aptitude research

The Agentic Moment:

HR'S OPPORTUNITY TO LEAD THE NEXT ERA OF WORK

Authored by:

Madeline Laurano

Sponsored by:

eightfold.ai

The Agentic Moment:

HR'S OPPORTUNITY TO LEAD THE NEXT ERA OF WORK

The AI conversation in HR has matured. Most large organizations have moved past debate about whether AI belongs in talent functions. The reality is that they are already using it. We found that 69% of companies are using AI in Talent Acquisition and 74% are using AI in HR.

69%

of companies are using
AI in Talent Acquisition

74%

are using
AI in HR

But the next shift is already underway, and it operates on a completely different level.

Agentic AI (systems that can reason across multiple steps, use tools autonomously, and execute decisions without waiting for human instruction) is beginning to move from research and early pilots into operational deployment. This is not a small improvement on existing AI capabilities. It is a change in what AI is responsible for doing. Previous systems produced recommendations. Agentic systems act on them.

For HR leaders, this distinction is significant. The talent function is defined by good or bad decisions including who gets hired, who gets developed, who gets opportunity. Agentic AI does not just accelerate those decisions. It begins to make them or enable how a human makes them. That creates potential but also requires clarity on where, how, and under what conditions agentic AI is appropriate.

This report is designed to provide that clarity. It examines where agentic AI stands today, what is deployable, what is still maturing, and what the path to move forward looks like. It is written for leaders who understand the opportunity and now need a rigorous framework for capturing it responsibly.

The organizations that move thoughtfully in this moment will not simply run more efficient HR functions. They will build a talent operating model that is faster, fairer, and more adaptive than anything traditional processes alone can deliver. That is the opportunity agentic AI puts in front of HR leadership right now.

SECTION 1:
WHY AGENTIC? WHY NOW?

The story of AI in HR is not one shift. It is four. Each wave expanded what was possible and revealed the ceiling of what that generation of technology could do on its own. Rules-based automation eliminated friction and standardized process. Predictive AI surfaced risk, ranked candidates, and introduced intelligence to decisions that had previously been made on instinct. Generative AI created content at scale and changed the speed at which organizations could communicate with candidates. In each case, the value was real. So were the limits.

Automation required humans to make every decision. The system moved things, people still chose. Predictive AI was advisory. It told humans what might happen, but humans still decided what to do. Generative AI was reactive. It responded to prompts but did not initiate, plan, or execute. Every wave augmented human work. None changed what the human was responsible for doing.

Agentic AI is the fourth wave, and it operates differently. An agentic system does not wait for instruction at each step. It is given a goal and pursues it — reasoning across multiple steps, connecting to and operating across external systems, making decisions, and executing actions autonomously. When an action does not produce the expected result, it recognizes that and adjusts. The distinction between systems that advise and systems that execute is not subtle. It changes the nature of what HR is responsible for managing, the architecture of HR technology, and the governance requirements organizations must have in place before they scale.

When asked how familiar they are with agentic AI as a concept, HR leaders in this survey said:

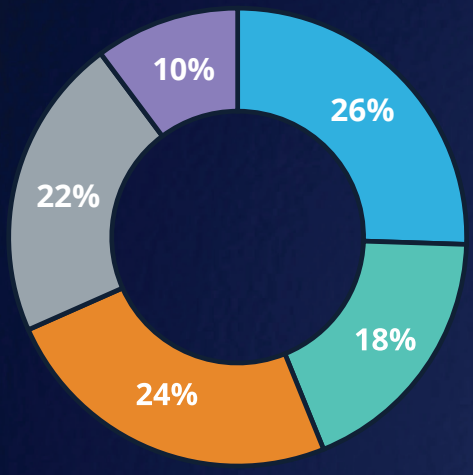

- **26%** Not familiar — this is the first time I'm seeing a clear definition
- **18%** I've heard the term but have not seen a working definition
- **24%** Somewhat familiar — I understand the concept broadly
- **22%** Very familiar — I can explain it clearly to others in my organization
- **10%** Expert — I am actively evaluating or deploying agentic systems in HR

The market is full of confusion about what agentic AI actually is. Vendors apply the term to chatbots. Analysts apply it to copilots. Some use it to describe any AI that sends an automated follow-up. A chatbot that answers HR policy questions is not agentic. It responds to inputs, it does not pursue goals. A copilot that drafts job descriptions is not agentic. It generates content, it does not execute workflows. A dashboard that surfaces attrition risk is not agentic. It reports, it does not act. The test is straightforward: does the system wait to be told what to do at each step, or does it determine that itself? If it waits, it is not agentic. Organizations that cannot make this distinction cannot evaluate vendors accurately, make sound adoption decisions, or set realistic expectations for what their investments will deliver. It is the first step and maybe the most important. We asked companies how familiar they are with agentic AI as a concept and only 31% are very familiar or an expert.

The familiarity gap matters. Organizations that cannot define agentic AI cannot evaluate it, govern it, or deploy it responsibly. Awareness without clarity produces misaligned vendor selection, underestimated implementation requirements, and governance frameworks that are designed for the wrong category of technology. Getting the definition right is not a semantic exercise. It is the starting point for every decision that follows.

Where the Market Stands Today

Aptitude Research has found that the foundational systems HR relies on today were not built for the intelligence era, let alone the agentic one. Only 28% of companies are currently satisfied with their ATS decisions. Eighty-two percent say their ATS has significant gaps. Investment continues: 34% of companies are increasing ATS investment in 2025, and one in four are replacing their system entirely. But spending more on the same category of tool does not close the capability gap. What changes the equation is the nature of the technology itself.

28%
of companies are currently satisfied with their ATS decisions

82%
say their ATS has significant gaps

34%
of companies are increasing ATS investment in 2025

Use of AI:

Most organizations in this survey describe themselves as active AI users but concentrated in earlier stages of the maturity curve. Predictive tools, generative content support, and workflow automation are the most common deployments. Agentic execution remains early-stage. The workflows where AI is most active today (content creation, sourcing support, scheduling assistance) reflect the generative and predictive focus of current investment. Only 12% are using Agentic AI. The workflows that require autonomous, multi-step execution across systems remain largely unaddressed. That is where the distance between current state and future state is greatest.

