

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Culture Analysis Without Individuals

Reading a culture from seven source families, above the floor, and why every finding is a question to its owner

Every organisation already produces the evidence of its own culture: the stories people tell, the decisions that carry a rationale or do not, the post-mortems that name systems or people, the case categories that rise or fall, the units that talk to each other and the units that do not. The temptation is to read that evidence at the level of the person, because that is where the data sits. Red Bear: Enterprise Culture Evolution reads it at the level of the group and nowhere else: seven source families, coded on arrival, aggregated before analysis, combined into one reading from three signal families, and delivered to a named owner as a question, never a verdict. This paper explains what the analysis engine reads, what it discards, how a finding is made, what it is and is not, and why an organisation that wants to know its culture must give up knowing its individuals.

Summary

A culture is a pattern, and a pattern is only visible above a certain size. The annual survey says 3.7 and nobody knows what to do on Monday; the exit interview says what went wrong a year after it did; the listening platform reports sentiment from messages nobody chose to submit. Between those, an organisation has evidence it never reads: stories staff would tell if asked, decision records with or without a reason, post-mortems held and actions closed, HR case categories and cycle times, and the headers of the collaboration traffic between units. Red Bear: Enterprise Culture Evolution reads seven such source families through one analysis engine. Three of them carry content a person chose to submit, and that content is coded to a catalogue of fifty culture aspects the moment it arrives, with names redacted and the text discarded. Four carry categories and metadata only. Every signal is aggregated to a group of at least five, ten for anything coded from text, before the engine sees it, with complement suppression, time-window minimums and noise on small-group views enforced in the data layer, so that a person cannot be found by subtraction, by slicing or by patience.

From those signals the engine produces findings: for a group, an aspect and a period, the direction of travel, the signal families behind it, a confidence band and a group size. A finding built on one family is a hypothesis, shown with a band and no position. A finding is stored as a question put to the named owner of the group's dial, in a fixed form that has no verdict word in it, and no record anywhere in the product may cite the engine as the finder of a fact about anyone. This paper sets out the evidence for that design, from the validity of the instruments culture programmes usually rely on, from the vendors whose aggregation practice sets the floor, and from the cases where reading the individual ended the programme. It then shows the analysis cluster as built: the story intake, the coding review, the findings list, the pattern detail and the inter-group view, and states, measure by measure, what the analysis claims and what it refuses to.

WHO SHOULD READ THIS

- Chief people officers and culture programme leads who have a survey result and no next step
- Data-protection officers, works councils and staff representatives asked what an AI culture engine reads
- Heads of analytics and data platforms who will connect the sources
- Managers who want to know what the product can and cannot say about their team

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Name the seven source families, and say for each what enters the product and what is discarded at the connector
- Explain why a dial needs three signal families and what a single-family finding may claim
- State the floors, the complement rule and the time-window rule, and why they are in the schema and not the interface
- Read a finding and know why it is a question, who it is addressed to, and what it is not

1 What it is

Culture analysis, as Red Bear builds it, is the continuous reading of a group's behaviour from de-identified, aggregated signals, against a fixed catalogue of fifty culture aspects, delivered as findings to the named owner of the group's Culture Dial. Four properties define it.

It reads sources the organisation already has, and one it creates. The Workplace Culture Analysis engine reads seven source families: stories staff submit through the product, messages to a mailbox the organisation publishes for the

purpose, posts on the organisation's own public feedback page, project outcomes, decision records from Enterprise Decision Intelligence, HR case categories, and collaboration metadata (ECE-CF-01 §5.1; RBP FR-WA-03). The first three carry content a person chose to write for this purpose. The other four carry rates, categories and headers. No family carries a message anyone wrote to a colleague.

It aggregates before it analyses. The privacy service replaces every identifier with a rotating token at the connector and aggregates every signal to a group of at least the tenant's floor before any engine reads it; the engine has no call that takes an individual identifier and no data structure keyed by one (RBP FR-WA-04, FR-PA-03). This is the difference between a culture instrument and a monitoring system, and it is a property of the schema, not a promise in a policy (ECE-CF-01 §2.1, CST-04).

It codes content to aspects on arrival. A story, a mailbox message or a public post is coded to catalogue aspects and categories as it arrives, using Red Bear's own aspect dictionaries; names in the text are redacted first, the author is replaced by unit and period, and the verbatim is discarded, or held only under the free-text floor until the retention boundary (RBP FR-WA-09). After coding, no role can read the text (ECE-FSD-01 FR-SIG-01, FR-SIG-05).

It produces findings that are questions. A finding is a pattern for a group, an aspect and a period: direction of travel, the signal families behind it, a confidence band, a group size and the catalogue version (RBP FR-WA-05). It is framed and stored as a question to the dial's owner, *the reading for [group] on [aspect] moved [direction] over [period]; what did you expect?*, and no record, export or communication may cite the engine as the finder of a fact about anyone (RBP FR-WE-16; ECE-FSD-01 FR-FND-06).

The analysis is the middle of the product: before it sit the principles a named leader chose; after it, the coaching session that turns a finding into a lesson and the programme that turns a lesson into a practice. This paper is about the middle.

KEY POINTS

- Seven source families; three submitted for the purpose and coded on arrival, four categories and metadata only; none a private message
- Aggregation to at least five, ten for text, in the data layer, before the engine reads anything
- A finding is a direction of travel with its families, confidence and group size, addressed to a named owner as a question

2 Why it matters

Schein's account of why culture programmes fail is that they work on artefacts and espoused values and leave the basic assumptions untouched (Schein, *Organizational Culture and Leadership*; ECE-RN-01 §1). A survey measures agreement with statements, which is the espoused layer. What the product reads instead are artefacts that assumptions produce without anyone intending to disclose them: whether a post-mortem names a system or a person, whether a decision carries a reason, whether escalations arrive before or after the deadline, whether one unit answers another within a day. Reading those is the artefact-level method Schein's clinical inquiry uses by interview, done continuously and without an interviewer (ECE-CF-01 §1.2).

