

INFORMATION MATTERS

ANALYST RESEARCH // EUROPEAN AGENTIC AI

Sizing the European Agentic AI Stack

Market sizing, regulatory architecture & sovereignty horizon

Thematic Report — May 2026

Fiona (AI Analyst) | Technology Analyst

Martin De Saulles | Principal Analyst

Information Matters | May 2026

Headline findings

- European agentic AI market sized at **\$5.5 billion** in 2026 (range \$4.0–\$7.5B), built bottom-up from 113 European-headquartered companies — the lowest published European estimate and the only one built from vendor revenue rather than enterprise spend.
- Europe sits at **13-17 per cent** of the \$40 billion global agentic AI market — a smaller share of the world than Europe's share of global GDP. The gap is the analytical pivot of this report.
- The August 2026 EU AI Act high-risk-system rollout is the regulatory date that matters most. Vendors that started compliance work early are seeing measurable RFP advantage in late 2025 and Q1 2026; the gap will widen rather than narrow through 2027.
- Eight European-headquartered companies lead the commercial watch list: **Mistral, Aleph Alpha, Helsing** (€12B as of June 2025), **Parloa, ElevenLabs, Veriff, Lawhive** and **Tines**.
- The European stack runs on US infrastructure: **70 per cent** of the European cloud market is held by AWS, Azure and Google Cloud, and the EU's Chips Act target of a 20 per cent global chip-manufacturing share is now forecast by the European Court of Auditors at **11.7 per cent by 2030**. Substitution will reach the application and enabler layers first.
- The sovereignty horizon for multinationals is a **2027 execution question**, not a 2030 planning question — vendor-substitution decisions are visible in Q2 2026 RFPs already.

Executive summary

We put Europe's agentic AI market at **\$5.5 billion in 2026**, with a range of \$4.0 to \$7.5 billion. Our figure is built bottom-up from 113 European-headquartered companies in the Information Matters library plus the European-attributable share of foundation-model revenues, and it places Europe at 13 to 17 per cent of the global agentic AI market we sized in our Q1-2026 report at \$40 billion. That is a smaller share of the world than Europe's share of global GDP, and the gap is the analytical pivot of this report. The diffusion of agentic AI in Europe is constrained by three things at once: a foundation-model layer concentrated in two providers; a regulatory architecture that adds real, material cost to deployment timelines; and an enterprise procurement cadence that has not yet adjusted to the new technology's economics. The interaction of those three is the European story.

Three positions sit behind the number.

The European foundation-model layer is small but commercially real. Mistral and Aleph Alpha together account for the substrate. Mistral is the larger commercial play with the more diversified customer profile and the closer working relationship with the European public sector; Aleph Alpha sits closer to a sovereign-AI utility for German and EU public-sector buyers. Both will continue to matter through 2030. Neither is on the trajectory to reach the revenue tier that OpenAI or Anthropic will occupy by then.

The application layer is where Europe is most active. Sixty-four European-headquartered agentic AI companies are doing useful commercial work across coding, customer-operations, vertical specialisation, and voice. Average funding stage runs earlier than the equivalent US cohort, but the segment leaders — Helsing in defence, Parloa in customer-operations, Lawhive in legal, ElevenLabs in voice — are already operating at scale. The story here is not that Europe lacks builders. It is that the builders are competing against a US application layer with a meaningful funding-and-timing advantage.

The deployment cadence is slower than the published narrative suggests. EU AI Act compliance work is genuinely consuming enterprise procurement timelines, and the high-risk-system rollout from August 2026 will compress vendor-side decision windows further. The sovereignty horizon claim from our published hypotheses — that multinational enterprises will be navigating which agentic AI stack to use in which jurisdiction within twenty-four months — is sharpened in the final two sections of this report, which together set out the structural sovereignty backdrop and the multinational decision sequence inside it. The decisions are not a 2030 planning question. They are a 2027 execution question.

This is Information Matters' first thematic deep-dive on the European stack. The methodology mirrors our quarterly approach: primary sources only, foundation-model attribution rate disclosed in full, scope reconciled against four external published sizings (McKinsey, the European Investment Bank, Atomico, the OECD). We will revisit

this thematic at six-monthly cadence.

The European agentic AI market in 2026

We size the European agentic AI market at \$5.5 billion in 2026, with a bear case of \$4.0 billion and a bull case of \$7.5 billion. The figure aggregates from three sub-totals.

Foundation-model layer (Europe-attributable): \$1.0-1.4 billion. Mistral's run-rate revenue at end-Q1 2026 sits in the \$400-600 million range based on disclosed enterprise contracts, public-sector commitments, and OpenRouter API throughput. We apply the same 25-35 per cent agentic-attribution rate we use globally; the European mix may run modestly higher because the enterprise demand profile skews toward more-agentic workloads (compliance, document automation, code), though we have not adjusted the rate explicitly for this report's v1.0. Aleph Alpha contributes a smaller revenue base but a higher share of public-sector and sovereignty-anchored deployments where the agentic attribution is clearer. Combined Europe-attributable foundation-model revenue is \$1.0-1.4 billion.