When asked to describe their organization's current overall use of AI in HR, respondents said:

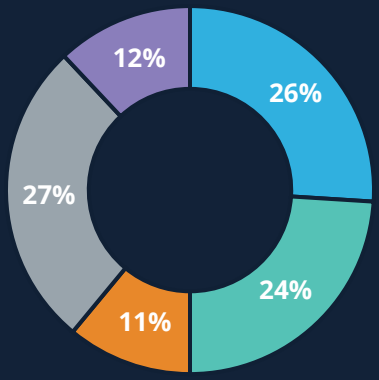

- **26%** No AI tools actively deployed in HR
- **24%** Basic automation only — scheduling, notifications, workflow triggers
- **11%** Predictive / analytics AI — attrition modeling, performance signals, sourcing rankings
- **27%** Generative AI — content creation, job descriptions, candidate communication
- **12%** Agentic AI — systems that reason, plan, and execute autonomously across workflows

When asked which HR workflows they are currently using AI in, respondents identified:

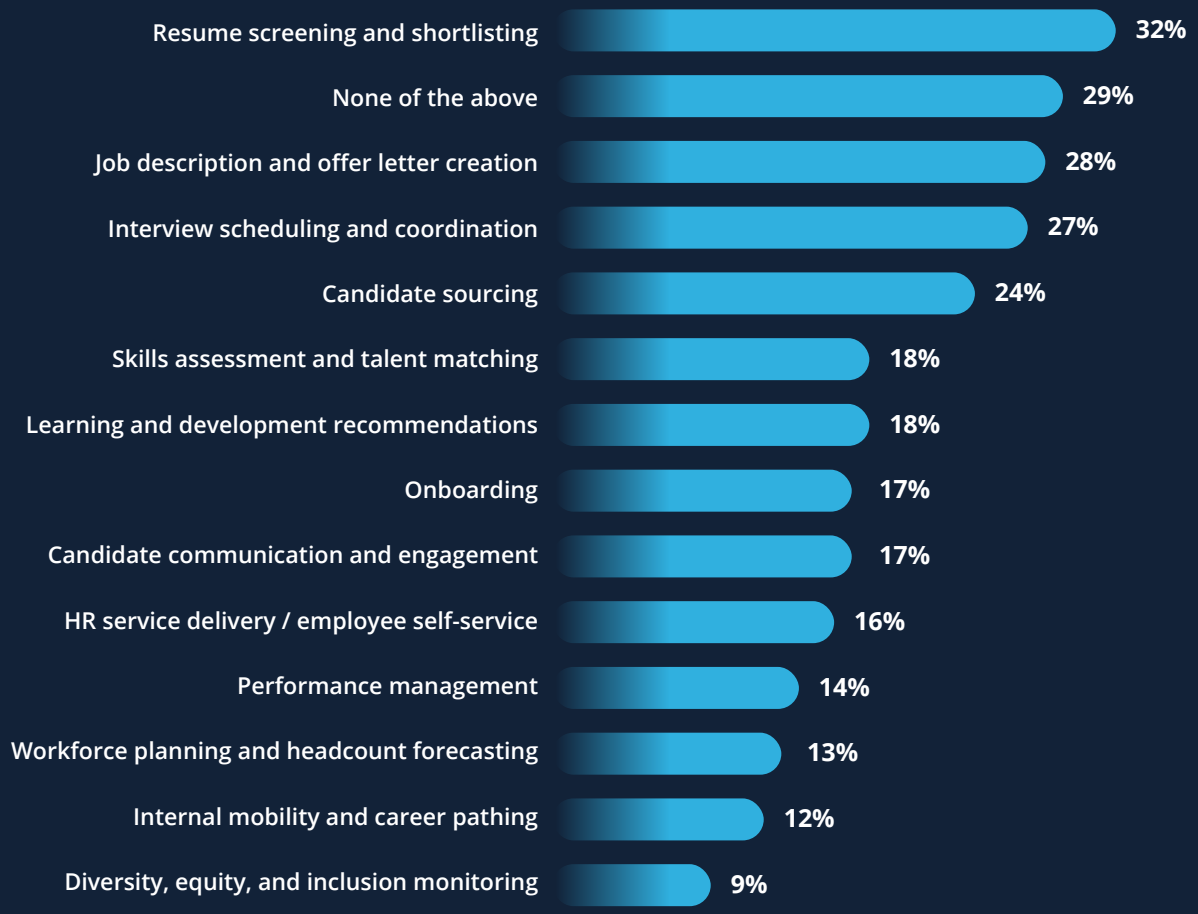

The workflow data reveals a consistent pattern: AI is most active where tasks are discrete and content-based, and least active where execution requires coordination across multiple systems and stakeholders. That gap between where AI is deployed today and where agentic execution is needed most defines the opportunity in front of HR functions right now.

Satisfaction with AI:

We asked about satisfaction and found that over half of companies are satisfied or very satisfied. Satisfaction with AI in HR remains measured. Most organizations see operational value but do not always describe it as transformative. Satisfaction and measurement is also basic at this point. Some companies are only looking at how frequently employees use AI rather than the impact they are seeing.

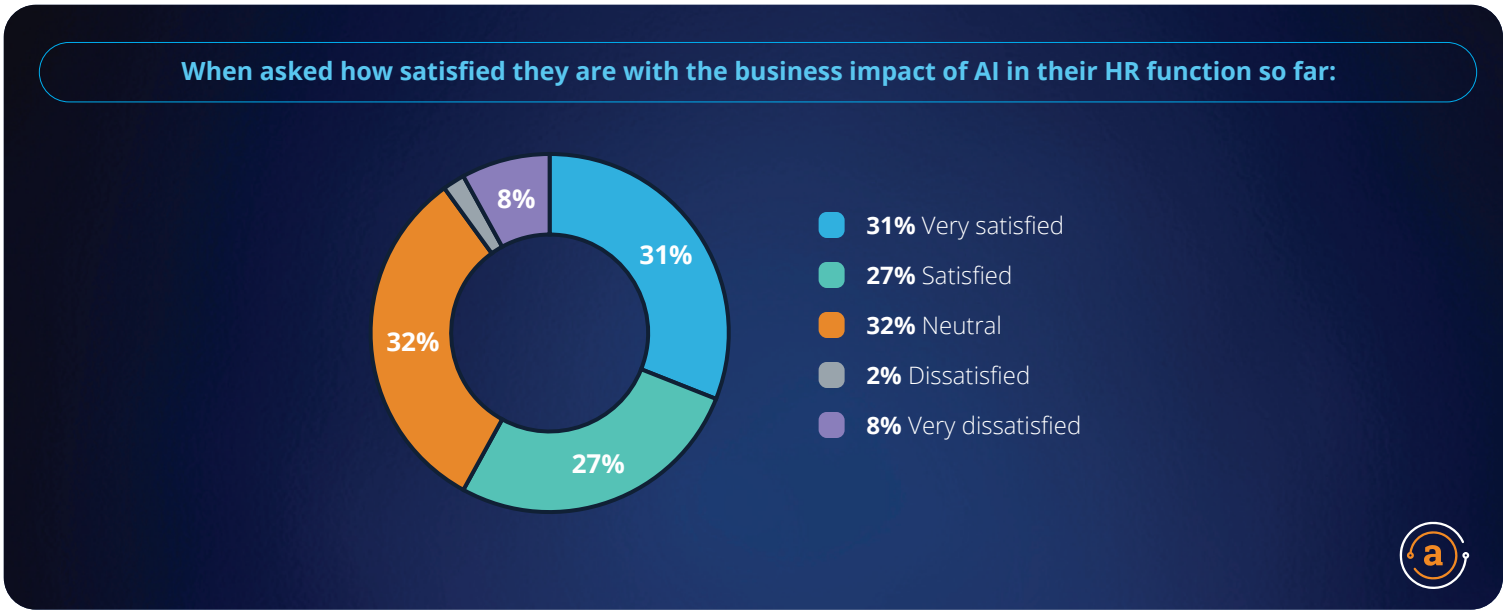

Vendor Landscape:

Vendor strategy is also revealing. Point solutions deliver value within a narrow workflow. Platforms create the conditions for orchestration and the ability to coordinate AI action across the full talent lifecycle. Agentic AI requires orchestration. A system that autonomously sources, screens, schedules, and communicates across an entire hiring pipeline cannot function if its components cannot talk to each other. Organizations managing fragmented, best-of-breed estates face a compounding readiness challenge that vendor selection alone cannot solve.

