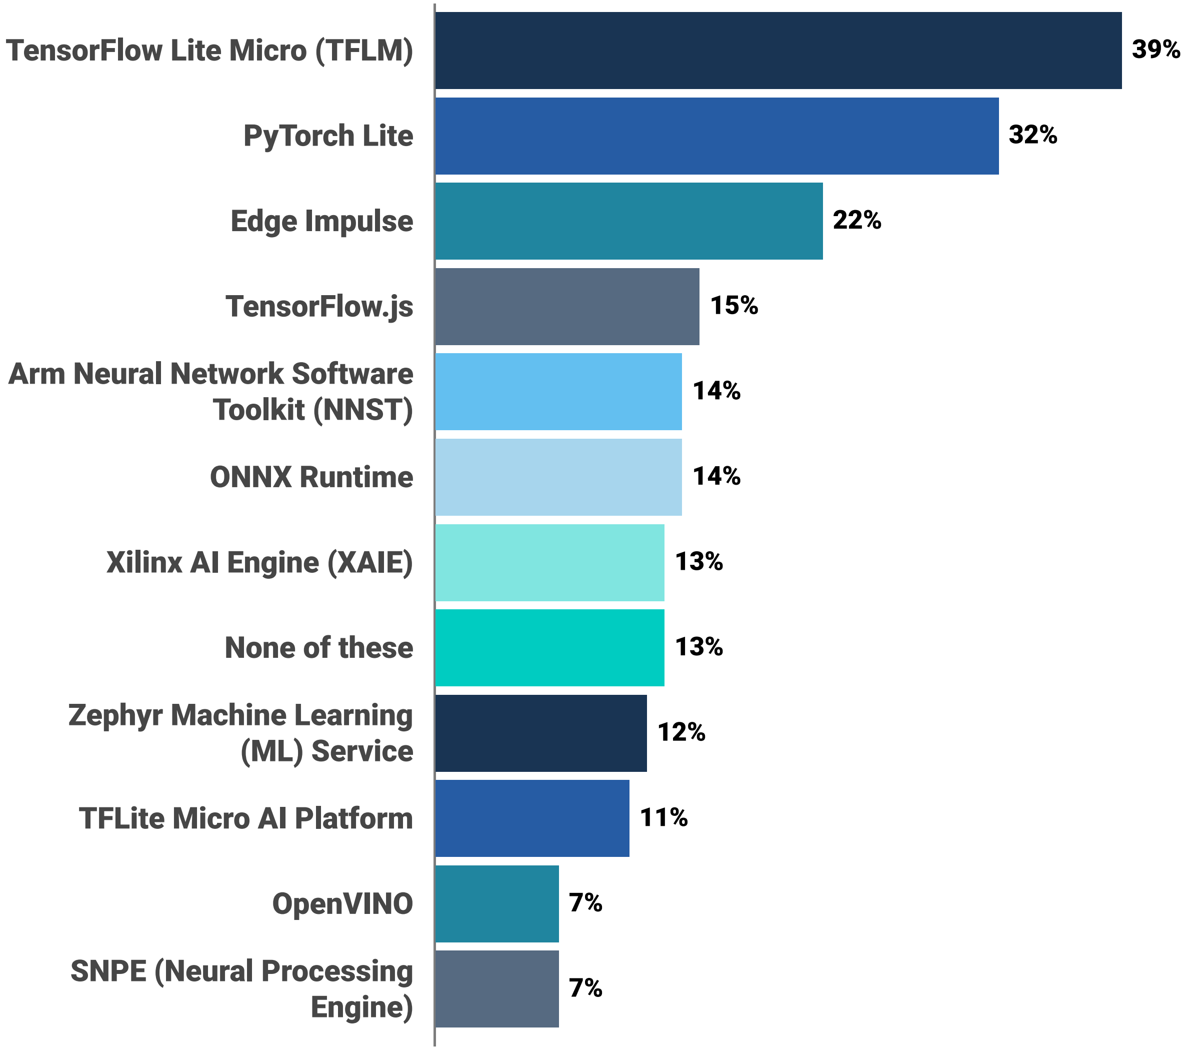

n=98

Developers rely heavily on TensorFlow Lite Micro and PyTorch Lite as their primary embedded AI frameworks, reflecting the dominance of these two ecosystems across both training and deployment.

- Edge Impulse has also gained meaningful traction, particularly among teams seeking rapid prototyping and model deployment tools tailored for IoT and low-power devices.
- Secondary use of TensorFlow.js, Arm NN Toolkit, and ONNX Runtime suggests a growing interest in flexible, hardware-agnostic options that simplify edge integration.
- About 13% report using none of the listed platforms, indicating that some developers rely on custom or proprietary frameworks.

The findings show that TensorFlow and PyTorch remain prominent in embedded AI development, while newer tools like Edge Impulse are helping engineers streamline workflows and cut time-to-deployment.

Tools & Techniques

Which of the following tools or techniques do you employ to optimize and compress AI models specifically for the constrained resources of embedded systems?

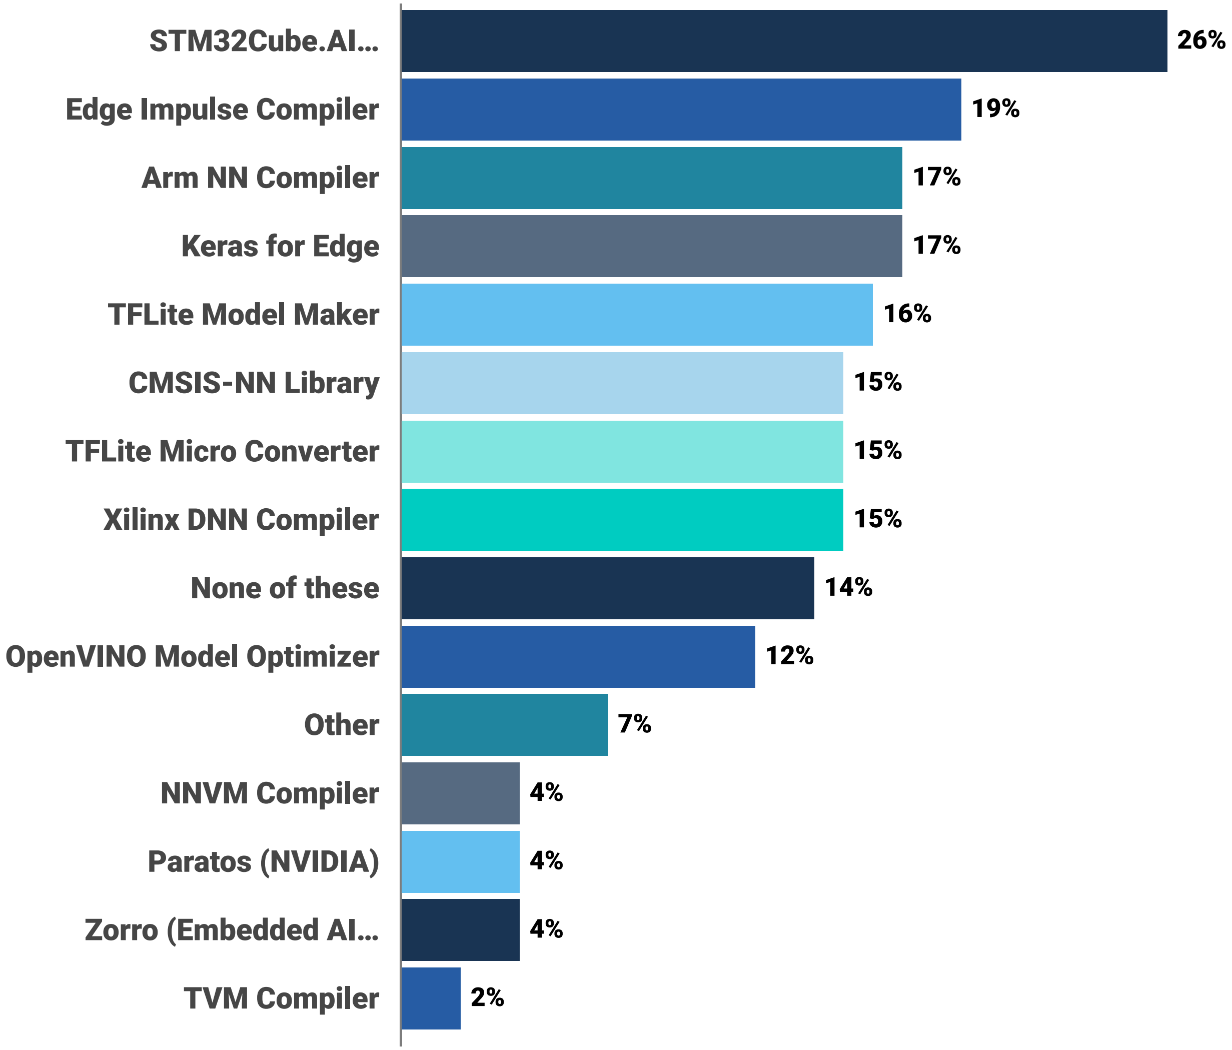

n=95

Optimization for constrained hardware is a critical step in embedded AI workflows, and engineers rely on a variety of vendor- and framework-specific tools to achieve it.

- STM32Cube.AI is the most widely used, reflecting STMicroelectronics' strong ecosystem presence among embedded developers.
- Edge Impulse Compiler, Arm NN Compiler, and Keras for Edge form a second tier of commonly used solutions that simplify model conversion and quantization for low-power devices.
- Roughly 15% use TensorFlow-based or Xilinx tools, showing diversity in optimization approaches depending on hardware platforms.
- 14% report using none of these tools, suggesting that some teams still rely on manual tuning or possible proprietary methods for model compression.

The results highlight that model optimization is a widely recognized but fragmented space, with no single dominant tool—developers select platforms that align most closely with their chosen hardware and deployment environments.

Open Source Model Optimization

In which (if any) of the following ways are you optimizing your open-source models?

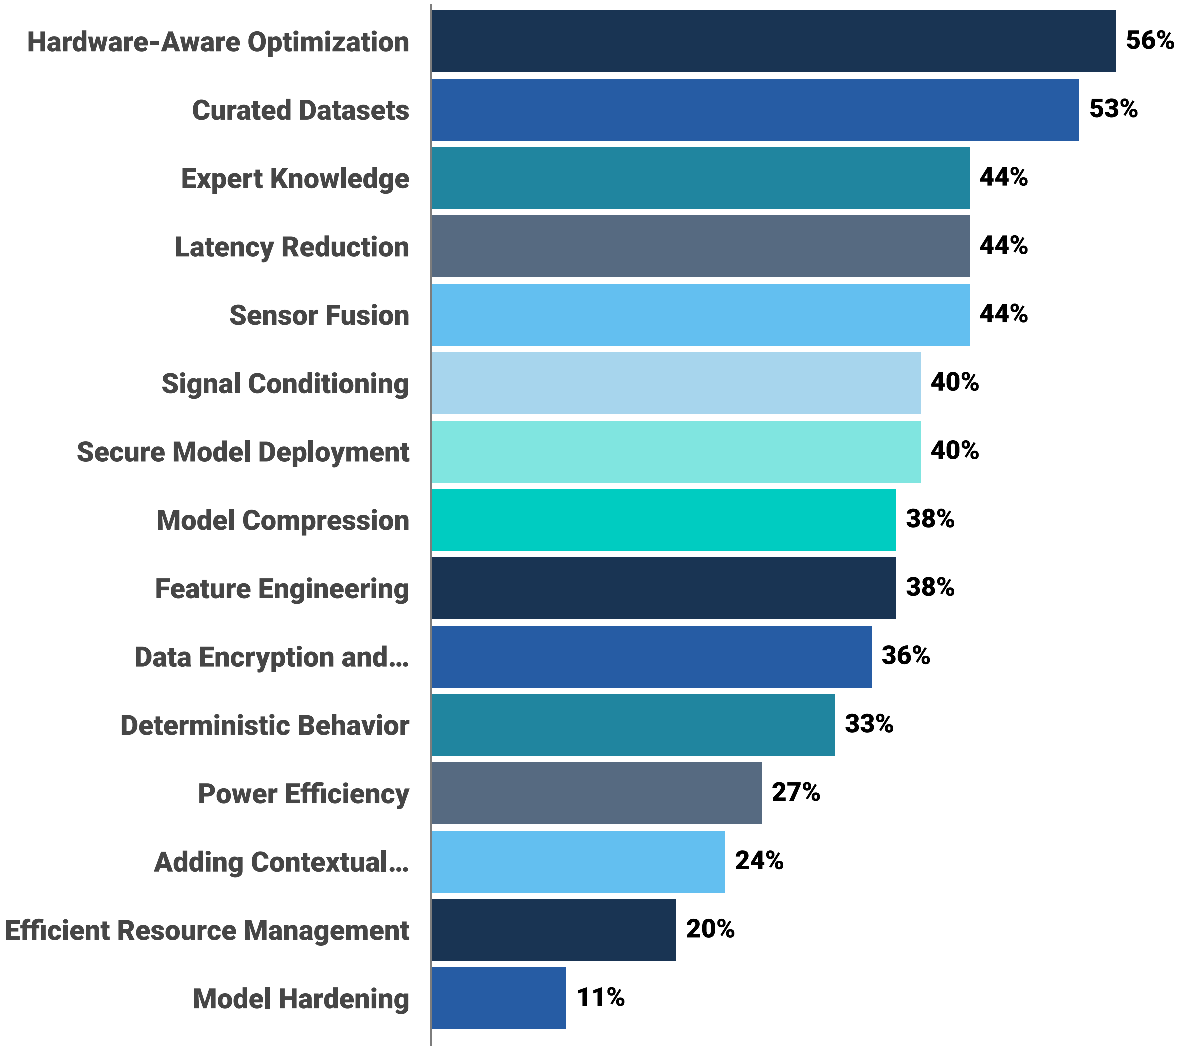

n=55

Engineers are applying a wide range of optimization methods to adapt open-source AI models for embedded environments, with a strong focus on hardware performance and data quality.

- The most common practices are hardware-aware optimization and curated datasets, showing that developers prioritize tailoring models to their specific devices and use cases.
- Roughly four in ten use expert knowledge, latency reduction, or sensor fusion, reflecting a practical focus on real-time responsiveness and integration with physical systems.
- Secondary priorities include signal conditioning, secure deployment, and model compression, all aimed at balancing model accuracy with reliability and efficiency.

Overall, the results show that optimization in embedded AI is multi-faceted, blending hardware tuning, data engineering, and security practices to make open-source models viable in resource-constrained environments.

Open Source Models

Do you use open-source models (foundation models) as a starting point for your embedded AI projects? (For example models found on Hugging Face)

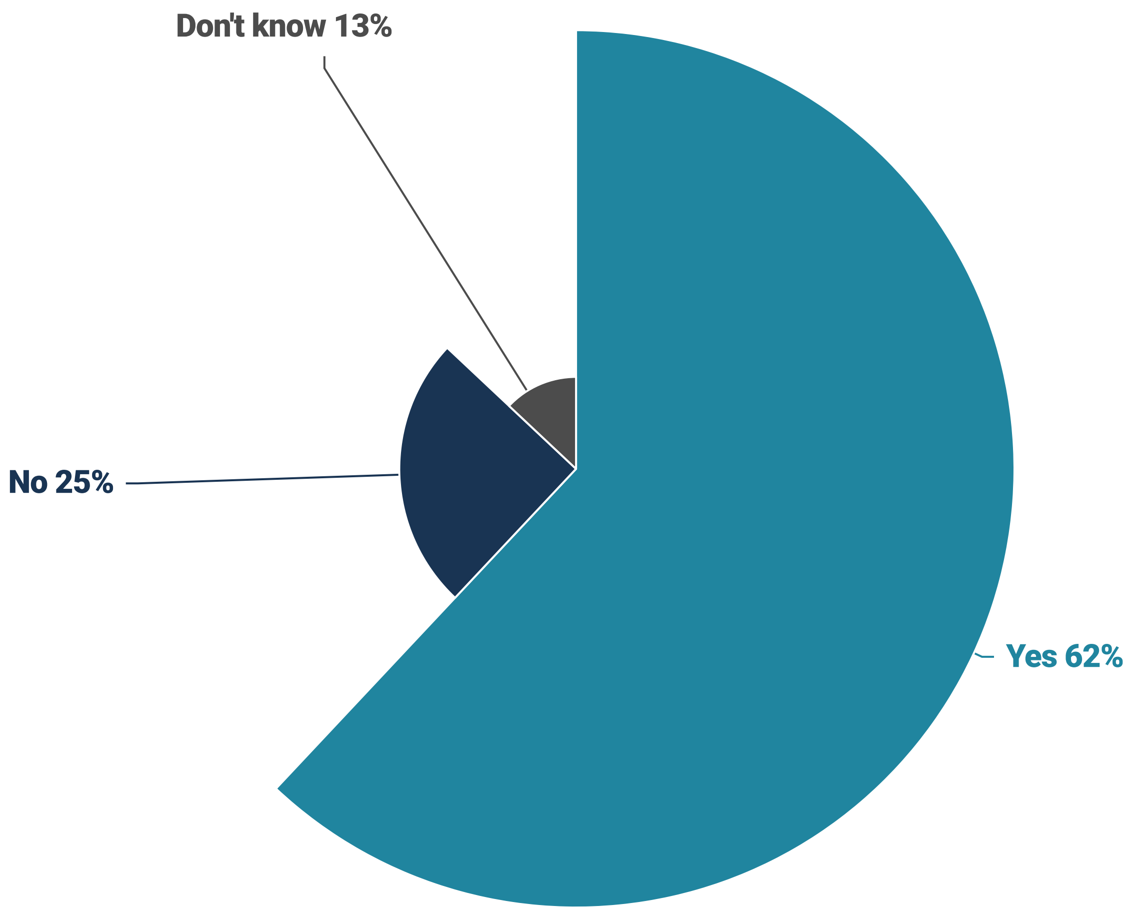

n=99

A clear majority of embedded developers are leveraging open-source foundation models as the basis for their AI projects.

- 62% report using publicly available models—such as those hosted on Hugging Face—to accelerate development, reduce training time, and take advantage of community-tested architectures.
- 25% do not currently use open-source models, while 13% are unsure, likely reflecting teams that rely on pre-integrated frameworks or vendor-provided tools without direct visibility into their origins.

The data confirms that open-source models have become a mainstream entry point for embedded AI, enabling faster prototyping and democratizing access to advanced machine learning capabilities.

Primary Open Sources

What are your primary sources for open-source or pre-trained models as a starting point for your embedded AI projects?

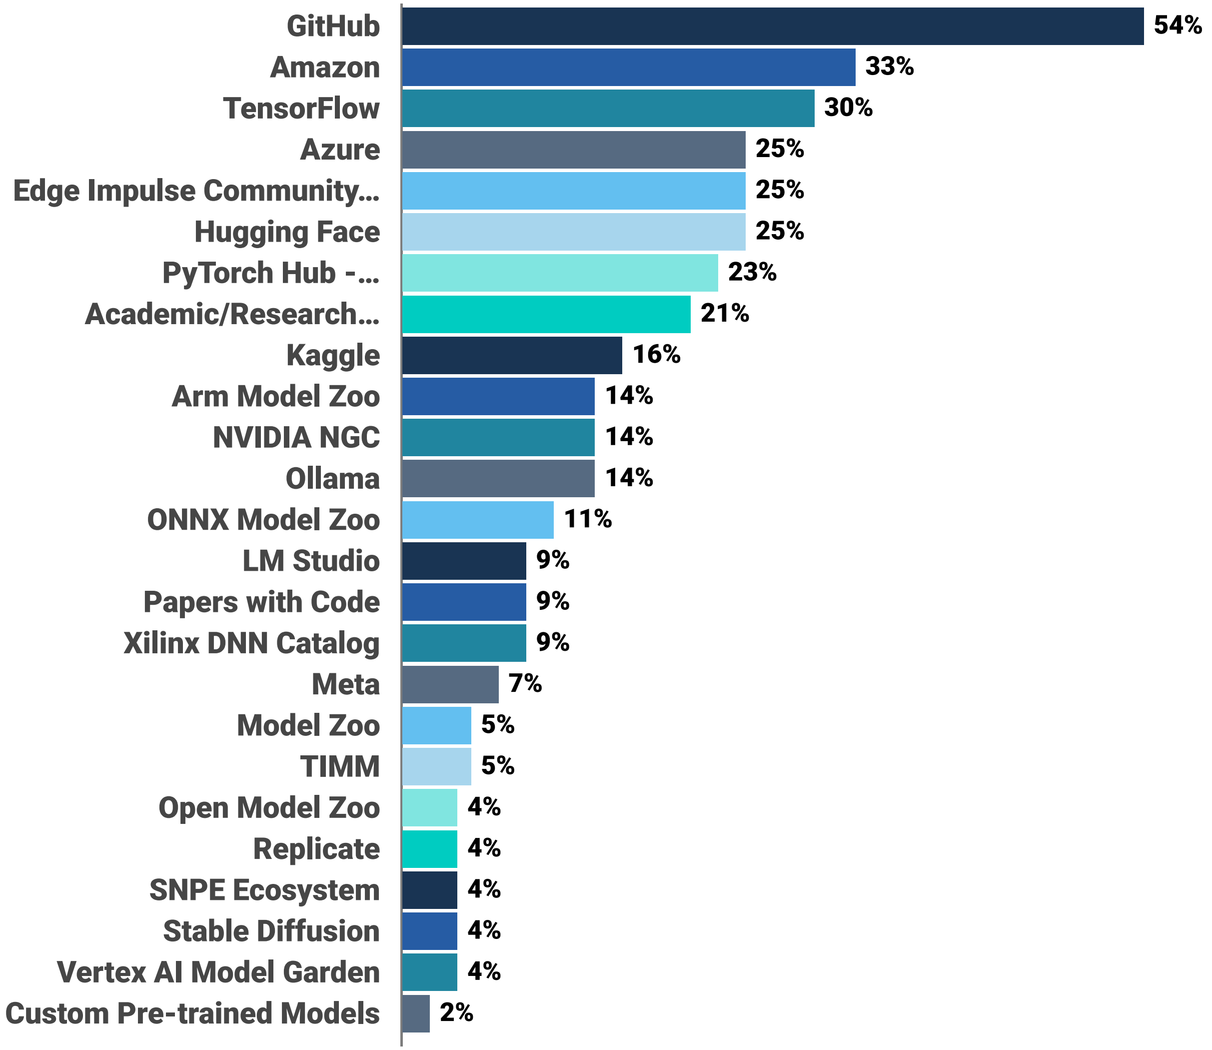

n=57

GitHub is by far the leading source for open-source or pre-trained models, reflecting its role as the central repository for code sharing and collaboration across the embedded AI community.

- Other widely used sources include AWS, TensorFlow Hub, and a trio of Azure, Edge Impulse, and Hugging Face, showing strong adoption of both cloud-based and community-driven ecosystems.
- A mix of academic and vendor-specific repositories (including PyTorch Hub, NVIDIA NGC, and Arm Model Zoo) demonstrates that engineers draw from diverse sources depending on performance requirements and preferred framework

The findings show that embedded AI work is highly collaborative, as engineers draw on open communities, cloud tools, and vendor platforms to accelerate development.

Prototyping & Testing

What hardware platforms or development boards do you commonly use for prototyping and testing your embedded AI solutions?

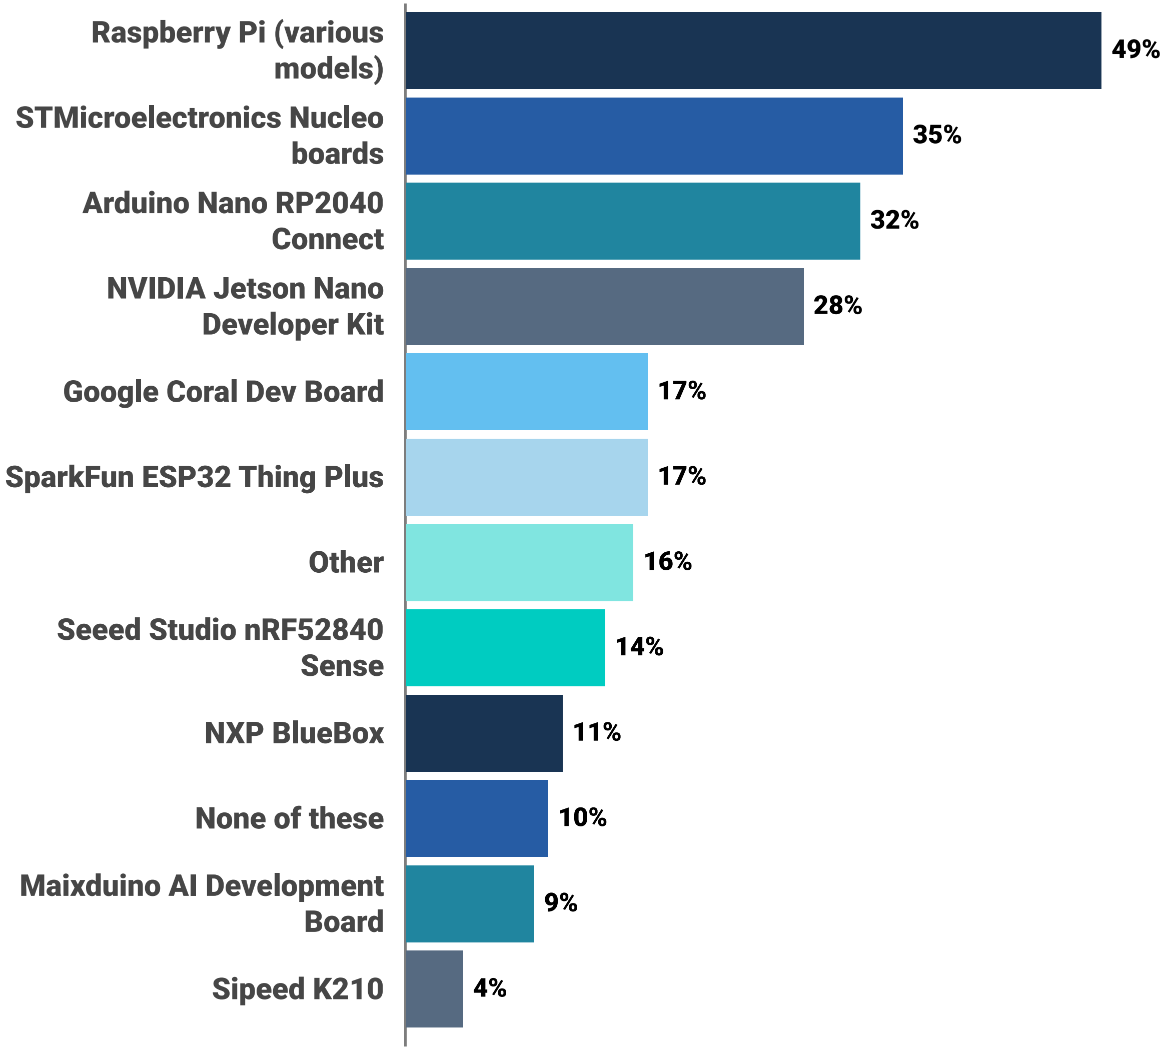

n=94

Raspberry Pi dominates as the preferred prototyping platform for embedded AI, valued for its flexibility, affordability, and extensive software ecosystem.

- STMicroelectronics Nucleo and Arduino Nano RP2040 Connect follow closely, confirming their popularity for rapid testing and hardware evaluation in academic and industrial environments.
- NVIDIA Jetson Nano is also represented strongly, particularly for projects requiring higher-performance edge AI processing.
- Smaller but notable use of Google Coral, SparkFun ESP32, and Seeed Studio nRF52840 highlights the variety of low-power and specialized platforms engineers choose for proof-of-concept work.

The diversity of hardware used for prototyping proves how embedded AI developers rely on a mix of general-purpose, vendor-specific, and AI-optimized boards to balance cost, performance, and accessibility during early-stage design.

Data & Sensors

How do you handle data collection, labeling, and preprocessing for your embedded AI projects, particularly when real-world sensor data is involved?

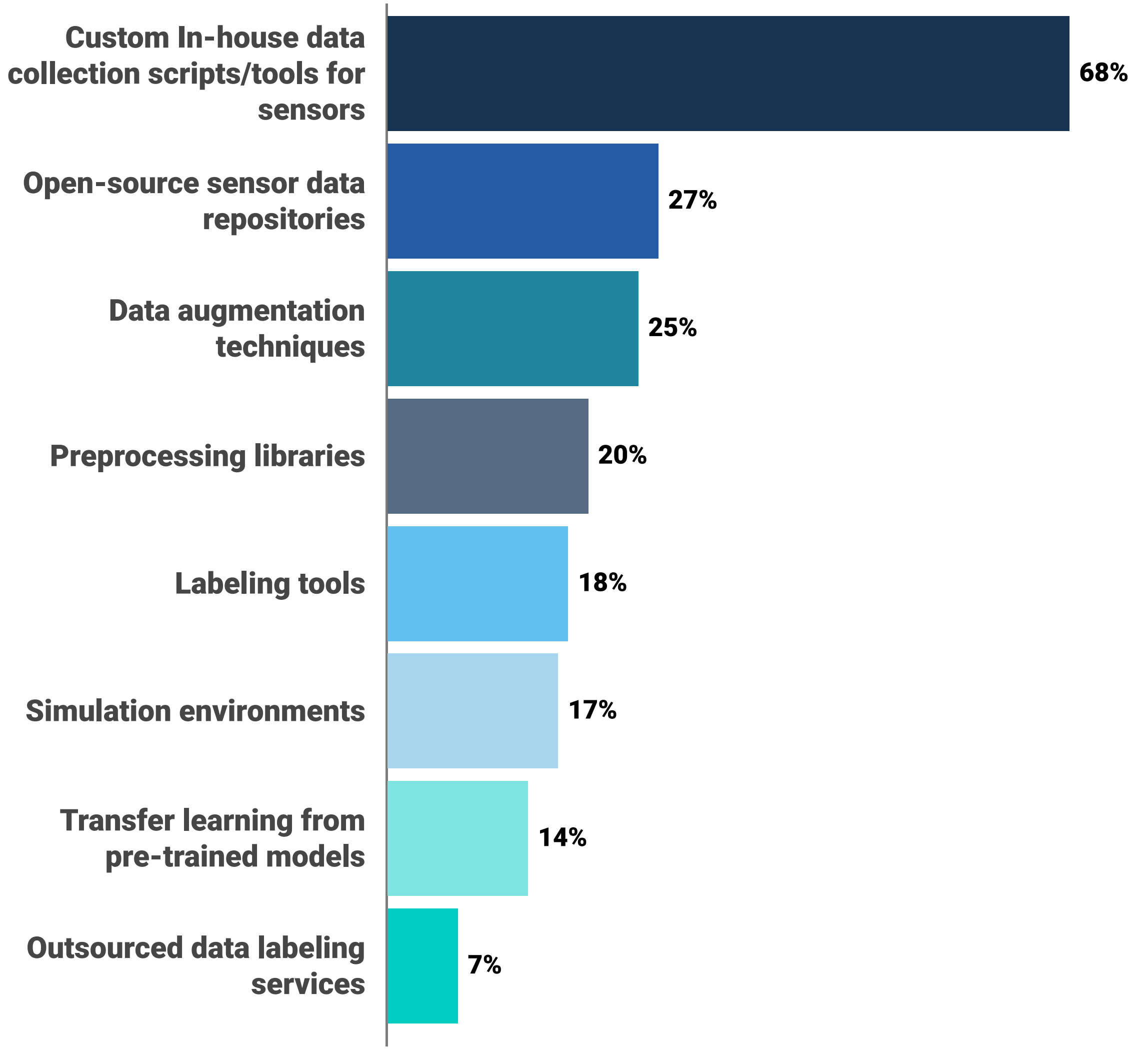

n=92

Data collection, labeling, and preprocessing in embedded AI remain largely an in-house effort, driven by the need for precise and application-specific sensor data.

- 68% of respondents build and manage custom data collection scripts and tools, reflecting the unique and hardware-dependent nature of embedded applications.
- 27% use open-source sensor repositories, and 25% apply data augmentation techniques, showing a growing reliance on hybrid datasets to supplement limited real-world samples.
- Only 18% use formal labeling tools and 7% outsource data labeling, suggesting that annotation and dataset preparation are still primarily internal, manual processes.
- Smaller adoption of simulation and transfer learning indicates that while these methods are known, they remain secondary to direct, hardware-based data acquisition.

These results highlight that data collection is a critical bottleneck in embedded AI development, with most engineers opting to build their own pipelines rather than rely on off-the-shelf solutions.

AI Tools & Frameworks - Key Takeaways

Embedded developers are working with a broad mix of frameworks and platforms to bring AI into production, showing a clear preference for accessible, lightweight tools that balance performance and usability. TensorFlow Lite Micro and PyTorch Lite are the most commonly used, followed by Edge Impulse and Arm Neural Network Software Toolkit. These results suggest that engineers are not only experimenting with AI but actively selecting tools that simplify deployment on constrained hardware.

Cloud integration remains an important part of the development workflow. Among those using external services, Amazon Web Services is by far the most common, with smaller shares using Google Vertex AI, NVIDIA AI Enterprise, and Tensorflow Hub. The cloud remains a valuable tool for model training and optimization, even as inference continues to move toward the edge.

Developers are also focused on optimization and reuse. Twenty-six percent rely on vendor compilers such as STM32Cube.AI, while others use Edge Impulse Compiler or Arm NN Compiler to streamline models for embedded environments. A majority report using open-source or foundation models, most commonly from GitHub, AWS Model Zoo, and TensorFlow Hub. Engineers are combining cloud flexibility, open-source collaboration, and hardware-aware tuning to accelerate their development cycles.

AI Tools & Frameworks - Implications for Marketers

For semiconductor and technology marketers, this section provides a clear signal: engineers value tools that make AI integration practical, transparent, and resource-efficient. Marketing should position frameworks, SDKs, and compilers not as advanced or complex features, but as productivity multipliers that help engineers reach deployment faster. Tutorials, workflow diagrams, and code examples that shorten time-to-inference will be far more effective than general claims of capability.

Open-source visibility plays a central role in how engineers choose frameworks. A strong brand presence on GitHub, TensorFlow Hub, and ONNX communities builds credibility and discoverability. Engineers often find solutions organically through repositories and technical forums rather than through advertising. Marketers who invest in high-quality open documentation, public repositories, and sample projects will reach engineers at the exact moment they are evaluating tools.

