

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Privacy, Trust and No-Attribution

The schema, the gate and the consultation view: how a culture product is built so that it cannot name a person, and how staff and their representatives can see that it cannot

A culture programme lives or dies on whether staff believe nothing in it can land on them. That belief cannot be created by a policy, a notice or an intent statement, because every programme that ended in a fine, a revolt or a retreat had all three. Red Bear: Enterprise Culture Evolution earns it differently: identifiable rows do not exist after the connector, so no role can export one; no source family reads a signal until workers' representatives have been consulted, the impact assessment exists and staff have been told; and a live consultation view shows representatives, at any moment, what is read, aggregated and retained. This paper sets out the schema, the gate and the view, the law they answer in the EU, Germany, the UK, the UAE and the US, the documented failures each is built to prevent, and the list of things Red Bear will not

Summary

Trust in a culture product is a property of its data, not of its documentation. CNIL fined Amazon France Logistique €32 million for per-worker tracking where aggregated data would have served; H&M Germany was fined €35.3 million because managers wrote down what staff told them and used it in evaluations; Barclays ended a desk-time trial after a staff revolt; Microsoft removed user names from Productivity Score within days; Aware's aggregated analytics were not the story, its identify-the-individual tier was. In every case the organisation had a privacy policy and a stated purpose, and in every case the employee's model of the tool was set by the most intrusive thing it could do. Red Bear: Enterprise Culture Evolution is built so that the most intrusive thing it can do is show a pattern for a group of five or more.

Three mechanisms carry that guarantee. **The schema:** identifiers are replaced by rotating tokens at the connector, message content is discarded before any pipeline stores it, submitted stories are coded to catalogue aspects on arrival and their text discarded, and every signal is aggregated above a floor before an engine reads it, so that no query by any role, including Red Bear's own operator and the hosting platform's administrator, can return a row about a person, because none exists. **The gate:** no connector ingests anything for a tenant until the consultation record, the data-protection impact assessment and the staff notice all exist and cover the configuration in force, and any change to a source family, a floor, a retention period or a role closes the gate again; the gate is in the connector service, not in a checklist. **The consultation view:** workers' representatives see, live and generated from the configuration that drives the connectors, what each family reads and discards, the floors, the retention, the roles, the lawful basis, the nudge budget, the coaching release in force, any dial flagged as used as a target, and the analyst query log, and they can annotate and question any of it in a record that cannot be deleted. Around the three sits a never-build list fixed in the platform, so that a product cannot describe an individual score, an emotion inference, a stored message or an identifiable export even by mistake. This paper traces the law in six jurisdictions and the documented failures to the rule and the platform requirement that answers each.

WHO SHOULD READ THIS

- Data-protection officers and privacy counsel asked to approve an AI culture engine
- Works councils, staff councils, unions and workers' representatives who will be consulted
- Chief people officers and chief information security officers who own the impact assessment
- Employees who want to know what the product can and cannot see

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- State what enters the product per source family, what is discarded at the connector, and why no role can export a person
- Explain the go-live gate, what opens it, what closes it, and where it is enforced
- Walk through the consultation view and know what a representative can see, annotate and question
- Name the legal duties in the EU AI Act, GDPR, BetrVG, the ICO guidance, the UAE PDPL and Labour Law and the US notice statutes, and which rule answers each

1 What it is

No-attribution, as Red Bear builds it, is the property that nothing in the product describes, scores, flags, nudges or names a person, and that the property holds in the data rather than in the interface. Three things make it true.

The schema. Every source family reaches the product through a connector that discards what must never be stored and replaces what must never be resolved. Collaboration signals are headers only: unit-level who-to-whom counts, timing and channel; subject lines, bodies and attachments are discarded at the connector and never stored by any signal pipeline (RBP FR-PA-02). Identifiers are replaced by rotating tokens before any engine reads a signal; the mapping is held by the privacy service alone, used for no analytical purpose, and never resolved to address a person (RBP FR-PA-03). Submitted content, a story, a message to the monitored mailbox, a post on the public page, is coded to catalogue aspects on arrival with names redacted and the text discarded (RBP FR-WA-09). Every signal is aggregated to a group of at least five, ten for anything coded from text, with complement suppression and time-window minimums, before the engine sees it (RBP FR-WA-04, FR-PA-01, FR-PA-06). The consequence is stated in one requirement: no query in the engine or the stores beneath it returns a row below the floor, for any role, including the tenant administrator, the Red Bear operator and the hosting platform's administrator, because identifiable rows do not exist after the connector (RBP FR-WA-11).

The gate. No connector ingests a signal for a tenant until three records exist and cover the configuration in force: the consultation record (date, body consulted, outcome, and the record that workers' representatives and affected workers were informed before use), the data-protection impact assessment, and the versioned staff notice (RBP FR-PA-09, FR-PA-11, FR-PA-05; ECE-CF-01 CST-19, §5.7). The gate opens by itself when the three exist; no role can open it by hand. A change to any source family, floor, retention period or role closes it again until the records are re-versioned, and every refused read while it is closed is audited (ECE-FSD-01 FR-GOV-08).

The consultation view. Workers' representatives hold a live view, generated from the same configuration that drives the connectors and the staff notice, of what the product reads, aggregates and retains; they may annotate any item or raise a question, and the record cannot be deleted (RBP FR-PA-09, FR-WA-13; ECE-FSD-01 FR-GOV-01, FR-GOV-02).

The three are backed by a never-build list. The platform offers no call, schema or connector for emotion or mental-state inference from any signal, for an individual score, ranking, flag, reading or nudge, for storage of message content by a signal pipeline, or for export of an identifiable row by any role; a product built on it cannot describe any of these (RBP FR-PA-08). And the platform has no sibling: no screen, export, API or tier of the product names a person against a signal, and Red Bear sells no companion product that does (ECE-CF-01 CST-20).

KEY POINTS

- Identifiable rows do not exist after the connector, so no role can export one; the rule is in the schema, not the interface
- Nothing is ingested until consultation, impact assessment and staff notice exist; a change to the configuration closes the gate again
- Representatives see what is read, live, from the configuration in force, and their annotations cannot be deleted

2 Why it matters

The product's purpose depends on the guarantee. Red Bear's position across its products is that fear stops information moving, and organisations that stop information moving decide badly (Red Bear Platform, *Platform Functional Specification*, §1.2). A culture programme exists to remove fear; a culture instrument that staff believe can name them adds it. The evidence that trust is decided by the data layer is direct. Microsoft's statement that Productivity Score "was never designed to score individual users" did not survive an interface that put a name next to a number (The Register, *Productivity Score*; ECE-RN-04 §1). Aware's aggregated analytics over billions of messages were not the scandal; the eDiscovery tier from the same vendor was, and critics saw only that tier (CNBC, *AI might be reading your Slack and Teams messages*; ECE-RN-02 §3). A coaching feature, a dial or a finding inherits the trust of the data beneath it (ECE-RN-04 §3, mechanism 1).

