


The Verification Gap

More agents than **evidence**.

What close to 12,000 company websites reveal about AI agent readiness

The verification gap is the distance between what a company says its AI agents do and what it can show. It is the first question a customer, an auditor or a regulator will ask. Close to half the companies mentioning agents have no public answer.

~12,000 company websites analyzed	33% show AI signals	771 publicly mention AI agents	43% of those cannot show how they verify them
---	-------------------------------	--	---



This report describes public website language, not internal practice.
Published September 2026 · progress.com

How to read this report

This report analyzes public website content from close to 12,000 organizations in one software ecosystem, spanning software, IT services, financial services, healthcare, manufacturing, government, retail, education, insurance, logistics, media, energy and legal services. Findings reflect public positioning and messaging. They do not measure internal practices, production deployments, or actual verification maturity.

AI signal

Public website language indicating AI activity. Capability signals – AI-powered products, agents, LLM tooling, machine learning – carry more evidentiary weight than communication signals such as AI content and AI hiring.

AI-active

At least one AI signal on the public website.

AI-native / adopter / curious

AI appears core to the business / publicly describes shipped AI capability / discusses AI without convincing public evidence of shipped capability.

Agent signal

Public language describing AI agents, agentic workflows or autonomous assistants.

Verification signal

Public language indicating steps to determine whether AI is working as intended: monitoring, evaluation, governance, testing, accuracy, observability or human oversight.

What this report measures

Public website language only. A company may verify its AI thoroughly without saying so publicly, and may mention agents it has not deployed. Absence of public language is not proof that a practice is absent.

How the analysis was done

Each site's public content was classified by a language model against the fixed definitions above, with a supporting verbatim quotation required for every positive signal and a confidence grade recorded per site. Sector, headquarters and verification signals come from a second, dedicated pass over the same content. No company is named or ranked anywhere in this report, and no non-public data was used. The full method, dataset description and limitations are in the methodology section at the end of the report.

Five findings

01 AI signals are widespread; AI-native depth is rare.

A third of organizations show AI signals on their public websites, and 74% of those show at least one capability signal. Only 0.7% appear to be AI-native businesses, and 15% of AI-active companies qualify on AI content or hiring language alone.

02 The agent wave is visible, led by the builders.

771 companies publicly mention AI agents or agentic workflows – one in five AI-active organizations. IT services and consulting firms lead at 28%; many are positioning around agent development for everyone else.

03 43% of companies with agent signals are silent on verification.

332 of 771 show no monitoring, evaluation, governance or human oversight language at all. In their agent-related messaging, “autonomous” appears five times more often than evaluation or testing. This is a gap in public evidence, not proof that verification is absent.

04 Governance language outpaces evaluation language.

Among companies with agent signals that do address trust, governance (32%) and human oversight (32%) outpace evaluation (25%). Only one in four mentions measuring whether outputs are accurate, useful or correct.

05 The regulated paradox: accountable sectors show less verification language.

Financial services (29%), insurance (30%) and government (25%) show less public verification language than software (48%). 90 regulated-sector companies mention agents with no public verification language at all.

33%

of organizations show AI signals

771

companies with agent signals

43%

of those show no verification language


25%

mention evaluation

All findings are drawn from public website content, classified per site against fixed definitions with a supporting quotation required for every signal recorded. Absence of public language is not proof that a practice is absent. Full method, dataset and limitations in the methodology section at the end of the report.


AI signals are widespread. AI-native depth is rare.

A third of the companies analyzed show AI signals on their public websites. The shape of that public positioning is a pyramid with a very thin top: only 0.7% appear to be AI-native businesses. Most signals come from established companies publicly describing shipped AI features (20% of all organizations analyzed), while a full third of AI-active companies talk about AI without convincing public evidence of shipped capability.



Categories assigned per site from public content: AI-native (AI appears to be the core business), adopter (publicly describes shipped AI capability), curious (AI talk without convincing public evidence of shipped capability).

Not all signals carry equal weight. **Capability signals** – AI products, agents, LLM tooling, machine learning – indicate a claim about what a company builds or runs. **Communication signals** – AI content and AI hiring – indicate attention rather than capability. 74% of AI-active companies show at least one capability signal; 15% qualify on communication signals alone, and 40% carry only a single signal of any kind.




Companies can show multiple signals; values are shares of AI-active companies, based on public website content. Blue: capability signals. Green: communication signals.

33% of organizations show AI signals	74% of those show a capability signal	0.7% appear AI-native	20% of AI-active companies mention agents
--	---	---------------------------------	---

One in five AI-active companies publicly mentions agents.