Because optimization is a recurring theme, marketers should also communicate hardware-specific performance advantages. Engineers want clear data on how a model behaves on particular boards, processors, or accelerators. Performance benchmarks, latency comparisons, and case studies showing quantifiable improvements create immediate engagement and trust.

Finally, this group’s continued use of cloud platforms like AWS and Azure highlights the opportunity to emphasize hybrid deployment flexibility. Campaigns that show how a solution bridges cloud training with edge inference will resonate strongly. Engineers want options, not restrictions, and they favor technologies that scale seamlessly from prototype to production. Marketers who provide that clarity—through examples, not slogans—will be the ones who turn interest into design wins.

Challenges, Data & Security Constraints

Technical Obstacles

What are the key technical obstacles hindering further adoption of AI in your embedded projects?

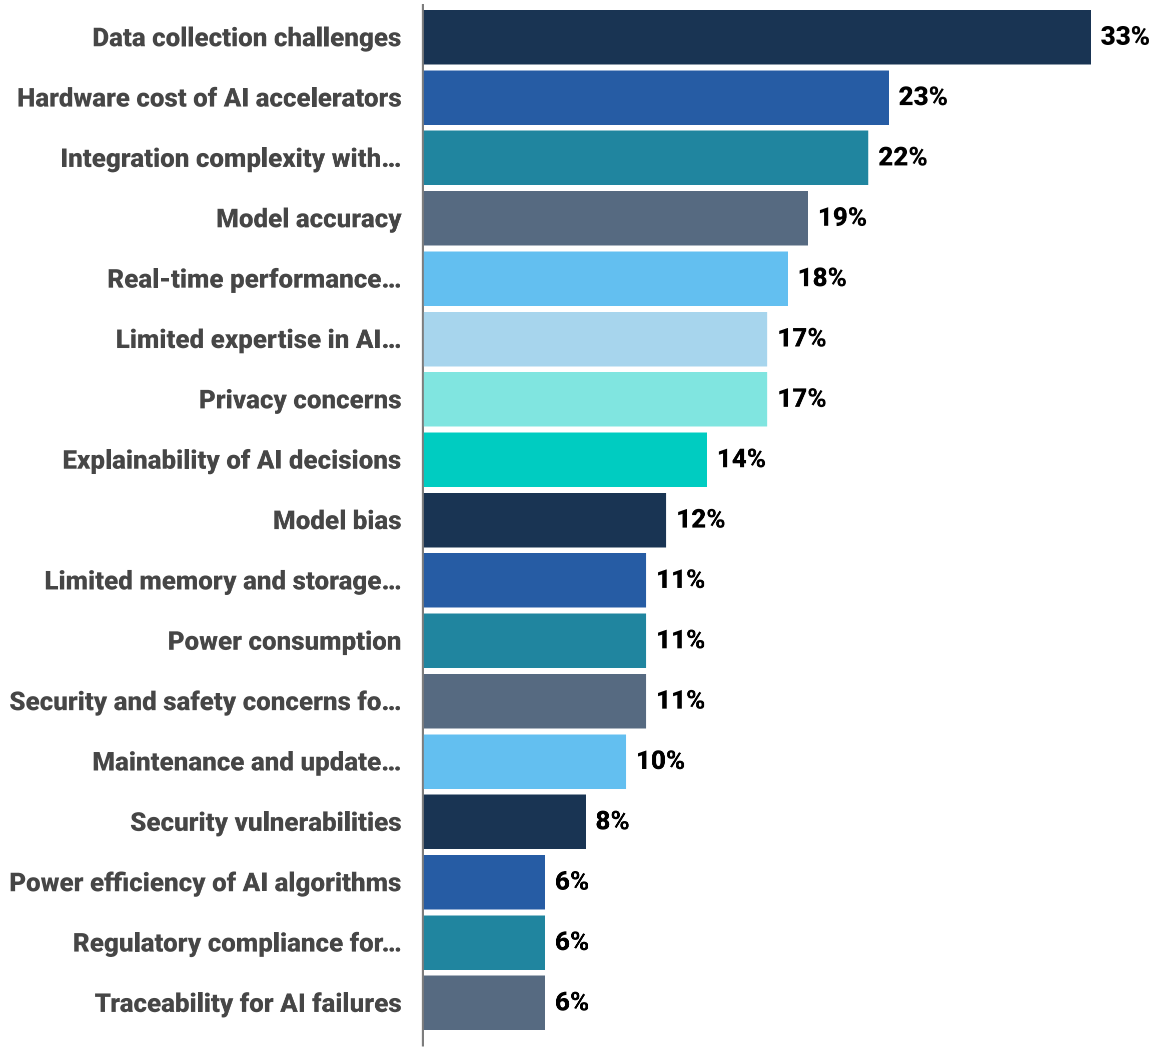

n=90

Engineers cite data and hardware barriers as the most significant obstacles to wider AI adoption in embedded systems.

- Data collection tops the list, reinforcing that access to high-quality, domain-specific datasets remains a critical challenge.
- Hardware cost and integration complexity follow, highlighting both economic and technical friction in scaling AI capabilities across platforms.
- Secondary issues, such as model accuracy, real-time performance and limited expertise, point to skill and optimization gaps that slow progress.
- Concerns like privacy, explainability, and potential model bias reflect growing awareness of responsible AI practices.

The data paints a clear picture of an industry advancing in capability but still constrained by data availability, system integration, and cost-performance tradeoffs.

Security Concerns

Which (if any) of the following are security concerns you have had when developing your current embedded systems?

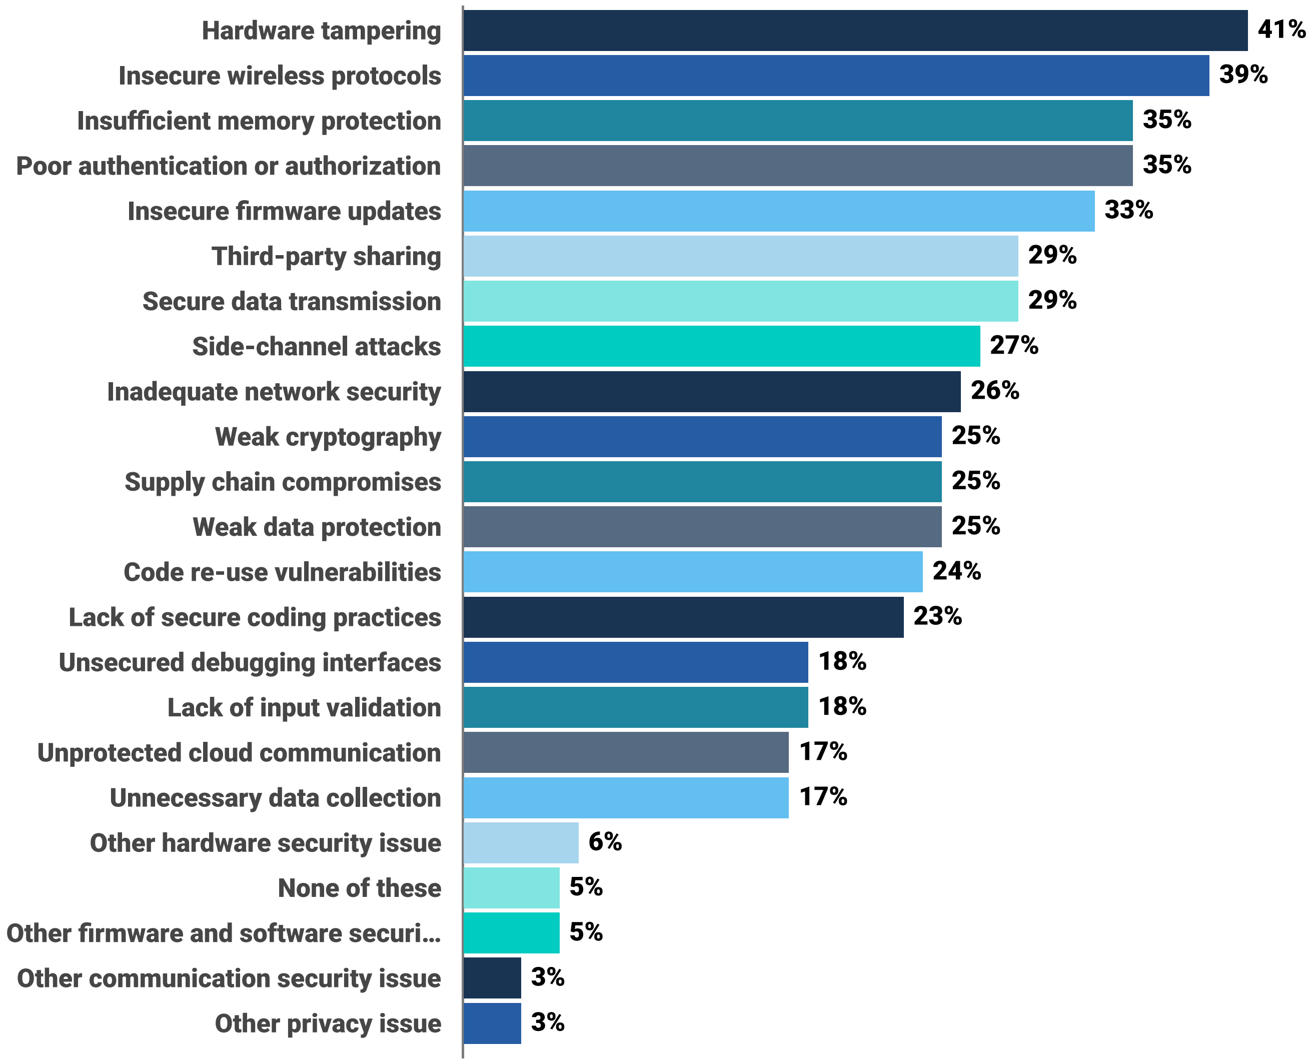

n=153

Security remains a core concern for embedded developers, with risks spanning hardware, firmware, and communication layers.

- The most common concerns are hardware tampering and insecure wireless protocols, both tied to the physical and connectivity vulnerabilities common in IoT and field-deployed devices.
- A cluster of software-related threats follows closely (memory protection, authentication , and firmware updates) highlighting persistent challenges in ensuring secure system design.
- Roughly one-quarter of respondents cite issues such as weak cryptography, supply chain compromises, and data protection, suggesting that software supply chain integrity and data privacy are equally important concerns.

Only 5% report having no security concerns, underscoring that embedded development now treats security as a universal design requirement, not a specialized consideration.

Explainability Frameworks

As AI plays a larger role in autonomous decision-making, which of the explainability frameworks do you anticipate being most important for public trust and regulatory compliance?

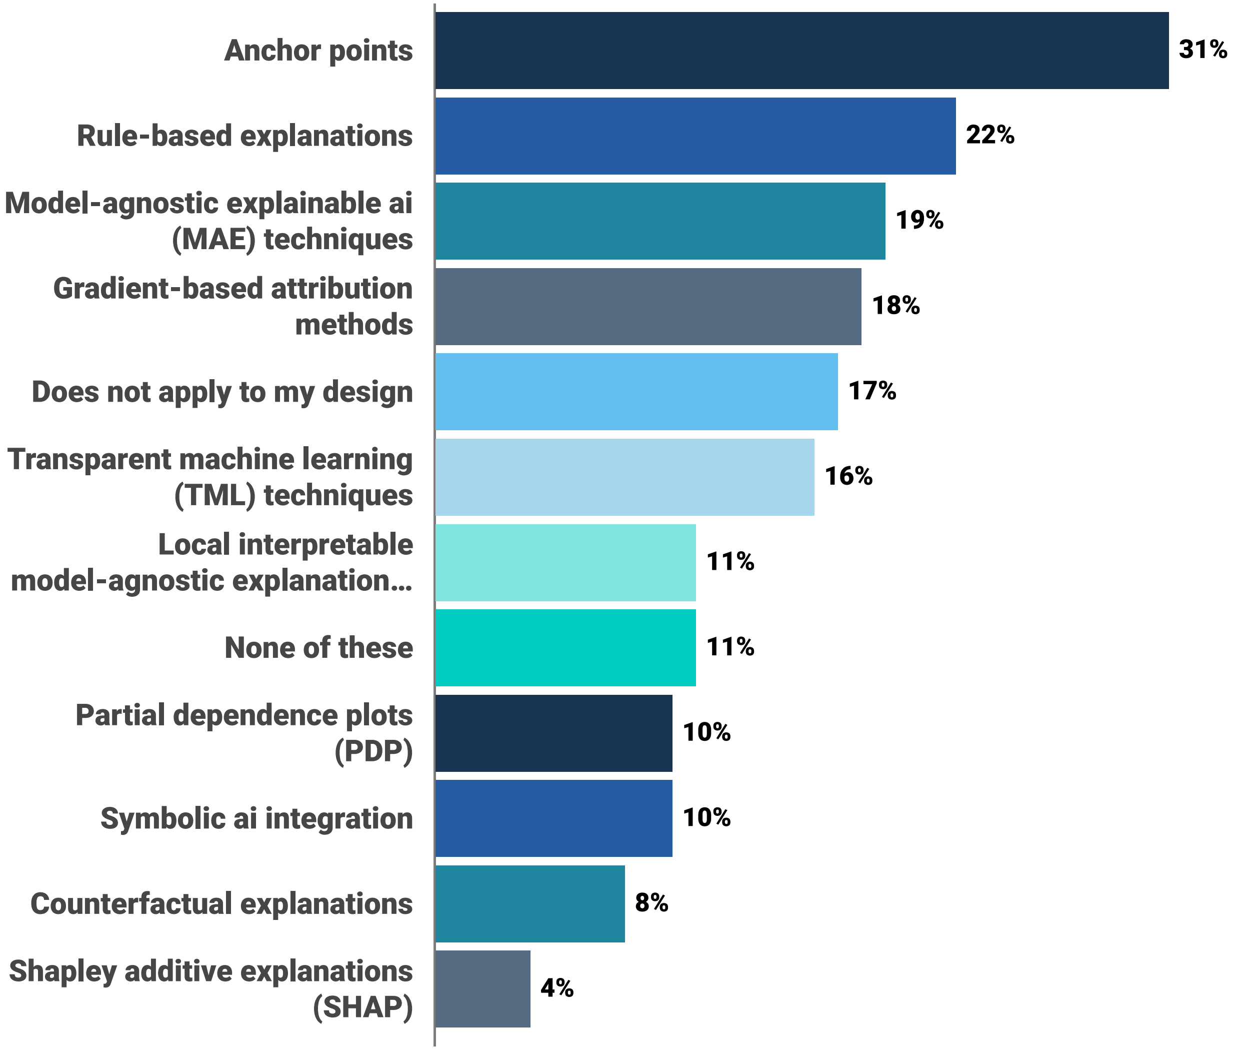

n=190

Engineers anticipate that human-readable and rule-based approaches will be most important for building public trust and meeting compliance standards in AI-driven systems.

- Anchor points lead as the top framework, reflecting the need for explanations that can be easily understood by end users, auditors, and regulators.
- Rule-based methods and model-agnostic techniques follow, showing a preference for transparent, flexible frameworks that can be applied across models and industries.
- A smaller segment cite gradient-based attribution methods, indicating continued reliance on visualization-driven interpretability tools for deep learning models.

The data underscores a clear industry direction: embedded engineers favor simple, interpretable, and model-agnostic methods that align with practical communication, certification, and compliance needs.

Data Collection & Labeling

To what extent is data collection and labeling for AI applications, particularly for edge cases and rare events, a problem issue to overcome in your embedded designs?

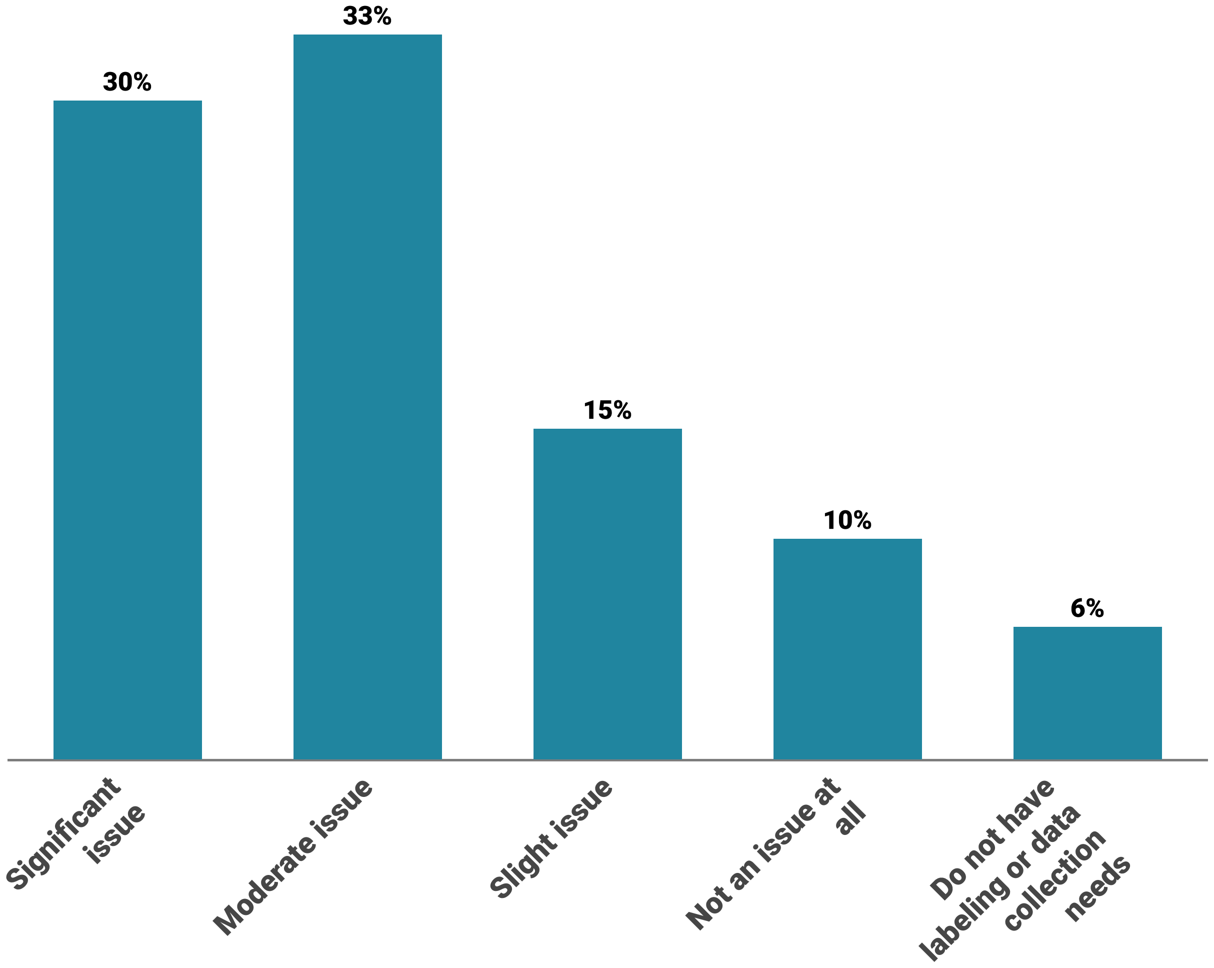

n=88

Data collection and labeling remain major challenges for embedded AI developers, especially when dealing with rare or unpredictable events.

- 63% of respondents identify it as a significant or moderate issue, confirming that obtaining diverse, high-quality datasets continues to be a leading barrier to AI performance.
- 15% see it as a slight issue, while only 10% say it's not an issue at all.
- A small group of 6% report having no data collection or labeling needs, likely due to projects that rely entirely on simulated data or pre-trained models.

These findings make it clear that data availability and labeling accuracy are still limiting factors in embedded AI development, particularly for edge-case detection and real-world validation.

AI + Cloud Integration

Have latency and connectivity limitations been an issue in your current embedded design?

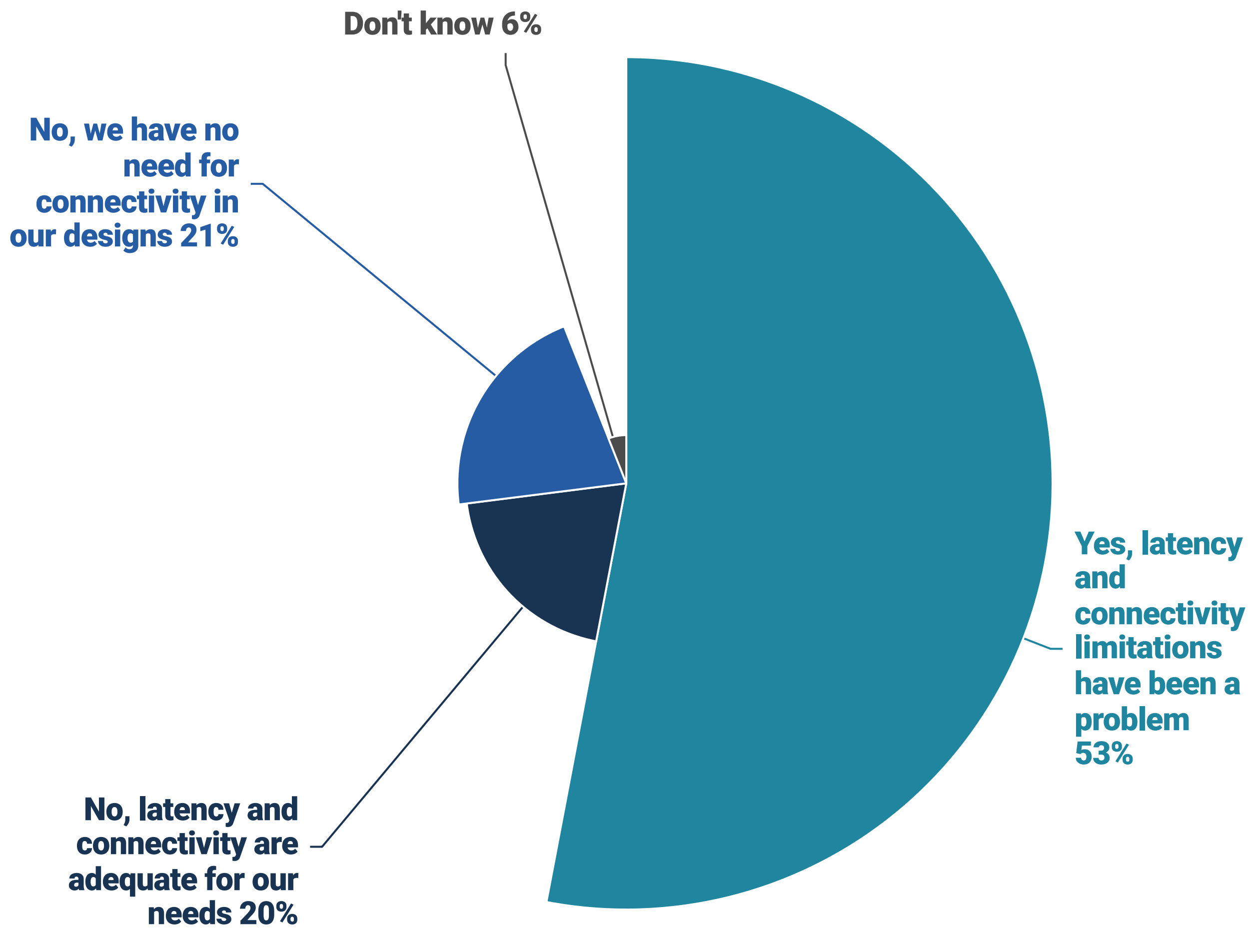

n=90

More than half of engineers report that latency or connectivity limitations have posed challenges in their current embedded designs.

- About 21% say that performance is adequate, while 22% indicate that connectivity is not required for their projects—often the case in fully self-contained or offline systems.
- The high incidence of reported issues underscores how network reliability and speed remain key design constraints, particularly for teams exploring edge–cloud integration or real-time AI processing.

These results reinforce the ongoing need for low-latency architectures and improved communication infrastructure to support next-generation embedded AI applications.

Edge vs. Cloud Balance

Thinking about the balance between cloud-based AI services vs. edge AI (deploying AI models directly on embedded devices), which approach do you currently take?

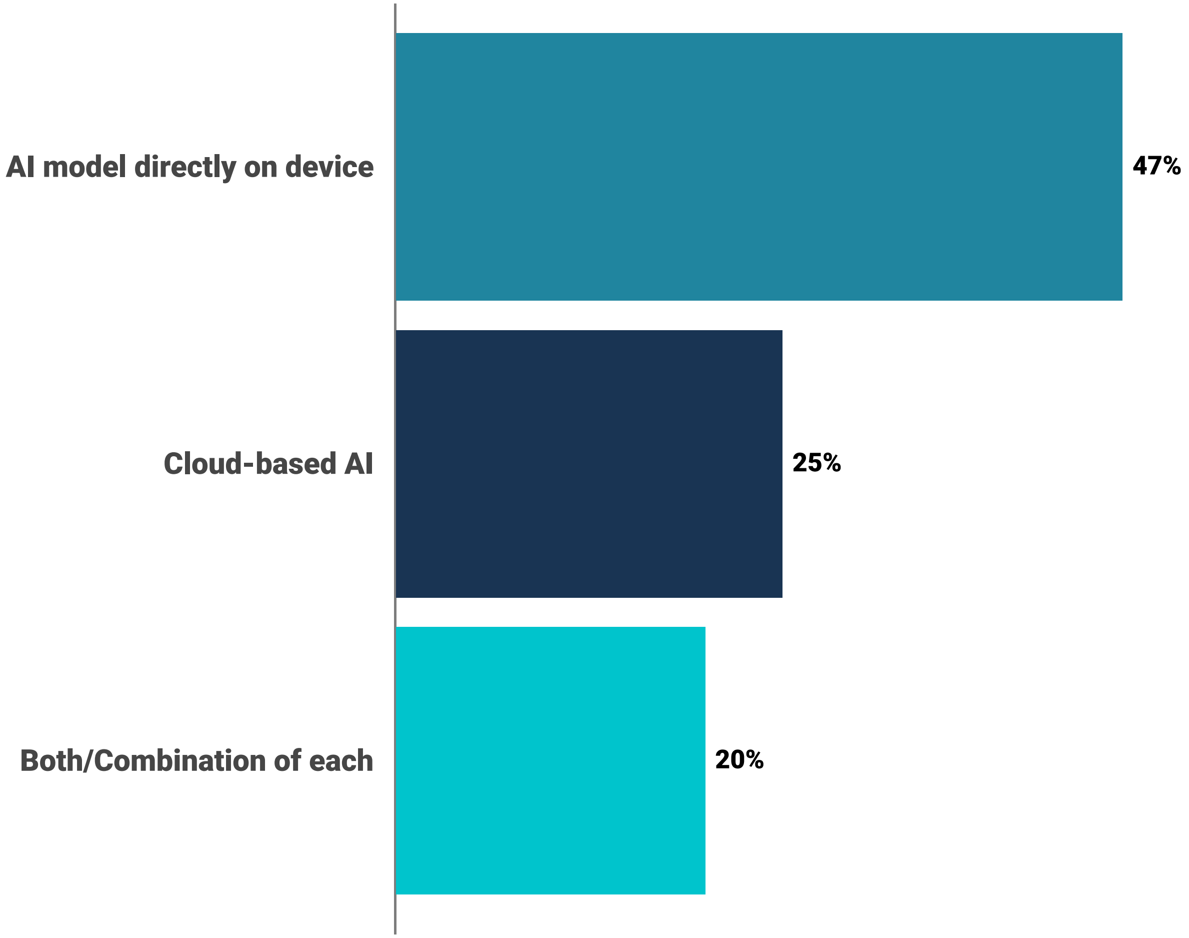

n=85

Edge AI leads the way as the dominant deployment strategy among embedded developers.

- Almost half (47%) currently deploy AI models directly on the device, favoring reduced latency, improved data privacy, and greater system autonomy.
- 25% rely primarily on cloud-based AI, while 20% use a hybrid approach, combining local inference with cloud training or analytics.

The data suggests that engineers are prioritizing real-time performance and independence from network connectivity, with cloud services playing a supporting role rather than serving as the default execution environment.

Future Balance

Thinking about the balance between cloud-based AI services vs. edge AI (deploying AI models directly on embedded devices), which do you anticipate will be used in the future?

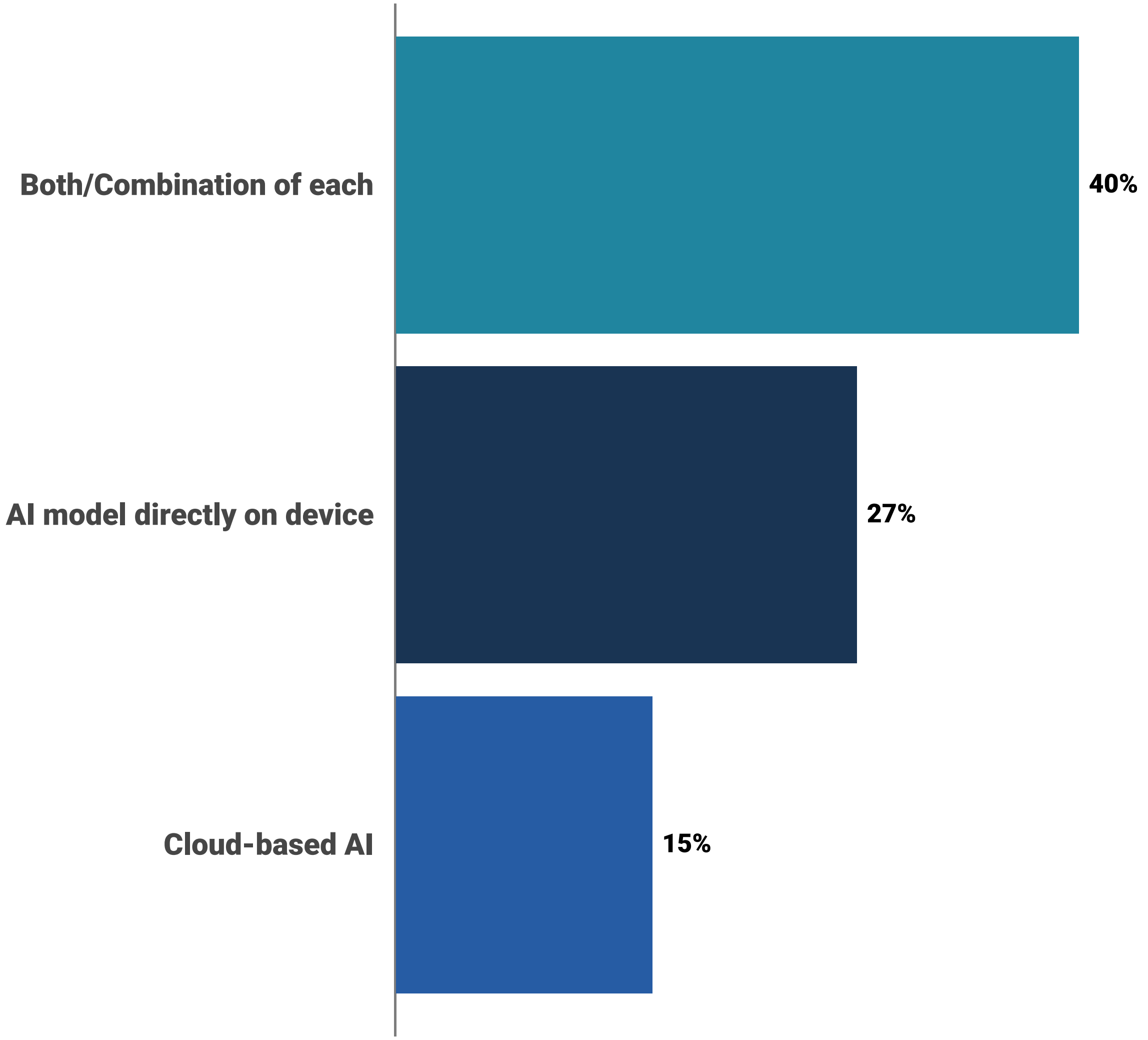

n=78

Engineers expect the future of embedded AI to be hybrid, blending the strengths of edge and cloud processing.

- 40% anticipate using a combination of edge and cloud approaches, signaling a shift toward distributed architectures that balance real-time responsiveness with scalable compute resources.
- 27% expect to continue focusing on on-device AI, underscoring the ongoing importance of autonomy and low-latency performance.
- Only 15% predict reliance on cloud-based AI alone, reflecting growing confidence in edge-capable hardware and optimization tools.

This outlook suggests that the industry is moving toward a collaborative edge–cloud ecosystem, where both environments play complementary roles in future AI deployments.

Shifting

In the next 1-3 years, what is the likelihood that 5G/6G technologies will shift your embedded AI designs to a more cloud-based AI approach?