Application layer (European-headquartered): \$2.5-3.8 billion. Sixty-four European-headquartered agentic AI companies in the IM library across the segments we cover. The funding-stage distribution is weighted earlier than the equivalent US cohort: of the 113 European companies in the library, only thirteen sit at Series C or later. We size revenue per-company by funding-stage band and segment archetype, using the same per-company revenue ranges applied in Q1-2026, then aggregate. The bear case assumes the early-stage tail produces materially less than the median; the bull case allows for the named-cohort outliers — Helsing, Parloa, Lawhive, ElevenLabs — carrying a larger share than the funding distribution implies.

Enabler and adjacent layers: \$0.5-1.3 billion. Tooling, orchestration, observability, evaluation, identity. Forty-seven European-headquartered companies in this band. The European enabler layer is genuinely strong — Tines, n8n, DRUID AI, deepset (Haystack), Langfuse, the broader observability tier — but smaller in aggregate than the equivalent US enabler layer. We size this category cautiously; the bull case allows for larger contributions from n8n's European deployment base and the security-automation cohort.

The reconciliation against external European-AI sizings is sparse but improved over our v0.5. Four external sources are worth addressing.

McKinsey's most recent Europe-and-Generative-AI note (May 2025) put European AI at \$50-80 billion across all of generative and agentic combined; applying a 25-35 per cent agentic share gives \$13-28 billion. That is materially higher than our figure. The gap is methodology rather than disagreement: McKinsey appears to size from enterprise spend on AI projects, double-counting the consulting and services attached to the underlying technology. Our \$5.5 billion measures vendor revenue.

The European Investment Bank’s 2024–25 Investment Report flagged AI as the fastest-growing investment category in Europe and a structural priority for European competitiveness, particularly in the chapter-three “scale-up gap” analysis.¹ The report does not size the agentic portion specifically, and the public chapter does not contain a clean cumulative European-AI VC figure of the kind we’d footnote here; we use the EIB qualitatively to support our framing rather than quoting a specific cumulative number, and treat the report as supportive of our claim that the European stack is smaller than its share of the GDP base would predict.

Atomico’s *State of European Tech 2025* reports approximately **\$14 billion** of European VC funding flowing into AI/ML companies during 2025 — the largest single category, with AI/ML representing roughly 31 per cent of all European VC capital raised in the year.² At current FX, that is around €13 billion. The cohort distribution Atomico publishes — heavy weighting toward the UK, France, and Germany, with notable showings from Israel-overlap players — broadly aligns with the IM library distribution.

The OECD’s AI Policy Observatory tracks European AI investment and adoption indicators but does not currently publish an agentic-specific cut. Their adoption-rate data — European enterprise AI adoption running approximately 60–70 per cent of US adoption rates, varying by sector — supports our claim that the European deployment cadence is structurally slower.

Across the four sources, our \$5.5 billion figure is the lowest published European agentic AI estimate we have seen, and the only one built bottom-up from vendor revenue rather than from enterprise spend or VC-investment proxies. We expect the figure to be challenged on its conservatism, and we welcome the challenge.

The regulatory architecture

The EU AI Act came into force in August 2024 and is being phased in over three years. By August 2026, the high-risk-system classification provisions take effect. That is the regulatory date that matters most for agentic AI deployment in Europe.

The General-Purpose AI obligations — in force from August 2025 — already require providers of general-purpose AI models, including foundation models served into the EU market, to publish detailed model documentation, copyright-compliant training-data summaries, and post-deployment incident reporting. Mistral and Aleph Alpha publish to this standard natively. OpenAI, Anthropic, Google, and Meta are also compliant for their EU-served products, with more variation in the depth of disclosure than the European-headquartered providers offer. The GPAI tier has not, in our reading, become the binding constraint for most deployments. It has, however, raised the bar

¹European Investment Bank, *EIB Investment Report 2024/2025: Innovation, Integration and Simplification in Europe*, chapter 3 (the “scale-up gap” analysis). <https://www.eib.org/en/publications/20240354-investment-report-2024>

²Atomico, *State of European Tech 2025*, “Startup Investment Trends” chapter. <https://www.stateofeuropeantech.com/chapters/startup-investment-trends>

for what counts as serious model documentation, and we expect that bar to compound through 2026 and 2027.

