Insights-First AI

Better and Explainable People Decisions

the joshbersin company

Overview

Over the last decade, we've invested heavily in the talent intelligence market—an ecosystem of AI tools that analyzes people data for recruiting, internal mobility, skills development, and workforce planning. These systems promised a kind of "Google for people," inferring skills and fit from billions of profiles, resumes, and job descriptions.

But this approach has plateaued. Public and internal data are messy, incomplete, and often outdated; most profiles list a job title but say little about real impact, growth, or context. Plugging large language models (LLMs) into this noise doesn't solve the problem.

The next era of AI in HR will not be won by better models but by better, expert-labeled data: clean, reliable people data with "expert labels"—tags that encode recruiter and manager judgment about what success looks like in specific roles and domains.

This report explores this "insights-first" approach, how companies can move to this model, and how HR leaders can drive better talent outcomes while preparing for a new HR tech infrastructure of superagents.

In This Report

From Model-First to Insights-First AI

Human recruiters and managers rely on instinct all the time. They read between the lines of a resume, notice patterns in a career story, and ask probing questions informed by years of experience. Instinct is not perfect, but it is grounded in judgment, lived experience, and accountability. If a hire fails, a human being has to explain why.

AI agents lack this kind of judgment or accountability. When they act, they are not drawing on lived experience. Instead, they are guessing from patterns in the data they have seen. That's the essence of what we call "model-first AI."

Model-first AI reflects the traditional notion of talent intelligence: feed resumes, profiles, and HR records into a model; infer patterns; and make confident recommendations—without showing the underlying evidence. It has enabled faster search and impressive results, broadening the talent pool and discovering hidden gems. But its usefulness has plateaued due to shaky data foundations, creating three problems:

- **Opacity.** AI systems often cannot explain why one candidate was ranked higher than another.
- **Fragility.** Biased, incomplete, or incorrect data—as HR data often is—leads the model to replicate and amplify errors.
- **False confidence.** Recommendations may appear certain even when based on superficial correlations.

Insights-First AI

"Insights-first AI" flips the logic. Rather than relying primarily on pattern recognition over unstructured data, this approach emphasizes designing structured, domain-informed representations before model application. Instead of throwing LLMs at raw resumes, profiles, and inconsistent HR data, insights-first AI initially designs a decision-ready people-data layer comprised of people, roles, teams, attributes, and relationships—and enriches it with expert labels about growth, scope, and impact. Only then does it bring in LLMs to search, match, and make recommendations.

That sounds subtle, but it's the difference between a clever black box and a system you can explain to a CHRO, a regulator, or a candidate. Using an insights-first AI approach in HR means:

- **Treating people data as business intelligence.** People data is more than exhaust from HR systems. The goal is a decision-ready insight, not more records.
- **Designing a clean, structured people-data model.** This includes entities (people, roles, teams), attributes (skills, industries, contexts), and relationships (career moves, reporting lines, project histories).
- **Enriching that model with expert labels.** Data tags encode recruiter and manager judgment about what success looks like in specific roles and domains.
- **Applying AI to this higher-quality layer.** Powerful AI models can now reason using curated, contextualized data instead of raw text.

In other words, insights-first AI produces trusted, explainable talent intelligence, and AI is the way to access and scale that intelligence. Findem is an example of a pioneering company embarking on building an attribute-based, expert-labeled talent graph and then letting recruiters and managers explore it with AI, rather than asking the model to guess everything from scratch.

This matters because every downstream objective—explainability, finding the best candidate, increased talent density, improved people decisions—depends on the structure and quality of the underlying data.

Model-first tools and insights-first tools can look similar on the surface—both show ranked lists of candidates—but under the hood they behave very differently. Model-first tools are essentially pattern-matching on whatever data you give them. Insights-first AI platforms focus first on defining the signals that matter, building an attribute-rich talent graph, and only then applying AI. This means recommendations are explainable and governable, not just impressive. See Figure 1 on the next page.

Figure 1: Model-First vs. Insights-First AI

Aspect	Model-First Talent Intelligence	Insights-First Talent Intelligence
Starting Point	Start with models and hope they learn enough from raw HR data.	Start with the decisions to make and design a clean, explainable people-data layer.
Data	Raw resumes, profiles, and inconsistent HR fields; mainly unstructured text.	Structured attributes (people, roles, teams, relationships) that can be governed.
Judgment	The model quietly infers skills and fit; human judgment lives outside the system.	Recruiter, HR, and manager judgment are encoded as expert labels in the data model.
"Instinct" vs. Proof	The model "acts on instinct" by guessing from correlations in the data.	Decisions are grounded in explicitly defined and reviewable success signals that can be reviewed and refined over time.
Output	Ranked candidate lists with limited explanation of "why" or ability to fine-tune.	Explainable recommendations with clear evidence and signals you can refine.
Risk & Governance	Sometimes hard to audit or explain, easy to overtrust or dismiss.	Auditable and governable; easier to align with fairness and compliance expectations.
Examples	**Recruiting:** candidate lists, sorted and ranked by match to skills and experience. **Internal mobility:** employees matched to opportunities based on titles or manager networks.	**Recruiting:** candidate lists, sorted and ranked by match to skills and experience, with specific criteria to explain reasoning **Internal mobility:** employees matched to opportunities based on shared attributes with high performers and signals that show which drove each recommendation.