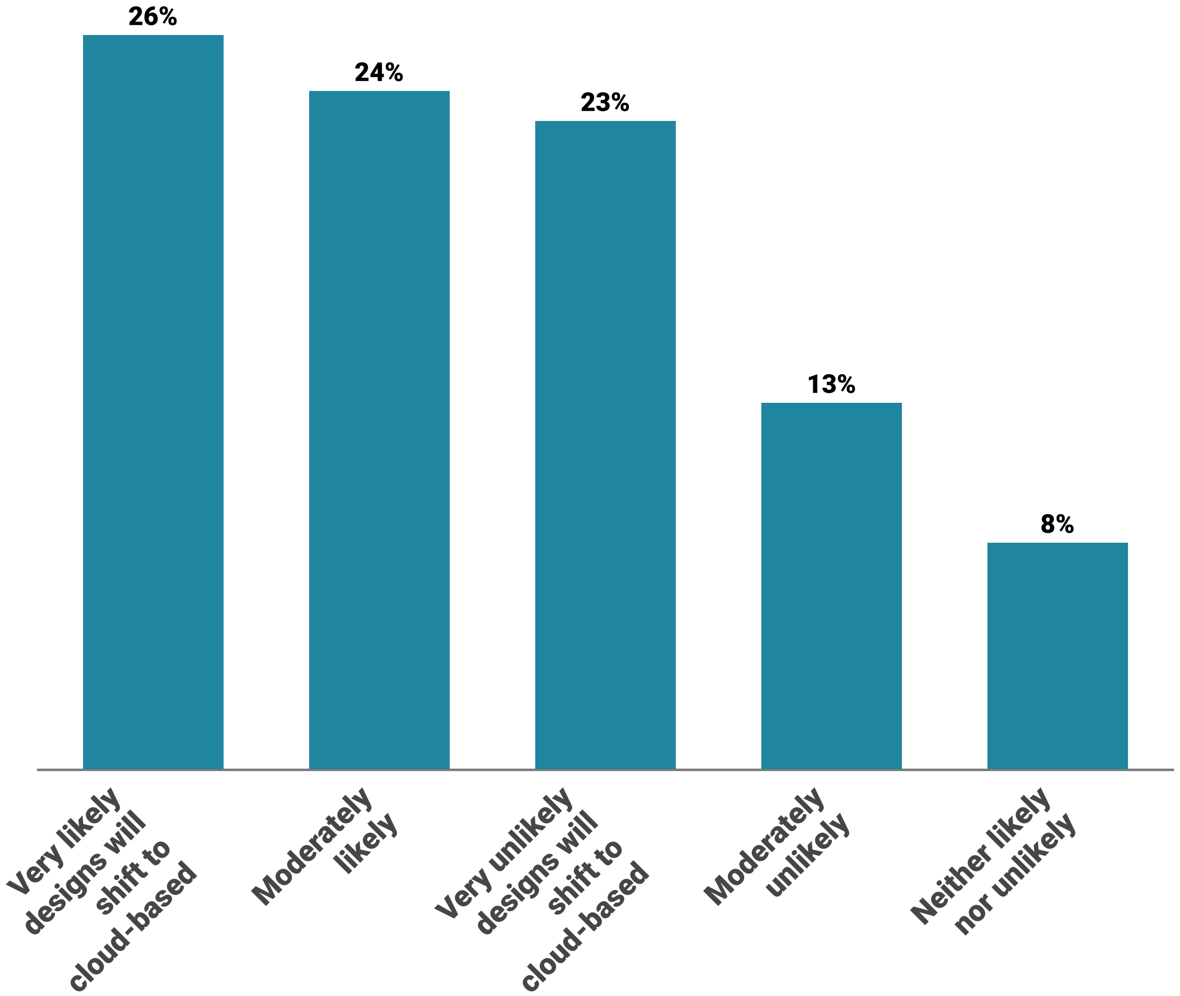

n=88

Engineers are divided on whether faster wireless technologies like 5G and 6G will drive a move toward more cloud-based AI designs.

- Half of respondents believe a shift is moderately or very likely, suggesting optimism that higher bandwidth and lower latency will make cloud integration more practical.
- At the same time, 36% consider such a shift unlikely, indicating that many expect to maintain edge-focused architectures despite connectivity improvements.
- Only 8% remain neutral, reinforcing that most developers already hold a clear view on where AI processing should occur.

The data highlights a maturing but cautious outlook—while 5G/6G will expand cloud potential, embedded engineers continue to value local control and responsiveness in system design.

Overcoming Obstacles

When do you anticipate solutions to the obstacles you selected above will be overcome?

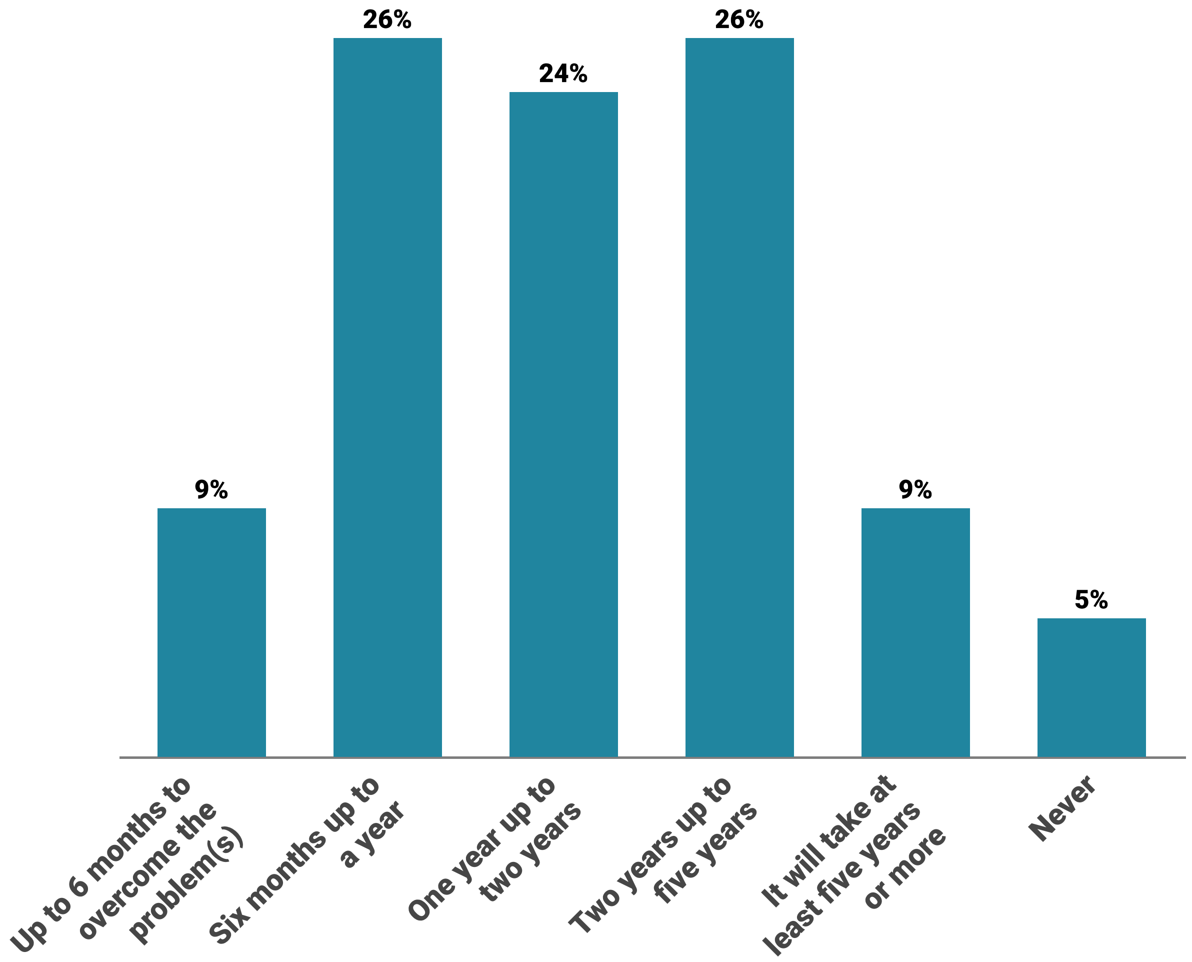

n=87

Most engineers (85%) believe that the technical and operational challenges limiting AI adoption in embedded systems will be resolved within the next five years.

- 59% expect meaningful progress in the short-to-mid term (within two years), while another 26% foresee solutions emerging within two to five years.
- A small but notable share of 9% think resolution could take five years or more, and 5% remain skeptical that current problems will ever be fully solved.
- Only 9% anticipate rapid progress in under six months, reflecting the complexity and multi-layered nature of the obstacles involved.

These results suggest a steady but realistic optimism among embedded engineers, who largely view current AI adoption barriers as solvable with continued innovation and ecosystem support.

Challenges, Data & Security Constraints - Key Takeaways

While AI integration in embedded systems is accelerating, developers face persistent technical and organizational challenges that limit broader adoption. The most cited barriers include data collection and quality, hardware cost of AI accelerators, and integration complexity. Engineers also identify real-time performance and limited AI expertise as ongoing pain points, revealing that while enthusiasm for AI is strong, practical implementation still requires tradeoffs between speed, accuracy, and available resources.

Security and data management are growing concerns. Forty-one percent of developers report hardware tampering as a risk, followed closely by insecure wireless protocols and insufficient memory protection. As more devices operate at the edge, exposure to physical and digital attack vectors increases. Only 5% of respondents said they have no current security concerns, indicating that protection and privacy are now part of every embedded design conversation.

Connectivity also presents obstacles. More than half of respondents say latency and connectivity limitations have been a problem in their current projects. Engineers remain divided on the best architecture: 47% currently deploy AI directly on the device, while 40% expect a hybrid edge-cloud balance in the future. This evolution suggests that while cloud platforms offer scalability, reliability and real-time responsiveness keep engineers focused on local processing. Many expect improvement soon: half believe today’s challenges will be resolved within two years, reflecting a pragmatic but optimistic outlook.

Challenges, Data & Security Constraints - Implications for Marketers

For marketers, this section underscores a critical truth about embedded engineers: they are not looking for innovation alone; they are looking for solutions that reduce friction in data handling, security, and deployment. Marketing that directly addresses these pain points, such as guidance on reducing latency, securing firmware updates, or optimizing for real-time inference, will stand out. Engineers value suppliers who anticipate and mitigate their obstacles before they result in failures in the field.

Security is a particularly powerful area for thought leadership. Developers want proof that a vendor's hardware, software, or development environment has been designed with protection in mind. Content that highlights secure boot processes, encryption standards, and vulnerability testing can shift a brand from "supplier" to "trusted partner." Positioning products as inherently secure helps marketers align with the priorities of teams designing for safety-critical and connected applications.

Because data availability is such a central challenge, marketers also have an opportunity to enable experimentation through accessible datasets, synthetic data tools, or edge-ready AI examples. Offering these resources, even in simplified form, proves goodwill and builds practical relevance. Similarly, messaging that bridges the edge and cloud conversation, showing how products maintain consistent performance across both, positions a brand as adaptive rather than prescriptive.

In short, engineers are eager for clarity, not complexity. The companies that provide transparent guidance, actionable resources, and security-backed performance data will be the ones that engineers depend on to overcome the constraints slowing their AI innovation.

Automotive AI Applications

AI + Automotive

In which specific automotive applications are you currently utilizing AI techniques?

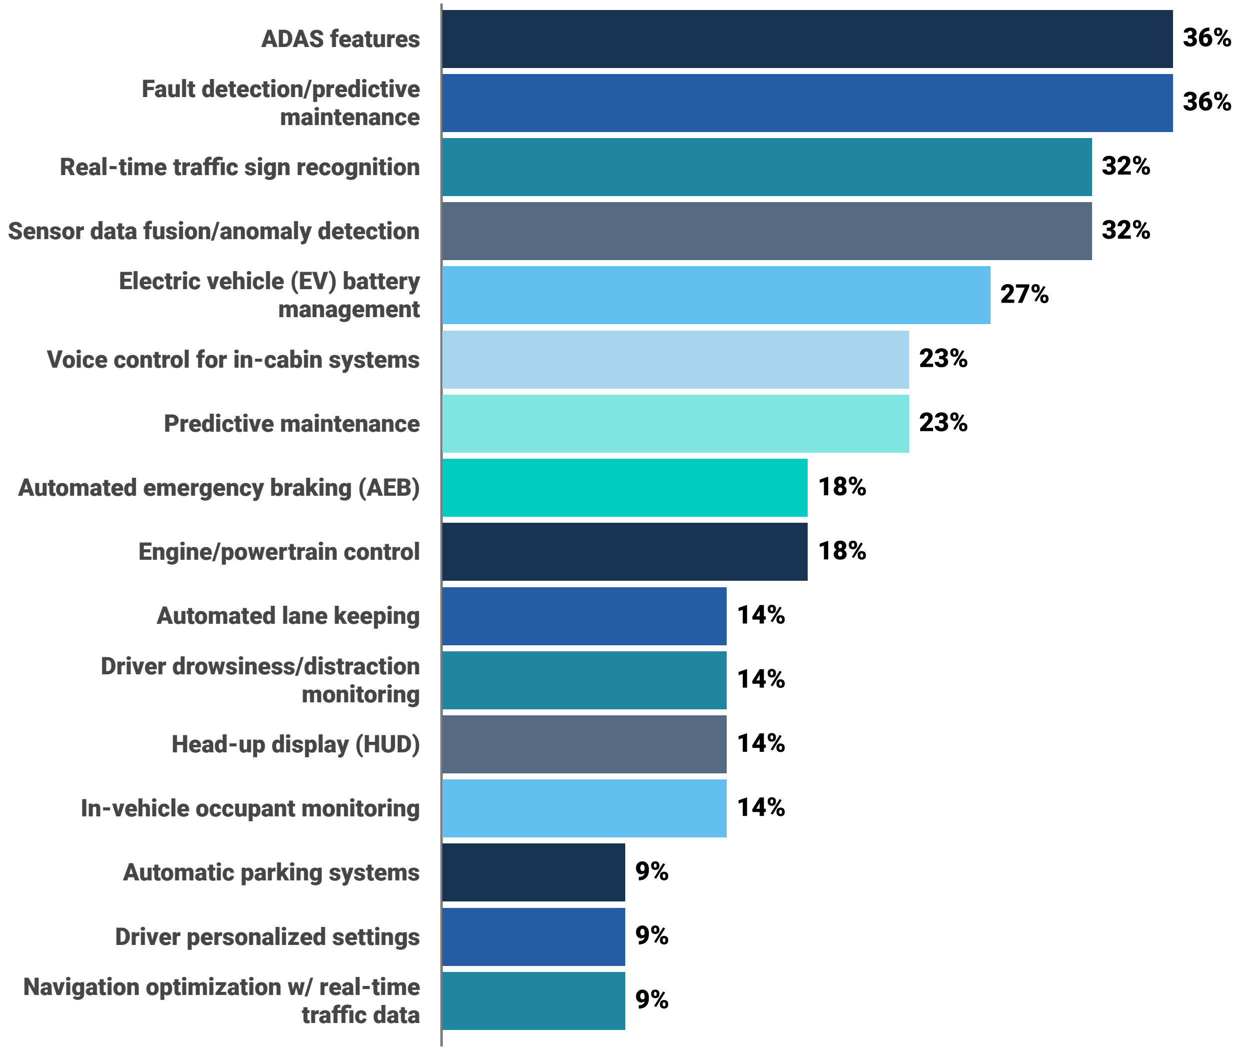

n=22

AI applications in automotive design are concentrated in safety, reliability, and intelligent perception systems.

- The top use cases—ADAS and fault detection or predictive maintenance—show that engineers are primarily focused on improving driver safety and vehicle reliability.
- Traffic sign recognition and sensor data fusion further demonstrate the growing role of AI in enabling situational awareness and decision-making for semi-autonomous systems.
- Secondary implementations such as EV battery management and voice control indicate how AI is also enhancing energy optimization and user experience inside the vehicle.
- Emerging areas like lane keeping assistants, driver drowsiness monitoring, and HUD interfaces illustrate AI's expanding footprint across both core driving functions and in-cabin intelligence.

The data reinforces that automotive AI is shifting from isolated features toward holistic vehicle intelligence that encompasses safety, performance, and user experience.

AI + Automotive

Which of the following sensor modalities are most crucial in your current automotive related embedded design?

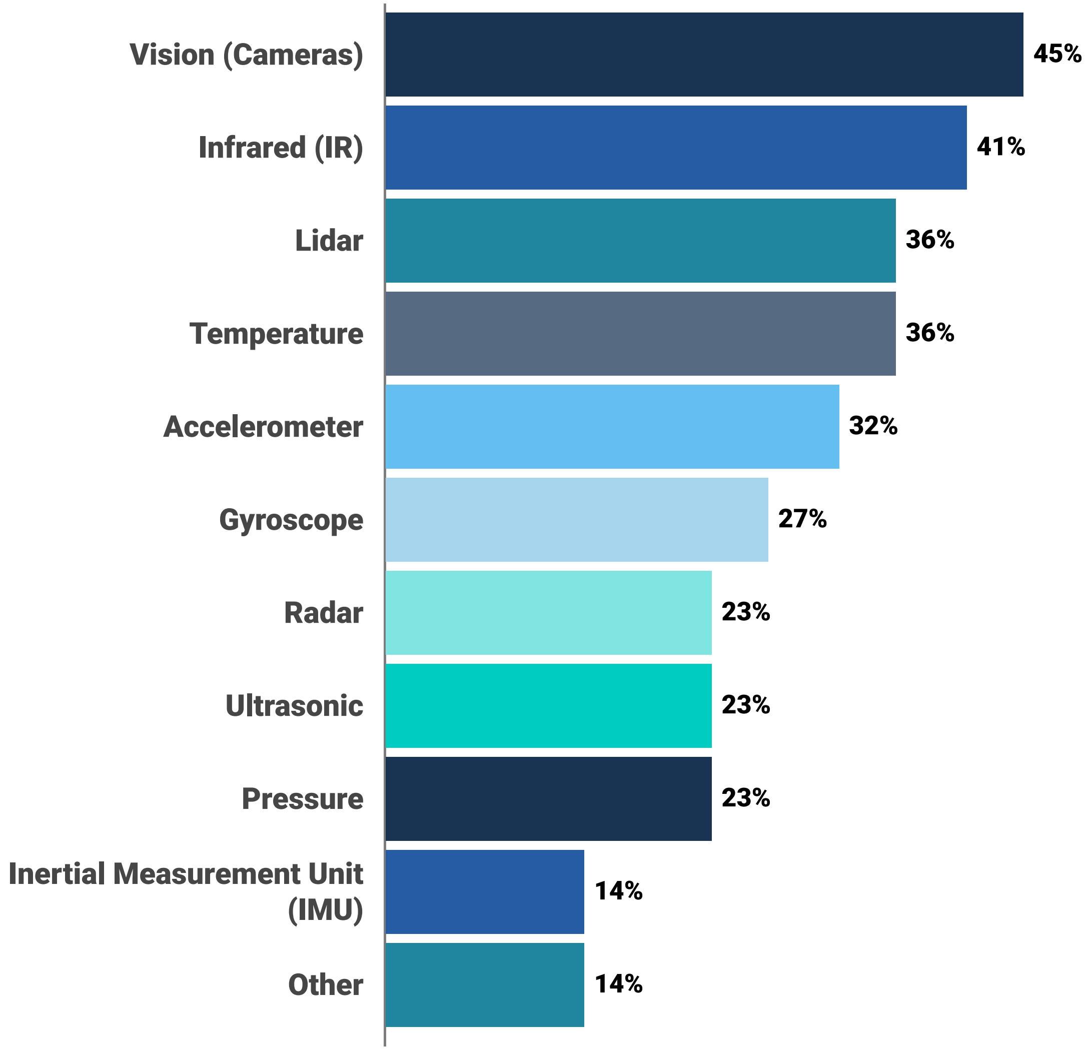

n=22

Automotive AI designs rely heavily on multi-sensor perception, with a clear emphasis on vision-based and environmental awareness technologies.

- Vision leads as the most crucial modality, underscoring the central role of camera data in advanced driver-assistance and situational awareness systems.
- Infrared, Lidar, and temperature sensing show that thermal and depth perception technologies complement visual inputs to improve reliability in varied lighting and weather conditions.
- Motion and stability sensors such as accelerometers and gyroscopes remain key for positioning and control.
- The inclusion of radar, ultrasonic, and pressure sensors highlights the layered approach required for modern automotive safety and automation.

These responses reflect a growing convergence of camera, thermal, and distance-sensing technologies, forming the foundation for perception, navigation, and environmental monitoring in next-generation vehicles.

AI + Automotive

How satisfied are you with the explainability and traceability of AI-based decision making in your automotive embedded systems, especially those directly impacting vehicle control or safety?

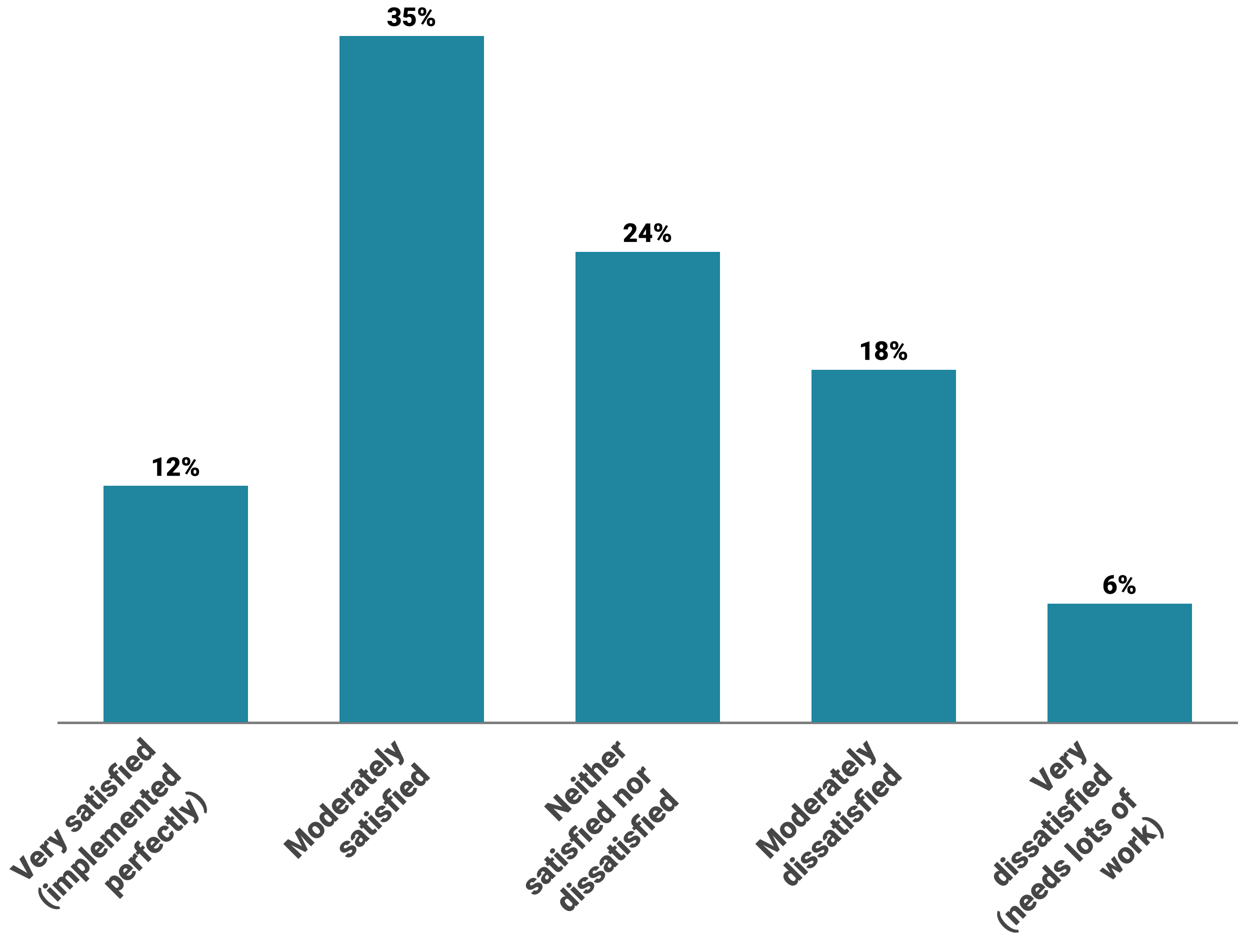

n=17

Satisfaction with the explainability and traceability of AI decisions in automotive systems is mixed, with most engineers acknowledging progress but recognizing room for improvement.

- Nearly half (47%) report being very or moderately satisfied, suggesting that explainability frameworks are becoming more integrated into safety-critical design workflows.
- 24% remain neutral, while 24% express moderate or strong dissatisfaction—reflecting ongoing challenges in meeting regulatory and transparency expectations.
- Only a small group (12%) say their systems achieve full explainability, underscoring that trust and traceability remain evolving priorities in automotive AI.

These responses indicate that while visibility into AI-driven vehicle decisions is improving, consistent, auditable frameworks are still needed to meet industry and safety standards.

AI + Automotive

Which of the following automotive categories do you anticipate being most impacted by AI-powered optimization?

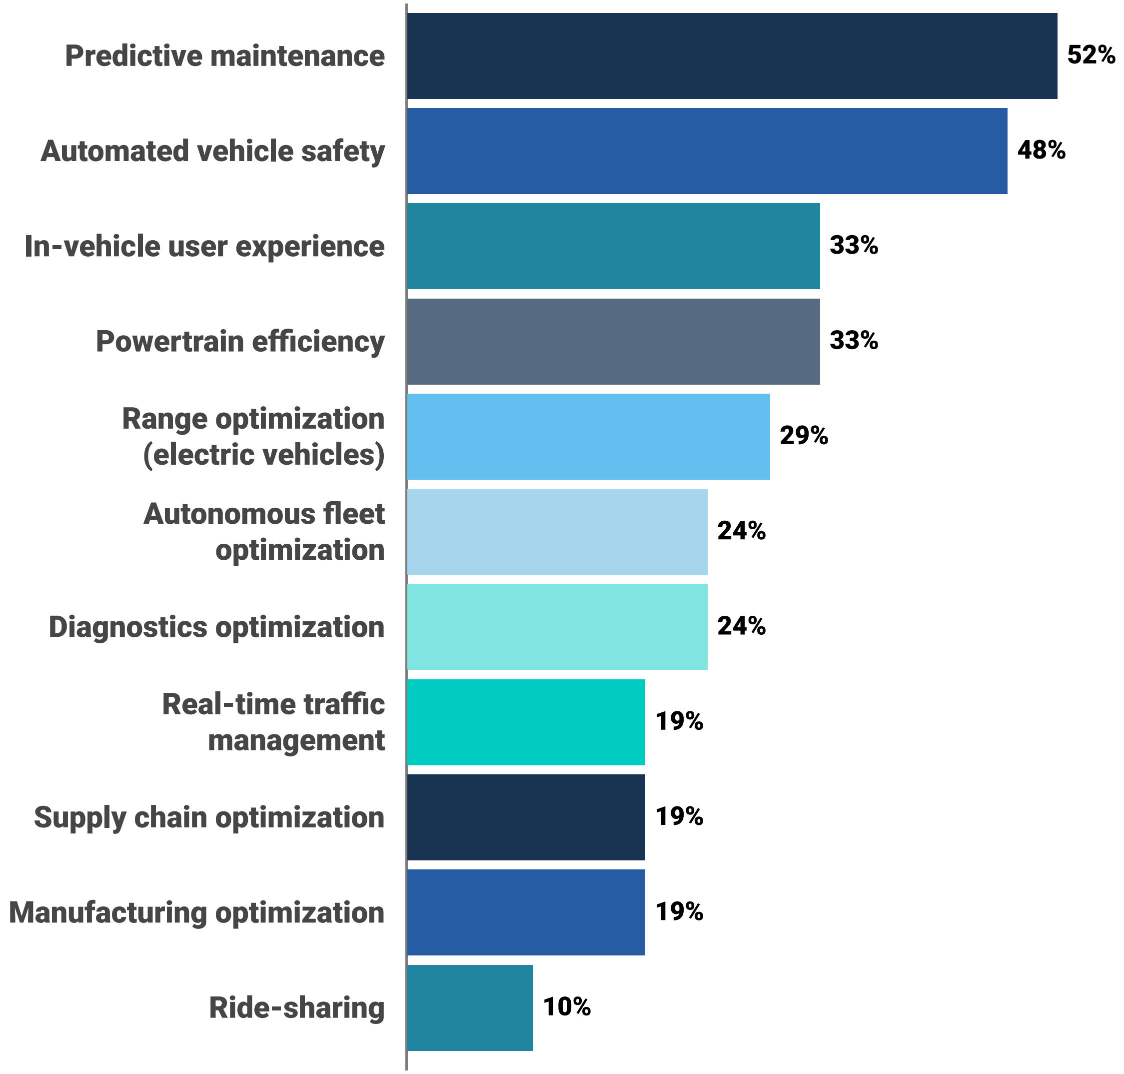

n=21

Engineers expect predictive maintenance and automated vehicle safety to be the areas most transformed by AI optimization in the automotive industry.

- These priorities highlight AI's growing role in preventing system failures and enhancing driver safety, both critical to reliability and brand trust.
- User experience and powertrain efficiency are next in importance, showing that AI's value extends beyond safety into comfort, personalization, and performance.
- Interest in range optimization underscores AI's impact on electric mobility, where intelligent energy management directly influences vehicle competitiveness.
- Broader ecosystem benefits such as autonomous fleet management, diagnostics optimization, and traffic management show emerging interest in system-level efficiency across connected mobility networks.

These results reveal a strong focus on AI-driven reliability, safety, and efficiency—core themes shaping the future of automotive innovation.

AI + Automotive

How likely do you think it will be that collaborative AI between vehicles and infrastructure (V2X) will significantly increase in the next five years?

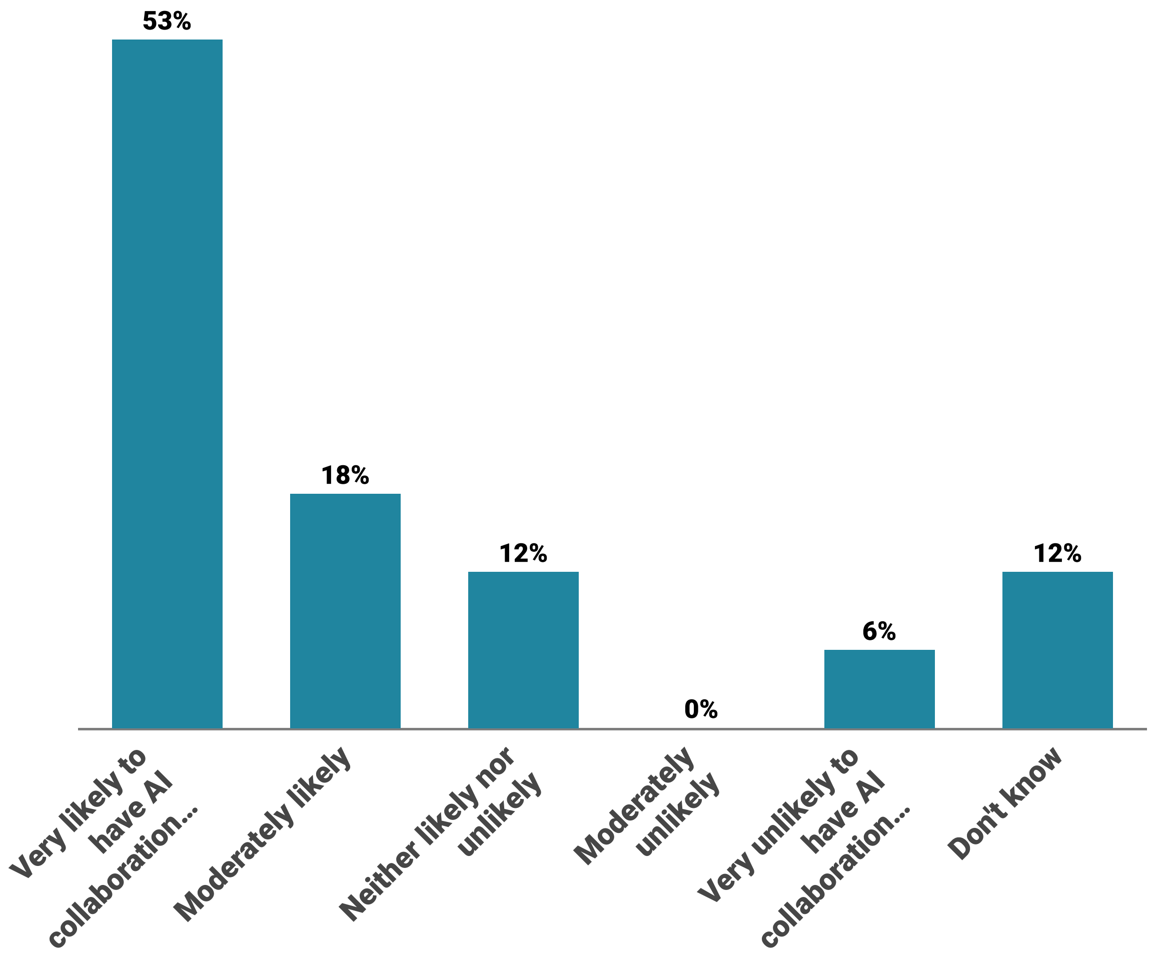

n=17

Embedded developers show strong confidence that collaborative AI and V2X communication will expand rapidly in the coming years.

- More than half (53%) believe it is very likely that AI collaboration between vehicles and infrastructure will increase within five years, while another 18% consider it moderately likely.
- Only a small share (6%) see this progress as unlikely, with the remainder either neutral or unsure.

This consensus underscores growing optimism that AI-enhanced vehicle-to-infrastructure interaction will advance safety, traffic efficiency, and autonomous coordination—positioning V2X integration as a key milestone in the evolution of intelligent transportation systems.

Automotive - Key Takeaways

AI adoption in automotive embedded systems is gaining meaningful momentum, driven by real-world use cases that enhance safety, performance, and efficiency. The most common applications include advanced driver-assistance systems, predictive maintenance, and real-time traffic sign recognition. Engineers are also leveraging AI for sensor data fusion, electric vehicle battery management, and voice control, showing that the technology's role extends well beyond autonomy into daily vehicle reliability and user experience.

Sensor technology underpins much of this progress. Vision systems and infrared sensors are the top modalities used in current designs, followed by lidar, temperature, and accelerometers. This reflects a continued emphasis on perception and environmental awareness as the foundation for AI-driven decision-making in vehicles. However, engineers express mixed satisfaction with AI transparency: only 12 percent are very satisfied with the explainability and traceability of AI-based decisions, while a combined 24% report some level of dissatisfaction.

Looking forward, the potential impact of AI in automotive is expansive. Over half of respondents expect predictive maintenance to see the greatest transformation, followed closely by automated safety systems and in-vehicle user experience. Moreover, 53 percent believe that collaboration between vehicles and infrastructure (V2X) will increase significantly within the next five years. Together, these findings paint a picture of an industry where embedded AI is moving steadily from augmentation to essential infrastructure.

Automotive - Implications for Marketers

For semiconductor and technology marketers, the automotive segment represents one of the most strategic opportunities for embedded AI adoption. Engineers in this space prioritize reliability, explainability, and integration—three areas where strong marketing and communication can have a measurable impact. Campaigns should focus on how products enable trusted performance under safety-critical conditions, such as sensor accuracy, thermal stability, and fault-tolerant processing.

Marketers can differentiate by demonstrating how their solutions address real-world automotive design constraints. This includes showcasing how hardware accelerates perception algorithms, how firmware supports deterministic behavior, and how AI frameworks comply with safety standards such as ISO 26262. Engineers will pay close attention to any vendor that connects marketing claims with verifiable data on reliability and compliance.