The high-risk-system tier — phasing in from August 2026 — is where agentic AI deployment in Europe will materially compress. Any agentic system used in employment decisions, education, essential public services, law enforcement, migration, justice, or significant private-sector decisions affecting consumers (creditworthiness, insurance pricing) falls into the high-risk classification. Once classified, the system is subject to mandatory pre-deployment conformity assessment, post-deployment monitoring, and human-oversight requirements that are demanding to implement and that no vendor can plausibly retrofit on a short timeline. Enterprises selling into European customers in any of these segments are running large compliance programmes now. The vendors that started early have a measurable RFP advantage in late 2025 and early 2026; the ones that didn't are visibly catching up.

National sovereignty programmes vary in size and direction. France's €1.5 billion AI investment programme is the most concrete, anchored on Mistral, the public computing infrastructure FR-Alliance, and the Choose France AI Summit cadence. Germany's federal AI strategy is more dispersed across the Länder and less commercially anchored, with Aleph Alpha as the focal commercial play and a stronger emphasis on the public-sector adoption side. The UK's post-Brexit position differs structurally — no EU AI Act compliance burden, a lighter-touch domestic AI policy, and a stronger emphasis on the foundation-model evaluation work that the AI Security Institute leads. The UK's strategic position is genuinely different from the European Union's, and this report's sizing treats the UK as part of the European stack while acknowledging the policy divergence.

The AI Liability Directive, separately tracked, would extend product-liability principles to AI systems and is in late drafting. If it passes in its current shape during 2026, it will create a presumption of fault for high-risk-system harms that materially raises insurance and indemnification costs across the European stack. We expect the legislative trajectory to crystallise during the second half of 2026 and will return to it in a dedicated brief at that point.

The implication for vendors — including non-European vendors selling into Europe — is that the regulatory architecture is not a future cost. It is a current cost. The August 2026 high-risk rollout is far enough out to be ignored in most pitch decks but close enough that responsible buying organisations are already specifying compliance in their procurement. Our reading is that the regulatory frame is the most under-priced variable in the European stack right now: vendors that have done the work are seeing measurable advantage, and the gap will widen rather than narrow through 2027.

The European stack: foundation models, enablers, agents, adjacent

Substrate

Two European-headquartered foundation-model providers in the library: Mistral AI (France) and Aleph Alpha (Germany). Both Series D+. Mistral is the larger commercial play, with a clearer enterprise distribution motion and the closer working relationship with the European public sector. Aleph Alpha sits in a narrower sovereignty-utility role; its commercial trajectory depends on the German federal procurement decisions through 2026 and 2027.

The substrate layer in Europe is not where the report's analytical interest sits. The European foundation-model layer is small in revenue terms relative to the application layer that depends on it, and the trajectory through 2030 is for that ratio to widen rather than close. The story is in what gets built on top of these two and on the imported foundation-model layer (OpenAI, Anthropic, Google, Meta) that European application-layer companies use without controversy.

Agents

Sixty-four European-headquartered agent companies. The distribution by segment archetype, using the IM Q1-2026 segment taxonomy:

- *Horizontal Enterprise Agents* — the largest segment, with cohorts in customer-operations, sales, and general-purpose enterprise automation. Notable: Parloa (Germany, customer-operations, Series D+), Pactum AI (Estonia, contract negotiation, Series B), Aizee (UK, retrieval-augmented enterprise agents).
- *Vertical and Specialised Agents* — Lawhive (UK, legal, Series B) is the standout commercial example; Helsing (Germany, defence, Series D+) is the largest by funding and the most strategically positioned for the sovereignty argument; Luminance (UK, legal) is the longer-established vertical AI play.
- *Capability-Specific Agents* — ElevenLabs (UK, voice) is the layer's commercial anchor in Europe; Gladia (France, voice and transcription, Series A) and Poolside (France, code generation) are the cohort to watch.

Enablers

Forty-three European-headquartered enablers covering orchestration, infrastructure, evaluation, observability, and security automation. The European enabler layer is one of the genuine bright spots in the stack: Tines (Ireland, security automation, Series B), n8n (Germany, workflow orchestration, Series C), DRUID AI (Romania, conversational orchestration, Series C), deepset / Haystack (Germany, retrieval orchestration), Langfuse (Germany, observability), PRISM Eval (France, evaluation), and DeepMind's Project Astra in the UK (a special case, given Google's ownership, but European-rooted in its research and a meaningful application-layer enabler regardless).

The compliance-and-evaluation tier is particularly active in Europe because the EU AI Act's high-risk-system tier creates a demand-pull for evaluation infrastructure that the rest of the world doesn't yet experience. This is a competitive advantage for European enabler vendors selling into European enterprises preparing for the August 2026 rollout. It does not yet translate into outsized US market share for the same vendors. We expect it to in 2027 and 2028 as US enterprises start absorbing similar regulatory architecture.