The cost of getting it wrong is measured in weeks, not quarters. CNIL's decision against Amazon France Logistique, Italy's Garante against Replika, the Nuremberg authority against H&M, and the ICO's intervention at Barclays each ended or reshaped a programme in weeks (ECE-RN-04 §3, mechanism 12). In Germany a works council can enjoin the roll-out of software capable of monitoring performance and contest evidence obtained without an agreement (Luther, *Software versus co-determination*; ECE-RN-02 §2). The EU AI Act now requires deployers to inform workers' representatives and affected workers before putting a high-risk system into service (Regulation (EU) 2024/1689, Art. 26(7)). Consultation and disclosure are preconditions, and a product that treats them as launch tasks will be launched without them the first time a date slips.

There is a third reason. A listening instrument that staff distrust measures nothing: Boeing's staff bypassed a speak-up channel they did not trust (Axios, *Boeing safety culture report*; ECE-RN-04 §2), and bossware provokes mouse jigglers and slowdowns (CSCW 2025, bossware study; ECE-RN-02 §3). The framework's measure of signal trust exists because a reading is only as good as the willingness of the people behind it to be read (ECE-CF-01 KPI-05).

3 The challenges

This section states the law the product must answer, jurisdiction by jurisdiction, and the documented failures it must prevent. Both are drawn from Red Bear's research note on continuous measurement, law and failures (ECE-RN-02 §2, §3, §5); the reading is Red Bear's own and is not legal advice. The legal review of the note is an open item in the product's ledger (ECE-FSD-01 §9, G-3).

3.1 The European Union

The EU AI Act. Article 5(1)(f) prohibits AI systems that infer the emotions of natural persons in the workplace, except for medical or safety reasons, and has applied since 2 February 2025; the Commission's guidance reads the prohibition as turning on biometric data, so that inference from written text is outside it, and the safety exception as covering life and health, not stress monitoring (Future of Privacy Forum, *Red lines under the EU AI Act*; ECE-RN-02 §2). Annex III point 4 classifies as high-risk AI used for monitoring and evaluating the performance and behaviour of workers; a tool reading collaboration metadata and case categories sits on that side even when it reports at group level, so the provider obligations of Articles 9 to 15 follow (CDT, *EU AI Act brief: AI at work*; ECE-RN-02 §2). Article 26 sets the deployer's duties, among them human oversight, retention of automatically generated logs for at least six months (26(6)), and in 26(7) informing workers' representatives and affected workers before use; 26(9) lets the deployer build its impact assessment on the provider's Article 13 information. Article 50 requires that people interacting with an AI system be told, from 2 August 2026. The May 2026 omnibus moved stand-alone Annex III obligations to 2 December 2027 and left the prohibitions untouched (Gibson Dunn, *EU AI Act omnibus agreement*; ECE-RN-02 §2).

GDPR. Article 35 requires an impact assessment for systematic and extensive profiling and for large-scale processing; under the nine European criteria, scoring, systematic monitoring, vulnerable subjects among them, a product combining several employee sources with new AI is a DPIA case on any reading, and employees count as vulnerable because of the power imbalance (ICO, *When do we need to do a DPIA*; ECE-RN-02 §2). Article 88 lets Member States set specific rules for employment data, naming monitoring systems at the workplace.

Germany. BetrVG §87(1) no. 6 gives the works council co-determination over technical devices intended to monitor behaviour or performance; the Federal Labour Court's objective-suitability test covers any tool *capable* of monitoring regardless of intent, Microsoft 365 itself has been held subject to it, and a works agreement is therefore a deployment prerequisite (Luther, *Software versus co-determination*; ECE-RN-02 §2).

3.2 The United Kingdom

The ICO's October 2023 guidance on monitoring workers requires a lawful basis chosen up front, expects a DPIA for monitoring even where not strictly mandatory, says monitoring must reflect workers' expectations and not be excessive, expects workers' or union views to be documented, treats monitoring of email *content* as likely special-category processing, and applies Article 22 to solely automated decisions with significant effects (ICO, *Employment practices: monitoring workers*; ECE-RN-02 §2). Consent is unreliable given the imbalance; the lawful bases in practice are legitimate interests with a documented balancing test, or legal obligation.

3.3 The United Arab Emirates

The Personal Data Protection Law, Federal Decree-Law 45/2021, applies outside the financial free zones; it lists exceptions to consent including contract and legal obligation (Art. 5), defines sensitive data to include health, religion, ethnicity and biometrics (Art. 1), requires a data-protection officer where processing involves systematic assessment including profiling (Art. 10), gives a right to object to automated decisions with serious effects, and requires an impact assessment before new technologies likely to result in high risk, in particular systematic automated evaluation of personal aspects (Art. 21) (ECE-RN-02 §2). The Cybercrime Law, Decree-Law 34/2021 Art. 2, criminalises interception of communications, so monitoring personal channels is a criminal rather than a civil risk. The Labour Law, Decree-Law 33/2021, prohibits discrimination (Art. 4), sets employer conduct-policy obligations (Art. 13), and prohibits sexual harassment, bullying and any verbal, physical or psychological violence against a worker (Art. 14); surfacing bullying *patterns* supports an Art. 14 duty and is a purpose to record in the impact assessment (GPA, *UAE bullying and harassment and the new Labour Law*; ECE-RN-02 §2). In the free zones, DIFC Data Protection Law 5/2020 requires an impact assessment for systematic observation of employees (Art. 20), and the ADGM Data Protection Regulations 2021 require one for high-risk processing and a data-protection officer for regular and systematic monitoring at scale, with fines up to USD 28 million.

3.4 The United States

New York Civil Rights Law §52-c requires prior written notice, an acknowledgement and a conspicuous posting for monitoring of telephone, email or internet use, with penalties of USD 500, 1,000 and 3,000 for successive offences, and exempts system processes not targeted at individuals; Connecticut §31-48d is broader with the same penalties; Delaware 19 Del. C. §705 accepts a one-time or daily electronic notice (Morrison Foerster, *New York enacts employee monitoring notification law*; ECE-RN-02 §2). NLRB General Counsel Memo GC 23-02 proposed that surveillance and algorithmic management tending to interfere with protected concerted activity be presumptively unlawful and that employers disclose what technology is used, why, and how data is used; since rescinded, it remains the clearest statement of the Section 7 risk (NLRB, *GC Memo 23-02*; ECE-RN-02 §2). The non-targeted exemption is the space a patterns-never-individuals product occupies; Red Bear gives the notice anyway.

