

Europe's new AI edge? The emerging application layer opportunity

September 2026



About this report

This report is a collaborative publication by McKinsey and Boardwave and is based on qualitative and quantitative research conducted from June 1 through August 31, 2026.

Table of contents

About the authors.....	1
Sounding board	2
Executive summary	3
About the research	4
Europe's AI position: The application layer shift.....	6
How sustainable is Europe's AI application momentum?	10
Europe's AI application drivers	12
The key strategic choices of Europe's AI application leaders	17
Getting inside legal workflows: How Legora achieves cross-border scale	21
The core operating practices of Europe's AI application leaders	23
Prove early, move fast: How Parloa scales adoption.....	28
Stakeholder takeaways: Seizing the AI application opportunity	29
Conclusion	33
Endnotes.....	34

About the authors



Julien Gagnon
Partner
McKinsey & Company



Nick Padgett
Partner
McKinsey & Company



Ruben Schaubroeck
Senior Partner
McKinsey & Company



Amy Tong
Associate Partner
McKinsey & Company



Karyna Iusupova
Consultant
McKinsey & Company



Kath Easthope
Cofounder, CEO
Boardwave



Phill Robinson
Cofounder, Chair
Boardwave

The authors wish to thank the more than 30 European tech founders, CEOs, investors, and ecosystem leaders who shared their insights, as well as members of the sounding board, acknowledged on the next page.

The authors also wish to thank Chris Llewellyn, Fritz Björklund, Jeremy Schneider, Katy Wakefield, Kim Baroudy, Lieven Van Der Veken, Martin Wrulich, Naveen Sastry, Nicolas Maechler, Siyeona Chang, Thomas Schumacher, and Tunde Olanrewaju for their contributions to this report.

This report was edited by Daniel Eisenberg, an executive editor in McKinsey's New York office.

Sounding board



Matthew Brockman

Matthew Brockman is chief investment officer and chair of Hg's Investment Committee, playing a central role in the firm's investment activities and strategy, with a particular focus on driving Hg's approach to AI, as a key pillar of future growth as this technology continues to transform the software and services industries. Matthew joined Hg in 2010 to lead the creation of the Mercury Fund, before serving as managing partner from 2017 to 2025. During this time, he helped guide Hg's evolution into a transatlantic platform and one of the world's largest technology investors. Prior to Hg, Matthew was a partner at Apax Partners.



Franck Cohen

Franck Cohen is an experienced technology executive and former president of SAP Digital Core & Industry Solutions, where he led 15,700 engineers and championed machine learning and AI integration. He joined SAP in 2009 as senior vice president and COO for Europe, Middle East, and Africa (EMEA), later advancing to president of EMEA and global chief customer officer. He began his career in 1987 by founding a French enterprise resource planning (ERP) company and later held senior executive roles at Intentia and Lawson Software. With a background in mathematics and electronic engineering, he currently advises HighRadius and C3 AI, and serves as chairman for CYE, GoCardless, and Nomadia.



Martina King

Martina King is the board chair at Luminance. She previously served as CEO and executive chair of Featurespace, leading the British AI company from its start-up phase in 2012 to its acquisition by Visa in 2024. Her extensive executive career includes roles as managing director of Aurasma, managing director of Yahoo Europe, managing director of London Stations at Capital Radio, and commercial leadership at Guardian Media Group. She has also held prominent non-executive and nonprofit board roles at Capita, Debenhams, and the Prince's Trust, and as president of the London Chamber of Commerce. Her accolades include being named Tech Business Woman 2024 at the UK Tech Awards and receiving the 2023 Inspiration Award from Women in Payments EMEA.



Jakob Mökander

Jakob Mökander is vice president, policy—Europe and UK at the General Catalyst Institute, where he works with founders to scale transformative companies and engages governments on building resilient innovation ecosystems. Prior to General Catalyst, Jakob served as director of science and tech policy at the Tony Blair Institute. He has also worked at McKinsey & Company, advising Fortune 500 firms on AI adoption, and at the Swedish Trade and Invest Council, supporting companies expanding internationally. Originally trained as an engineer, Jakob holds a PhD from the University of Oxford focused on AI governance. He has been a visiting scholar at Princeton University and remains affiliated with Yale's Digital Ethics Center.



Fredrik Nylander

Fredrik Nylander is chairman of the board at Quinyx, Truesec, and SITS Group. He also serves as a non-executive director at Paragin Group and the Swedish Chamber of Commerce UK, as well as a board adviser to both public and private companies. He is a seasoned enterprise software executive with over 25 years of regional and global leadership experience spanning cloud technology, human capital management (HCM), fintech, and software as a service. He previously spent several years at Workday, where he served as senior vice president of service and senior vice president of north and emerging markets. Prior to Workday, he held executive roles at Microsoft, SAP, and Oracle.

Executive summary

Europe's technology ecosystem is facing a new horizon of AI maturity. As AI adoption grows, the era of one-size-fits-all software is beginning to shift, suggesting that the greatest economic opportunities may lie in transforming complex, everyday business processes, something European companies have been particularly adept at in recent decades. Because some of the most potentially valuable enterprise AI applications demand deep integration into highly regulated and heterogeneous environments, those same European companies may be well positioned for this coming era.

This report, produced jointly by McKinsey and Boardwave,¹ examines where Europe could thrive in the next phase of AI value creation and what that means for European AI-native start-ups, software incumbents, and investors. It is based on McKinsey proprietary funding analysis of data from Dealroom, PitchBook, and Gain, as well as interviews with more than 30 European technology founders, CEOs, investors, and ecosystem leaders (see sidebar “About the research”).

The research shows that this transition is starting to be felt in the funding landscape. Europe does still lag other regions in overall venture funding; North America attracted 71 percent of the €835 billion raised between January 2025 and June 2026, compared to Europe's 10 percent and Asia's 16 percent. But while Europe also trails the US in total AI venture funding, its pool of start-ups focused on the application layer approached 44 percent of comparable US funding levels in the first half of 2026, up from just 11 percent in 2018. Within this layer, there is notable growth in vertical or industry-specific AI: European start-ups developing vertical applications now attract 67 percent of the funding raised by their US peers, up from 9 percent in 2018. This momentum is further supported by several structural factors—including a reversal in AI talent flows, more proactive commercialization from universities, strategic co-investment from corporate incumbents, and an increase in patient sovereign capital—that collectively reinforce the region's position.

Amid these shifting conditions, an emerging new cohort of European AI champions may turn historical market complexities into the foundation of renewed global competitiveness. To succeed, founders are making three deliberate strategic choices, all built on the kind of specificity that is increasingly key to scaling. They embed products deeply into enterprise systems of record and complex workflows, specialize in Europe's historically strong industrial ecosystems, and take a global, multimarket approach from the outset.

Setting the right strategic direction, however, isn't sufficient on its own. Our research indicates that Europe's AI start-ups navigating a capital-constrained environment and value-conscious buyers have been forced to execute with a high degree of efficiency across their business, commercial, and operating models. High-performing companies rely on five operating practices to achieve sustainable scale: accelerating R&D velocity with AI, utilizing an adoption-centric go-to-market approach, implementing a value-based commercial model, leveraging AI-enabled internal operations, and deploying an AI-focused talent strategy that includes workforce upskilling. While the practices in and of themselves aren't unique to European companies, certain aspects of the region's business environment make them especially critical for success.

To capture the value of this new horizon, stakeholders across the ecosystem are adapting their approaches. Founders are targeting high-value workflows and leveraging sector ecosystems to scale globally from day one. Established software leaders are drawing on their deep customer understanding to rewire legacy processes and partner with AI natives. Simultaneously, leading investors are adapting their capital deployment models to support the forward-deployed implementation and cross-border scaling velocities required to build sustainable European champions.