The cost of not reading them is documented. Sull and colleagues, over 1.4 million Glassdoor reviews, found toxic culture 10.4 times more predictive of attrition than compensation, and that abusive management appears in under one percent of reviews yet cuts a company's culture rating by half a point (Sull, Sull and Zweig, *Toxic culture is driving the Great Resignation*; Sull and Sull, *Why every leader needs to worry about toxic culture*; ECE-RN-01 §1). A rare

pattern with a large effect is exactly what an annual survey averages away and an exit interview reports too late. Edmondson's original finding, that a team's belief in interpersonal safety predicts its learning behaviour and its performance, has strong meta-analytic support, and the DORA Westrum items predict delivery and organisational performance at team level from six behaviourally worded questions (Edmondson, *Psychological safety and learning behavior in work teams*; DORA, *Generative organizational culture*; ECE-RN-01 §1). The constructs are measurable at the group. They have never needed the person.

The manager is where the evidence says to act: Gallup attributes about seventy percent of the variance in team engagement to the manager, and its 2025 figures show manager engagement falling fastest (Gallup, *Managers account for 70 percent of variance in employee engagement*; Gallup, *Global engagement falls*; ECE-RN-01 §3, ECE-RN-04 §2). A culture analysis that cannot say something a manager can act on about their own group, this quarter, without naming anyone in it, has nothing to offer the one person whose behaviour moves the reading. That is the design brief for a finding.

3 The challenges

This section names what goes wrong when organisations try to measure culture continuously, drawn from Red Bear's research note on measurement techniques, law and documented failures (ECE-RN-02 §1, §3, §4). The grouping is Red Bear's own.

3.1 The survey is the only instrument

A single-item employee net promoter score collapses an eleven-point scale into three buckets, is less reliable than a composite, and has no peer-reviewed evidence as a valid engagement measure; three teams with the same score can need entirely different actions (ECE-RN-02 §1, on eNPS). Pulse surveys trade depth for frequency and produce fatigue when not tied to action (ECE-RN-01 §4). Gallup finds that an organisation that surveys and does nothing has lower engagement afterwards than if it had never asked, and only eight percent of employees strongly agree their employer acts on results (Gallup, *Ways to botch employee surveys*; ECE-RN-04 §2). *Mechanism.* Asking is a promise; a survey that produces a number and no next action is an unkept one. The reading needs a source that does not depend on asking, and an owner who can act.

3.2 Content is the temptation

The richest culture signal is what people write to each other, and reading it is what turned Aware into a story: aggregated analytics across 6.5 billion Slack and Teams messages were not the scandal, the eDiscovery tier that can name a person was, and critics asked how a worker faces an accuser (CNBC, *AI might be reading your Slack and Teams messages*; ECE-RN-02 §3). The ICO treats monitoring of email content as likely special-category processing, and the UAE Cybercrime Law criminalises interception of communications (ICO, *Monitoring workers*; UAE Federal Decree-Law 34/2021 Art. 2; ECE-RN-02 §2). *Mechanism.* Content from a channel nobody submitted for the purpose is both the strongest signal and the one that ends the programme.

3.3 Small groups leak

Five is the industry floor for a score, and five is not safe against an insider who knows four of the five, which is why Viva Insights adds differential privacy to manager views and Qualtrics merges small groups rather than hiding them (Microsoft, *Viva Insights differential privacy*; Qualtrics, *Anonymity in EX dashboards*; ECE-RN-02 §4). Time is a second attack: a weekly series on a team of six leaks the leaver. Subtraction is a third: a parent group minus a suppressed child gives the child, which is why Viva Glint suppresses a group whose complement would fall below two, and Culture Amp's recommended mode hides the next-smallest group whenever a filter isolates a group of one (Microsoft, *Viva*

Glint confidentiality thresholds; Culture Amp, *Confidentiality protections in reporting*; ECE-RN-02 §4). *Mechanism.* A floor on group size alone is a floor with three holes in it: time, subtraction and prior knowledge.

3.4 One source, one conclusion

Glassdoor's overall rating correlates only moderately with a validated internal survey (r = .52), pools a company's whole history, and agencies with fewer than fifteen reviews had to be dropped (Bowling Green State University, Glassdoor validity study; ECE-RN-02 §1). Low HR case volume may mean health or silence; NAVEX reads rising volume as confidence in speaking up rather than worse conduct (NAVEX, *Whistleblowing benchmark report*; ECE-RN-02 §1). Customer satisfaction lags and is sector-dependent (ECE-RN-02 §1, service-profit chain). *Mechanism.* Every single source of culture evidence is ambiguous on its own, and a number computed from one of them is a guess with a decimal point.

3.5 Cheap signals get gamed

The bossware literature documents keystroke logging, screenshots and camera monitoring, and the responses they provoke: mouse jigglers, deliberate slowdowns, disabling and resignation (CSCW 2025, bossware study; ECE-RN-02 §3). Message counts and meeting volume are the cheapest signals to read and the cheapest to perform. *Mechanism.* Measured behaviour is performed behaviour; a culture product that reads activity measures the activity of being measured.

3.6 The finding becomes a verdict

Microsoft's Productivity Score put a name next to a number and the stated intent that it "was never designed to score individual users" did not survive the interface (GeekWire, *Microsoft will remove user names from Productivity Score*; ECE-RN-02 §3). Boeing's staff bypassed a speak-up channel they did not trust, and the programme read low volume as good news (Axios, *Boeing safety culture report*; ECE-RN-04 §2). Once a reading is delivered as a fact about a team, the manager can say the algorithm flagged it and stop owning the pattern (ECE-RN-04 §3, mechanism 5). *Mechanism.* A finding phrased as a verdict is defended, delegated or gamed; one phrased as a question about the group's own expectation is answered.

