


White paper

# Building agentic AI in the ATS

Roadblocks, realities, and strategic paths



# Executive summary

The convergence of human recruiters with AI agents is rapidly changing how work gets done, and the hiring industry is at an inflection point.

Clients of Applicant Tracking Systems (ATS) are no longer satisfied with basic automation – they now expect **agentic AI** capabilities that actively assist or even autonomously handle recruiting tasks. Recent surveys show that 99% of hiring managers are using AI in some capacity. Furthermore, 95% plan to increase investment in AI to optimize recruitment.

This surge in adoption, fuelled by the viral mainstreaming of tools like ChatGPT - which reached 100 million users just two months after launch, the fastest adoption in history - has dramatically raised the bar for ATS products.

Yet, for many product and engineering leaders, delivering sophisticated AI-driven features fast enough feels **nearly impossible**. Legacy architectures, compliance constraints, and finite resources create a perfect storm of obstacles.

Despite AI being named a top strategic priority for 2025, most ATS providers are struggling to keep pace with innovations like agentic AI—systems that can plan and act autonomously toward user-defined goals. The result is a growing gap between what clients expect and

demand and what traditional development cycles can deliver.

This white paper explores the reality behind that gap. We outline the key roadblocks ATS teams face and examine three strategic paths to close it:

- **Build in-house** – develop agentic AI capabilities internally from scratch.
- **Integrate point solutions** – plug in specialized AI tools to extend your platform.
- **Embed a white-label AI layer** – partner with an AI platform that integrates deeply into your ATS, under your brand.

Each approach carries distinct pros and cons in terms of time-to-market, cost, control, and risk. So in this paper we'll provide a side-by-side comparison to help strategize the best path forward.

Crucially, we affirm that doing nothing is not an option. With almost all companies investing in AI but only 1% feeling they've fully matured their use of it, the competitive window is open now for ATS providers to lead - or risk being left behind.

This document concludes with a clear case for why acting fast on AI isn't just smart—it's essential for ATS leaders—and outlines how to move forward with confidence.

# In this white paper:

- Client expectations: The rise of agentic AI 4
- Why fast innovation feels impossible in an ATS 6
- Bridging the gap: Strategic paths to deliver agentic AI 9
  - Build in-house 9
  - Integrate point solutions 11
  - Embed a white-label AI layer 13
- Conclusion: The strategic imperative for moving now 18

# Client expectations: The rise of agentic AI

Clients of ATS platforms – recruiters, HR departments, and talent acquisition leaders – are experiencing a shift in their expectations due to the rapid advancements in AI.

Agentic AI has moved from buzzword to baseline expectation. Gartner defines agentic AI as “autonomous AI that can plan and take action to achieve goals set by the user”, essentially creating a virtual workforce of agents to assist and augment human work.

In practical terms, an agentic AI layer in recruiting could automatically screen candidates, schedule interviews, send follow-up communications, or even adjust hiring workflows on the fly – all with minimal human prompting.

These AI agents combine large language models, machine learning, and automation to not only follow rules but to understand context, make decisions, and continuously learn. The promise is a hiring process that is faster, smarter, and more proactive, with mundane tasks offloaded to tireless AI helpers.

This vision is no longer sci-fi; it is becoming the new normal that clients anticipate. The

explosive debut of generative AI in 2023–2024, epitomized by ChatGPT’s overnight popularity, made AI-assisted work mainstream.

Recruiters have seen or experimented with AI writing job descriptions, sourcing candidates, and chatting with applicants. In fact, a recent survey of 1,005 hiring managers found nearly all (99%) are already using AI in some sort of way, and another study indicates 79% of organizations have integrated some form of AI or automation into their ATS so far, reflecting how pervasive the expectation has become.

Beyond statistics, high-profile moves in the market reinforce this trend. For example, LinkedIn, the main platform recruiters use alongside their ATS, launched its first AI Hiring Assistant in late 2024 - An AI agent that automates a wide array of recruitment tasks, from drafting job descriptions to sourcing and engaging candidates.

This was described as a milestone in LinkedIn’s AI journey and a direct response to customer demand. Notably, LinkedIn’s team acknowledged that the rise of generative AI left them scrambling to bring the product up to speed on the front-end.

If even a tech giant had to rush to meet the new AI expectations, it's clear that ATS providers, large and small, are under pressure to elevate their offerings. Clients now envision their ATS not just as a database or workflow tool, but as an intelligent collaborator – an AI colleague that works with them.