The data also reveals that explainability remains a concern. This creates an opportunity for brands to lead with transparent, education-driven content grounded in real data and applications. Providing insight into how AI models interpret sensor data, make decisions, or fail safely in uncertain conditions can build trust between engineers and OEMs.

Finally, with over half of respondents expecting growth in vehicle-to-infrastructure (V2X) collaboration, marketers should begin positioning products within the broader mobility ecosystem, not just at the component level. Highlighting interoperability, scalability, and cross-domain performance can help position a brand as an enabler of connected mobility rather than a single-point supplier.

Chip Type & Vendor Usage

Chip Architectures

Which of the following chip architectures are you using in your current embedded design?

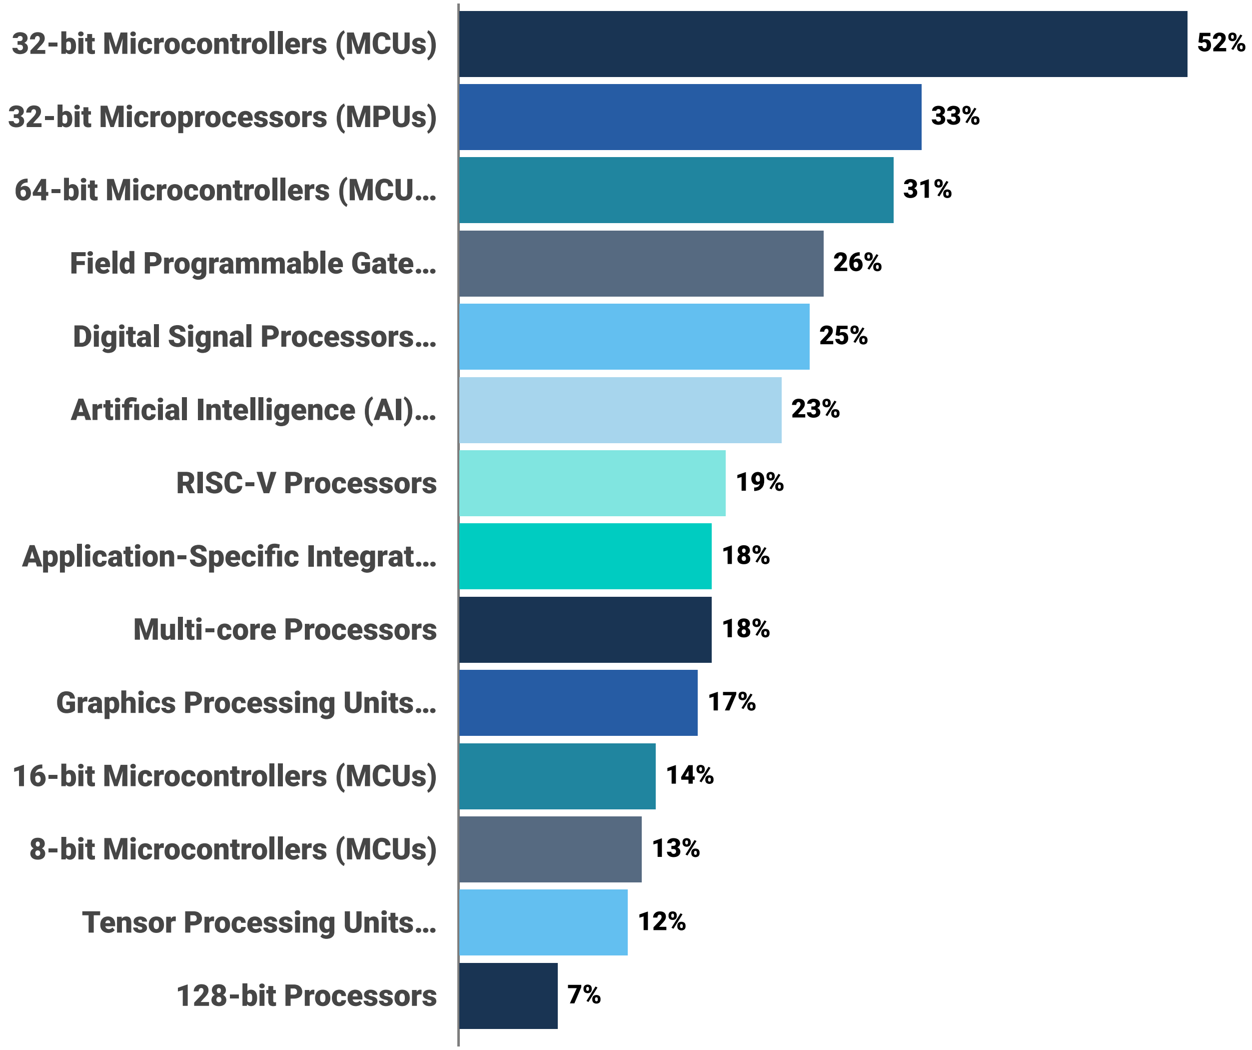

n=141

Embedded AI development is dominated by 32-bit architectures, though diversification toward higher-performance and AI-specific hardware is accelerating.

- 32-bit MCUs remain the standard for embedded designs, with 32-bit MPUs and 64-bit processors reflecting broader scalability for AI workloads.
- FPGAs and DSPs continue to play key roles in custom signal processing and rapid prototyping, while AI accelerators signal the growing integration of specialized compute engines directly on embedded boards.
- Emerging architectures such as RISC-V, ASICs , and TPUs show that engineers are exploring new paths to efficiency and performance.

The data illustrates an industry in transition—from general-purpose microcontrollers toward heterogeneous, AI-optimized processing environments capable of handling complex edge workloads.

Microprocessors

How many microprocessors are in your current embedded project?

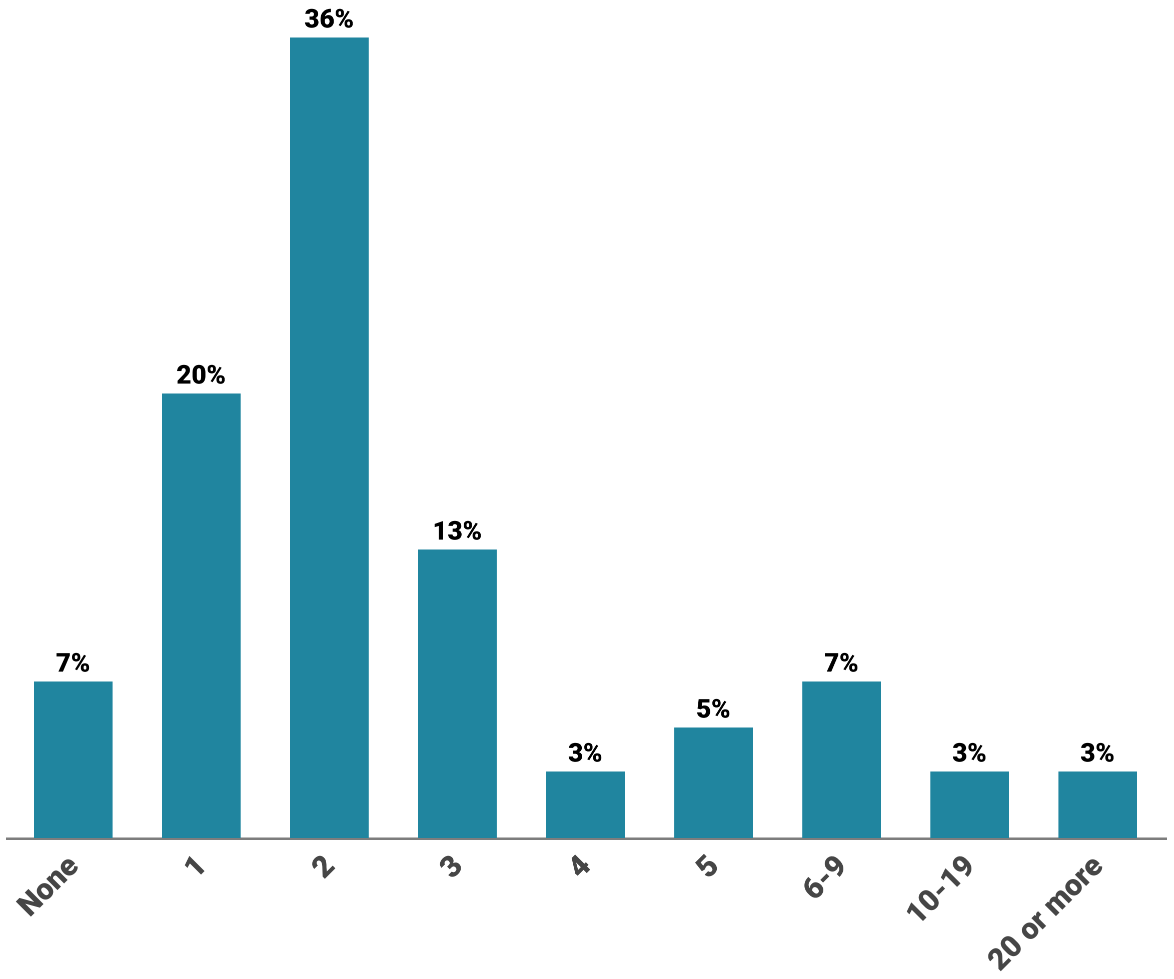

n=61

Most embedded projects feature a small number of microprocessors, reflecting the compact and specialized nature of typical designs.

- The majority of respondents use between one and three processors, with two-processor designs being the most common configuration.
- A small share report systems with no microprocessor, likely referring to MCU-only or ASIC-based implementations.
- Only 18% of engineers work with highly parallel or complex architectures (five or more processors), which are more typical in advanced compute or multi-sensor systems.

The data indicates that dual-processor architectures have become the norm, balancing performance needs with cost and power efficiency in embedded AI and control applications.

Chip Families

Which of the following chip series or chip families are you currently using in your embedded design?

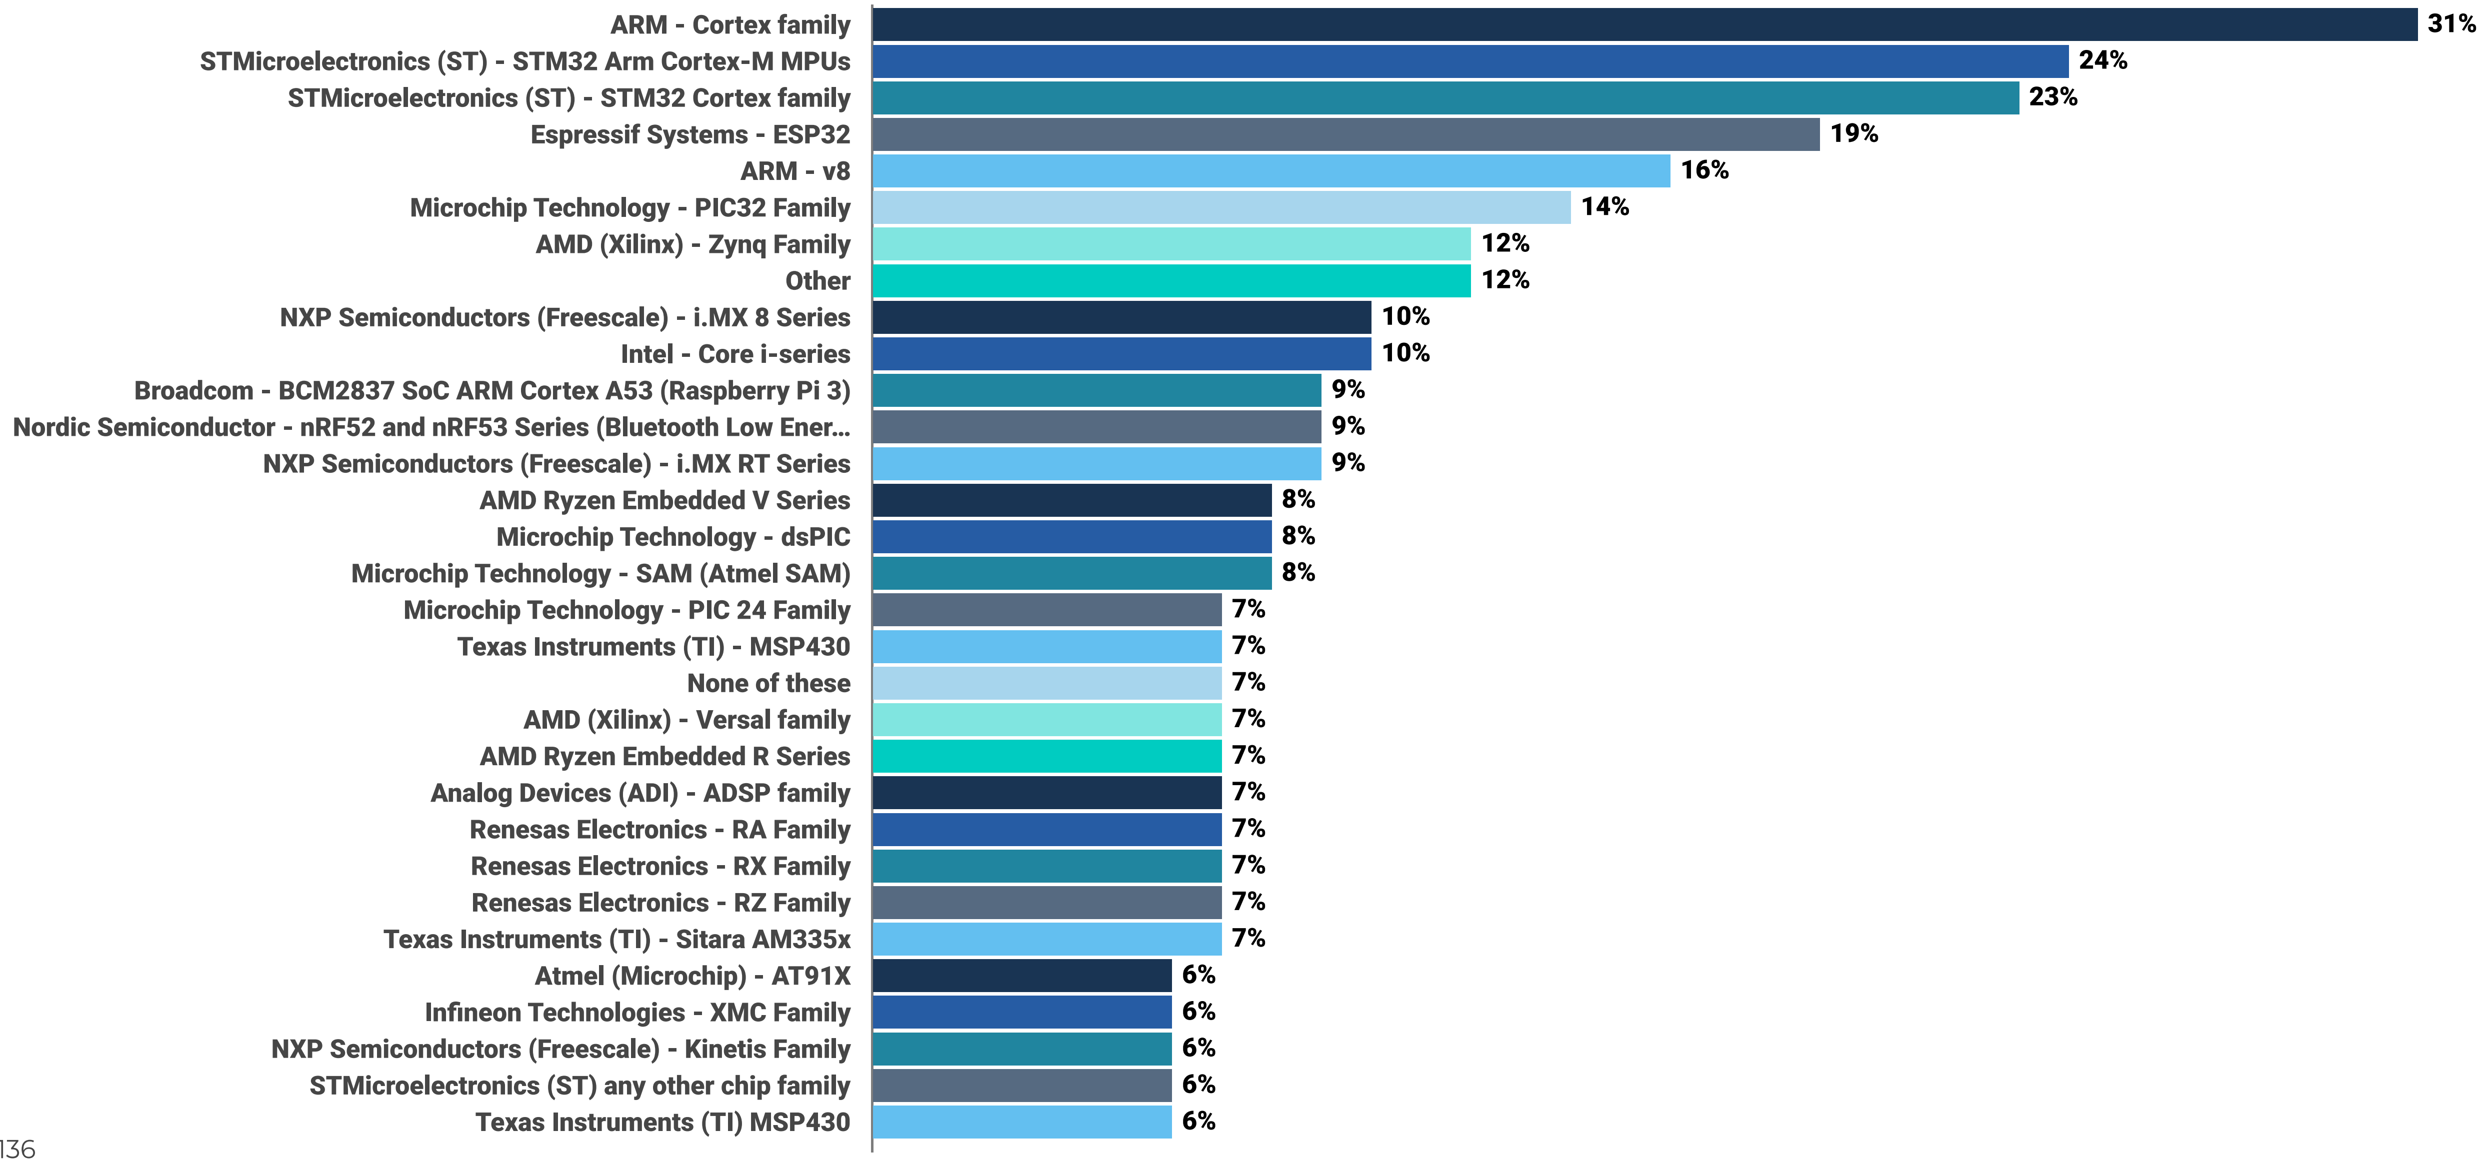

n=136

Answer options at/below 5% were removed - for the full list, please see the Appendix

Chip Families

Which of the following chip series or chip families are you currently using in your embedded design?

Embedded development remains centered on ARM-based architectures, with widespread adoption across both general-purpose and vendor-specific platforms.

- ARM Cortex leads as the most used architecture, followed closely by STMicroelectronics STM32 families, underscoring the deep integration of ARM cores in mainstream embedded projects.
- Espressif's ESP32 continues to gain relevance for IoT and connected applications, valued for its low cost and strong wireless capabilities.
- Other major ecosystems—Microchip, NXP, and AMD/Xilinx Zynq—demonstrate that engineers leverage a wide mix of platforms to balance performance, cost, and availability.
- The presence of Intel, Renesas, and TI families reflects the continuing relevance of established vendors across specific niches.

These results confirm that ARM-based systems remain the dominant foundation for embedded design, supported by a robust ecosystem of vendors and development tools spanning industrial, IoT, and AI-driven applications.

Chip Vendors

Which of the following chip vendors do you currently use for you embedded designs?

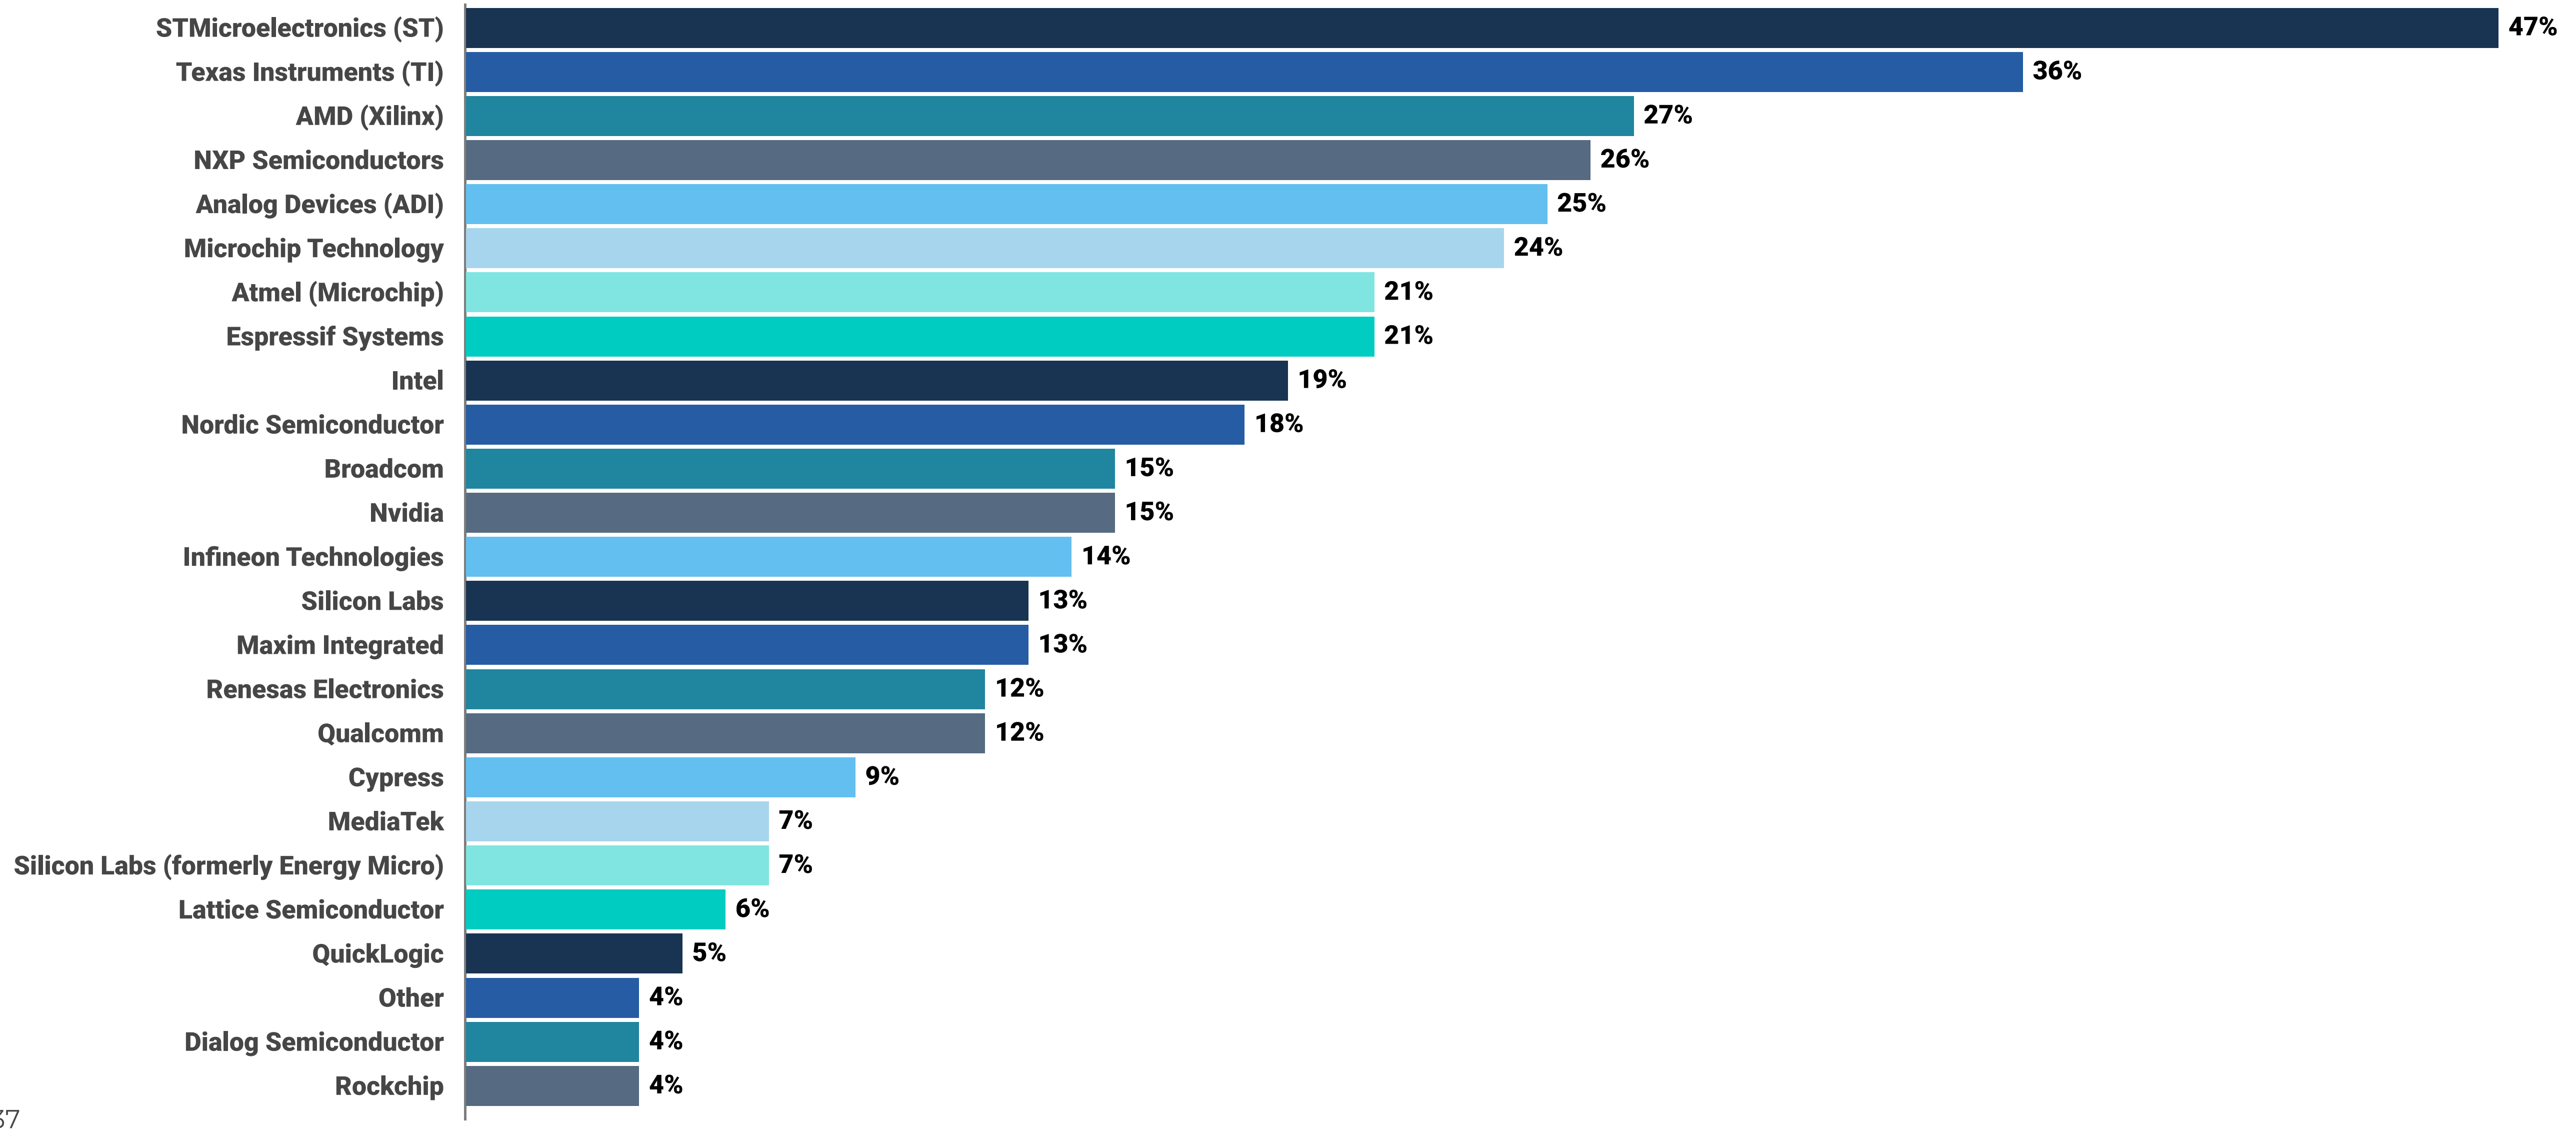

n=137

Answer options at/below 3% were removed - for the full list, please see the Appendix

Chip Vendors

Which of the following chip vendors do you currently use for you embedded designs?

The embedded design landscape remains highly diversified, though a few leading vendors dominate usage across projects.

- STMicroelectronics leads by a clear margin, confirming its position as the most widely used supplier for embedded and AI-enabled designs.
- Texas Instruments, AMD/Xilinx, and NXP form the next tier of major vendors, each serving broad application categories ranging from industrial control to high-performance edge processing.
- Analog Devices and Microchip also hold strong positions, particularly in sensing, mixed-signal, and MCU-based solutions.
- A long list of specialized vendors—including Espressif, Nordic, Infineon, and Renesas—shows the continued fragmentation of the market as developers tailor components to specific performance, connectivity, and power requirements.

The data illustrates a competitive supplier ecosystem where STMicroelectronics and TI set the pace, but numerous other players remain essential to specialized segments of embedded and AI hardware design.

The Best Overall

Which ONE of the following chip vendors do you consider best OVERALL as a vendor for embedded designs?

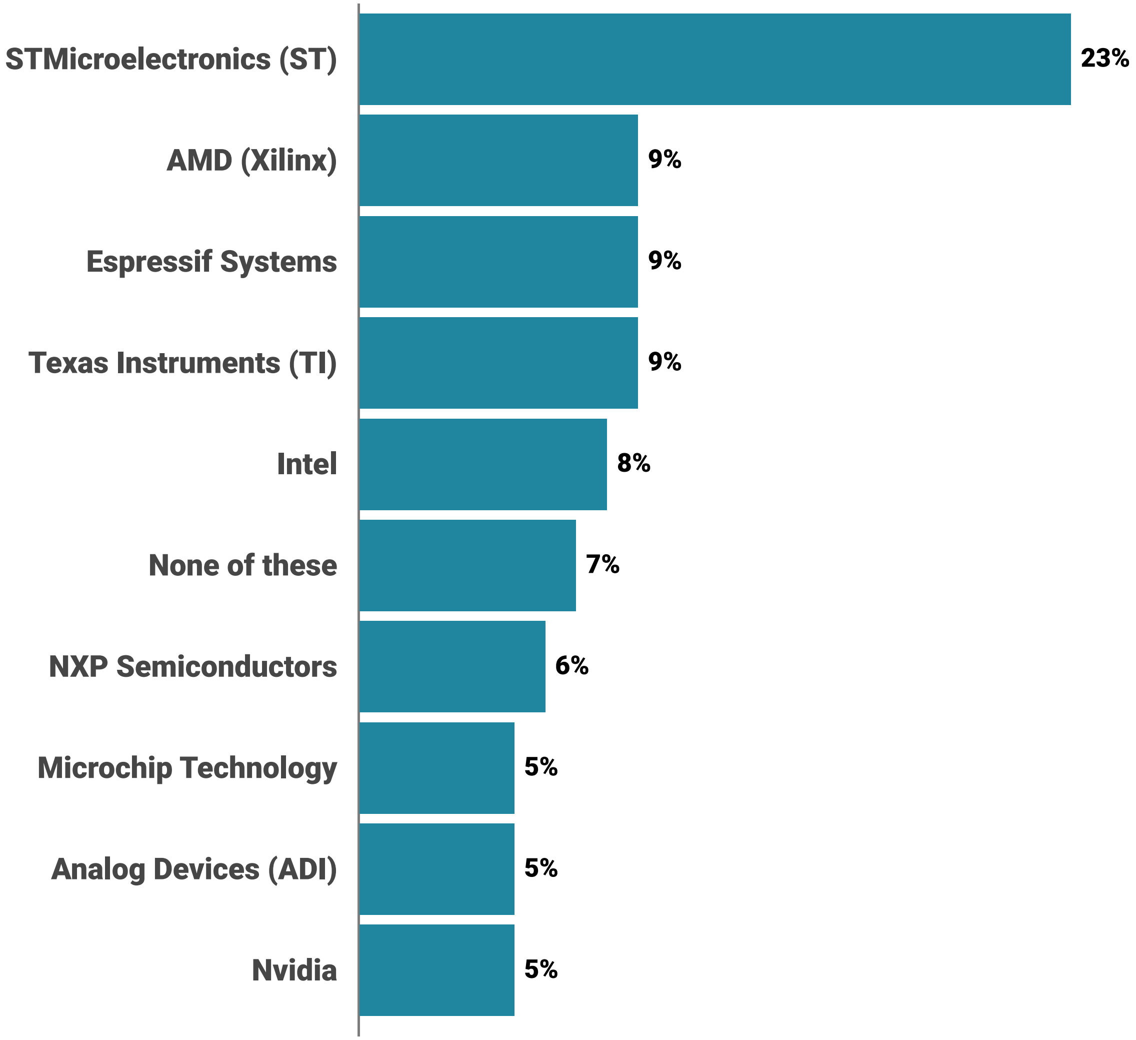

n=133

Answer options at/below 4% were removed - for the full list, please see the Appendix

STMicroelectronics stands out as the clear leader in overall brand perception for embedded design, ranking well ahead of all other chip vendors.

- A group of three competitors—AMD/Xilinx, Espressif Systems, and Texas Instruments—share the second tier, reflecting strong reputations for innovation and reliability across different application domains.
- Intel and NXP follow, while Microchip, ADI, and Nvidia maintain presence among engineers focused on specific niches like sensing, signal processing, and AI acceleration.
- The 7% who selected "none of these" suggest that some engineers view vendor performance as application-dependent rather than brand-driven.