3.5 The documented failures

Amazon France Logistique, CNIL, €32 million, decided December 2023. Warehouse scanners tracked scan speed and idle time per worker; raw data was kept thirty-one days and used for scheduling and appraisal. CNIL found breaches of minimisation, lawfulness, information and security, said aggregated data would have served the purpose, and weighed the effect on morale in the proportionality test (CNIL, *Amazon France Logistique*, as reported by Herbert Smith Freehills Kramer; ECE-RN-02 §3). *Lesson.* Aggregate at source, keep raw events only as long as aggregation takes, and let the impact assessment say why the aggregate is enough.

H&M Germany, €35.3 million, 1 October 2020. Managers at the Nuremberg service centre recorded holiday, illness, family and religious details from welcome-back talks on a drive readable by about fifty managers and used them in evaluations; a configuration error exposed it (Inside Privacy, *H&M receives record fine for employee surveillance*; ECE-RN-02 §3). *Lesson.* The harm was note-taking, not software: ingest categories, never narratives.

Barclays and Sapience, 2020. Individual desk-time tracking told staff they had spent "not enough time in the Zone"; the trial ended after a staff revolt, the bank self-referred, and the ICO said monitoring must be clear about purpose and bring real benefit (The Register, *Barclays employee monitoring*; ECE-RN-02 §3). *Lesson.* An individual nudge from monitoring is surveillance however it is framed; coaching goes to leaders about groups.

Aware, February 2024. Aggregated, anonymised sentiment and toxicity across some 6.5 billion Slack and Teams messages, and a separate eDiscovery product that can name individuals when a customer sets extreme-risk flags; critics called it thought crime and asked how a worker faces an accuser (CNBC, *AI might be reading your Slack and Teams messages*; ECE-RN-02 §3). *Lesson.* The aggregate product and the identify-the-individual mode from one vendor are one product in the employee's mind.

The listening platforms' own warnings complete the file. Viva Glint and Viva Pulse both warn that privileged users may, in some circumstances, identify respondents (Microsoft, *Viva Glint survey privacy*; *Viva Pulse privacy and data protection*; ECE-RN-02 §1). *Lesson.* An administrative back door is a back door whether or not anyone has used it; the notice that admits one is honest and insufficient.

KEY POINTS

- Three duties recur in every jurisdiction: assess the impact before deployment, tell staff and their representatives before use, and never rely on consent
- Every failure had a policy; what it lacked was a schema that made the harm impossible
- An identify-the-individual mode anywhere in the vendor's range is the product, in the employee's mind

4 How Red Bear: Enterprise Culture Evolution helps

Each duty and each failure has an answer, and each answer is a rule enforced by a platform requirement. The rules are the Red Bear Culture Framework's constraints CST-01 to CST-13 and CST-19 to CST-20 (ECE-CF-01 §5.6, §5.7) and the platform's privacy and culture-analysis modules (Red Bear Platform, *Platform Functional Specification*, v0.3).

Duty or failure	What Red Bear does	Rule
Emotion inference in the workplace (AI Act Art. 5(1)(f); EDPB on voice)	No emotion, personality or mental-state inference from any signal; no biometric input; no individual-level sentiment analysis, in any market; where voice coaching is enabled, audio is transcribed in the stream, discarded, and no prosody computed	RBP FR-WA-08, FR-DC-03, FR-PA-08; ECE-CF-01 CST-01

Duty or failure	What Red Bear does	Rule
Annex III high-risk and the Arts. 9 to 15 file	The culture engines are treated as high-risk now and the technical file is maintained; the product's DPIA record names the file version it relies on and the Art. 13 information it supplies	RBP NFR-13, FR-WA-13; ECE-CF-01 CST-08; ECE-FSD-01 FR-GOV-04
Impact assessment before deployment (GDPR Art. 35; PDPL Art. 21; DIFC Art. 20; ADGM; ICO)	A DPIA template ships per product, filled from the configuration documentation and the gateway's residency statement; no source family can be enabled until the record exists and covers it; a change to families, floors, retention or roles requires a new record	RBP FR-PA-11, FR-WA-13; ECE-CF-01 CST-07
Consultation and co-determination before use (BetrVG §87(1) no. 6; AI Act Art. 26(7); ICO; NLRB)	A consultation pack is generated per tenant: what is read, aggregated and retained per family, floors, roles, purposes, DPIA summary, staff notice, a works agreement reference where BetrVG applies, and the record that representatives and affected workers were informed before use; no connector ingests until it exists	RBP FR-PA-09; ECE-CF-01 CST-09, CST-19
Tell staff before collection (NY §52-c; Connecticut; Delaware; ICO; PDPL Art. 8)	A plain-language, bilingual, versioned staff notice: sources and fields read, floors and windows, retention, who sees what, that no one can see an individual, the lawful basis, that the coach is an AI; a change to any fact produces a new version before it takes effect	RBP FR-PA-05; ECE-CF-01 CST-11; ECE-FSD-01 FR-GOV-03
Lawful basis, never consent (ICO; PDPL Art. 5; CNIL's Art. 6 finding)	Legitimate interests with its balancing test, or legal obligation, recorded per family and shown in the notice and the pack; the notice asks for no consent and records no acknowledgement	RBP FR-PA-12; ECE-CF-01 CST-10
Disclosure of AI (AI Act Art. 50; ICF; EMCC)	The coaching agent's first utterance in every session says it is an AI, that nothing is scored or reported, and that a named human is one request away	RBP FR-DC-08; ECE-CF-01 CR-02
Logs for at least six months (AI Act Art. 26(6))	Automatically generated logs of the culture engines and the coaching agent are retained for at least six months under the tenant's control and hold no message content and no transcript	RBP NFR-15, FR-PA-10
Per-worker raw data kept for weeks (Amazon)	Raw events kept only as long as aggregation takes, in days; aggregates for the trend with 27 months as the ceiling; expired data deleted, not archived, and the deletion logged	RBP FR-PA-04; ECE-CF-01 CST-06
Narratives about people in a shared store (H&M)	HR cases, grievances and post-mortems enter as categories only: type, volume, cycle time, outcome; parties, narrative and evidence are never ingested	RBP FR-WA-06, FR-PA-02; ECE-CF-01 CST-05
An individual nudge from monitoring (Barclays)	No score, reading, flag or nudge for an individual, including to the person themselves; coaching goes to the named owner of a dial or decision about a group, never derived from a signal	RBP FR-WA-10, FR-PA-03, FR-DC-01; ECE-CF-01 CST-02
An identify-the-individual tier (Aware) and a privileged back door (Glint, Pulse)	No query returns a row below the floor for any role, including Red Bear's operator and the hosting administrator; the tenant administrator is limited to configuration; no screen, export, API or tier names a person against a signal, and Red Bear sells no companion product that does	RBP FR-WA-11, FR-PA-08; ECE-CF-01 CST-12, CST-20
A name next to a number (Productivity Score)	Floors of five and ten, complement suppression, the next-smallest group hidden, time-window minimums and noise on small-group views, applied in the data layer; every role sees *suppressed* and no count	RBP FR-PA-01, FR-PA-06, FR-PA-07; ECE-CF-01 CST-04

