synq

The Blueprint for Agentic HR

Moving beyond chatbots to agents that can reason about your workforce

A guide for HR technology leaders



Executive summary

Every HR technology vendor now claims to have “AI agents.” Most of them are lying, or at best confused about what the term means.

Here’s the uncomfortable truth: hooking a large language model up to your HRIS API and wrapping it in a chat interface does not make an agent. It makes a slightly smarter chatbot that can look things up and push buttons. It can answer “what’s my PTO balance?” and maybe submit a time-off request. That’s useful. But it’s not what’s going to transform how your organization manages its workforce.

The difference between an AI that executes transactions and an AI that makes intelligent workforce decisions comes down to **two architectural choices that most vendors are getting wrong:**

01

What the agent knows

Without a Workforce Context Engine (a living, connected understanding of people, skills, roles, careers, and organizational dynamics) an agent has no basis for judgment. It can retrieve data. It can execute commands. But it cannot reason about your workforce. And reasoning is where the value is.

02

Where the agent lives

HR technology has historically suffered abysmal adoption because it lives in separate portals that employees visit reluctantly, if at all. Agents that live in vendor portals will meet the same fate. The agents that actually transform work are the ones that live inside Teams, Slack, and Microsoft Copilot, the tools people already have open all day.

This guide explains why these two choices matter more than anything else, what the technology actually looks like under the hood, and how HR technology leaders should evaluate the flood of “agentic” offerings hitting the market.

Anatomy of an agent

Synq is an AI-native HCM built around a Workforce Context Engine. It runs on top of the HR systems you already have, or as your HCM. Here is what sits underneath every answer and every action.

Knowledge graph

A living map of your workforce. People, skills, roles, and relationships stored as connected entities, so Synq can follow connections no relational database can answer.

Intelligent tools

Purpose-built capabilities Synq uses to act: matching, skills clustering, career trajectory modeling, skill inference, workforce embeddings, and market intelligence.

Memory and personalization

Remembers what happened across conversations and over time. Recommendations, tone, and priorities adapt, so no conversation starts from zero.

Retrieval and embeddings

Models trained on workforce semantics: career trajectories, skill adjacencies, and organizational context. Not generic embeddings trained on web text.

Business rules and permissions

Your policies, workflows, and approval chains, encoded as rules Synq has to follow. Every action runs with the permissions of the person asking, so a recommendation never sees, or acts on, more than that person could.

Contents

01 The three tiers of HR agents 5
And why only one changes everything

02 The Workforce Context Engine 9
The missing layer

03 The surface problem 15
Why HR tech adoption has always failed

04 The real architecture of agentic HR 20

05 Governance, risk, and the EU AI Act 25

06 Evaluating vendors 29
The questions that matter

07 Building your roadmap 32



The three tiers of HR agents

And why only one changes everything

Josh Bersin, whose research has shaped how the industry thinks about HR technology for two decades, has outlined a progression that HR technology leaders need to internalize. Not because it's academic theory, but because it reveals where real value lives, and where vendors are trying to sell you something that sounds like the future but is actually the past with a new label.

1

Informational

Answers questions, retrieves data, summarizes documents.

2

Transactional

Executes actions and moves data between systems.

3

Intelligent

Reasons about your workforce and makes calls that need judgment.

1 Informational agents

Answer questions | Retrieve data | Summarize documents

> *"What's our parental leave policy in Germany?" "Show me headcount by department." "Summarize this candidate's resume."*

These agents connect an LLM to your knowledge base and data systems. They're essentially search engines with natural language interfaces. They're genuinely useful: they save time, reduce HR service desk tickets, and give employees faster answers.

What they require technically: an LLM, a retrieval layer (RAG over policy documents and FAQ databases), and read access to your HRIS.

The honest assessment: this is a solved problem. Every major HCM vendor, every HR service delivery platform, and most IT service management tools can do this today. If a vendor is selling you an "AI agent" that mostly answers questions, they're selling you 2024 technology at 2026 prices.

2 Transactional agents

Execute actions | Push buttons | Move data between systems

> *"Submit my time-off request." "Schedule this interview for Thursday." "Update this job posting and publish it." "Kick off the onboarding workflow for this new hire."*

These agents can call APIs, trigger workflows, and complete multi-step processes. They combine an LLM's ability to understand intent with tool use to take action across systems. Bersin's research describes this as "process automation," and it's where most vendors are racing to build.

What they require technically: an LLM with tool use, API integrations to your HCM, ATS, and LMS, workflow orchestration, and permission management.

The honest assessment: transactional agents are valuable. They eliminate manual steps, reduce processing time, and handle high-volume operational work. ServiceNow, which acquired Moveworks, excels here. Workday's $1.1 billion acquisition of Sana is largely a play for this tier, letting employees ask questions, run reports, and invoke transactions through a conversational interface.

But here's the thing: transactional agents don't make decisions. They execute decisions that have already been made. Scheduling an interview is not the same as deciding which candidate should be interviewed. Submitting a transfer request is not the same as knowing which role would best develop an employee's career. Publishing a job posting is not the same as understanding whether the role should exist at all, or whether the work could be shared across existing team members with adjacent skills.

3 Intelligent agents

Reason about your workforce | Make recommendations that need judgment | Orchestrate multi-factor decisions

"Given the restructuring of the EMEA sales organization, which of the 340 affected employees can be redeployed to open roles based on their skills, development trajectory, location, compensation, licensing, and career preferences, and what reskilling would bridge the gaps?"

