AI agents inside the HR department

Whether an agent judges a person, what it can read, and who decides

IN BRIEF

01 HR is the department where an agent's output is about a person. A candidate who was not seen. An employee who was rated. A leaver who was flagged. That is why the question to ask first is not where an agent runs. It is whether the agent judges a person, and who decides.

02 The public material answers that question differently. One published workflow template scores every application and emails the candidates, "all without human review". One vendor's screening template says its scores "must be reviewed by a human recruiter". One self-service agent hands "legal or personnel decisions" to HR specialists. One HR platform says every action it proposes "requires explicit human confirmation before execution". Hosting separates them less: three of the vendor pages we read offer an on-premises option in so many words.

03 The record explains why the first question matters. A United States court has let a collective action over a vendor's applicant screening tools go ahead, and the case is still open. A tutoring company paid $365,000 to settle a charge that its software rejected older applicants. A recruiting chatbot's supplier left a test account open, and applicants' names and phone numbers were exposed.

04 This paper is about the HR department's own work, from employee questions and recruiting administration to pay, employee relations and compliance. It sets out what an agent can read, how it reaches HR systems, what it may do without a person, and what the rules on AI at work and on decisions about people already say. It ends with seven questions to put to any supplier, and with a place to start.

05 The position throughout is simple. An agent prepares. A named person decides. No agent hires, rejects, ranks or scores a candidate. None rates an employee, predicts who will leave, or decides pay, promotion, discipline or dismissal. None reads emotion. No agent builds a profile of a person. What reaches a candidate or an employee has been read and sent by a person. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where HR departments are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 7
06 A worked example: one retailer's HR team in three countries 9
07 Seven questions before an agent touches a person's record 10
08 Where to start 10

PART 01

Where HR departments are today

$365,000

to be paid by a tutoring company to settle a charge that its software automatically rejected older applicants

64M

records exposed, including applicants' names, email addresses and phone numbers, by a weak password on a hiring chatbot supplier's test account

The suppliers have arrived, and many began with the CV. Vendor pages now describe agents for employee questions, HR cases, scheduling, onboarding, screening and performance reviews. Workday announced HR agents in September 2025, one of which "Analyzes performance data from various enterprise applications to streamline reviews and recommend actions" [6].

They disagree on whether the agent judges, and on who decides. This is the clearest difference in the published material.

- One published workflow template sets out to "Automatically score every job application with AI" and to email the candidates, "all without human review". A candidate below the pass mark goes straight to "Rejection Email to Candidate" [1].
- A second template gives each CV a score and a status of accepted or rejected. "When they submit, the system handles scoring and sends an email instantly. HR teams can review all results directly in Google Sheets" [1]. The review comes after the candidate has been told.
- UiPath's marketplace template works out "a fit percentage to recommend approve or reject". It then says: "Candidate scores and recommendations must be reviewed by a human recruiter before any hiring decision is made" [2].
- Microsoft's employee self-service agent "enforces boundaries by escalating complex or sensitive requests (such as legal or personnel decisions) to HR specialists" [3].
- Rippling answers the question on its own page. "Can Rippling AI take actions without approval? No. Any action Rippling AI proposes requires explicit human confirmation before execution" [4].
- Salesforce describes agents that "complete routine requests without manual intervention" [7].

These are different answers, and a buyer should know which one it is buying. Two things are in play. One is whether a person approves before something is sent or changed. The other is whether the agent has already judged the person by the time anyone looks.

Hosting separates them less than buyers expect. Leena AI's page lists "VPC or on-prem" [5]. IBM's page for its agent product says it "Supports multiple clouds (IBM Cloud, AWS and more) or on-premise deployment" [5]. Microsoft says its self-service agent "is adaptable to on-premises, hybrid, or fully cloud-based environments" [3]. Other pages we read do not say. So where it runs is still a question to ask. It is not the question that tells suppliers apart.

The court record is about screening. In Mobley v. Workday, in the federal court for the Northern District of California, the judge granted preliminary certification of a collective action on 16 May 2025 [8]. A law firm's note describes the group as job seekers aged 40 and over who applied through the vendor's system and were allegedly denied employment "recommendations" [8]. In June 2026 the court refused to dismiss the claims brought under California law [8]. The docket showed new filings in November 2026. Nothing has been decided on the merits, and the vendor contests the claims. We note only that the dispute is about software said to have recommended for or against people.

A regulator has already settled one case. The United States Equal Employment Opportunity Commission announced in September 2023 that a tutoring company, iTutorGroup, would pay $365,000. According to the lawsuit, the company "programmed their tutor application software to automatically reject female applicants aged 55 or older and male applicants aged 60 or older" [9]. The software was simple. It still decided.