4.1 The schema

The schema is the reason the product does not need to be trusted. Three transformations happen at the connector, before any engine, screen or role: content is discarded, so that only unit-level counts, timing and channel remain (RBP FR-PA-02); identity is tokenised, with the mapping held by the privacy service alone and never resolved to address a person (RBP FR-PA-03); and submitted text is coded, names redacted first, the author replaced by unit and period or never received under the anonymous option (RBP FR-WA-09; ECE-FSD-01 FR-SIG-02). Then the privacy service aggregates: floors, complement suppression, time-window coarsening, noise (RBP FR-WA-04). The engine has no call that takes an individual identifier and no data structure keyed by one, which is what makes RBP FR-WA-11 true rather than promised: no query returns a row below the floor because there is no such row to return.

4.2 The gate

The gate is the consequence of one lesson: Koko disclosed its use of a language model after four thousand messages had been sent, and the benefit became a scandal; Replika's undisclosed change became a harm event (ECE-RN-04 §1). Consultation before the first signal is therefore enforced where a launch under time pressure cannot skip it, in the connector service (ECE-CF-01 §5.7, CST-19). Three records open it: the consultation record with date, body and outcome and the fact that representatives and affected workers were informed before use; the impact assessment with its date, owner and the Annex III file version it relies on; and the staff notice, versioned and covering every family enabled (RBP FR-PA-09, FR-PA-11, FR-PA-05). While it is closed, every connector refuses to ingest, every source-family enable is refused with the missing record named, and each refusal is written to the audit log (ECE-FSD-01 FR-GOV-08, FR-SIG-06). It closes again on any change to a source family, a floor, a retention period or a role, until the three records are re-versioned. A tenant that wants to add the mailbox family after launch consults on the mailbox first; the product will not read a message until it has.

4.3 The consultation view

The view answers the ICO's expectation that workers' views be documented, the works council's right under BetrVG, and the AI Act's Article 26(7), by making what the product does visible to the people it is done to (ECE-FSD-01 FR-GOV-01). It is generated live from the configuration in force, the same configuration that drives the connectors and the staff notice, so the three cannot drift (RBP FR-WA-13; ECE-FSD-01 FR-SIG-09). It shows the source families enabled and the fields each reads and discards; the floors and time windows; the retention per family; the roles and what each may read; the purposes and lawful basis per family; the DPIA summary and date; the nudge budget default; the health-monitor thresholds; the coaching release in force with its regression result; any dial flagged *used as a target*; and a summary of the analyst query log by role and scope, without content (RBP FR-PA-10, FR-WE-15). It shows no finding, dial reading, story or lesson. A representative may annotate any item or raise a question; annotations and questions are logged, shown to the HR partner and the sponsor, and cannot be deleted (ECE-FSD-01 FR-GOV-02).

4.4 The never-build list

> **WHAT THE PLATFORM CANNOT DO**
>
> The Red Bear Platform offers no call, schema or connector for: emotion or mental-state inference from any biometric or behavioural signal; an individual score, ranking, flag, reading or nudge, including to the person themselves; storage of message or document content by a signal pipeline; or export of an identifiable row by any role. A product built on the platform cannot describe any of these, because the capability does not exist (RBP FR-PA-08).

The list is the product's answer to Aware. It is not enough that this product has no identify-the-individual mode; the vendor must have none, because the employee's model of the tool is set by the most intrusive thing the vendor sells (ECE-RN-04 §3, mechanism 1). CST-20 fixes it: no sibling product, tier, export or API names a person against a signal, and a customer whose requirement is one of these is a customer for a different vendor (ECE-CF-01 CST-20).

> **STANDING RULE**
>
> Patterns, never individuals. Behavioural signals are aggregated above the minimum group size; message content is never stored; nothing in the product attributes a lesson, a finding or a reading to a person, including the person themselves. CST-02, CST-03 and CST-04 are the product form of this rule and are enforced in the data schema, not in the user interface (ECE-CF-01 §5.6).

5 How it is used

The governance flow runs before the first signal and again on every change; the people in it are the HR partner, the workers' representative, the sponsor and the system administrator, and the product refuses to move until each has done their part.