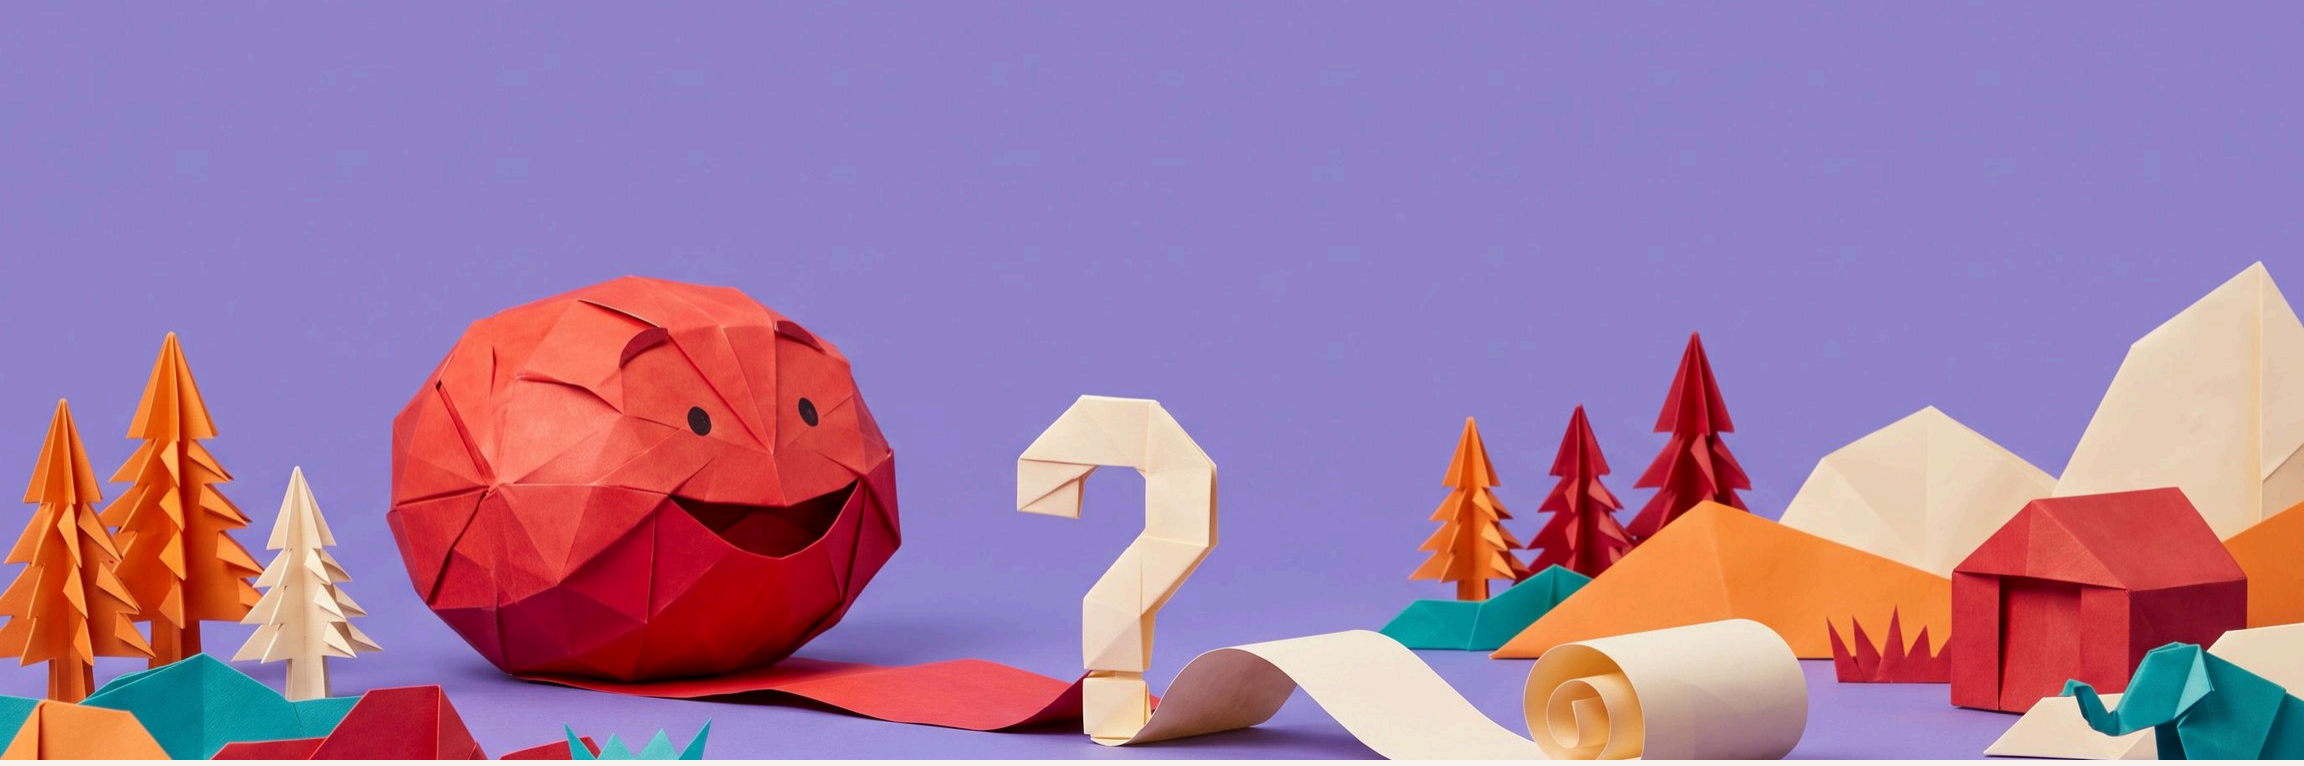
"This engineering manager has a 73% retention risk. Based on their skill profile, the career goals they've shared, internal opportunities that match their trajectory, and comparable market pay, here are three interventions ranked by likelihood of success."

"Based on how AI is reshaping the actuarial function, here's how the role architecture should evolve over 18 months, which skills the current team needs to develop, and which adjacent talent pools could fill the emerging gaps."

This is what Bersin calls **Superagents**: autonomous, outcome-focused systems that coordinate several specialized capabilities toward complex workforce goals. He uses the autonomous vehicle analogy, the shift from driver-assistance features to fully self-driving systems. The intelligence isn't in any single function. It's in the orchestration, the context, and the judgment.

What they require technically: everything from Tiers 1 and 2, plus something most vendors don't have and can't easily build: **a Workforce Context Engine.**

And this is where the architecture diverges.



The Workforce Context Engine

The missing layer

Why LLMs alone can’t reason about your workforce

Large language models are extraordinary at understanding language, generating text, and following instructions. But they know nothing about your organization. They don’t know that Maria in finance has been building data engineering skills through three internal projects over the past year. They don’t know that the skills for your “Senior Product Analyst” role overlap 70% with your “Growth Strategy Manager” role. They don’t know that your Singapore office has a cluster of people with emerging AI skills who aren’t being used. They don’t know that when you lost your last VP of Engineering, the two best internal successors were in completely different functions.

An LLM without workforce context is like a brilliant consultant on day one: articulate, logical, capable of impressive analysis, but disconnected from the relationships, history, patterns, and dynamics that decide whether a workforce decision is good or catastrophic.

This is why most “AI agents” in HR will plateau at Tier 2. They can retrieve and transact because those operations need system access, not understanding. But reasoning about workforce decisions, the ones that actually move talent outcomes, needs a connected, continuously updated, deeply structured understanding of how people, skills, roles, and organizational dynamics relate to each other.

That understanding is what a Workforce Context Engine provides.

What a Workforce Context Engine actually is

It’s not a database. It’s not a data warehouse. It’s not a skills taxonomy sitting in a spreadsheet. It’s a living model of your workforce with five interlocking parts.

1

A knowledge graph of your workforce

A connected map of people, skills, roles, teams, projects, career paths, and org structure, held not as flat records but as a graph where relationships come first. The graph captures not just “Maria has Python skills” but “Maria built Python skills on Project X, which is adjacent to the data engineering skills Role Y needs, which reports into the team that’s expanding in Q3.”

That’s fundamentally different from a traditional HRIS. An HRIS stores records. A knowledge graph stores relationships. And in workforce decisions, **the relationships are where the intelligence lives.**

2

Workforce-specific AI models

Generic LLMs are trained on internet text. They understand language. They don’t understand career trajectory patterns, skill adjacency patterns, or mobility patterns. A Workforce Context Engine runs on models trained on career trajectories and skill adjacencies across millions of employees: models that know a move from financial analysis to product management is a well-worn path, that certain skill combinations predict success in leadership roles, and that skills decay at different rates in different domains.

This matters in an evaluation: **a vendor that started with workforce data has a fundamentally different AI foundation than one that started with a generic LLM and bolted on HR features.** It shows in the quality of every recommendation, match, and prediction.