PART 01 · CONTINUED

An audit rule has proved hard to enforce. New York City's Local Law 144 bars an employer from using an automated employment decision tool "unless the tool has been subject to a bias audit within one year of the use of the tool" [10]. A law firm's note on the State Comptroller's audit of December 2025 reports that the city reviewed 32 published bias audits and found one case of non-compliance, where the Comptroller's auditors found at least 17 in the same set [10]. An audit of a scoring tool is not the same as a tool that is fair.

Candidate data has been exposed. In July 2025 Krebs on Security reported that researchers guessed the password to a test account at Paradox.ai, the supplier of a hiring chatbot used by McDonald's franchisees. The weak password "exposed 64 million records, including applicants' names, email addresses and phone numbers". The supplier said the researchers viewed only five chats, and that the account "had not been logged into since 2019 and frankly, should have been decommissioned" [11].

We found no named, dated, independently measured result for an HR agent, and this paper quotes none as proof.

PART 02

The work nobody has packaged

Both

workflow templates give the candidate a number and a status, and send the result to the candidate before anyone in HR has read the application

Hosted

services are what both are built from: model, form, mailbox, spreadsheet

We listed the jobs an HR department does and looked for each among the templates and prebuilt agents we opened. They are thick around the CV and the employee's question. We did not find them where the work is assembling a file for a named person to decide on.

WHERE TEMPLATES OR PREBUILT AGENTS EXIST	WHERE WE FOUND NONE
A CV scored against the job and the candidate emailed	**The documents of a grievance put in date order, with the gaps listed**
A fit percentage with a recommendation to approve or reject	**Pay gap figures worked out by category of worker, with the method shown**
An employee's policy question answered from the handbook	**The information pack for a consultation with staff representatives**
Routine HR requests completed in the HR system	**An employee's request for their own data gathered for the privacy officer**
Performance data analysed and actions recommended	**Ratings already given, set out by team and grade for a calibration meeting**

PART 02 · CONTINUED

Two findings matter more than the list. First, the screening templates we opened judge. Both workflow templates give the candidate a number and a status, and both send the result to the candidate before anyone in HR has read the application [1]. The pass mark is a setting. Whoever sets it has written a selection rule, whether or not the head of HR knows.

Second, both are built from hosted services: a hosted model, a hosted form, a hosted mailbox, a hosted spreadsheet. That settles where applicants' CVs go before anyone has asked [1].

The helpdesk templates show the better pattern. One answers from the handbook and, "If it finds no clear answer or the question concerns a personal case, it hands the question to the HR team with the conversation attached" [7]. That is an agent that informs and then stops.

The gaps sit where the reading and assembling is slow and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In an HR department that means three kinds of thing.

Documents.

The employee handbook, HR policies, collective agreements, the pay structure, role profiles, letter and contract templates. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the clause it relied on and says when the documents are silent.

Systems of record.

Employee records, leave balances, vacancies, training records, HR cases. Most are reached through an interface the company already runs. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Figures for reports.

Headcount, joiners and leavers, absence, pay by band, ratings already given. These are held in the data warehouse or a reporting database, and read only. Reports are about groups, and a group too small to hide a person is left out.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

Where the company's interface checks the caller, the agent acts as the person using it and cannot claim another. An employee sees their own record, and a manager their own team.

SIX WAYS INTO AN HR DEPARTMENT'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A job brief, an application, a contract, a witness statement, a survey export	Nothing is connected. The safest start
02	**A knowledge source**	The handbook, policies, collective agreements, role profiles, templates	Permissions are kept per reader
03	**An imported description of an interface**	Employee records, leave balances, vacancies, training records, HR cases	You choose the operations. Changes wait for approval
04	**A reporting database**	Headcount, absence, pay by band, ratings already given	Read only. Small groups left out
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**Public sources**	New laws, regulators' guidance, published collective agreements	Reads only, from the publisher

75

public interface descriptions read, from 23 publishers; 73 import

3,217

of the 5,957 operations in the 73 that import are sent as changes

What is published. We read 75 public interface descriptions from 23 publishers and ran each through the platform's importer. 73 import. None is refused, and 2 are over the size limit. The 73 carry 5,957 operations, of which 3,217 are sent as changes, and the method alone cannot tell which only read [12].

What that exercise taught us. More than half of the operations are sent as changes.

What is not published. The 23 publishers are mid-market HR and payroll systems, employer-of-record platforms, connectors, and identity, signing, background check, scheduling and messaging services. No large HR suite publishes a description a file address reaches. For those, the description comes from the company's own system, or the agent reads from the data warehouse.