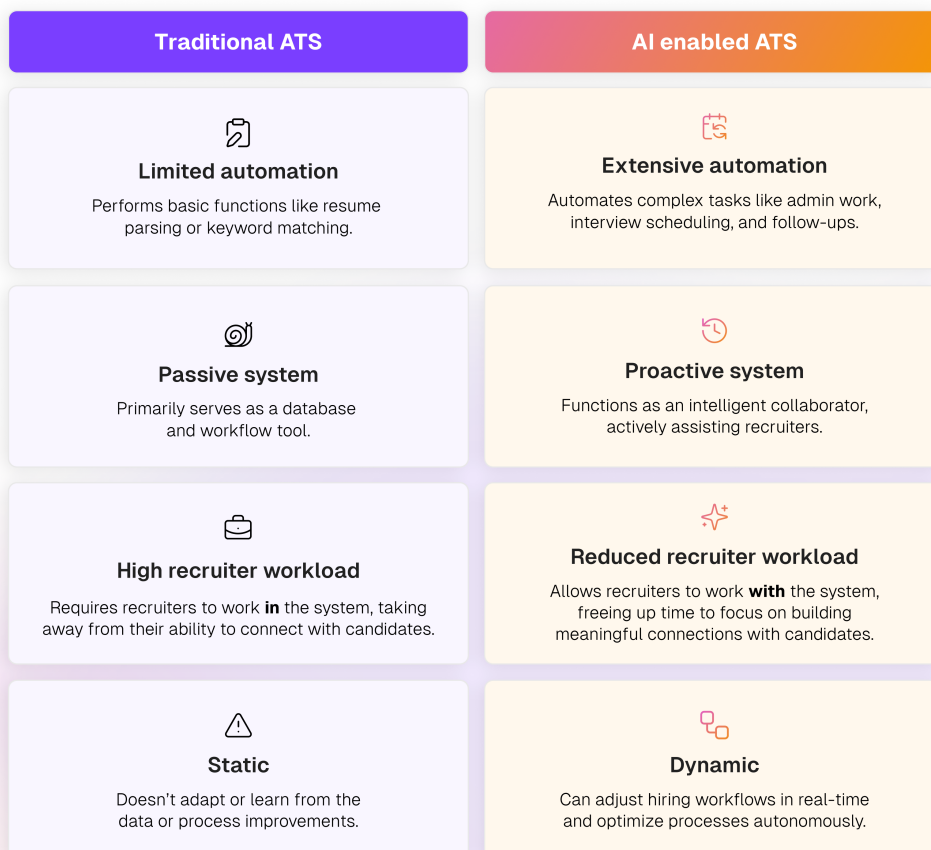
Put simply, the market now views agentic AI capabilities as the next must-have.

Recruiters want an ATS that proactively takes action: automatically nurturing candidates, optimizing job postings, flagging hiring process bottlenecks, and more. This is a dramatic leap from the ATS of the past, which primarily acted as a passive system of record.

Embracing this “AI colleague” model could unlock huge gains for end-users—from eliminating “wasted” candidates through continuous re-engagement of past applicants (zero waste), to drastically reducing the administrative burden on recruiters (zero admin).

The opportunity for ATS vendors is to meet these expectations head-on – harnessing agentic AI to transform their platforms from static applicant tracking systems to dynamic talent acquisition assistants.

However, delivering on this vision is easier said than done. As the next section explores, many ATS teams face internal realities that make such rapid innovation challenging.



# Why fast innovation feels impossible in an ATS

If the rise of agentic AI is the opportunity, the reality inside most ATS product teams is a maze of constraints that make seizing that opportunity difficult. It's not for lack of vision – ATS leaders are well aware of what modern AI could do for their platforms.

The frustration is that knowing and doing are very different when you're carrying the weight of a mature software product and an established customer base.

Several factors contribute to why fast-paced AI innovation can feel nearly impossible.

## 1. Legacy architectures and technical debt

Many leading ATS platforms have been around for a decade or more, with code bases and infrastructures that pre-date the AI boom. Core workflows and data models are tightly coupled and meticulously tuned for reliability and compliance.

Introducing complex AI algorithms or real-time decision-making agents into such an environment can require fundamental refactoring. Moreover, the “pre-committed

roadmap” is often already filled with essential upgrades, client-specific customizations, and quality improvements that cannot be dropped.

In short, product teams find themselves with big AI dreams but limited bandwidth. As quarters pass with minimal progress on this front, it reinforces the sense that the innovation “engine” is stuck in low gear.

## 2. Resource and talent constraints

Implementing cutting-edge AI is resource-intensive. It's not just about coding an integration to an API; it requires data science expertise, model training and tuning, ongoing ML ops, and experimentation.

ATS firms – especially mid-sized ones – may not have teams of PhD-level AI researchers on staff. Hiring such talent is fiercely competitive in 2025; “there's a gold rush on for AI talent,” as one industry VP put it.

The demand for generative AI skills exceeds the available supply, with job postings for AI roles far outpacing qualified candidates.

Even if an ATS can hire a few AI specialists, they must work within a larger organization that may not be versed in agile AI development practices. Budget is another aspect: building AI features from scratch can be costly in terms of cloud computing for model training, experimentation time, and opportunity cost of diverting engineers from other projects.

To sum it up, limited specialized talent and budget make it hard to spin up an AI team capable of rapid breakthroughs.

### 3. Compliance, risk, and trust considerations

Nowhere is the adage “move fast and break things” less applicable than in hiring technology. Recruiters and employers rightly demand that their tools comply with equal opportunity laws, data privacy regulations, and ethical standards. Any AI that makes decisions about candidates comes under scrutiny for potential bias or discrimination.

A vivid example is Amazon’s infamous experiment with an AI recruiting tool that had to be scrapped after it was found to systematically downgrade female candidates. The lesson is clear: rushing an inadequately vetted AI feature into an ATS could have dire consequences for clients and for the vendor’s reputation.