The results confirm STMicroelectronics' dominant reputation in the embedded sector, with competitive but more fragmented brand strength among other leading chipmakers.

Distributors

Which of the following embedded systems component distributors have you used in the last 12 months?

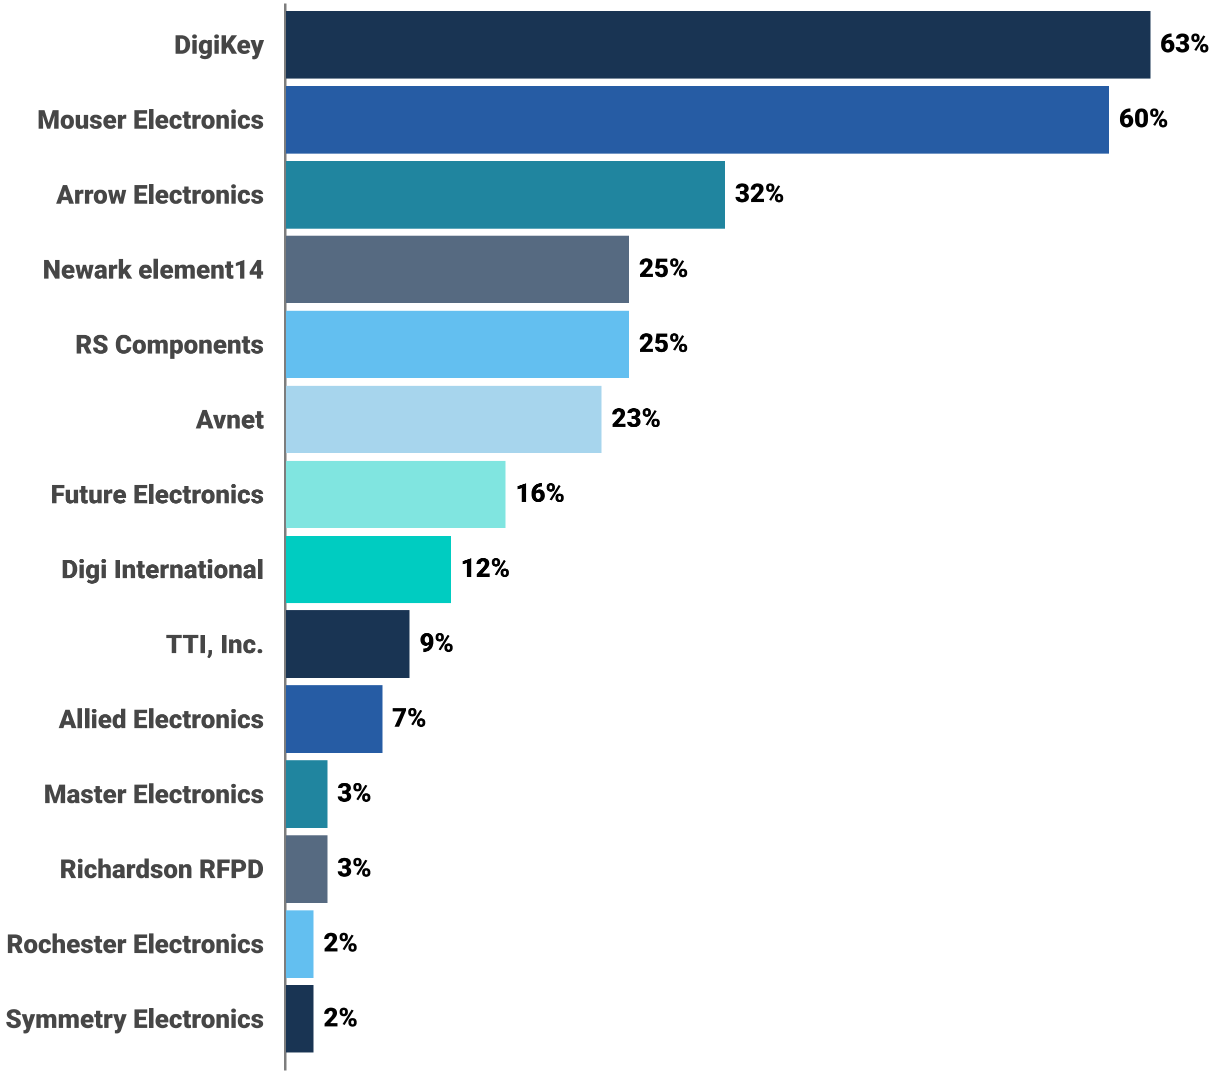

n=131

DigiKey and Mouser dominate as the most-used distributors for embedded systems components, serving as the go-to sources for many engineers across industries.

- A second tier of major distributors includes Arrow, Newark, RS Components, and Avnet, indicating strong but less universal usage compared to DigiKey and Mouser.
- Smaller but notable shares for Future Electronics, Digi International, and TTI reflect the diversity of sourcing channels engineers leverage to meet design and supply needs.

These results highlight the central role of broadline distributors—particularly DigiKey and Mouser—in supporting embedded design workflows through availability, selection, and speed of delivery.

AI Chip Vendors

Which of the following Artificial Intelligence (AI) chip accelerator vendors (if any) are you currently using in you embedded design?

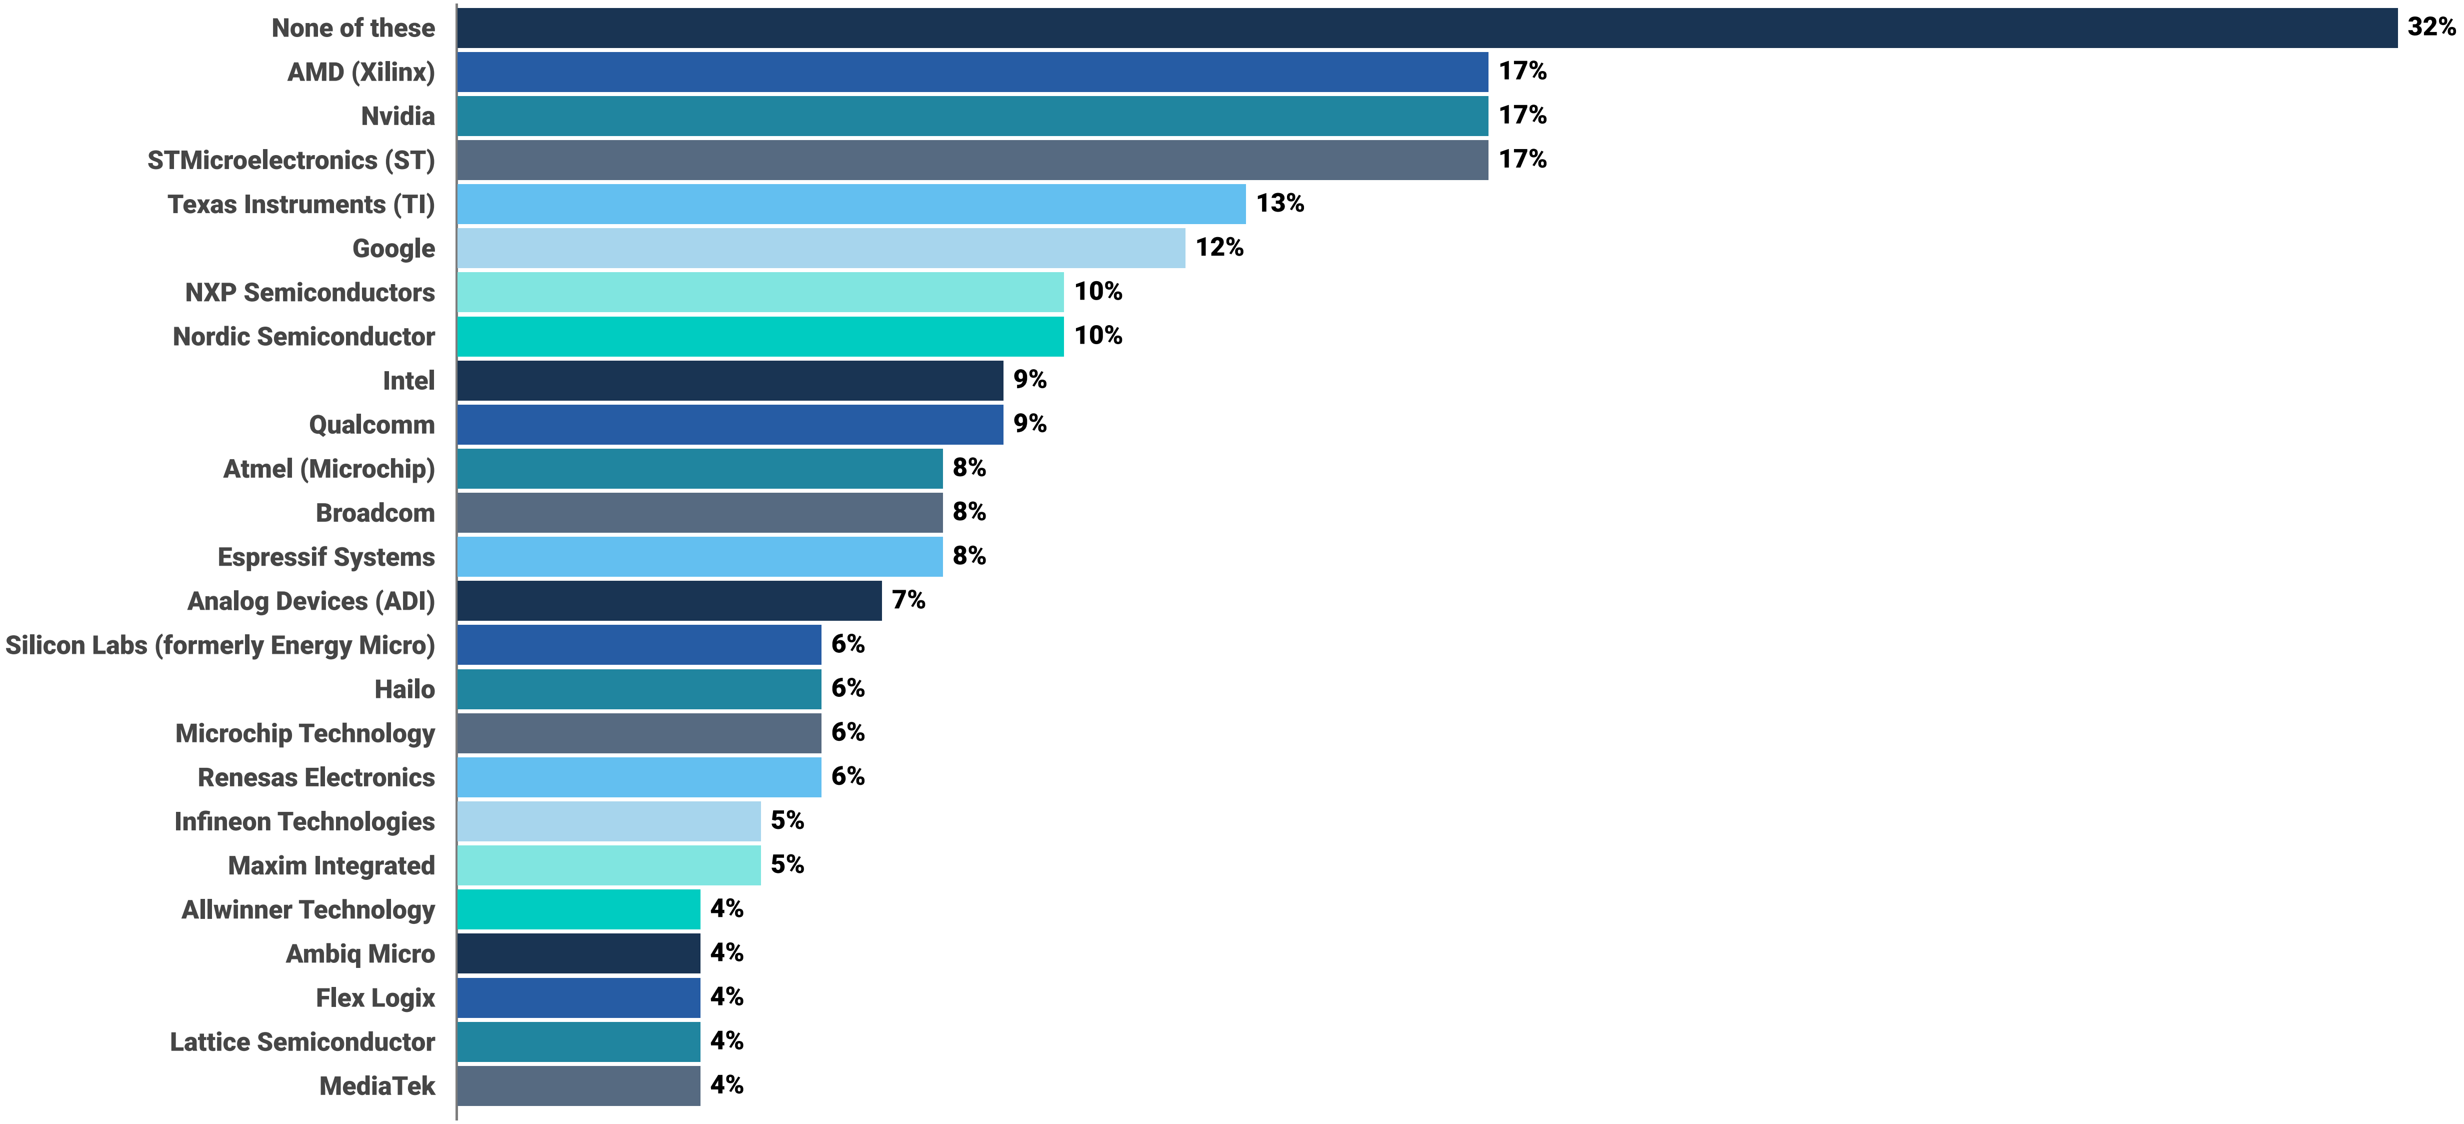

n=126

Answer options at/below 3% were removed - for the full list, please see the Appendix

AI Chip Vendors

Which of the following Artificial Intelligence (AI) chip accelerator vendors (if any) are you currently using in you embedded design?

Adoption of AI chip accelerators in embedded design remains fragmented, with a mix of established semiconductor leaders and emerging edge-AI specialists in use.

- About one-third of respondents report using none of the listed vendors, indicating that many embedded projects still rely on standard CPUs or MCUs without dedicated AI acceleration.
- Among active users, AMD/Xilinx, Nvidia, and STMicroelectronics lead adoption, showing strong engagement with vendors that offer flexible FPGA- or MCU-based AI solutions.
- Texas Instruments, Google, and NXP form the next tier, hinting at growing reliance on integrated AI toolchains within established hardware ecosystems.
- The appearance of niche players such as Hailo, Ambiq Micro, and Flex Logix signals early adoption of dedicated low-power AI accelerators tailored for edge devices.

These results suggest that AI acceleration in embedded systems is still in early growth, with engineers actively exploring both mainstream and specialized platforms as hardware innovation continues.

Software

Which of the following software vendors do you currently use in your embedded designs?

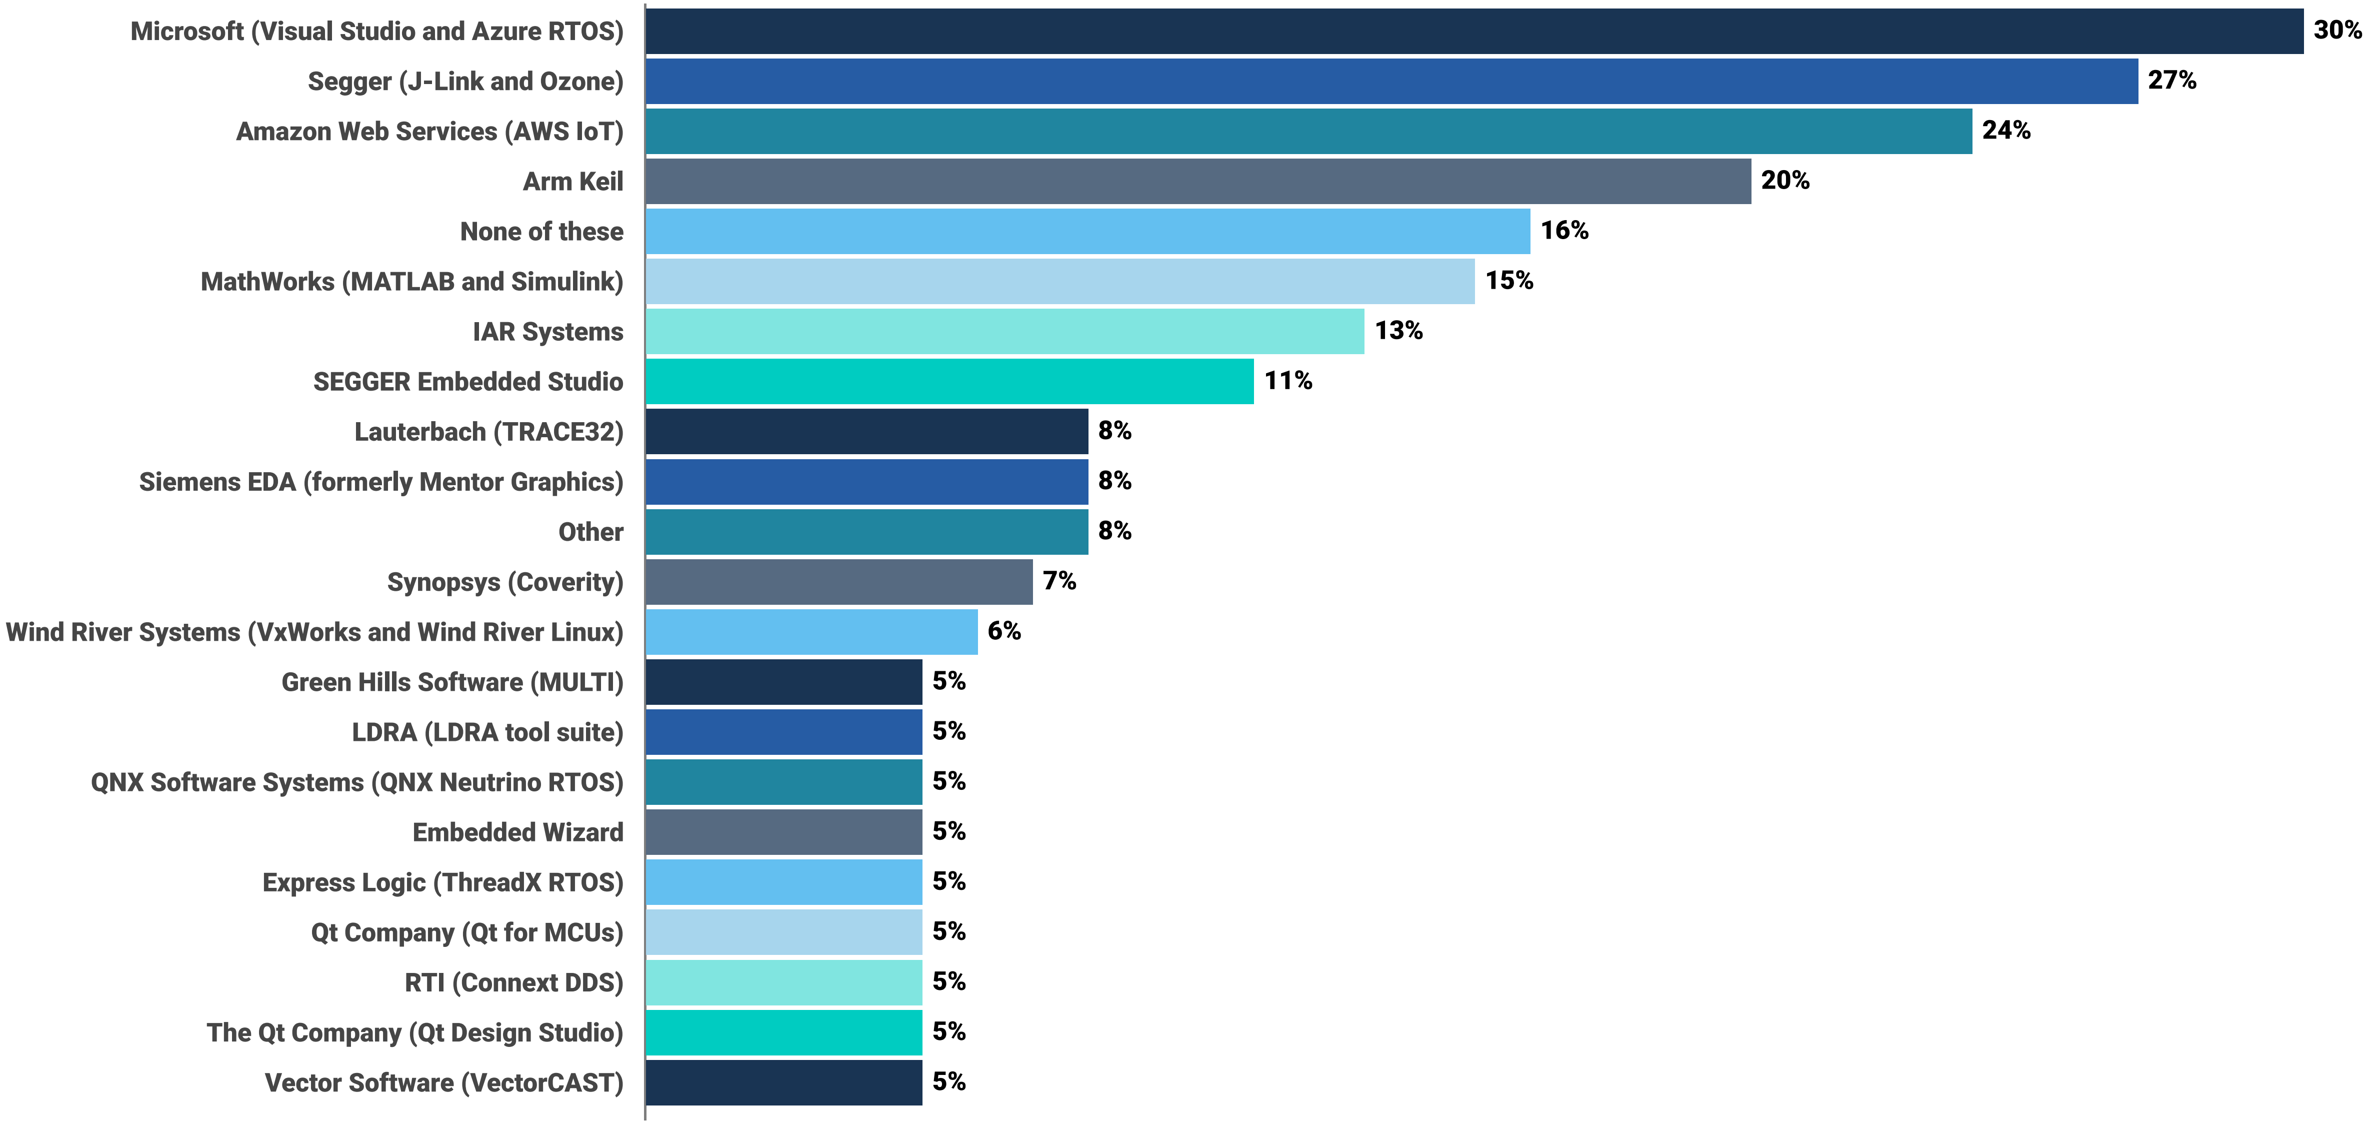

n=132

Answer options at/below 4% were removed - for the full list, please see the Appendix

Software

Which of the following software vendors do you currently use in your embedded designs?

Embedded software development is supported by a broad ecosystem of commercial and open-source vendors, with a handful of leaders dominating across toolchains and runtime environments.

- Microsoft tops the list, driven by strong adoption of Visual Studio for development and Azure RTOS for connected applications.
- Segger and AWS follow, underscoring the growing role of debugging tools and cloud integration in embedded workflows.
- Arm Keil and MathWorks remain essential tools for developers focused on real-time systems and model-based design.
- A diverse mix of smaller vendors—IAR Systems, Synopsys, Wind River, and Green Hills—reflects a diverse ecosystem catering to safety-critical, industrial, and AI-driven use cases.

The findings show that while Microsoft and Segger lead adoption, embedded engineers continue to mix tools across ecosystems to meet project-specific needs for development, debugging, simulation, and cloud connectivity.

Chip Type & Vendor Usage - Key Takeaways

Embedded engineers rely on a wide range of architectures and vendors, resulting in an ecosystem defined by diversity and competition rather than dominance. Fifty-two percent of respondents use 32-bit microcontrollers, making them the most common architecture for current embedded designs. Other widely used architectures include 32-bit microprocessors, 64-bit devices, and FPGAs, reflecting the growing need for flexible compute power that enables AI and signal processing tasks. Smaller but notable adoption of DSPs and AI accelerators indicates that performance specialization is increasingly part of mainstream design.

When it comes to vendors, the embedded community demonstrates a healthy mix of loyalty and pragmatism. STMicroelectronics leads as the most frequently used chip vendor, followed by Texas Instruments, AMD/Xilinx, and NXP Semiconductors. Secondary players such as Analog Devices, Microchip, and Espressif Systems also maintain strong representation across various application areas. When asked which vendor they consider best overall, STMicroelectronics topped the list, reinforcing its strong positioning in the embedded design space.

Distributor preferences show similar concentration among trusted names. DigiKey and Mouser dominate as preferred sources for component procurement, followed by Arrow, Newark, and RS Components. These findings highlight the importance of availability, support, and speed in supply chain relationships. On the software side, engineers most often use Microsoft, Segger, and Amazon Web Services, with additional usage of tools from Arm Keil, MathWorks, and IAR Systems, indicating that design ecosystems encompass hardware, software, and cloud infrastructure.

Chip Type & Vendor Usage - Implications for Marketers

For semiconductor and technology marketers, the data suggests that engineers rarely commit to a single brand—they choose reliability, support, and ease of integration. The most successful vendors are those that create consistent, well-documented experiences across architectures and development environments. Marketing should reinforce how a brand's ecosystem simplifies decision-making, accelerates prototyping, and reduces cost and risk for design teams under pressure to deliver.

Vendor preference data shows that engineers tend to trust suppliers that combine strong product performance with high availability and dependable distribution partners. Marketers should emphasize how their components are supported by global inventory networks, responsive technical teams, and accessible development tools. Partnerships with major distributors like DigiKey or Mouser have to be highlighted as part of a seamless customer experience rather than a separate procurement step.

For emerging or challenger brands, differentiation lies in transparency and enablement. Providing detailed performance documentation, long-term supply guarantees, and accessible online support resources can close the trust gap with established competitors. Engineers are pragmatic buyers: they adopt the solution that works best within their existing workflow and has credible evidence of long-term reliability.

Finally, this section illustrates that the traditional boundary between hardware and software has largely disappeared. Engineers now expect vendors to provide full ecosystems—hardware, SDKs, simulation environments, and reference designs—that work together without friction. Marketing strategies should communicate how each product fits within a complete development path rather than as a stand-alone component. The brands that make the engineer's life simpler at every step of the design process will be the ones that earn repeat design-ins and long-term customer loyalty.

Learning Methods & Media Consumption

Skill Development

Which of the following general methods do you use to acquire and develop skill sets to effectively work with AI in an embedded design context?

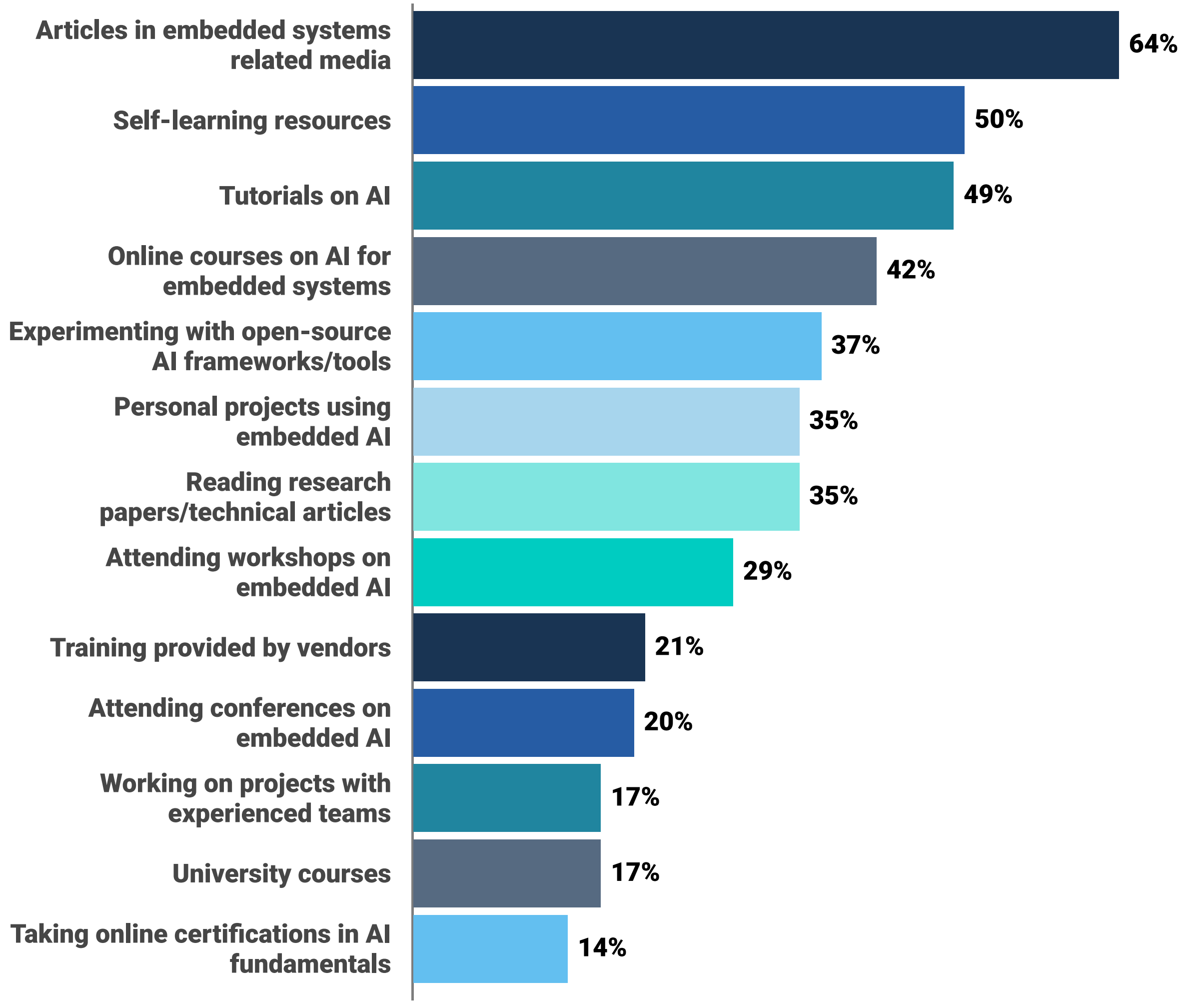

n=133

Engineers are building AI expertise in embedded design primarily through independent, media-based, and self-directed learning.

- The top channels (industry media, self-learning resources, and online tutorials) show that engineers prefer on-demand, informal education tailored to their immediate design challenges.
- Structured training options like online courses and workshops also play a significant role, indicating a continuing demand for applied, hands-on instruction.
- Nearly four in ten developers strengthen AI skills through open-source experimentation and personal projects, highlighting the embedded community's strong maker and testing culture.
- Formal education pathways rank lower, reflecting the fast-moving and self-taught nature of embedded AI skill development.

The data makes clear that media-driven insights and experiential learning are the dominant ways embedded engineers stay current in the evolving field of AI.

Media

Which of the following specific media do you currently use to learn about embedded systems and the semiconductor industry?

