

The Agent Race

How real are fincrime vendors' AI agents?

AI agent claims and products across 333 financial-crime software vendors, mostly from Europe, the UK and the US

Authored by Derrick Cramer, Founder, Gossamer

November 2026

Contents

Foreword 3

Key findings 4

More than four in five vendors that claim agents have a released product 5

Ten vendors name a customer using their agents 6

US vendors lead at every step, but European claimers keep pace on work agents 7

Agents mostly do compliance groundwork 9

The platforms are building agents too 11

Conclusion 11

Data and methodology 13

Foreword

Of the 333 financial-crime software vendors we track, 129 say they offer AI agents. When we checked the evidence behind those claims, most vendors could point to a concrete product, only 10 named a customer using one in production, and US vendors stayed ahead of European vendors at every step.

Every figure rests on public sources that were rechecked against their live pages on 4 November 2026; the data and methodology section at the end explains how.

Our main report, *How Fincrime software vendors talk to the market after raising a round*, tracks 333 software vendors that sell financial-crime tools to banks, fintechs, insurers and other regulated firms. That report found a wide Atlantic gap in AI agents. In the past year, 38% of US vendors launched a product that mentions AI agents, against 7% of European vendors. In both reports, "European" includes the UK.

A launch that mentions agents shows what a vendor is promoting. It does not say whether a customer can use the product today, or whether any customer already does. For founders and marketers deciding how to present their own agents, and for buyers trying to tell one pitch from another, that difference matters. So this companion report asks a narrower question: how real are the agents?

To answer it, we started from every vendor that claims AI agents in its homepage headline or opening line, in what it says it offers, or in a product launch from the past year. We kept only claims that we could confirm on a live page on 4 November 2026, which left 129 vendors. For each of them we graded the evidence behind the claim against three questions. Is there a concrete product? Is it on general release? And does a named customer use it in production?

The report does not test how well any agent performs; it measures only what public evidence shows. Because most customers of compliance software are not named in public, its count of named customers is a floor. The data and methodology section at the end sets out the definitions, the checks and the limits in full.

The report is written for founders and marketers at fincrime vendors, for investors, and for compliance buyers weighing what stands behind an agent pitch.

Key findings

10

of 129 vendors that claim AI agents have a named customer using a work agent or assistant in production

1 Few vendors name a customer using their agents. Only 10 have a named customer using a work agent or assistant in production. Most customers of compliance software are not named in public, so this is a floor.

119

of the 129 vendors that claim AI agents have a concrete work agent, assistant or agent-trust product

2 Most agent claims point to a concrete product. Of the 129 vendors that claim AI agents, 119 have a concrete work agent, assistant or agent-trust product, and 110 have one on general release.

40%

of US vendors

17%

of European vendors

3 The Atlantic gap holds after checking. 40% of US vendors have a work agent or assistant on general release, against 17% of European vendors. In relative terms the gap is widest in agent-trust products, which 18% of US vendors sell, against 3% of European vendors.

45

of the 48 vendors whose homepage headline or opening line mentions agents have a concrete product

4 Vendors that put agents on their homepage can back the claim up. Of the 48 vendors whose homepage headline or opening line mentions agents, 45 have a concrete product.

47

know your customer (KYC), know your business (KYB) and due-diligence research

45

alert and case work

5 Agents mostly do compliance groundwork. The most common jobs are know your customer (KYC), know your business (KYB) and due-diligence research (47 vendors) and alert and case work (45). European vendors with a work agent use it for research and screening at least as often as US vendors do. Measured against all tracked vendors, Europe trails most in alert and case work, fraud operations and agent-trust products, where the US share is about four times the European share or more.