Indeed, regulators are stepping in. New York City’s AI in Hiring law (Local Law 144), for instance, now requires that automated

employment decision tools be subjected to bias audits and that candidates be notified of their use. Similar regulations are emerging elsewhere.

For ATS providers, this means any AI-driven functionality might need extensive validation, audit trails, and transparency measures before release. Such due diligence inevitably slows down the development cycle.

Moreover, product leaders may be risk-averse to deploying AI that they don’t fully understand, fearing unintended outcomes. They must be confident that an AI feature will do no harm – a high bar that can delay deployment or lead to very conservative designs, limiting the “autonomy” of the AI to avoid mistakes.

### 4. User adoption and change management

Even if an ATS builds a clever AI assistant, there’s the question of how users will adopt it. Recruiters are busy professionals; if an AI feature is not deeply integrated and seamless, it may go unused.

Ensuring a smooth user experience often means tightly coupling AI features into existing workflows – which can conflict with the modular experimentation that AI development prefers. It’s a catch-22: to get the AI right you want to iterate quickly in a sandbox, but to get users to actually use it, you need to polish and embed it within the product.

This tension can also slow down development. Additionally, ATS companies have to account for the fact that their clients vary in tech-savviness and openness to change.

A drastic new “autonomous” system might confuse or intimidate some customers unless it’s rolled out carefully—with education, feature flags, and support. All of this adds friction to delivering AI updates in the product.

These realities explain why, despite the hype, most ATS suites have only taken small steps with AI so far – things like basic resume parsing, keyword matching, or chatbots powered by third-party AI.

While nearly four out of five organizations say they’ve added AI automation to their ATS, these are usually narrow point solutions rather than a holistic “AI colleague” embedded throughout the platform.

Achieving the latter can feel like turning a battleship - slow and arduous. As one CTO

noted, launching your own AI agent is a forever project, not a one-time build – it demands constant attention from a talented team, ongoing model updates, and non-trivial work on data privacy and security.

In other words, it’s a major commitment that some ATS providers aren’t sure they can afford while keeping the lights on.

Acknowledging these roadblocks is important. It’s not that ATS leaders lack ambition; it’s that the deck seems stacked against rapid, radical innovation.

However, “impossible” does not have to remain so. With the right strategy, even a legacy-bound product can deliver new AI-driven value to customers – if you choose the approach that aligns with your constraints and strengths.

In the next section, we explore three strategic paths to bringing agentic AI capabilities to an ATS platform, each addressing the challenges above in different ways.

## What’s blocking AI innovation in ATS’s?

### 1. Legacy systems

Outdated infrastructures make it difficult to integrate advanced AI without major system overhauls.

### 2. User adoption

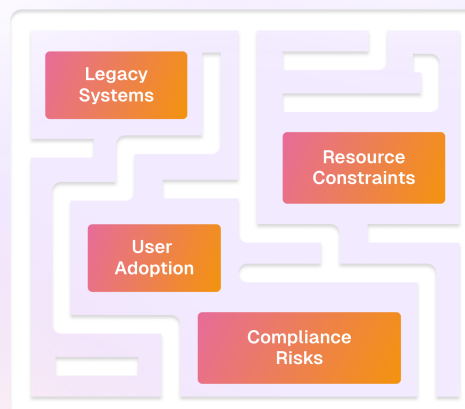
Users may resist new AI features if they disrupt established workflows or seem too complex to use.

### 3. Resource constraints

AI development requires specialized talent and significant investment, which many ATS providers, especially smaller ones, lack.

### 4. Compliance risks

AI powered recruitment software must comply with strict regulations to avoid bias, adding delays due to necessary audits and transparency measures.



# Bridging the gap: Strategic paths to deliver agentic AI

Given the high client expectations and the internal challenges, ATS product leaders must decide how to bring agentic AI to their platform in a practical, sustainable way.