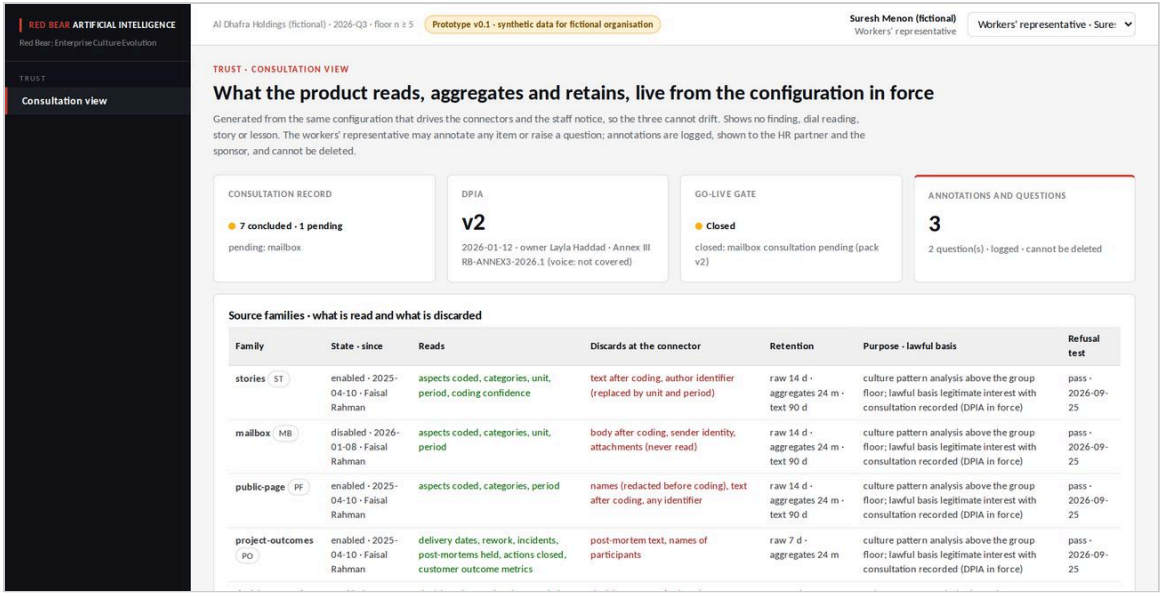
Step	Who	What happens
1 Configure without content	System administrator	A named C-level sponsor recorded (a function or committee is refused); personas mapped to Entra groups; floors at or above the platform minimum; retention within the maxima; the named human coach; the nudge budget default at or below two; every setting recorded with who set it and when (RBP FR-WE-13, FR-ID-04, FR-PA-04, FR-DC-13)
2 Generate the pack	HR partner	The consultation pack: what is read, aggregated and retained per family, the floors, the roles, the purposes and lawful basis, the DPIA summary, the staff notice, a works agreement reference where BetrVG applies (RBP FR-PA-09)
3 Consult	Workers' representative, HR partner	The representative reviews the pack on the consultation view, annotates and questions; the HR partner records the consultation: date, body, outcome, informed before use (ECE-FSD-01 FR-GOV-02)
4 Assess the impact	HR partner	The DPIA record is completed from the pre-filled template, with its owner, date and the Annex III file version it relies on (RBP FR-PA-11; ECE-FSD-01 FR-GOV-04)
5 Publish the notice	HR partner	The staff notice, in English and Arabic, versioned; it asks for no consent and records no acknowledgement (RBP FR-PA-05; ECE-FSD-01 FR-GOV-03)
6 The gate opens	Privacy service	When the three records exist and cover the configuration, the gate opens by itself; no role can open it by hand (RBP FR-PA-09; ECE-FSD-01 FR-GOV-08)
7 Enable families	System administrator	Each source family is enabled only while the gate is open for it; the administrator sees connector health, counts and errors and no content (ECE-FSD-01 FR-SIG-06, FR-SIG-08)
8 Watch and question	Workers' representative	The view stays live: the query log summary, any *used as a target* marker, the coaching release in force, the last refusal-test run; annotations at any time (RBP FR-PA-10, FR-WE-15)
9 Change	Any role with authority	A change to a family, floor, retention or role closes the gate; steps 2 to 6 run again on the change before it takes effect; a new notice version shows staff what changed (RBP FR-PA-05, FR-PA-11)

Roles and the absence of a back door. The tenant administrator configures connectors, floors and retention and cannot run analysis; the dial owner reads dials and patterns for their scope; the analyst reads de-identified aggregates across scopes (ECE-CF-01 §5.5). No role can export an identifiable row because none exists after the connector; the split follows Microsoft's separation of the Insights administrator from analysis and is the answer to the Aware case (ECE-RN-02 §4). HR holds configuration, consultation and the impact assessment, and receives only the escalations a person chooses to make and the health monitor's proposals on abuse and harassment; it cannot own a dial, edit a trigger or author a communication (ECE-CF-01 CST-18).

Refusals are recorded, not hidden. Every refused action is written to the append-only audit log with the actor, the authority and the object: below-floor reads under any role, reads while the gate was closed, refused floor changes, refused exports (RBP FR-AU-01; ECE-FSD-01 FR-GOV-07).

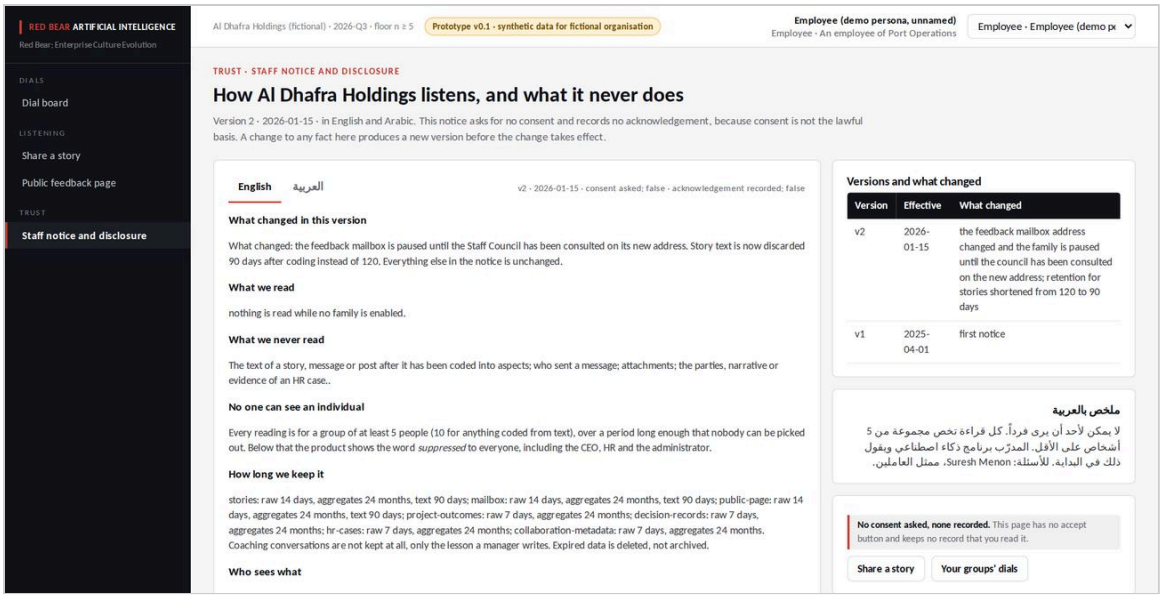
6 The user experience

The governance cluster has four screens, and the source configuration screen belongs with them; all five are shown as the prototype renders them, on synthetic data for a fictional organisation.