When asked which statement best describes their organization's approach to AI vendor strategy in HR:

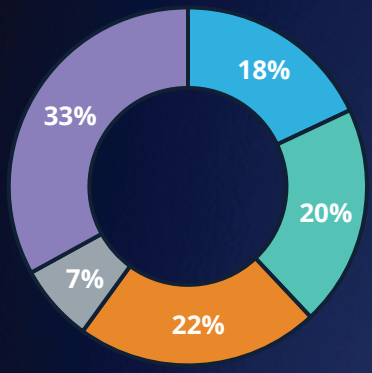

- **18%** Single platform — we have consolidated onto one AI-enabled HR platform
- **20%** Best-of-breed — we use multiple point solutions that each do one thing well
- **22%** Hybrid — a core platform plus selected point solutions
- **7%** Build internally — we are building proprietary AI capabilities
- **33%** No clear strategy yet

Ownership of AI:

Ownership of AI strategy within the HR function shapes how fast organizations move and whether governance keeps pace with deployment. Organizations where the CHRO or CPO owns AI strategy directly tend to deploy faster and operate with clearer governance frameworks. Where ownership is undefined or sits outside HR, adoption is slower and more reactive. It is also important to note that for 44% of companies, there is no clear ownership or unsure ownership of AI.

When asked who has primary ownership of AI strategy within their HR function:

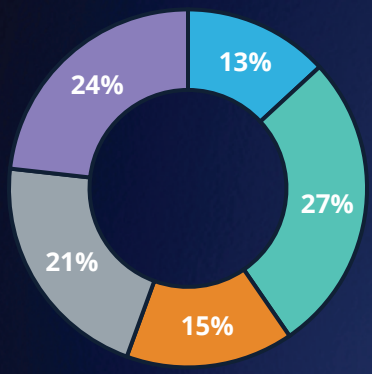

- **13%** CHRO / CPO owns it directly
- **27%** Dedicated HR technology or People Analytics team
- **15%** Shared with IT / CIO's office
- **21%** No clear ownership yet
- **24%** Unsure

The Opportunity Ahead

The case for agentic AI in HR is not primarily an efficiency argument, even though efficiency is a real and measurable outcome. The deeper argument is structural. Agentic AI changes what is possible, not just how fast organizations can do what they already do. Modern talent operations are running at a scale and speed that human-mediated processes were not designed to handle. High-volume hiring environments process thousands of applications for a single role. Internal mobility programs need to match vast numbers of data points (skills, aspirations, role requirements, team dynamics) across large employee populations. Workforce planning requires real-time modeling of variables that change daily. The answer to a scale problem cannot be adding headcount. Agentic AI provides a structurally different path.

A recruiter who no longer coordinates interviews, chases hiring manager feedback, and manages pipeline status updates is not a faster recruiter. They are a different kind of professional. They are one whose time is available for the work that actually requires human judgment: building relationships, assessing fit, making consequential decisions about people. Aptitude Research found that one in three companies are already considering or would consider agentic AI for talent acquisition. That number will accelerate as early deployments produce visible and measurable results. When asked where companies would like to use agentic, screening and interviewing was top of the list (34%) followed by high-volume and frontline hiring (31%) and internal mobility (21%).