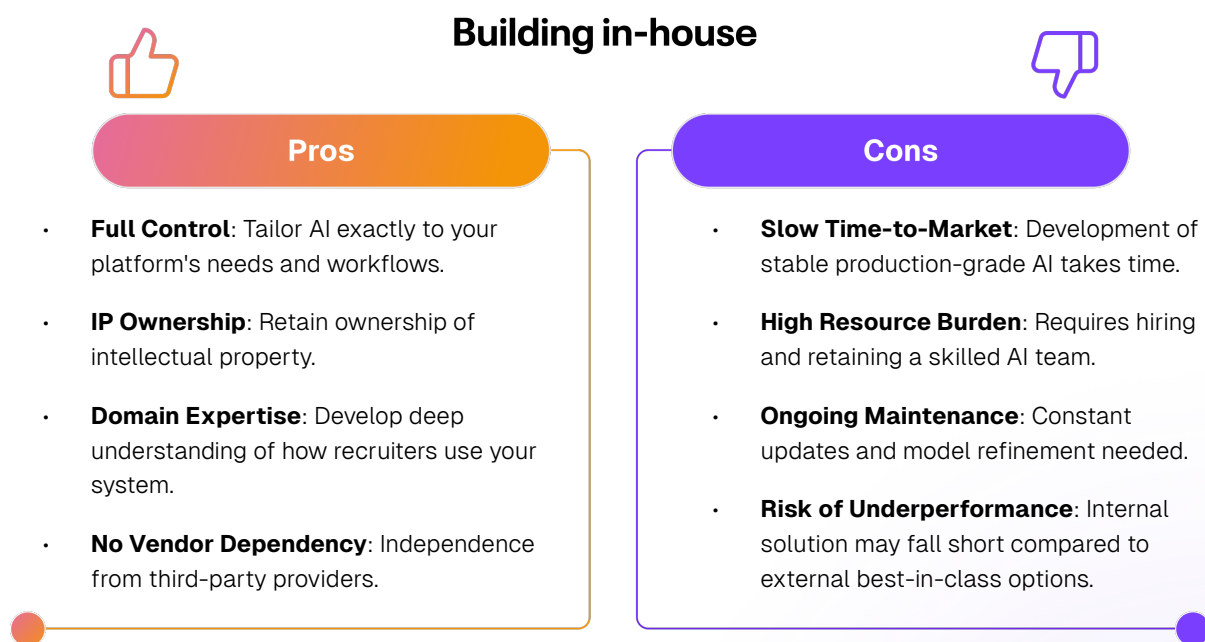
Broadly, there are three paths forward: Build in-house, integrate point solutions, or embed a white-label AI layer.

There is no one-size-fits-all answer – each option comes with trade-offs in speed, cost, control, and risk. In this section, we break

down each strategy and then compare them to inform your decision-making.

## Build in-house

This approach means developing the AI capabilities internally using your own engineering and data science resources. Essentially, you commit to becoming an AI builder - designing, training, and maintaining your own models or AI systems that deliver agentic functionality within your ATS.



## Pros of this approach

Building in-house offers the highest degree of control and differentiation.

The AI solution can be tailored exactly to your platform's workflows and your clients' unique needs, potentially giving you a proprietary advantage. You own the intellectual property end-to-end, which could become a competitive edge.

Over time, an in-house AI team builds deep domain expertise (e.g., understanding how recruiters use your system) that off-the-shelf products might lack. If AI is strategically core to your long-term value, owning the capability might make sense.

Also, there's no dependency on third-party vendors; you set the roadmap and you're not exposed to a partner's business risks or changes.

## Cons of building in-house

The drawbacks of this scenario are significant. First, time-to-market is slow; building a production-grade AI, especially an autonomous agent that is reliable, can take years.

Your developers would need to go through cycles of research, prototyping, user testing, and iteration. Meanwhile, competitors or new entrants might leap ahead with ready solutions.

Second, resource burden is very high. As discussed, AI talent is scarce and expensive; you must recruit and retain a skilled team. Even then, launching your own advanced AI is a forever project, not a one-time build – the team must continuously refine models, update for new data, and handle maintenance issues like uptime, scaling, and security.

All this represents a substantial ongoing cost. It's easy to underestimate the total cost of ownership: beyond initial development, think of continuous model training, computing infrastructure, data labeling efforts, compliance audits, and so on.

Third, the risk of falling short is real. If your in-house solution ends up achieving only 40% of task automation, where a best-in-class external solution could achieve 60+%, the ROI might actually be worse despite all the investment.

There is a pitfall in reinventing the wheel but doing it less effectively; customers won't care that you built it if it doesn't perform as well as other AI they've seen.

In short, building your own solution requires constant iteration, substantial resources, and typically delivers lower ROI than buying a pre-built solution, as industry experts have already observed.

### When should you build in-house?

Building in-house can be the right choice if AI is viewed as a core competency you want to develop and if you have patient leadership and funding. If you go this route, it's wise to start with contained pilot projects and gradually expand, proving value along the way.

Also, ensure you're ready to invest in MLOps and governance from day one, so you can manage the models responsibly at scale. Be aware: building AI internally is not just a feature decision – it's an organizational commitment.

In practice, you should expect to expand your product and design team by at least 40%, and your engineering team by 60%, to support the necessary R&D, infrastructure, and iterative release cycles. These hires will often command salaries roughly 25% above your current workforce average, due to high