Source: The Josh Bersin Company, 2026

Where Explainability Fits In

When AI operates on unlabeled resumes, profiles, and HR records, it often assigns opaque scores and rankings that even experts struggle to unpack. Recruiters and managers see outputs—lists of names—but not the reasons behind them. That is increasingly unacceptable for decisions that affect careers, pay, and access to opportunity—especially given increased legal and compliance requirements (e.g., EU AI Act or NYC Local Law 144).

By contrast, insights-first AI can "show its work" in human terms. Instead of saying "this person is an 87% match," an insights-first system can say: "this candidate looks like your top performers because they have scaled from IC to director in similar companies, built zero-to-one products in your domain, and led cross-functional teams in regulated environments."

In other words, explainability is not about exposing the math of the model but about surfacing the evidence—growth, scope, impact, and context that drove a recommendation.

Insights-first AI matters because:

- **It builds trust.** Recruiters and hiring managers can see and challenge the signals behind each suggestion rather than blindly accepting or rejecting a black box.
- **It supports fairness and compliance.** When success criteria are explicit and traceable, organizations can audit whether AI recommendations are grounded in job-relevant factors instead of hidden proxies.
- **It enables learning.** When people can see which labels and patterns are driving strong outcomes, they can refine those labels over time, improving both the data and the AI.

In essence, the labels and success signals are the *why*; insights-first AI is how recruiters, managers, employees, and leaders see that "why" at the moment of decision (see Figure 2).

Expert-Labeled Data in HR

If data and insights are at the center, what counts as insights in talent? Traditional systems extract generic skills and keywords from profiles and job descriptions. These signals are helpful but shallow, rarely capturing the deeper questions that good recruiters ask, such as:

- What does a great enterprise sales leader look like in a specific business model and stage of growth?
- How do we recognize a standout clinical operations leader in a complex, regulated diagnostics environment?
- How do we decode a veteran's military career into civilian-ready signals like leadership under pressure, secure communications, or logistics at scale?
- Where is the right marketing executive with experience scaling a company in our industry and growth cycle?

These are judgment-heavy questions. Historically, the answers have lived in recruiters' heads and managers' networks. Expert labeling is about turning those answers into structured, machine-readable insights.

To understand why this matters in HR, it helps to look at how expert labeling has become central in other domains.

The Rise of the Data Labeling Economy

The market for data labeling and enrichment has shifted from a niche technical requirement into a multibillion-dollar