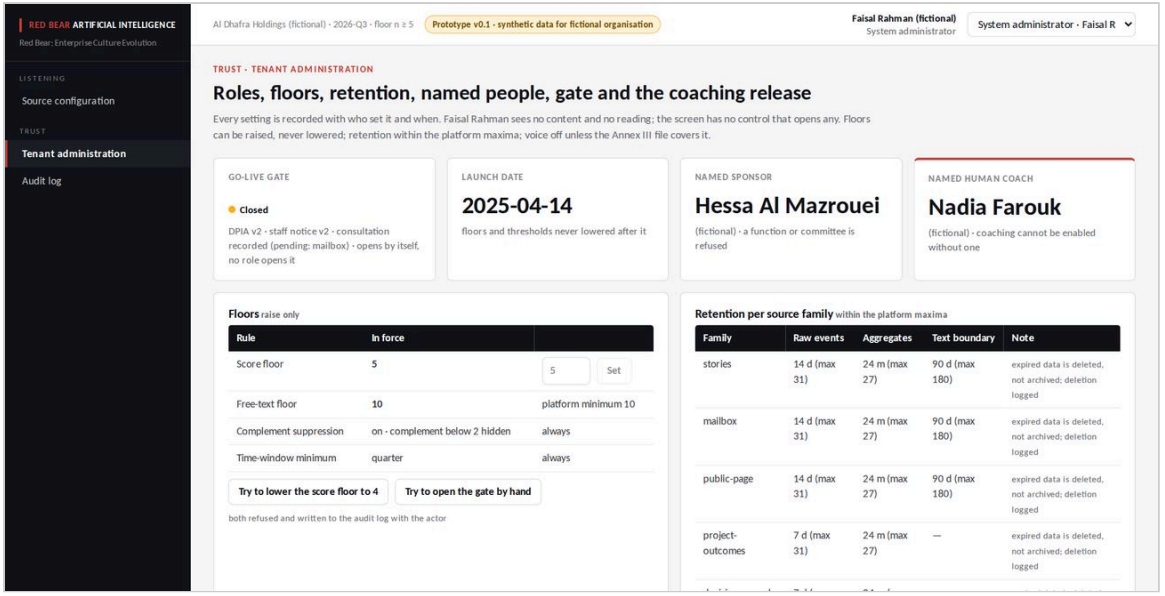


S-24 Consultation view: S-24 Consultation view. What the workers' representative sees, live from the configuration in force: the consultation record with its concluded and pending families, the DPIA version, the go-live gate state, and the count of annotations and questions. Beneath, every source family with what it reads, what it discards at the connector, its retention, its purpose and lawful basis, and the last refusal-test result. The representative's own questions and annotations sit inline, logged and undeletable. Prototype render on synthetic data for fictional entities.

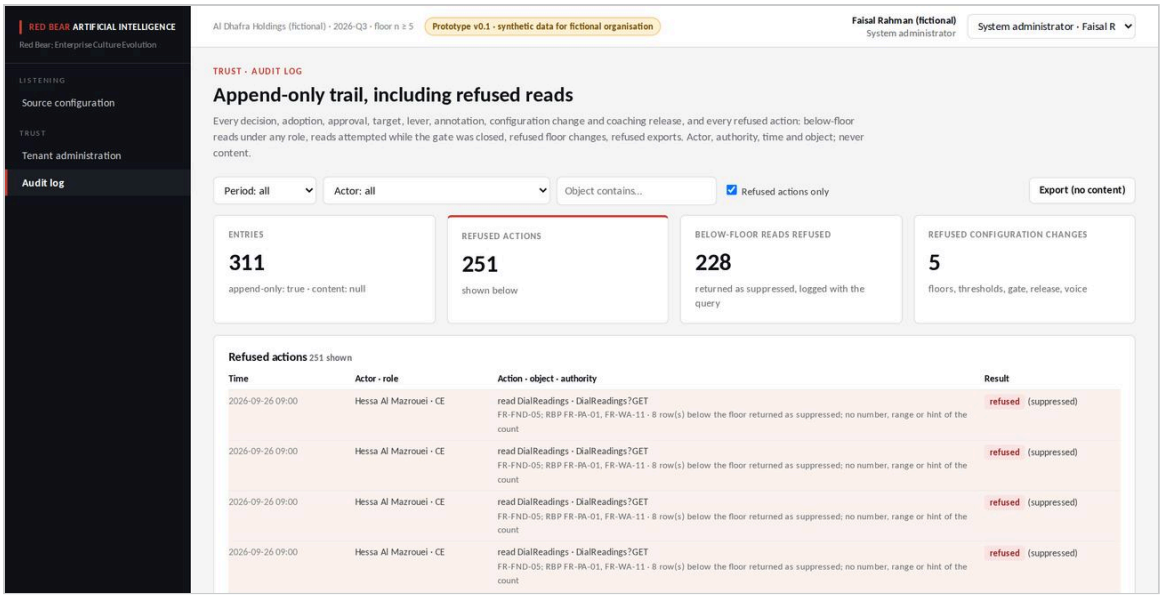
The view is a document, not a dashboard: it holds no reading, no finding and no dial, because the representative's question is not *how is the culture?* but *what are you reading, and can it reach a person?*.



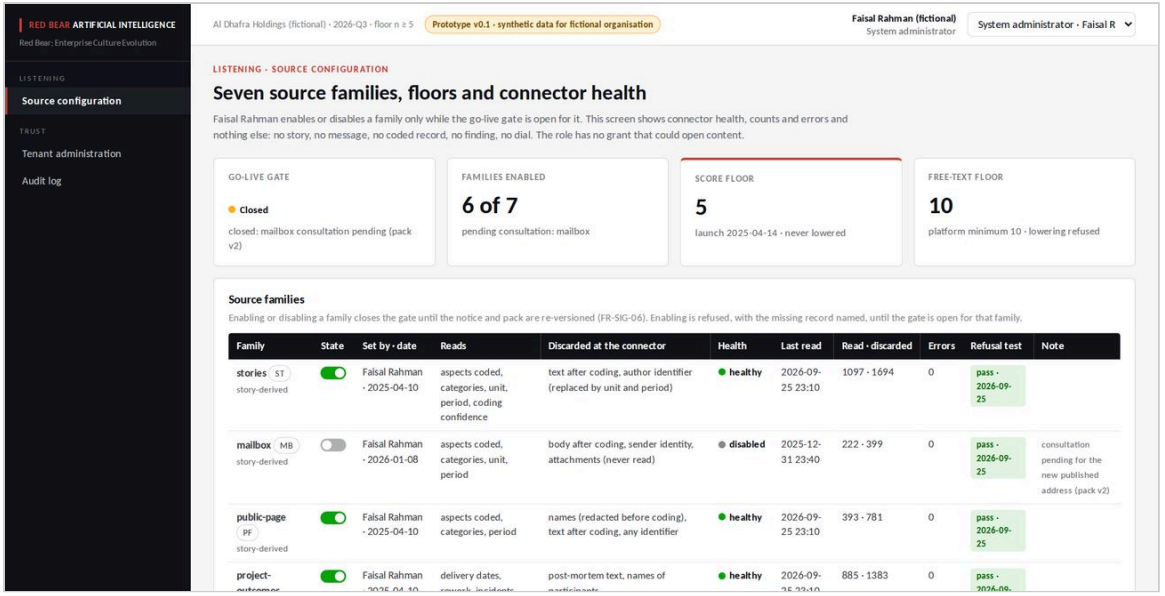
S-25 Staff notice and disclosure: S-25 Staff notice and disclosure. What every employee sees, in English and Arabic: what changed in this version, what is read, what is never read, that no one can see an individual and why, how long each family is kept, who sees what, that the coach is an AI, the lawful basis, and how to reach the workers' representative. The version table shows each notice and what changed. The page has no accept button and records no acknowledgement. Prototype render on synthetic data for fictional entities.