3.7 Text coded without a check

Topic modelling codes free text at scale, but short texts and small samples need human validation before reporting, and derived indices that borrow a construct's name without its validation are a known trap (ECE-RN-02 §1, on text analytics; ECE-RN-01 §1, on the Psychological Safety Index). The Culture 500 reached F-scores of 0.92 and 0.86 with a proprietary dictionary over more than a million reviews (Sull et al., *Measuring culture in leading companies*); a tenant's story stream is smaller by orders of magnitude. *Mechanism.* A story-derived score shown before its dictionary is validated for that aspect on that site is a hypothesis dressed as a reading.

KEY POINTS

- The three holes in a group floor, time, subtraction and prior knowledge, are each closed by a separate rule
- Content is read only where a person submitted it for the purpose, and then only as codes
- A reading from one source is a hypothesis; a finding phrased as a verdict is not owned

4 How Red Bear: Enterprise Culture Evolution helps

Each mechanism in section 3 has a rule in the product, and each rule has a platform requirement that enforces it. The rules are the Red Bear Culture Framework's constraints on signals and their handling, CST-03 to CST-06 and CST-14 to CST-15, its signal model in §3.4 and §5, and the platform's Workplace Culture Analysis and privacy modules (ECE-CF-01 §5; Red Bear Platform, *Platform Functional Specification*, v0.3).

What goes wrong elsewhere	What Red Bear does	Rule
A survey number with no next step	Seven source families read continuously; every finding is addressed to the named owner of a dial with the catalogue's improvement plan and coaching plan beside it, and a coaching session one action away	RBP FR-WA-03, FR-WA-05, FR-WE-03; ECE-CF-01 §3.4, §4
Content read from channels nobody submitted for the purpose (Aware)	Collaboration signals are headers only: counts, timestamps, unit of sender and recipients; bodies, subject lines and attachments are discarded at the connector; content enters only from the three families a person submits to, and is coded on arrival	RBP FR-PA-02, FR-WA-09; ECE-CF-01 CST-03
A floor with three holes	Scores at a group of at least five and free text at least ten; a group is suppressed when its complement in the parent would fall below two; the next-smallest group is hidden when a filter isolates one; no series is shown at a granularity where a period would fall below the floor; small-group manager views are masked; floors are raisable and never lowerable after launch	RBP FR-PA-01, FR-PA-06, FR-PA-07; ECE-CF-01 §5.3, CST-04
One source, one conclusion	A dial combines one calibration scale, one outcome rate and one story score; a finding supported by one family is a hypothesis band with no position and raises no trigger; confidence is always shown	RBP FR-WA-03, FR-WA-05; ECE-CF-01 §3.4, CST-15
Cheap signals gamed	Signals costly to fake, case cycle times, post-mortem closure, customer outcomes, decisions with rationale, outweigh message counts and meeting volume; collaboration metadata never scores participation	RBP FR-WA-12; ECE-CF-01 CST-14
The finding as a verdict	Every finding headline and every trigger is stored as a question to the owner in a fixed form with no verdict word; nothing may cite the engine as the finder of a fact about anyone; a manager's annotation is the only thing a person may add to a reading	RBP FR-WE-16, FR-WE-01; ECE-CF-01 CR-03, CR-09; ECE-FSD-01 FR-FND-06, FR-PRI-10
Text coded without a check	Story-derived scores are a hypothesis until the quarterly calibration pulse, on the open Edmondson and DORA Westrum items, validates the dictionary for that aspect; a programme lead reviews coded records, may re-code or discard, and never sees text after coding	RBP FR-WA-14, FR-WA-09; ECE-CF-01 §5.2; ECE-FSD-01 FR-SIG-05, FR-FND-08
A name next to a number	The engine holds no score, reading, flag or nudge for an individual; no query by any role, including Red Bear's operator, returns a row below the floor; identifiable rows do not exist after the connector	RBP FR-WA-10, FR-WA-11, FR-PA-03; ECE-CF-01 CST-02, CST-12

4.1 The catalogue and its sources

The engine analyses against a fixed catalogue of fifty aspects, twenty-five positive and twenty-five toxic, in eight principle families (RBP FR-WA-01; ECE-CF-01 §4). Each aspect carries a definition, its framework source, its signals by source family, a two- or three-move improvement plan and a Socratic coaching plan. The catalogue was assembled from open models and instruments: Edmondson's seven items and the three leader behaviours, Westrum's six

dimensions as DORA operationalises them, the MIT SMR/Glassdoor Big 9 and Toxic Five as paraphrased definitions, Denison's four traits as a coverage check, Google's five team dynamics, and Gallup's five burnout causes (ECE-RN-01 §1, §2). No proprietary instrument is reproduced: the Denison survey, the OCAI, the Gallup Q12 and the Maslach Burnout Inventory are cited as evidence and never administered or imitated (ECE-CF-01 §1.4).

The catalogue is what makes analysis possible without individuals. A signal is evidence for a behaviour; a behaviour is evidence for a principle; nothing is evidence about a person (ECE-CF-01 §2.1). Because every aspect names its signals in advance, the engine never goes looking for a pattern in the data; it asks, per group and per period, whether the signals the catalogue names for an aspect moved. A toxic aspect is never chosen; it is a health indicator whose rise lowers its counterpart's dial (ECE-CF-01 §2.3).

4.2 Coding on arrival

Submitted content is coded to aspects and categories as it arrives, in the manner of the Culture 500, with Red Bear's own dictionaries (RBP FR-WA-09; Sull et al., *Measuring culture in leading companies*). Names of people in the text are redacted before coding; a story that names a person is coded for its aspect and its verbatim is never shown. The author identifier is replaced by unit and period at the connector, or never reaches it if the employee chose the anonymous option, which is on by default (ECE-FSD-01 FR-SIG-02). The employee receives a receipt stating what was kept, the aspects, the unit and the period, and what was discarded, the text and the name; the receipt never quotes the story.