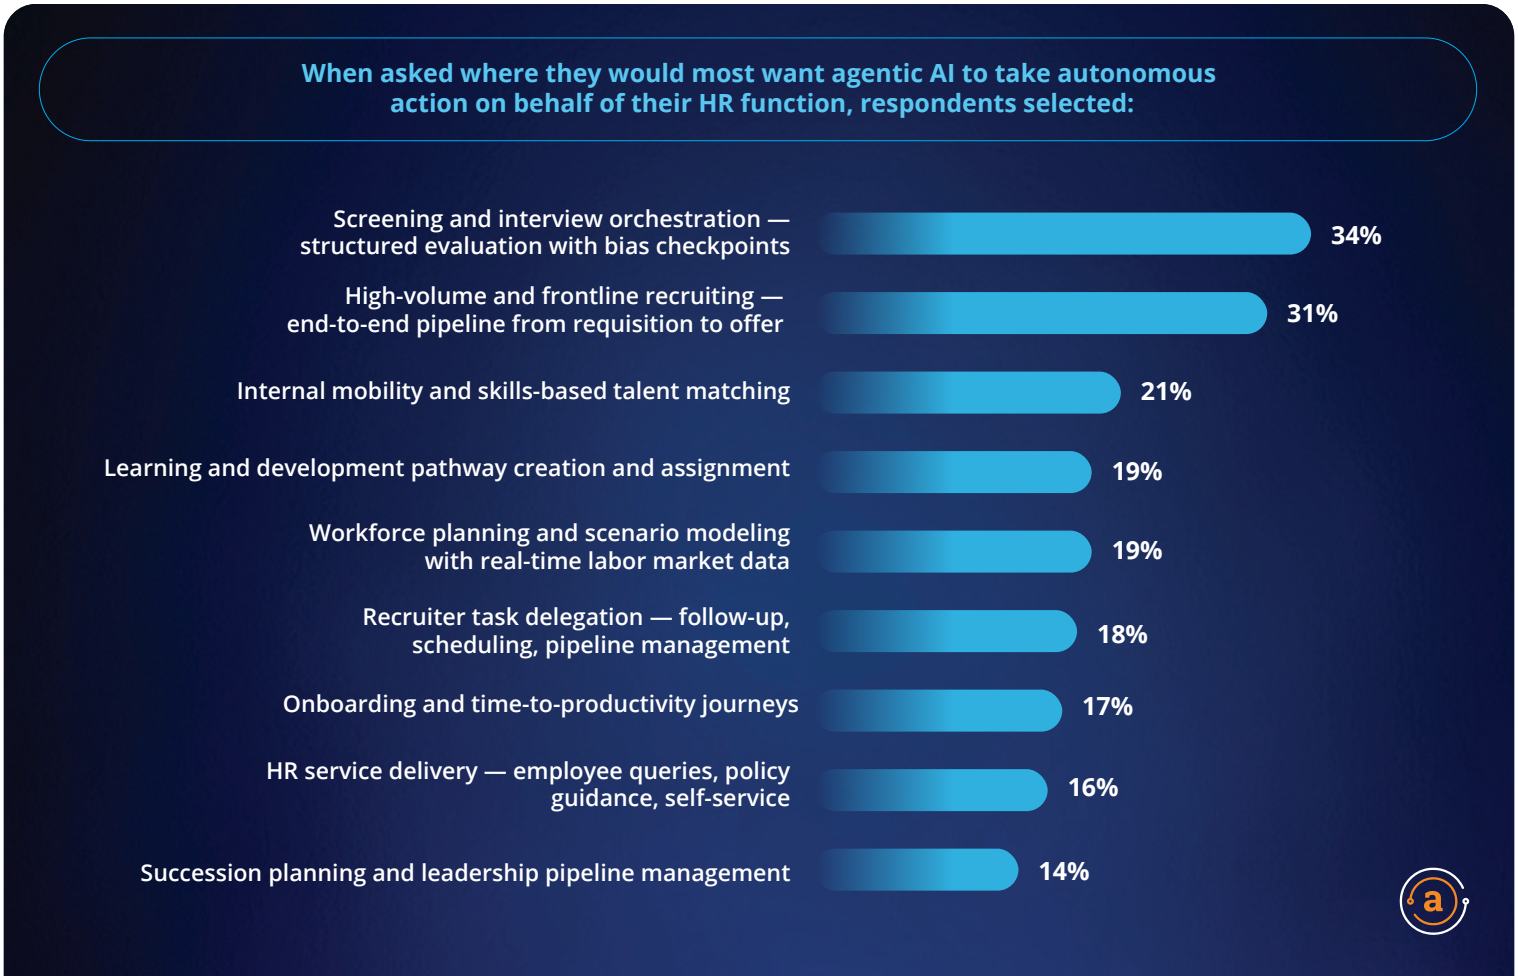

When asked why more companies are not adopting an agentic strategy, we found that only 38% of respondents agree that their HR team has the skills needed to work alongside agentic systems. The organizations that move thoughtfully and decisively now will not simply run more efficient HR functions. They will build a talent operating model that is faster, fairer, and more adaptive than anything human-mediated processes alone can deliver and that advantage will compound. Better data produces better decisions. Better decisions produce better outcomes. Better outcomes feed back into systems that get smarter over time.

Part Two of this report examines exactly where that value materializes by use case, with a consistent focus on what the agent does, what humans stop doing, and what improves as a result.

PART 2:
WHERE AGENTIC AI DELIVERS VALUE
THE USE CASES

Understanding what agentic AI is represents the first step. Understanding where it delivers value is what enables HR leaders to build a deployment strategy that is specific, sequenced, and grounded in real outcomes. Each use case in this section is examined through a consistent lens: the problem it solves, what agentic AI does and what that means for human work, the outcomes it delivers, and the governance requirements specific to that use case. Each is designed to stand alone as a reference for HR leaders focused on a particular part of the talent lifecycle.

High-Volume and Frontline Hiring

THE PROBLEM

High-volume hiring is where the gap between current capability and operational need is most visible. Retail, logistics, healthcare, manufacturing, hospitality — any organization managing large frontline workforces faces the same structural challenge. Application volumes that no human team can process at speed without sacrificing quality. Candidate drop-off that increases with every hour of delay. Time-to-fill numbers that directly affect business operations, customer experience, and revenue. The traditional response has been to add headcount, layer on point solutions, or accept that speed and quality are in permanent tension. None of these responses scale.

Aptitude Research found that 48% of companies cite scheduling as their top ATS gap. In high-volume environments, that gap compounds — every coordination delay is a candidate who accepted an offer elsewhere. The organizations most exposed to this problem are also the ones least able to absorb it: frontline-heavy industries where turnover is high, hiring is continuous, and the cost of a vacant shift is immediate and measurable.

48%
of companies cite scheduling as their top ATS gap.

WHAT AGENTIC AI DOES

A well-architected agentic system manages the full pipeline from requisition creation through offer generation — sourcing candidates, screening applications, scheduling interviews, communicating with candidates, and routing decisions — autonomously and simultaneously across thousands of open roles. It does not simply move candidates through predefined workflow stages. It reads context, adjusts outreach based on candidate behavior, surfaces candidates from existing talent pools before posting externally, flags pipeline gaps before they become critical, and escalates to a human recruiter only when judgment is genuinely required.

What humans stop doing:
manually reviewing every application, chasing candidates who have gone dark, tracking pipeline status across disconnected systems, coordinating interview logistics for hundreds of candidates simultaneously.

What humans focus on instead:
evaluating finalists, advising hiring managers, managing exceptions, and refining the criteria and feedback loops that make the system more effective over time.

THE OUTCOMES

When asked what would have the greatest impact on hiring outcomes, 1 in 3 companies stated screening and ranking. Organizations deploying agentic AI in high-volume hiring report significant reductions in time-to-fill, lower cost-per-hire, and improved candidate experience — not because the process is faster in a narrow sense, but because it is more consistent, more responsive, and less dependent on recruiter bandwidth at any given moment. The candidate who applies at 11pm receives the same quality of engagement as the one who applies Monday morning. The pipeline does not stall because a recruiter is out of office. Screening decisions are not different depending on who happens to be reviewing.

When asked how they are currently managing high-volume or frontline hiring at scale:

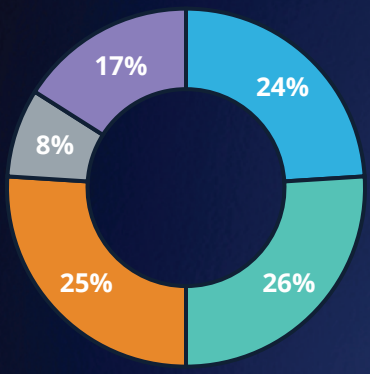

- **24%** Primarily manual — recruiters handle most steps
- **26%** Partially automated — some workflow automation in place
- **25%** AI-assisted — AI surfaces candidates but humans decide and act
- **8%** Agentic — AI orchestrates the full pipeline with minimal human intervention
- **17%** We do not do high-volume or frontline hiring

Fairness is an underappreciated outcome in high-volume environments. When criteria are defined clearly and applied consistently, the idiosyncratic variation that enters human review processes is reduced. That is not an argument for removing humans from hiring decisions. It is an argument for ensuring the criteria humans have defined are applied the same way every time, at every volume level.

Screening

THE PROBLEM

Screening is one of the most time-intensive activities in talent acquisition and one of the most inconsistent. Aptitude Research found that 72% of recruiters identify screening and evaluation as where they spend the most time. It is also where bias is most likely to enter the process — not because recruiters intend it, but because manual review of large applicant volumes under time pressure produces variable, inconsistent judgments. The resume has remained the primary screening artifact for decades despite abundant evidence that it is a poor predictor of job performance. Keyword matching surfaces candidates who know how to write a resume, not candidates who can do the job.

The consequences compound. Qualified candidates who do not use the right terminology are screened out before a human sees them. Candidates who are skilled at resume presentation advance regardless of actual fit. And the recruiters doing this work spend the majority of their time on a task that is high-volume, low-consistency, and deeply susceptible to fatigue-driven shortcuts.

WHAT AGENTIC AI DOES

An agentic screening system evaluates candidates against a skills-based model of what the role actually requires — drawing on the job description, historical hiring data, performance outcomes of past hires, and real-time labor market intelligence. It surfaces candidates who would have been overlooked by keyword matching, flags candidates whose profiles suggest high potential beyond their current title, and applies the same evaluation criteria to every applicant regardless of volume. The agent does not make the hire. It ensures that the humans who make the hire are working from a shortlist that is representative, skillsbased, and consistently generated.

What humans stop doing:

manually reviewing every application against a checklist, making inconsistent pass-fail decisions based on resume format or title, spending the majority of screening time on the lowest-value part of the process.

What humans focus on instead:

evaluating the shortlist the agent surfaces, refining the criteria that drive screening decisions, and engaging with candidates who have already been validated as relevant.