S-26 Tenant administration: S-26 Tenant administration. What the system administrator sees: the gate state with the record still pending, the launch date after which floors and thresholds are never lowered, the named sponsor and the named human coach, floors that can be raised and not lowered, retention per family within the platform maxima, personas mapped to Entra groups, settings including the nudge budget default and the voice toggle held off until the Annex III file covers it, and the DPIA record. Two controls on the screen exist only to be refused: lowering the score floor and opening the gate by hand. Prototype render on synthetic data for fictional entities.



S-27 Audit log: S-27 Audit log. The append-only trail with refused actions filtered to the top: below-floor reads under the sponsor persona returned as suppressed, each with the requirement that refused it; refused configuration changes; and the counts of entries, refused actions and below-floor reads. Actor, role, object and authority on every row; content on none. Export carries no content. Prototype render on synthetic data for fictional entities.



S-08 Source configuration: S-08 Source configuration. The seven source families with their state, who enabled each and when, what each reads and discards at the connector, connector health, records read and discarded, and the last refusal test per family. The mailbox family is disabled with the reason: consultation pending on its new address. Floors have raise-only controls; the panel on the right lists the refusal tests from the last pipeline run. Prototype render on synthetic data for fictional entities.

7 Where the data comes from, and how it is checked

Per source family, what enters the product and what the connector discards, as the consultation view and the staff notice show it (ECE-CF-01 §5.1).

Source family	What enters	Kind	What the connector discards
stories	Aspects coded, categories, unit, period, coding confidence	Content, submitted	Text after coding; author identifier, replaced by unit and period, or never received

Source family	What enters	Kind	What the connector discards
mailbox	Aspects coded, categories, unit, period	Content, submitted	Body after coding; sender identity; attachments never read
public-page	Aspects coded, categories, period	Content, submitted; public	Names redacted before coding; text after coding; any identifier
project-outcomes	Delivery dates, rework, incidents, post-mortems held and actions closed, customer outcome metrics	Categories and rates	Post-mortem text; names of participants
decision-records	Decision class, rationale recorded, level, number of owners, outcome versus expectation	Metadata of decisions	Decision content, unless the owner opens it to a session
hr-cases	Case category, volume, cycle time, outcome	Categories only	Parties, narrative, evidence: never ingested
collaboration-metadata	Counts, timestamps, unit of sender and recipients	Metadata only	Subject lines, bodies, attachments, individual identifiers, which become unit-level counts at the connector

7.1 How it is checked

- **Refusal tests on every pipeline run.** The below-floor read under all seven personas, the tenant administrator, the Red Bear operator and the hosting administrator; content refusal per family; a second mailbox connect, which has no call; gate-closed ingest. The last run's result and date are shown on the source configuration screen and the consultation view (RBP NFR-03; ECE-FSD-01 NFR-01, FR-SIG-09).
- **One configuration, three documents.** The fields read and discarded, the floors, the retention and the roles are generated from the configuration that drives the connectors, into the consultation view and the staff notice, so that what representatives are shown is what the connector does (RBP FR-WA-13).
- **Floors that cannot fall.** After launch, any attempt by any role to lower a floor, relax the complement rule or lower a health threshold is refused in the data layer, shown as refused, and written to the audit log with the actor (RBP FR-PA-01; ECE-FSD-01 FR-GOV-06).
- **Retention.** Raw events in days; aggregates up to 27 months; coded stories lose their text at the configured boundary; coaching sessions as the lesson record alone; audio never; expired data deleted, not archived, and the deletion logged (RBP FR-PA-04; ECE-FSD-01 FR-GOV-05).

7.2 What is never stored

Message bodies, subject lines and attachments from any channel. Audio from a coaching session, or a transcript. The parties, narrative or evidence of an HR case. The names of post-mortem participants. Any attribute of a person. Any score, reading, flag or nudge for an individual. Any emotion or mental-state inference. The platform offers no call, schema or connector for any of these (RBP FR-PA-08, FR-DC-03, FR-DC-04; ECE-CF-01 CST-01, CST-03, CST-05).

8 Benefits realised

The governance cluster is judged on whether staff and their representatives trust the instrument enough to feed it. The measures are the framework's, each with a *Never* clause (ECE-CF-01 §9.1).

Benefit	Measured by	Never
Staff trust the instrument	ECE-CF-01 KPI-05 signal trust: response share to the quarterly pulse; stories submitted per hundred staff per window, across the organisation	Per unit below the floor; a target for managers to drive
Speaking up is read honestly	ECE-CF-01 KPI-06 speak-up balance: case-category volume paired with the Edmondson and DORA speak-up items for the same units	Read alone; a claim that conduct improved or worsened without the paired items
The programme survives its first legal question	The consultation record, DPIA and notice exist before the first signal and are re-versioned on every change	A launch date that overrides the gate
Representatives can verify, not only be told	Annotations and questions on the consultation view, logged, shown to the HR partner and the sponsor, undeletable	A view that shows a reading, a finding, a story or a lesson to a representative

What is explicitly not claimed. The product does not claim that a group of five is proof against every inference; it claims that the floor is raisable to ten at launch and never lowerable, and that complement suppression, hiding and coarsening close the holes the literature names (ECE-RN-02 §4; ledger OD-3). It does not claim that consultation makes a deployment lawful in every jurisdiction; the records the law requires exist before the first signal, and the legal review of the research note is still open (ECE-FSD-01 §9, G-3). And a notice does not create trust: it states what the schema already makes true, and the audit log is where a representative checks it.