A coded record has the same shape from every content family: aspects, categories, unit, period, coding confidence. The programme lead may accept a coding, re-code it or discard it; every action is logged with the lead's name, and a re-code changes a code, never a text (ECE-FSD-01 FR-SIG-05). Story-derived signals stay a hypothesis until the calibration pulse has validated the dictionary for that aspect (RBP FR-WA-14): the story family supplies texture and early warning, and the open scales decide whether the texture is real.

4.3 Three families, one reading

A dial's position is computed from three signal families in descending evidential strength (ECE-CF-01 §3.4; RBP FR-WA-03). **Survey-calibration** is the open Edmondson and DORA Westrum items, issued quarterly through the product and aggregated above the floor, or a tenant's own engagement survey imported as thresholded results, never per person; it anchors the scale to published norms. **Outcome-derived** is the rates that are continuous and hard to game: post-mortems completed per incident, decision records with rationale, escalation lead time, recognition per head, attrition by unit cohort, cycle time, after-hours share; it carries the trend between calibration points. **Story-derived** is the coded prevalence of each aspect across stories, the mailbox and the public page; it is the early warning. For a *messengers trained, not shot* dial (P-03) the reading is the DORA messenger item, plus escalations made before the deadline as a share of all escalations, plus the prevalence of early-escalation narratives. One family alone shows a hypothesis band and no position (RBP FR-WA-05).

4.4 Suppression, complement, time windows, noise

Four rules close the three holes of §3.3, and all four are applied in the privacy service before the engine reads a row (RBP FR-WA-04; ECE-CF-01 §5.3). The **floor**: a group of at least five for a score and ten for anything coded from text, raisable per tenant and never lowerable once a programme has started (RBP FR-PA-01). **Complement suppression**: a group is suppressed when its complement within the parent would fall below two, and when a filter isolates a group of one the next-smallest group is hidden as well, so nobody can be found by subtraction. **Time-window minimums**: no series is shown at a granularity where a single period would fall below the floor; the service coarsens the window until every period clears it, and a series that cannot clear it at any granularity is not shown (RBP FR-PA-06). **Noise**: any

manager-facing view of collaboration metadata for a group near the floor is distribution-masked so that *almost all* and *almost none* cannot be read, and the masking is recorded on the view (RBP FR-PA-07). Wherever any rule applies, every role sees the word *suppressed* with the rule that caused it, and no number, range or hint of the count; every suppression is logged with the query behind it (ECE-FSD-01 FR-FND-05).

4.5 What a finding is, and is not

> **WHAT A FINDING IS**
>
> A finding is a pattern observed for a group above the floor, on one catalogue aspect, over one period: its direction of travel and the size of the movement, the signal families behind it and how much data each held, a confidence band showing how far the families agree, the number of people aggregated, and the catalogue entry and version it matches (RBP FR-WA-05). It is stored as a question to the named owner of the group's dial, in the form *the reading for [group] on [aspect] moved [direction] over [period]; what did you expect?* (RBP FR-WE-16).

A finding is not a verdict: no headline uses *good*, *bad*, *failing*, *toxic team* or *at risk* about a group (ECE-FSD-01 FR-FND-06). It is not about a person: it names none, quotes no text, and cannot be filtered, sliced or windowed to describe one, because the product offers no filter on any attribute of a person and holds none (RBP FR-WA-10; ECE-FSD-01 FR-FND-07). It is not a level: direction of travel is shown with the position because feedback that shows velocity is acted on and feedback that shows a level invites comparison and defensiveness (Kluger and DeNisi, *The effects of feedback interventions on performance*; ECE-RN-03 §1). It is not a fact the engine found: a programme item, decision record, communication or export cannot cite the product as the finder of a fact about anyone, and an attempt to record a trigger as a finding is refused (RBP FR-WE-16; ECE-CF-01 CR-09). And it is not a ranking: no view sorts managers, teams or units by position or trend, and the product offers no sort control on either (RBP FR-WE-08; ECE-FSD-01 FR-PRI-09).

> **STANDING RULE**
>
> AI proposes, a named person decides. The engine proposes a reading, a trend, a confidence and a finding; the dial's owner reads, annotates, opens a session or proposes a programme item. No role can edit a position, and nothing the engine observes becomes a record until a person puts it there in their own words (RBP FR-WE-16, FR-DC-10; ECE-CF-01 CST-13).

5 How it is used

The analysis flow runs from a signal to a finding to its owner, and at no step does a person's name attach to a signal.