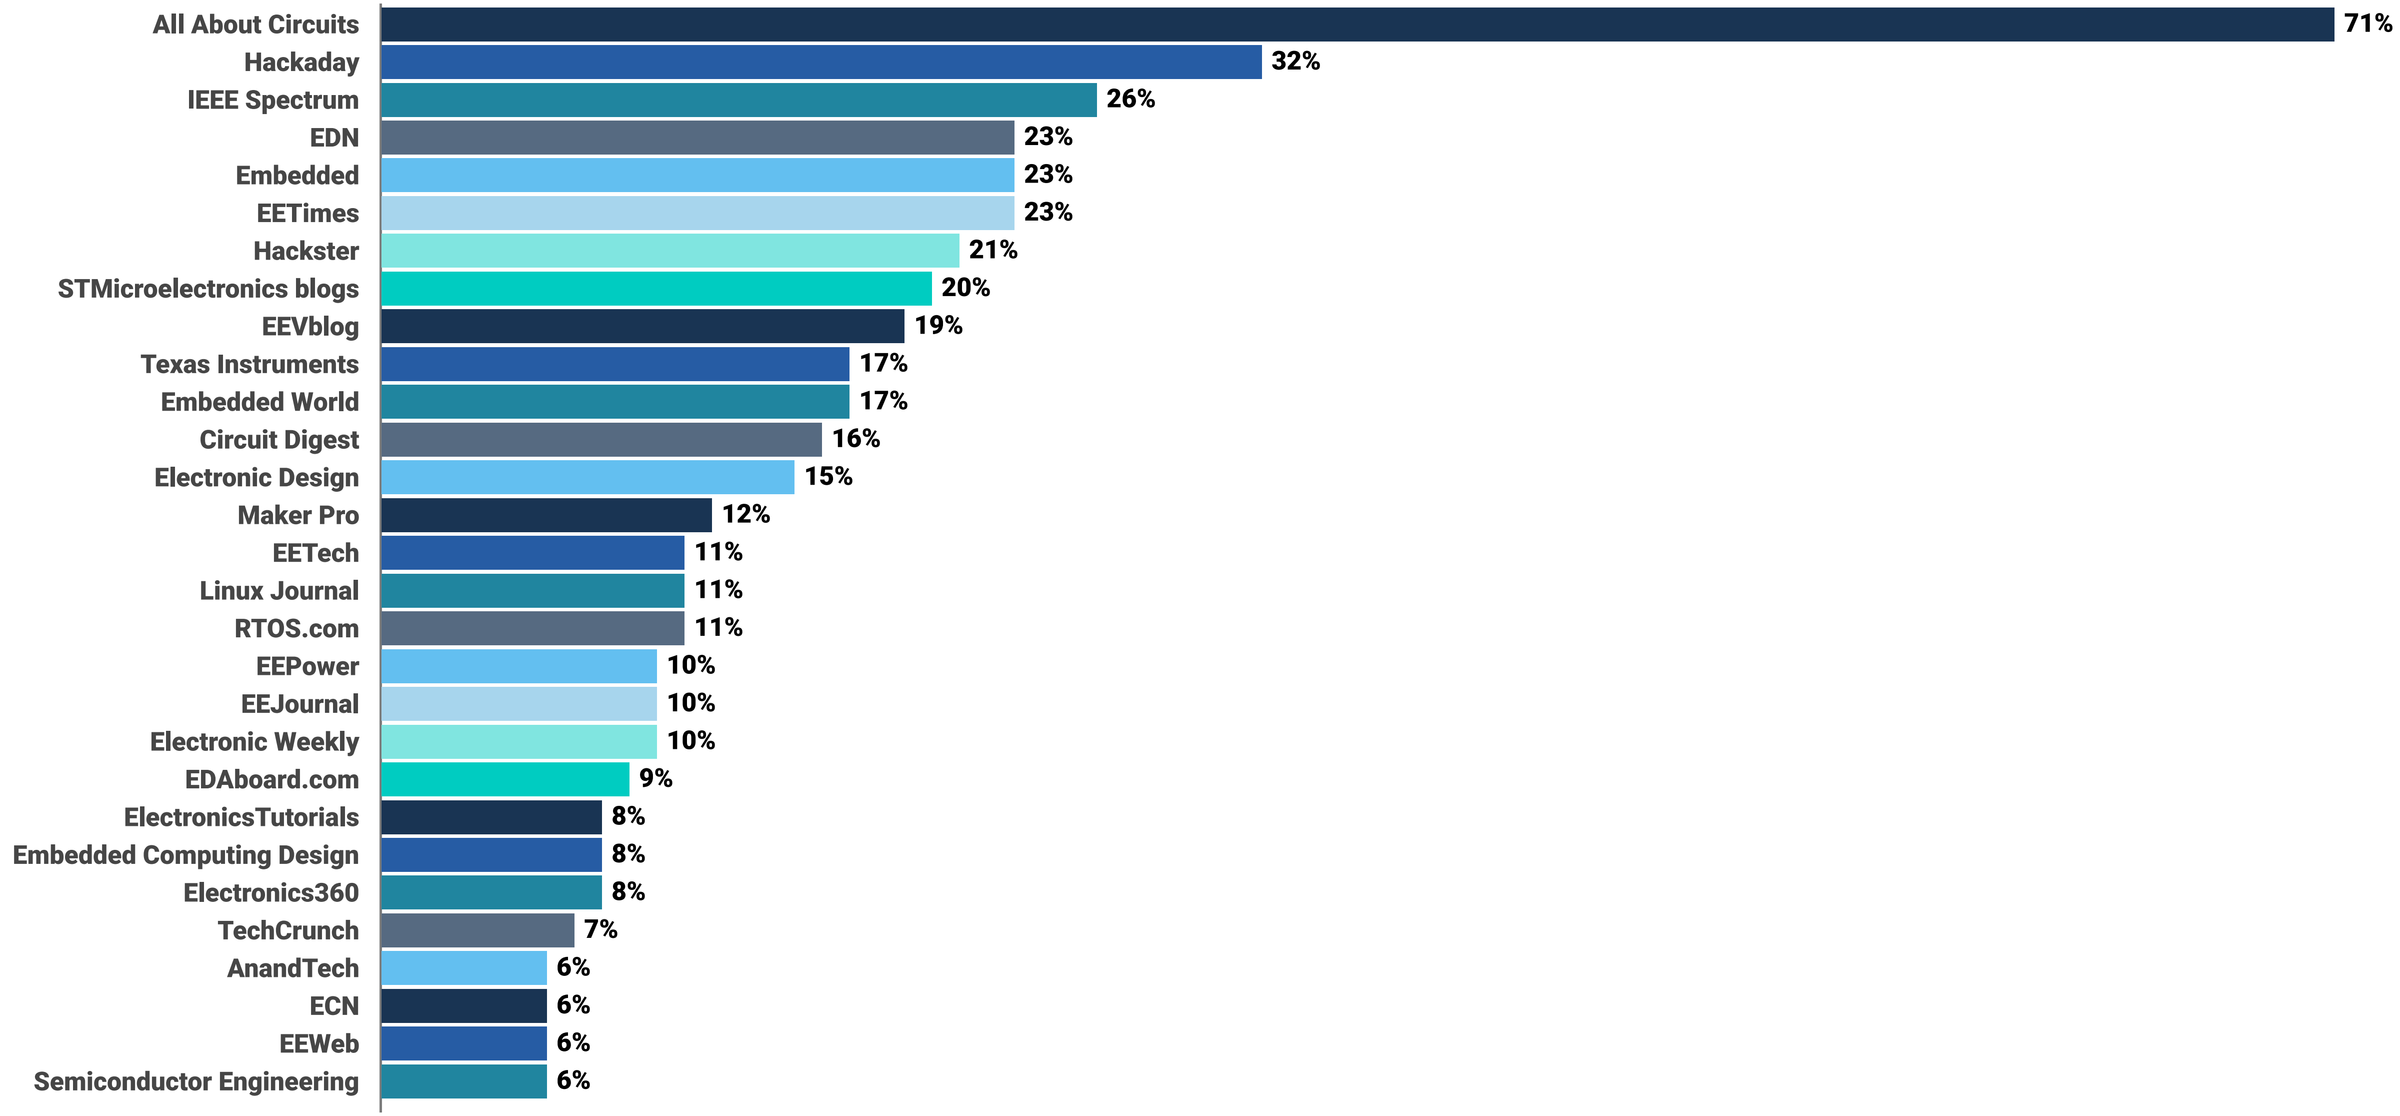

n=133

Answer options at/below 5% were removed - for the full list, please see the Appendix

Media

Which of the following specific media do you currently use to learn about embedded systems and the semiconductor industry?

All About Circuits is the undisputed leader among media sources used by embedded engineers—outpacing all other publications by a wide margin.

- A secondary group of go-to outlets includes Hackaday, IEEE Spectrum, and EDN, Embedded, and EETimes, which collectively represent the next most influential set of information sources.
- Hackster and STMicroelectronics blogs also show strong representation, highlighting the blend of community-driven learning and vendor-supported content engineers reference for guidance.
- EETech's brands together demonstrate a powerful reach across the embedded audience, confirming their leadership position as a trusted ecosystem for technical education and design insights.

These results emphasize that All About Circuits anchors the media landscape for embedded professionals, supported by a mix of independent, community, and manufacturer platforms that engineers use to stay updated on technology and design trends.

Media

Which of the following media is your "go to" source for information on AI related embedded systems topics?

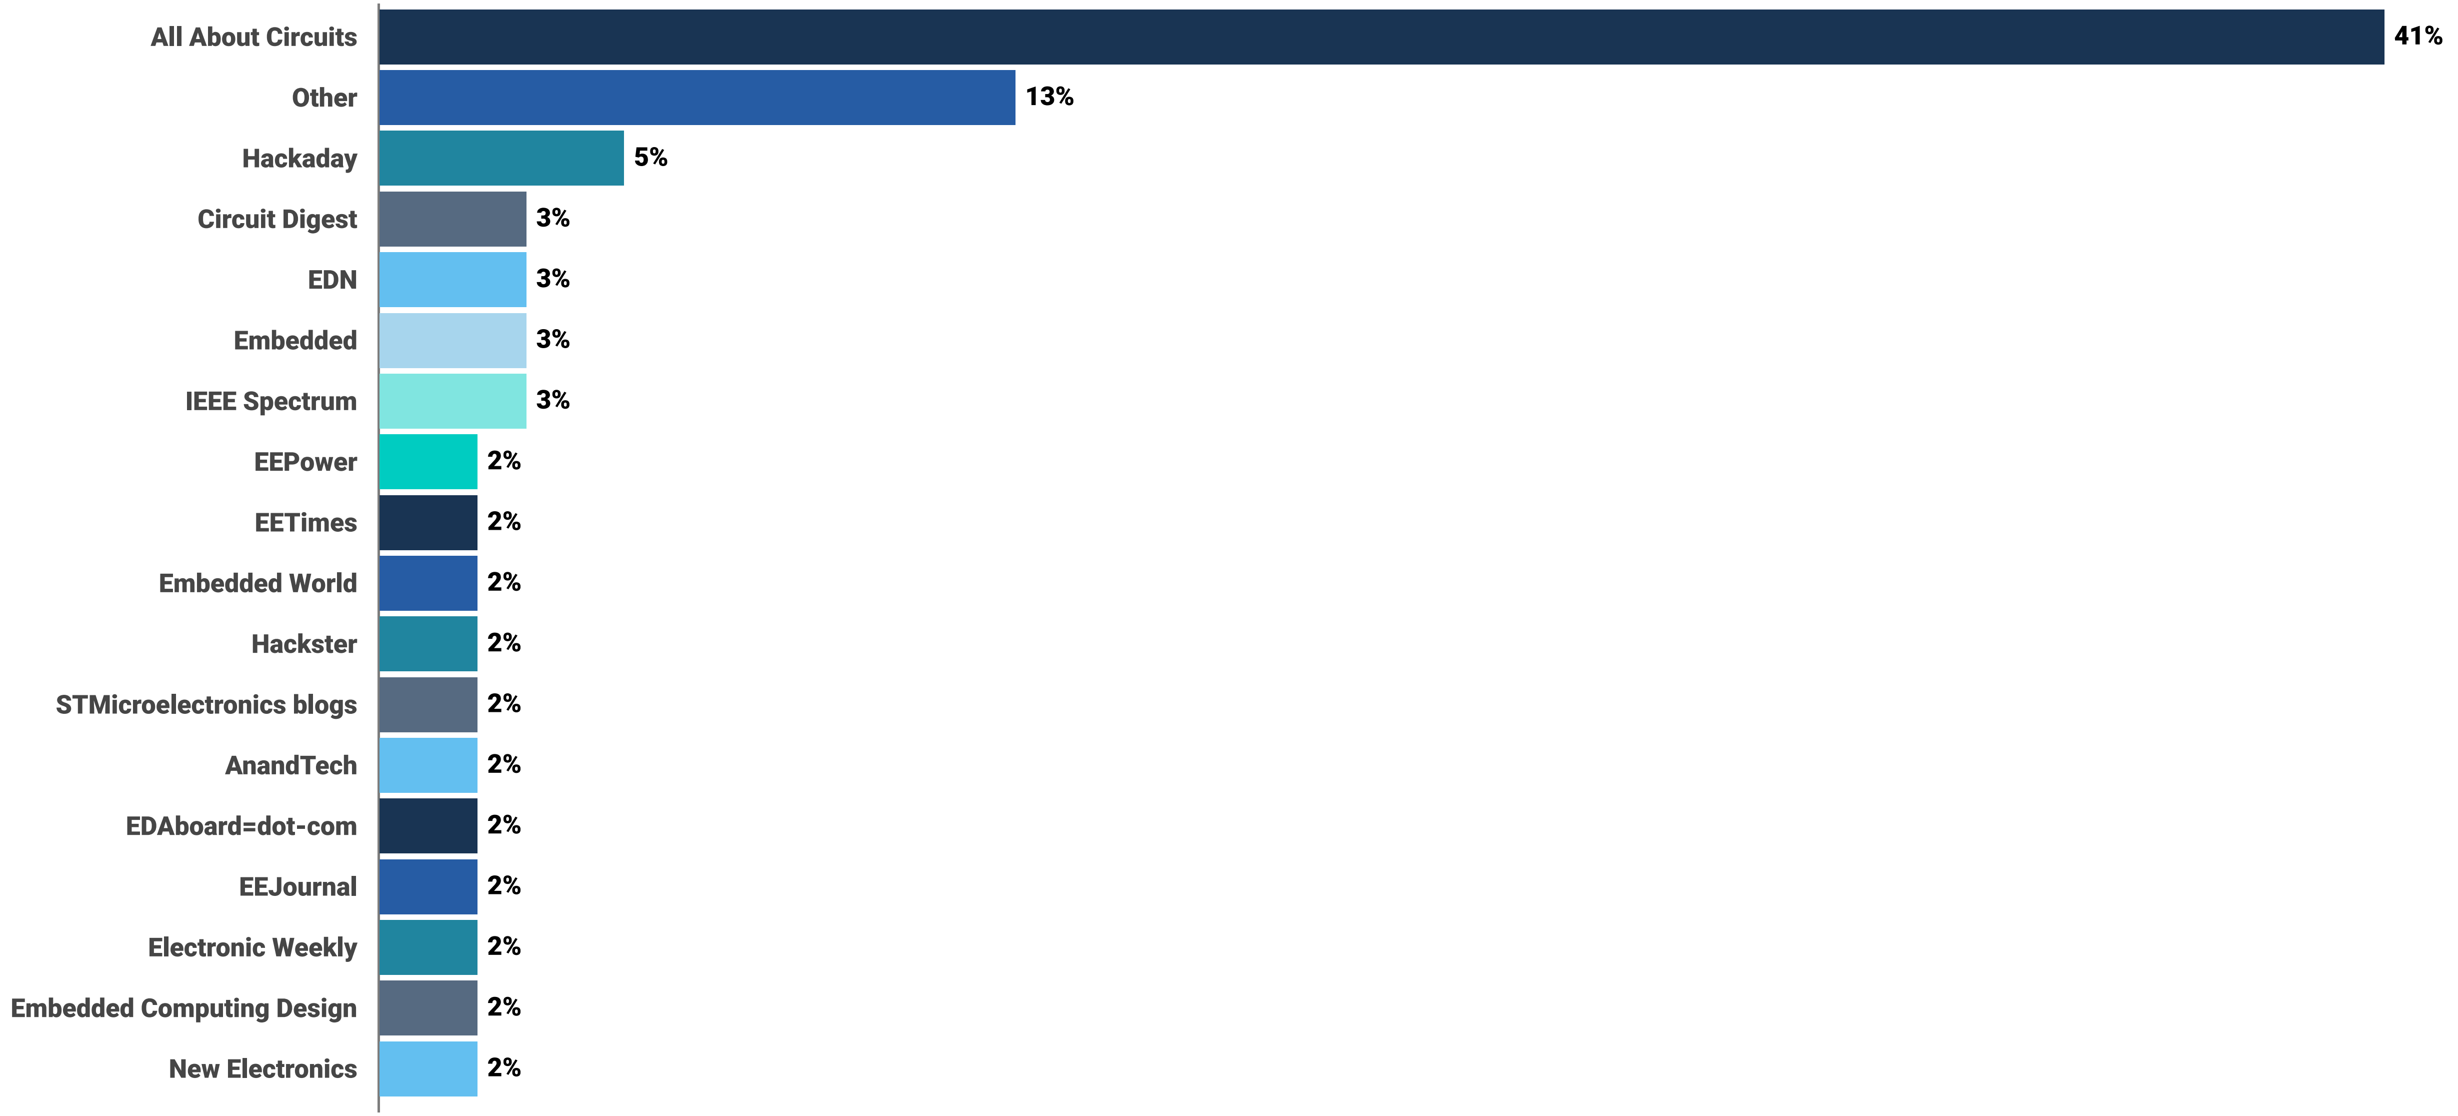

n=123

Answer options at/below 1% were removed - for the full list, please see the Appendix

Media

Which of the following media is your "go to" source for information on AI related embedded systems topics?

All About Circuits is the dominant “go-to” source for AI-related embedded systems information—far ahead of any other publication or platform.

- No single other outlet exceeds 5%, with Hackaday being the closest competitor, followed by a broad set of niche and vendor-affiliated media, each cited by 2–3% of respondents.
- 13% selected “Other,” suggesting that some engineers turn to independent outlets, specialized forums, vendor documentation, or AI-focused technical communities outside traditional media.

The data reinforces that All About Circuits is the clear leader for embedded AI education and insights, commanding a uniquely trusted position among engineers exploring the intersection of AI and hardware design.

Media

Which ONE resource do you consider the most influential OVERALL within the embedded systems market?

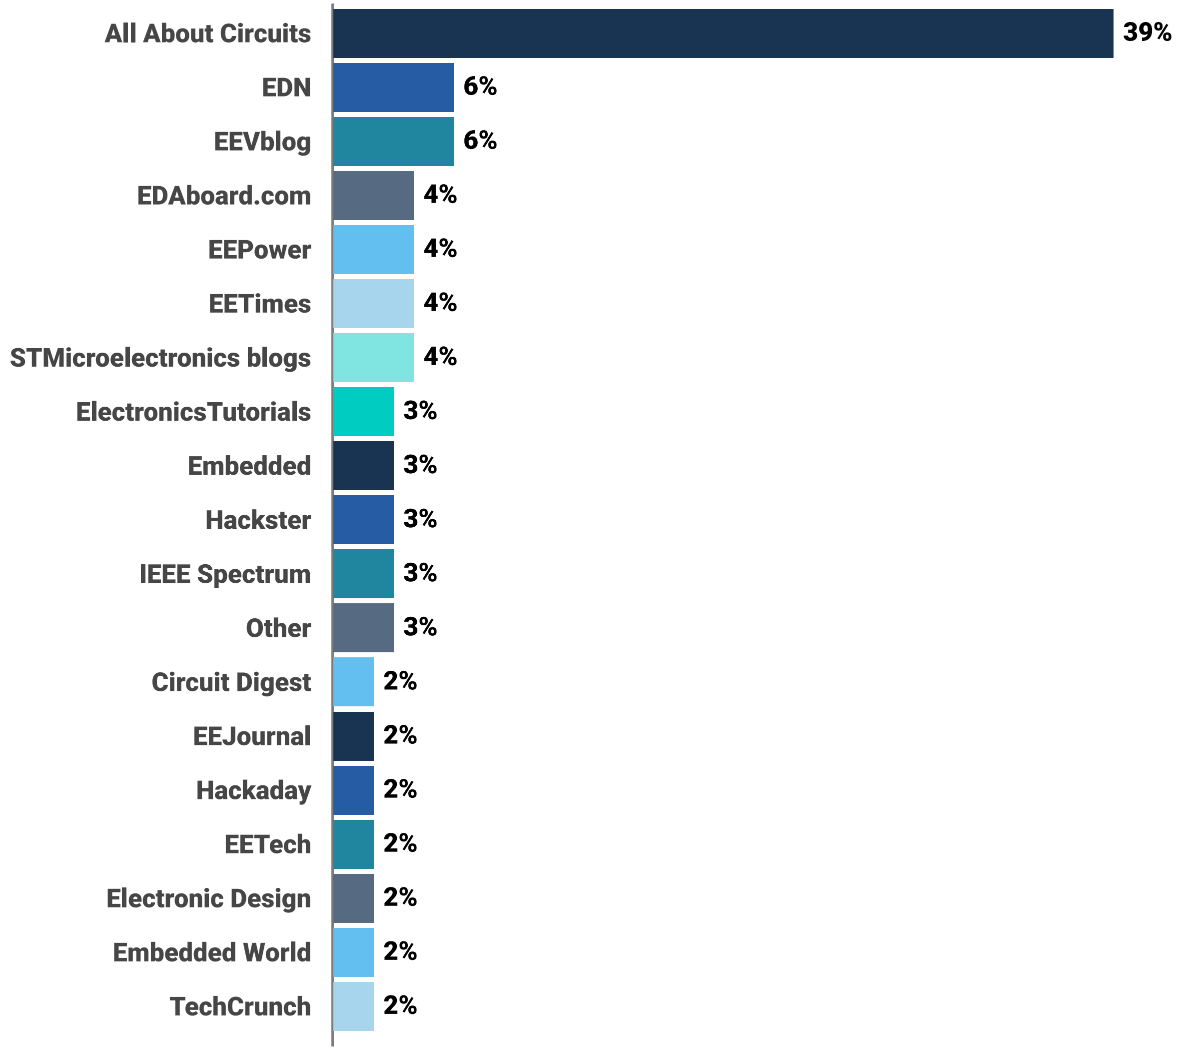

n=124

Answer options at/below 1% were removed - for the full list, please see the Appendix

All About Circuits is viewed as the most influential resource in the embedded systems market, far surpassing all other media outlets.

- The next tier (EDN, EEVblog, and EDAboard.com) illustrates that while several legacy and community-driven sources retain relevance, none come close to AAC's influence.
- EEPower and STMicroelectronics blogs also register notable mentions, reflecting the growing importance of both power systems content and vendor-led education in the embedded ecosystem.

The wide dispersion among the remaining outlets reflects how engineers coalesce around a few trusted sources, with All About Circuits serving as the clear authority and central hub for industry insight and technical learning.

Media Overview

Across every measure, All About Circuits (AAC) stands out as the dominant media resource within the embedded systems ecosystem.

- 71% of respondents report using All About Circuits for general embedded and semiconductor news—more than double any other outlet.
- When focused specifically on AI-related embedded topics, 41% identify AAC as their primary source, far ahead of Hackaday (5%) and other niche platforms.
- When asked to name the most influential media resource overall, 39% again chose All About Circuits, with the next most-cited outlets—EDN and EEVblog—trailing at just 6% each.

Because this study was fielded among members of the All About Circuits community, these findings likely reflect some degree of brand familiarity and audience affinity.

Even so, the consistency and magnitude of AAC's lead across multiple questions suggest that its influence extends well beyond its own readership—positioning it as one of the most trusted and widely recognized information sources in the embedded engineering market. The strength of All About Circuits as both a learning hub and trusted reference point underscores its value as a primary engagement platform for reaching embedded and AI-focused engineers.

Even accounting for sample bias from the AAC community, the platform's consistent leadership across awareness, usage, and influence indicates a deep level of brand trust and technical credibility that few competitors match.

Learning Methods & Media Consumption - Key Takeaways

Embedded engineers are self-directed learners who rely heavily on community-based and technical media for professional development. 64% of respondents report learning through articles in engineering-focused publications such as Embedded.com, EETimes, and All About Circuits. Self-learning remains the second most common method, with half of respondents using blogs, books, or online tutorials to build skills independently. Another 49% turn to video-based learning, most often through YouTube tutorials, while 42% participate in structured online courses focused on embedded AI or related technologies.

Formal education plays a relatively small role. Only 17% of respondents rely on university courses, and 14% pursue professional online AI certifications. Instead, most engineers emphasize practical experience and experimentation, often through open-source projects, online forums, and peer collaboration. This indicates that embedded professionals view learning as a continuous, application-driven process rather than a structured academic pursuit.

When it comes to where engineers consume industry information, All About Circuits stands out. 71% of respondents cite it as a source for staying informed, making it the top media platform among all options. It is also the most frequently named "go-to" source for AI-related embedded topics and the most influential publication overall. Other popular sources include Hackaday, IEEE Spectrum, and EDN.

However, because this survey was deployed to the All About Circuits audience, its dominant position is likely influenced by sample bias. Even so, the results reinforce that All About Circuits has built a high level of trust and engagement within the embedded engineering community.

Learning Methods & Media Consumption - Implications for Marketers

For semiconductor and technology marketers, these findings reaffirm that engineers trust educational, technically credible media far more than promotional channels. Engineers turn to platforms like All About Circuits because they are looking for insight, not advertising. Successful marketing strategies should build credibility within these trusted environments, positioning brands as contributors to engineers' learning journeys rather than as advertisers competing for attention.

While the strong representation of All About Circuits is partly explained by the survey audience, the pattern reflects a broader truth: engineers form long-term relationships with media ecosystems that consistently deliver practical value. Marketers should invest in content partnerships, sponsored articles, and webinars that provide engineers with actionable takeaways—tutorials, design walk-throughs, and data-backed performance insights. These formats perform best because they acknowledge how engineers prefer to learn: hands-on, precise, and peer-relevant.

Marketers can also use these insights to better target their educational investments. Because formal training plays a minor role, engineers depend on brands and media to fill the gap between theory and implementation. Providing technical resources, open-access training modules, or on-demand demos positions a company as an enabler of professional growth. This type of content builds brand affinity and trust that extends beyond individual campaigns.

Finally, the clear preference for community-driven learning means marketers should treat engineering media as ecosystems of engagement rather than single publishing outlets. Brands that participate consistently—through articles, event sponsorships, and interactive learning formats—gain long-term visibility and credibility. In the embedded world, trust is cumulative. The companies that teach engineers something new every time they engage are the ones engineers will turn to when it's time to design, specify, and buy.

Recommendations

Recommendations for Marketers

Lead with Enablement, Not Promotion

Embedded engineers are problem solvers first and buyers second. Marketing should help them design, test, and implement solutions more efficiently. Focus on educational content—white papers, webinars, reference designs, and how-to videos—that provides practical benefits and answers real-world engineering questions. The brands that make complex technology easier to understand build lasting credibility and preference.

Speak the Language of Proof

Engineers make decisions based on measurable and repeatable results, not marketing claims. Use benchmark data, power-consumption figures, and latency comparisons to show exactly how your products perform in real applications. Replace aspirational language with technical clarity. Case studies showing verified performance gains resonate far more than general statements about innovation or leadership.

Participate Where Engineers Learn

Engineers trust familiar technical media ecosystems like All About Circuits, Hackaday, and IEEE Spectrum because these platforms provide value, not noise. Invest in sponsored educational content, community-driven Q&As, and co-branded webinars that deliver insight and offer solutions rather than promotion. Consistent participation in these trusted spaces builds recognition and positions your brand as part of the engineering community.

Recommendations for Marketers

Support Open-Source Ecosystems

The study shows that most embedded projects incorporate open-source components. Marketers should meet engineers where they work—on GitHub, TensorFlow Hub, and developer forums. Contributing example code, SDKs, and tutorials creates organic visibility and trust. Open, documented engagement signals technical transparency—a quality engineers reward with loyalty.

Align Messaging with Design Reality

Engineers face daily constraints in power, latency, cost, and data access. Marketing should directly acknowledge those trade-offs and show how their technology mitigates them. Position your solutions realistically, highlighting how they reduce integration effort or simplify compliance. Avoid overclaiming, support claims with verifiable data and be transparent about any limitations, focusing on how your product helps achieve specific design goals.

Bridge the Edge and Cloud Conversation

Developers increasingly use hybrid architectures that combine on-device AI with cloud-based training or updates. Marketers should emphasize interoperability, flexible deployment options, and performance consistency across environments. Content that shows how products operate seamlessly between edge and cloud contexts will appeal to teams planning for long-term scalability.

Recommendations for Marketers

Elevate Security and Reliability as Brand Values
Security is now a defining feature of engineering trust. Highlight secure boot processes, encryption standards, and design-for-safety capabilities in your messaging. Publish technical validation data or third-party testing results where possible. Reliability and protection are differentiators—especially in automotive, industrial, and medical applications where failure is unacceptable.

Showcase Complete Ecosystems, Not Individual Parts
Engineers prefer suppliers that provide everything they need to move from concept to prototype. Present your hardware, software, and development tools as an integrated ecosystem. Highlight how your solutions reduce time-to-market and minimize friction across the design process. A cohesive narrative of enablement will stand out in a crowded component marketplace.

Build Authentic, Lasting Relationships with Engineers
Embedded engineers value clarity, reliability, and authenticity. Localization, empathy, and trust all reinforce the same goal: to become part of the engineer's long-term workflow. Share region-specific success stories, provide translated resources, and maintain consistent support after the design win. Every touchpoint—from datasheet clarity to post-purchase service—builds brand equity. Trust is cumulative; it is earned through accuracy, consistency, and respect for the engineer's intelligence.

Conclusion

The embedded systems community is advancing rapidly toward a future defined by intelligent, efficient, and explainable design. Engineers are integrating AI pragmatically, balancing performance with reliability, and seeking vendors who empower rather than overwhelm.

For marketers, the opportunity lies in education, integration, and credibility—bridging the gap between innovation and application. Brands that communicate transparently, demonstrate practical value, and participate in the engineering learning ecosystem will lead not just in visibility, but in lasting influence across the embedded design landscape.

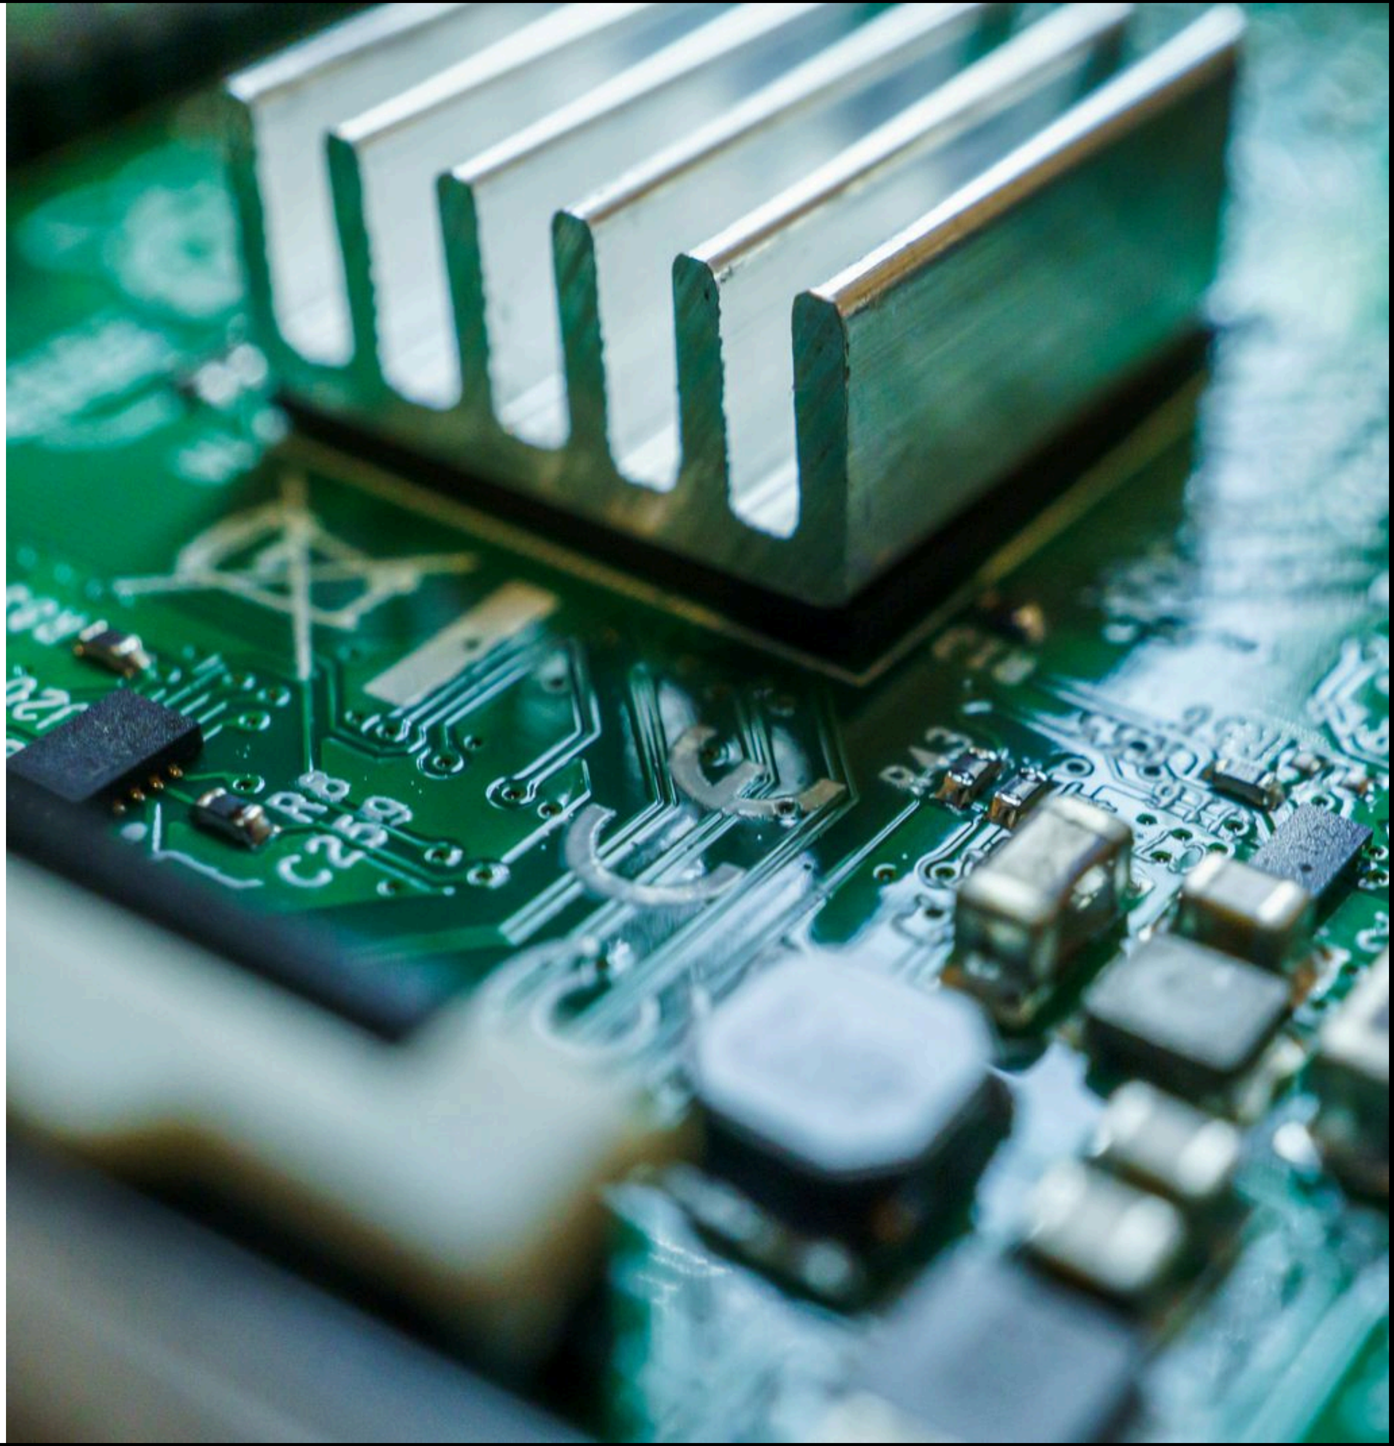

Disclaimers

By purchasing the Wilson Research Group/EETech AI Trends and Innovations Within Embedded Systems, you agree to use this data in a confidential manner internal to your organization only, and to take reasonable and appropriate care to ensure that the data and report are distributed to individuals in your organization on this same non-disclosure basis, and you and they shall not in any way disclose or distribute the data itself, results, summaries of the data or other conclusions from the data, to entities external to your organization including individuals, companies, associations, institutes, publications, media, electronic systems, online or Internet sites, or other organizations and entities. Such disclosure might cause financial harm to Wilson Research Group, LLC or EETech Media, LLC for which you or your organization could be liable, if such disclosure is considered to undermine the market value of this study to others or otherwise limit the value of the study. In short, you agree that this study is for your organization's internal use only.

