



H2 2026

The State of Agentic AI

From pilots to incident logs. 1,030 people who build, buy, or influence AI systems on what their agents did last quarter, who could prove it, and where trust moved.

nylas.com/agentic-ai-report-h2-2026

Table of contents

Introduction	3
H2 2026 at a glance	4
Research methodology	5
What changed since December, and what we had to ask instead	7
Agent autonomy: Agents are running, and they are making mistakes	9
Trust is graduated, and it moved	10
Agent security and identity: Can you prove what your agents did?	12
How agents get access to email and calendar	14
Adoption and buying decisions: Where agents work	15
What buyers want, and what is in the way	16
The vendor question	18
Roadmap and strategy	19
Conclusion	20
FAQ: about this survey	21
FAQ: terms in this report	22
Glossary	23
Agent Accounts by Nylas	24

Introduction

When we first ran this survey, most of the people answering it were getting started. They were deciding whether to let an agent act on its own, which tasks they would hand over, and what "agentic AI" even meant. The questions matched: would you trust an agent to send a customer email? What is making this hard?

That is not where the category is anymore. In this wave, more than half of respondents run 6 or more agents, and more than half have already cleaned up after one. The premise of the research changed with it. Asking whether people would trust an agent stopped being useful once the agent was already in production, so we asked what it did wrong last quarter, whether anyone could prove what it did last week, and whose credentials it was using.

The first wave fielded in December 2025 with 1,026 respondents. This one fielded 9 months later with 1,030. From here the survey runs twice a year, and the next wave fields in early 2027. A category that rewrites its own vocabulary between quarters is not one an annual review can describe.



7.5%

let agents act with no human approval. Up from 4.2% in December.

nylas

nylas



54%

had an AI agent make a mistake in the past 90 days.

H2 2026 at a glance

7 numbers from 1,030 people who build, buy, or influence AI systems, fielded August 26 to September 12, 2026.

AGENT AUTONOMY

7.5%

act without human approval

54%

had an agent error in 90 days

53%

can fully audit their agents

26%

give agents their own identity

ADOPTION AND BUYING

67%

may replace a vendor with an agent

81%

replaced a tool with in-house build


1 in 3

say agentic AI is already expected

The 3 groups follow the structure of the report: agent autonomy (pages 9 to 11), agent security and identity (pages 12 to 14), and adoption and buying decisions (pages 15 to 19).

Research methodology

Nylas surveyed 1,030 U.S.-based professionals between August 26 and September 12, 2026, through the Centiment research panel. Respondents work in engineering (33%), IT and internal systems (28%), product management (20%), data, AI, or ML (12%), and founder or executive roles (7%). We asked how involved they are in AI decisions: **42% build or implement AI systems**, 37% influence AI or product decisions, 9% provide input without deciding, and 12% describe themselves as end users.



How this wave compares to December

This wave drew on a broader mix of companies than the first. Software and SaaS companies make up 48% of respondents (61% in December), and more respondents came from consulting and professional services and from companies with more than 1,000 employees. We also added questions on what teams now face in production: agent incidents, auditability, agent identity, and vendor replacement.

To keep comparisons fair, this report compares the two waves on only the 8 questions asked identically both times, shown together in the next section. Every other figure describes September 2026 on its own. Details on sample changes, question changes, and how we measured agent identity are in the FAQ under "About this survey" on page 21.


Percentages are rounded to whole numbers except where a tenth matters to a comparison. Bases are noted wherever a question was shown to fewer than the full sample.

Where the sample sits: adoption and auditability by industry

Software and SaaS companies, half the sample, run the most agents and audit them most completely: 60% run 6 or more agents and 61% can produce a full audit log. CX teams run even more agents (69%) but only 46% can fully audit them. Consulting, healthcare, and education sit below average on adoption. The link holds at the respondent level too: across the sample, organizations running 6 or more agents are more likely to fully audit them (60% against 48%). CX is the exception, where 38% of teams running 6 or more agents can fully audit them against 65% of those running fewer, on a small base (42 and 17).

Agent adoption and auditability, by industry

Horizontal: share of organizations running 6 or more agents. Vertical: share that can fully audit agent activity. Bubble size = respondents. Dashed lines = sample average. n=1,030; legal (10) and HR (21) are not plotted because their samples are too small to read.



What changed since December, and what we had to ask instead


8 questions were asked with identical wording and identical options in December and September. They are shown once, below, and they are the only basis for trend claims in this report. Because the sample mix shifted, we treat movements of 5 points or more as meaningful.