3

Matching and analysis tools

Purpose-built tools an agent can call: skills clustering, career trajectory modeling, semantic workforce search, compensation benchmarking, retention risk scoring, skills gap analysis. These aren't generic LLM features. They're specialized tools refined on workforce data over years.

When an agent needs the best internal candidates for a role, it doesn't run a keyword search. It calls a matching engine that understands skill adjacencies, development velocity, career goals, team dynamics, and dozens of other signals.

4

A governance engine

Policies, approval chains, compliance requirements, and business rules, encoded not as documentation an agent might or might not follow, but as constraints that are enforced in the system. An agent working inside a governance engine cannot recommend a pay change that breaks pay equity policy, or suggest a transfer that ignores licensing rules, because the constraints are part of how it runs.

This is critical for EU AI Act compliance (covered on page 25). Policy-as-code is not optional when regulators require documented, auditable decision constraints for AI used in employment.

5

Persistent memory

A career coaching agent that forgets every conversation is useless. A Workforce Context Engine keeps interaction history, preferences, development progress, and context across sessions and over time. When an employee comes back six months later, the agent remembers what was discussed, what they were working toward, what's changed in the organization since, and how their skills have grown.

The context gap in the current market

Here’s the landscape as it stands.

Legacy HCM suites (Workday, SAP, Oracle)

Rich employee data, but optimized for record-keeping and transactions, not the connected intelligence that powers reasoning. Workday’s acquisition of Sana adds a conversational, agent-building layer, but the data model underneath is still records, not relationships.

Horizontal AI platforms (Microsoft Copilot Studio, ServiceNow)

Powerful agent infrastructure and broad enterprise reach, but no native understanding of workforce dynamics. They can orchestrate tools and manage conversations. But when the agent has to judge whether a skill gap is bridgeable or a career move is viable, the platform has nothing to draw on except what it can retrieve from other systems. Retrieval is not understanding.

LLM providers

Massive general intelligence, zero organizational context. OpenAI, for example, has announced a jobs platform with partners like Indeed, Walmart, and Accenture. These tools can match keywords in resumes to keywords in job descriptions. They can’t reason about your specific workforce.