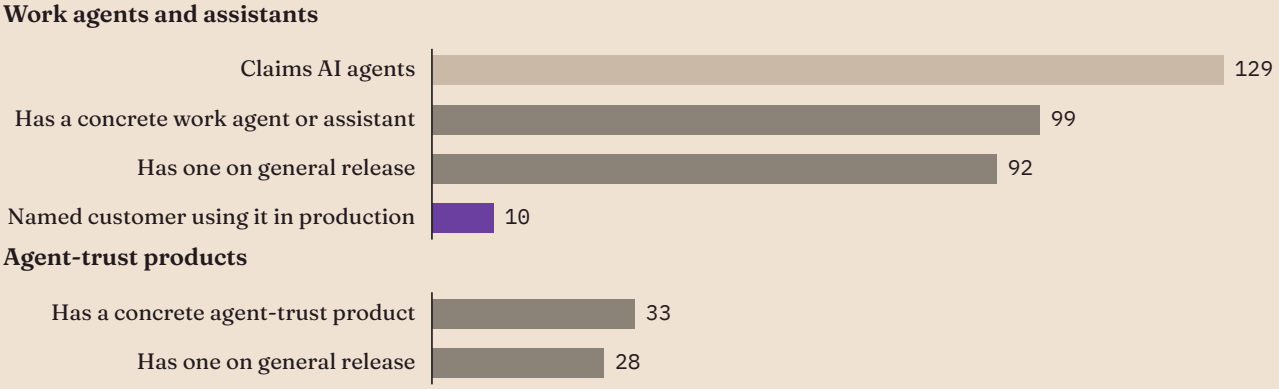
More than four in five vendors that claim agents have a released product

A vendor counts as claiming AI agents if it does so, on a live page, in at least one of four places. Two of them are themes in what the vendor says it offers: 91 vendors, the largest group, describe investigation and review agents (AI agents that investigate alerts, review cases or prepare decisions), and 36 describe identity products for AI agents. In addition, in the past year 66 vendors launched a product whose description mentions agents, and 48 mention agents in their homepage headline or opening line. Many vendors make the claim in more than one place.

We sorted what stands behind the claims into two kinds of product. Work agents are AI agents that do financial-crime, fraud or compliance work themselves, such as clearing alerts or researching a business; we count AI assistants and copilots with them. Agent-trust products verify, detect or secure other firms' AI agents, such as "know your agent" checks or bot and agent traffic controls.

Most claims survive the check (Chart 1). Of the 129 vendors, 99 have a concrete work agent or assistant, and 92, or seven in ten, have one on general release. Leaving assistants out, 88 have a work agent on general release. On the agent-trust side, 33 vendors have a concrete product and 28 have one on general release.

Chart 1. Most agent claims lead to a product; few lead to a named customer.
Number of vendors at each step, among the 129 of the 333 tracked vendors that claim AI agents.



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

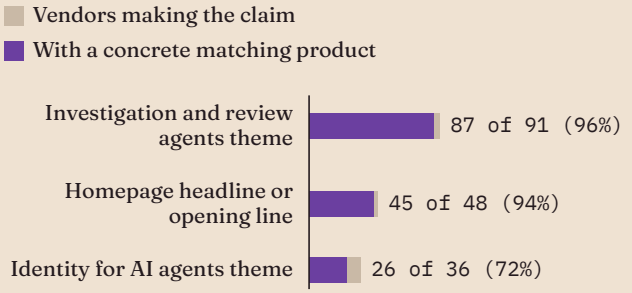
Put the two kinds together and 110 of the 129 vendors have a work agent, assistant or agent-trust product on general release, which is more than four in five. The two groups overlap: 10 vendors have both a work agent and an agent-trust product on general release. At the other end, 10 have no concrete work agent, assistant or agent-trust product.

Where a vendor makes the claim says a good deal about what stands behind it (Chart 2). Of the 91 vendors that describe investigation and review agents, 87 have a concrete work agent or assistant. Homepage claims hold up just as well: 45 of the 48 vendors whose headline or opening line mentions agents have a concrete product of one kind or the other, and all 45 have one on general release. Of the other three, two lack

a confirmed product page and one has no agent product.

Claims about identity for AI agents lead to a product less often. Of the 36 vendors making that claim, 26 have a concrete agent-trust product.

Chart 2. Homepage and investigation-agent claims almost always lead to a product.
Vendors making each claim, and those with a concrete product of the matching kind.



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

Ten vendors name a customer using their agents

The step from product to public proof is where the numbers fall away. Only 10 vendors have a named customer using a work agent or assistant in production, according to public sources checked live on 4 November 2026. The bar is high: the customer must be said to use the agent product in production, and pilots, design partners and plain logo strips do not count.