You further agree that your and your organization's use of this data is entirely at your own risk. Wilson Research Group, LLC and EETech Media, LLC, Inc. may not be held responsible for nor liable for any conclusions, projections, applications, comparisons, investment decisions, promotions, financial considerations, marketing strategies, policies implemented by you or your organization that may result in financial loss, lowered market perceptions, or less than optimum results of policies or administrative decisions based on the use of this data. Wilson Research Group, LLC or EETech Media, LLC will not be held responsible nor liable for any inadvertent errors, tolerance levels or interpretations of the data itself and makes no warranty as to the overall trustworthiness of the data in relation to various markets for which the study was aimed, based on the methodology that all data are represented as approximations with varying statistical tolerances as to confidences in the data as is standard practice in the conduct of market research.

Appendix

Chip Families

Which of the following chip series or chip families are you currently using in your embedded design?

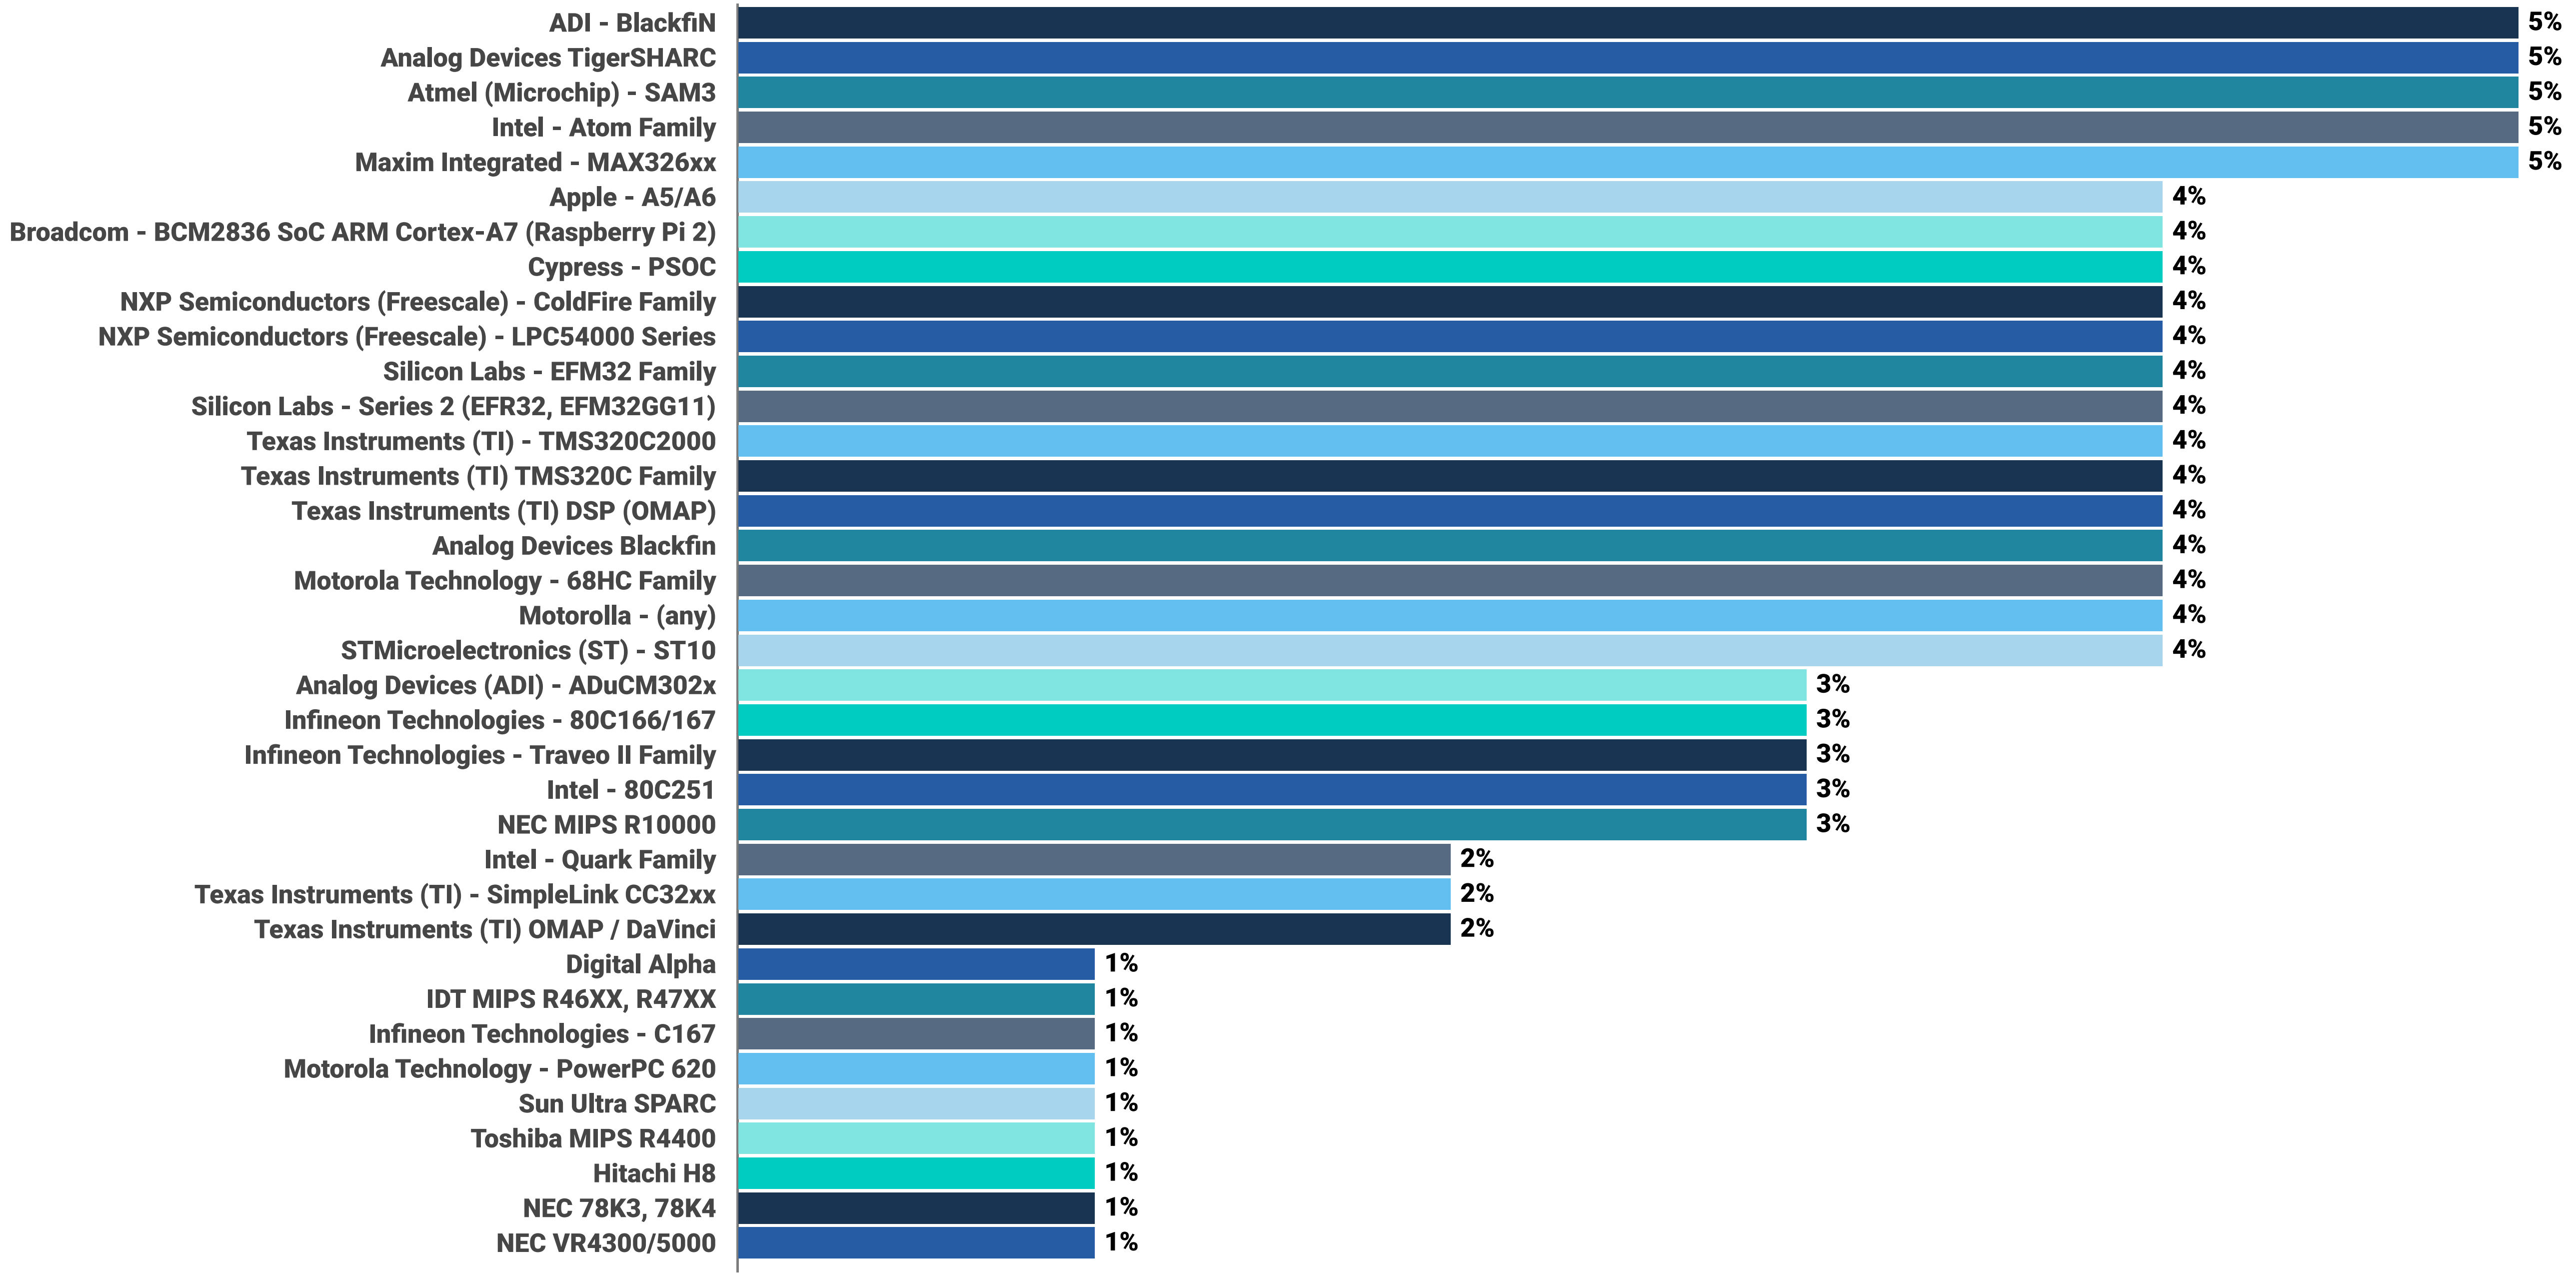

Chip Vendors

Which of the following chip vendors do you currently use for you embedded designs?

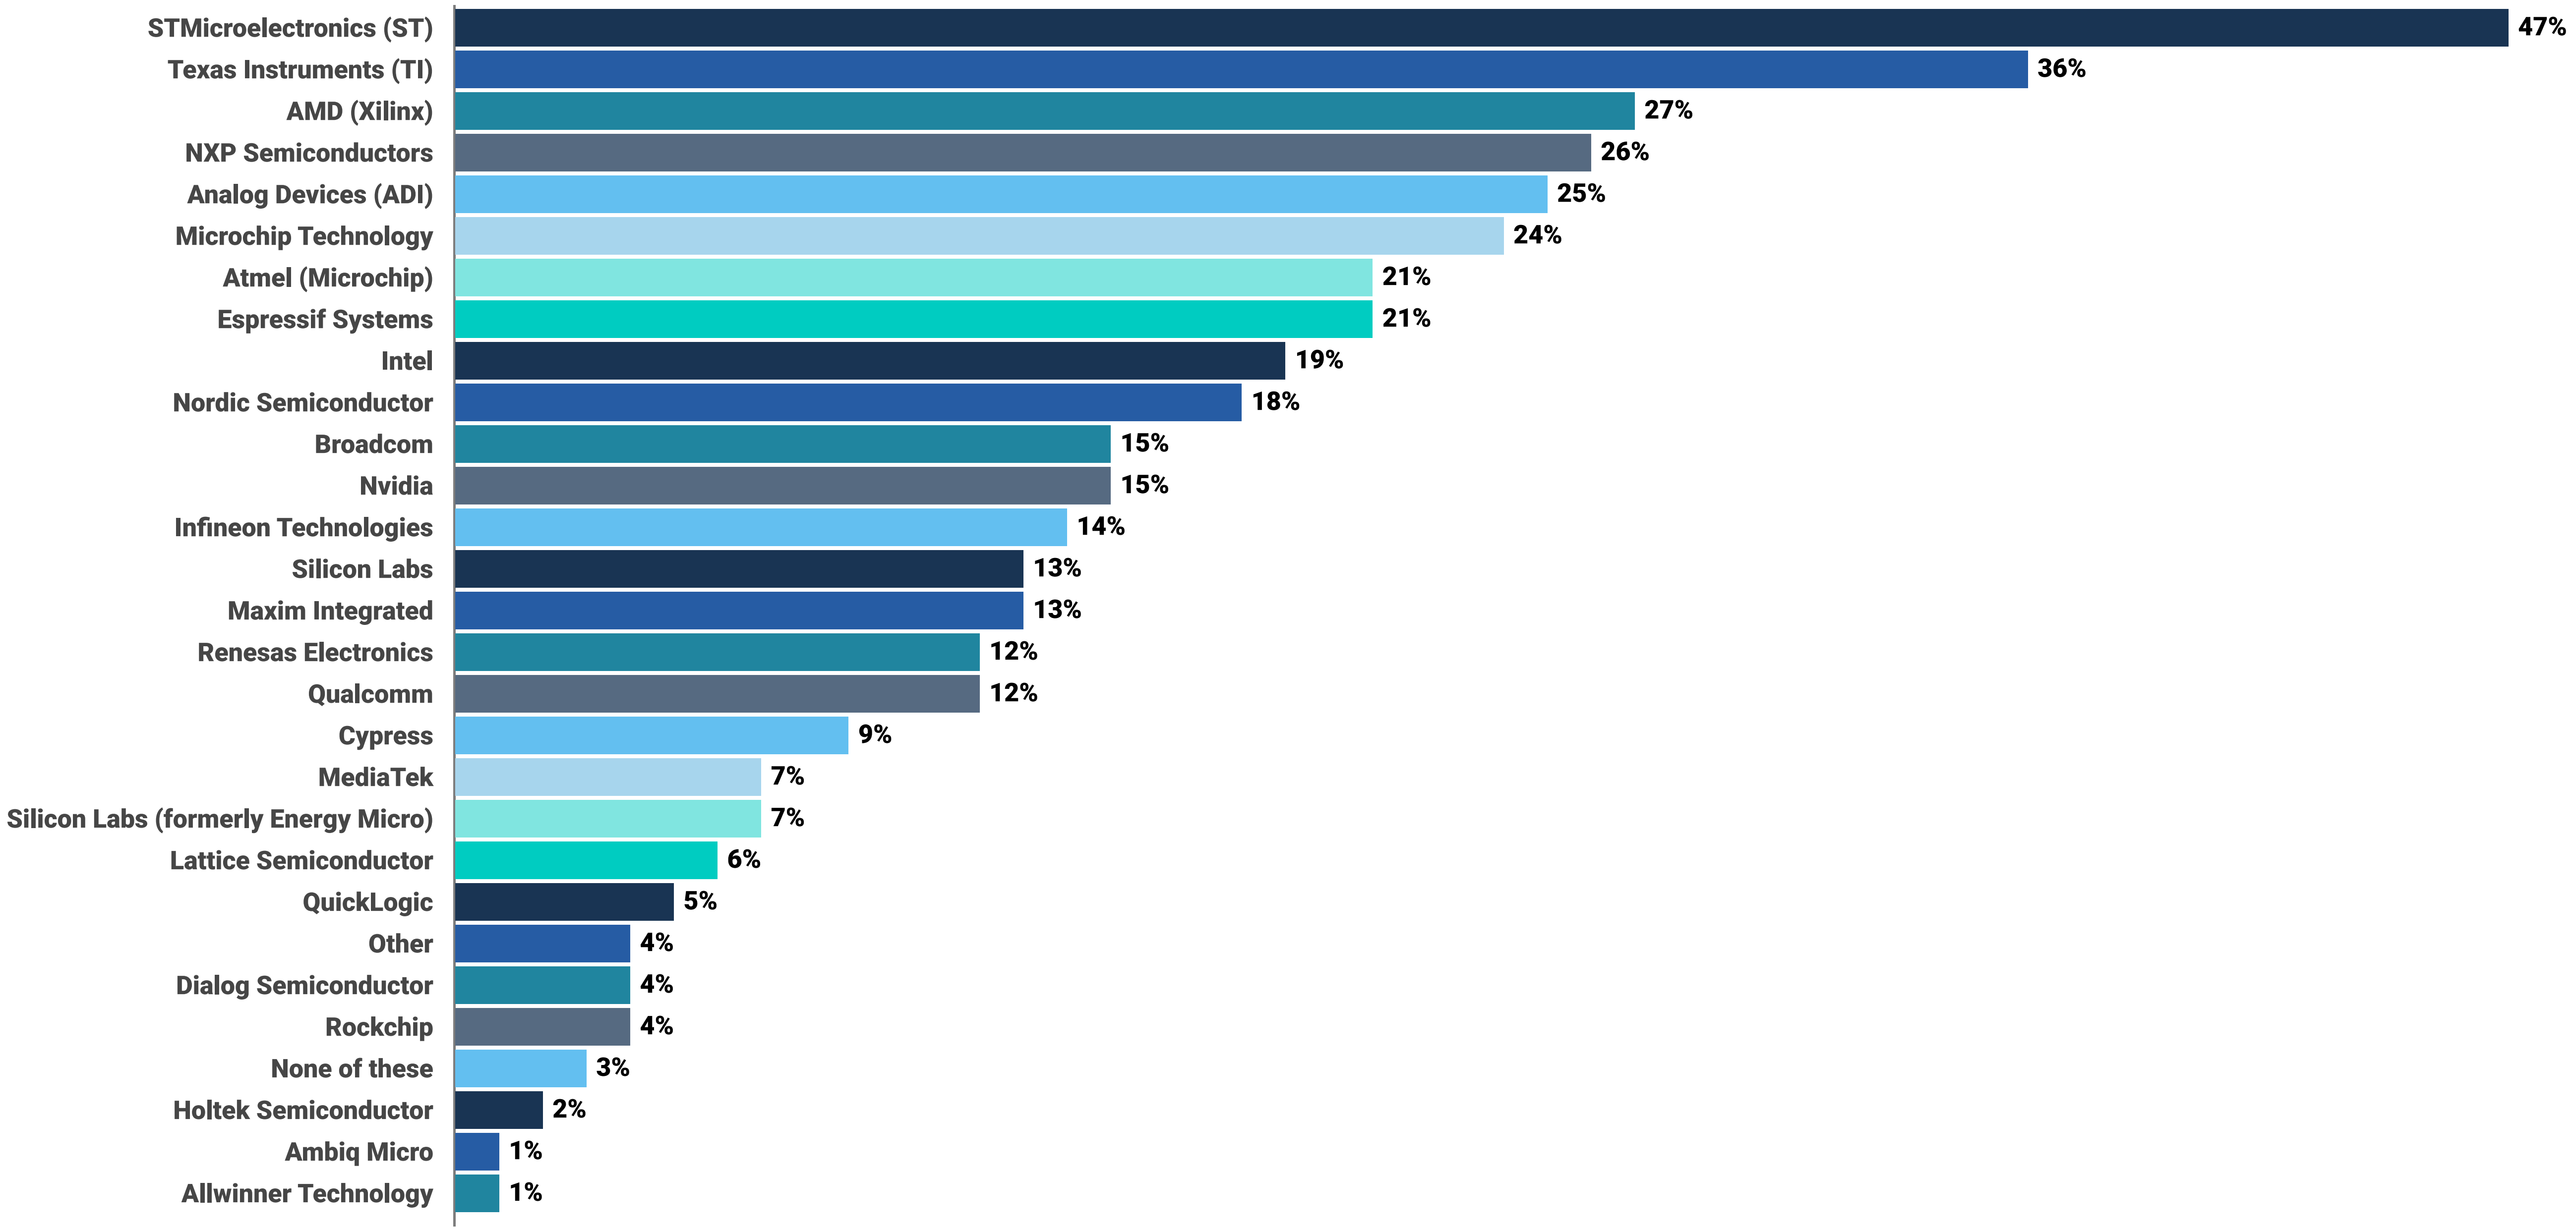

The Best Overall

Which ONE of the following chip vendors do you consider best OVERALL as a vendor for embedded designs?

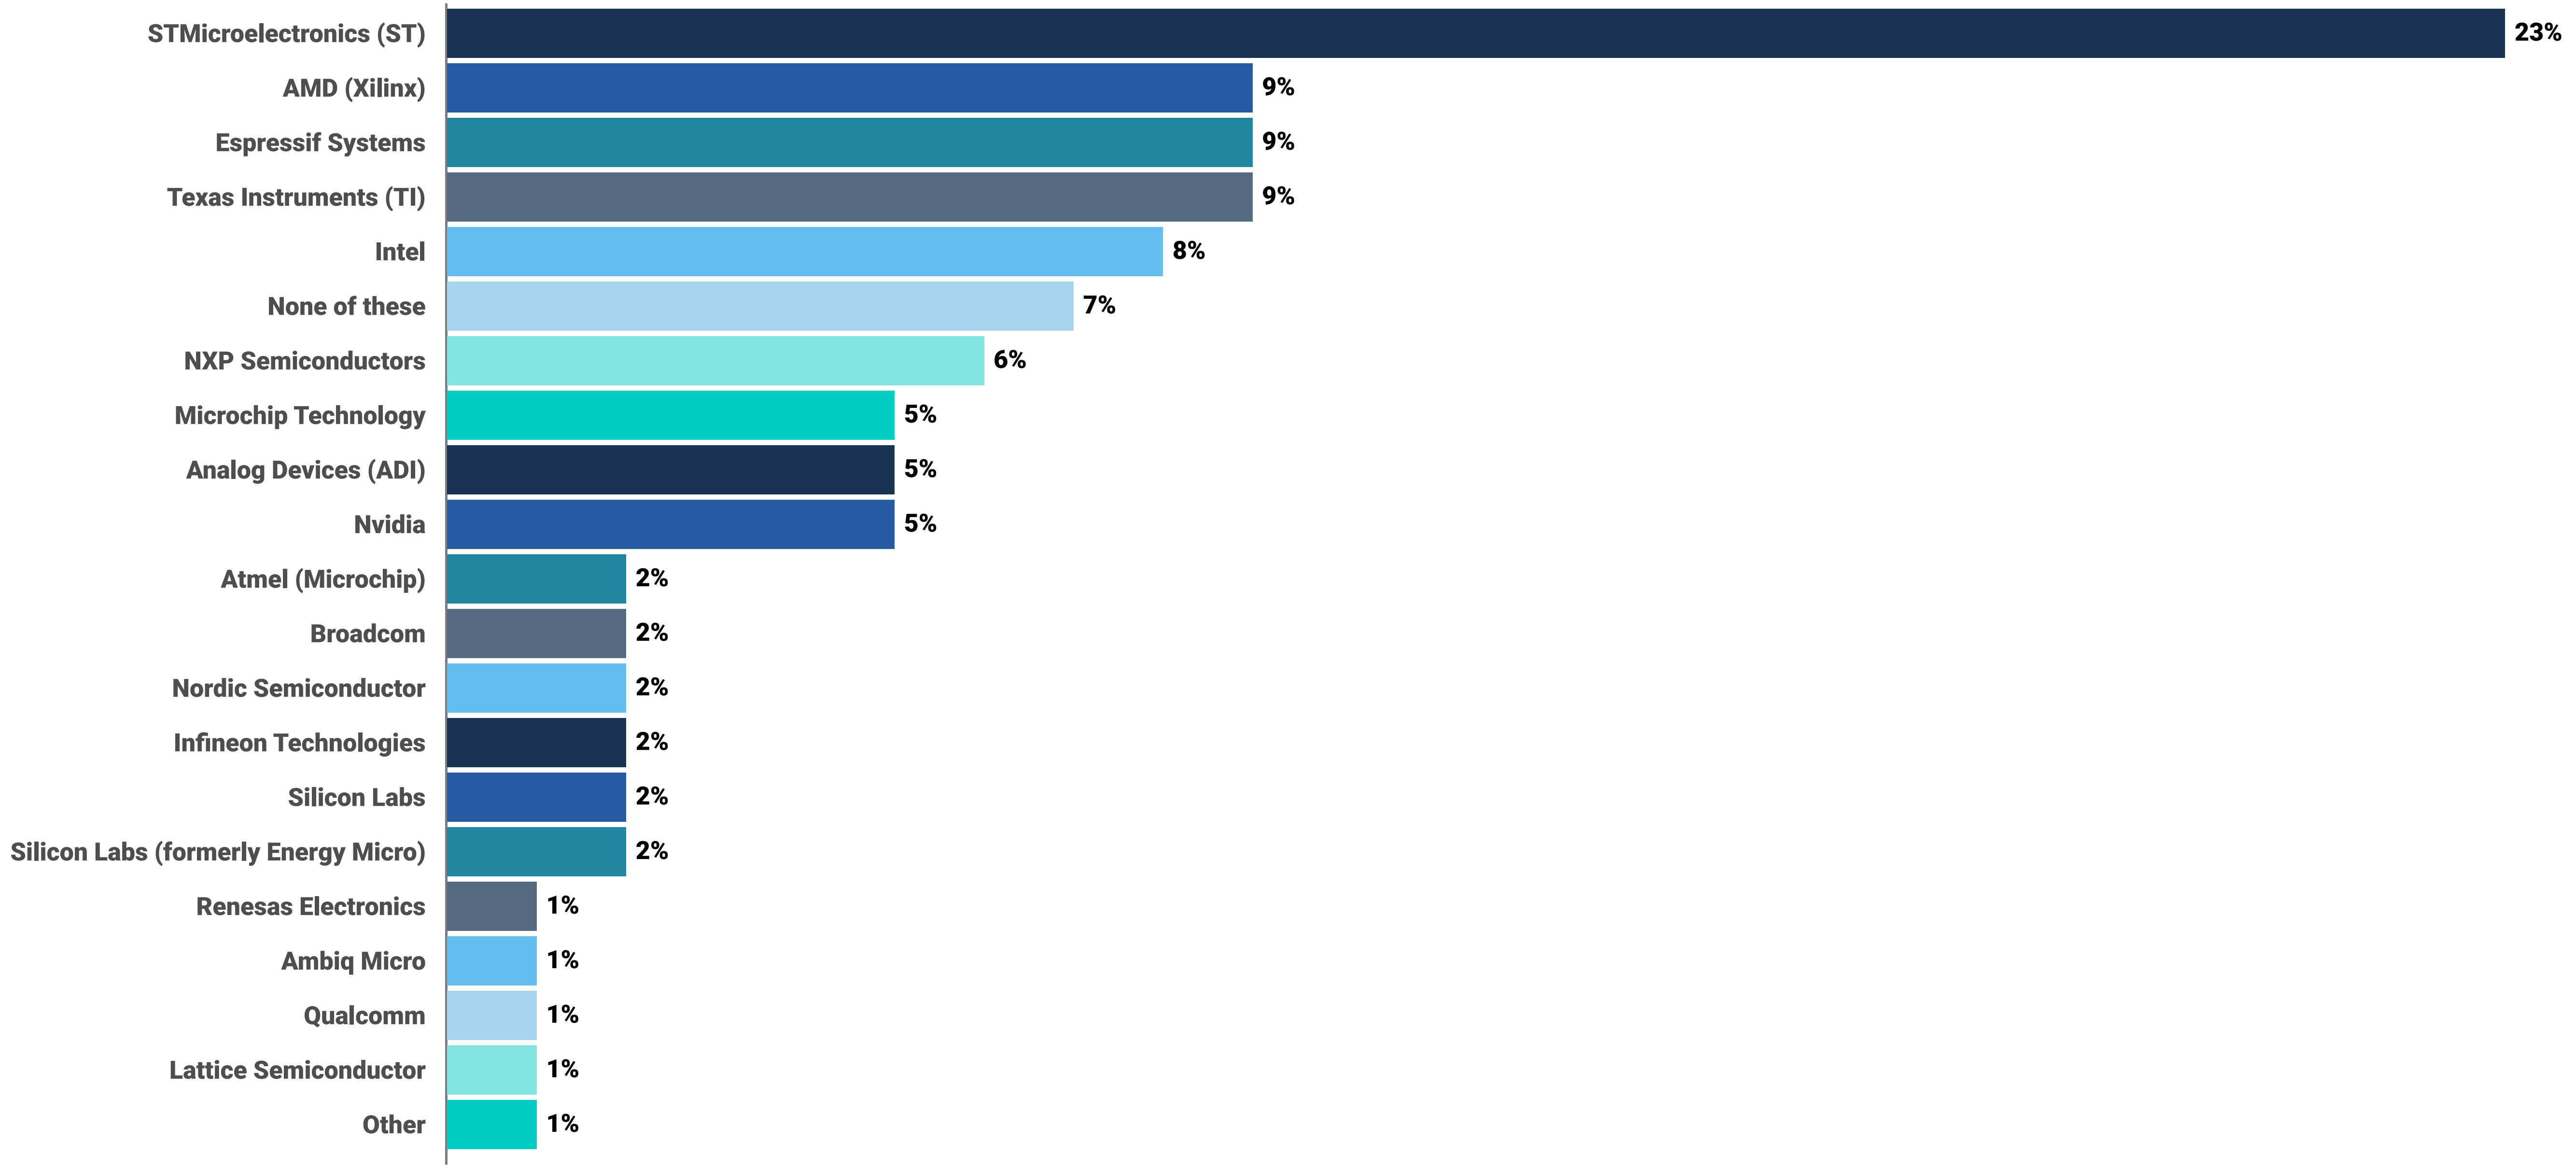

AI Chip Vendors

Which of the following Artificial Intelligence (AI) chip accelerator vendors (if any) are you currently using in you embedded design?