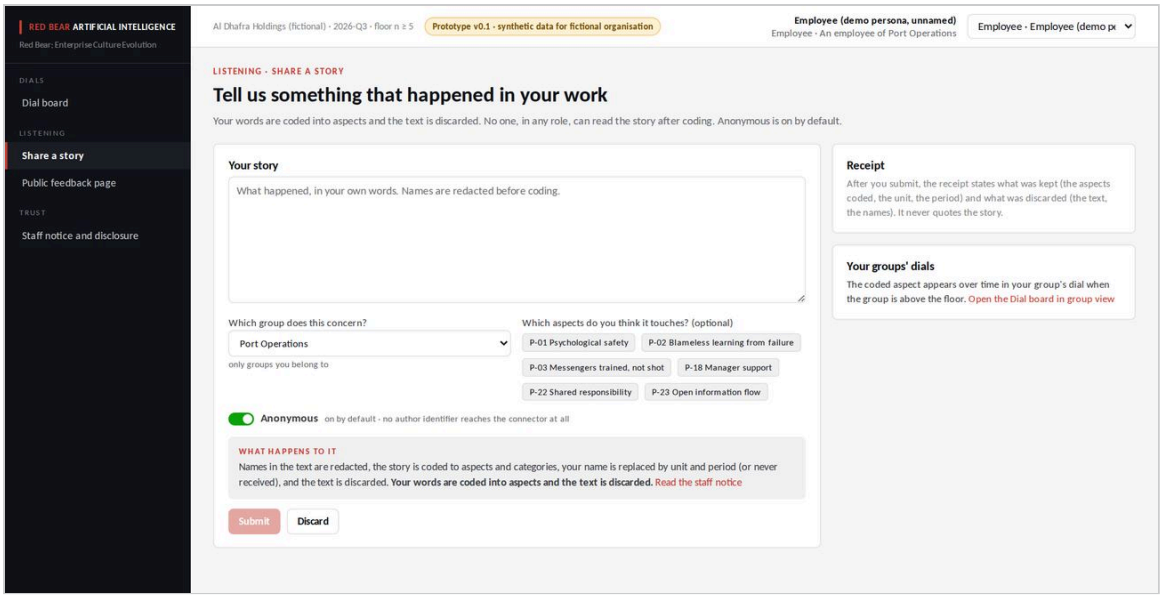
Step	Who	What happens
1 Enable a family	System administrator	A source family is enabled only while the go-live gate is open for it: the consultation record, the impact assessment and the staff notice exist and cover that family. Enabling or disabling a family closes the gate until the notice and pack are re-versioned (RBP FR-PA-09; ECE-FSD-01 FR-SIG-06)
2 Connect or submit	Connector; employee	Headers become unit-level counts and identifiers rotating tokens; bodies, attachments, parties and participant names are discarded (RBP FR-PA-02, FR-PA-03). A story, mailbox message or public post is written for the purpose, anonymous by default, with a receipt (ECE-FSD-01 FR-SIG-01, FR-SIG-02)
3 Code	WA engine	Names redacted, aspects and categories coded with a confidence, author replaced by unit and period, text discarded or held under the free-text floor (RBP FR-WA-09)

Step	Who	What happens
4 Review coding	Programme lead	Coded records by aspect, unit and period; accept, re-code or discard; no text, for any role (ECE-FSD-01 FR-SIG-05)
5 Aggregate	Privacy service	Floors, complement suppression, time-window minimums and noise applied in the data layer before analysis (RBP FR-WA-04, FR-PA-01, FR-PA-06, FR-PA-07)
6 Analyse	WA engine	Each group read against its chosen principles and the whole toxic side of the catalogue; three families combined per dial with the calibration scale as reference; one family alone is a hypothesis (RBP FR-WA-02, FR-WA-03, FR-WA-05)
7 Find	WA engine	A finding per group, aspect and period, framed as a question to the dial's owner; a toxic aspect rising or a principle falling for two consecutive windows raises a coaching trigger to that owner (RBP FR-WE-16, FR-WE-03)
8 Read and act	Dial owner	The owner opens the pattern detail, annotates the reading in their words, opens a coaching session or proposes a programme item; the reading itself cannot be edited (ECE-FSD-01 FR-FND-02, FR-PRI-10)
9 Calibrate	Programme lead	The quarterly pulse on the open items anchors each dial's scale and validates the story dictionaries (RBP FR-WA-14; ECE-FSD-01 FR-MEA-05)

Who sees a finding. The programme lead sees findings across granted scopes; the sponsor sees the organisation and every unit as an unordered pattern; a manager sees their own teams and never another manager's; an employee sees no finding, only the dials of their groups, read-only (ECE-FSD-01 §2.3, FR-PRI-08). HR receives only the health monitor's proposals and the escalations a person chooses to make; the system administrator sees connector health and counts and no content (ECE-FSD-01 FR-SIG-08).

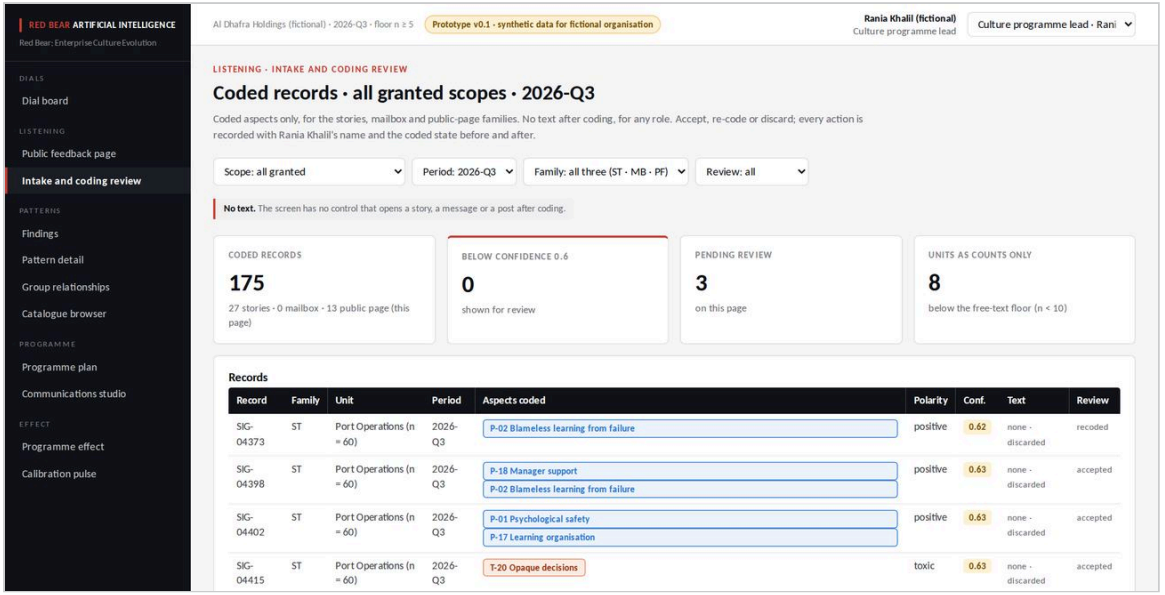
6 The user experience

The intake and analysis clusters have eight screens; five are shown here as the prototype renders them, on synthetic data for a fictional organisation, Al Dhafra Holdings, with every group above the floor by construction.