Most customers of compliance software are not named in public, so the 10 is a floor. Buyers looking for public reference customers will therefore find few. The 10, and the customers their public sources name, are these:

Vendor	Agent product and job	Named customer
Diligent (godiligent.ai)	AI agents: merchant KYC and anti-money laundering (AML) review	Scalapay, Teya
Fenergo	Document Agent: KYC document review	ACT Group
MouseCat	AI investigation agents: fraud investigation	Coinbase
Pindrop	Protect Fraud Assist (an AI assistant): fraud investigation	First National Bank of Omaha (an early customer, per Pindrop)
Roe AI	Investigation and review agents: fraud and AML	eBay, Affirm, Crossmint (per Roe AI)
Socratix AI	AI coworkers: check-fraud review	First Community Bank and Trust
Sphinx	AI Agents: alert review	Conduit
Unit21	AI Agent: first-level alert triage	Green Dot, Nexo, Uphold, Sunrise Bank, Underdog, Zepz
Vivox AI	Screening and due-diligence agents	TELF AG
WorkFusion	Tara: transaction screening review	Crown Agents Bank

Six of the 10 are in the independent venture-backed segment, meaning their latest round was below USD 40M or undisclosed, and five of those six raised that round in 2025 or 2026. Three of the 10 are European: Diligent (godiligent.ai), Fenergo and Vivox AI.

US vendors lead at every step, but European claimers keep pace on work agents

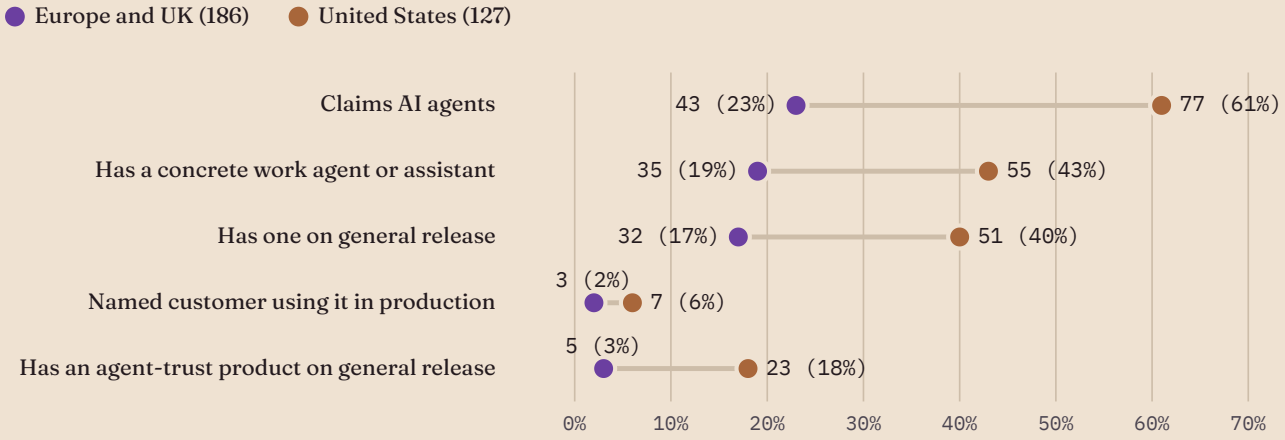
How Fincrime software vendors talk to the market after raising a round found US vendors far ahead on agent launches. Checking the products behind the claims does not close that gap (Chart 3). In the US, 61% of vendors claim AI agents, against 23% in Europe.

The difference carries through to products. 43% of US vendors have a concrete work agent or assistant, against 19% of European vendors, and 40% have one on general release, against 17%. A named production customer is rare on both sides, at 6% of US vendors and 2% of European vendors.

In work agents, most of the gap is already there at the claim stage. Among vendors that claim agents, the share with a work agent or assistant on general release is similar on both sides, at 74% in Europe

and 66% in the US (Chart 4). If products of any status are counted, the shares are 81% in Europe and 71% in the US. Once a European vendor makes the claim, it is about as likely as a US vendor to have a released work agent behind it.

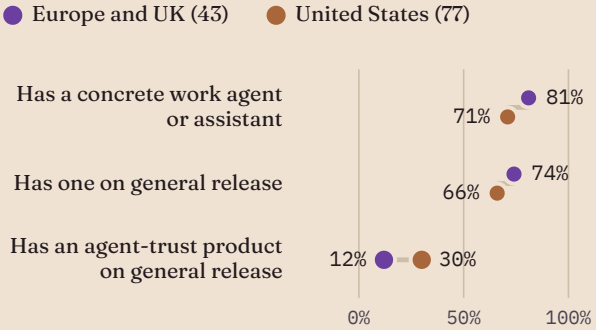
Chart 3. US vendors lead at every step from claim to production.
Share of all tracked vendors at each step, by headquarters: Europe and UK (186), United States (127).