The 8 questions asked identically in both waves

December 2025 (n=1,026) vs. September 2026 (n=1,030)

● Dec 2025 ● Sep 2026

TRUST MODES



AUTONOMY TRUST



INTERNAL AUTOMATION



HOW AI TOOLS ARE USED



WORKFLOWS INFLUENCED



MOTIVATIONS



TASKS TRUSTED WITHOUT APPROVAL



VENDOR SWITCHING



Read together, the 8 rows describe a category getting more careful, not more excited. Autonomy is spreading at the edges: the share of teams letting agents act without any approval nearly doubled, and the share gating every action fell 8 points. At the same time, the confident claims softened. Fewer respondents say they fully trust an agent, fewer would trust one to send a customer email, and fewer say they would very likely switch vendors for stronger agentic capabilities. “Meet leadership expectations” nearly doubled as a stated motivation: leadership is asking for autonomy, and the teams running agents are asking for guardrails. Teams that have lived with agents since December are extending them more rope and promising less on their behalf.

The more important change is the one the table cannot show. 5 lines of questioning from December were replaced, because the answers they were designed to produce had become the wrong ones to ask for.

What December asked, what September asked, and why

DECEMBER ASKED	SEPTEMBER ASKED	WHY
Would you trust an agent to perform this task without approval?	In the past 90 days, which of these did an agent do wrong at your organization?	Agents are running. The question is what happened.
What is making agentic AI development difficult?	Can your team produce a full audit log of what agents did last week, and identify which agent took which action?	Accountability replaced difficulty.
Which platforms or frameworks are you using?	How do your agents authenticate to send email or manage calendars?	Identity became an infrastructure decision.
How concerned are you about competitors shipping agentic AI first?	Are you considering replacing a vendor or tool with an AI agent? How many have you replaced already?	Competitive pressure became a build-versus-buy decision.
What do you think “agentic AI” means?	Which customer-facing channels do your agents operate in, and do they have their own visible identity?	Definition debates gave way to deployment facts.

If the survey questions had to change this much between December and September, an annual review would have described a category that no longer existed by the time it published.

AGENT AUTONOMY


Agents are running, and they are making mistakes

Do approval gates actually prevent errors?

54% of respondents say an AI agent did something wrong at their organization in the past 90 days. The most common failure was an agent triggering a workflow the team had to reverse or correct, reported by 31%. A quarter had an agent update the wrong record or data field. 23% had one schedule a conflicting or incorrect meeting, 18% had one send an incorrect or inappropriate email, and 18% had one respond inappropriately to a customer. 42% reported none of these, and 4% weren't sure.

In the past 90 days, what did an agent do wrong at your organization?


Multi-select, n=1,028



The incident rate rises with scale. 45% of organizations running 1 to 5 agents reported an incident in the past 90 days, 64% of those running 6 to 20, and 70% of those running more than 20. 36% run 1 to 5, another 36% run 6 to 20, 14% run 21 to 50, and 4% run more than 50. Only 4% run none.

How many distinct AI agents does your organization run?

n=1,030



AGENT AUTONOMY

One relationship in the data is worth stating carefully. Teams that require human approval for every agent action report incidents at a slightly higher rate (58%) than teams in an intermediate mode (54%) or teams that let agents act freely (52%). The difference is small and the direction could be selection: teams that have been burned may add gates afterward. What it does not show is approval gates preventing errors. It is at least as likely that a review step is where errors get noticed. The flip side is that teams with fewer checkpoints may have incidents slipping by in the background, uncounted.

By September 2026, the category had moved past pilots and roadmaps to incident logs.

Trust is graduated, and it moved

Asked how they build trust with their AI tools today, **59% of respondents describe an intermediate mode in which read operations are auto-approved and write operations require a human.** 30% require approval for every action. 7.5% run what the survey calls YOLO mode, where agents act without approval, up from 4.2% in December. 4% use no trust model at all.

How are you building trust with your AI tools today?

n=1,030




AGENT AUTONOMY

The tasks respondents would let an agent do without a sign-off follow a clear line between internal and external consequences. **66% would let an agent send internal reminders unsupervised**. 58% would let it create tickets or tasks, 56% schedule meetings, and 46% update internal records. The number drops to 44% for sending customer emails and to 27% for changing system settings. 3% would not trust an agent with any of these.

Which tasks would you trust an agent to do without approval?


Multi-select, n=1,030



The risks people name explain the line. Asked the single biggest risk of giving an agent access to a person's email, **35% chose exposure of PII or confidential information** and **30% chose the agent sending something incorrect or unauthorized**. Prompt injection was named by 14%, loss of an audit trail by 11%. 8% said none of these concerned them.