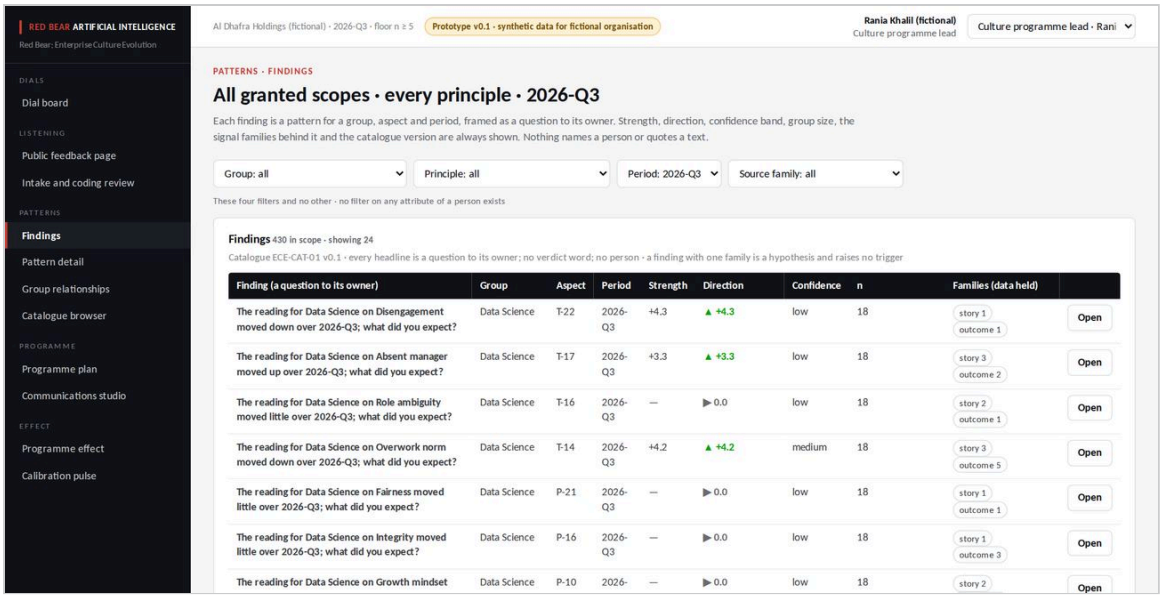


S-05 Share a story: S-05 Share a story. What an employee sees: a text box, the choice of group from the groups they belong to, optional aspect tags, and the anonymous option on by default. Beneath the box, in plain words, what happens to the story: names redacted, coded to aspects, the author replaced by unit and period or never received, the text discarded. The receipt panel states what is kept and what is not, and never quotes the story. Prototype render on synthetic data for fictional entities.

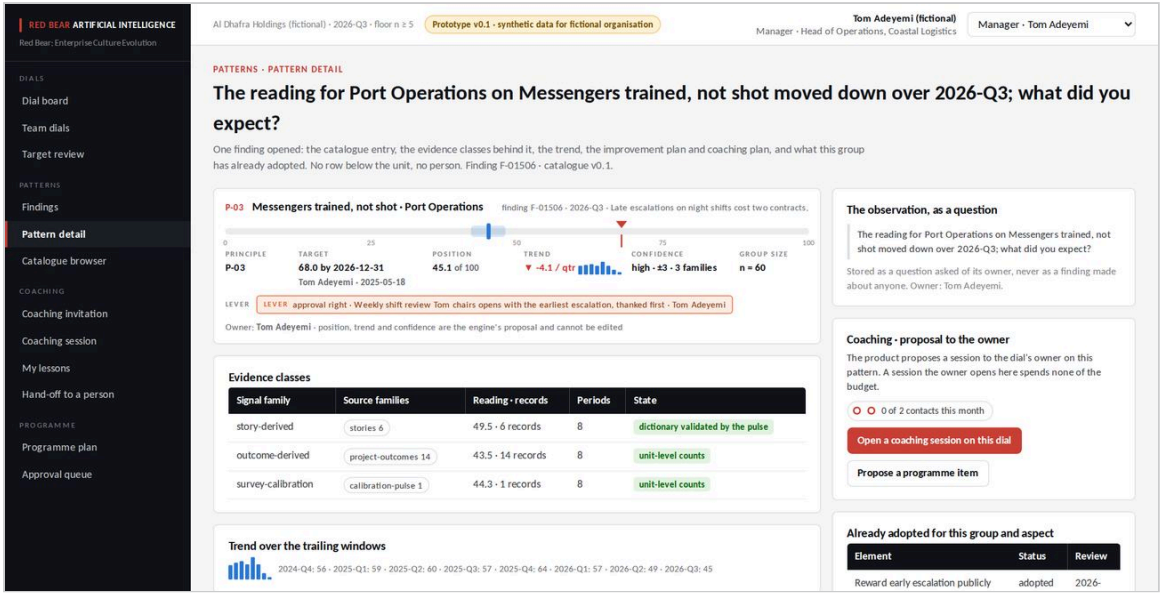
The intake screen is where the product earns the trust the analysis depends on. The employee reads, before submitting, that no one in any role can read the story after coding, and the receipt afterwards says the same in the past tense. There is no reaction count and no visibility of who else has submitted; the story is a contribution to a group's reading.



S-07 Intake and coding review: S-07 Intake and coding review. The programme lead's view of coded records for the stories, mailbox and public-page families: record id, family, unit with its group size, period, the aspects coded with polarity and confidence, and the review state. The text column reads *none*, *discarded* on every row. Units below the free-text floor appear as suppressed counts. Accept, re-code and discard are the only actions, each logged with the lead's name. Prototype render on synthetic data for fictional entities.