Workforce intelligence built in from the start

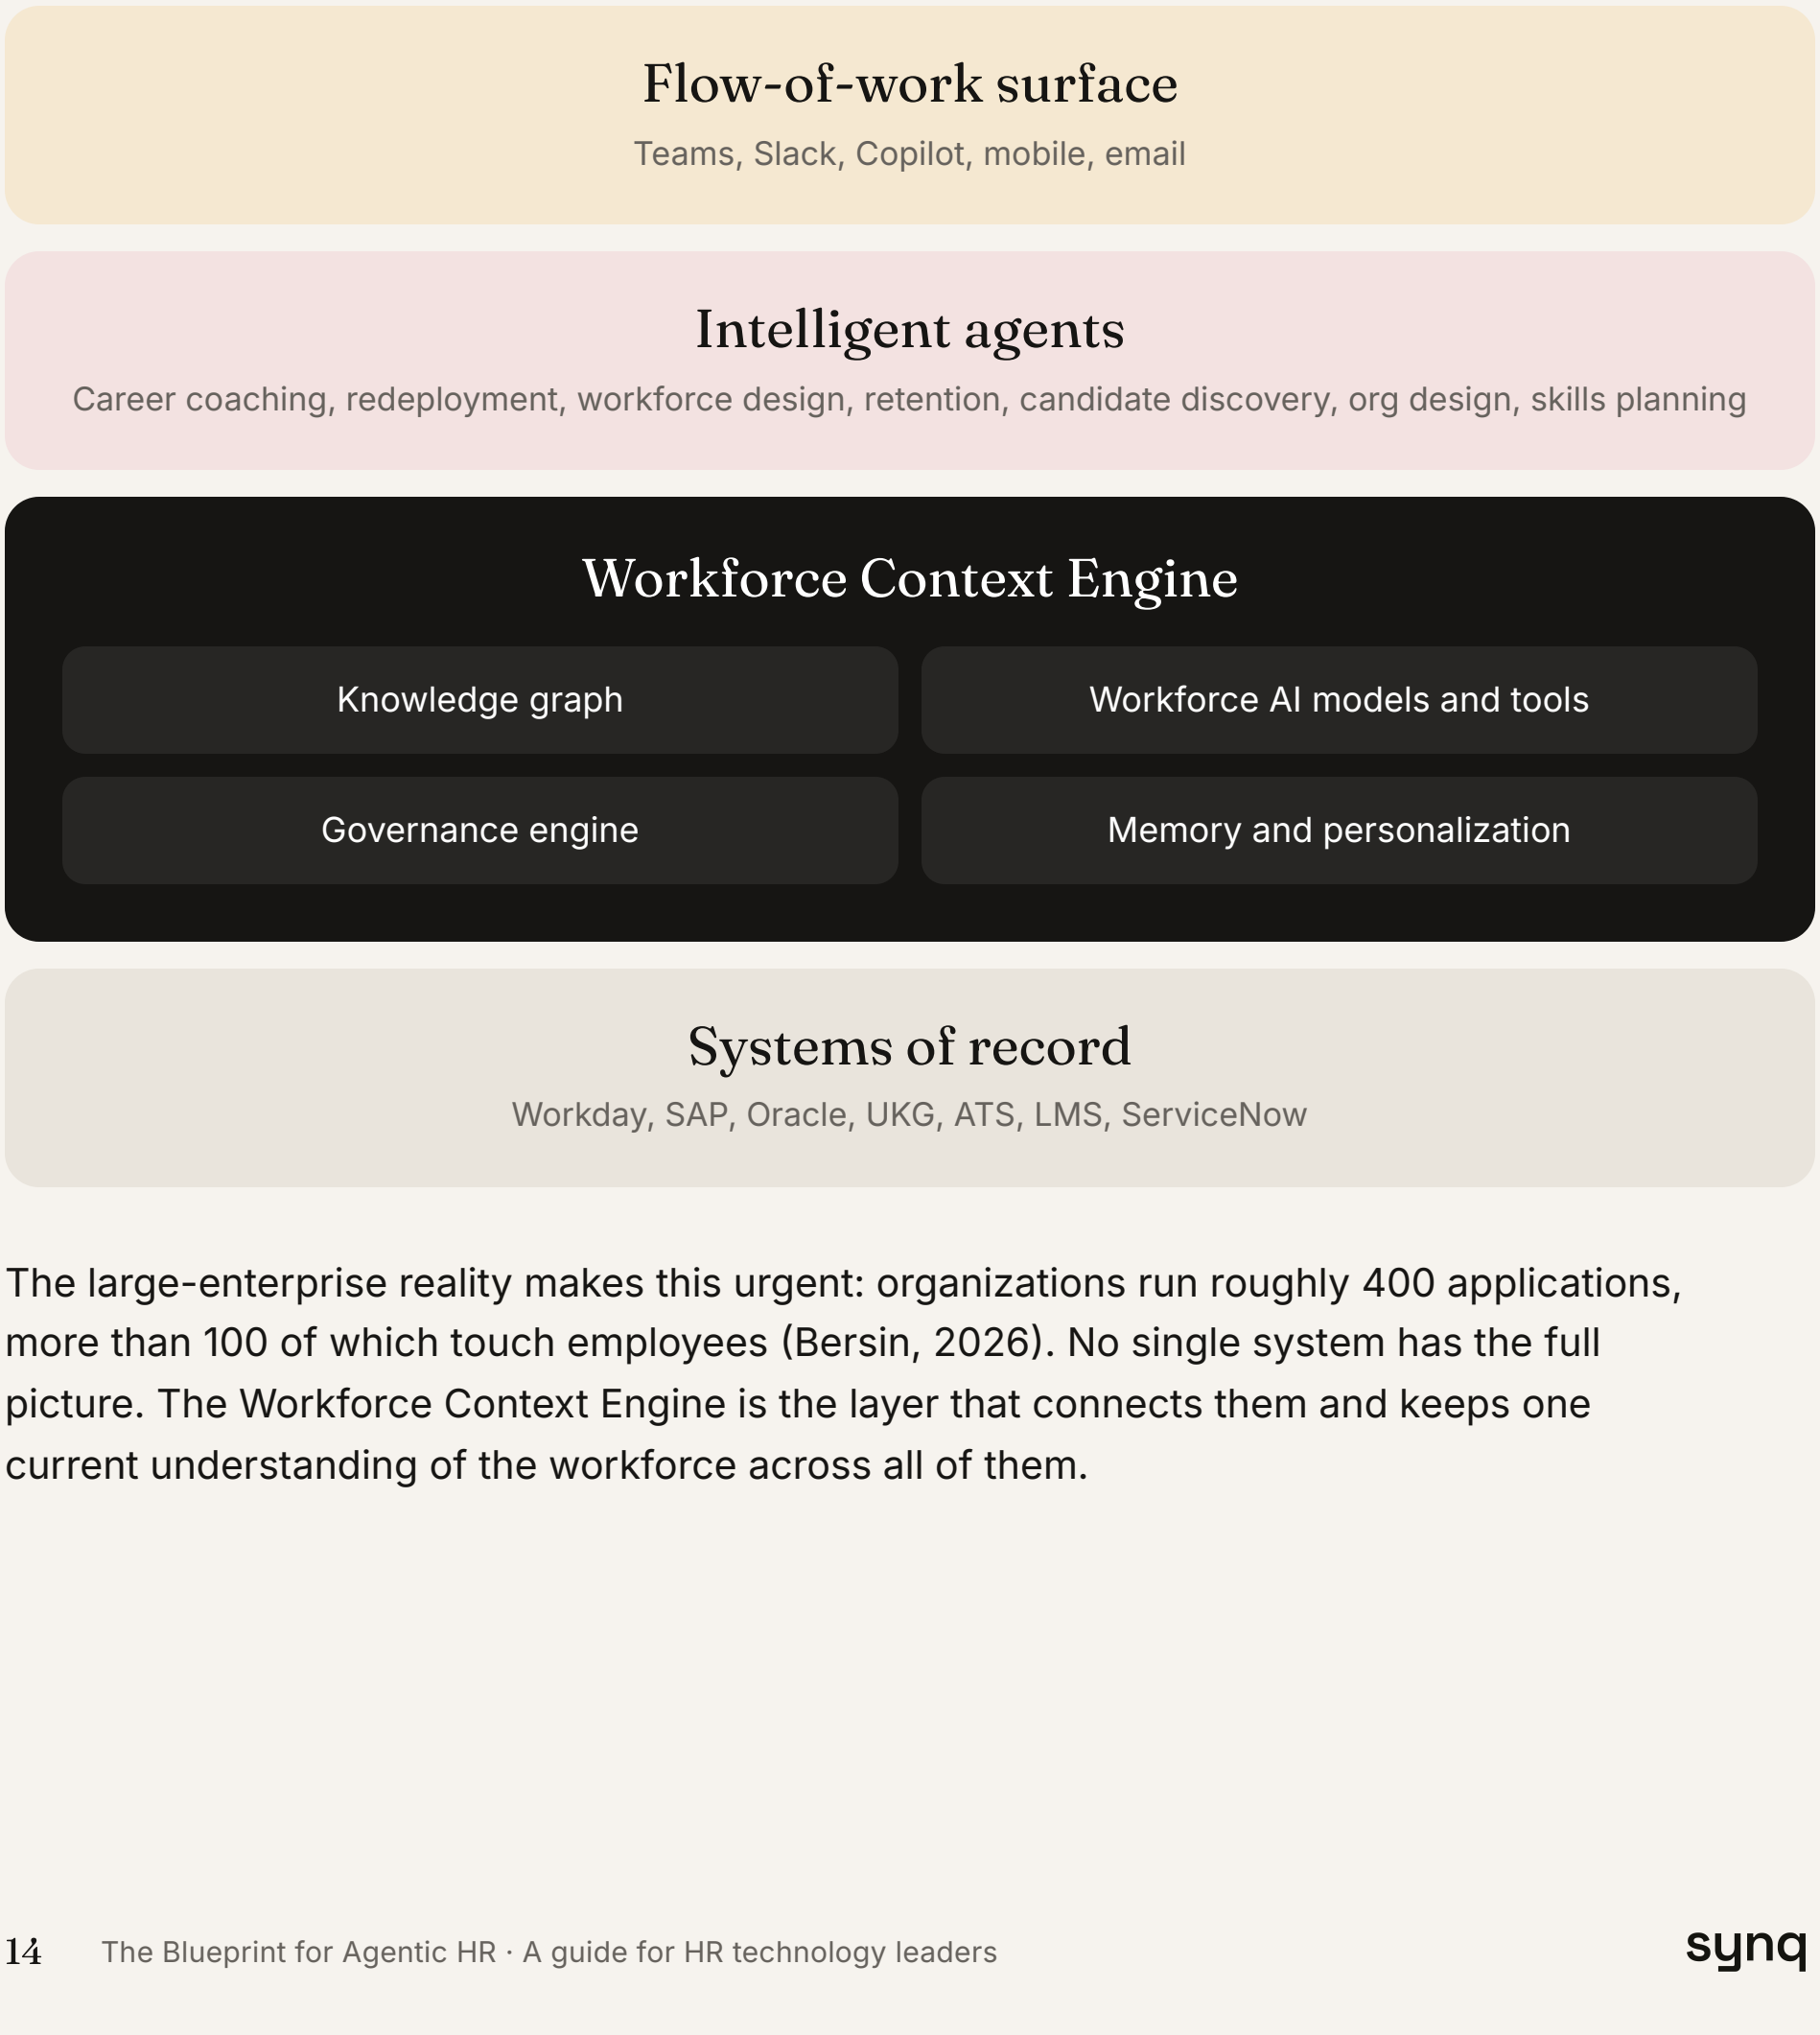
The few platforms that have spent years building exactly this: the connected workforce data, the specialized models, the matching engines, the graph. Synq is built this way. Its Workforce Context Engine rests on a decade of enterprise workforce AI across 5 million+ employees in 112 countries, with models trained on career trajectories and skill adjacencies rather than generic web text. That isn't a feature you get by adding an LLM to a database. It takes years and millions of data points to build.