Adjacent

Four adjacents in the library: identity, fraud, image generation, and financial-crime infrastructure that intersects with agentic deployments without being agentic-native. Veriff (Estonia, identity, Series D+), Salv (Estonia, financial crime), Black Forest Labs (Germany, image generation, Series B), Helsing's adjacent operations.

Companies to watch

Eight European-headquartered agentic AI companies worth tracking through 2026 and into 2027. The selection prioritises commercial viability — revenue base, growth trajectory, market position — rather than spread across layers or segments. The full thirty-eight-company cohort is in the appendix.

1. **Mistral AI** (France, substrate, Series D+). Europe's leading foundation-model provider. Revenue base of \$400–600 million ARR at end-Q1 2026 is the commercial floor under the entire European foundation-model layer. Watch enterprise-deal disclosures, the FR-Alliance public computing relationship, and any move to a paid API tier with service-level agreements specific to high-risk-system deployments.
2. **Aleph Alpha** (Germany, substrate, Series D+). The sovereign-AI utility play. Smaller revenue base than Mistral, higher share of public-sector and strategic-enterprise deployments. Watch the German federal procurement decisions through 2026 and the public-sector adoption rate. The commercial fortunes of the company depend on a small number of large procurement outcomes.
3. **Helsing** (Germany, defence agent, Series D+). Defence AI valued at **€12 billion** as of the June 2025 Series D, more than double the prior round at approximately €4.95 billion in 2024.³ Defence-AI revenue is opaque by sector convention; we track contract-award disclosures and any expansion beyond Germany. Helsing is the most strategically positioned company in the European stack for the sovereignty argument and the most likely to be cited in any 2027 multinational procurement decision involving sensitive workloads.

³Helsing newsroom, "Helsing raises €600M to invest in European technological sovereignty," 17 June 2025. <https://helsing.ai/newsroom/helsing-raises-eur600m-to-invest-in-european-technological-sovereignty>. Cross-referenced via PitchBook, "Defense unicorn Helsing now among Europe's most valuable startups with €600M raise," 17 June 2025.

4. **Parloa** (Germany, customer-operations agent, Series D+). The European customer-operations agent leader. Series D+ funding gives Parloa the commercial runway to compete with US customer-operations specialists at scale. Watch contact-centre platform partnerships and ARR disclosure.
5. **ElevenLabs** (UK, voice infrastructure agent). The European voice-infrastructure leader and one of the fastest-growing companies in the global agentic AI sector. Funding-stage data is undisclosed in our library, but commercial signals — pricing-tier expansion, enterprise customer mix, partnership announcements — place ElevenLabs in the late-private commercial tier. Watch the enterprise pricing model and any move into agentic-voice composition rather than pure synthesis.
6. **Veriff** (Estonia, identity adjacent, Series D+). Identity verification and fraud detection at scale, with a meaningful agentic-flow tie-in as enterprise agents increasingly need verified-user authentication. Veriff is a large company by European standards and one of the few Estonian Series D+ companies in the cohort.
7. **Lawhive** (UK, vertical legal agent, Series B). The most commercially visible vertical agentic AI play in Europe. Series B funding and ARR growth signals place Lawhive in the upper tier of European vertical AI, and the legal market's specific compliance and data-residency requirements give Lawhive a structural advantage against US competitors that have not localised. Watch ARR disclosures and category coverage.
8. **Tines** (Ireland, security automation enabler, Series B). Security automation orchestration with a strong US enterprise base — Tines is the European enabler most likely to grow into a US-scale business through 2027. Watch US enterprise penetration and any platform expansion beyond security automation into broader agentic workflows.

Digital sovereignty and the European AI stack

The European agentic AI sector does not run on European infrastructure. It runs on American infrastructure, accessed from Europe. Three US providers — Amazon Web Services, Microsoft Azure, and Google Cloud — control approximately 70 per cent of the European cloud market, with European cloud providers' share of their own home market having fallen from 29 per cent in 2017 to roughly 15 per cent by 2022, where it has held steady since.⁴⁵ The substrate beneath those clouds — the chips for training and inference — is dominated by Nvidia, with the EU's own Chips Act target of a 20 per cent share of global chip manufacturing now forecast by the European

⁴European Council on Foreign Relations, *Get over your X: A European plan to escape American technology* (ECFR/616, December 2025). <https://ecfr.eu/publication/get-over-your-x-a-european-plan-to-escape-american-technology/>

⁵Synergy Research Group, "European cloud providers' local market share now holds steady at 15%" (cited in the ECFR brief above). <https://www.srgresearch.com/articles/european-cloud-providers-local-market-share-now-holds-steady-at-15>

Court of Auditors to land at 11.7 per cent by 2030.⁶ The phone operating systems on which mobile-distributed AI runs are American. The foundation-model layer European application-layer companies use in production is dominated by OpenAI, Anthropic, and Google. Mistral and Aleph Alpha, sized earlier in this report, are the European exceptions rather than the norm.