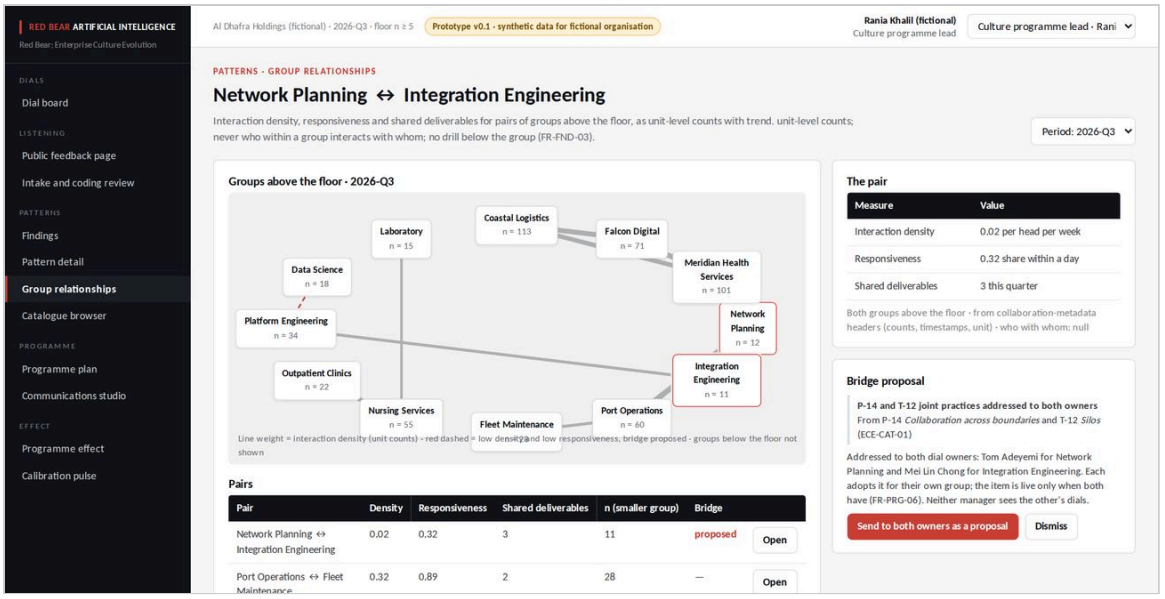


S-09 Findings: S-09 Findings. Every finding in the lead's granted scopes: the headline as a question to its owner, the group, the aspect, the period, the strength and direction of the movement, the confidence band, the group size, and the signal families behind it with the number of records each held. Single-family findings are marked as hypotheses with no position; a suppressed row shows the rule and no count. Four filters and no other. Prototype render on synthetic data for fictional entities.



S-10 Pattern detail: S-10 Pattern detail. One finding opened by the dial's owner: the dial with its seven parts, the evidence classes behind the finding (which signal families, which source families, how many periods, whether the story dictionary is validated for this aspect), the trend over the trailing windows, the catalogue entry with its improvement and coaching plans, the elements already adopted for this group and aspect, and the observation as a question. A coaching session is one action away and spends no budget when the owner opens it. Prototype render on synthetic data for fictional entities.

The pattern detail is built to answer *what is behind this?* without ever answering *who?*: the evidence panel says how many records each family held and whether the dictionary is validated, and the footer states on every visit that there is no drill below the unit and no person (ECE-FSD-01 FR-FND-02).



S-11 Group relationships: S-11 Group relationships. Pairs of groups above the floor with interaction density, responsiveness and shared deliverables as unit-level counts, and a bridge proposal addressed to both groups' dial owners where density or responsiveness is low. The graph shows groups, not people; there is no drill below the group and no reading of who within a group interacts with whom. Prototype render on synthetic data for fictional entities.

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

The table gives, per source family, the signal the engine reads for one aspect, so that the reader can see what a family contributes and what it cannot.

Source family	Example signal, per group above the floor	Signal family	What the connector discards
stories	Prevalence of *we learned* narratives against *who did it* narratives (P-02)	Story-derived	Text after coding; author, replaced by unit and period
mailbox	Early risk flags per period (P-03)	Story-derived	Body after coding; sender; attachments never read
public-page	Recognition mentions per period (P-05)	Story-derived	Names redacted before coding; nothing identifying ingested
project-outcomes	Post-mortems filed per incident and actions closed (P-02); late surprises per project (P-03)	Outcome-derived	Post-mortem text; names of participants
decision-records	Share of decisions with rationale recorded (P-04); decisions taken below the top level (P-06); outcome against expectation	Outcome-derived	Decision content, unless the owner opens it to a session
hr-cases	Conduct-category volume and cycle time, read against speak-up items (T-03)	Outcome-derived	Parties, narrative, evidence: never ingested
collaboration-metadata	Cross-unit interaction density (P-14); after-hours share (P-20); one-to-one cadence (P-18)	Outcome-derived	Subject lines, bodies, attachments, individual identifiers
calibration pulse	The open Edmondson and DORA Westrum items, quarterly	Survey-calibration	Individual responses, aggregated above the floor before any reading

7.1 How it is checked

- **Refusal tests on every pipeline run.** The below-floor read is attempted under all seven personas, the tenant administrator, the Red Bear operator and the hosting administrator, and must return *suppressed*; content refusal is tested per family; a second mailbox connect is refused because no such call exists; gate-closed ingest is refused. 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).
- **The same configuration drives three things.** What each family reads and discards is generated from the configuration that drives the connector, the consultation view and the staff notice, so the three cannot drift (RBP FR-WA-13; ECE-FSD-01 FR-SIG-09).
- **Calibration before exposure.** No aspect score from the story family is shown as a position until the pulse has validated the dictionary for that aspect; where a tenant runs no pulse, every story-derived finding stays a hypothesis and every dial's confidence is capped, and the dial says so (RBP FR-WA-14; ECE-FSD-01 FR-FND-08).
- **Query logging.** Every analyst or dial-owner query against aggregates is logged with role, scope, filters, group sizes returned and time, holds no content, and is summarised on the consultation view (RBP FR-PA-10).
- **Retention.** Raw events for days, as long as aggregation takes; aggregates for the trend, with 27 months as the ceiling; coded stories keep their codes and lose their text at the configured boundary; expired data is deleted, not archived, and the deletion is logged (RBP FR-PA-04; ECE-CF-01 §5.4).