And because HR can describe a new process, field, or workflow in plain language and Synq builds it on the spot, the context grows with the organization instead of waiting on a vendor roadmap.

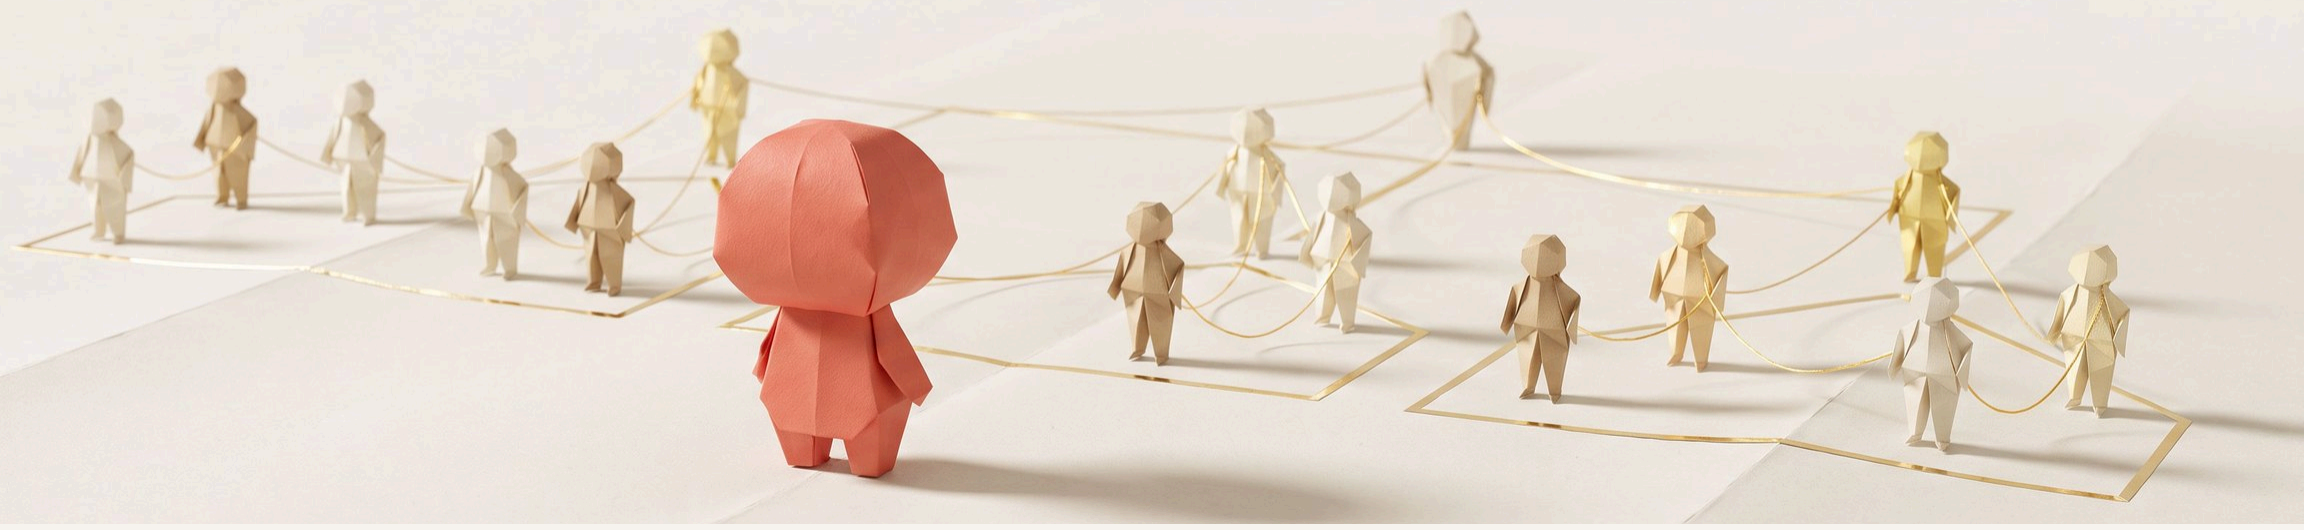
Bersin named the distinction in his March 2026 analysis: **the differentiator is the ability to auto-discover entities, workflows, and business rules from HR systems and keep a living, connected model of the workforce that agents can reason over.** Without that layer, you get Tier 1 and Tier 2 agents. With it, you get Tier 3.

What this means for HR technology architecture

The Workforce Context Engine sits between your systems of record and your agents.



The large-enterprise reality makes this urgent: organizations run roughly 400 applications, more than 100 of which touch employees (Bersin, 2026). No single system has the full picture. The Workforce Context Engine is the layer that connects them and keeps one current understanding of the workforce across all of them.



The surface problem

Why HR tech adoption has always failed

The secret of HR technology

HR technology has a chronic, industry-wide adoption problem that almost nobody talks about honestly.

Organizations spend millions on sophisticated HR platforms: talent suites, learning platforms, career tools, internal mobility portals. Then they measure adoption and find that a fraction of employees ever log in, and fewer still come back.

The reason is simple: these systems live in their own portals, and employees don’t.

An employee's workday happens in email, in Slack or Teams, in their productivity tools. Asking them to switch to a separate HR portal to explore career opportunities, update their skills, or find development resources is asking them to break their flow. Most won't. The ones who do are self-motivated outliers who would probably have found those opportunities anyway.

This isn't a training problem or a communication problem. It's an architecture problem.

Why this gets worse with agents

Here's the irony: agentic HR could be the biggest shift ever in how employees interact with HR systems. But if agents live in the same portals nobody visits, the shift never happens.

A career coaching agent that gives brilliant, personal guidance is worthless if the employee has to remember to log in to a vendor portal to talk to it. A redeployment agent that finds perfect-match opportunities is worthless if affected employees never see them because they were posted in a system they don't check.

> *The adoption problem doesn't just reduce ROI. It makes the entire investment invisible. And invisible HR technology gets cut in the next budget cycle.*

The flow-of-work imperative

The principle is non-negotiable: **intelligent agents must live where employees already work.** That means:

1. **Microsoft Teams and Slack**

   Not as a link that opens a vendor portal, but as a native experience inside the collaboration tool. The agent is a conversation partner in the tool the employee already has open. Career coaching happens in a Teams chat. Opportunity recommendations arrive as a Slack message. Validating a skill is a quick reply in a tool they use 50 times a day.

2. **Microsoft Copilot**

   As an extension of the AI assistant that's becoming the default interface for knowledge work. When an employee asks Copilot about career growth, the answer draws on the Workforce Context Engine. When a manager asks Copilot to help plan their team's development, the agent has full workforce intelligence behind it.

3. **Email and mobile**

   Because frontline workers and field employees don't live in Teams or Slack. The surface has to adapt to how each population actually works.

What “in the flow of work” actually requires

This is harder than it sounds, and HR technology leaders should push vendors on the details.

Native integration, not iframe embedding

Embedding a vendor portal in a Teams tab is not flow of work. It’s a portal in a frame. Real flow of work means the interaction feels native to the platform: Teams cards, Slack blocks, Copilot extensions, native mobile notifications. The employee should never feel they’ve left their main tool.