Figure 2: Example of Explainable Recommendations in Candidate Assessment

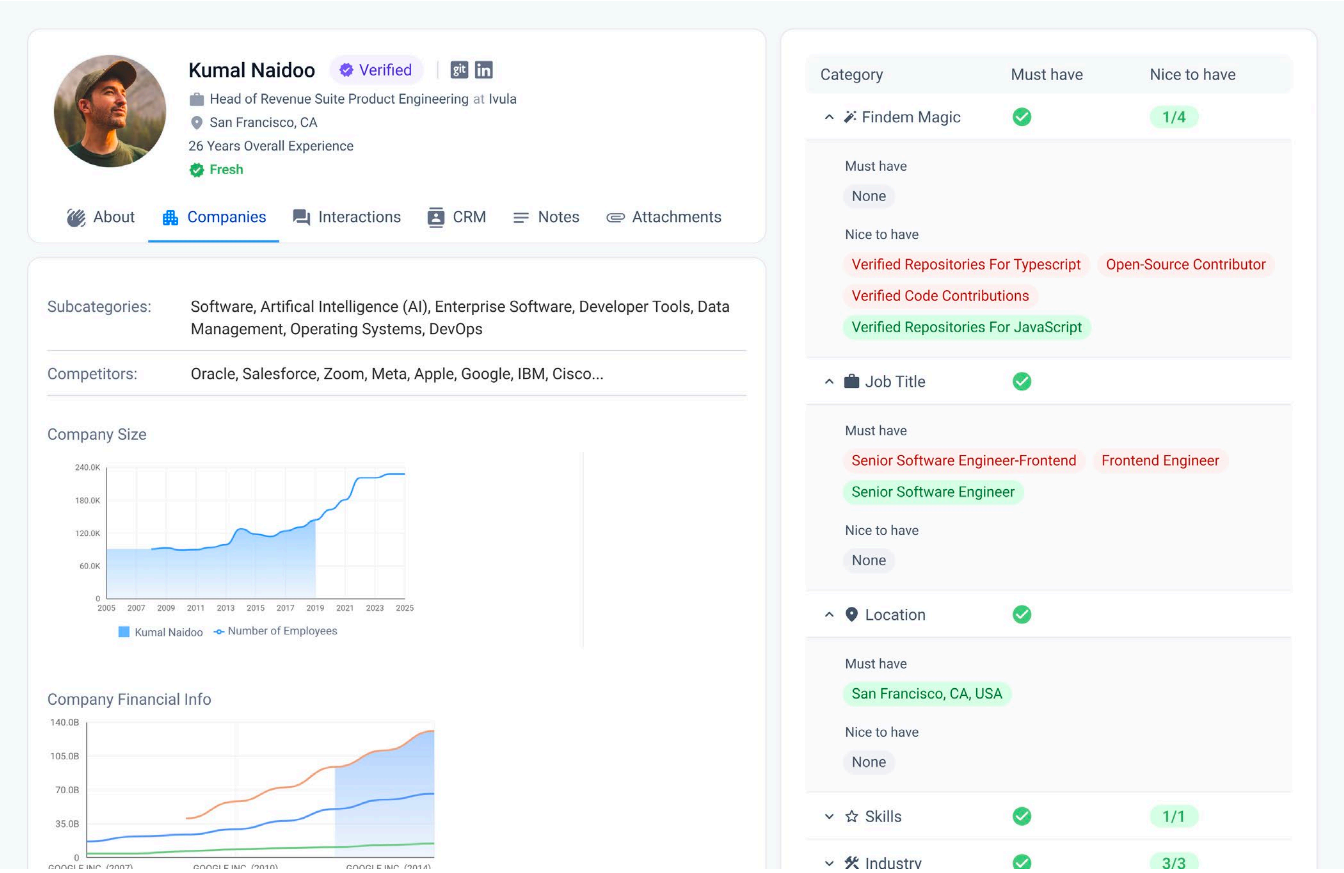

Source: Findem, 2026

cornerstone of the global AI economy. The combined market for data labeling, annotation, and enrichment is estimated to be between $5 billion and $22 billion globally, with examples in consulting, healthcare, and other domains. Data annotation tools alone are projected around $2.5 billion in 2026, growing at roughly 26% to 27% CAGR, and managed labeling services—outsourced teams that provide labeled data as a service—are even larger, projected at around $24.2 billion in 2026, with around 22% growth.[1]

Several structural shifts are driving this:

- **Generative AI and reinforcement learning from human feedback (RLHF).** LLMs depend on RLHF, meaning humans rank, correct, and refine model outputs and provide high-quality reasoning traces and preferences. This involves expert feedback on complex text, code, and decision-making.
- **Specialized domain expertise.** In healthcare, legal, and financial services, annotators need advanced degrees or deep practical experience. Radiology labels, contract clause classifications, and risk assessments cannot be farmed out to generic crowd workers but require subject-matter experts (SMEs), raising the value of each labeled example.
- **Autonomous systems.** Automotive and robotics remain huge consumers of labeled data (e.g., video for self-driving and perception systems), accounting for roughly a quarter of the total labeling market.

Despite advances in autolabeling and "AI labeling AI," human-in-the-loop remains the gold standard. Manual annotation still accounts for over 70% of the market because ground-truth accuracy is critical. Around 80% of labeling work is outsourced, creating a global network of "digital factories" run by providers such as Scale AI, Labelbox, and Appen.[2]

Importantly, the market is pivoting from quantity to quality. The emphasis is shifting from "big data" to "smart data," with a growing subsector focused on data enhancement—cleaning, debiasing, and enriching existing datasets with metadata to make them fit for specific enterprise use cases.

Highly qualified professionals like consultants, radiologists, and lawyers are now labeling data to be ready for specialized tasks. Across domains, the pattern looks similar:

- **Consulting and strategy.** Professionals pose complex, multistep business problems to AI models—for example, how an airline should respond to a safety crisis or how a manufacturer should handle a product recall. They critique the model's reasoning, correct errors, and provide model-ready examples of strong analyses. Over time, the model learns not just answers but also the structure of expert reasoning.
- **Healthcare.** Radiologists and pathologists label images, segment tumors, and classify pathologies, but they also provide free-text rationales and severity assessments. The labels encode expert judgments that would be impossible for a generic model to infer reliably from pixels alone.
- **Legal and compliance.** Lawyers classify clauses, flag risks, and annotate where a contract deviates from a firm's preferred templates. Over time, this creates a labeled corpus that can support clause suggestion, risk summarization, and negotiation support.
- **Autonomous systems.** Engineers and trained annotators label lanes, pedestrians, signs, and edge cases in video and other data. As systems become more advanced, labels become richer: degree of occlusion, likely intent of other actors, or contextual factors like weather and lighting.

