




HIREWAND · RESEARCH WHITE PAPER

Fragmentation and Integration

How AI is changing the economics of each category of recruitment software — and why the resulting tools increasingly need to work together as one system.

Author

Shiva Maran

Field

Recruitment & HR Technology

Reading time

~18 minutes

Audience

TA leaders · founders · operators

ABSTRACT

Artificial intelligence is changing not only what software can do, but how software is built, sold, and made to work together. This paper examines that shift in the context of recruitment technology.

The analysis proceeds in three parts. The first distinguishes the main categories of software and the distinct value each provides, then describes how AI affects each one differently. The second section focuses on two of these categories — the broad platform, which ties many separate functions together into one system, and the hyper-specialist, which goes exceptionally deep in a single area — and on what it takes for the two to work together. The third turns to the customer and to what now creates the most value: how to choose among all four categories given the challenge each carries, and how an integrated architecture — a platform that ties hyper-specialists together — lets a customer draw on the best of all of them at once. Figures draw on industry research published in 2025 and 2026; sources are listed in full at the end.

PART 01

How AI Reshapes the Software Categories

To understand AI's effect on software, it helps to look past individual products at how software is built and sold. Most software falls into four categories. Each has its own economics, and each provides a distinct form of value.

The four are **custom software** (built by an organisation for itself), **broad platforms** (one wide product sold to many), **narrow specialists** (one task, performed in depth), and **tech-enabled services** (software paired with human judgment). The figure below sets out each, alongside the value it provides and an example.

Each category exists because it serves a need the others do not. The relevant question is not which category is best, but how AI changes each one.

FIG. 1

Four categories of software, and the value each provides




Figure 1. The four categories differ on who the software is built for, how wide a span it covers, and whether value is delivered as code alone or as code together with human service. Each provides value the others do not.

AI works through two distinct mechanisms

AI affects software in two separate ways, and keeping them apart is useful. It lowers the cost of *building* software, and it lowers the cost of adding intelligence *inside* software. Both mechanisms are widely available to anyone building or buying technology.

FIG. 2

AI's two mechanisms




Figure 2. The first mechanism lowers the cost of making software; the second lowers the cost of putting intelligence inside it. Because both are widely accessible, differentiation increasingly rests on what a product does with them.

Running these two mechanisms through the four categories produces a different effect in each — set out in the table below. In every case the category is reshaped, not removed.

The most consequential shift is to custom software. For years it lost ground to off-the-shelf products on cost, and its one advantage was exact fit. As building gets cheaper, that fit becomes affordable for far more organisations, and the old choice between buying something general and building something precise begins to even out.

FIG. 3 How AI reshapes each category

CATEGORY	CORE VALUE	HOW AI CHANGES IT
Custom software built for one org	Exact fit to requirements	Cheaper to build, so the market for custom grows
Broad platform wide span	One unified system	Prices fall, and built-in AI keeps increasing
Narrow specialist one task, deep	Best-in-class depth	Many more, and more varied, tools appear — mostly AI ones
Tech + services software + people	Software plus judgment	Routine parts automated; human judgment stays

Figure 3. Each category retains its distinct value while AI changes it in a different way – none is removed, and each is affected on its own terms.

The competitive pressure on the specialist category is worth illustrating, because it is sometimes misread as a verdict on specialists in general. It applies most to tools whose differentiation rests largely on access to a widely available model; specialists with genuine depth are in a different position.

FIG. 4 Where competitive pressure concentrates — undifferentiated tools




Figure 4. Where a tool's value rests mainly on access to a model that is widely available and falling in cost, margins and durability are pressured. Specialists built on depth, proprietary data, and integration are not in this position. Sources: TechTimes/Crunchbase analysis 2026; IdeaProof failure database 2023-26.

Across the market, the overall effect of these shifts is a larger number of tools in active use. That is, in part, a positive development: more capability is available to organisations than before. It also raises a practical question that the next section examines — how the chosen tools work together as one system.