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

Agent trust is a different story. In relative terms, this is where the gap is widest. Almost one US vendor in five, 18%, sells a product to verify or detect other firms' AI agents, against 3% of European vendors. The gap persists among vendors that claim agents: 30% of those in the US have an agent-trust product on general release, against 12% of those in Europe.

Chart 4. Among vendors that claim agents, the regions are close on work agents.
Share of vendors that claim AI agents, by headquarters: Europe and UK (43), United States (77).



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

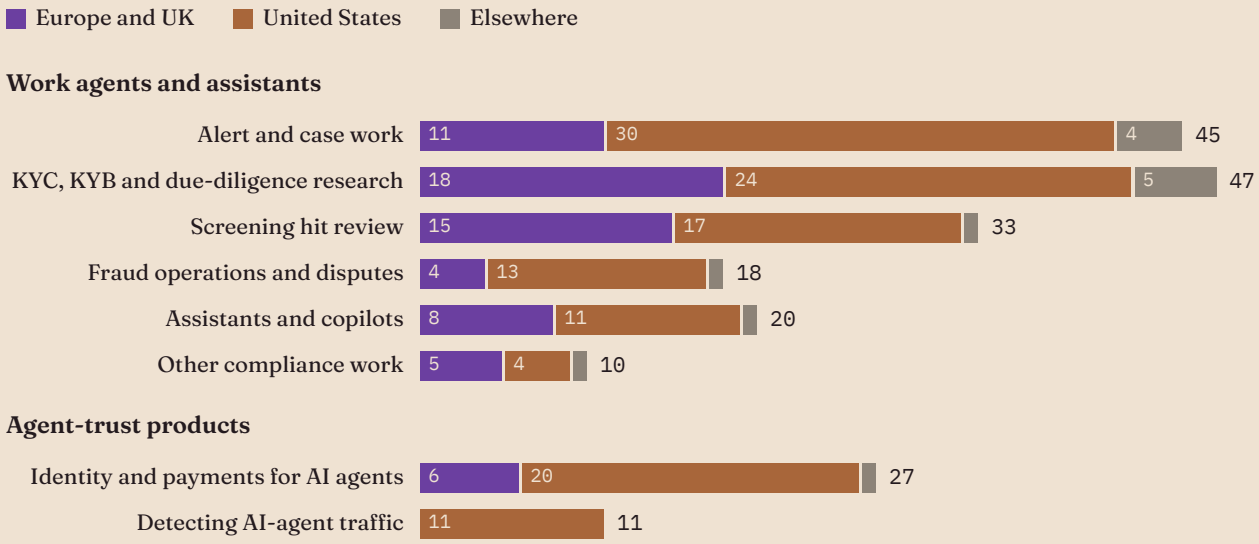
By segment, the picture is flatter. Large, listed or acquired vendors and independent venture-backed vendors are about level on work agents: 31 of 120 large vendors (26%) have one on general release, and so do 47 of 153 venture-backed vendors (31%). Three large and six venture-backed vendors have a named production customer.

Agents mostly do compliance groundwork

KYC, KYB and due-diligence research is the most common job for the products behind the claims, done by 47 vendors (Chart 5). Alert and case work, such as first-level alert review, case summaries and suspicious activity report (SAR) drafting, follows closely at 45.

Screening hit review for sanctions and politically exposed person (PEP) matches comes next, at 33 vendors. Smaller groups are assistants and copilots that work on the vendor's own data (20), and fraud operations and disputes (18). Ten vendors have agents for other compliance work such as quality assurance (QA), regulatory change and reporting. Among agent-trust products, 27 vendors offer identity and payments for AI agents and 11 detect AI-agent traffic on websites. A vendor can do several jobs, so the counts overlap.

Chart 5. Agents mostly do compliance groundwork.
Number of vendors, of the 129 that claim AI agents, with a concrete product doing each job; a vendor can do several jobs.