Context in both directions

The agent needs to understand what the employee is doing in their work tool. A manager reviewing a team member’s performance should be able to ask for development recommendations without re-entering context. An employee finishing a project should get relevant opportunities based on the skills they just showed, without navigating anywhere.

Proactive, not just reactive

The most powerful flow-of-work agents don’t wait to be asked. They raise the right thing at the right moment: a new internal role that matches a recently stated career interest, a course that closes a skill gap before an upcoming project, a retention risk a manager should act on. That turns HR technology from something employees have to seek out into something that comes to them.

Zero-friction interactions

Every tap, click, or step between the employee and the value cuts adoption. “Click here to open the portal to see the full recommendation” is an adoption killer. The interaction has to finish inside the flow-of-work surface, or at most one click away.

The adoption evidence

The data backs this principle overwhelmingly.

Bersin's research consistently names the end-user experience layer as the critical differentiator, specifically calling out agents embedded in *"tools employees already open 50 times a day."*

Organizations that have deployed HR capabilities inside collaboration platforms see dramatically higher engagement than those using standalone portals.

The $6.5 trillion frontline work market (per UKG's analysis) is almost entirely unreachable through traditional portal-based HR technology.

Synq is built on this commitment. It works natively in Teams, Slack, and Microsoft Copilot, and through MCP in AI assistants like Claude, so the employee never leaves their main work tool. The career conversation, the opportunity, the skills check: it all happens where the employee already is. That isn't a feature. It's the design decision that decides whether the investment delivers value or sits unused.



The real architecture of agentic HR

Putting it all together

With the two critical layers in place, the Workforce Context Engine that enables intelligence and the flow-of-work surface that enables adoption, the full architecture comes into focus.

Bersin's March 2026 research described a five-layer model. The architecture here refines it, with sharper detail on what each layer must contain and why.

1 Systems of record

Your HRIS, ATS, LMS, and payroll

These aren't going anywhere. Workday, SAP SuccessFactors, Oracle HCM, and UKG remain the authoritative source for employee records, payroll, compliance documentation, and core transactions. They're excellent at what they do.

What they are not: systems that understand workforce dynamics. They store records. They process transactions. They maintain compliance. **They were never designed to be the intelligence layer for agent-based decisions, and bolting an agent framework onto a record-keeping system doesn't change the data architecture underneath.**

2 The Workforce Context Engine

This is the layer most “agentic HR” architectures are missing, and the one that decides whether your agents are intelligent or just fast. The Workforce Context Engine:

- Brings in data from every system of record and connects it into one model
- Auto-discovers entities, workflows, and business rules from HR systems and stays in sync as they change
- Keeps a knowledge graph of people, skills, roles, teams, career paths, and organizational relationships
- Provides workforce-specific AI, with models trained on career trajectories, skill adjacencies, and organizational patterns rather than generic internet text
- Encodes governance: policies, approval chains, and compliance requirements as enforced constraints
- Carries context across sessions, so agents build on what came before instead of starting from zero

Without this layer, agents reach straight into systems of record, get flat data back, and try to make sense of it with generic LLM reasoning. That works for looking up PTO balances. It doesn’t work for deciding which of 340 affected employees should be redeployed to which roles.

3 Intelligent agents

With the Workforce Context Engine as their foundation, agents can work across the full tier spectrum.

Tier 1 agents (informational) answer questions, now with workforce intelligence behind them. “What skills are growing fastest in our engineering organization?” becomes answerable, because the Context Engine tracks how skills change across the graph, not just keyword counts in profiles.

Tier 2 agents (transactional) run processes, with intelligent guardrails. An agent handling a transfer doesn’t just submit the paperwork. The governance engine checks the move against pay policy, skills requirements, and regulatory constraints before anything happens.

Tier 3 agents (intelligent) reason and orchestrate, which is only possible because the Context Engine gives them the connected understanding they need. These are the agents that deliver transformative value:

- **Career development:** personal career coaching grounded in the employee’s actual skills, trajectory, internal opportunities, and context, across months of conversation
- **Workforce redeployment:** matching displaced employees to new roles across skills, location, pay, licensing, and career preferences
- **Workforce redesign:** analyzing how AI and automation are reshaping specific roles, and how the organization should evolve its role architecture and skills strategy
- **AI impact analysis:** modeling the impact of emerging technology on the workforce, function by function
- **Candidate discovery:** deep semantic matching of internal and external talent to opportunities on skills, potential, and fit, not keyword overlap
- **Organizational design:** data-driven recommendations for how teams, roles, and reporting lines should evolve

Bersin's description of building a redeployment agent shows why the Context Engine matters: the agent has to coordinate skill assessments, job requirements, coaching, training availability, location constraints, licensing, union rules, and pay all at once. **No amount of LLM intelligence makes up for missing connected workforce data to reason over.**

4 The flow-of-work surface

Where employees, managers, and HR professionals work with agents: inside Teams, Slack, Microsoft Copilot, email, and mobile. The surface is not an afterthought or a "UI layer." **It's an architectural commitment that decides adoption, and with it, whether the value is ever realized.** The surface must support:

- **Conversation:** chat-based coaching, Q&A, guided workflows
- **Proactive messages:** opportunity recommendations, risk alerts, nudges to act
- **Quick actions:** approve, apply, give feedback, without leaving the tool
- **Rich content:** career paths, skill gap analyses, team dashboards, rendered natively in the platform

How the layers interact

A concrete example

An engineering manager opens Teams on a Monday morning. A message from the career development agent is waiting:

> *"Two people on your team, Priya and James, have been building skills that closely match what the new AI Platform team needs in Q3. Based on the career interests they've shared in earlier coaching sessions, their skill trajectories, and how fast they're developing, both would be strong candidates. Want to look at development plans that could prepare them, or bring it up in your next 1:1s?"*

What happened behind the scenes

01 **Systems of record** provided current role data, team structure, and the Q3 hiring plan.

02 **The Workforce Context Engine** found the skill adjacency in its knowledge graph, checked the match against workforce models trained on career trajectory patterns, checked the governance engine for constraints (neither employee is on a retention hold or in a restricted transfer period), and drew on memory of both employees' stated career interests from earlier coaching conversations.

03 **The intelligent agent** orchestrated the analysis, put the reasoning together, and wrote a recommendation a manager could act on.

04 **The Teams surface** delivered it as a native message in the tool the manager already had open, with quick actions to explore further. No portal, no login, no context switch.

That's the architecture. Each layer is necessary. Remove the Context Engine, and the agent can't make the recommendation. Remove the flow-of-work surface, and the manager never sees it.



Governance, risk, and the EU AI Act

This is not optional

Employment AI is in regulators’ crosshairs around the world. HR technology leaders who treat governance as a later-phase concern are building on a foundation of risk.

The EU AI Act: HR is explicitly high-risk

The EU AI Act (Regulation (EU) 2024/1689), the world’s first comprehensive AI law, explicitly classifies employment AI as high-risk. It names “AI tools for employment, management of workers and access to self-employment (e.g. CV-sorting software for recruitment)” as needing enhanced compliance.