What is the biggest risk of giving an agent access to a person's email?

n=1,029



One caution on the trust figures. **32% of respondents say they fully trust an agent to run autonomously**, yet 136 of those same respondents also say every action requires human approval. Stated trust and configured trust are different measurements. Throughout this report, the trust-mode question is treated as the description of what teams actually do.

AGENT SECURITY AND IDENTITY


Where an agent acts as a person, the log shows the person.

Can you prove what your agents did?

53% of respondents say their team can produce a full audit log of what their AI agents did last week and identify which agent took which action. 37% can do so partially. 4% cannot, and 6% do not audit agent activity at all.

Can your team produce a full audit log of what agents did last week?


n=1,030



Auditability varies with how the agent gets into systems in the first place. Among teams whose agents authenticate with a dedicated identity of their own, **66% report full audit capability**. Among teams whose agents use API credentials with no communications identity, 62% do. The figure falls to 53% for dedicated service accounts, 52% for shared team inboxes, and 51% where the agent operates through OAuth on a human employee's account. Where an agent acts as a person, the log shows the person.

How agents get access, and whether anyone can see what they did

Each dot is one way agents authenticate. Horizontal: share of organizations using it. Vertical: share of those organizations that can fully audit agent activity. n=1,029




AGENT SECURITY AND IDENTITY

This is consistent with how respondents rank what matters to their agentic AI strategy. Security and access controls rank first, with a mean score of 4.76 on a 7-item ranking, followed by reliability and failure handling (4.46), ease of integration (4.17), and compliance readiness (3.98). Observability and debugging (3.68) and cost (3.67) sit in the middle. Native access to communications data such as email, calendar, and messaging ranks last at 3.29.

What matters most to your agentic AI strategy?


Mean score on a 7-item ranking, n=1,030



Communications data access is the item closest to what Nylas sells, and it ranked last. We are publishing the ranking as it came back. The ranking measures what respondents consider strategic, and access to communications data reads to most of them as a given rather than a decision. The audit numbers above suggest it is a decision. How an agent authenticates to a mailbox determines whether anyone can later say what it did there.

26%

of teams say their agents have an identity of their own



42%

say security approvals are the #1 blocker

35%

say exposing PII is the top risk of agent email access

Security is the gate agents have to get through


AGENT SECURITY AND IDENTITY

How agents get access to email and calendar

There is no standard yet for how an agent authenticates to a mailbox or a calendar. 26% of respondents say the agent has its own dedicated identity with an email address and calendar. 21% use a dedicated service account. 14% use API credentials with no communications identity. 12% route the agent through a shared team inbox, and 11% through OAuth on a human employee's account. 8% say no agent has this access yet, and 7% said not applicable or weren't sure.

How do your agents authenticate to send email or manage calendars?

n=1,029




Put differently, at roughly **1 in 4 organizations**, agents operate under credentials that belong to a person or a group of people. For those teams the audit gap described earlier is a design consequence, not an oversight.

Among the 551 respondents whose organizations have not given agents a dedicated identity, the blocker is clear. **56% cite security review**. 37% cite technical complexity, 22% see no clear use case, 17% worry about vendor lock-in, and 15% had not considered it. 16% say nothing has stopped them.

What has stopped you from giving agents a dedicated identity?

Multi-select; organizations without a dedicated agent identity, n=551



ADOPTION AND BUYING DECISIONS


Can the customer tell it's an agent?

Where agents work

Agent activity remains weighted toward internal operations. Asked which workflows they have seen agentic AI influence, **75% of respondents named IT and internal operations**. Customer support followed at 54%, sales workflows at 51%, and project and delivery management at 48%. Meeting follow-ups and documentation follow at 44%. Finance and reporting (33%) and recruiting and HR (27%) trail.

Which workflows have you seen agentic AI influence?


Multi-select, n=1,030



Where agents face customers, email is the dominant channel. Among the 813 respondents asked about customer-facing channels, 78% named email, 74% chat, 44% SMS, and 38% phone. 7% said none yet.

Which customer-facing channels do your agents operate in?

Multi-select; respondents whose agents are customer-facing, n=813



Those customer-facing agents are split on whether the customer can tell. Of 756 respondents with a customer-facing agent, **39% say it has its own visible identity**, such as its own email address. 26% say it acts invisibly through a human employee's account, and 32% describe a mix of both.

ADOPTION AND BUYING DECISIONS

Does your customer-facing agent have its own visible identity?