771 companies (20% of the AI-active group) publicly describe building, selling, or using AI agents and agentic workflows. Among AI-native companies it is near-universal: 79% mention agents. The builders lead the wave: IT services and consulting firms mention agents at the highest rate of any sector – many are positioning around agent development for everyone else – with software close behind. At the other end, insurance and education barely touch the word.



Share of each sector's AI-active companies that publicly mention agents or agentic workflows. Sectors with sufficient sample size.

771 companies mention AI agents	20% of the AI-active group	28% IT services – highest rate	79% of AI-natives mention agents
---	--------------------------------------	--	--

THE BUILDERS · IT SERVICES & CONSULTING

Agents as a service

28% mention agents – the highest rate of any sector. They construct agents for clients, so one firm's approach may shape many deployments.

THE PRODUCT COMPANIES · SOFTWARE & SAAS

Agents inside the product

22% mention agents, usually as autonomous or agentic features layered onto an established product for existing customers.


THE REGULATED ADOPTERS

Agents under supervision

27% of companies mentioning agents operate in regulated sectors. They have the most to prove and, as the regulated paradox section shows, say less publicly about verification.

What companies say about agents – and what they don't.

The public language around agents leans heavily toward autonomy. In the agent-related messaging of these 771 companies, “autonomous” appears in 23% – five times more often than any mention of evaluation or testing (4%). Company websites describe digital workforces that run 24/7 and operate at scale; the words that would indicate how those systems are evaluated or supervised – accuracy, testing, oversight – trail far behind.



Share of companies with agent signals whose extracted site content contains each term family.


23% include "autonomous"	4% mention evaluation or testing	5x autonomy over evaluation	3% mention human oversight
------------------------------------	--	---------------------------------------	--------------------------------------

Agent messaging is far more likely to emphasise autonomy than evaluation.

Five times more likely, across the 771 companies with agent signals.

43% of companies with agent signals are silent on verification.

For every company with agent signals, we checked its site for four kinds of verification language: monitoring, evaluation, governance, and human oversight. 332 of 771 (43%) show none of the four. Companies with agent signals do mention verification more than the average AI-active company – autonomy raises the stakes, and some of the market appears to recognise it. But the composition is telling: governance and oversight language outpaces evaluation. It is more common to publish responsible-AI language than to describe how agent accuracy is evaluated. This is a gap in public evidence, not proof that verification practices are absent.

**332**

agent-signal companies, no verification signal

25%

mention evaluation – the rarest

32%

mention governance

27%


mention monitoring

Only 1 in 4 companies with agent signals mentions evaluation.

Evaluation – language about measuring whether an agent's outputs are accurate, useful or correct – is the least common verification signal (24.8%). In this ecosystem's public messaging, policy language beats proof language.

More mature AI companies say more about verification.

Verification language tracks maturity closely. 88% of AI-native companies mention it – for them, trust is often part of the product story. Adopters drop to 49%. Among the AI-curious it falls to 23%.



Share of each maturity category showing at least one verification signal.

88% of AI-natives talk verification	49% of adopters	23% of the AI-curious	1 in 5 organizations are AI adopters
---	---------------------------	---------------------------------	--

The more AI does, the more companies need to explain how it is verified.

The AI-curious and AI-adopter companies may be earlier on the same path the natives have walked: as their AI moves from feature to product to autonomous system, verification is the language they will need to learn – because customers, auditors, and increasingly regulators are likely to ask for it.

AI-CURIOUS · 23% TALK VERIFICATION

Talking about AI

Blog posts, hiring, and marketing language. Nothing shipped yet, so nothing to verify – and no vocabulary for it.

AI ADOPTERS · 49%

Shipping AI features

Copilots and AI-powered features inside existing products. Governance language appears; accuracy or evaluation language is less visible.

AI-NATIVE · 88%


AI is the product

Trust is often part of the sales argument, so monitoring, evaluation and oversight are more likely to appear on product pages themselves.

The most accountable sectors talk about verification the least.


Healthcare aside, some sectors that answer to regulators are not the ones saying most about AI verification. Financial services (29%), insurance (30%) and government (25%) all sit well below software (48%). And the gap has teeth: 90 regulated-sector companies mention agents with no public verification language at all.

Why this is a paradox, and what could explain it. These sectors carry the highest accountability for automated decisions, so we would expect them to say the most. Two explanations compete. Their AI may simply be earlier-stage, and the maturity mix below supports this: government, insurance and financial services are the three sectors most weighted toward AI-curious companies, and none of them contains a single AI-native business in this dataset. Or their verification practices may exist but sit in internal documentation, audit files and regulatory submissions rather than on marketing pages. The data cannot separate the two; it can only show that the public evidence is thinner where the stakes are highest.