When asked which screening or interview tasks they would be most comfortable delegating fully to an AI agent:

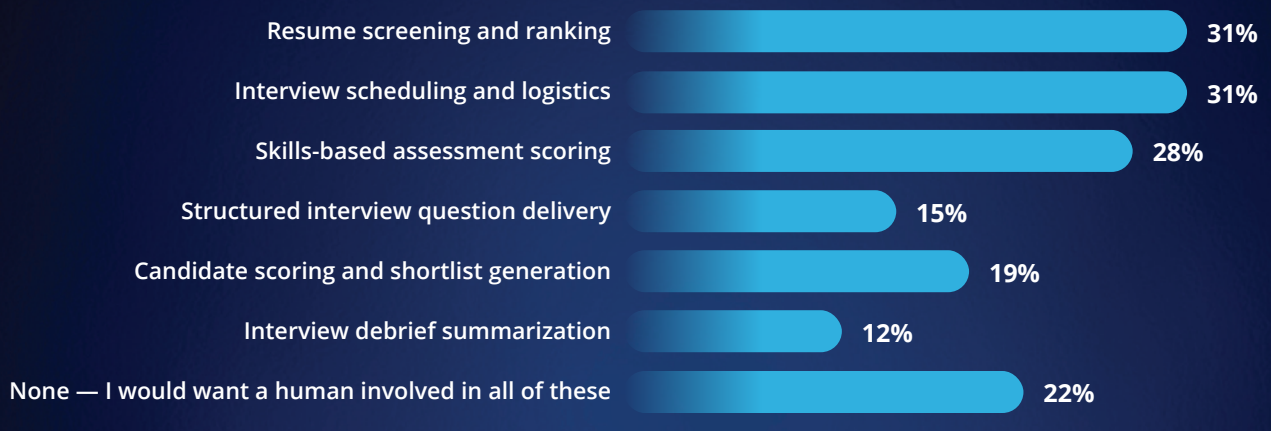

aptitude research

THE OUTCOMES

Agentic screening reduces the time recruiters spend reviewing unqualified applications, improves the consistency and quality of shortlists, and surfaces talent that keyword-based systems miss. Organizations report improvements in diversity of shortlists when skills-based screening replaces resume-based review because the evaluation is grounded in capability rather than credentials, titles, or institutional names. The downstream effect is a hiring process that produces better decisions, not just faster ones.

AI Interviewer

THE PROBLEM

The initial screening interview is one of the most resource-intensive steps in the hiring process and one of the least strategic. A recruiter spending thirty minutes on a first-round screen with an unqualified candidate is thirty minutes unavailable for one who is. Multiply that across hundreds of open roles and thousands of applicants, and the math becomes unsustainable. The problem is compounded by speed: research consistently shows that top candidates are off the market within days of beginning their search. An organization that takes three to five days to schedule a first screening call is not competing effectively — regardless of how strong the role or the employer brand.

Consistency is the second structural problem. When multiple recruiters conduct first-round screens for the same role, they do not ask the same questions, weight responses the same way, or arrive at scores through the same logic. The candidate's experience and the organization's ability to compare candidates meaningfully varies based on which recruiter happened to take the call. The third problem is capacity. High-volume environments face a ceiling that human teams cannot break through without adding headcount — a constraint that does not scale and does not recover quickly when surges hit.

WHAT AGENTIC AI DOES

An AI interviewer conducts an initial structured interview autonomously — engaging candidates immediately after they apply, asking role-specific questions, evaluating responses against defined criteria, and delivering a scored, ranked shortlist to the recruiter without a human conducting the first conversation. Candidates complete the interview on their own time, on their own device, by video, audio, or text. The recruiter receives evaluated, prioritized results when they log in — not a pile of applications requiring manual triage. What makes this genuinely agentic is the combination of autonomous execution, structured evaluation, and adaptive interaction. The system does not play back a recorded question and collect a response. It conducts a conversation, evaluates against a skills-based model of the role, and produces an output that is explainable and auditable.

What humans stop doing:

conducting every first-round screening call, manually reviewing every application before a human conversation occurs, losing candidates to competitors during multi-day delays between application and first contact.

What humans focus on instead:

reviewing scored shortlists, engaging with candidates who have already been evaluated and prioritized, and spending interview time on the questions that genuinely require human judgment.

When asked whether their organization is currently using or evaluating AI-powered interviewing technology:

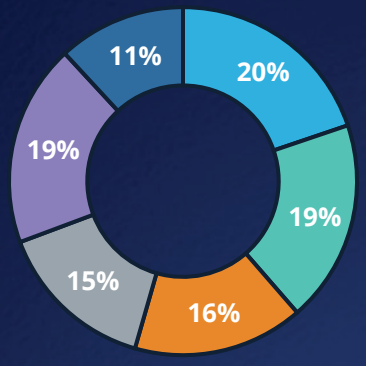

- **20%** Yes — actively deployed and in use
- **19%** Yes — currently in pilot
- **16%** Currently evaluating vendors
- **15%** On our roadmap but not yet in evaluation
- **19%** No — and we have no current plans to adopt it
- **11%** No — we have considered it and decided against it

When asked which conditions would need to be true for them to trust an AI interviewer to conduct and score candidate interviews autonomously:

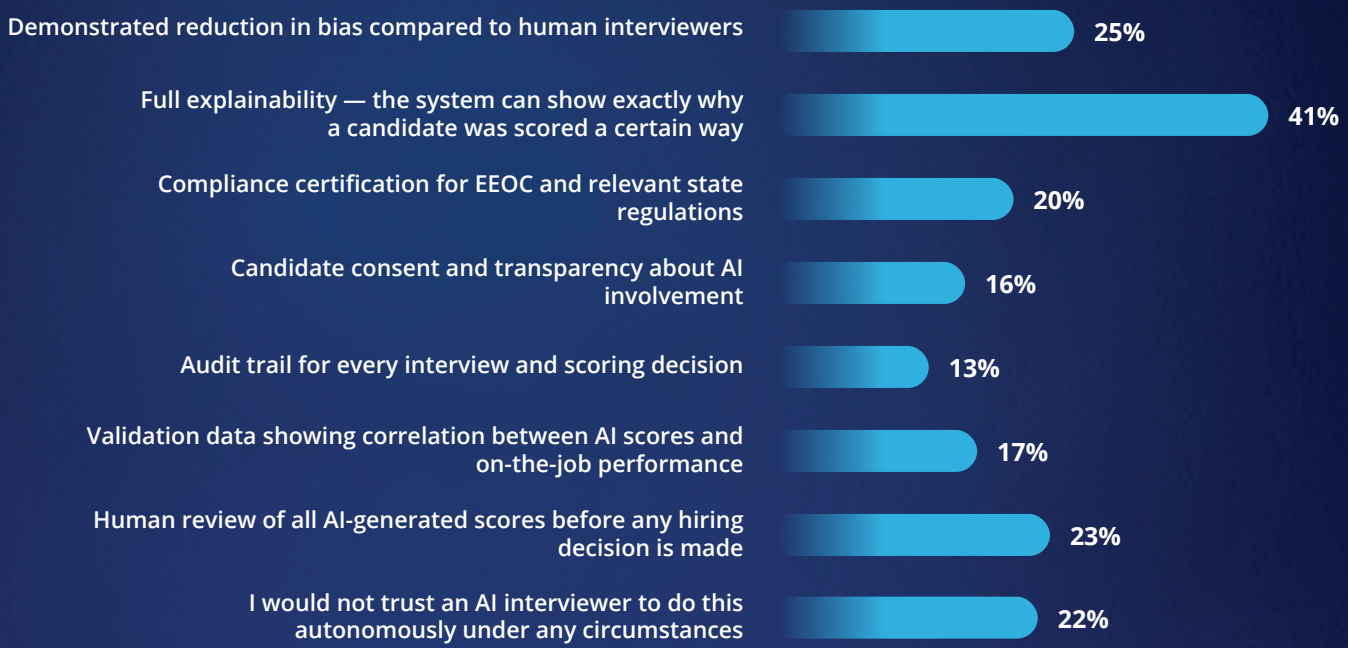

THE OUTCOMES

Organizations deploying AI interviewers at scale report an 85% reduction in time-to-hire, time to first interview under two hours, and four hours saved per position through automated screening conversations and summaries. Beyond speed, consistency improves significantly — every candidate for a given role answers the same questions, evaluated against the same criteria, with no variance introduced by individual recruiter behavior. Candidate experience improves too: immediate engagement signals organizational responsiveness, and candidates report higher satisfaction with hiring processes where they can complete initial interviews on their own schedule.

Internal Mobility

THE PROBLEM

Internal mobility is one of the most consistently underperformed areas in talent management. Organizations invest heavily in external hiring while significant talent sits inside the organization — invisible to the systems and leaders who could deploy it. Most internal mobility programs rely on manager networks, informal referrals, and occasional internal job postings. These mechanisms systematically favor employees who are already visible, already connected, and already in the right conversations. They leave substantial talent undiscovered and create structural inequality: employees with supportive managers have mobility options, employees with protective managers do not.

The consequences are measurable. Employees who cannot see a path forward leave. Organizations that cannot redeploy internal talent efficiently pay for external hiring they did not need. Skills gaps widen not because the organization lacks capability but because it lacks visibility into the capability it already has. Internal mobility is not a retention program. It is a core talent strategy that most organizations are executing poorly because they do not have the data infrastructure or decisionmaking tools to do it well.

WHAT AGENTIC AI DOES

Agentic AI changes internal mobility by making skills visible and acting on that visibility. A skills-based talent intelligence system continuously maps what every employee knows, what they are capable of, what they are interested in, and what opportunities exist across the organization. An agentic layer on top of that intelligence proactively surfaces matches — recommending employees for open roles, stretch assignments, or development opportunities — without waiting for a manager to nominate someone or an employee to find a posting on their own. The agent does not gatekeep. It surfaces. Humans decide.

What humans stop doing:

relying on manager networks to identify internal candidates, manually reviewing employee profiles against open roles, losing talent to competitors because internal opportunity was invisible.

What humans focus on instead:

evaluating agent-surfaced matches, having career conversations with employees who now have visible pathways, and using talent intelligence to make better workforce planning decisions.

When asked what the biggest barrier to better internal mobility is in their organization:

THE OUTCOMES

Organizations with strong agentic internal mobility programs report higher retention, lower external hiring costs, faster time-tofill for internal roles, and improved employee engagement — because employees can see that the organization is invested in their growth. The diversity impact is significant: when opportunity surfaces through skills-based matching rather than manager networks, employees from underrepresented groups who have been historically less visible in informal networks gain equitable access to internal opportunity. Skills-based internal mobility is also the foundation of workforce agility — the ability to redeploy talent in response to changing business conditions without waiting for external hiring cycles.

Organizations that have deployed recruiter workflow optimization AI report significant reductions in time-to-fill, improvements in hiring manager satisfaction, and higher recruiter retention because the job becomes more sustainable and more meaningful. The ROI case is straightforward: the inputs are measurable (hours per recruiter per week on administrative tasks), the outputs are measurable (how that time is redirected and what it produces), and the improvement is visible quickly. This is one of the most mature and immediately deployable use cases for agentic AI in HR, the tasks recruiter workflow AI handles are also the tasks where agentic AI performs most reliably today.

Onboarding

THE PROBLEM

Onboarding is the most underdiscussed use case for agentic AI in HR — and one of the clearest opportunities for measurable impact. The relationship between onboarding quality and new hire outcomes is well established. Employees who experience structured, personalized onboarding reach full productivity faster, report higher engagement, and are significantly less likely to leave in their first year. Most organizations know this. Most still deliver a largely standardized onboarding experience designed around administrative completion rather than individual acceleration.

The gap between what onboarding should be and what it typically is comes down to capacity. Building a genuinely personalized experience for every new hire — one that adapts to their role, location, prior experience, and specific skills gaps — requires more human attention than most HR and manager teams can consistently provide. The result is a generic process that feels transactional, a ramp time that is longer than it needs to be, and a first-ninety-days experience that is the highest-risk period for new hire attrition yet receives the least individualized support.

WHAT AGENTIC AI DOES

An agentic onboarding system assesses each new hire's profile at the point of offer acceptance — their skills, role requirements, location, and start date — and builds a tailored onboarding journey before their first day. It sequences content, surfaces the right connections, assigns role-specific learning, checks in at key milestones, and adjusts the journey based on how the new hire is progressing. The manager is freed from logistics and prompted for the moments where their involvement actually matters: the first conversation, the thirty-day check-in, the first meaningful project assignment.

What humans stop doing:

manually building onboarding checklists for each new hire, tracking completion across disconnected systems, delivering the same generic content to every employee regardless of role or experience, following up on incomplete tasks.

What humans focus on instead:

welcoming new hires into the team, having meaningful development conversations, identifying early engagement signals, and accelerating the point at which the new hire is contributing at full capacity.

When asked how much of their current onboarding experience is personalized to the individual:

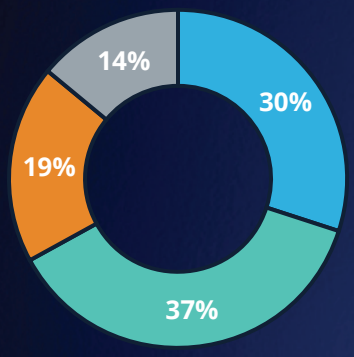

- **30%** Fully personalized — every new hire gets a tailored experience
- **37%** Partially personalized — role-based tracks exist but limited individual tailoring
- **19%** Mostly standardized — same process for everyone with minor variations
- **14%** Fully standardized — one-size-fits-all regardless of role or background

THE OUTCOMES

Organizations with agentic onboarding programs report faster time-to-productivity, higher ninety-day retention, and improved new hire engagement scores. Because the onboarding journey is structured and logged, the relationship between onboarding interventions and productivity milestones becomes visible and actionable — creating a feedback loop that improves the program over time. The early warning capability is one of the most valuable outcomes: an agentic system that monitors engagement signals during the first ninety days can surface indicators before a recoverable situation becomes an exit, triggering a manager conversation or HR intervention at the moment it can still make a difference.