About the research

This report combines quantitative venture capital analysis with qualitative expert interviews—a mixed-methods approach designed to show not only where markets are placing bets but also why. Our objective was to identify segments of the AI market where Europe is demonstrating emerging promise.

The funding analysis draws on a McKinsey proprietary database aggregating data from Dealroom, PitchBook, Gain, and McKinsey private sources, covering venture capital investments from 2018 through the first half of 2026 in Europe and the United States. Europe comprises companies headquartered across the European continent, including the United Kingdom and non-EU markets such as Switzerland, Norway, and Ukraine (consistent with standard venture databases such as PitchBook).

Each company was classified using a consistent taxonomy. We analyzed the market as follows:

- **Software stack versus non-software.** Software stack companies are those whose primary deliverable is software or software infrastructure that customers integrate and operate to run their own workflows. Non-software companies are those whose core offering is a service, marketplace, or physical product, even if software enabled.
- **Elements of the software stack.** We grouped software stack companies into three layers. The first, infrastructure, covers compute, chips, data center capacity, networking, and cloud primitives—the substrate on which models train and run. Platform covers foundation models, model-serving APIs, developer tooling, and horizontal AI platforms—the general-purpose intelligence layer. And application, the third layer, covers software that packages technology into end user or enterprise workflows.
- **Within software, AI versus non-AI.** AI companies are those that have material AI and machine learning (AI/ML) in the product today or that position AI/ML in primary product marketing. Non-AI companies are those where the product contains limited AI/ML and makes no such claims in positioning.
- **Within AI applications, horizontal versus vertical.** Horizontal applications are used across industries (for example, coding assistants or customer service agents). Vertical applications are built for a specific sector as target customer.

- **Vertical industries.** Vertical application companies were grouped into 13 industry categories, including agriculture and food systems; defense and security; education and research; financial services and insurance; healthcare and life sciences; HR, recruiting, and workforce; industrial, energy, and climate; legal and compliance; media, marketing, and sales; real estate and proptech; retail, commerce, and consumer; transport, mobility, and logistics; and other/unclassified.

Classifications involved judgment at category boundaries. We applied consistent rules throughout, and these calls do not materially change our aggregate findings.

The clearest signals of Europe's emerging strength appeared in AI applications, especially vertical applications. We tested these patterns in more than 30 confidential interviews: 12 with start-up operators (including seven from among Europe's top ten funded start-ups), eight with investors across venture capital and software-focused private equity active in Europe and the United States, and ten with technology leaders (senior executives, board members, and advisers at global software, AI, and hyperscaler companies). These conversations explored the drivers behind the funding data, including Europe's rising share of application layer investment.

Europe's AI position: The application layer shift


As enterprises increasingly move from experimenting with models to redesigning legacy business processes, more of the competitive focus is shifting toward the application layer of the AI tech stack: software that can understand specific workflows, operate inside existing enterprise systems, and automate work in complex environments. Our research suggests this shift could matter disproportionately for Europe's technology competitiveness. The region has long combined deep industrial expertise, regulatory rigor, and heterogeneous markets, factors that could become more valuable as AI evolves from general-purpose intelligence toward specialized deployment.

Early funding patterns indicate that investors are beginning to tap that opportunity. In 2018, European AI application funding stood at just 11 percent of US levels. Eight years later, that ratio has climbed to 44 percent, with absolute European application funding reaching €6.8 billion in the first half of 2026 alone; that's as much as the total of 2024 and nearly matches the €8.1 billion raised across all of 2025 (Exhibit 1).

Exhibit 1

Europe’s AI application start-ups are narrowing the funding gap with the US, while US top model makers attract most platform funding.

European venture capital funding amount by AI tech stack layer,¹ % of US amount




¹Excludes companies whose main offering is a product or a service that is not AI software or software infrastructure, even if software enabled, eg, digital banks, robotics, hardware-heavy defense.
²Platform, intelligence, and infrastructure layers.
³OpenAI, Anthropic, xAI.
Source: Dealroom; Gain; PitchBook; McKinsey analysis

Within vertical AI, this convergence is even sharper—in fact, the sharpest in the entire tech stack. The amount of venture capital funding attracted by European start-ups focused on vertical applications is now 67 percent of the amount raised by their US counterparts, up from 9 percent in 2018 (Exhibit 2). While these absolute figures remain a fraction of total global AI funding, the trend lines reveal a positive structural pattern. This momentum coincides with the rise of a new cohort of European AI leaders, such as Helsing, Legora, Lovable, and Wayve, which average just five years since founding yet already rank among the top 100 European tech companies by valuation.² As one former global enterprise software executive observes, “When it comes to vertical software, complex environments, and deep, specific solutions, Europe can build leaders that grow and compete globally.”

Exhibit 2
Europe’s biggest AI tech stack funding surge has been for vertical AI application start-ups, with levels reaching two-thirds of US peers.

European venture capital funding amount for AI applications by application purpose,¹ % of US amount by application purpose²



¹Vertical applications serve a specific industry need, while horizontal applications can be used across multiple industries.
²Compares Europe's vertical AI application funding to a US baseline of vertical AI application funding, and Europe's horizontal AI application funding to a US baseline of horizontal AI application funding.
Source: Dealroom; Gain; PitchBook; McKinsey analysis


McKinsey & Company

This momentum is highly targeted: Around 94 percent of this European vertical AI investment is concentrated in six sectors: industrial, energy, and climate; defense and security; transport, mobility, and logistics; legal and compliance; healthcare and life sciences; and financial services and insurance (Exhibit 3). In Europe, these sectors tend to be defined by deeply entrenched incumbents, affiliated research institutions and ecosystems, and regulatory rigor. By navigating these industry complexities, European AI start-ups can deliver the specialized, compliance-ready solutions that enterprise buyers require—solutions that may be less vulnerable to competition from larger, general-purpose wrappers.

Exhibit 3

Six complex and regulated sectors attract 94 percent of European vertical AI application start-up funding.

European vertical AI application funding, 2025–H1 2026



Note: Figures may not sum to 100%, because of rounding.

¹Total European vertical AI application funding of €9.1 billion.

Source: Dealroom; Gain; PitchBook; McKinsey analysis, rounds >€10 million

McKinsey & Company

Despite these positive developments, Europe still lags significantly in overall start-up capital. The region captured just 10 percent of the €835 billion in total global venture funding raised between January 2025 and June 2026, compared to 71 percent for North America and 16 percent for Asia. Over that same time period, European start-ups were recipients of 14 percent of all US tech venture capital funding (€82 billion versus €595 billion).³ In the heavily concentrated, foundational layer of platforms and infrastructure, European funding sits at just 4 percent of US levels as capital consolidates around leading US frontier labs; the top three players in this space—Anthropic,

OpenAI, and xAI—raised €229 billion in the first half of 2026, receiving 85 percent of US platform and infrastructure funding.⁴

Some industry observers believe that large language models (LLMs) could eventually obviate the need for a separate application layer altogether, though how that plays out remains very uncertain. Despite the funding gap, the region has not given up on its own development of foundational models. At the same time, there are compelling reasons, including some historical precedent, to support the view that the region's focus on applications may align with how the next phase of the AI cycle may ultimately unfold (see sidebar “How sustainable is Europe's AI application momentum?”).

How sustainable is Europe's AI application momentum?