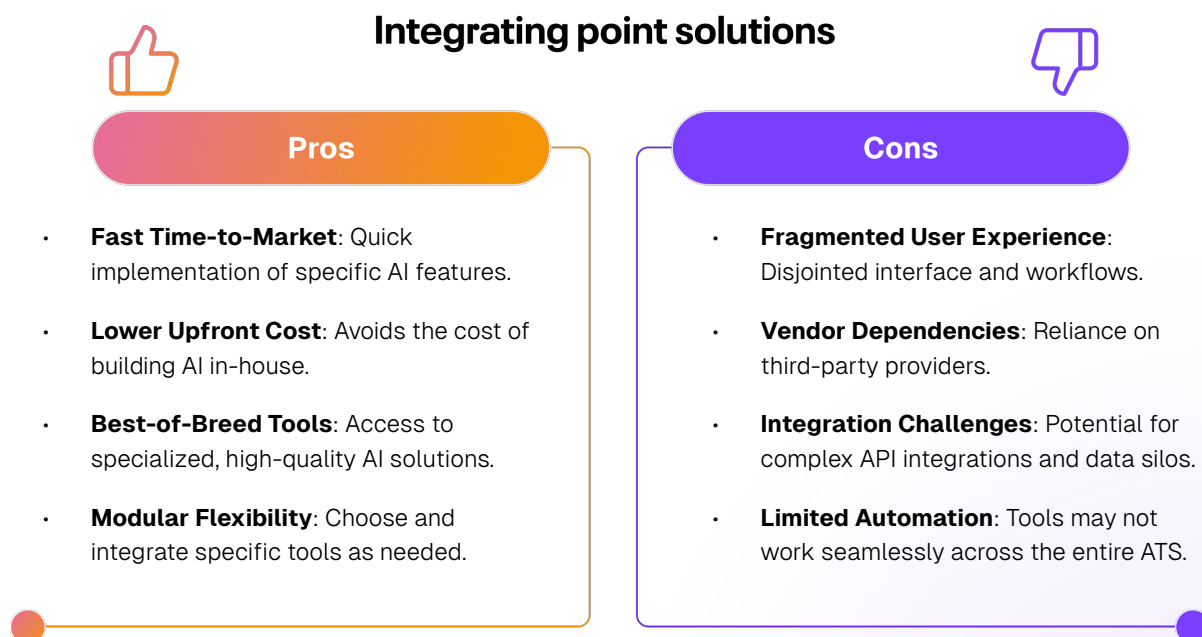
competition for AI and machine learning talent.

Building in-house is a long haul – it can eventually yield a highly integrated AI that is yours – but it requires vision, patience, and executive-level commitment to see it through.

### Integrate point solutions

The second strategy is to leverage third-party AI point solutions – in other words, partner with specialized AI tool providers and integrate their capabilities into your ATS.

Many ATS vendors have already been doing this to add specific features. Examples include integrating an AI resume parsing service, plugging in a chatbot for candidate Q&A, or using a scheduling assistant to automate interview coordination. These usually come via APIs or marketplace apps.



## Pros of this approach

The point solution approach generally offers the fastest initial time-to-market for each capability. Instead of building from scratch, you tap into an existing solution that is often ready-to-use. This can allow you to roll out new features in a matter of months or even weeks.

It's also a lower upfront investment – you don't need to hire a big AI team; the solution provider has already done the heavy lifting on the technology. You can cherry-pick “best of breed” tools for each function: one vendor might excel at chatbot conversations, another at candidate rediscovery in the resume database, and so on.

By integrating them, you give your customers access to advanced functionality without having to develop it all in-house. This modular approach can be pragmatic and cost-effective.

Additionally, since these solutions are focused on one problem, they are often quite polished in that domain. For example, a dedicated interview scheduling AI might handle calendar edge cases better than a generic system you'd build.

For a product strategy, integrating point solutions can be a way to test the market's response to AI features – if a particular AI add-on gets great adoption, it validates

demand without you having sunk huge R&D cost.

## Cons of integrating point solutions

The biggest downside however is the resulting fragmented user experience (UX) and architecture. Each point solution may have its own interface elements, data stores, and ways of working.

From the recruiter's perspective, using these features might feel like using separate tools rather than one cohesive ATS platform. They might have to jump between screens or deal with inconsistent UI patterns, which can hurt adoption.

For the ATS provider, managing multiple integrations can become a major maintenance headache: different APIs, different update cycles, and potential data silos. The more external pieces, the more potential points of failure.

There's also a vendor dependency issue – if one of those solution providers is acquired, changes their API, or goes out of business, your product roadmap can be thrown off. Moreover, simply adding disparate tools doesn't necessarily give you a true agentic AI that spans the whole hiring workflow; it might solve individual tasks but not coordinate between them.

Each point solution addresses a slice of the process, but who ties it all together? Without

careful integration design, you risk ending up with a “Frankentool” – an assemblage of features that don’t feel unified.

Another consideration is the limited differentiation: these same third-party AI tools could be (and often are) integrated by your competitors as well. If every ATS is plugging in the popular AI resume screener of the day, it’s hard to stand out.

In essence, you’re outsourcing innovation, which can be fine, but it means your product’s unique value must come from elsewhere (like superior workflow design or service).

Finally, costs can add up in the form of subscription/license fees to these vendors or revenue-sharing agreements, which might eat into your margins over time.

### **When should you choose point solutions?**

Integrating point solutions is attractive when you need quick wins in functionality and when you’re comfortable treating AI as additive, not central, to your platform. It works best if you maintain a strong core product and use integrations to fill gaps.

Many up-and-coming ATS or HR tech platforms took this route initially – for example, partnering with a known AI scheduling assistant to offer automated interview scheduling, rather than building it.

If you choose this path, aim to mitigate the fragmentation: work with your partners on

deeper integration (e.g., single sign-on, embedding their UI components into your interface so it feels native).

Also, have a plan for data – ensure the data from these tools flows back into your ATS database so you’re not blind to what the AI did. Point solutions can be a component of your AI strategy, but orchestrating them into a smooth symphony is the challenge.