For most of the 2010s and the early 2020s this dependency was an academic policy question. It is now a procurement and strategy question, and public sentiment has shifted faster than the policy machinery has. A YouGov European Political Monthly Survey conducted 6–16 March 2026 across France, Germany, Spain, Italy, and Poland (more than 1,000 respondents in each country) found that just over 60 per cent thought replacing US tech services in data storage, payments, video-conferencing, and email with European alternatives was a good idea. Sixty-two per cent said they favoured or had considered the substitution for data storage and payment services; fifty-nine per cent said the same for video-conferencing. The same poll found public opinion divided on whether such a shift is realistic — forty-one per cent thought it unrealistic, forty per cent thought it realistic, nineteen per cent did not know.⁷ The political headroom for substitution exists. The operational pathway is less obvious.

The policy machinery has begun to move regardless. France committed in January 2026 to generalising the domestic Visio platform across all state services by 2027, replacing Microsoft Teams and Zoom; the Ministry of Public Service announcement on 26 January 2026 named 200,000 agents as the deployment target, with CNRS, Assurance Maladie, the DGFIP, and the Ministry of Armed Forces among the Q1 2026 first-movers.⁸ The state government of Schleswig-Holstein in Germany completed its email migration off Microsoft Exchange and Outlook to Open-Xchange and Thunderbird in October 2025, with a LibreOffice rollout to roughly 30,000 workstations underway and projected savings above €15 million per year from 2026.⁹ The European Commission’s Industrial Accelerator Act, proposed on 4 March 2026, sets “Made in EU” requirements across public procurement in strategic sectors with specific quotas

⁶European Court of Auditors, *Special Report 12/2025: The EU’s strategy for microchips*, 28 April 2025 (citing the Commission’s July 2024 forecast of an EU global semiconductor value-chain share rising from 9.8 per cent in 2022 to 11.7 per cent by 2030, against the Chips Act’s 20 per cent target). <https://www.eca.europa.eu/en/publications/sr-2025-12>

⁷YouGov European Political Monthly Survey, March 2026 (fieldwork 6–16 March 2026), as reported by Tech Policy Press, “Almost Two-Thirds of Europeans Back Replacing US Tech, Poll Finds.” <https://www.techpolicy.press/almost-two-thirds-of-europeans-back-replacing-us-tech-poll-finds/>. The full topline tables for this wave were not publicly accessible at time of publication; the figures cited above are taken from the Tech Policy Press write-up.

⁸Direction interministérielle du numérique (DINUM) / Ministère du Service public, “Souveraineté numérique : l’État généralise « Visio »,” 26 January 2026. <https://www.numerique.gouv.fr/sinformer/espace-presse/souverainete-numerique-etat-visio-solution-visioconference-agents-publics/>

⁹Coverage in heise online and Nextcloud’s interview with Schleswig-Holstein state CIO Sven Thomsen documenting the Microsoft Exchange / Outlook → Open-Xchange / Thunderbird migration completed October 2025 and the wider LibreOffice rollout. <https://www.heise.de/en/news/Goodbye-Microsoft-Schleswig-Holstein-relies-on-Open-Source-and-saves-millions-11105459.html> and <https://nextcloud.com/blog/digital-sovereignty-in-action-interview-with-sven-thomsen/>. We have not located a single canonical schleswig-holstein.de press release consolidating all milestones; the citation here is to the strongest secondary record.

for low-carbon steel, aluminium, and cement.¹⁰ None of these are AI-specific. Each, however, redirects a sizeable public-sector spending base toward European vendors, with direct second-order effects on the agentic AI companies that sell into the same procurement frames. The enterprise side is moving too: a Gartner survey of 241 Western European CIOs and IT leaders, conducted May–July 2025, found that 61 per cent intend to shift more workloads to local or regional providers, 53 per cent plan to restrict their use of global hyperscalers, and 44 per cent have already started.¹¹

For the European agentic AI sector specifically, the implication runs in two directions at once. The demand-pull is real and will compound through 2027. European agentic AI vendors selling into European public-sector buyers, regulated industries, and the larger multinationals running EU-resident-data programmes will find themselves on shortlists they would not have made on capability alone two years ago. The named cohort in this report — Mistral, Aleph Alpha, Helsing, Parloa, Lawhive, Tines — is structurally well-positioned to capture that demand. The supply-side dependency is less easily addressed. European agentic AI companies still train and deploy on US hyperscaler infrastructure, still use US-headquartered foundation models in the majority of their production stacks, and will continue to do so for the foreseeable future. The 2026–2030 sovereignty story for the European AI sector is not one of substitution at every layer. It is substitution at the layers where the political, regulatory, and procurement pressure aligns — application layer first, then orchestration and observability, then foundation-model layer in a narrower set of public-sector and high-sensitivity verticals, and the cloud and chip layers last and least.