Given the heavy concentration of AI investment in foundational models, it's understandable that many would assume today's LLMs will eventually take over other parts of the tech stack, including the application layer, where European start-ups have recently had some success. After all, US and European companies in the AI platform and infrastructure layer, comprised largely of the top model makers, have a substantial capital advantage, having raised six times as much total funding as AI applications (€420 billion versus €67 billion between 2025 and the first half of 2026). They have also already shown interest in moving up the stack: Anthropic launched Claude Design and Claude Legal earlier this year, and xAI acquired Cursor in August 2026. They also have massive reach to tap, and expanding tech products into new markets by bundling has a proven track record. OpenAI's ChatGPT has one billion weekly active users as of August 2026,¹ and Anthropic is the most widely adopted AI model with US businesses, commanding an estimated nearly 40 percent of the market.²

Still, there are valid reasons why AI applications could remain a distinct stronghold in the future. Applications, for starters, require a different talent profile than platforms and infrastructure, which tend to focus on compute, research, and training. By contrast, large teams of deployment engineers, knowledge experts, and other specialists need to be embedded with customers to provide the domain specificity and workflow depth that successful enterprise applications depend on.

In addition, if the AI cycle follows previous technological revolutions, Europe's momentum in the application layer could prove to be very valuable and long-lasting. Historical patterns suggest that major technological shifts follow a distinct trajectory: An initial, capital-intensive installation phase focused on foundational infrastructure eventually gives way to a broader deployment phase where value creation shifts to applications deeply embedded across the economy.³ During the internet/cloud cycle (1995–2020), application companies captured roughly 70 percent of the ecosystem, whereas during mobile (2007–2024), applications accounted for roughly 85 percent.⁴

¹“A milestone in expanding access to AI,” OpenAI, August 31, 2026.

²“2025: The state of generative AI in the enterprise,” Menlo Ventures, December 9, 2025.

³Promise and Pressure, JPMorgan Chase & Co, November 2025.

⁴The return of the Consumer Gold Rush? What Chi-Hua Chien is betting on next, Strictly VC Download, June 16, 2026.

However the application business plays out, Europe isn't planning to concede the model layer. Europe's push for technological sovereignty, including the European Commission's June 2026 "tech sovereignty package"⁵ and a €10 billion initiative to establish seven AI gigafactories,⁶ is actively working to narrow the infrastructure gap, even as it also fuels the application momentum. Europe is already demonstrating its capacity to compete in this stack, driven by the success of sovereign-focused players like Mistral, Nebius, and a widening base of sovereign compute programs. These sovereign programs, if anchoring demand while keeping provider eligibility open across the region, could help scale both layers.

Furthermore, the current LLM paradigm may not be the last. A widening body of research points to a next frontier extending beyond text-based reasoning into world models, physical AI, scientific foundation and other domain-specific, "small" models. Early examples are already visible: Yann LeCun's world model program at Advanced Machine Intelligence, Google DeepMind's Genie series, Wayve's driving world models, Isomorphic Labs' biological foundation models, and PhysicsX's industrial physics AI. While the long-term trajectory remains unclear, this emerging model evolution indicates that Europe could still credibly compete for a share of the next platform era.

⁵Strengthening Europe's tech sovereignty, European Commission Directorate-General for Communication, June 3, 2026.

⁶EU launches AI Gigafactories call to boost Europe's computing capacity and unlock more than €30 billion in investment, European Commission, July 30, 2026.

When it comes to vertical software, complex environments, and deep, specific solutions, Europe can build leaders that grow and compete globally.

—Global enterprise software executive

Europe's AI application drivers

Last year's McKinsey and Boardwave report, *Europe's moonshot moment*, highlighted both the foundational strengths of European tech and the specific challenges hindering its ability to scale.⁵ While the region boasts deep talent and research assets, obstacles such as fragmented markets, institutional constraints, and a lack of late-stage funding have historically limited growth. Several distinct tailwinds are now accelerating the ecosystem, reinforcing Europe's emerging strength in building and scaling workflow-specific AI within complex, heavily regulated sectors. By addressing historical gaps in talent, commercialization, and capital, these shifts support the enterprise-grade, industry-embedded model that offers Europe its best current AI opportunity.

Reversal in transatlantic talent flows

Historically, Europe faced talent attrition, with some promising homegrown entrepreneurs eventually relocating their operations or headquarters to the US. As a deep-tech venture capitalist observed, “Before, there was always that fear that we would be losing our best minds eventually to the US.” Today, that trend is showing signs of reversing. In 2025, there was a net inflow of AI professionals from the United States to Europe for the first time in years.⁶

Germany now ranks third globally for highly skilled AI talent, behind the United States and India, while Switzerland leads in per capita metrics.⁷ Simultaneously, the United Kingdom maintains its position as a global talent magnet, driven by a sprawling, highly concentrated ecosystem anchored in institutions such as Google DeepMind. This concentration supports powerful cluster effects, offering proximity to elite research groups and industry partners.

Furthermore, the historic transatlantic pay gap is narrowing. Leading European AI labs now offer top-tier research and engineering packages in the €300,000 to €500,000 range, putting them closer to US peers.⁸

Universities as venture builders

Sectors like advanced manufacturing and life sciences require deep foundational research, making academic spinouts essential to build a solid position in vertical applications. Europe is already a leader by some measures of research quality, capturing 19.5 percent of global AI citations.⁹ This research strength is translating into substantial investor backing and highly valued spinouts, such as Neural Concept, Synthesia, and Wayve.¹⁰ University spinouts now claim 12 percent of all VC capital invested in Europe—totaling €3.1 billion in deal value year-to-date—and Europe boasts roughly 14,000 active private university spinouts, with 226 launched between 2025 and 2026.¹¹

To sustain this momentum, institutions are removing long-standing barriers to scale by reducing equity demands. Although this trend predates the AI boom, recent frameworks, such as the 2023 TenU USIT, have accelerated the shift to meet founder-friendly global benchmarks.¹² Consequently, average UK university stakes, which peaked at 28 percent in 2017, fell to 16.1 percent in 2024.¹³ Leading institutions highlight this new approach: While Oxford maintains a 10 to 20 percent stake, Cambridge has lowered it to just 8 to 9 percent, and ETH Zurich has dropped its baseline equity stake to just 2 percent.¹⁴

In parallel, tech transfer offices are shifting from passive intellectual property (IP) administrators into proactive venture builders. Hubs like TUM's UnternehmerTUM, Oxford Science Enterprises, and Imperial's DT Prime now combine incubation with the corporate networks necessary to scale spinouts in regulated verticals. As a senior executive at a global venture capital firm noted, “We have incredible universities in Europe, like Oxford and Cambridge, who are steering research in the right directions. It's a structural advantage, and a large portion of the global talent coming from those places is now being channeled into highly commercialized, scalable innovation ecosystems.”

We have incredible universities in Europe, who are steering research in the right directions. It's a structural advantage.

—Senior executive at a global venture capital firm

Corporates as strategic co-investors and partners

Building industry-embedded AI solutions requires access to complex workflows and proprietary data. While European corporates have historically not been as active venture investors or partners as their US peers,¹⁵ more European corporates are stepping up to help start-ups gain entree to those two critical assets alongside powerful distribution channels that enable enterprise-grade deployments. European AI application venture capital deal value featuring at least one corporate investor grew from 14 percent of US levels in 2018 to 39 percent in the first half of 2026.