### **Embed a white-label AI layer**

A newer strategic option is to embed a white-label AI “layer” into your ATS – effectively partnering with an AI platform provider that offers a broad suite of capabilities which you integrate wholesale and present as part of your own product.

This differs from point solutions (which are usually narrow in scope) in that a white-label AI layer is modular and comprehensive, designed to sit underneath your UI and workflows across the board. It’s “white-label” in that it’s meant to be branded as your product’s AI – the end-user may not know an external partner is powering it.

### **Pros of this approach**

This approach can significantly accelerate time-to-market for a wide array of AI features. Instead of integrating 5–6 different tools or building for years, a single integration with a robust AI layer could infuse capabilities like intelligent automation, conversational

assistants, smart recommendations, and analytics all at once.

Because the solution is designed to plug into existing systems (often via APIs and embeddable components), it tends to require minimal engineering lift on your side. For example, you might add a few API endpoints and drop in pre-built UI modules, and suddenly your ATS has an AI copiloting every step – all within months, not years of effort.

Another advantage is that these platforms are often built with enterprise considerations in mind – compliance, security, scalability – since they aim to serve as an AI backbone for multiple vendors.

That means a lot of the heavy work around data privacy, bias mitigation, and reliability is handled by specialists, relieving your team of that burden. It's a way to quickly offer "production-grade" AI that's been tested and refined elsewhere.

White-label integration also lets you maintain your brand experience; you control how the AI features are presented to users.

To the customer, it looks like your ATS just got smart everywhere overnight. Internally, this can be far more cost-effective than building – typically a partnership or licensing fee structure – and you don't need to expand an entire AI division to support it.

In summary, it's the fastest and arguably lowest-risk way to attain a comprehensive agentic AI capability in your platform, effectively outsourcing the innovation but in a tightly integrated fashion.

### **Cons of the white-label route**

The primary trade-off is reliance on an external partner for a core part of your product. You are entrusting a chunk of your user experience and innovation roadmap to another company. It's important to have clear agreements (e.g., data ownership, service level uptime, update frequency, confidentiality of your user data, etc.).

Another consideration is loss of some control over feature differentiation. While white-label solutions are often customizable to a degree, you might not be able to tweak the AI's behavior as deeply as if you had built it.

Additionally, while initial integration is easier than building from scratch, it's not hands-off – your engineers will still need to work closely with the partner's APIs and possibly adjust your workflows to fully leverage the AI.

It's also key to manage the branding and messaging carefully: you want customers to value the new AI features as part of your platform, not view them as an external add-on. That usually means tight UI integration and consistent support (your support teams will need to understand the AI layer to assist users).

In terms of cost, a white-label solution might involve revenue sharing or a usage-based fee; depending on how you price your product, you'll need to ensure the economics still make sense.

Lastly, consider long-term strategy – if down the line you decide to build your own AI, transitioning away from a deeply embedded partner could be complex (though not impossible).

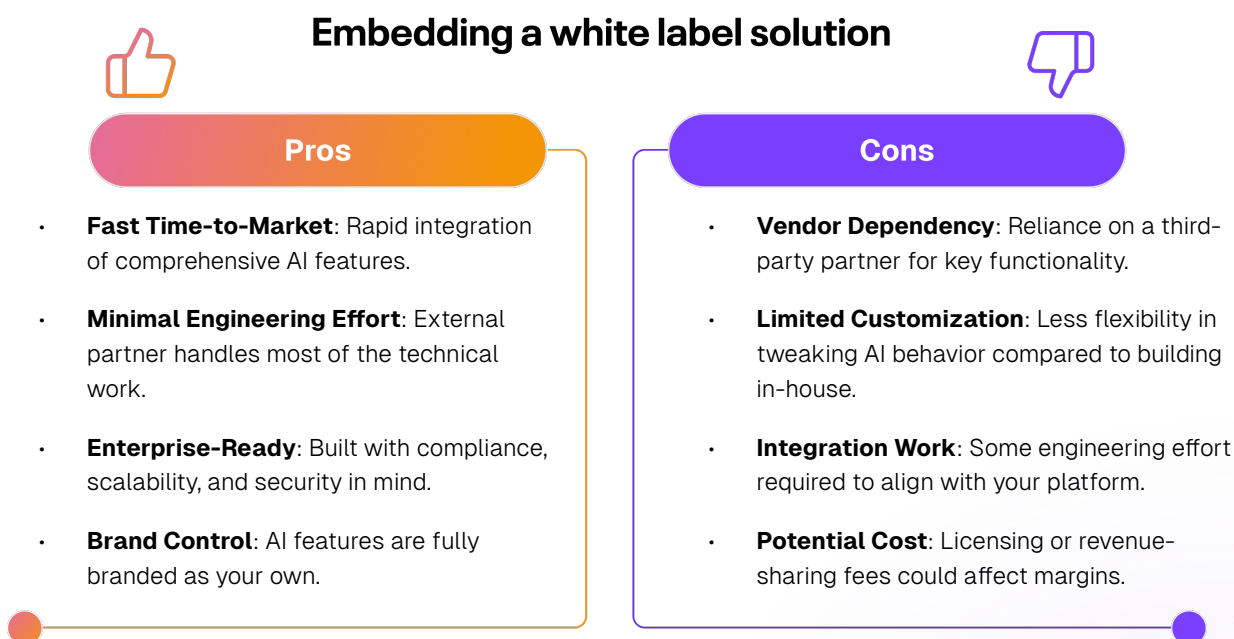
### When should you opt for white-labeling?