Health, union membership and the log. Sick notes and union membership are among the most sensitive data HR holds. The safe design keeps them out of free text and out of the record of calls wherever it can.

So choose the operations at import, read what each one does, start with the ones that only read, and add nothing that changes a person's record until the approval step has been seen working.

PART 04

What an agent may do

Two rules run through every agent in this paper. It prepares, and a named person decides. And it builds no profile of a person.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the letter, the job description, the checklist, the pack, with each fact tied to its source.
03	**It lists** what is missing, what is out of date, what does not fit, and what falls due.
04	**It stops** before anything that changes a record or reaches a candidate or an employee. A tool that changes something waits for approval.
05	**It does one job.** An agent is given the tools for one duty. The agent that sets out an application holds no tool that sends a rejection.
06	**It judges nobody.** It gives no score, no rank and no recommendation on a candidate or an employee. It works on documents for a task, not on a view of a person's traits, performance or behaviour.
07	**People are told.** The company tells candidates and employees when an agent has prepared something about them, and informs staff representatives before an agent is switched on.

That gives 35 agents in nine areas: employee questions; recruiting administration; joining and leaving; HR records; learning; the review cycle; pay, benefits and time; employee relations and reporting; compliance, safety and consultation. The full list is a separate document. Four lines are drawn tighter than the rest.

Candidates.

An agent may draft a job description and mark wording that could put off a group. It may set out each application in the same format against the requirements the vacancy states, with the passage for each. It may propose interview times and prepare the same pack of questions for every candidate. It gives no score, no order and no recommendation, and it rejects no one. A recruiter reads every application and decides who goes forward. The list has no screening, ranking or scoring agent. That is on purpose. A score that a busy recruiter leans on becomes the decision in practice. Section 5 shows a court reasoning that way about a credit score.

Reviews and pay.

An agent may gather what the employee and the manager have already written into the review form. It may check that a draft goal is specific and dated. It may assemble a calibration pack from ratings managers have already given, and a pay review pack from the bands and the budget. It writes no assessment, proposes no rating and no award. The manager writes the review. The panel calibrates. The reward team approves every award.

Cases and leavers.

An agent may put the documents of a grievance or a disciplinary case in date order, with a neutral summary of each and the gaps. It may summarise exit interview notes into themes by department. It draws no conclusion about anyone, predicts nobody's departure, and holds no conversation with a leaver. The investigator investigates. The decision-maker decides.

Records and letters.

An agent may list records with missing or inconsistent fields, list permits that expire soon, and draft a standard letter from the template and the record. It changes no record and sends nothing. An HR administrator corrects the record. An HR officer signs and sends every letter.

No agent in the list infers mood or emotion from a voice, a face or a text.

PART 05

What the rules already ask

This part needs an employment lawyer's read, and a privacy officer's, before it is relied on. It reports what the texts say; it is not advice. It does not say how any agent would be classified under any law. That is for counsel, system by system. Nothing in it says that a platform makes a company compliant.

AI at work in the European Union

Recruiting and managing workers are on the high-risk list. Annex III of the AI Act lists, under employment:

- "(a) AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates" [13];
- "(b) AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships" [13].

A preparatory task is treated differently. Profiling never is. Article 6(3) says a listed system "shall not be considered to be high-risk where it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision making". One of its conditions is that the system "is intended to perform a preparatory task to an assessment". The same paragraph then closes the door: a listed system "shall always be considered to be high-risk where the AI system performs profiling of natural persons" [13]. Whether any one agent meets these words is a question for counsel, not for this paper.

Reading emotion at work is prohibited. Article 5(1)(f) bans "the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions", except for medical or safety reasons [13]. The prohibitions have applied since 2 February 2025 [14].

Workers are told first. Article 26(7): "Before putting into service or using a high-risk AI system at the workplace, deployers who are employers shall inform workers' representatives and the affected workers that they will be subject to the use of the high-risk AI system" [13].

The dates have moved. Regulation (EU) 2026/1744 says the high-risk duties "shall apply from: (i) 2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" [14]. Employment sits in Annex III, so its date is 2 December 2027.

Decisions about a person

A person decides about a person. Under the GDPR, a person "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her", with stated exceptions [15].

A score can be the decision. In December 2023 the Court of Justice of the European Union ruled on a credit agency's score. It held that producing the score is itself "automated individual decision-making" where a third party "draws strongly on that probability value to establish, implement or terminate a contractual relationship with that person" [16]. The case was about credit, not employment. Whether the same reasoning reaches a recruiter who leans on a CV score is for counsel. It is the reason no agent in this paper produces a score.