FIG. 5

Tools in use per organisation, over time




Figure 5. The average organisation ran about 10 HR technology modules in 2020 and roughly 26 by 2024 — more capability in use, and more pieces that must work together. **Source:** Sapient Insights Group, 2024–25 HR Systems Survey.

SUMMARY SO FAR

AI reshapes each category differently. The shared result is more capable tools in use — and a larger question of how they work together.

PART 02

Platforms and Specialists: Two Design Philosophies

Two of the four categories — broad platforms and hyper-specialists — represent different answers to the same question: how should a domain's many tasks be served? Each answer is internally consistent, each is valid, and each carries a structural consequence. Understanding both is useful before considering how they work together.

How broad platforms behave

A platform that covers a wide span of a domain distributes its resources — roadmap, research, engineering — across many capabilities. By design, it optimises for breadth and coherence rather than maximum depth in any single area. This is not a shortcoming; it is the nature of the choice. The platform's value is the unified system: shared data, a consistent experience, and a workflow that holds together end to end. The corresponding trade-off is that any one capability may be less specialised than a tool dedicated solely to it.

How hyper-specialists behave

A hyper-specialist concentrates all of its resources on a single capability — a hyper-specialisation — and can therefore reach a depth a generalist rarely matches. That depth is its value, and for an organisation with a demanding need in one area it can be considerable. The corresponding trade-off is that it addresses only one part of the workflow. An organisation that wants best-in-class depth across several areas will, by definition, use several hyper-specialists.




FIG. 7

Each choice has a value and a trade-off

Broad platform

+ GAINS

Unity, coherence, one record, end-to-end workflow

- TRADES

Some depth in any single area

Narrow specialist

+ GAINS

Best-in-class depth in its single area

- TRADES

Coverage of the wider workflow

Figure 7. Both are sound choices for different needs. Neither is strictly better; each optimises for a different objective.


The practical question: who connects the tools?

Because specialists cover one area each, an organisation that wants depth across several areas uses several of them — often alongside a platform. For those tools to function as one workflow, they have to be connected: data has to move between them, records have to stay consistent, and the handoffs between steps have to work. Today, that connecting work commonly falls to the organisation itself. Its own team, or contracted help, builds and maintains the integrations, reconciles data across systems, and absorbs the effort when something does not line up.

This connecting work has a measurable cost, and surveys consistently find that buying tools and connecting them are far from the same thing.

FIG. 8


Tools owned versus tools connected



62% of organisations run 2–4 paid HR tools from different vendors.


Across those stacks, a minority are reported to be usefully connected; about one in five rate their integration weak.

Figure 8. Owning several tools is common; having them work together as one system is less so. The gap is the connecting work. **Source:** HR.com, State of Today's HR Technology & Integrations, 2025.



How this varies by organisation size

The capacity to do this connecting work is not evenly distributed, as the figure below sets out. Large enterprises can fund it directly; small organisations sidestep much of it by standardising on fewer tools. Mid-sized organisations sit between the two: they often need both depth and coverage, yet have the least spare capacity to build the integration by hand.



In each case the underlying question is the same: when an organisation uses several tools, who carries the work of making them function as one system? The next section describes an arrangement in which that work shifts away from the organisation.

Getting the Most Value

For any organisation, the categories examined so far are not only a way to describe the market — they are a set of options. The practical question is which of them now delivers the most value, given that AI has reshaped each. Because all four have changed at once, the answer is no longer what it was a few years ago.

The four options, and the challenge of each

Custom software is more attainable than it was: as AI lowers the cost of building, an exact-fit system comes within reach of more organisations. The challenge is that the organisation must build and maintain it, and may end up recreating capabilities that hyper-specialists already provide. A **broad platform** makes the opposite trade — one coherent system covering the whole workflow, with little to assemble, but rarely the best version of any single capability. A **hyper-specialist** offers the deepest version of one capability, but covers only that area; an organisation wanting depth in several areas needs several of them. A **tech-enabled service** handles work pure software cannot, by pairing it with people, at the cost of that human layer and some direct control.