Respondents with a customer-facing agent, n=756




Asked where an agent identity would be most valuable, respondents chose **internal operations and reporting (48%)** and **customer support (41%)** ahead of meeting scheduling (34%), executive or administrative assistance (32%), and sales outreach (25%).

What buyers want, and what is in the way

Asked what matters most when choosing an agentic AI platform or API, with up to 4 selections, respondents put reliability first at **65%** and **security controls second at 61%**. Performance followed at 50%, ease of integration at 46%, and cost at 41%. Compliance readiness (32%), observability or debugging (27%), integration coverage (27%), and vendor roadmap or support (14%) rounded out the list.

What matters most when choosing an agentic AI platform or API?

Multi-select, up to 4, n=1,030




ADOPTION AND BUYING DECISIONS

The blockers are the same 2 themes seen from the other side. **Security approvals (42%)** and **integration complexity (40%)** lead the list of what is preventing implementation today, followed by organizational hesitation (28%), unclear ROI (25%), and engineering bandwidth (24%). Capacity to build is no longer the main constraint; approvals and integration are. Fear of vendor lock-in was named by 15% and lack of a strong use case by 13%. A new option this wave, “Nothing, we aren’t blocked,” was chosen by 23%.


What is preventing implementation today?

Multi-select, n=1,030



Motivations are practical. Improving response time and speed leads at **67%**, then improving data quality across systems at 56% and reducing manual work at 53%. Reducing costs was named by 40%. Launching differentiated features (26%) and meeting leadership expectations (25%) are nearly tied.

Which best describes your relationship with agentic AI right now?



ADOPTION AND BUYING DECISIONS


Is agentic AI a reason to leave a vendor?

The vendor question

Agents have become a reason to leave a vendor. 67% of respondents are considering replacing a vendor or tool with an AI agent: 26% say they are considering replacing many, and 41% say 1 or 2. 17% say they extend existing tools instead, 11% have no plans, and 5% weren't sure. Separately, 81% have already replaced at least one vendor or tool with in-house tooling, and 32% have replaced 6 or more; that question did not ask whether the replacement was an agent.

Are you considering replacing a vendor or tool with an AI agent?

n=1,030



Asked whether stronger agentic AI capabilities would persuade them to switch tools or vendors, **38% say they are very likely to switch and 55% say they would possibly switch**, which signals openness rather than intent. Among those open to it, reliability is the leading reason to switch (23%), followed by security and compliance (21%) and better integrations (19%). Lower cost (13%) and more autonomy (10%) trail. Better UX (8%) and vendor trust and support (6%) barely register.

Would stronger agentic AI capabilities persuade you to switch tools or vendors?

n=1,030



The “very likely” share is 8 points lower than in December, and we would rather report that than combine the 2 positive answers into a single figure. The willingness to switch is broad. The urgency has cooled.


ADOPTION AND BUYING DECISIONS

Roadmap and strategy

Among the 924 respondents who are involved in product or AI strategy at their company, **68% say agentic AI is on the product roadmap and 29% say it is under consideration**. 3% say no. Asked how important it is to overall product strategy, 21% say critical and 59% say very important; 17% say somewhat, and 4% say low priority or not relevant.

How important is agentic AI to your product strategy?

Respondents involved in product or AI strategy, n=924



On timing, a third of respondents say agentic AI has already become a baseline expectation in their market. 37% expect it within 12 months, 21% within 2 to 3 years, and 4% within 5 to 10 years. 5% do not expect it to reach that point. As individuals, 42% describe themselves as actively building agentic AI, 35% as evaluating it seriously, and 17% as experimenting informally; 6% are watching, skeptical, or not interested.

When will agentic AI become a baseline expectation in your market?

n=1,030




The “already” option is new this wave, so the timing figures are not comparable to December’s. In this wave, **95% of respondents expect agentic AI to become a baseline expectation in their market at some point**; only 5% say it never will. The disagreement is only about when.

Conclusion

The category has moved from “will this work” to “what do we do when it doesn’t.” More than half of the organizations in this survey have already answered that question at least once, by reversing a workflow, correcting a record, or apologizing to a customer. About half can fully audit what their agents did. 7.5% trust their agents enough to skip the approval step entirely, up from 4.2% in December.

Extending autonomy safely depends on being able to answer 3 questions about any agent: who it acted as, what it did, and under which approval rule. Those answers depend on infrastructure choices made before the agent ever sends an email, and the audit figures in this report suggest many teams have not yet made them deliberately.

The next check-in fields in early 2027. The questions will change again.



FAQ: about this survey

How did the sample change from December?