The EU AI Champions Initiative, which launched last year, illustrates this shift. The initiative, an ambitious, €200 billion public–private joint effort to help accelerate the region's innovation and competitiveness, moved beyond broad capital commitments to launch 18 specific corporate scale-up partnerships, such as Allianz with Parloa and Deutsche Telekom with PhysicsX, unlocking €1 billion in co-investments.¹⁶

The region's industrial incumbents also are accelerating adoption by acting as investors, customers, and procurement partners in European verticals. Recent industry surveys indicate that successful European AI deployments now move from strategy to shipped feature in just four to six months, down from 12 to 18 months previously.¹⁷

Patient capital from sovereign and state-backed funds

Scaling applications in specialized, heavily regulated verticals like advanced manufacturing, defense, and healthcare typically takes longer and often involves larger capital requirements, and state-backed investors are stepping in to help to fill that funding gap. In the first half of 2026, such institutions participated in 24 percent of the region's AI application venture capital deal value, backing scale-ups such as PhysicsX and Wayve.¹⁸ This rate was up from a recent high of 14 percent in 2022 but only single digits each of the last two years. The uptick was also felt in overall European AI funding: Sovereign funds played a role in 36 percent of deal value in the first half of this year.

Coordinated initiatives—including the European Commission's €200 billion InvestAI initiative, which blends public and private capital to fund regional AI leaders, as well as national funds like France's Bpifrance, the UK's Sovereign AI Unit, and Germany's KfW Capital—are driving this shift. Crucially, these efforts operate alongside multilateral initiatives like the €1 billion NATO Innovation Fund and the NATO DIANA accelerator network, which provide DARPA-style¹⁹ backing for dual-use technologies.


US investor funding and late-stage private capital

US investors also are fueling Europe's application layer momentum. The share of European AI application deals involving at least one US-based investor rose from 49 percent in 2018 to 84 percent in 2026. Much of this activity concentrates in Europe's vertical AI layer, where domain strength, compliance-native architecture, and global scaling offer attractive growth opportunities. As the founder of a Spanish AI start-up noted, this capital allows European companies to “compete head-to-head against the US.”

In parallel, private markets are driving an increase in late-stage capital. Historically, Europe has performed well in seed and early-stage funding but suffered from relative weakness in the late-stage capital critical for scaling global champions. However, recent dynamics have suggested that may be starting to change. While the United States still leads in late-stage (and overall) AI funding value, European late-stage venture capital funding in AI applications has accelerated rapidly from 6 percent of US levels in 2018 to 53 percent in 2026 (Exhibit 4). This momentum in the application layer has grown even as broader European software (including AI) funding has continued to experience challenges in recent years.

Exhibit 4

Europe's momentum in AI application funding has been particularly strong in late-stage rounds, long an area of relative weakness.



Source: Dealroom; Gain; PitchBook; McKinsey analysis

McKinsey & Company

Furthermore, Europe's average late-stage round size in AI applications overtook that of the United States in 2022, climbing to nearly €200 million in the first half of 2026 (versus €167 million for the United States). This application surge extends across all funding stages, where average European rounds doubled year over year to €68 million in the first half of 2026, outpacing (for the first time) the US average of €60 million.²⁰ As an executive at a US venture capital firm noted, "Right now, if I look at the AI application ecosystem in Europe, we don't feel that there is a lack of growth capital."

While a short-period spike does not guarantee a durable trend, the steady, multiyear expansion suggests that the region is building the necessary capital foundation to support sustainable scale-ups. The fact that some European scale-ups (such as Legora and Lovable) have recently been able to generate at least €100 million in annual recurring revenue (ARR) in less than two years is also an encouraging signal for the region; on average, European software start-ups have taken 15 years to reach that milestone, five years longer than their US peers.²¹

However, despite the positive indicators, these structural tailwinds alone do not automatically produce success. Founders must actively convert them into deliberate strategic choices.

The key strategic choices of Europe's AI application leaders

Europe brings valuable structural assets to the application layer, such as domain expertise and dense sector ecosystems, but successful founders are using three deliberate strategic choices to translate these regional strengths into market growth: embedding deeply into enterprise workflows, specializing in Europe's strongest verticals, and taking a global, multimarket approach from the outset.

Deep workflow integration: The shift from standardized software to tailored automation

The SaaS era historically rewarded companies that scaled off-the-shelf solutions as widely as possible. In contrast, the AI era rewards deep workflow integration. While incumbent software giants have always provided bespoke configurations, traditional off-the-shelf software can struggle to accommodate the strict compliance, domain expertise, and workflow depth many AI applications require, as customers increasingly expect software not just to support tasks, but to automate and execute them.

The next wave of enterprise automation targets tasks that are intrinsically complex and highly customized, where a generic, one-size-fits-all approach may not be up to the task. By contrast, European companies, having evolved in heterogeneous markets with divergent B2B practices and complex localized needs, are used to building adaptable, deeply tailored software. As the founder of a Southern European AI start-up put it, “We’re more ready for this unglamorous, messier work that AI requires—where you need to be embedded in complex, slow-moving businesses.”

To successfully automate these bespoke tasks, such as resolving a nuanced billing dispute, updating a service case, or executing multistep fulfillment, leading companies recognize that AI cannot simply be laid on top of an existing process. Instead, they are embedding these solutions deeply into the enterprise’s system of record, where the technology can understand live operational context, executing permitted actions and writing outcomes directly back into the workflow. Crucially, Europe’s comprehensive regulatory environment, which many public and private sector participants believe can hamper innovation, could also help foster this level of integration. By architecting AI products that sit deep inside the enterprise’s existing perimeter, founders can simultaneously solve for compliance. Data, audit trails, and governance (mandated by frameworks like GDPR, the EU AI Act, and DORA) remain secure within the incumbent stack that procurement teams already trust. By making this deliberate commitment to deep, tailored integration, European start-ups trade fast out-of-the-box deployments for potentially valuable enterprise moats that include higher switching costs, faster procurement in regulated industries, and products that genuinely change how work gets done.

Parloa, the Berlin-founded contact center AI platform (with €427 million raised in 2025–26), illustrates the durable value of integrating deep into the core of customer operations. The platform embeds voice and chat agents directly into customer service workflows through integration with contact-center-as-a-service (CCaaS), customer relationship management (CRM), and ERP applications, reading live customer and case data and writing outcomes back into the system of record.

Sector specialization: The rise of vertical AI in Europe's strongest ecosystems

As mentioned earlier, European vertical AI funding is concentrated in six sectors: industrial, energy, and climate; defense and security; legal and compliance; transport, mobility, and logistics; healthcare and life sciences; and financial services and insurance. In these sectors, European start-ups are converting long-standing structural assets into commercial momentum by building within emerging ecosystems (clusters of universities, national labs, incumbents, talent, and capital). As a policy executive at a global venture capital firm noted, "In some sectors, we are much better positioned to win and outcompete. If you take industrial data for robotics or healthcare, we possess longitudinal, rich data sets that few other countries have."

Last year's *Europe's moonshot moment* report argued that Europe's tech-scaling constraint was not a lack of raw ingredients (talent, institutions, and capital) but the difficulty of setting a coordinated innovation flywheel in motion. Yet the foundation remains promising: Europe is home to 42 of the world's top 100 emerging ecosystems, nearly half the global total, bringing together founders, investors, universities, and technical talent in clusters.²² For AI applications, there are early signs that these ingredients are starting to produce such a flywheel in the strongest verticals: Institutions supply frontier research; specialist talent translates it into products; industry giants provide proprietary data, workflows, customers, and credibility to help bring it to market; and capital and co-innovation follow as corporates and institutions become investors and strategic partners.

In some sectors, we are much better positioned to win and outcompete. If you take industrial data for robotics or healthcare, we possess longitudinal, rich data sets that few other countries have.