This is the harder version of the sovereignty argument and the version that will determine whether the European agentic AI market sized in this report grows materially faster than the central case through 2027 and 2028. If the substitution wave reaches the application and enabler layers — and the procurement and policy signals suggest it will — the upside sits in the bull range we already published at \$7.5 billion. If the substitution stays mostly at the public-sector edges, the central \$5.5 billion case is the right one. Our reading is that application-layer substitution is already in train and that 2026 is the first year it shows up in vendor revenue numbers. We will track that signal in the six-monthly revisits to this report.

¹⁰European Commission, IP/26/515, “Industrial Accelerator Act” proposal press release, 4 March 2026, and proposal text COM(2026) 100 final. https://ec.europa.eu/commission/presscorner/detail/en/ip_26_515. The Act is a Commission proposal entering ordinary legislative procedure; final adoption requires Council and Parliament agreement.

¹¹Gartner, “Geopolitics Will Drive 61% of CIOs and Information Technology Leaders in Western Europe to Increase Reliance on Local Cloud Providers,” 12 November 2025 (survey conducted online May–July 2025 among 241 Western European CIOs and IT leaders). <https://www.gartner.com/en/newsroom/press-releases/2025-11-12-gartner-survey-reveals-geopolitics-will-drive-61-percent-of-cios-and-information-technology-leaders-in-western-europe-to-increase-reliance-on-local-cloud-providers>

The sovereignty horizon for multinationals

Beyond the structural picture set out in the previous section, the multinational decision sequence is the part of the sovereignty story that is most actionable in the next eighteen months. The About page commits Information Matters to treating AI sovereignty as central rather than peripheral, and makes a specific claim: that multinational enterprises will be navigating which agentic AI stack to use in which jurisdiction within twenty-four months. This thematic sharpens that claim.

The sovereignty horizon is not “European stack vs US stack” in the abstract. It is a sequence of concrete jurisdiction-specific decisions that multinational enterprises hit in a defined order through 2027.

The earliest of those decisions is already happening. Any enterprise selling into the EU with agentic AI in a high-risk-system category needs a deployed conformity-assessment programme in time for the August 2026 rollout. This is procurement-level visible in Q2 2026 RFPs. Vendors that have done the work, or that operate in segments outside the high-risk classification, are seeing the advantage we discussed in the regulatory architecture section. By the second half of 2026, the multinational enterprises that started this work in 2025 will have selected their first European agentic AI vendors; the ones that delayed will be running compressed selection processes with weaker leverage.

The second wave of decisions sits in 2027. The same multinationals will be making concrete vendor-substitution decisions where their existing US-anchored agentic AI tooling lacks an EU-resident-data-and-compliance variant. The substitution mostly favours either US vendors that have done the EU compliance work — which is most of the larger ones — or European-headquartered alternatives where the European vendor has reached commercial maturity. The named cohort in this report flags which European companies are positioned for that substitution wave. Mistral, Aleph Alpha, Parloa, Helsing, Lawhive and Tines all have a structural fit to it.

The third wave is the AI Liability Directive, if it passes in its current shape during 2026. By 2027, it will be entering enforcement, and insurance and indemnification costs for high-risk-system deployments will materially shift. Multinational enterprises will be running parallel deployment stacks — one for EU jurisdictions, one for US jurisdictions, one for Greater China where applicable — with the cost of the parallel architecture being the visible decision pressure on board agendas.

We hold the view that the parallel-stack architecture is a 2027 reality for the multinationals that started the compliance work in 2025, and a 2027–2028 catching-up project for the ones that did not. The concrete sovereignty decisions are not a 2030 planning question. They are a 2027 execution question. That is the sharpened version of our published hypothesis, and it is the version we will be applying in our coverage of the agentic AI sector through the coming quarters.

Methodology

The European agentic AI sizing in this report is built bottom-up from primary sources. For the foundation-model layer Europe-attributable revenue, we use disclosed enterprise contracts, public-sector commitments, and API-throughput indicators (OpenRouter monthly volume) for Mistral and Aleph Alpha. We apply the same 25–35 per cent agentic-attribution rate used in our Q1-2026 quarterly report. The European enterprise mix may run higher in some segments because of the compliance-driven demand profile, but we have not adjusted the rate explicitly for v1.0 of this report; that adjustment is the most likely revision when we revisit at six-monthly cadence.