Share of each sector's AI-active companies showing any verification language. Dark blue: regulated sectors trailing software.

Maturity mix of the sectors discussed




Composition of each sector's AI-active companies. Government (70% AI-curious), insurance (59%) and financial services (52%) carry far more AI-curious companies than software (19%) or healthcare (21%), and none has a single AI-native company in the dataset.

Either way, as AI regulation reaches enforcement, the distance between what these sectors publicly describe and what they publicly evidence is likely to close.

The gap is wider in Europe than in North America.


Among companies with agent signals whose website states a headquarters, the share showing no verification language differs by region: 51% in Europe against 41% in North America. That is the opposite of what the regulatory environment might suggest, and it holds across the largest national samples – Dutch and Italian companies show verification language most often, German companies least among the major European markets.



Regions assigned from headquarters stated on the company's own site.



Verification language by country



Share of each country's AI-active companies showing any verification language. Countries with sufficient sample size. This measures public language, not regulatory compliance.

What we measured, and what we didn't.

- **Source.** Public website content from close to 12,000 organizations in one software ecosystem, analyzed and classified in 2026. Findings describe public website signals across this dataset – not the global market, and not internal company practices.
- **Dataset.** The organizations span software, IT services, financial services, healthcare, manufacturing, government, retail, education, insurance, logistics, media, energy and legal services. 82% of AI-active companies state a headquarters on their own site; North America and Europe are represented in near-equal measure.
- **Method.** Each site's public content was classified by a language model against fixed signal definitions, with a supporting verbatim quotation required for every positive signal and a confidence grade recorded per site. Sector, headquarters and verification signals come from a second, dedicated classification pass over the same content.
- **Definitions.** AI signals: public website language about AI-powered products or features, AI agents and agentic workflows, LLM usage or tooling, machine learning and data science (capability signals), and AI content or AI hiring (communication signals). AI-active: at least one AI signal. AI-native: AI appears central to the core business. AI adopter: publicly describes shipped AI capability. AI-curious: discusses AI without convincing public evidence of shipped capability. Agent signal: public language describing AI agents, agentic workflows or autonomous assistants. Verification signal: public language indicating steps to determine whether AI is working as intended – monitoring, evaluation, governance, testing, accuracy, observability or human oversight.
- **Vocabulary.** Term-family frequencies are measured over the site content analyzed for each company with agent signals, and are therefore floors rather than exact site-wide rates.

Limitations, stated plainly

This report does not measure internal AI practices. It does not verify production deployment. It does not rank or evaluate individual companies, and no company is named or identifiable in it. Absence of public language is not proof that verification is absent; presence of public language is not proof of mature practice. The verification gap described here is a gap in **public evidence**: the distance between agent-related messaging and public language about monitoring, evaluation, governance and human oversight.

Closing the verification gap

The companies ahead on this are not doing more. They are able to show what they do. Three moves close most of the distance, and none of them requires a purchase.

1. Define

Write down what “working correctly” means for one agent or workflow: five to ten concrete criteria covering expected outputs, unacceptable behaviors, edge cases, and when a human should review the result.

2. Measure

Check the agent against those criteria on a fixed set of real or representative cases, and repeat it over time. That turns “it seems fine” into evidence you can track, improve and show a customer.

3. Make it visible

Explain how agents are monitored, evaluated, governed or overseen where customers, auditors and internal stakeholders can find it. One sentence on your website is the minimum. Closing the gap means making verification visible where trust decisions are made.

See where you stand

A two-minute verification readiness check: seven questions on how your team defines, measures, monitors and oversees agent behavior – and which move comes first.

Start free

Progress Agent Engineering: see where agent workflows break, what they cost, and whether outputs are good enough to ship. Built for teams working in .NET, Python and JavaScript. telerik.com/ai-observability-platform

Agents will keep multiplying. Evidence is the choice.



About Progress

Dedicated to propelling business forward in a technology-driven world, Progress (NASDAQ: PRGS) helps businesses drive faster cycles of innovation, fuel momentum and accelerate their path to success. As the trusted provider of the best products to develop, deploy and manage high-impact applications, Progress enables customers to build the applications and experiences they need, deploy where and how they want and manage it all safely and securely. Learn more at www.progress.com

2026 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved. Rev 2026/09

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA
01803, USA
Tel: +1-800-477-6473

facebook.com/progresssw

twitter.com/progresssw

youtube.com/progresssw

linkedin.com/company/progress-software
[progress_sw_](#)