—Policy executive at a global venture capital firm

Advanced industries (including manufacturing, energy, and climate) offer one of the clearest illustrations. Research institutions like EPFL, ETH Zurich, TUM's UnternehmerTUM, and Fraunhofer Institute supply a spin-out pipeline in simulation, materials, and manufacturing. Talent pools combine physics PhDs, computational engineers, and computer-aided design and engineering (CAD/CAE) practitioners with operators from the sector. Global industry giants such as Siemens and automotive OEMs shape demand and increasingly act as customers, co-developers, and investors. "The historic enterprises that we have in Europe are really driving adoption," said one senior leader at a European engineering AI company. "They have their own venture arms and own R&D centers; that's a very large competitive advantage."

London-based PhysicsX (with €400 million raised in 2025–26) shows the flywheel in motion. Its platform embeds AI across simulation, optimization, and manufacturing, with founders bringing F1, automotive engineering, and AI pedigrees from Renault, Audi, and Mercedes-Benz. Siemens, which has been at the forefront of industrial foundational models, was a strategic investor in its €120 million Series B and co-develops with the start-up on an AI-powered extension to its own Simcenter X product. By integrating directly into the engineering environments of partners such as Siemens and General Motors, PhysicsX accesses data that horizontal vendors are hard-pressed to replicate; each deployment and more in-production use generates more training signals and continuously improves model performance.

Multimarket blueprint: The push for global scale built on specificity from the start

AI presents a distinct opportunity to lower the fragmentation obstacles that have inhibited Europe's tech growth, a potential shift consistently highlighted by founders and investors interviewed for this report. For decades, Europe's divergent languages, tax codes, legal systems, and buying behaviors imposed a strategic dilemma. As a senior European tech executive observed, this reality forced European companies to navigate "more complexity from day one in their design process" than US companies face in a larger domestic market.

European AI-native architectures could help overcome these historical scaling challenges. This is particularly the case with applications, which have historically been the most constrained by the different interface and business rule requirements of market localization. As one European tech executive and investor remarked, "Previously, localizing a product meant heavy costs for translating software, rewriting educational materials, and hiring local language support. Today, translation and multilingual support cost almost nothing, allowing European start-ups to think globally immediately." However, these new capabilities go far beyond just overcoming language barriers. AI can now more easily ingest and adapt to divergent regulatory frameworks, local compliance standards, and bespoke business rules

that previously made multimarket deployment costly. What once took lengthy market-by-market engineering can now be built into the product from launch, compressing the path to cross-border scale.

Synthesia, a London-based AI video generation platform founded by Danish, German, and Spanish engineers, that raised €357 million between January 2025 and June 2026, illustrates how a global mindset can accelerate AI scaling. Founded in 2017, Synthesia spent three years building computer vision to re-animate faces for multilingual dubbing, but it was only able to scale quickly when generative models matured. Language coverage compounded: it went from roughly 60 languages in 2021 to 140+ by May 2025,²³ and 160+ by 2026, alongside a voice library that grew from 400+ to over 1,000.²⁴ The commercial result is a genuinely distributed business: around half of revenue from the United States and half from Europe and Asia, with the DACH region (Germany, Austria, and Switzerland) being its third-largest market by ARR globally.²⁵ Another example of this global approach is Legora, a legal AI assistant that accelerates research, contract review, and legal document drafting. Founded in Stockholm in 2023 and having raised €755 million, the corporate law AI platform went beyond the Swedish domestic market from day one. (see sidebar, “Getting inside legal workflows: How Legora achieves cross-border scale”).

Setting the right strategic direction is only part of the equation. In the next section, we examine the specific operating practices high-performing European AI application start-ups use to execute their vision and scale effectively.

Getting inside legal workflows: How Legora achieves cross-border scale

Legora illustrates how a European AI application company can scale when founders make deliberate decisions about where to compete. Stockholm-founded and backed by €755 million, the company focused on complex legal workflows, targeted the legal sector at a moment of openness to collaboration, and planned for global markets from the start. Those three choices it made, explored below, were instrumental in building a business that now reaches more than 1,800 enterprise customers in some 50 countries.

Developing highly customized solutions for even the most niche legal workflows. Instead of creating a generic legal overlay, Legora embeds directly into how lawyers work. To get there, the team worked alongside Mannheimer Swartling, the largest law firm in the Nordics, for nine months, co-creating the product with the practitioners who would use it. “The only way to build product is to know what your customers actually say that they want, but maybe also actually should say that they want”, said Vilgot Kledzik, Legora’s Vice President of Operations. To scale that depth across traditionally risk-averse enterprises, Legora deploys “legal engineers,” AI-native lawyers who help Legora’s customers redesign daily workflows with agentic AI, including those with only a relatively small user base, such as certain M&A workstreams.

Building inside a sector where customers are open to rapid innovation. The legal sector is one of Europe's strongest verticals given the region's history of regulatory rigor. At the moment, leading law firms are willing to collaborate on AI to address long-standing productivity bottlenecks. Legora leaned into that window, treating incumbent law firms that were willing to invest time with a vendor as active development partners.

Competing globally from day one and investing accordingly. Most legal software has tended to scale slowly, market by market. Because legal frameworks, sources of law, and court procedures vary from jurisdiction to jurisdiction, legal SaaS had typically anchored around a single common-law jurisdiction and grown across borders through acquisition, often leaving lawyers to bridge jurisdictions manually from a searchable database. Legora charted a different path: invest in proprietary international legal datasets, structure fragmented legal data across geographies from the start, and open local offices to serve firms where they operate. That significant upfront investment in architecture allowed the product to serve multiple markets at launch.

The core operating practices of Europe's AI application leaders

Capitalizing on workflow depth in complex industries with AI applications requires an efficient operating model that manages complexity without inflating customer costs. Successful European AI start-ups achieve this through five core practices: R&D velocity, an adoption-centric go-to-market approach, a value-based commercial model, AI-enabled internal operations, and an AI-focused talent strategy.

While these fundamentals apply globally, distinct European dynamics raise the stakes. First, European enterprises spend roughly half as much on software per employee as their US peers (\$167,000 versus \$349,000 annually).²⁶ This lower-budget environment and the region's generally risk-averse culture heavily favor start-ups that use an adoption-centric go-to-market approach and value-based commercial models. Second, despite making some progress of late, European start-ups still face more limited late-stage growth capital, as well as high labor costs in major hubs.²⁷ These constraints demand particularly high efficiency. As a senior leader at a European engineering AI company put it, the gap shows up in how companies deploy: "In Europe, you're operating under a survival constraint. A US company out there with potentially five times the funding can afford to be less effective and throw forward-deployed engineers at every customer."

R&D velocity

While speed to market has always been important, the rapid pace of AI adoption is widening the gap between leaders and laggards. In software development, top-performing organizations utilizing an AI-powered product development life cycle are already seeing up to a 30 percent faster time to market.²⁸ McKinsey's research suggests that as companies mature into a third horizon of AI transformation, moving beyond basic tool adoption to fundamentally redesign their operating models, the impact on engineering productivity could scale even higher.²⁹ Across industries, adopting end-to-end digital engineering can accelerate the entire R&D life cycle by 20 to 80 percent, with the potential to double the innovation rate in some IP-heavy sectors.³⁰

For European AI companies operating with leaner budgets and higher labor costs than US peers, superior R&D velocity serves as a critical structural response to the growth capital gap. AI-enabled simulation and automated experimentation can deliver more iteration per euro invested, substituting capital-efficient engineering for head count growth, especially since R&D loops are costly and lengthy in many of the complex European industries that hold the most application layer potential.