In all of these cases, the same dynamic is at work:

- **Human experts provide labels** that turn those patterns into insights aligned with real-world tasks, standards, and risks.
- **Powerful models are used as agents** and assistants to run on top of these models.
- **Over time, the models can handle more of the routine work,** and experts move upstream—to design tasks, supervise quality, and tackle new problems.

The direction is clear: powerful AI systems will increasingly depend on high-quality, expert-labeled data.

---

1 *Data Collection And Labeling Market (2025–2030)*, Grand View Research, 2025.

2 *AI Data Labeling Market Size & Share Analysis - Growth Trends and Forecast (2026–2031)*, Mordor Intelligence, 2026.

Expert Data Labeling in HR

The same pattern emerges in HR. It is no longer enough to have large volumes of resumes and profiles. The question is whether those records are labeled and enriched with the right signals to support high-stakes decisions about people and work.

For insights-first AI in HR, we need expert-labeled data before applying AI models to translate various signals into talent insights (see Figure 3).

These labels do not automatically appear in resumes or profiles but must be designed and curated, often in collaboration with experienced recruiters, hiring managers, and domain experts. Once in place, however, they can be applied at scale, enriching millions of records with human-grade insight.

A useful way to think about insights-first AI is:

> Insights (labeled data) + organizational context = verifiable insights about people in your environment
>
> Verifiable truth + action (AI agents) = trusted, better outcomes in hiring, mobility, and development

This progression—from raw data to labeled evidence to verifiable insight—is where the "insights-first" aspect becomes real. It means treating people data as an intelligence asset to design, govern, and continuously refine.

Finding the Best Fit: Talent Density in Action

"Talent density," which refers to fewer, higher-performing people doing more impactful work, is emerging as a central theme for recruitment and broader talent management. Leaders know they cannot simply hire their way out of every problem, especially in a world of budget constraints, AI automation, and scarce skills. The question is not "How do we get more candidates?" but "How do we consistently hire and develop people who truly raise the bar?"[3]

This approach is drastically different from the integrated talent management approach of the past (see Figure 4 on the next page).

Figure 3: Examples of Signals and Talent Insights

Type of Signal	Sample Talent Insights
Growth	▪ Moved from individual contributor to leading a team of 30 within four years ▪ Repeatedly took on broader scope and complexity across roles ▪ Built zero-to-one products or functions, not just operated existing ones
Context	▪ Worked in hypergrowth SaaS, oncology diagnostics, or highly regulated domains ▪ Navigated global, matrixed organizations ▪ Operated in specific stages of a company lifecycle (early-stage startup, post-IPO, turnarounds)
Impact	▪ Led a successful turnaround of a struggling unit ▪ Launched new business lines or products that reached scale ▪ Took products through regulatory approvals such as FDA or EMA
Domain-Specific	▪ For veterans, translating military occupational specialty (MOS) codes and deployment histories into civilian attributes such as team leadership, logistics coordination, or secure communications ▪ For early talent, converting academic projects, internships, and extracurricular experiences into signals of potential.

Source: The Josh Bersin Company, 2026

3 *The Superworker Organization: AI Goes Enterprise*, Josh Bersin/The Josh Bersin Company, 2026.

Figure 4: Talent Management vs. Talent Density

	Integrated Talent Management	Steadily Increasing Talent Density
What It Is	Hire into a job; set goals; evaluate performance; rank and rate; promote the performers; eliminate the underperformers	Hire for deep capabilities and culture fit; enable people to learn and add value in their job; move people to new roles as needed; hire only when new capabilities are needed
Concept	A supply chain that expands and contracts based on business demand	A team of high performers who move toward high-value opportunities and away from low-performing areas
Hiring	Rapid hiring; some workforce planning; periodic layoffs	Slow, selective hiring for quality and multiplicative impact
Skills	Bring in new skills through hiring	Hire for critical skills but assume most people can learn; invest in internal growth and mobility
Mobility	Hiring and talent siloed within business units, functions, geographies	High-frequency mobility to share talent, develop people, and facilitate agility
Performance Management	Forced distributions, five-point scales, clustered in the middle	Expect high performance from everyone; hire for culture and learning agility; develop or move people to continually increase performance

Source: The Josh Bersin Company, 2026

The gap between model-first AI and insights-first AI in accomplishing talent density is striking. Most recruiting systems today are optimized for volume, increasing the number of applicants, expanding the top of the funnel, and filling requisitions faster. What they rarely do is systematically separate “qualified” from “best.”

- **Qualified:** A candidate who can do the job as described; they add incremental capacity.
- **Best:** A candidate whose performance, attributes, and context suggest they will multiply the value of the role; they raise the bar for the team.