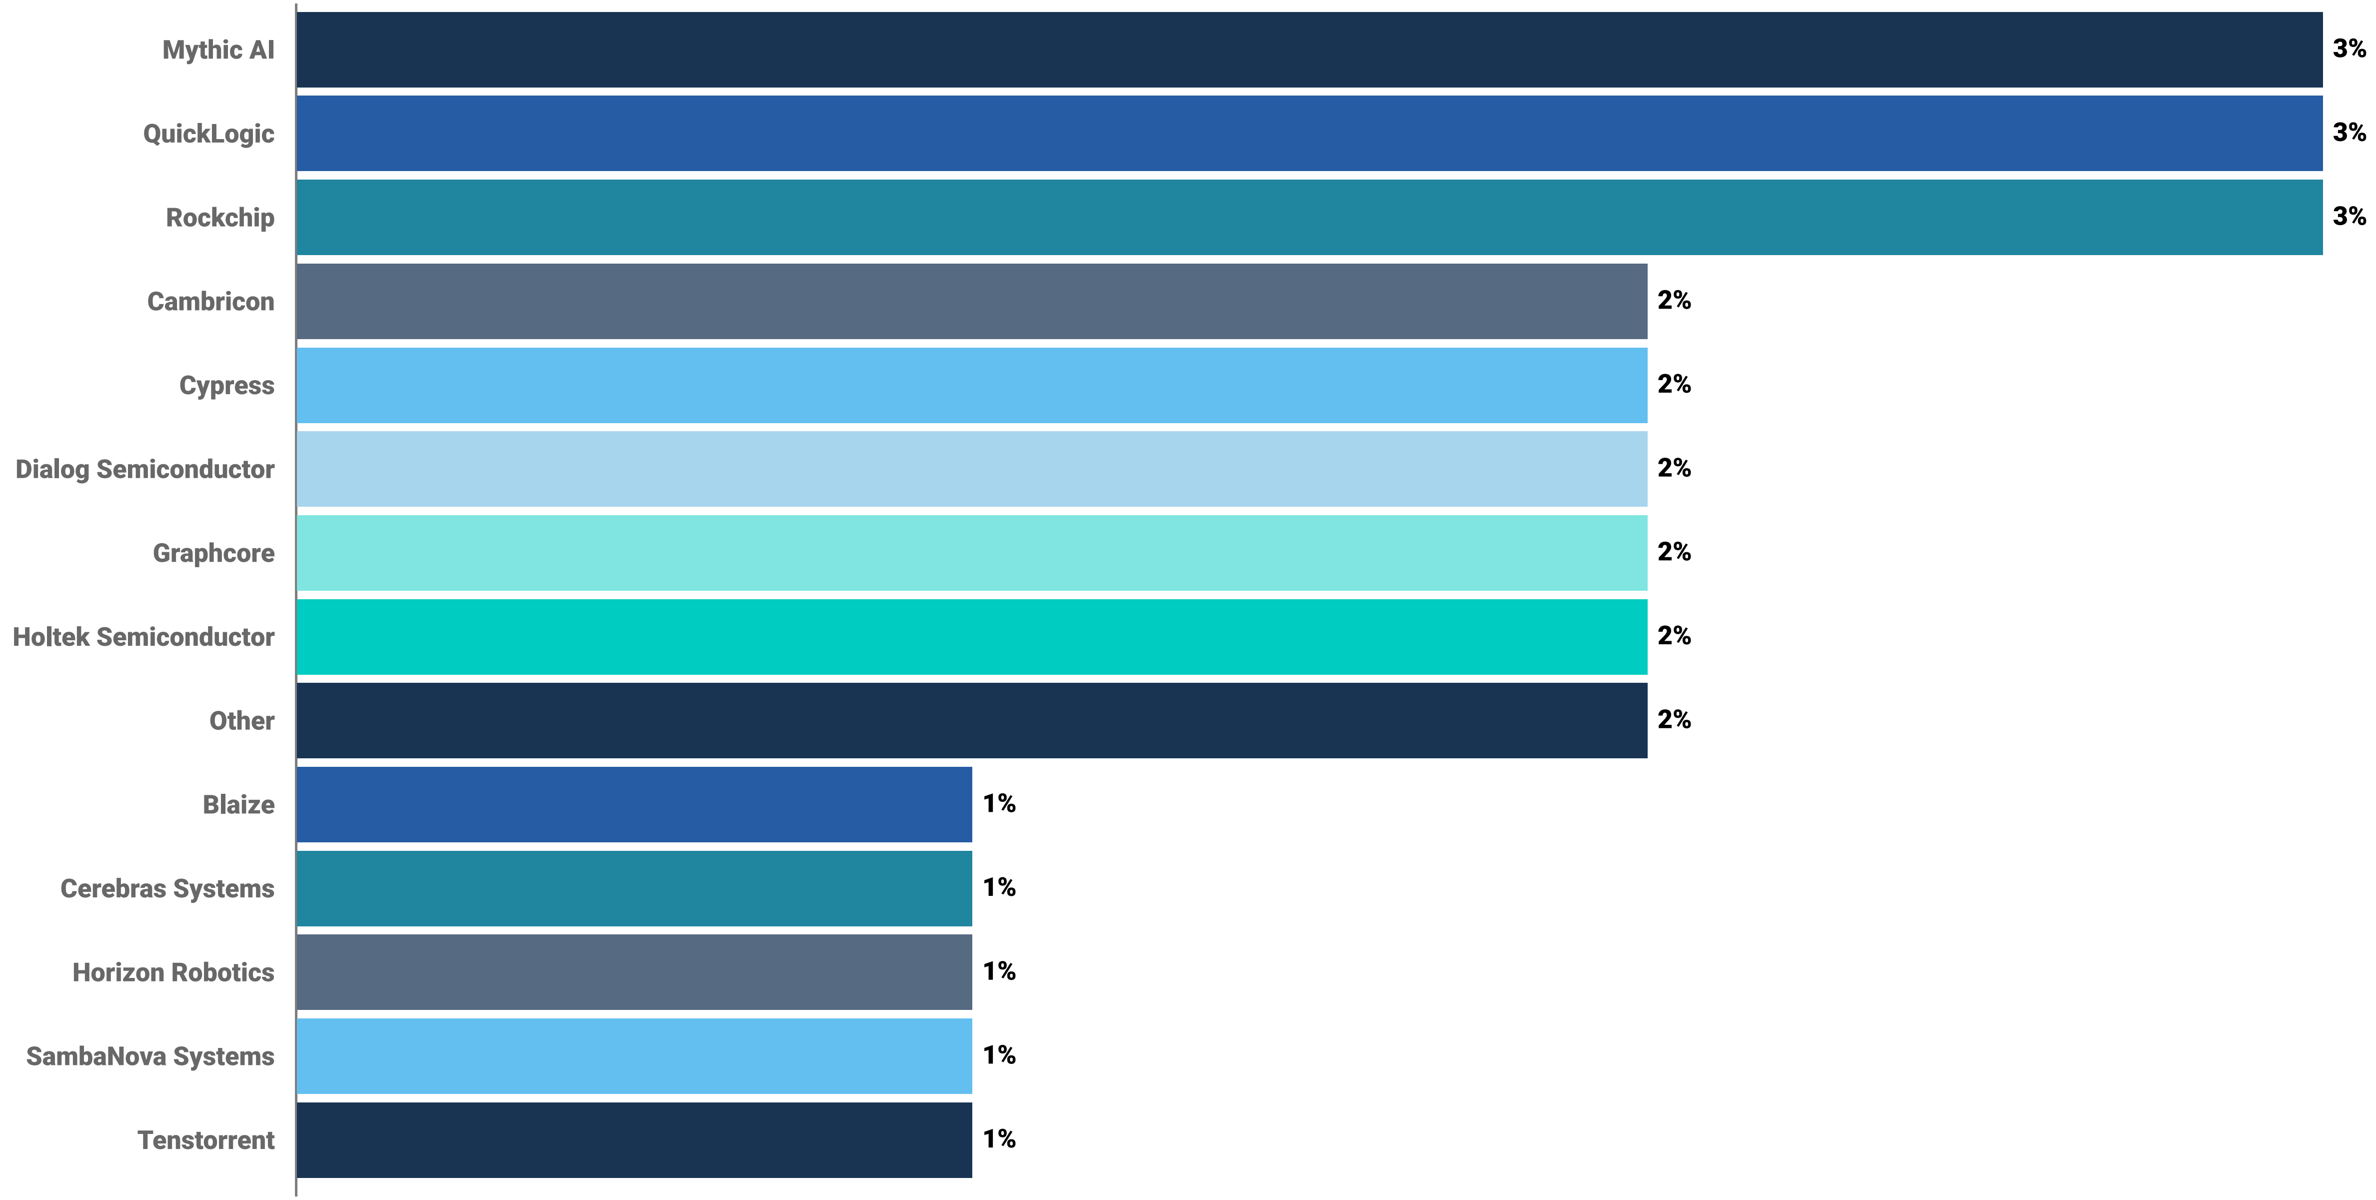

Software

Which of the following software vendors do you currently use in your embedded designs?

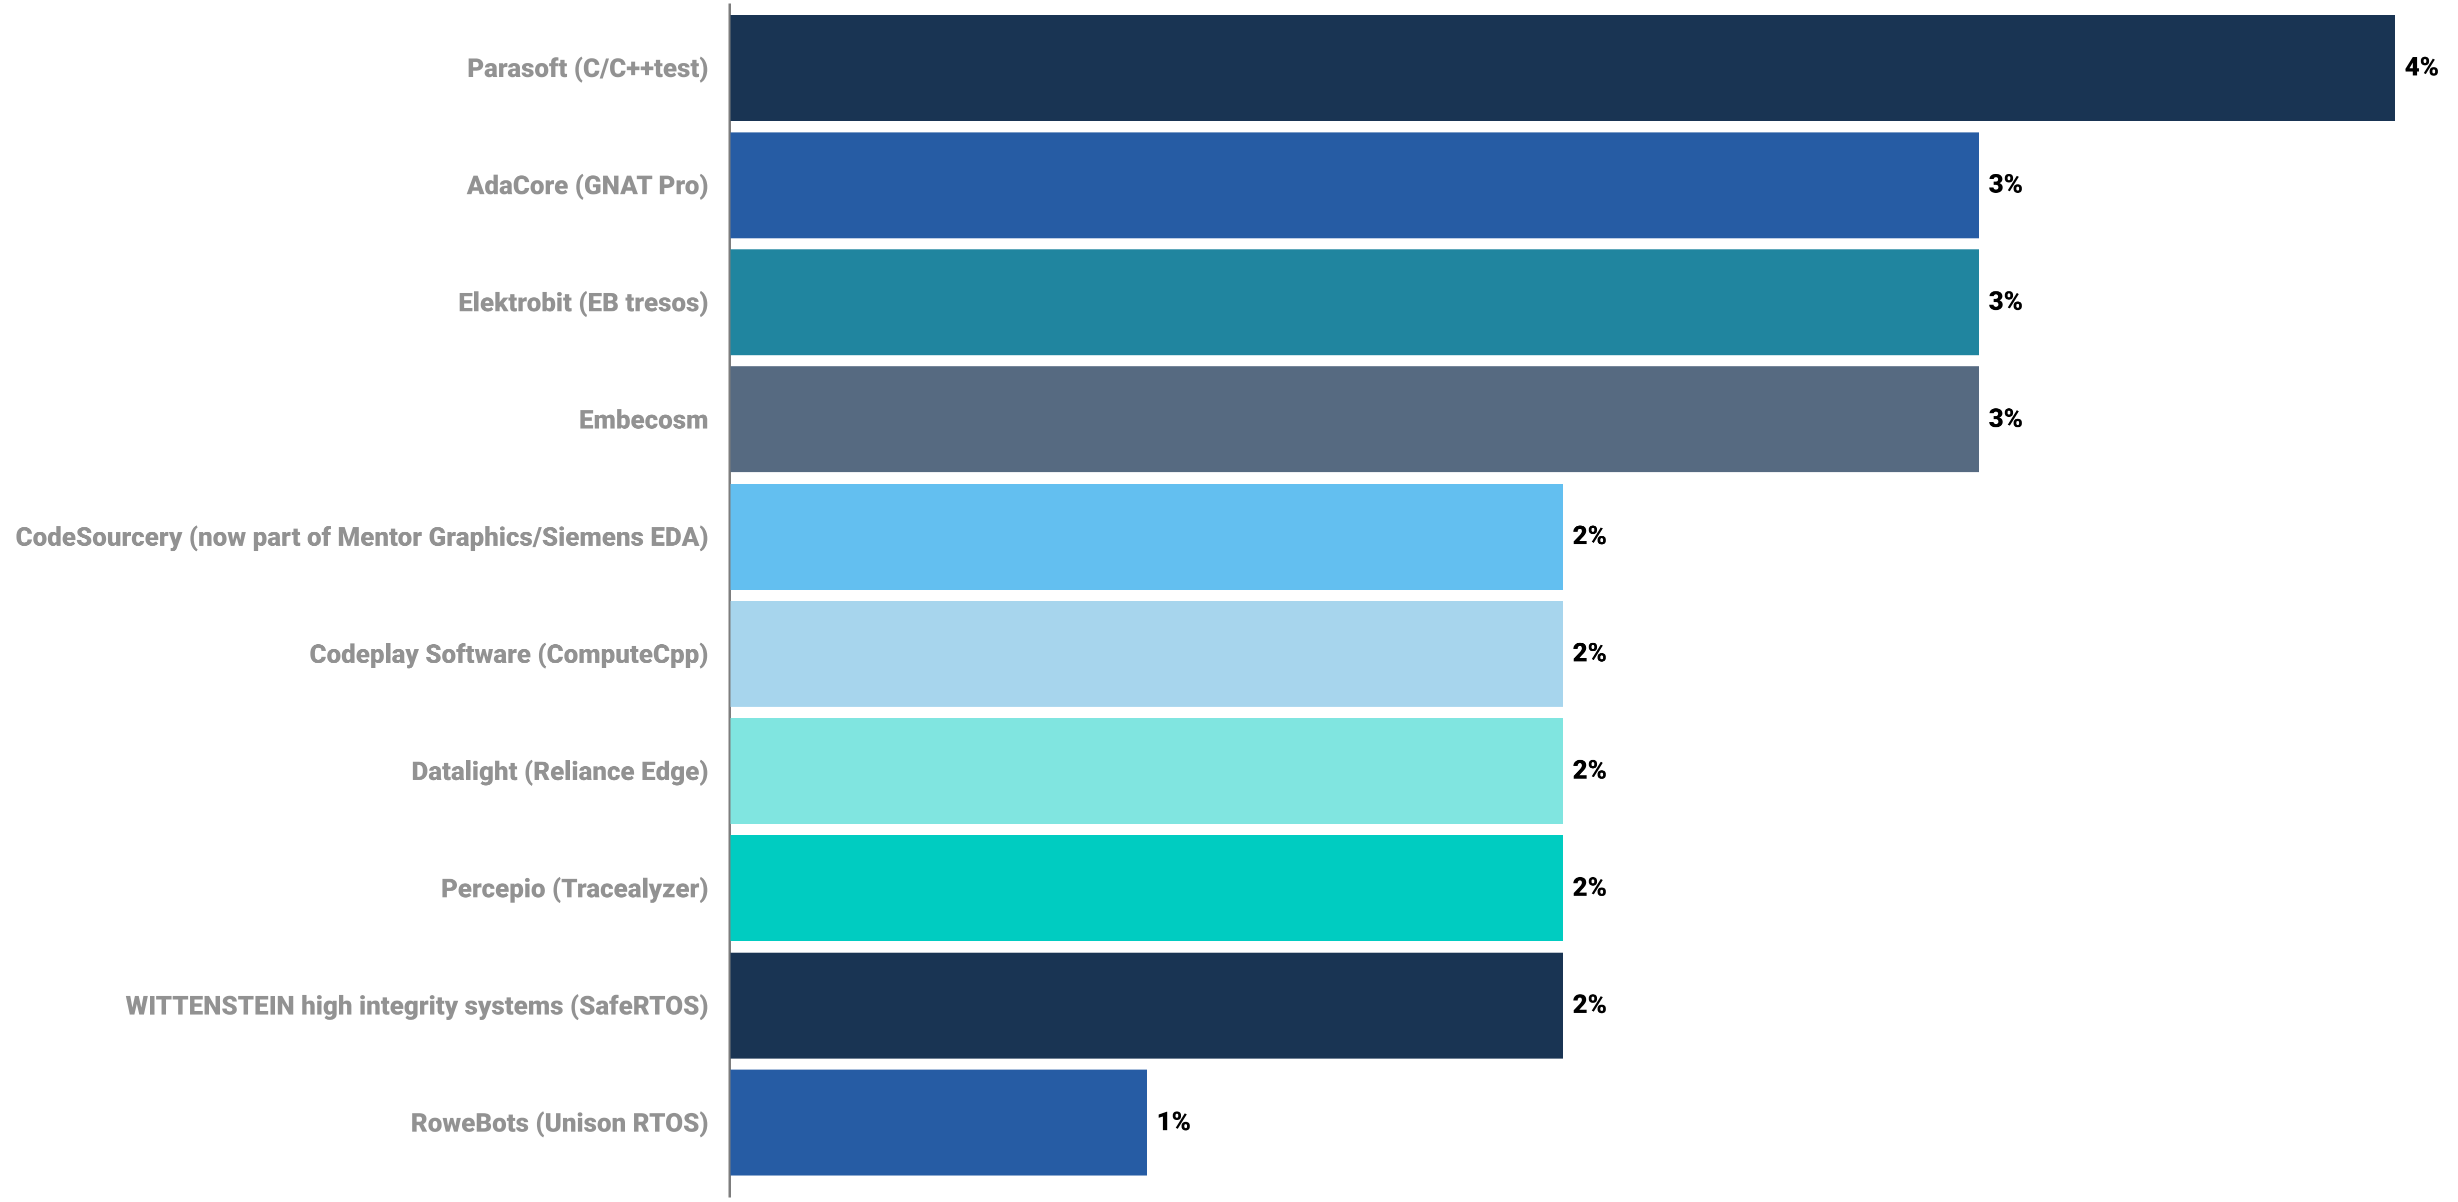

Media

Which of the following specific media do you currently use to learn about embedded systems and the semiconductor industry?

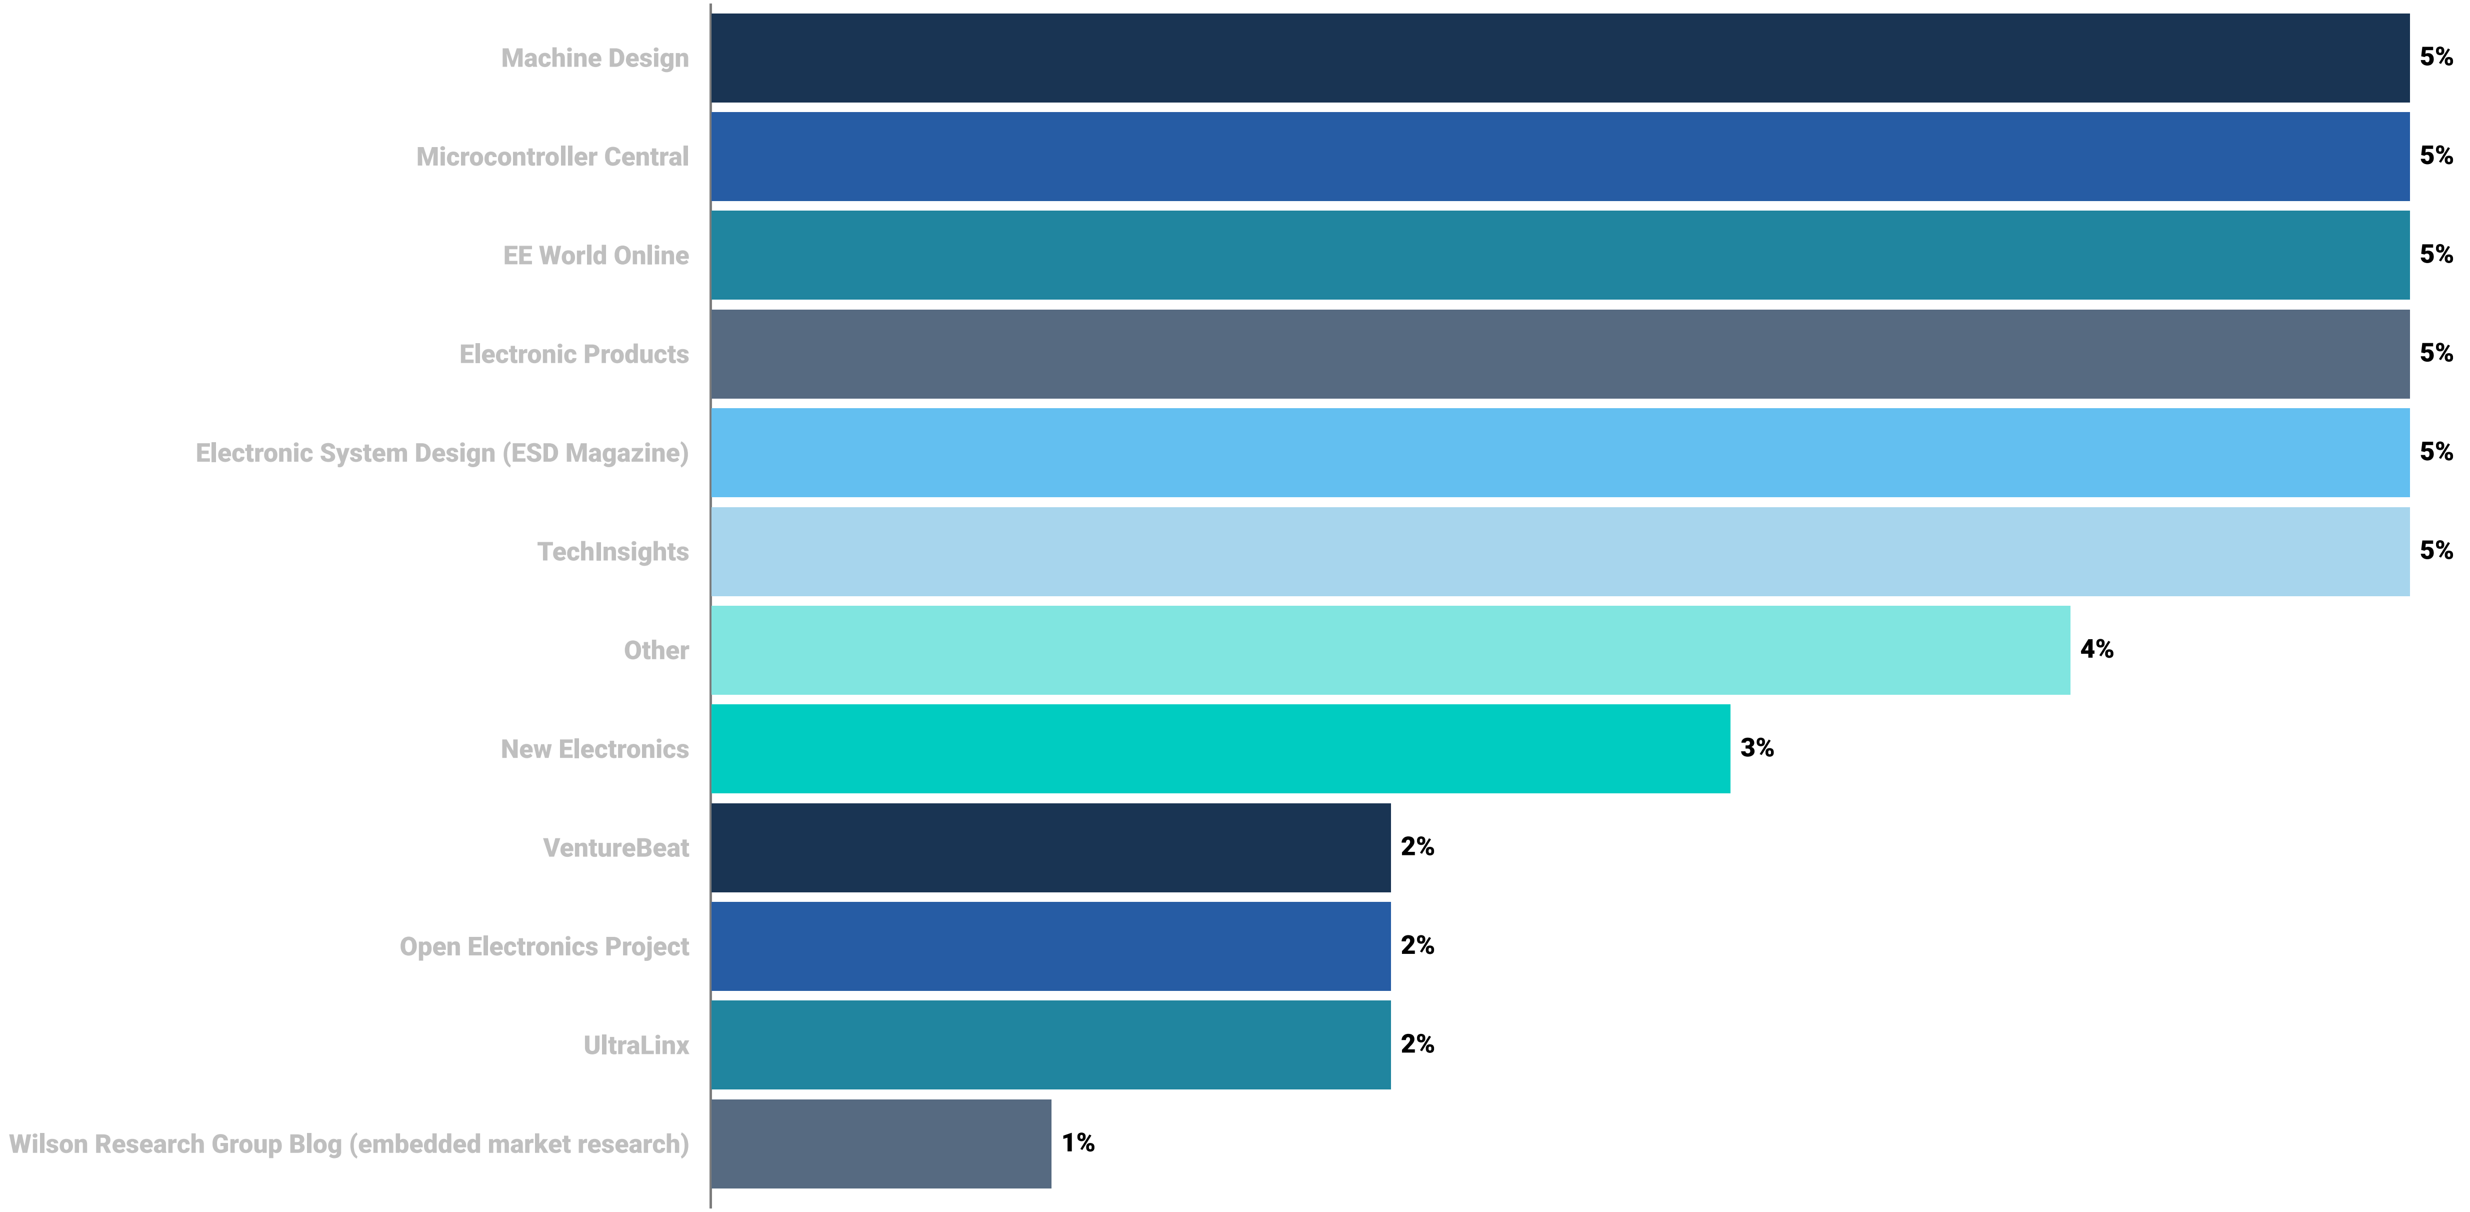

Media

Which of the following media is your "go to" source for information on AI related embedded systems topics?

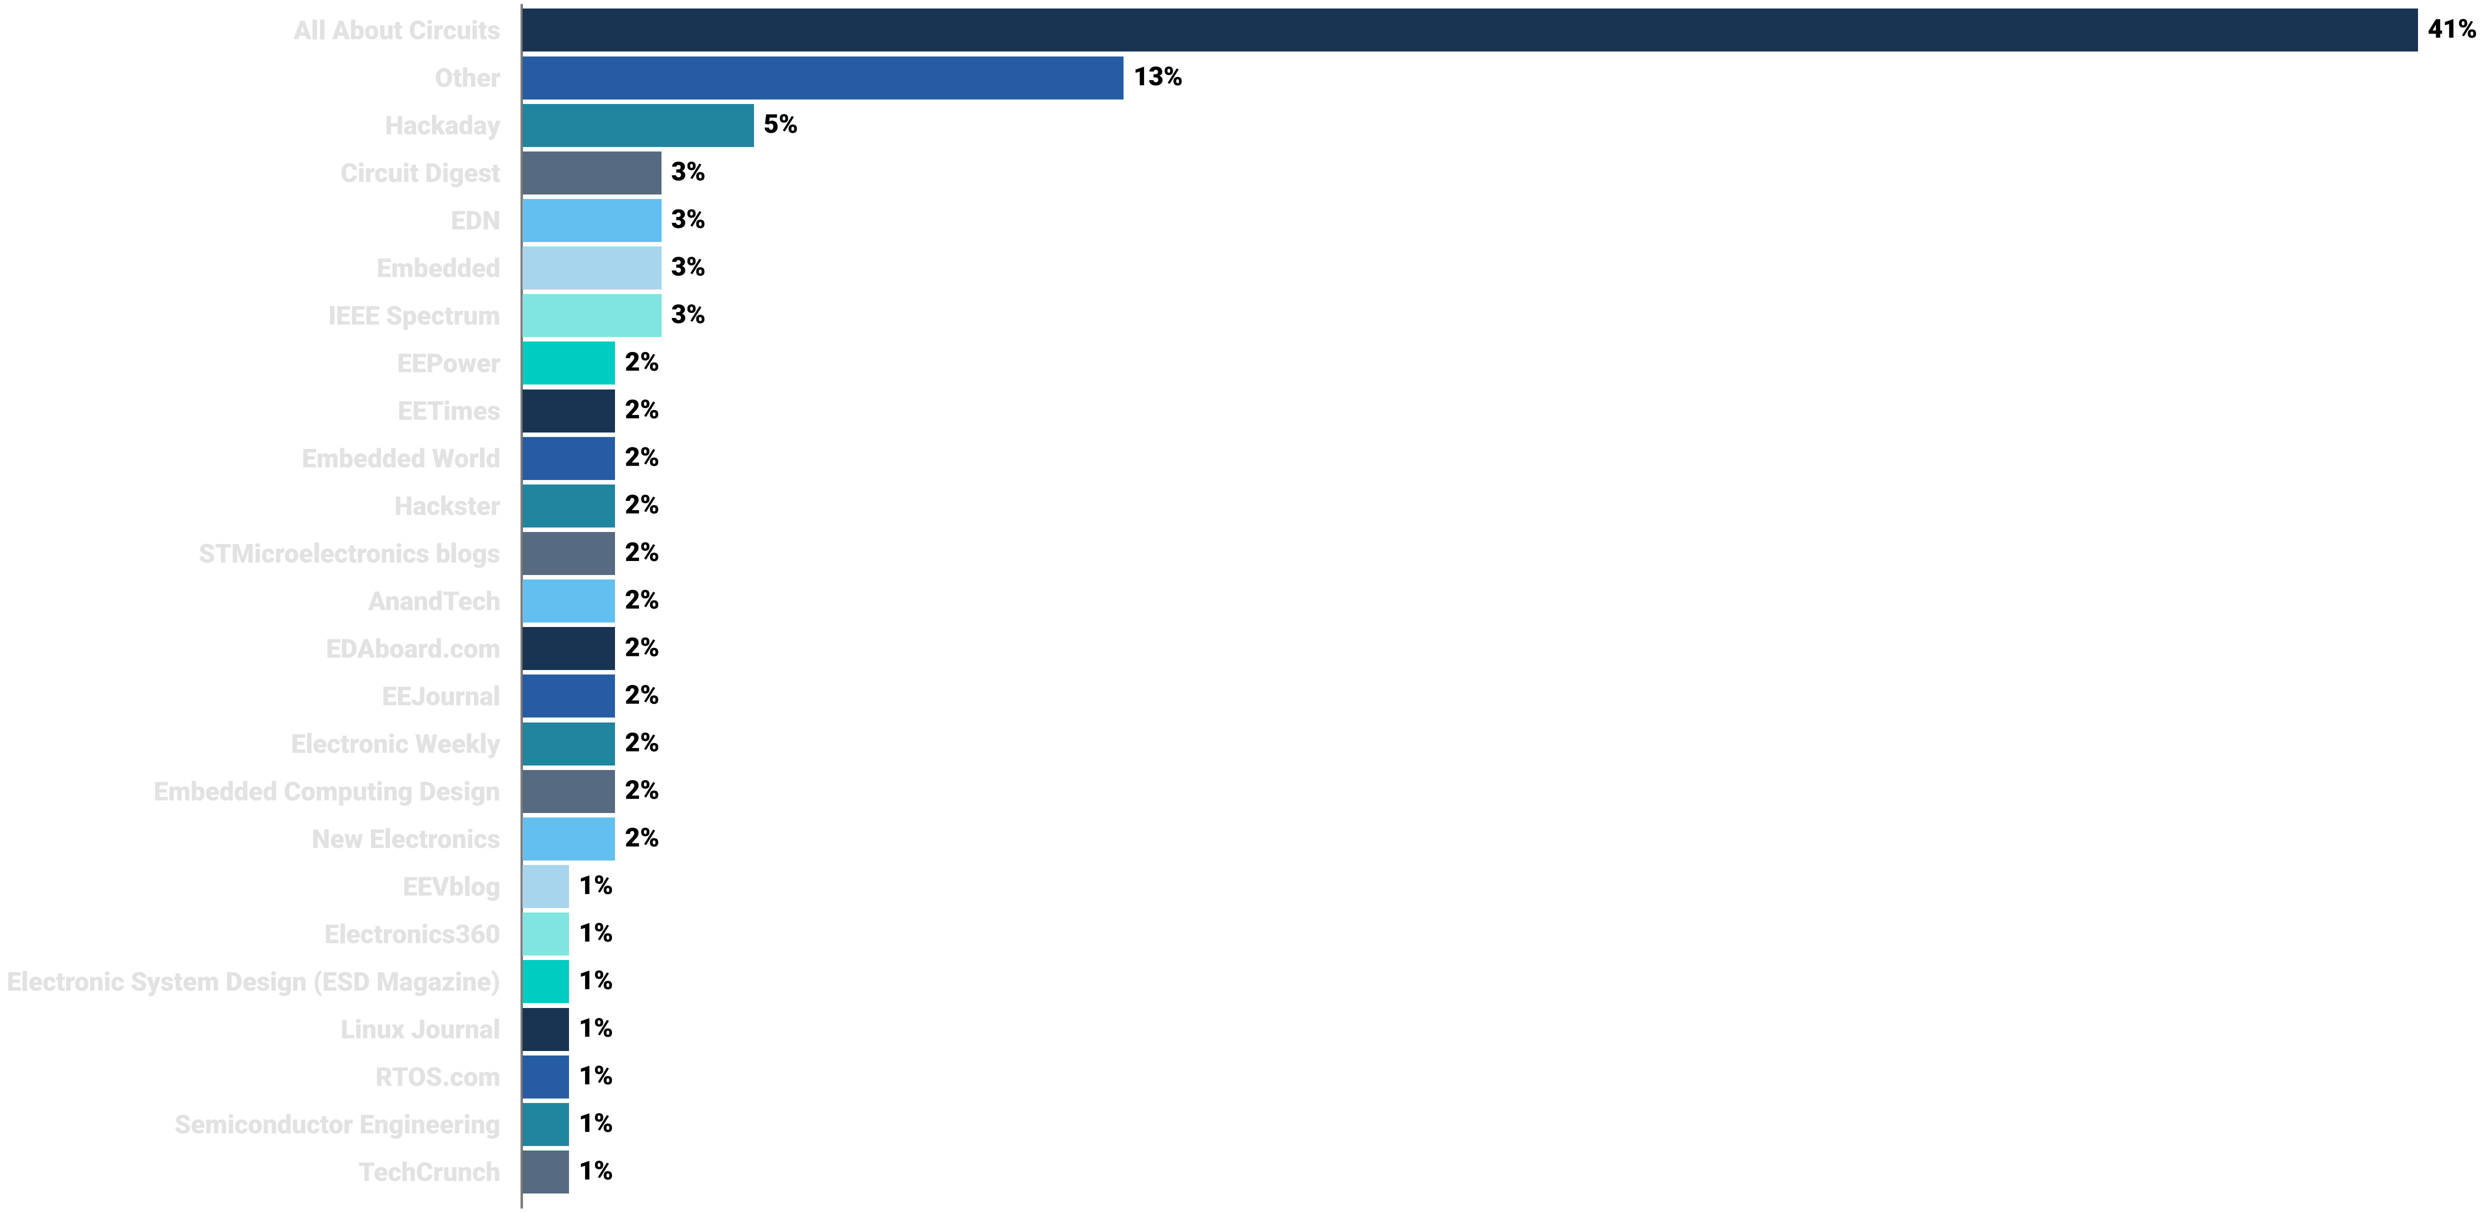

Media

Which ONE resource do you consider the most influential OVERALL within the embedded systems market?