Lovable, a Stockholm-based software creation platform, has redesigned its engineering workflow around AI agents to increase the speed at which it can develop and ship its own product. Engineers can now orchestrate six or seven coding agents in parallel, while an AI-based risk system triages their output: routine changes can pass through automated review, while sensitive changes involving areas such as infrastructure or authentication require human review. One Lovable engineer reported that this model increased his weekly output from 20 to 30 accepted code changes in January 2026 to more than 150 by June.³¹ The broader shift is from using AI simply to write code faster to redesigning the development loop itself, concentrating scarce engineering attention on higher-value architecture and design decisions.

Just offering a solution and hoping people will change their ways of working to deploy it is very dreamy. You will not capture the gains of AI if you don't redesign your processes alongside the technology.

—Founder of applied AI lab

Adoption-centric go-to-market (GTM)

An adoption-centric GTM shifts the focus from selling licenses to embedding AI inside customer workflows, typically through forward-deployed engineers (FDEs) and technical account teams who can work alongside enterprise customers to prove value and drive usage. For complex enterprise AI, this high-touch approach is increasingly critical. As one applied AI lab founder noted, “Just offering a solution and hoping people will change their ways of working to deploy it is very dreamy. You will not capture the gains of AI if you don't redesign your processes alongside the technology.”

In Europe, this model is a strategic necessity and potentially a strong suit. Risk-averse enterprise buyers demand deep technical validation before scaling beyond a proof of concept. And European vendors, operating with leaner budgets and higher labor costs, cannot scale GTM through traditional sales head count alone. But because European founders have long navigated heterogeneous markets by building highly tailored, deeply integrated solutions, they may already be well disposed to such a high-touch deployment model.

Neural Concept, an EPFL spin-out that uses AI to test and predict design performance, focuses on customer adoption from day one in the sales process. Based in Lausanne and having raised €115 million in 2025–26, the company bypasses self-serve trials altogether;

instead it deploys specialists to redesign legacy workflows by running proofs of concept on proprietary customer data during the presales phase. These surrogate solutions train on historical simulation data to prove measurable impact before any enterprise licensing is discussed. This high-touch approach can overcome industrial skepticism with hard metrics: At Airbus, predictions shrank from one hour to just 30 milliseconds, unlocking 10,000 times more design variants, and at General Motors, crashworthiness simulations plummeted from days to seconds.³²

Value-based commercial model

Leading AI companies are increasingly adopting hybrid commercial models, blending per-seat licensing with usage-, consumption-, or value-based pricing to tie software costs directly to delivered results. Modern enterprise buyers increasingly demand measurable ROI before scaling AI across their organizations; in Europe, the bar is higher still. Smaller local software budgets, tighter IT budgets, and “investors who scrutinize burn rates more heavily than US peers,” as one executive noted, all force start-ups to shift away from pure seat-based deals with real urgency.

Hybrid pricing that links revenue to documented outcomes lowers the adoption hurdle, helping European vendors gain and expand accounts without relying on up-front license scale while preserving the cost predictability that buyers require. As the CEO of a European AI company noted, “From the get-go, we always anchored our conversations on outcomes: How much time are you negotiating something, and how does this reduce the time to negotiation?”

n8n, a Berlin-based workflow automation platform, illustrates this approach. The company helps technical teams connect apps, APIs, and data to automate core business processes at scale. It drives adoption by pricing on completed workflow executions rather than per-step operations or per-seat licenses.³³ Customers pay when an entire automation runs from trigger to finish, regardless of workflow complexity. Enterprise plans scale on execution volume, and n8n positions this as tying cost directly to operational value delivered.³⁴

AI-enabled internal operations

AI-enabled operations involve redesigning internal functions to automate routine execution, thereby freeing human capital for higher-order strategic decision-making. Globally, this lets tech companies decouple revenue growth from linear head count expansion. In Europe, the case is sharper still. Since employment laws make rapidly adjusting head count to match short-term market fluctuations structurally more complex, companies are especially incentivized to maximize workforce efficiency from the outset. AI-assisted automation across finance, support, delivery, and internal tooling preserves scarce talent for customer-facing and product work. This discipline is not new. “Because European start-ups have always operated in capital-constrained environments, they have developed a culture of extreme capital efficiency,” one European tech investor noted.

ElevenLabs, a London-based voice AI company, runs much of its own internal operations on its agent platform. Initial sales qualification calls are handled entirely by AI, using the company's own gen AI speech and audio models. Job candidates interact with a virtual recruiter that explains the business and handles scheduling logistics. A Slack bot answers routine contract, service-level, and policy questions from hundreds of account executives, leaving the legal team to handle exceptions. Because each agent runs on the same underlying platform, their knowledge bases compound: the legal bot drew immediately on material already assembled for the sales and support agents, lowering the cost of automating each additional function.

AI talent strategy

An effective AI talent strategy goes beyond hiring elite researchers. Recruitment targets individuals with a unique hybrid of technical excellence, product thinking, and rapid experimentation. The demand for AI fluency, such as the practical ability to use and manage AI in day-to-day work, increased five times in Europe between 2023 and 2025, compared with just 1.7 times for technical AI skills.³⁵ As a sales director at a vertical AI start-up stated, hiring increasingly focuses on a growth mindset: "In the interview context, what I would try to understand is how adaptable are they? Are they using the latest tools?" Recent job postings emphasize adaptability: comfort with ambiguity, rapid experimentation, continuous learning, and a willingness to take on roles that evolve as the company scales. For example, in advertising for a sales leader, Lovable specifies that applicants should have "genuine fluency with AI tools in your day-to-day workflow,"³⁶ PhysicsX expects communications hires to use AI as a tool to streamline projects and sharpen thinking,³⁷ and 9fin has sought an AI fluency lead to help embed AI into its work processes through tool trials, role-relevant training, and many-to-many learning via lunch-and-learn sessions.³⁸

Because scaling only through recruitment is challenging, 82 percent of executives see upskilling as the best route to securing scarce resources.³⁹ Structured gen AI training and coaching can help developers complete coding tasks up to two times faster.⁴⁰ Crucially, this succeeds only when upskilling is paired with workflow redesign, not tool rollout alone.⁴¹ In traditional companies, successful upskilling is typically driven from the CEO and board levels. As a European board chairman noted, "If you set the tone at the board level and inspire the whole organization to go for it, the results are enormous. It has to start in the boardroom." For European start-ups operating under tighter capital and labor cost constraints, targeted hiring combined with broad workforce upskilling is critical to building lean teams that can compete globally. To institutionalize this adoption efficiently, companies are increasingly relying on peer-driven learning models. As the CEO of a European AI start-up explained, "We pick one champion per team that needs to rethink the whole department and the way how they work."

By embracing these operating practices, European AI start-ups may be better positioned to overcome regional capital constraints to build highly efficient, competitive businesses that can scale effectively. Many fast-growing, European AI application start-ups are already adopting most of these practices to get compounding effects. (See sidebar, "Prove early, move fast: How Parloa scales adoption"). The following section outlines how founders, operators, and investors can tap into this momentum.


Prove early, move fast: How Parloa scales adoption

Parloa's recent growth story illustrates the value of moving very quickly in the AI application era. Backed by €427 million in funding since 2025, the Berlin-founded contact center AI platform more than tripled annual recurring revenue in one year, to \$50 million at end of 2025. Three decisive moves the company made help to sustain that rapid pace: prove return on investment early, deploy fast once live, and leverage AI to run the company.