9 Risks managed

Residual risk	Mitigation
A launch under time pressure skips consultation	The gate is in the connector service; nothing is ingested until the three records exist; an attempted read is refused and audited (RBP FR-PA-09; ECE-FSD-01 FR-GOV-08)
A privileged administrator reads a row below the floor	No such row exists after the connector; the query path returns *suppressed* for every role, the attempt is logged, and the refusal test runs on every build (RBP FR-WA-11; ECE-FSD-01 NFR-01)
A free-text story identifies its author or a colleague	Names redacted before coding; anonymous by default; the public page moderated once, pre-coding, with content that belongs with HR routed and never coded (RBP FR-WA-09; ECE-FSD-01 FR-SIG-02, FR-SIG-03)
Voice brings biometric data into scope	Voice off by default, offered only once the Annex III file covers it and the pack lists it; audio never stored; no prosody (RBP FR-DC-11, FR-DC-03; EDPB, *Guidelines on virtual voice assistants*)
The coaching agent's behaviour changes without the tenant's knowledge	Prompt set, generation mode, model mapping and evaluation set versioned as one release, gated by the regression suite; the release in force is on the consultation view (RBP FR-DC-14, NFR-18)

What Red Bear will not build, even if asked. No individual score, reading, flag or nudge. No emotion inference from any signal, in any market. No export of an identifiable row for any role. No investigation mode, eDiscovery tier or sibling product that names a person against a signal (RBP FR-PA-08; ECE-CF-01 CST-01, CST-02, CST-12, CST-20).

Sources

Every claim attributed to a source in this paper names it in the text in the form *author*, *work*, *section*; the full reference and URL for each is held in the product's research notes, which are the citation source for the product. Statements without a citation are Red Bear's own position, developed for Red Bear: Enterprise Culture Evolution. Quotations are limited to fifteen words. Nothing in this paper is legal advice.

Cited as	Source	Authority
ECE-RN-02; ECE-RN-04	Red Bear research notes: *Continuous culture measurement: techniques, law, failures and design constraints*; *Why AI culture coaching has failed: cases and lessons* (research/monitoring-and-law.md, failures.md)	Red Bear research notes
ECE-CF-01; ECE-FSD-01; Platform FSD	Red Bear, *The Red Bear Culture Framework* v0.2 (§5, §9); Red Bear, *Functional Specification* ECE-FSD-01 v0.1 (modules SIG, GOV; §9); Red Bear Platform, *Platform Functional Specification* RBP-FSD-01 v0.3 (§1.2; modules WA, DC, PA, AU)	Red Bear standard
EU AI Act; omnibus	Regulation (EU) 2024/1689, Arts. 5(1)(f), 6(3), 9 to 15, 26, 50 and Annex III; Future of Privacy Forum, *Red lines under the EU AI Act*; CDT, *EU AI Act brief: AI at work*; Gibson Dunn, *EU AI Act omnibus agreement* (2026)	Regulation; legal commentary
GDPR; ICO	GDPR Arts. 35 and 88; ICO, *When do we need to do a DPIA*; ICO, *Employment practices: monitoring workers* (2023), as reported by Norton Rose Fulbright	Regulation; regulatory guidance
BetrVG	German Works Constitution Act §87(1) no. 6; Luther, *Software versus co-determination*	Statute; legal commentary
UAE law	Federal Decree-Law 45/2021 (PDPL) Arts. 1, 2, 5, 6, 10, 12 to 15, 21; Decree-Law 34/2021 (Cybercrime) Art. 2; Decree-Law 33/2021 (Labour) Arts. 4, 13, 14; DIFC Data Protection Law 5/2020 Art. 20; ADGM Data Protection Regulations 2021; Securiti and GPA commentaries	Statute; legal commentary
US statutes; NLRB	New York Civil Rights Law §52-c; Connecticut §31-48d; Delaware 19 Del. C. §705, as summarised by Morrison Foerster; NLRB, *GC Memo 23-02* (2022)	Statute; regulatory memorandum
CNIL; H&M	CNIL, *Amazon France Logistique* (2023), as reported by Herbert Smith Freehills Kramer; Inside Privacy, *H&M receives record fine for employee surveillance* (2020)	Regulatory decision; legal commentary
Barclays; Productivity Score; Aware; Boeing; bossware	The Register, *Barclays employee monitoring* (2020); The Register, *Productivity Score* (2020); CNBC, *AI might be reading your Slack and Teams messages using tech from Aware* (2024); Axios, *Boeing safety culture report* (2024); CSCW 2025 bossware study	Press reporting; academic reference
Koko; Replika	NBC News, *ChatGPT AI experiment in mental health tech* (2023); OECD AI incidents, Replika and the Garante (2023)	Press reporting; incident repository
Microsoft	*Viva Insights privacy*; *Viva Glint confidentiality thresholds*; *Viva Glint survey privacy*; *Viva Pulse privacy and data protection*	Vendor reference
EDPB; ICF; EMCC	EDPB, *Guidelines on virtual voice assistants*; ICF, *AI Coaching Framework and Standards* (2025); EMCC Global, *Ethics guidance for coaching using technology and AI* (2025)	Regulatory guidance; professional standard

Glossary

Term	Meaning in this paper
No-attribution	The property that nothing in the product describes, scores, flags, nudges or names a person, held in the data rather than the interface
Connector	The boundary at which a source family enters the product and where content, identity and narrative are discarded or replaced
Floor	The minimum group size below which nothing is shown: five for scores, ten for free text; raisable, never lowerable after launch
Go-live gate	The condition, enforced in the connector service, that the consultation record, the DPIA and the staff notice exist and cover the configuration before any signal is ingested
Consultation view	The live screen for workers' representatives generated from the configuration in force, with annotations and questions that cannot be deleted
Staff notice	The plain-language, bilingual, versioned notice every employee sees before collection; asks for no consent
Never-build list	The platform's fixed absence of any call, schema or connector for individual scoring, emotion inference, content storage or identifiable export
Annex III file	The EU AI Act Arts. 9 to 15 technical file Red Bear maintains for the culture engines, whose version the tenant's DPIA record names

The series

This paper is the third of five on Red Bear: Enterprise Culture Evolution. Each stands alone and each follows the same nine sections, so a reader who knows one knows where to look in the next.

Paper	Title	Theme	Status
ECE-WP-01	Culture Coaching Without Surveillance	What a decade of failed culture technology taught us	Draft
ECE-WP-02	Culture Analysis Without Individuals	Reading a culture from seven source families, above the floor	Draft
ECE-WP-03	Privacy, Trust and No-Attribution	The schema, the gate and the consultation view (this paper)	Draft
ECE-WP-04	Breaking Down Silos	Inter-group relationships and bridge building	Draft
ECE-WP-05	Culture Dials: Measuring What Leaders Choose	Principles, targets, levers and the honest measure of a programme	Draft

Red Bear Artificial Intelligence · Abu Dhabi, United Arab Emirates · info@redbearai.net · www.redbearai.net