Workforce Readiness

THE PROBLEM

Technology readiness and workforce readiness are not the same problem. An organization can have the right platform, the right data infrastructure, and the right governance framework in place and still fail at agentic AI adoption because the people responsible for working alongside these systems do not yet know how to. HR leaders consistently underestimate this gap. Internal AI literacy and skills scored 3.27 out of 5 in this survey's readiness assessment, above governance and executive alignment, but still well short of what sustained deployment requires. And unlike infrastructure, which can be procured, workforce AI readiness has to be built.

The gap is not primarily technical. Most HR practitioners do not need to understand how large language models work. What they need is something more specific and more learnable: the ability to evaluate AI outputs critically rather than accepting them at face value, to recognize when a model's recommendation reflects a pattern in data rather than a judgment about a person, to design the human review processes that sit around agentic systems, and to communicate to candidates and employees what role AI played in a decision and why. These are skills that require deliberate development and not a one-time training module, but an ongoing practice that evolves alongside the systems themselves.

WHAT AGENTIC AI CAN DO

The organizations advancing most confidently on agentic AI are investing in AI literacy as a core HR competency, not an IT adjacency. They are embedding it in recruiter onboarding, in performance expectations for people operations roles, and in the leadership development curriculum for HR business partners. They are building internal communities of practice where practitioners share what they are learning about working with specific systems. And they are creating feedback loops between the people using agentic tools and the people configuring them — because the practitioners closest to the work are often the first to notice when a model's outputs are drifting from what the business actually needs.

Only 38% of respondents in this survey agree that their HR team has the skills needed to implement and manage AI effectively. That number will not improve through vendor selection alone. The organizations that close the workforce readiness gap will do so by treating AI capability as a talent development priority — and by recognizing that the competitive advantage of agentic AI is only as durable as the human judgment applied around it.

Part Three examines how to build toward the former and avoid the latter.

ONLY

38%

of respondents in this survey agree that their HR team has the skills needed to implement and manage AI effectively.

PART 3: WHAT ORGANIZATIONS NEED TO CONSIDER

Most HR leaders reading this report are not starting from zero. They have AI in their stack. They have vendor relationships. They have pilots underway or on the roadmap. The question is not whether to pursue agentic AI it is how to sequence it, what to build toward, and what to put in place before scaling. This section provides a practical framework for moving from awareness to action.

Assess Where You Are

Before selecting a use case or evaluating a vendor, HR leaders need an honest assessment of organizational readiness across the dimensions that determine whether agentic AI can be deployed successfully.

When asked to rate their organization's readiness for agentic AI across key dimensions

(1 = Not ready at all, 5 = Fully ready)

The dimension with the lowest score is the starting point. Not because it is the most interesting problem, but because it is the one most likely to limit everything else. Data quality limits matching accuracy. Integration gaps limit orchestration. Governance gaps limit trust. Addressing the foundation is not a delay. It is the fastest path to sustainable deployment.

Barriers

The barriers HR leaders identify are consistent: data readiness, integration complexity, governance maturity, and unclear ROI. Each has a specific remedy.

Data readiness.

Audit what you have before buying what you think you need. Most organizations have more skills data than they realize — it is fragmented, not absent. Start by mapping where employee and candidate data lives and what it would take to connect it.

Integration complexity.

Agentic AI requires systems that talk to each other. If your HR technology estate is fragmented, the highest-value investment before any agentic deployment is the integration architecture that enables orchestration. A platform that connects data across the talent lifecycle is the prerequisite, not the optional add-on.

Governance.

You do not need a comprehensive AI governance framework before you start. You need a clear policy for the specific use case you are deploying — what the agent can decide, what requires human review, how decisions are documented, and who is accountable. Build governance use case by use case and let it grow into a framework.

ROI clarity.

The ROI case for agentic AI is not complicated — it is just underspecified in most organizations. Time-to-fill, cost-perhire, recruiter hours on administrative tasks, new hire ninety-day retention: these are the metrics that change. Define the baseline before deployment so the improvement is visible and attributable.