- **Proving ROI early.** Parloa avoids seat-based licenses that commit enterprises upfront for capacity they may not use. Cost rises only with usage, so buyers can pilot, add use cases, and expand without renegotiating seats. As founder and CEO Malte Kosub explained, “We price per minute and per conversation. It’s visible, and it’s familiar.” Before the sale, he noted, the company often builds “mock integrations to show the value”. Because connecting to live backends before closing a sale would take too much time, Parloa simulates client systems to show how deeply integrated agents execute tasks.
- **Deploying fast once live.** For one insurance customer, pre-sale, proof-of-concept models went straight to production, going live in four months and pushing call wait times down 33 percent in four weeks. Parloa’s engagement managers and forward-deployed engineers typically work with customers as an internal product team, consistently shipping new use cases, languages, and regions based on a roadmap. As a result of this approach, Parloa reports 150 percent net revenue retention.
- **Leveraging AI for its own operations.** Parloa has enabled the same speed of innovation internally by investing in a central AI platform rather than having hundreds of employees build different tools in parallel. As Kosub put it, “When everyone is doing their own things, it won’t scale. So we brought together a team and told them, ‘it’s your job to scale AI across the organization.’” These dedicated employees design AI workflows team by team, helping to make the adoption curve turn “exponentially higher,” as Kosub noted. Some examples of the platform include “the kitchen,” a purpose-built environment for engineers featuring more than 500 tools; “Rev-ops brain” for go-to-market queries; and a “Value Studio” that auto-generates client business cases.

Stakeholder takeaways: Seizing the AI application opportunity

As the AI landscape evolves, the shift away from generic scale toward workflow-specific depth presents a unique opportunity for the European tech ecosystem. The heterogeneous markets and complex, regulated sectors that have long made scaling difficult could evolve into a foundation for new tech resilience. However, tapping into this momentum requires successfully driving demand and adoption by Europe's risk-averse enterprise customers, which has often been a stumbling block for the region's scale-ups. The following lessons highlight how founders, established software leaders, and investors are navigating this shift to overcome that challenge and build sustainable value.



Opportunities for founders

AI is creating entirely new markets—some in categories that barely existed three years ago and that are, for a short period, genuinely up for grabs. Founders who wait for market definitions to settle may end up competing with someone who has already built a first-mover advantage. Especially in this early stage of the AI cycle, where founders choose to play has outsize implications. Successful founders seize the opportunity through three deliberate moves:

- **Targeting high-value workflows.** The workflows worth pursuing carry a visible, quantified cost today: rework, manual reconciliation, scarce specialist time, compliance overhead, delayed revenue. Driving value for customers requires that companies make deep investments to embed into customer operations, workflows, and system of record. Success depends on becoming a deeply integrated solution with high switching cost. As one European engineering AI leader described it, “We run the proof of concept on the customer’s own historic data before anyone talks about a license. Show an engineering director their own simulation collapsing from an hour to milliseconds, and the sale is largely over.”
- **Leveraging Europe’s sector ecosystems.** Europe’s strongest verticals offer assets that generic AI may be hard-pressed to create: strategic capital tied to industry advancement, proprietary data, regulatory insight, and buyers who trust local operators. Partnering with industry associations and corporate players can provide access to the data and domain expertise needed to train successful vertical models. As one AI leader noted, “A €7 million check from an incumbent doesn’t look like much next to a US round, but it arrives with the customer, the data, and the credibility attached.”
- **Planning for global scale from inception.** Localization once made it difficult for European companies to expand beyond national borders. AI has the potential to change that. By automating the management of divergent rules that have made multimarket deployment expensive, cross-border architecture can be built in from launch. Integrating cross-border architecture, regulatory readiness, and go-to-market strategies from day one gives first entrants a chance to thrive in multiple markets before local challengers scale.



Opportunities for software incumbents

The market devaluation of software companies early this year has put established software companies on notice: If AI isn't a clear part of your revenue story, your future is uncertain. Established software leaders prove their AI credentials by innovating, distributing, and resetting.

- ***Innovating by building or partnering with AI natives.*** A board member of software companies encourages established software companies to get started now to figure out the high-value workflows for which customers will spend on AI solutions. “Yes, the technology is new and unproven beyond a few use cases. But if you’re not experimenting now, you will wake up in three or five years’ time, and the world will have moved forward.” A potentially faster, less risky approach is to partner with new AI start-ups, many of which already possess the AI solutions the market is seeking.
- ***Leveraging a strong customer base.*** Established software companies have distribution, customer relationships, and perhaps most importantly, customer context. Markets have questioned incumbents’ ability to maintain leadership in the AI era, but workflow depth often favors those already embedded in the customer system.
- ***Resetting operations for the AI era.*** AI transformation is a fundamental business process change. Treating it as an IT upgrade can hinder value creation and strain existing operations by accelerating rather than redesigning legacy processes. Customers built decades of process around existing products; new commercial models and adoption-centric go-to-market strategies are what can turn technology into new revenue. To capture true productivity gains, leaders are setting ambitious targets to rethink end-to-end workflows from the ground up. As one board executive noted, “We put AI on the board agenda every other week, and anyone can present what they built. Once the board is watching, good ideas spread instead of dying in one team.”



Opportunities for investors

Global software valuations continue to be much lower than historical multiples. In July 2026, the median public software company traded at 3.8 times forward revenue, 56 percent below its long-run median of 8.7 times. However, this discount is not spread evenly. Among companies with the same rule of 40 score⁴², those weighted toward growth traded at eight to ten times revenue, while those weighted toward efficiency traded at 3.5 to four times.⁴³ To help create growth, leading investors are focusing on three main strategies:

- ***Adapting capital deployment to AI scaling velocity and multimarket reality.*** European AI leaders often expand across borders earlier than US peers, which spend more time scaling in a large single domestic market. Many investors in Europe are now adapting their capital deployment strategies to support cross-border go-to-market efforts and growth-stage scaling, accelerating expansion before domestic proof points alone would justify it under traditional SaaS timelines. As one investor put it, “It is far too easy here to win your national market and then stall. Fund that pattern, and you have financed a national champion, not a global one.”
- ***Orchestrating industrial connections and strategic co-investments.*** Europe’s potential strength lies in mature sectors (such as advanced manufacturing, automotive, and healthcare) that feature proprietary data and long-standing relationships. At the same time, procurement in those same enterprises can be as much of a constraint as demand. As one private equity leader mentioned, “In Europe, an enterprise might say, ‘That sounds interesting. Let me come back to you in eight months, and I will send you a 50-page RFP [request for proposal].’” Leading investors work to overcome this obstacle by opening doors for young AI application companies, fostering closer ties with industrial customers and syndicating rounds with corporate partners. Bringing corporates directly onto the cap table can provide vital growth capital and unlock rich proprietary data—both of which are essential for accelerating the path from pilot to production.
- ***Providing sustained support through global scaling.*** The widest capital gap between US and European start-ups opens at the breakout and late stages, which is precisely when European application layer companies need funding to compete globally. Active ownership plays a critical role here: Private equity sponsors and sovereign co-investors who stay committed through global scaling help portfolio companies navigate operational changes, sustain momentum, and avoid premature exits.

Conclusion

As AI moves deeper into enterprise workflows, Europe has an opportunity to reset its competitive standing in the rapidly changing technology landscape. The region still trails significantly in overall AI funding and investment, especially in platforms and models, and many of the structural constraints that have historically slowed European technology companies remain. But the AI market's evolution from its initial, primary focus on foundational models to the development of highly tailored applications could play to Europe's strengths. Founders there have long navigated the complexities of heterogeneous markets and regulated, entrenched industrial incumbents. The recent growth in venture capital for AI application start-ups, and the rapid scaling of some of the most promising start-ups, suggests that Europe could be well positioned to seize this new opportunity. Whether early momentum translates to sustainable, long-term success will depend on how quickly founders, incumbents, and investors can turn those structural strengths into global scale.