For the application and enabler layers, we use per-company revenue estimates by funding-stage band and segment archetype, drawn from the IM library of 632 companies (113 European-headquartered). Per-company estimates are not disclosed individually; the aggregation method is consistent with our quarterly approach and is described in full in our public methodology page.

For reconciliation, we addressed four external published sizings: McKinsey’s May 2025 Europe-and-Generative-AI note (\$50–80 billion total, implied agentic share \$13–28 billion), the European Investment Bank’s 2024–25 Investment Report (used qualitatively; see the reconciliation discussion above), Atomico’s *State of European Tech 2025* (approximately \$14 billion of European VC funding into AI/ML during 2025), and the OECD AI Policy Observatory’s adoption-rate indicators. None of the four sizes the agentic market specifically; the reconciliation explains where our figure differs and why. Our \$5.5 billion is the lowest published European agentic AI estimate we have seen and the only one built bottom-up from vendor revenue rather than enterprise spend or VC-investment proxies. We expect the figure to be challenged.

The library coverage of European companies is incomplete in two known ways. Fifteen of the 113 European entries do not have a populated funding-stage field — a coverage gap noted in our 25 April taxonomy refactor — and the regulatory-fintech and govtech segments may be under-represented relative to their actual market activity. Both gaps will be addressed in the next full library audit, scheduled for the Q2-2026 reporting cycle.

The figures cited in the digital-sovereignty section are sourced separately and are not part of the sizing reconciliation: cloud-market-share data from the ECFR’s December 2025 *Get over your X* policy brief and Synergy Research Group’s underlying data series; Chips Act numbers from the European Court of Auditors’ *Special Report 12/2025* (citing the Commission’s July 2024 forecast); public-opinion numbers from YouGov’s March 2026 European Political Monthly Survey as reported by Tech Policy Press; the Visio commitment from the French DINUM’s January 2026 announcement; the Schleswig-Holstein migration from heise online and the state CIO’s recorded interview; the Industrial Accelerator Act dated to its 4 March 2026 Commission proposal; and the Western European CIO indicators from Gartner’s November 2025 release of its May–July 2025 survey of 241 CIOs and IT leaders.

If you spot a methodological flaw or a figure you can't reproduce, write to us. We have changed our methodology in response to readers before and will do so again.

Methodology note

This report was produced through a combination of human expertise and oversight supported by an AI research agent specifically designed to carry out detailed market research and forecasts using established methodologies and with access to the latest data and AI models. This research consists of the opinions of Information Matters' research team, human and AI, and the information contained within it should not be considered as statements of fact. None of the information presented here should be taken as investment advice. Reproduction or distribution of this research without written permission from Information Matters Ltd is prohibited. © 2026 Information Matters Ltd. All rights reserved.

Appendix — The European Agentic AI Cohort

Thirty-eight European-headquartered companies worth tracking. The first eight repeat from the body's "Companies to watch" section in the same order; the remainder are listed by funding stage where disclosed, then alphabetically. Each entry is a one-paragraph take giving the commercial position, the segment, and what we are watching.

Headline cohort (eight)

Mistral AI — France, substrate, Series D+. Europe's leading foundation-model provider, with \$400-600 million ARR at end-Q1 2026. The commercial floor under the entire European foundation-model layer. Watch enterprise-deal disclosures and the FR-Alliance public computing relationship.

Aleph Alpha — Germany, substrate, Series D+. The sovereign-AI utility play. Smaller revenue base than Mistral; higher share of public-sector and strategic-enterprise deployments. Trajectory depends on a small number of large German federal procurement decisions through 2026 and 2027.

Helsing — Germany, defence agent, Series D+. Valued at €12 billion as of the June 2025 Series D, more than double the prior round at approximately €4.95 billion in 2024. Defence-AI revenue is opaque by sector convention. The most strategically positioned company in the European stack for the sovereignty argument; the most likely to be cited in 2027 multinational procurement decisions involving sensitive workloads.

Parloa — Germany, customer-operations agent, Series D+. The European customer-operations agent leader with the funding runway to compete with US specialists at scale. Watch contact-centre platform partnerships and ARR disclosure.

ElevenLabs — UK, voice infrastructure agent. The European voice-infrastructure leader and one of the fastest-growing companies in the global agentic AI sector. Watch the enterprise pricing model and any move into agentic-voice composition.

Veriff — Estonia, identity adjacent, Series D+. Identity verification and fraud detection at scale, with a meaningful agentic-flow tie-in as enterprise agents increasingly need verified-user authentication. One of the few Estonian Series D+ companies in the cohort.

Lawhive — UK, vertical legal agent, Series B. The most commercially visible vertical agentic AI play in Europe. Legal-market compliance and data-residency requirements give Lawhive a structural advantage against unlocalised US competitors.

Tines — Ireland, security automation enabler, Series B. The European enabler most likely to grow into a US-scale business through 2027. Watch US enterprise penetration and platform expansion beyond security automation.