In a traditional candidate slate, the majority of applicants are “qualified,” and a small minority are truly “best,” but they are hard to identify. In an insights-first candidate slate, the overall pool may be smaller, but the share of best candidates is much higher because of the ability to match specific success signals, not just keywords or titles.

Insights-first systems don’t just find more people but more of the right people, with clear reasons that can be examined and refined.

That is what makes talent density real: defining “best” explicitly, searching and matching on those definitions, and measuring how well hiring and talent decisions actually reflect them.

CASE IN POINT

Amplitude Rediscovers Hidden “Best” Talent with Insights-First AI

High-growth tech company Amplitude illustrates the “qualified vs. best” problem. After years of rapid hiring, the company accumulated thousands of candidates in its ATS—silver medalists, people who exited the process, and strong profiles that weren’t quite right at the time. Traditional tools treat this database as a graveyard: hard to search, difficult to separate high-signal candidates from noise, and quickly abandoned in favor of fresh sourcing on platforms like LinkedIn.

Continued →

With expert-labeled, attribute-based search provided by Findem, the picture changes. Instead of filtering on old requisitions and keywords or using model-first AI, Amplitude can rediscover candidates based on special attributes and fine-tuned criteria like:

- Demonstrated experience in product-led growth environments
- Domain expertise in analytics or observability
- Evidence of thriving in similar company stages and business models

Amplitude describes this as a shift from "more profiles" to "more insight." Plugging AI into weak data changes little; what matters is having the right signals to search on in the first place.

As Brett Coin, VP of Talent Acquisition at Amplitude explained:

"We don't need AI to show us more of the same candidate profiles. What changed the game for us was seeing why someone was a great fit—growth, context, impact—and being able to search on those signals directly."

In practice, recruiters can build slates of candidates who already know the space, have shown they can grow with a fast-moving business, and may have engaged with Amplitude in the past—all without running a new campaign to source external talent.

Once "best fit" candidates were identified, Findem enables recruiters to do personalized outreach, increasing the engagement with them exponentially. This approach both reduces external spend and increases the share of best candidates in each slate, because the search is anchored in outcome-relevant labels rather than generic keywords.

The Impact of Insights-First AI

For many HR leaders, the conceptual case for evidence-first AI is intuitive. The challenge is proving its value in measurable terms. How do you know that the right data, labeled and cleaned, with powerful AI agents creates better results?

There are four practical levers where you can see the impact of insights-first AI:

- Controlled experiments
- Applicant tracking system (ATS) rediscovery
- Hard-to-fill roles
- Inclusion of underrepresented talent pools

Controlled Experiments

One way to determine the power of labeled data is to run controlled experiments comparing three approaches:

- A labeled, evidence-first talent intelligence engine
- Traditional tools operating on unlabeled resume databases or public platforms
- General-purpose LLMs applied directly to raw profiles

Start by selecting a defined set of roles or candidate pools and defining success signals upfront (e.g., prior success in similar contexts, growth trajectory, or specific domain depth). Then:

- Generate candidate slates from each approach
- Blind-review them with experienced recruiters and hiring managers
- Score the slates on relevance, quality, diversity of backgrounds, and presence of best candidates

You can also track downstream outcomes: which candidates move to interviews, offers, and hires; how quickly decisions are made; and whether hiring managers express greater confidence in the recommendations.

For example, the head of talent at a venture capital firm working with Findem reported significant value from combining expert-labeled data with AI agents to source candidates for their investment companies. Using deeper criteria (e.g., growth stage of the company, specific role responsibilities, and cultural fit) produced stronger results than traditional tools or generic LLMs.

ATS Rediscovery and Reduced External Spend

Another tangible metric set involves ATS rediscovery and reduced external spend. Rather than sourcing candidates externally, expert-labeled data help identify the best fit candidate already in the database and reconnect with them in a highly personalized way.

Companies like Amplitude have approached this in three steps:

- **Step 1: Establish a baseline.** Determine what percentage of hires in a given period come from ATS rediscovery versus new external sourcing, and how much budget is spent on job boards and LinkedIn.
- **Step 2: Apply expert-labeled, evidence-first search.** Use it on the ATS for a defined set of roles.
- **Step 3: Measure impact.** Track changes in the share of hires from rediscovery, quality of hire indicators, time-to-fill, and external spend.

This approach is particularly compelling in environments where past candidate pools are large and untapped.

Hard-to-Fill Roles and Specialized Domains

For organizations with highly specialized, hard-to-fill roles, the value case is often qualitative but powerful. Comparing traditional and insights-first approaches (see Figure 5 on the next page) can make the difference clear.