Embedding a white-label AI layer is a compelling choice when speed is of the essence and you want a broad AI upgrade across your platform without waiting for incremental rollouts. It's especially useful for mid-size ATS providers who cannot afford to

fall behind in the AI race, or larger providers that want to leapfrog current piecemeal efforts and deliver a big bang of AI features.

For instance, an ATS that wants to transform into an "AI-first" system in the eyes of customers within the next year could partner with a white-label AI platform that provides end-to-end recruiting intelligence – from sourcing to onboarding – under the hood.

When pursuing this, do thorough due diligence: evaluate the partner's technology (does it truly handle the tasks as advertised?), talk to reference clients if possible, and possibly pilot in a sandbox environment.



Structure the partnership so that you retain ownership of your data and have contingencies.

After outlining these options, it's helpful to see them side-by-side. The table below compares key factors for each strategic path:

If executed well, this strategy can turn your ATS into an AI-powered recruiting assistant very quickly, effectively turning your system into one that “works with” recruiters rather than just being a system they work in.

Strategy	Time to Market	Upfront Effort & Cost	Control & Differentiation	Risks/Challenges
<b>Build In-House</b>	<b>Slow.</b> Multi-quarter or multi-year development for advanced AI features.	<b>High.</b> Requires hiring AI talent, investing in R&D, infrastructure, continuous development.	<b>High.</b> Full control over design, unique IP potential if successful.	<b>Development risk</b> (outcome may underperform), <b>long-term maintenance burden</b> , need to manage compliance and AI ops internally.
<b>Integrate Point Solutions</b>	<b>Fast</b> (per feature). Can add specific capabilities in weeks or months via APIs.	<b>Moderate.</b> Lower dev cost upfront per integration, but effort multiplies with each additional tool. Ongoing vendor fees.	<b>Low.</b> You get good functionality but competitors can integrate the same tools; experience can feel less unified.	<b>Fragmented UX</b> if not well integrated, <b>vendor dependencies</b> (multiple), potential data silos, limited overall automation (features not “talking” to each other).
<b>White-Label AI Layer</b>	<b>Fast (broadly).</b> A comprehensive AI upgrade can be rolled out in a few months with one integration.	<b>Moderate.</b> One major integration effort; partner provides the heavy tech. Cost is in partnership/license fees rather than big internal R&D.	<b>Medium.</b> Solution is customized and branded as yours, but underlying tech is shared. Differentiation via how you integrate and utilize it.	<b>Dependency on partner</b> for core functionality, must ensure alignment and reliability. Less fine-grain control over AI behavior. Need to handle partnership management and integration tweaks.

Each ATS provider's situation will influence which path makes the most sense.

- **Build in-house** might suit a well-resourced market leader aiming for long-term proprietary AI capabilities.
- **Point solutions** fit a more conservative, incremental improvement strategy or niche needs.
- **White-label AI layer** appeals to those who want a rapid, transformative leap in capabilities with manageable investment.

In some cases, a hybrid approach could emerge – for example, integrate a white-label core for general AI functions but also build a couple of bespoke AI features in-house that are unique to your platform's domain.

What's crucial is to have a clear strategy; given how fast the landscape is evolving, standing still is the one option that is guaranteed to leave you behind.

# Conclusion: The strategic imperative for moving now

The advent of agentic AI in talent acquisition is a paradigm shift that ATS leaders cannot afford to ignore or endlessly defer. The window of opportunity to delight your customers – and reassure them that your platform will meet their emerging needs – is open now, but it may narrow quickly.

Clients are already experiencing AI's benefits in piecemeal ways, and they are hungry for their primary systems (their ATS) to bring it all together. If their current provider doesn't move fast enough, they will seek alternatives, whether that's supplementing with external tools or even switching to a more AI-forward competitor down the line.

From a competitive standpoint, we are in a phase where nearly everyone is investing in AI, yet almost no one feels they've mastered it. As noted, almost all companies invest in AI, but just 1% believe they are at maturity in using it effectively. This includes ATS vendors and their customers alike.

In other words, the playing field is still being leveled. New entrants with AI-centric products are trying to disrupt, while established players are ramping up AI initiatives.

This creates a strategic opening: decisive action now can set you apart before AI capabilities become ubiquitous table stakes.

Moving now does not mean rushing in blindly – it means choosing a thoughtful path (build, integrate, or partner) and executing with urgency. It means securing leadership buy-in that AI features are not “nice to have” gimmicks but core to the future value proposition of an ATS.

It also means communicating to your customers that you're on this journey with them. For example, even as you work behind the scenes – perhaps partnering with a white-label AI provider or piloting integrations, start educating clients on what's coming.

Frame new AI capabilities not as moonshots, but as concrete solutions to their daily pains: less data entry, faster candidate matching, richer insights, etc. This will build anticipation and goodwill that you are proactively addressing the roadblocks and realities they know so well.