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

Europe's agents lean towards research and screening

Europe has fewer vendors with work agents, 35 against 55 in the US, but those it has lean towards research and screening (Chart 6). Just over half of them, 51%, use agents for KYC, KYB and due-diligence research, against 44% of US vendors with a work agent. In screening hit review the shares are 43% in Europe and 31% in the US.

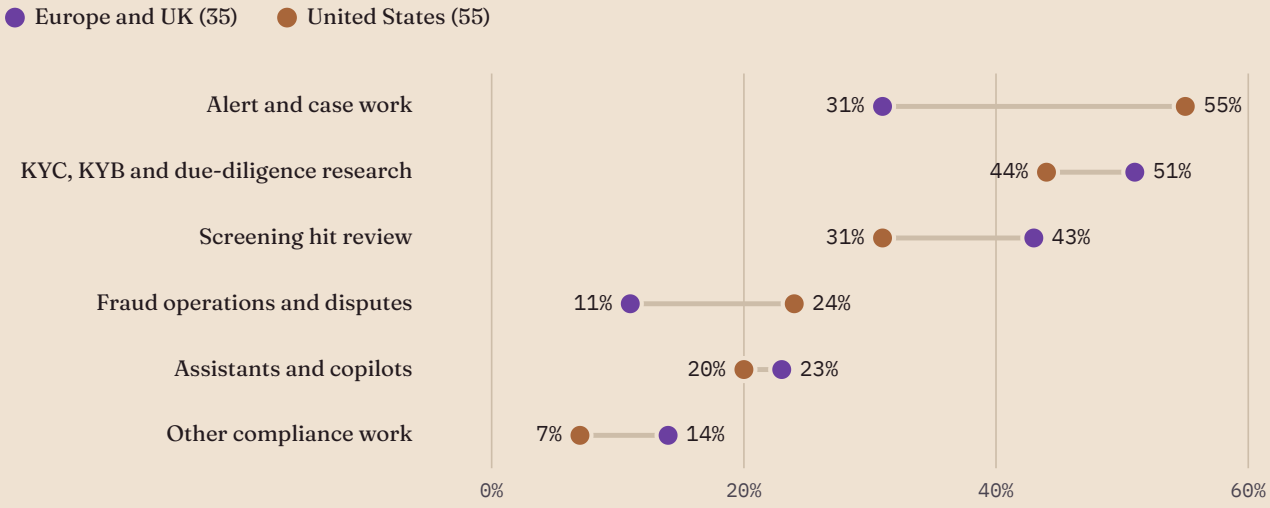
The pattern reverses in alert and case work. Of the European vendors with a work agent, 31% do alert and case work, against 55% in the US. That gap is at the edge of what chance alone could produce. The differences in research and screening could easily be chance. In other words, Europe shows neither a shortfall nor a lead in those jobs.

Chart 6 also shows the smaller jobs. Of European vendors with a work agent, 11% handle fraud operations and disputes, against 24% in the US. In assistants and copilots the European share is the higher of the two, at 23% against 20%. In other

compliance work, Europe's share is 14%, against 7% in the US.

Chart 6. Europe's work agents lean towards research and screening.

Share of vendors with a concrete work agent or assistant that do each job, by headquarters: Europe and UK (35), United States (55); a vendor can have several.



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

Customer due diligence is a core duty under the EU's Anti-Money Laundering Regulation (AMLR), which applies from 10 July 2027 (see the section on rules in *How Fincrime software vendors talk to the market after raising a round*). This data cannot show whether the rule shapes what vendors build.

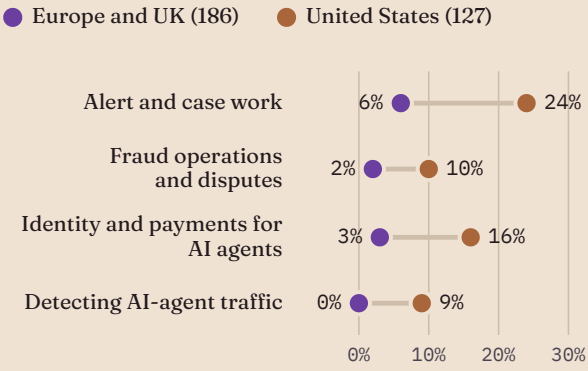
Europe trails most in alert work, fraud and agent trust

Measured against all tracked vendors, including those without agents, Europe trails most in four jobs, where the US share is about four times the European share or more (Chart 7). In alert and case work, 6% of European vendors have a concrete product, against 24% of US vendors. In fraud operations and disputes, the shares are 2% in Europe and 10% in the US.