The United Kingdom now defines the test. Its data protection law, as amended in 2025, says "a decision is based solely on automated processing if there is no meaningful human involvement in the taking of the decision" [18]. That text has been in force since 5 February 2026.

PART 05 · CONTINUED

Staff representatives

In Germany the works council has a say. The Works Constitution Act gives the works council a right of co-determination over the introduction and use of technical equipment designed to keep track of employees' behaviour or performance. It also says that its rules on selection guidelines apply where artificial intelligence is used to draw them up [17]. We read the German text; the English here is ours. We did not open the French, Dutch or any other national rules.

The United States, in brief

We opened New York City's audit law, and the settlement, the open case and the audit note described in section 1 [8][9][10]. We did not open the federal statutes or any other state's law, and we say nothing about them here.

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	AI Act, Annex III point 4, with Regulation 2026/1744	Recruitment, selection and worker management listed; duties from 2 December 2027 [13][14]
European Union	GDPR, Article 22(1)	No decision "based solely on automated processing" with a significant effect [15]
United Kingdom	UK GDPR, Article 22A	"meaningful human involvement in the taking of the decision" [18]
Germany	Works Constitution Act, sections 87 and 95	The works council's say over such equipment, and over selection guidelines drawn up with AI [17]

What the rules have in common

01	**Judging a person is the regulated act,** whether it is called screening, evaluating or scoring.
02	**A person's review must be real.** A score that is leaned on can count as the decision.
03	**A profile of a person removes the lighter treatment** the European text gives to preparatory work.
04	**Workers and their representatives are told** before the system is used.

PART 06

A worked example: one retailer’s HR team in three countries

Take a made-up retailer: shops in three countries, about 9,000 staff, and an HR team of 46. Nothing here is a customer’s story, and every figure is illustrative; it is section 4 put to work.

01

An employee asks how much parental leave the policy gives in her country. The agent answers from that country’s policy, cites the clause, and says the policy is silent on one point she raised. That point goes to an HR adviser.

02

A bank asks for confirmation of employment for a mortgage. The agent drafts the standard letter from the template and the record: name, role, start date. An HR officer checks it, signs it and sends it.

03

A store manager needs a job description for a new role. The agent drafts it from her brief and the role profile, and marks two phrases that could put off older applicants. The recruiter and the manager edit and approve it before it is posted.

04

For the vacancy, 140 applications arrive. The agent sets out each one against the six requirements the vacancy states, with the passage for each, in the order they arrived. It gives no score. A recruiter reads all 140 and decides who is invited.

05

In a grievance, the investigator uploads 31 documents. The agent puts them in date order, summarises each in neutral words, and lists four gaps: two undated statements and two meetings with no note. The investigator investigates. A separate manager decides.

Every one of these reads documents or records the company already has, changes no record by itself, judges no one, sends nothing to a candidate or an employee without a person, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches a person's record

01 **Does it judge a person?**
Does any step score, rank, filter, rate or recommend on a candidate or an employee? If so, who sees the score, and what do they do with it?

02 **Who decides, and is there a pass mark below which nobody looks?**
Show us every threshold, who set it, and who can change it.

03 **Does anything reach a candidate or an employee unseen?**
Which messages, letters or rejections go out without a person reading them first?

04 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Does any step send candidate, employee, pay or health data elsewhere, and which?

05 **What can it read and change?**
Does it act as the person using it, with their permissions, or as one all-seeing account? Which tools change a record, and does each wait for a person?

06 **What is recorded, and who has been told?**
Who asked, which tool ran, what was decided, how long is that kept, and what does the record itself hold about people? What were candidates, employees and staff representatives told, and when?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and nothing about one person is judged.

HR policy questions and answers
Employees ask what the handbook says and get the answer with the clause cited. No record is touched.

HR letter drafting
The agent drafts the standard letter from the template and the record. An HR officer signs every one.

Job description drafting
The agent drafts and marks excluding wording. The recruiter and the hiring manager approve.