The high-risk rules for employment were set to apply from August 2, 2026. The EU's Digital Omnibus proposal would tie that date to the availability of standards, with a backstop of December 2027. Either way, the systems you choose now are the ones that will be audited. What's required:

Obligation	What it means
Risk assessment and mitigation	Document the risks of every AI system used in employment decisions
High-quality training data	AI must be trained on data that minimizes discrimination
Activity logging and traceability	Complete audit trails of every AI-influenced employment decision
Technical documentation	Architecture, capabilities, limits, and intended use, documented
Human oversight	Defined processes for human review and override
Robustness and accuracy	Testing and monitoring, including under adversarial conditions
Transparency	Employees must know when they're interacting with an AI system

Already prohibited (since February 2025): emotion recognition in the workplace, social scoring, and AI that exploits people's vulnerabilities.

EEOC: you're liable even for vendor AI

In the United States, the EEOC has made clear that **employers are liable for the disparate impact of AI tools, even when a third-party vendor provides them.** The four-fifths rule applies, and adverse impact analysis is expected. Recent litigation against Workday and Eightfold, which Bersin references, shows this is being tested in court.

Why the governance engine matters

This is where the governance part of the Workforce Context Engine goes from useful to essential.

Most agent platforms handle governance through prompting: telling the LLM to follow certain rules. That's inadequate for high-risk employment AI. Prompt-based constraints can be applied inconsistently, are hard to audit, and don't give the deterministic guarantees regulators expect.

A governance engine encodes policies, approval chains, and compliance requirements as **enforced constraints: rules the system applies every time, not suggestions made in conversation.** Every agent decision passes through it, and it:

- Checks the decision against policy before any action is taken
- Logs the full decision trail: data accessed, reasoning applied, alternatives considered, constraints checked, action taken
- Enforces human-in-the-loop checkpoints for high-stakes decisions
- Produces the auditable documentation regulators and legal teams need

In Synq, every action also runs with the permissions of the person asking. Access is decided before any data is read, down to the row and the field, so the agent can never see or change more than that person could.

> *As Bersin warns: "Scaling automation creates systemic risk. Single algorithmic errors can affect thousands instantly." The governance architecture has to match the scale of deployment.*

Human in the loop

Where to draw the lines

The EU AI Act requires human oversight for all high-risk employment AI. But “human in the loop” is a design spectrum, not a switch.

High stakes

Hiring decisions, terminations, compensation changes. **A human approves before anything happens.**

Medium stakes

Interview scheduling, development plan recommendations. **Human review within a set timeframe; routine cases run on their own.**

Low stakes

Policy Q&A, calendar coordination, looking up information. **Autonomous, with full logging.**

The architecture has to make these boundaries explicit, enforceable, and auditable, not dependent on whether the LLM “remembers” to check with a human.



Evaluating vendors

The questions that matter

Every HR technology vendor now claims AI agents. Here’s how to separate architectural substance from marketing.

On workforce context

“What is your agent reasoning over?”

If the answer is “we connect to your HRIS and retrieve data,” that’s Tier 1 or Tier 2. The agent is a retrieval layer, not an intelligence layer. Ask specifically:

- Do you maintain a knowledge graph of workforce relationships, or query flat records?
- Are your models trained on workforce-specific data (career trajectories, skill adjacencies), or do you use generic LLMs?
- How do you handle skill adjacency and inference? Can you find people who don’t have the listed skills but have shown they can build them?

- How is workforce context kept current: batch sync or continuous?
- How many years of enterprise workforce data have your models been trained on? Across how many employees and organizations?

“Show me a Tier 3 use case end to end.”

Ask to see a redeployment or career coaching scenario that shows real reasoning, not just retrieval and display. Watch whether the agent makes judgment calls that depend on relationships, or just returns a sorted list of matches.

On surface and adoption

“Where do employees interact with the agent?”

If the answer involves a vendor portal, even a well-designed one, you’re accepting HR technology’s historical adoption problem. Push on:

- Is the interaction native to Teams, Slack, and Copilot, or an embedded frame?
- Can the agent act proactively (messages, contextual suggestions), or only respond when asked?
- Can employees finish the whole interaction without leaving their work tool?
- How do you reach frontline workers who don’t use collaboration platforms?

On governance

“How do you enforce policy constraints?”

If the answer involves prompt engineering or system instructions to the LLM, that’s a soft governance layer. For employment AI under the EU AI Act, ask:

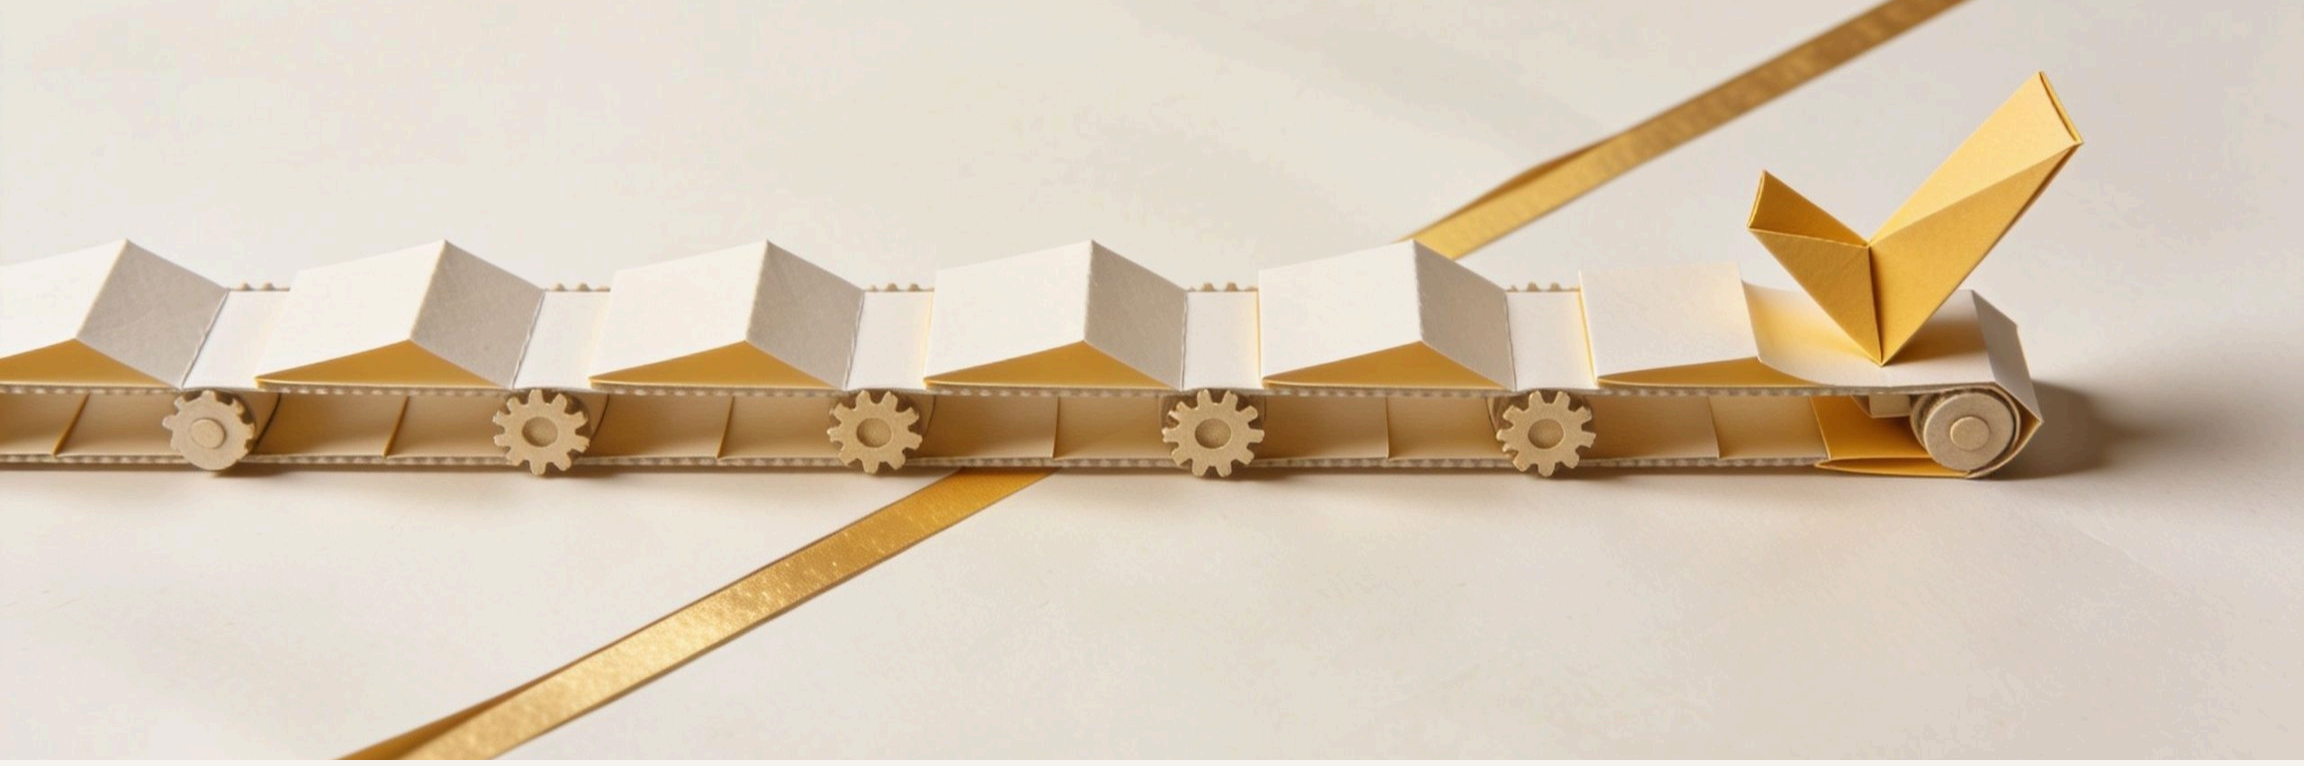
- Are policies enforced as constraints, or written as LLM instructions?
- Can you produce a complete audit trail for any agent decision?
- Does the agent act with the permissions of the person asking, or with a system-wide account?
- How do you implement and enforce human-in-the-loop checkpoints?
- What adverse impact monitoring is built in?
- Can you show how you meet the EU AI Act requirements for high-risk employment AI?

On architecture

“What happens when I have agents from multiple vendors?”

Large enterprises will end up with agents from their HCM vendor, their service management platform, and specialized platforms. Ask:

- How does your platform work with other agent systems?
- Can your Workforce Context Engine serve as the intelligence layer for agents running on other platforms?
- How do you handle overlapping capabilities and conflicts between agent systems?



Building your roadmap

The market context

The data is clear on direction, even if the timeline is debated.

33%

of enterprise software applications will include agentic AI by 2028, up from less than 1% in 2024

Gartner

4%

of organizations have a clear AI strategy for HR today

Bersin

30–40%

of HR roles are automatable with relatively low effort, while HR job postings rose 60% over five years

Bersin

85%

of HR executives plan a skills-flexibility redesign within three years, but only 10% classify skills effectively today

Deloitte

72%

of workers want work shaped around their skills; only 12% experience it

Deloitte

The gap between ambition and infrastructure is where HR technology leaders create value.

PHASE 1
Foundation
Months 1–4

PHASE 2
Expansion
Months 5–10

PHASE 3
Orchestration
Months 11–18

PHASE 4
Transformation
Ongoing

1 Foundation · months 1–4

Establish the Workforce Context Engine.

This is the single most important architectural decision. Without it, every later investment in agents tops out at Tier 1 and Tier 2.

- Evaluate platforms using the vendor questions on page 29
- Prefer platforms that auto-discover and sync from the HR systems you already run over ones that need a manual data migration
- Deploy skills intelligence first to check data quality and land quick wins
- Set up a governance framework aligned with the EU AI Act's high-risk requirements

Deploy a Tier 3 use case in a bounded domain.

Internal mobility or career coaching for a single business unit. The goal is to show intelligent-agent value, not just transactional automation, in a controlled setting. Measure adoption, recommendation quality, and employee satisfaction.

Get the surface right from day one.

Deploy in Teams, Slack, or Copilot. Don't start with a portal and plan to "add" flow of work later. The surface isn't a feature you add; it's a foundation you build on. If employees' first experience of your agentic HR investment is a portal, you'll be fighting that adoption pattern for years.

2 Expansion · months 5–10

- Expand to more Tier 3 use cases based on organizational priorities
- Scale across business units and geographies
- Put bias monitoring and adverse impact analysis in place
- Build in-house skill in configuring and managing agents
- Turn on proactive agent capabilities (recommendations and nudges in context)

3 Orchestration · months 11–18

- Deploy Superagent patterns that coordinate several agents toward complex outcomes (workforce redeployment, org redesign)
- Use workforce intelligence for strategic planning (AI impact analysis, skills forecasting)
- Set up continuous improvement loops based on outcome data
- Extend to frontline and field populations through mobile and other surfaces

4 Transformation · ongoing

- Move from deploying agents to redesigning work itself, based on workforce intelligence
- Shift the HR operating model from transactional support to strategic workforce orchestration
- Use the Workforce Context Engine as the foundation for skills-based practices across the organization
- Enable Bersin's "Superworker" model, where every role is augmented by intelligent agents

The decision that matters most

Josh Bersin has described what's happening as "The Great Reinvention of Human Resources." The organizations that get through it well won't be the ones that deploy the most agents or tick the most vendor boxes. They'll be the ones that get the architecture right:

Intelligence grounded in workforce context. A surface built into the flow of work.

Everything else follows from those two decisions.

The HR technology leaders who make them now, with clear eyes about what's architecturally real and what's vendor theater, will shape how their organizations compete for, develop, and deploy talent for the next decade.

synq

Synq is the AI-native HCM built for the agentic era.

Describe what you need and Synq builds it: a new process, field, workflow, or report, in minutes. It does the HR work in Slack, Teams, and Copilot, stands up in weeks on top of the systems you already run or as your HCM, and is governed by default.

See it at synq.ai

This guide is an educational resource for HR technology leaders. It reflects publicly available research, analyst reports, and vendor information current as of Q1 2026, with regulatory timelines updated November 2026.