The two agent-trust jobs show the same pattern. Identity and payment products for AI agents are offered by 3% of European vendors and 16% of US vendors. No European vendor detects AI-agent traffic on websites, while 9% of US vendors do.

Chart 7. Europe trails most in alert work, fraud and agent trust.

Share of all tracked vendors with a concrete product that does each job, by headquarters: Europe and UK (186), United States (127).



Source: vendor records, 333 vendors, checked against live pages on 4 November 2026.

The platforms are building agents too

Banking-technology providers now build and distribute agents themselves. Two examples follow, both from the section on platforms in *How Fincrime software vendors talk to the market after raising a round*.

Fiserv launched agentOS, which will offer its own agents (including one for AML triage) and a marketplace of agents from other companies. FIS launched an AML investigations agent, built with support from Anthropic, and made it available to any US bank, including banks that are not FIS clients.

For smaller vendors, these platforms can be both a route to banks and a new kind of competitor.

Conclusion

Most vendors that claim AI agents can point to a product, but we found only 10 with a named customer using one in production. That count is a floor, since most customers of compliance software are not named in public. In our view, public proof is core to winning the agent race, at least in early messaging, although the data cannot show whether a named customer helps a vendor win business. We also see the regional differences in this report, and the profile of the 10 vendors with a named customer, as a blueprint for vendors planning an agent launch.

If you are planning an agent launch

A product on general release, a generous level in this study, is the norm, at 110 of the 129 vendors that claim agents. We would make the first reference customer part of the launch plan, asking early which customer could be named for which agent and job.

The 10 vendors with a named production customer are not limited to large or US vendors: 6 are independent venture-backed vendors and 3 are European. Their agents do alert, KYC, screening and fraud work, such as first-level alert triage, KYC document review and check-fraud review.

Most of the Atlantic gap in work agents opens at the first step, the claim itself. Once a European vendor claims agents, it is about as likely as a US vendor to have a released work agent behind it. This study starts from claims, so it cannot see agents that vendors do not mention, or say how many such vendors there are. In our view, a European vendor with a released agent that it does not yet mention has reason to say so plainly, and to plan for a named customer; we found such a customer in public for only 3 European vendors.

Agents mostly do compliance groundwork, led by KYC, KYB and due-diligence research. In research and screening, European vendors with a work agent show no shortfall against US ones; Europe's higher shares there could be chance. Europe's thinner areas are alert and case work, fraud operations and agent trust, where the US share of

all tracked vendors is about four times the European share or more. Whether those gaps are an opening or a sign of weaker demand is worth asking before choosing which job an agent will do; the data cannot say.

In relative terms, the Atlantic gap is widest in agent-trust products, sold by 18% of US vendors against 3% of European ones. In our view, a European vendor entering agent trust should expect to explain the category to buyers as well as its own product.

Platforms such as Fiserv and FIS now build agents of their own; we would decide early whether such a platform is a route to banks, a competitor or both.

Whether a vendor should launch agents to its existing customers first, or use them to win new customers, is beyond the scope of this report.

If you are buying

A claim on a vendor's homepage is usually backed by a concrete product, as it was for 45 of the 48 vendors that mention agents in their headline or opening line. Public reference customers are scarce, so a buyer will usually need to ask for evidence of production use. We would put four questions to any vendor that claims agents:

- Is there a concrete product, and is it a work agent, an assistant or an agent-trust product?
- Is it on general release, or only in beta or announced?
- Which job does it do, such as alert review, KYC research or screening hit review?
- Can the vendor name a customer using it in production, as opposed to a pilot, a design partner or a logo?

In our view, the first question matters most for claims about identity for AI agents, which lead to a concrete agent-trust product for 26 of the 36 vendors that make them.

If you invest

Independent venture-backed vendors are about level with large ones on work agents: 31% have one on general release, against 26% of large, listed or acquired vendors. In our view, that makes it fair to hold a portfolio company that claims agents to the same bar as a large vendor: the four questions a buyer would ask, and which customer it expects to name.