FIG. 10 The customer’s four options, in this new environment

OPTION	WHAT IT OFFERS	ITS MAIN CHALLENGE
Custom software build it yourself	Exact fit — and now cheaper to build	You build and maintain it yourself
Broad platform one wide system	One system across the whole workflow	Rarely the best at any single task
Hyper-specialist one area, deep	The deepest version of one capability	Covers only that one area
Tech + service software + people	Handles what software cannot, using people	Cost of people; less direct control

Figure 10. Each option delivers real value, and each carries a challenge. Most organisations draw on more than one of them.

No single option is complete

Few organisations rely on only one of these. The exact fit of custom, the coverage of a platform, the depth of a hyper-specialist, and the judgment of a service each answer part of the need, and most organisations end up using a mix. The most value comes not from picking one option, but from getting that mix to work together.

That points to a different way of assembling the stack: an **integrated architecture**, in which one platform ties the chosen tools together so the organisation can use the best of each without holding the whole thing together itself. It is easiest to see as a picture.

FIG. 11 The integrated architecture, across a workflow




Figure 11. One platform runs the workflow end to end and owns the data and experience; specialist tools attach to the relevant stage and can be swapped when a better option appears. The organisation works with a single system rather than a stack of separate ones.

Why the model makes sense

The logic is one of focus. A single product that tries to be both broad and deep tends to be average at most things; a set of deep tools with no centre leaves the organisation to wire everything together. An integrated architecture avoids both. The platform holds the workflow, the data, and the connections together; each specialist goes as deep as it can in its own area. Each part does less, and therefore does it better — and the organisation gets coverage and depth at the same time.

What each part is responsible for

The architecture works only when the boundaries are clear. The platform owns the connective layer. Each specialist owns depth in its own area. The two meet at a shared, open standard that lets tools connect and disconnect without disturbing the rest.

FIG. 12

What the platform owns, what the specialists own




Figure 12. The platform owns the connective layer; each specialist owns depth in its area; the two meet at a shared, open standard. Each keeps its own focus while the whole holds together.

The platform must lead on AI and automation

Holding a whole workflow together is not a passive role, and it cannot be done by hand at any real scale. For that reason the platform has to lead on AI and automation rather than treat them as add-on features.

In concrete terms, the platform uses AI to connect tools that were never designed to align: it maps each tool's data into one shared record, keeps that record consistent as tools write to it, and automates the handoffs between stages so work moves through the process without manual stitching. On that shared record it runs analytics across the entire workflow, and adapts the process to each organisation through customisation. The connecting, the automation, and the analysis are the platform's product — and AI is what makes them possible at scale.

FIG. 13

Where the connecting work sits — the organisation, or the platform




Figure 13. The same specialist tools can be used either way. On the left the recruitment team is the integration point and carries that work; on the right the platform performs the connecting — feasible because AI can map data between systems, keep records consistent, and manage handoffs that once required custom engineering.

The specialists go deep — and connect

Each specialist has the opposite job: to be the best possible version of one capability. It uses AI too, but pointed inward, to reach a depth in its single area that a generalist will not match. How it does so depends on its type. A **pure-software specialist** embeds AI directly in its function — an AI screening tool, for instance, evaluates candidates more thoroughly than a broad platform would attempt. A **software-and-services specialist**, such as a verification or fraud-review service, uses AI to handle the routine, data-heavy work and to support the people who make the judgment calls, so the human effort is spent only where it adds the most.

The two work together because neither tries to do the other's job. The specialist does not try to become the whole system: it connects to the platform through the open standard, contributes its depth at the right stage, and can be removed when something better appears. The platform does not try to out-specialise it: it carries the connecting work so the specialist can stay focused on depth. Three things keep that connection native rather than brittle.