Endnotes

- ¹ As of August 2026, Boardwave's network includes more than 2,000 European software leaders and investors.
- ² "Europe's Top 100 Most Valuable Tech Companies," Multiples, updated August 2026.
- ³ McKinsey proprietary data; Dealroom; Gain; PitchBook, accessed July 2026. Venture capital funding to tech start-ups in Europe and the United States, 2025–26. Includes companies whose main offering is a product or service that is not AI software or software infrastructure (for example, digital banks, robotaxis, and hardware-heavy defense).
- ⁴ McKinsey proprietary data; Dealroom; Gain; PitchBook, accessed July 2026. Platform and infrastructure venture capital funding, United States, H1 2026.
- ⁵ "Europe's moonshot moment: Fueling its tech ecosystem for scale," McKinsey and Boardwave, June 2, 2025.
- ⁶ *Talent in, talent out: The shifting geography of the global AI workforce*, Interface, April 29, 2026.
- ⁷ *Artificial intelligence index report 2026*, Stanford Institute for Human-Centered Artificial Intelligence (HAI), April 13, 2026.
- ⁸ "The European startups offering the highest salaries," *Sifted*, July 24, 2026; "Top AI Companies in France in 2026: Mistral, Hugging Face, and the Paris AI Ecosystem," Talenlio, April 1, 2026; "ML Research Scientist vs ML Applied Scientist Jobs UK 2026: Which Pays More?," Machine Learning Jobs, June 2026.
- ⁹ *Artificial Intelligence Index Report 2026*, Stanford Institute for Human-Centered Artificial Intelligence (HAI), April 13, 2026.
- ¹⁰ Company funding announcements: Neural Concept, Series C (\$100 million, led by Goldman Sachs Alternatives Growth Equity), December 2025; Synthesia, \$4 billion valuation, January 2026; and Wayve, Series C (\$1.2 billion), February 2026.
- ¹¹ "European Spinouts Report 2025," Dealroom, November 20, 2025; Leah Hodgson, "VC hunger for deep tech sees uni spinouts claim record share of European investment," PitchBook, May 1, 2026.
- ¹² TenU USIT (University Spinout Investment Terms) is a set of best-practice guidelines and recommendations designed to turn academic research into commercial start-up companies more efficiently and make negotiations faster and fairer between universities and investors by standardizing terms for equity, licensing, and royalties. It was created in 2023 by a group of leading European universities and investors in affiliation with the organization TenU, an international collaboration of technology transfer offices from leading universities in the US, UK, and Europe.
- ¹³ *Spotlight on spinouts: UK academic spinout trends*, Royal Academy of Engineering and Beahurst, March 2025.
- ¹⁴ "Spinout Equity Policy," University of Oxford, accessed August 2026; "Spotlight on spinouts: UK academic spinout trends," Royal Academy of Engineering and Beahurst, March 2025; "FAQ: Spinning out of ETH," ETH Zurich, accessed August 2026.
- ¹⁵ "How corporate-scale-up partnering can boost Europe's tech competitiveness," McKinsey, June 10, 2026.
- ¹⁶ *18 new industry—AI partnerships join the growing wave of collaborations*, European AI Champions Initiative, November 2025.
- ¹⁷ "The state of European AI in 2025: Insights from 150 tech leaders," blog post by Badr Eddial, Lleverage, May 29, 2025.
- ¹⁸ McKinsey proprietary data; Dealroom; Gain; PitchBook, accessed July 2026. Share of European AI application venture capital deal value in the first half of 2026 with at least one sovereign, state-backed, or public-institutional investor participating in the round.
- ¹⁹ The US Defense Advanced Research Projects Agency
- ²⁰ McKinsey proprietary data; Dealroom; Gain; PitchBook, accessed July 2026. Venture capital funding to AI Application start-ups in Europe and the United States, 2025–26.
- ²¹ "Europe's moonshot moment: Fueling its tech ecosystem for scale," McKinsey and Boardwave, June 2, 2025.
- ²² "The Global Startup Ecosystem Report 2026," Startup Genome and Global Entrepreneurship Network, June 2026.
- ²³ Maya Lekhi, "Synthesia Business Breakdown & Founding Story," Contrary Research, accessed August 2026.
- ²⁴ "Make Videos in 160+ Languages and Accents," Synthesia, accessed August 2026; and "AI Actors," Synthesia, accessed August 2026; and Lekhi, "Synthesia Business Breakdown & Founding Story."
- ²⁵ "Synthesia Revenue, Valuation & Funding," Sacra, accessed August 2026.
- ²⁶ "Average SaaS spend per employee in 2026: The definitive benchmark," blog post by Brad van Leeuwen, Cledara, March 25, 2026; and "UK vs US vs Europe: How SaaS spending differs by region," blog post by Harald Meyer-Delius, Cledara, March 25, 2026.

- ²⁷ “Europe’s moonshot moment: Fueling its tech ecosystem for scale,” McKinsey and Boardwave, June 2, 2025.
- ²⁸ Charlotte Relyea, Martin Harrysson, and Matt Linderman, “Unlocking the value of AI in software development,” McKinsey, November 3, 2025.
- ²⁹ Aaron De Smet, Drew Goldstein, Holly Price, and Tanguy Catlin, “From adoption to impact: Three horizons of AI transformation,” *McKinsey Quarterly*, July 8, 2026.
- ³⁰ Alex Singla, Alexander Sukharevsky, Elia Berteletti, Lareina Yee, and Michael Chui, “The next innovation revolution—powered by AI,” McKinsey, June 20, 2025.
- ³¹ “\$85,000 in tokens later: What I learned from scaling agentic coding at Lovable,” Lovable, July 3, 2026.
- ³² “Renewed Collaboration between Airbus and Neural Concept,” Neural Concept, September 15, 2021; “Revolutionizing Pedestrian Safety with AI: General Motors with Neural Concept,” Neural Concept, August 5, 2025.
- ³³ “Pricing,” n8n, accessed July 2026.
- ³⁴ “Leveraging automation in cybersecurity” and “Delivery Hero saved 200 hours/month with a single IT ops workflow,” n8n case studies, both accessed July 2026.
- ³⁵ “Agents, robots, and us: How AI reshapes work and skills in Europe,” McKinsey Global Institute, May 12, 2026.
- ³⁶ “Founding sales leader, NYC” (job posting), Lovable, accessed July 2026.
- ³⁷ “Content and communications strategist,” PhysicsX job posting on Greenhouse, accessed July 2026.
- ³⁸ “AI fluency lead,” 9fin job posting on Flexa, accessed July 2026.
- ³⁹ Sandra Durth, Anna Lena Kalinna, and Henning Soller, “Closing the tech talent gap: Adopting the right mindset,” McKinsey, March 20, 2023.
- ⁴⁰ Begum Karaci Deniz, Chandra Gnanasambandam, Martin Harrysson, Alharith Hussin, and Shivam Srivastava, “Unleashing developer productivity with generative AI,” McKinsey, June 27, 2023.
- ⁴¹ Charlotte Relyea, Martin Harrysson, and Matt Linderman, “Unlocking the value of AI in software development,” McKinsey, November 3, 2025.
- ⁴² The Rule of 40 states that the sum of a software company’s growth rate and profit margin should be at least 40 percent.
- ⁴³ “Meritech Software Pulse,” Meritech Capital, July 17, 2026; Capital IQ; median enterprise value to next-twelve-months revenue across 100+ public software companies.