One country, one process, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/HR

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 n8n workflow templates 14138 (an AI resume screener from Google Forms to shortlisting and email; "Last update 8 months ago" on the page) and 7903 ("Automated resume screening and response system with Gemini AI, Gmail and Sheets"; "Last update 9 months ago" on the page). Neither carries a publication date. https://n8n.io/workflows/14138 ; https://n8n.io/workflows/7903-automated-resume-screening-and-response-system-with-gemini-ai-gmail-and-sheets/

02 UiPath Marketplace, "Match candidate skills to job requirements" (agent template listing), undated, read 22 November 2026. https://marketplace.uipath.com/listings/match-candidate-skills-to-job-requirements

03 Microsoft Learn, Employee Self-Service Custom Engine Agent, overview page, "Last updated on 2026-11-03" as shown on the page. https://learn.microsoft.com/en-us/copilot/microsoft-365/employee-self-service/overview

04 Rippling, "Rippling AI" product page (questions and answers section), undated, read 22 November 2026. https://www.rippling.com/ai

05 Leena AI, home page, undated, read 22 November 2026. https://leena.ai/ ; and IBM, watsonx Orchestrate product page, undated, read 22 November 2026. https://www.ibm.com/products/watsonx-orchestrate

06 Workday, "Workday Illuminate Expands with New AI Agents for HR, Finance, and Industry", press release, 16 September 2025. https://newsroom.workday.com/2025-09-16-Workday-Illuminate-TM-Expands-with-New-AI-Agents-for-HR%2C-Finance%2C-and-Industry

07 Salesforce, Agentforce employee support use case page, undated, read 22 November 2026. https://www.salesforce.com/agentforce/use-cases/employee-support/ ; and Make, workflow template "Answer employee policy questions", undated, read 22 November 2026. https://www.make.com/workflows/answer-employee-policy-questions

08 Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal.), filed 21 February 2023. CourtListener docket: entry 128, 16 May 2025, "ORDER by Judge Rita F. Lin granting 106 Preliminary Collective Certification"; latest filing shown 20 November 2026. https://www.courtlistener.com/docket/66831340/mobley-v-workday-inc/ Fisher Phillips, note on the certification order, 20 May 2025. https://www.fisherphillips.com/en/news-insights/discrimination-lawsuit-over-workdays-ai-hiring-tools-can-proceed-as-class-action-6-things.html Duane Morris, Class Action Defense blog, note on the order of 22 June 2026 (ECF No. 360), 24 June 2026. https://blogs.duanemorris.com/classactiondefense/2026/06/24/california-federal-court-grants-in-part-and-denies-in-part-workdays-motion-to-dismiss-in-mobley-v-workday/

09 U.S. Equal Employment Opportunity Commission, "iTutorGroup to Pay $365,000 to Settle EEOC Discriminatory Hiring Suit", press release dated "09-11-2023" on the page (11 September 2023 in United States date order). https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit

SOURCES continued

10 New York City Department of Consumer and Worker Protection, “Automated Employment Decision Tools (AEDT)” page, undated, read 22 November 2026. https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page DLA Piper, “Critical audit of NYC AI hiring law signals increased risk for employers”, 30 January 2026. https://www.dlapiper.com/en-ca/insights/publications/2026/01/critical-audit-of-nyc-ai-hiring-law-signals-increased-risk-for-employers

11 Krebs on Security, “Poor Passwords Tattle on AI Hiring Bot Maker Paradox.ai”, 17 July 2025. https://krebsonsecurity.com/2025/07/poor-passwords-tattle-on-ai-hiring-bot-maker-paradox-ai/

12 Templates and prebuilt agents of agent and workflow platforms, as opened for sources 1 to 7; and 75 interface descriptions from 23 publishers fetched 22 November 2026 and run through the platform’s importer (Katonic research).

13 Regulation (EU) 2024/1689 (AI Act) of 13 June 2024, Article 5(1)(f), Article 6(3) and (4), Article 26(7), Annex III point 4. http://publications.europa.eu/resource/celex/32024R1689

14 Regulation (EU) 2026/1744 of 8 July 2026 (Official Journal 24 July 2026), the replaced Article 113 points (a) and (c). http://publications.europa.eu/resource/celex/32026R1744

15 Regulation (EU) 2016/679 (GDPR) of 27 April 2016, Article 22(1). http://publications.europa.eu/resource/celex/32016R0679

16 Court of Justice of the European Union, Case C-634/21, SCHUFA Holding (Scoring), judgment of 7 December 2023, the operative part. http://publications.europa.eu/resource/celex/62021CJ0634

17 Germany, Betriebsverfassungsgesetz (Works Constitution Act), section 87(1) no. 6 and section 95(2a), current text read 22 November 2026. https://www.gesetze-im-internet.de/betrvg/__87.html ; https://www.gesetze-im-internet.de/betrvg/__95.html

18 United Kingdom, Data (Use and Access) Act 2025, section 80, inserting Article 22A into the UK GDPR; commencement note “S. 80 in force at 5.2.2026”. https://www.legislation.gov.uk/ukpga/2025/18/section/80