What to watch

The EU's AMLR applies from 10 July 2027 and makes customer due diligence a core duty, although this report cannot show whether the rule shapes what vendors build. It is also worth watching how far Fiserv's agentOS marketplace carries other companies' agents to banks. Finally, the count of named production customers is a floor resting on public evidence alone, and it is the figure to recheck first.

Data and methodology

Scope

The report covers all 333 vendors tracked in *How Fincrime software vendors talk to the market after raising a round*. Of these, 129 qualified because they claim AI agents in their coded themes (our record of what each vendor says it offers), in a launch description from the past year or in their homepage headline or opening line, with the claim confirmed on a live page on 4 November 2026. The research records showed 151 such vendors; for 22 of them the agent claim could not be confirmed on the cited page, and they are left out.

A vendor qualifies through any of these four places:

- the coded theme "investigation and review agents", AI agents that investigate alerts, review cases or prepare decisions (91 vendors);
- a launch whose description mentions agents, made in the past year (66);
- the coded theme "identity for AI agents" (36);
- the homepage headline or opening line (48). One vendor that mentions agents only in a label above its headline is not counted among the 48. The opening line is the first paragraph under the headline.

In this report, "European" includes the UK. Vendors that list both a US and a non-US headquarters are counted as US, as in *How Fincrime software vendors talk to the market after raising a round*. The 20 vendors headquartered outside Europe and the US (9 of which claim agents) are included in the totals but not in the regional figures, so the European and US figures do not add up to the totals.

Definitions

- **Work agent:** AI software that carries out a task or a series of steps itself. In this report, work agents are AI agents that do financial-crime, fraud or compliance work themselves, such as clearing alerts or researching a business.
- **Assistant:** a chat assistant or copilot that helps a person. Assistants are counted with work agents.
- **Agent-trust product:** a product that verifies, detects or secures other parties' AI agents, such as "know your agent" checks or bot and agent traffic controls.
- **Named production customer:** a customer that a public source says uses the agent product in production. Pilots, design partners and plain logo strips do not count.
- **General release:** the vendor's own page, or a public report of the vendor's own statement, offers the agent now or says it has launched. A product page counts even when it does not state availability, so this level is generous.

Segments

As in *How Fincrime software vendors talk to the market after raising a round*, "large, listed or acquired" (120 vendors) covers listed vendors, vendors that have changed owner, and vendors whose latest equity round was USD 40M or more. "Independent venture-backed" (153) covers vendors whose latest funding event is an equity round below USD 40M or undisclosed.

How the evidence was gathered and checked

The first pass, verification, tracing, production-customer check, spot check and second coding described below were run on the 151 vendors in the research records, before the live recheck left 129, so the counts given for those steps refer to that larger pool.

- **First pass.** A first coder graded each of the 151 vendors in the research records. The coder used the source-checked records from *How Fincrime software vendors talk to the market after raising a round*, the vendor's homepage, and one fresh web search per vendor on 4 November 2026 for its agent products, their availability and named customers.

- **Verification.** A separate verifier then re-read stored excerpts of every evidence page behind the vendors that the first pass graded as having a concrete product: 129 of those 151, a first-pass count that is a different measure from the 129 vendors whose claims were confirmed live. The excerpts, up to about 1,900 characters each, came from 264 pages. The verifier also read search excerpts and ran one extra search where these did not show the product (36 searches in all). For each product, the verifier recorded its kind (work agent, assistant, agent-trust product or not an agent) and its status (general release, beta, announced or not shown), using the page's own wording. The figures in this report come from this verification.
- **Tracing.** Where a vendor's grade rested on a quote that could not be found in the stored text, the product was checked again against a fresh copy of the vendor's page, or a stored search copy where the page did not load. All 11 products at the eight vendors checked this way were confirmed. Other quotes that could not be traced did not affect any vendor's grade.
- **Source rule.** As in *How Fincrime software vendors talk to the market after raising a round*, social-media posts do not count as evidence. Two vendors whose only agent evidence was a social-media post are counted as having no concrete work agent.
- **Outside financial crime.** Three vendors whose only agent on general release does work outside financial crime (credit spreading, payment exceptions and model development) are counted as having a concrete work agent but, for that reason, none on general release.
- **Production customers.** A third check compared every production-customer claim with its source. Six were removed: one named user was a training centre and not a production customer; in four, the customer quote did not refer to an agent product; and in one, the customers appeared only in a logo strip. Several entries were narrowed to the product the source confirms.
- **Live recheck of every source (4 November 2026).** Every product and customer source was fetched again live on 4 November 2026 and checked against its claim. A product counts only if its page confirmed it that day, through a fresh text copy or, where automated readers were blocked, a capture in a real browser. Every page behind a named company in this report was also captured as a dated screenshot, with the supporting text highlighted where the capture tool could find it. These captures and copies form the evidence file that accompanies this report.
- **Spot check.** After verification, an independent checker graded a random 25 of the vendors with a product on general release against fresh copies of their pages. It agreed on 24 of 25 work-agent grades and all 25 agent-trust grades. The one vendor it could not confirm was graded as having no work agent, and its claim was later not confirmed live. One vendor it confirmed only from a social-media post is also counted as having none, under the source rule.