A consultative, transparent approach is key. If you're integrating AI, ensure you're also providing guardrails and assurances around it – such as documentation on how the AI makes decisions, options to override or double-check AI actions, and compliance certifications.

Showing that you've thought through the ethical and practical implications will make clients more comfortable embracing the changes. Remember Gartner's note that agentic AI "requires robust guardrails to ensure alignment with intentions".

Bringing AI to your ATS is as much about trust as it is about technology. The good news is that success breeds success: early wins with AI (even small features that save recruiters time) will build momentum for further innovation.

In conclusion, ATS product and technical leaders should see agentic AI not as a daunting impossibility, but as a strategic imperative with attainable paths forward. The roadblocks are real – legacy tech, limited resources, compliance concerns – but as we've discussed, there are concrete strategies to navigate them.

Whether you choose to cultivate an in-house AI powerhouse, curate a smart ensemble of specialized tools, or turbo-charge your platform through a white-label AI partnership, the important thing is to make a choice and act.

The worst move would be standing still and watching the market evolve without you.

The next generation of ATS is being defined today. It will be characterized by systems that are agentic, adaptive, and deeply collaborative with humans.

For the first time, recruiters will not only use software but work alongside software that has a degree of "agency" – handling drudge work, surfacing recommendations, and continuously learning from data.


Achieving this in your product will ensure you remain a relevant, even indispensable, part of your clients' tech stack in the years ahead. Failing to do so risks your platform becoming obsolete, a relic of a previous era of talent tech.

As you plot your course, keep the end vision in mind: an ATS that truly operates as an AI colleague to the recruiter.

Each step you take – a pilot here, an integration there – should be measured against that vision. And remember, you are not alone in this journey. Engage your client advisory boards, talk to forward-thinking customers about what they'd value most, and lean on the growing ecosystem of AI solution providers who specialize in this domain.

The path to agentic AI in ATS may have obstacles, but it is navigable and well worth the effort. Those who start now, with a clear strategy and a bias for action, will shape the future of recruitment technology.

In a space that's all about connecting people with opportunities, let's ensure that AI becomes a powerful ally in delivering better, faster, and more humane hiring outcomes. The time to build that future is now.



# Sources

**Gartner (2024)** – Top Strategic Technology Trends for 2025. Gartner identifies Agentic AI as the #1 trend, defining it as AI that can plan and act to achieve user-defined goals autonomously.

**LinkedIn (2024)** – Announcing LinkedIn's AI Hiring Assistant. Blogpost describing the launch of LinkedIn's autonomous AI recruiter assistant and the internal race to keep up with generative AI expectations.

**ResumeBuilder.com (2023)** – Survey: 99% of hiring managers use AI in hiring. Includes stats on AI adoption, perceived productivity benefits, and trust in AI for job matching.

**Eightfold.ai (2023)** – AI Adoption in Talent Acquisition. Research summary indicating widespread use of AI in enterprise ATS and forward-looking investment trends.

**Harvard Business Review (2024)** – Why Building an AI Agent In-House Is Hard. Breaks down the difficulty of building agentic systems internally, even for tech-forward companies.

**Intercom (2024)** – Buy vs Build: The Real Cost of AI Product Development. Discusses cost, ROI, and maintenance burdens of developing AI features internally.

**New York City Department of Consumer and Worker Protection (2023)** – Local Law 144 Compliance Guide. Regulations for

automated employment decision tools, including audit and transparency requirements.

**The Verge (2018)** – Amazon scrapped a recruiting tool that showed bias against women. Case study of biased AI in hiring and the reputational risks involved.

**Shopify (2025)** – Company statistics page showing over 4 million online stores powered by Shopify, most of them under white-label storefronts.

**McKinsey & Company (2023)** – The State of AI in 2023. Includes data on talent shortages in AI, job market competitiveness, and resource constraints in AI development.

**Aptitude Research (2023)** – The ATS Market: A Changing Landscape. Includes buyer dissatisfaction data, intent to switch ATS, and friction in adoption.

**SelectSoftware Reviews (2025)** – ATS comparison tools and pricing benchmarks across platforms.

**HR Executive / Employ (2022–2024)** – Commentary on M&A trends in the ATS market and the shift in innovation post-acquisition.

**Sense HQ / Crelate (2023–2024)** – Research and whitepapers on fragmented tech stacks, user experience issues, and integration pain points in recruitment software.

# About Carv

Headquartered in Amsterdam, Carv is the #1 AI platform for recruiters, designed to take over all repetitive admin tasks, so recruiters can focus on what matters: Building human connections.

Carv's AI automates every stage of the recruitment process, from application handling and talent engagement to scheduling and post-interview evaluation.

With hands-free workflows and automated ATS updates, Carv helps recruiters cut up to 80% of admin tasks, reduce cost per hire by 70%, and achieve 98% ATS data accuracy.

The customizable software comes with ready-to-use integrations, paving the way towards all-encompassing autonomous recruitment solutions.

For more information on Carv, visit <https://www.carv.com> or follow us on [LinkedIn](#).

Ready to bring AI into your ATS?

# Take the first step today with Carv's white-label ATS integration

[Request a demo](#)

or visit [our website](#)