In hard-to-fill roles and specialized domains, metrics beyond just raw time-to-fill alone will be key:

- Time to first high-quality slate
- Number of credible options per role
- Manager satisfaction with candidate quality

CASE IN POINT

A Healthcare Innovator Finds Scarce, High-Impact Talent

For a high-growth healthcare innovator focused on oncology diagnostics and precision medicine, the talent challenge is not volume; it is extreme scarcity and complexity. A single role may require:

- Deep scientific credentials and experience with specific technologies
- Proven ability to navigate clinical trials, regulatory submissions, and payor dynamics
- Comfort working at the intersection of research, clinical practice, and commercial execution

On paper, many candidates can look similar, especially when viewed through generic titles and skills lists. In reality, only a small fraction has the combination of context and impact needed to truly succeed. Identifying them requires nuanced understanding of the experiences and attributes that define “best.”

The company worked with Findem’s insights-first AI to encode these nuances such as:

- Specific diagnostic or biomarker experience
- Prior involvement in taking products from lab bench to clinical deployment
- Cross-functional leadership in clinical, regulatory, and commercial teams

Instead of repeatedly returning to the same well-known names, recruiters can now search and rank on these deeper signals, surfacing candidates who may have been invisible in traditional search but are exceptionally well suited to the work.

The result is smaller but much stronger slates, better use of scarce recruiting capacity for roles where every hire matters, and higher manager satisfaction with candidate quality.

Figure 5: Traditional vs. Insights-First AI Approaches

Traditional	Insights-First AI
Repeated, relationship-driven outreach to a small, known pool	Holistic coverage of people with specific scientific, regulatory, and cross-functional signals
Big talent pool, long time-to-hire	Smaller but richer talent pools
Inconsistent coverage of the full relevant market	Reduced time and effort for hiring managers

Source: The Josh Bersin Company, 2026

Inclusion of Underrepresented Talent Pools

For underrepresented talent pools, like veteran and early-career pipelines, impact can be measured in:

- Increases in the number of pools of surfaced candidates for relevant roles
- Improvements in conversion rates at each stage of the funnel (screen, interview, offer, hire)
- Positive feedback from candidates and hiring managers regarding fit and clarity of role-mapping

Because structured labels surface relevant patterns and make assumptions explicit, insight-first systems can open doors for individuals previously overlooked due to resume format, jargon, or nontraditional paths.

CASE IN POINT

RecruitMilitary Connects Veterans and Corporations with Insights-First AI

Veteran-hiring highlights another dimension of the talent density challenge: a highly capable, often underleveraged talent pool that doesn't fit neatly into civilian categories.

Organizations frequently struggle with where to place veterans. Military roles and employment codes do not map cleanly to corporate job architectures; titles feel unfamiliar; experiences like deployments and command responsibilities are difficult to translate into business roles.

"Veterans can feel like a square peg in a round hole, even when they bring exactly the leadership and operational capabilities companies actually need," as Tim Best, CEO of RecruitMilitary, explains.

An expert-labeled, insights-first approach tackles this issue by converting decades of veteran recruiting experience into structured attributes. Veteran-focused platforms can enrich millions of profiles and express military experience in civilian-ready terms—leadership under pressure, operational excellence, secure communications, logistics at scale—so employers search for what actually matters rather than trying to decode military resumes.

The impact is twofold:

- Employers gain access to a larger share of top candidates from the veteran pool—not just those whose resumes happen to align with familiar job titles.
- Veterans are evaluated based on genuine capabilities—not on whether a hiring manager recognizes a specific unit or acronym.

Using Findem, RecruitMilitary has facilitated more than 600,000 connections between veterans and corporations through its Veteran Talent Source solution. This scale is made possible by expert-labeled data combined with AI agents designed to match capability to opportunity.

Superagents and Insights-First AI

The next wave of AI in HR will not be defined by chatbots or isolated point solutions. Instead, organizations are deploying "superagents"—AI systems that operate across platforms and help leaders, HR teams, and employees make better decisions in real time.

Rather than optimizing individual tools, superagents connect previously siloed systems systemically— creating a coordinated intelligence layer across the enterprise. The difference is similar to moving from driver-assist features to a self-driving car. Power steering and park assist are helpful, but an integrated system that navigates end-to-end is transformative.

An HR superagent might:

- Suggest internal candidates for an open role or critical project
- Recommend development paths based on emerging skill demands
- Identify teams with fragile talent density or overreliance on a few linchpin contributors
- Simulate the talent implications of entering a new market, automating a function, or restructuring a division
- Audit promotion and pay decisions for fairness and job-relevant criteria

This vision is compelling, but it cannot be achieved responsibly with fragmented, low-signal data.

If a superagent draws from inconsistent ATS records, static org charts, ambiguous skills tags, and unstructured performance notes, it will generate shallow insights—and sometimes execute the wrong actions with unwarranted confidence.

This shifts the HR tech infrastructure from an HCM-centric system to a systemic, employee experience–first stack, powered by AI (see Figure 6).