Software and SaaS companies fell from 61% to 48% of respondents. Consulting and professional services rose from 9% to 13%, and companies with more than 1,000 employees rose from 21% to 26%. Because the mix shifted, we treat only changes of 5 points or more as meaningful in the December comparison.

Which figures can be compared to December?

Only the 8 questions in the "What changed since December" section, which were asked with identical wording and options in both waves. Trend claims elsewhere in this report draw only on those questions. 11 questions are new this wave. Several others carried over but were reworded or shown only to people who gave a certain earlier answer, so their December figures are not comparable and we do not show them.

Why is the agent identity figure 26%?

Two questions touch on agent identity, and they measure different things. When asked how their agents authenticate to send email or manage calendars, 26% said the agent has its own dedicated identity. When asked whether their organization gives agents their own email address and calendar, 46% said yes. Many in that 46% described a service account, API credentials, OAuth on an employee's account, or a shared inbox when asked how the agent authenticates. Those setups give an agent an address but not an identity of its own. We report 26% because the authentication question describes how agents actually connect, and the identity breakdowns in this report are built on it.

Did Nylas write the identity questions with Agent Accounts in mind?

Nylas makes Agent Accounts, and we drafted the identity questions before it launched in June 2026. We published the reasoning behind them between the two waves, in [The infrastructure race behind agentic AI](#) (February 18), [The production stack for agentic AI](#) (February 26), and [How agentic AI adoption is scaling from the inside out](#) (March 4), all on the Nylas blog.

FAQ: terms in this report

What is an agent identity?

A set of credentials, typically an email address and calendar, that belongs to the agent rather than to a human employee, a shared inbox, or a generic service account. In this survey, 26% of organizations say their agents authenticate this way.

What is a trust mode?

The rule an organization applies to agent actions. Approval mode requires a human sign-off on every action. Intermediate mode auto-approves read operations and requires approval for writes. YOLO mode lets the agent act without approval. 59% of respondents run intermediate mode.

What counts as an agent incident in this report?

Any of 5 events in the prior 90 days: an agent sent an incorrect or inappropriate email, scheduled a conflicting or incorrect meeting, updated the wrong record, responded inappropriately to a customer, or triggered a workflow the team had to reverse. 54% reported at least one.

What does an agent audit log contain?

At minimum, a record of every action an agent took, with a timestamp and the identity of the agent that took it. The survey asked whether teams could produce this for the prior week and attribute each action to a specific agent; 53% said they fully could.

Who ran the survey?

Nylas commissioned and analyzed the survey. Fielding was handled by Centiment between August 26 and September 12, 2026, with 1,030 completed responses from U.S.-based professionals in engineering, IT, product, data, and leadership roles.

Glossary

Agentic AI. AI systems that plan and execute multi-step tasks with some degree of autonomy, using external tools and APIs to act inside operational systems.

Agent identity. Credentials that belong to an agent rather than to a person: its own email address, calendar, and authorization grants.

Approval mode. A trust configuration in which every agent action requires human approval before it executes.

Intermediate mode. A trust configuration in which read operations are auto-approved and write operations require approval.

YOLO mode. A trust configuration in which agents act without human approval. The term is the survey's and is common among practitioners.

Audit log. A record of the actions an agent took, when, and under which identity, sufficient to reconstruct its activity after the fact.

Service account. A non-human account created for software to authenticate with, typically shared across systems or agents rather than dedicated to one.

Communications data layer. The infrastructure that gives software and agents normalized access to email, calendar, and contacts across providers.

NYLAS AGENT ACCOUNTS

Give every agent an identity of its own

Half of this report is about one gap: agents are running, and teams cannot always say which agent did what, or whose credentials it used. At 1 in 4 organizations, agents act through an employee login or a shared inbox. Where they do, the audit trail shows a person, not the agent.

Nylas Agent Accounts give your AI agents their own email and calendar through a single API: a complete communication identity owned by your software, not borrowed from a person. Each agent gets its own address and calendar, created with one API call, with rate limits, allowlists, and send quotas enforced per account and Workspaces to group agents under one set of policies. In this survey, teams whose agents have a dedicated identity were the most likely to fully audit what those agents did.

66%

of teams whose agents have a dedicated identity can fully audit them, vs. 51% where agents use an employee's login and 62% on API credentials alone.

1 in 4

organizations run agents on borrowed credentials today.

56%

of teams without a dedicated agent identity say security review is what stopped them.

See how it works at nylas.com/products/agent-accounts



Agent Accounts is generally available. The identity questions in this survey were written before its launch; the reasoning behind them is in the three posts cited in the methodology.