FIG. 14 Three foundations of native integration




Figure 14. A single shared record keeps data coherent; open standards make connection routine rather than bespoke; and treating removal as a first-class capability lets a component be replaced without rebuilding the workflow.

The value of an integrated architecture

Put together, this architecture changes what an organisation can reasonably expect from its software. The value shows up in several distinct ways.

FIG. 15 The value of an integrated architecture

1	Coverage and depth at once	One system across the workflow, plus the best tool where depth matters.
2	The plumbing is not your job	The platform builds and maintains the connections — not your team.
3	One coherent record	Consistent data end to end — the basis for analytics and audit.
4	Freedom to upgrade	Swap any tool for a better one without rebuilding the workflow.
5	Value for the whole market	Organisations, platforms and specialists each gain.

Figure 15. Five ways the value of an integrated architecture shows up in practice, expanded below.

Coverage and depth at once. The long-standing trade-off — a broad platform that is average at everything, or deep specialists that do not work as one — falls away. The organisation runs the whole workflow on one system and still uses the best available tool wherever depth matters.

The plumbing is no longer the organisation's job. The platform builds and maintains the connections, reconciles the data, and manages the handoffs. The team's time goes to the actual work — hiring, in this case — rather than to keeping tools talking to each other.

One coherent record. Because every tool reads from and writes to the same record, data stays consistent across the whole workflow. That is what makes trustworthy analytics, clean auditing, and a single experience possible — and it is the asset that compounds in value over time.

Freedom to upgrade. The best tool in any area rarely stays best for long. Because a tool plugs out as easily as it plugs in, the organisation can adopt a better option in any single area without rebuilding the workflow around it. The stack keeps improving instead of ageing in place, which also makes it more resilient as AI itself advances.

Value for the whole market. The arrangement is positive-sum. Organisations get the best of every category; platforms hold a durable role as the connective layer and system of record, without needing to win every feature comparison; and specialists reach a wider market, adopted within many workflows and competing on depth rather than on whole-system breadth.

IN SUMMARY

The increase in available tools is not, in itself, a problem to be reversed; it reflects more capability reaching organisations than before. The open question is one of architecture — how separately built tools work together as one system, and who carries the work of connecting them. An integrated architecture, in which one platform ties specialists together through native integration, is one well-supported answer to that question.

Hirewand

Hirewand is built around the approach described in this paper. To learn more, reach out to the Hirewand team.

hirewand.com

SOURCES & FURTHER READING

Figures above draw on the following industry research. Statistics are reproduced as factual data points; interpretation is the author's.

1. **Sapient Insights Group** – 2024–25 HR Systems Survey (average HR modules per organisation: ~10 in 2020, ~26 in 2024).
2. **The Josh Bersin Company** – 2026 HR Technology Market Report (large enterprises run multiple HR systems; per-employee HR-tech spend).
3. **HR.com** – State of Today's HR Technology & Integrations, 2025 (62% run 2–4 tools; ~39% usefully integrated).
4. **Gartner Digital Markets** – 2025 HR Software Trends Survey (tool counts; AI adoption expectations).
5. **Second Talent** – HR Technology Stack Statistics, 2026 (higher administration cost; integration overruns).
6. **Torii** – 2025 SaaS Benchmark Report (share of inactive SaaS applications).
7. **EY** – Estimating HR Labor & Non-Labor Costs, 2025 (cost of manual data entry across systems).
8. **TechTimes / Crunchbase** and **IdeaProof** failure database, 2026 (margin patterns; inference-cost decline; commoditisation of undifferentiated tools).
9. **a16z**, **IDC**, **Bessemer State of AI 2025** – differentiation shifting toward proprietary data and workflow integration.
10. **Gartner / Deloitte** – projections on point-product tasks performed by AI agents over the coming years.

AUTHORING MODEL FOR THIS WHITE PAPER

Primary content and multiple reviews by the author, Shiva Maran.

Cleaned, structured, and source collection by Claude.ai.

Hirewand

Fragmentation and Integration · Research White Paper · 2026 · hirewand.com