Figure 6: Future HR Technology Stack

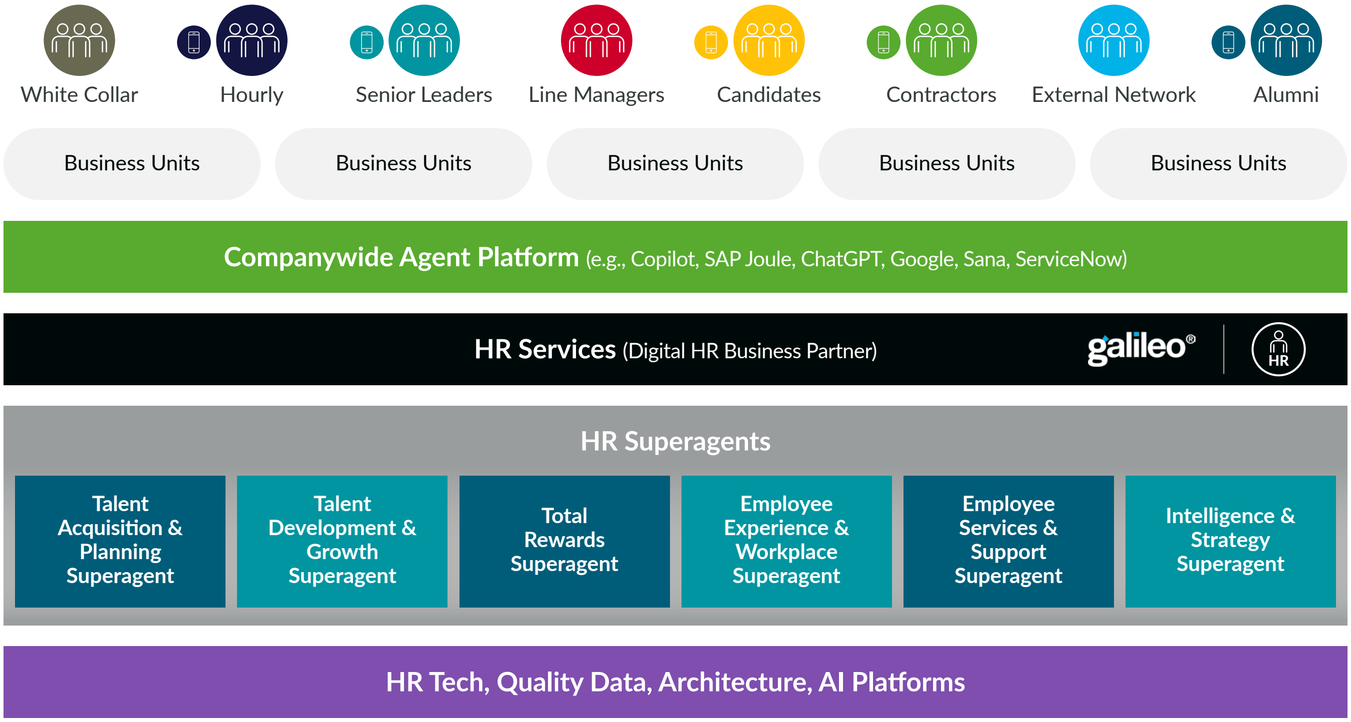

Source: The Josh Bersin Company, 2026

For superagents to operate systemically and responsibly, organizations need a unified, high-quality data layer (see purple layer in Figure 6) positioned as a foundation for AI agents. This layer includes:

- A **structured talent graph** that represents people, roles, teams, organizations, and relationships across the enterprise
- **Expert-labeled attributes** that capture growth trajectory, context, impact, and domain-specific signals—not just static skills or titles
- **Clear governance and data lineage,** so every label is traceable to who defined it, how it was applied, and how it influenced decisions.
- **Embedded risk management,** including bias audits, performance monitoring, explainability standards, and meaningful human oversight

The same insights-first, explainable intelligence layer that improves sourcing also supports higher-quality decisions across internal mobility, learning, workforce strategy, and DEI (see Figure 7).

Rather than optimizing isolated processes, it creates a consistent, transparent framework across the talent lifecycle.