Series-funded and IPO

SAP — Germany, agent, public. Established enterprise software with agentic features in Joule. Not agentic-native, but the largest revenue base in any European agentic-relevant company by some distance.

DRUID AI — Romania, orchestration enabler, Series C. Conversational AI orchestration with strong CEE and Western European enterprise traction. Watch the platform expansion into agentic flows beyond conversation.

n8n — Germany, orchestration enabler, Series C. Workflow orchestration with a large open-source community and growing commercial-tier deployment base. The European agentic-orchestration play with the broadest developer mindshare.

Edstronomy — Romania, vertical education agent, Series C. Vertical agentic AI in education with significant CEE commercial traction. One of the few Series C companies outside the German / French / UK core.

Black Forest Labs — Germany, image generation, Series B. Sits at the boundary between agentic and generative use cases; we treat the company as adjacent for sizing purposes. Notable European AI alumni network and a strong technical reputation.

Pactum AI — Estonia, contract negotiation agent, Series B. Vertical agentic AI for procurement negotiation with named-customer disclosures from Walmart and other large enterprise buyers. The strongest commercial signal from the Estonian cohort outside Veriff.

Gladia — France, voice and transcription, Series A. Voice infrastructure with API-first distribution. Direct competitive overlap with ElevenLabs in some segments; complementary in others.

Langfuse — Germany, observability enabler, Seed. Open-source LLM observability with rapid commercial adoption. Funding stage understates the commercial footprint.

Established commercial players

Luminance — UK, vertical legal agent. The longer-established vertical legal AI vendor in Europe. Established enterprise base across UK, US and ANZ legal markets. Less obviously agentic-native than Lawhive but with a deeper moat in document analysis.

PolyAI — UK, voice agent. Voice-driven customer-operations agents with a deep enterprise customer base. Established commercial position; competitive overlap with Parloa in some customer-operations segments.

DeepMind (Project Astra) — UK, application enabler. Google-owned and therefore partially out of scope for “European stack” sizing in a sovereignty sense, but European-rooted in research and a meaningful application-layer enabler regardless. Project Astra’s productisation through 2026 will affect what European competitors can plausibly build.

deepset (Haystack) — Germany, retrieval orchestration enabler. One of the better-known open-source retrieval-orchestration frameworks with European enterprise commercial uptake. Strong technical reputation in the European AI engineering community.

Dust — France, orchestration enabler. Agentic orchestration platform with a clean product surface and growing European enterprise customer base.

Synthflow AI — Germany, voice agent. Voice-agent platform with API distribution; competitive overlap with ElevenLabs in some segments and a more agentic-orchestration-led distribution model.

Poolside — France, code generation agent. European code-generation play with a research-heavy origin. Watch enterprise distribution and any Mistral or sovereign-AI integration.

Holistic AI — UK, governance enabler. AI governance and risk management platform with strong fit to the EU AI Act compliance demand-pull. Commercially well-positioned for the 2026–2027 high-risk-system rollout.

Humanloop — UK, observability enabler. LLM observability and evaluation platform with established enterprise customers. Direct overlap with Langfuse in some segments and with the broader observability tier.

PRISM Eval — France, evaluation enabler. French evaluation and red-teaming platform with strong fit to the EU AI Act compliance demand. Specialist player rather than generalist.

Notable specialists

Refact — UK, code agent. Open-source coding-agent platform with a growing self-hosted enterprise tier.

Quivr — France, retrieval agent. Open-source RAG framework with rapid community adoption and a commercial tier.

Salv — Estonia, financial-crime adjacent. Financial-crime infrastructure with agentic-flow tie-ins for large European banks. Strong commercial traction in Northern European financial services.

DeepOpinion — Austria, document intelligence enabler. Document-processing automation with agentic deployment patterns; small but commercially active.

Halion Fusion Group — Switzerland, horizontal enterprise agent. Swiss enterprise AI with regulated-industry positioning and cross-border European customer base.

Fetch.ai — UK, orchestration enabler. Agent-orchestration platform with a research origin and growing commercial tier.

Twin — UK, infrastructure enabler. UK enabler with a developer-tooling focus.

AutoGPT — UK, open-source orchestration. The original open-source agentic orchestration project with an evolving commercial tier. More notable historically than commercially.

iGent AI — UK, vertical code agent. UK code-agent specialist with an enterprise tier.

FairMind — Italy, vertical agent. Italian vertical AI with a regulated-industry positioning; small but representative of the southern European cohort.

Converflow — Spain, orchestration enabler. Spanish orchestration platform; representative of the smaller Iberian cohort.

LangWatch — Netherlands, observability enabler. Dutch observability platform with strong fit to evaluation and governance demand.