How the jobs were coded

Each concrete product was coded by the job it does, and a vendor counts for a job only if verification confirmed a product that does it. A vendor can have several. The full counts, with the US and European columns behind Charts 5 to 7, are these:

Job	Vendors	US	Europe
Work agents			
Alert and case work: first-level alert review, case summaries, SAR drafting	45	30	11
KYC, KYB and due-diligence research	47	24	18
Screening hit review (sanctions and PEP)	33	17	15
Fraud operations and disputes	18	13	4
Assistants and copilots over the vendor's data	20	11	8
Other compliance work (QA, regulatory change, reporting)	10	4	5
Agent-trust products			
Identity and payments for AI agents	27	20	6
Detecting AI-agent traffic on websites	11	11	0

Document review agents (14 of the 151 vendors in the first pass) are left out of the job table, because the coders agreed too little on them (Cohen's kappa 0.46).

Reliability

- A second coder repeated the first-pass coding blind for a random 30 vendors. The two agreed on the evidence level for all 30. The published levels differ from the first pass for 6 of these 30, so this agreement measures the first pass, not the published levels.
- Agreement on job groups was good: Cohen's kappa of 0.70 to 1.00 for the groups shown, though several groups had only one to six vendors in the sample, so these values are imprecise.
- Before verification, an initial spot check of 25 first-pass grades against fresh pages confirmed 10. That check led to the full verification described above.

Statistical tests

The job shares among vendors with a work agent rest on 35 European and 55 US vendors with a work agent. The gap in alert and case work is at the edge of what chance would produce (Fisher's exact test, $p = 0.05$). The differences in research and screening are well within chance ($p = 0.52$ and 0.27), so read them as no European shortfall, not as a European lead.

Limits

- One search per vendor will miss some products and customers, so the counts at each evidence level above the claim (concrete product, general release and named production customer) are a floor.
- "General release" is a generous level, for the reason given under the definitions above.
- "In production" needs public evidence that a named customer uses the agent, and most customers of compliance software are not named in public. The count of 10 is therefore a floor.

- Pages change. A product or customer that is no longer shown on its source page, or whose source page could not be loaded on 4 November 2026, counts as not confirmed, even if it was shown earlier. This is why two of the 48 vendors with agents on their homepage are counted without a concrete product; a third has no agent product.

Reference table: the Atlantic gap at each step

Step	Europe (186)	US (127)
Claims AI agents	43 (23%)	77 (61%)
Has a concrete work agent or assistant	35 (19%)	55 (43%)
Has one on general release	32 (17%)	51 (40%)
Named customer using it in production	3 (2%)	7 (6%)
Has an agent-trust product on general release	5 (3%)	23 (18%)

In Chart 4, the shares among vendors that claim agents are the counts in this table divided by the 43 European and 77 US vendors in its first row. All shares in the charts and text are rounded to whole percentages.

How this report was made

The research, coding, verification and analysis were carried out with AI research agents working from public sources, under the author's direction and review.

About this report

The author

Derrick Cramer is the founder of Gossamer, a marketing consultancy in Amsterdam that works with B2B software companies in fintech and AI.

How to cite

Cramer, D. (2026). The Agent Race: how real are fincrime vendors' AI agents? Gossamer, November 2026.

Contact

info@gossamergrowth.com
gossamergrowth.com



Gossamer
gossamergrowth.com