Our research on people analytics reinforces this pattern: organizations that combine high-quality people data with business outcomes—and make that data easily accessible to managers—are significantly more likely to achieve strong financial performance, innovation, and adaptability.[4]

Figure 7: How Insights-First AI Supports Talent Decisions

Lifecycle area	Sample Question	How Insights-First AI Helps
Talent Acquisition	Who are the strongest candidates, not just the first to apply?	Searches and ranks candidates using explicit growth, scope, and domain attributes, making it clear why each candidate surfaced and enabling faster focus on the top of the slate
Internal Mobility & Careers	Which internal people resemble our highest-performing external hires for this role?	Matches employees to opportunities based on shared high-impact attributes, not just titles or manager networks—and shows the signals driving each recommendation
Learning & Skills Development	Which skills and experiences actually improve performance?	Tracks how attributes evolve after learning journeys, rotations, or stretch assignments, and identities which correlate with measurable outcomes
Workforce Strategy & DEI	Where is talent density strong or at risk, and who is being overlooked?	Maps high-impact attributes across teams and segments, identifies capability gaps and audits who is considered for critical roles using transparent, job-relevant criteria

Source: The Josh Bersin Company, 2026

4 *The Definitive Guide to People Analytics: Systemic Business Analytics*, Stella Ioannidou and Kathi Enderes, PhD/The Josh Bersin Company, 2024.

Next Steps

Moving from today's model-first experiments to insights-first talent intelligence is not a one-off project. It is an architectural shift. Organizations can begin with a focused set of actions:

- **Prioritize and pilot a small number of high-value use cases** (e.g., critical roles, ATS rediscovery, veteran hiring) and define concrete success signals for each.
- **Design or refine a systemic people data model** that can store expert-labeled attributes consistently across core HR systems, integrating internal and external insights as needed.
- **Extend the labeled evidence layer** into internal mobility, learning, and workforce planning, ensuring its recommendations surface its underlying signals.
- **Upskill HR and people leaders** to leverage insights-first AI technologies to strengthen talent density and improve quality of people decisions.
- **Establish governance** to regularly review, validate, and adjust labels and AI-supported decisions for fairness, accuracy, and business impact, with specific focus on human oversight, bias mitigation, and drift monitoring.

Key Takeaways

- Traditional, model-first AI uses noisy, unlabeled HR data—making it opaque, fragile, and often overconfident in high-stakes people decisions.
- Insights-first AI treats people data as business intelligence, builds an expert-labeled talent graph, and then applies AI—making recommendations explainable and governable.
- Globally, data-labeling and enrichment have become multibillion-dollar markets, reflecting that expert-labeled data—not just models—is the core success factor in effective AI systems.
- In talent acquisition, insights-first AI converts ATS "graveyards" and undertapped pools into high-density pipelines by surfacing the candidates based on labeled patterns of growth, context, and impact.
- The labeled evidence layer and insights-first AI enable better decisions in internal mobility, learning, workforce strategy, and DEI by making success signals and trade-offs transparent across the lifecycle.
- HR can immediately start by defining success signals for a few critical use cases, refactoring people data to support expert labels, piloting insights-first AI, and deploying governance around labels and AI-supported decisions.

About the Author

Kathi Enderes, PhD

Kathi is the senior vice president of research and global industry analyst at The Josh Bersin Company, supporting clients and the market with evidence-based insights on all areas of HR, learning, talent, leadership, culture, analytics/AI, and HR technology. Kathi has more than 20 years of global experience as a human capital consulting leader with IBM, PwC, and EY, and as a talent and organizational performance executive with McKesson and Kaiser Permanente. Before joining The Josh Bersin Company, Kathi led talent and workforce research at Deloitte, empowering clients with advisory and insights. She is a frequent keynote speaker, author, and thought leader. Her passion is to make work better and more meaningful. Originally from Austria, Kathi has worked in Vienna, London, San Francisco, and Spain and now lives in Palo Alto, California. Kathi holds a doctoral degree in mathematics and a master's degree in mathematics from the University of Vienna, Austria.

The Josh Bersin Company Membership

The Josh Bersin Company provides a wide range of research, tools, and advisory services to help HR leaders and professionals address the ever-evolving needs and challenges of today's workforce. We cover all topics in HR, HR technology, talent, and corporate learning, with a special focus on the professional development of HR teams.

Our corporate membership program provides HR leaders and teams with the skills, strategies, benchmarks, and insights to build cutting-edge HR and people strategies through research, assessments, professional development, exclusive events, and community. Corporate membership also includes access to Galileo®, the world's first AI-powered expert assistant specifically developed for HR. Trained on 25 years of The Josh Bersin Company's research, insights, and expertise and enriched by carefully curated material from our trusted content partners, Galileo unlocks information from over 50,000 verified assets to answer any HR-related question with timely and meaningful answers. Galileo Suite also includes Galileo Learn, which provides hyperpersonalized, up-to-the-minute learning in the most dynamic, interactive HR professional development solution available today. Built on the legacy of Josh Bersin Academy—and its mission to build skills and capabilities through the power of shared stories—Galileo Learn brings advanced AI to the world's most comprehensive and ever-growing set of business insights, best practices, industry trends, and case studies.

To license Galileo Suite for individuals and small teams, please go to getgalileo.ai, and visit joshbersin.com/membership for more information about our corporate membership featuring Galileo and advisory support from Josh Bersin Company experts.