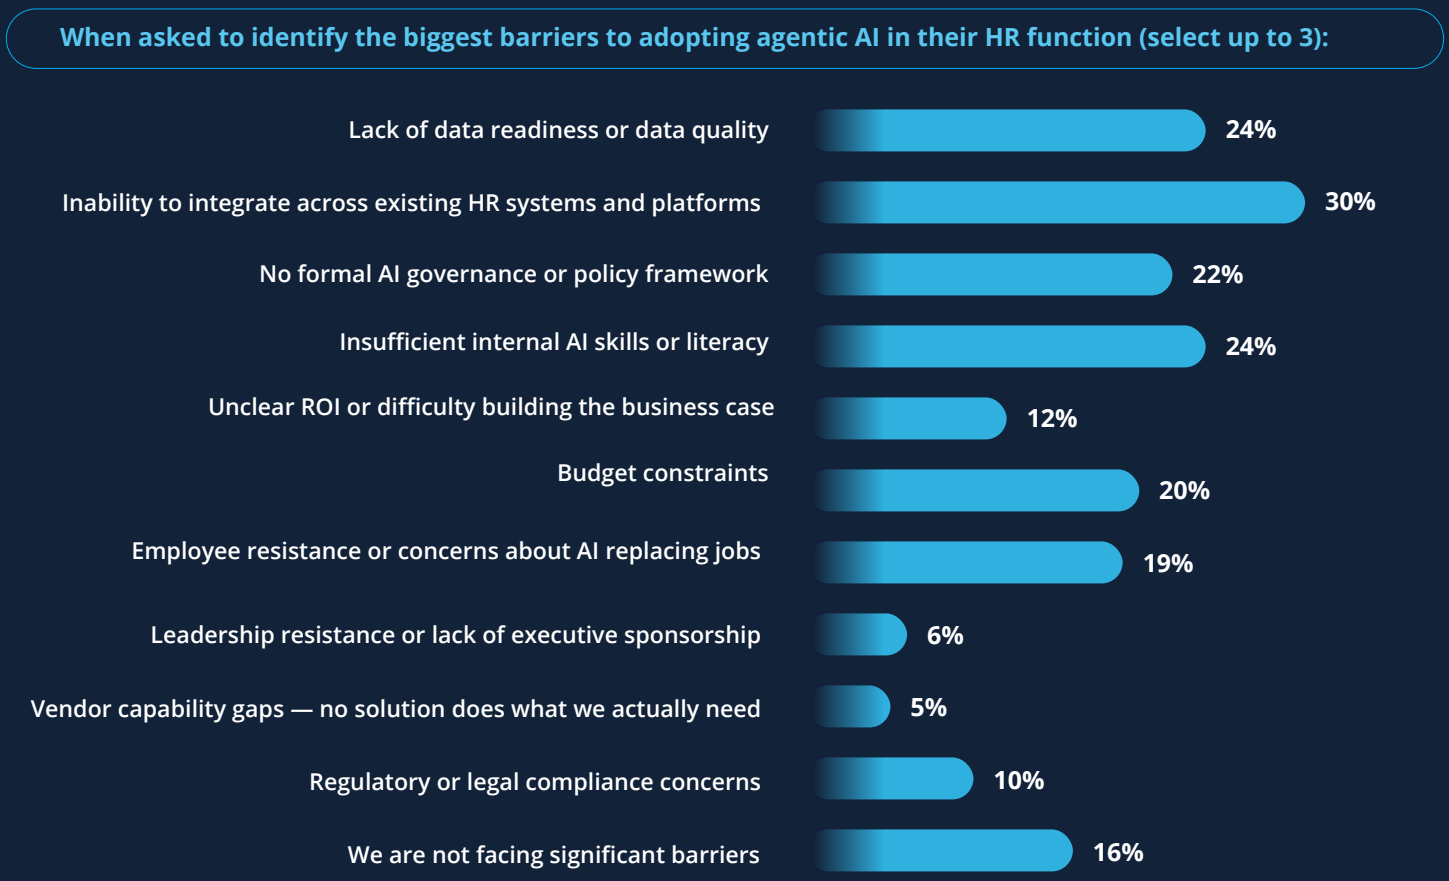

Build for Orchestration, Not Features

The most common mistake organizations make in agentic AI adoption is selecting point solutions for individual use cases and expecting them to add up to a capability. They do not. High-volume screening that does not connect to the AI interviewer. Internal mobility matching that does not share a skills model with workforce planning. Onboarding that starts the day someone shows up rather than the day they accept an offer. Each of these delivers partial value. None of them delivers the compounding benefit of a connected talent operating model.

The vendor strategy question is therefore not which tool is best for a given task. It is which platform creates the conditions for orchestration across the talent lifecycle. The right question to ask any vendor is not what their agent can do in isolation — it is how it connects to the rest of the stack. Organizations that answer that question before they buy will build something that compounds. Organizations that buy feature by feature will build something they eventually have to replace.

Govern Before You Scale

Governance is not a constraint on agentic AI adoption. It is what makes adoption sustainable. Organizations that deploy agentic AI without governance frameworks in place will face one of two outcomes: a high-profile failure that sets back adoption, or inconsistent and undocumented decision-making that creates regulatory exposure. Neither is acceptable for decisions that affect people's careers and livelihoods.

Four requirements apply across every agentic AI deployment in HR. The agent's decisions must be explainable — someone must be able to describe why a specific output was produced. They must be auditable — a record of every consequential decision must exist and be retrievable. Human override must be built in — for decisions above a defined threshold of consequence, a human must be able to intervene, and that intervention must be logged. And bias monitoring must be continuous — not an annual audit, but an ongoing process that flags systematic errors before they compound. 1 in 2 companies do not have a governance framework in HR.

Organizations do not need a perfect governance framework before they start. They need a clear one for the use case they are deploying.

The SaaS Reckoning: Build, Buy, or Partner

The arrival of agentic AI is not just a capabilities shift. It is a business model disruption. The SaaS era was built on a simple premise: buy workflow software, configure it to your process, and pay per seat. That model worked when software was a container for human work. It starts to break when the software begins doing the work itself. Organizations are already asking a question that did not exist five years ago: if an agentic system can source, screen, interview, and onboard without human involvement at every step, what exactly are they paying the per-seat fee for? The answer, increasingly, is unclear — and that unclarity is accelerating a rethinking of how HR technology is acquired, owned, and built.

This survey found that 7% of HR organizations are already building proprietary AI capabilities internally — a small number today, but a meaningful signal. The organizations most likely to build are those with established data science or engineering functions, distinctive talent processes that off-the-shelf software cannot replicate, and strategic conviction that their talent operation is a competitive differentiator worth owning. For these organizations, the question is not whether to build but how to build without creating a maintenance burden that outpaces the capability it delivers. The risk of the build path is not ambition — it is underestimating the ongoing investment required to keep a proprietary system competitive as the underlying models evolve.

For most HR organizations, the more practical path is a structured partnership model — one that sits between full SaaS dependency and full proprietary build. This means selecting a platform provider whose architecture supports configurability at the agent layer: the ability to define the criteria the agent uses, the thresholds at which it escalates, the data sources it connects to, and the outputs it produces. It means negotiating data portability and model transparency as contract requirements, not afterthoughts. And it means treating the vendor relationship as a co-development partnership rather than a software subscription — one in which the organization's talent data and process knowledge actively shape what the system becomes.

The Three-Horizon Roadmap

Most organizations will not move from their current state to a fully orchestrated agentic talent operating model in a single deployment cycle. The practical path is sequential — building capability and confidence horizon by horizon.

Horizon 1: Automate. Deploy agentic AI in the workflows where it is most mature, where the ROI case is clearest, and where governance requirements are most straightforward. Recruiter workflow optimization and high-volume screening are the most common starting points. Establish baselines, measure outcomes, and build internal confidence in what agentic systems can deliver.

Horizon 2: Augment. Expand into use cases that require richer data and more complex integration — AI interviewing, internal mobility, and workforce planning. Use the data and governance infrastructure built in Horizon 1 as the foundation. Begin connecting use cases so that outputs from one feed inputs to another.

Horizon 3: Orchestrate. Move from connected use cases to a coordinated talent operating model. The agent layer coordinates across the full lifecycle — from sourcing through onboarding — with humans governing outcomes rather than executing steps. This is the compounding advantage: a system that gets smarter over time because every deployment feeds data back into a shared intelligence layer.

The organizations that move on agentic AI now will not simply run more efficient HR functions. They will make structurally better talent decisions at scale — faster, more consistent, and more equitable than human-mediated processes alone can deliver. And that advantage will compound. Better data produces better decisions. Better decisions produce better outcomes. Better outcomes feed back into systems that get sharper over time.

The agentic moment is not coming. It is here. The question for HR leaders is not whether to act. It is how fast and how well.

The readiness gap is real. Most HR functions are not yet prepared to deploy agentic AI at scale across the talent lifecycle. But the gap is not a reason to wait. It is a roadmap. Every dimension of readiness — data, integration, governance, skills, alignment — is a specific and addressable investment. Organizations that begin that work now will be positioned to capture the compounding advantage. Organizations that wait for perfect conditions will find that the conditions never arrive — and that the gap between them and early movers has grown in the meantime.

ABOUT US

aptitude research

Aptitude Research is a leading human capital management research and advisory firm. Our in-depth research and vendor assessments help HR leaders develop a deep understanding of the HCM technology landscape, including talent acquisition and engagement, to ultimately make better purchase decisions.

Our proprietary research delivers a comprehensive look at talent acquisition trends and technology. This knowledge, combined with our consulting and advisory services, enables companies to save time, money, and improve the recruiting, hiring, and workforce management experience.

Founded by Madeline Laurano and based in the Boston area, Aptitude Research provides HCM expertise to companies, vendors, and investors. View recent and upcoming research at **www.AptitudeResearch.com**.

eightfold.ai

Eightfold is the agentic talent intelligence company that gives every HR leader, recruiter, and people manager the intelligence to see, develop, and deploy the full depth of talent around them. Our AI Native Talent Intelligence Platform powers a suite of talent agents, that execute alongside your people across the full talent lifecycle. The first flagship talent agent is the AI interviewer, the only talent agent built specifically to give every candidate the nine o'clock interview, evaluating one or one million with identical rigor and zero bias. By combining enterprise data with insights derived from over one billion career profiles, Eightfold's Talent Intelligence Platform powers the world's most complete picture of talent. We read adjacent skills and forecast trajectories to reveal the potential legacy systems miss.

Talent intelligence, human-led