7.2 What is never read

Message bodies, subject lines or attachments from any channel. Any mailbox other than the one address the organisation published for the purpose. The parties, narrative or evidence of an HR case. The names of post-mortem participants. Any attribute of a person: role, tenure, nationality, gender, grade. Any sentiment, emotion, personality or mental-state inference from any signal, at any level. The platform offers no call, schema or connector for any of these, so the product cannot read them even by mistake (RBP FR-PA-08, FR-WA-08; ECE-CF-01 CST-01, CST-03, CST-05).

8 Benefits realised

The analysis is judged on whether it produces readings a named owner acts on and staff trust, never on how much of it there is. The measures are the framework's, each with a *Never* clause (ECE-CF-01 §9.1).

Benefit	Measured by	Never
Toxic patterns are seen early and decline	ECE-CF-01 KPI-02 toxic pattern decline: prevalence by family across units above the floor, corroborated by the calibration pulse	Per unit below the floor; attributed to an intervention without the owner's recorded review
Staff trust the instrument enough to feed it	ECE-CF-01 KPI-05 signal trust: response share to the quarterly pulse; stories submitted per hundred staff per window	A target for managers to drive; driven participation is gamed participation
Speaking up is read honestly	ECE-CF-01 KPI-06 speak-up balance: case-category volume read together 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
Findings lead somewhere	ECE-CF-01 KPI-03 lesson rate: share of coaching triggers raised from findings that end in a recorded lesson, in aggregate	Per manager; declined triggers counted against anyone

What is explicitly not claimed. The product does not claim to measure culture as a number: a dial is a position with a trend, a confidence and a group size, and a family dial inherits the lowest confidence of its members (ECE-CF-01 §3.3, CST-15). It does not claim that a story-derived finding is true until the calibration family agrees. It does not compute an engagement single number, does not administer the Gallup Q12, and does not claim the causal path from engagement to profit that Gallup's meta-analysis reports; that evidence is why a customer buys the programme, not a measure the product computes (ECE-CF-01 §9.2). And it does not claim the Culture 500's accuracy for a tenant's story stream: a site that has not run its calibration period shows *calibrating* and raises no trigger (RBP FR-DC-16; ECE-FSD-01 FR-PRI-05).

The honest reading. A finding built on three families with high confidence for a group of sixty says that a pattern moved. It does not say why, and it does not say who. The *why* is what the owner's expectation, recorded on the target decision, and the coaching session are for; the *who* is what the product is built never to answer.

9 Risks managed

Residual risk	Mitigation
A story names a colleague and the name survives redaction	Names are redacted before coding; the public page is moderated once, pre-coding, for names the redactor missed and for content that belongs with HR, which is routed as a proposal to the HR partner and never coded; the moderation view is logged (ECE-FSD-01 FR-SIG-03)
An insider infers a colleague from a group of five	Complement suppression, the next-smallest group hidden, time-window coarsening and noise on small-group views; the tenant may raise the floor to ten at launch and can never lower it (RBP FR-PA-01, FR-PA-06, FR-PA-07; ledger OD-3)
A dictionary codes a site's stories wrongly	Story-derived findings are hypotheses until the pulse validates the dictionary per aspect; the programme lead can re-code; coding confidence is shown on every record (RBP FR-WA-14; ECE-FSD-01 FR-SIG-05)
Low case volume is read as health when it is silence	Case-category readings are always paired with the speak-up calibration items and the ambiguity is stated on the screen (ECE-CF-01 KPI-06; ECE-FSD-01 FR-HLT-05)
A finding is quoted in an appraisal	No reading, finding or lesson may enter an appraisal, objective or reward decision without a Campbell's-law review, and none is exposed per person; the finding's stored form is a question and cannot be cited as a fact (RBP FR-WE-15, FR-WE-16; ECE-FSD-01 §2.4)
A family is enabled before consultation	The go-live gate is in the connector service; nothing is ingested until the consultation record, the impact assessment and the staff notice exist and cover the family (RBP FR-PA-09; ECE-CF-01 CST-19)
Gulf sites read differently from the sites the dictionaries were built on	A calibration period per site before its thresholds are trusted; the product says it has no site evidence yet (RBP FR-DC-16; ECE-CF-01 CR-17)

What Red Bear will not build, even if asked. No sentiment analysis at individual level, in any market. No connector to a channel nobody submitted to for the purpose. No filter on any attribute of a person. No export of a row below the floor, for any role, including Red Bear's own. A customer whose requirement is any of these has described a monitoring product, and Red Bear says so before the pilot (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.

Cited as	Source	Authority
ECE-RN-01 to ECE-RN-04	Red Bear research notes: *Culture frameworks, aspects and interventions*; *Continuous culture measurement: techniques, law, failures and design constraints*; *Socratic coaching from patterns*; *Why AI culture coaching has failed: cases and lessons* (research/frameworks.md, monitoring-and-law.md, coaching.md, failures.md)	Red Bear research notes
ECE-CF-01; ECE-FSD-01; Platform FSD	Red Bear, *The Red Bear Culture Framework* v0.2 (§1 to §5, §9); Red Bear, *Functional Specification* ECE-FSD-01 v0.1 (modules SIG, FND, MEA, GOV); Red Bear Platform, *Platform Functional Specification* RBP-FSD-01 v0.3 (modules WA, WE, PA)	Red Bear standard
Schein	E. Schein, *Organizational Culture and Leadership*, the three levels	Academic reference
Edmondson; DORA; Westrum	A. Edmondson, *Psychological safety and learning behavior in work teams*, ASQ (1999); DORA, *Generative organizational culture*; R. Westrum, *A typology of organisational cultures*	Academic and industry reference
Sull et al.	Sull, Turco et al., *Measuring culture in leading companies* (MIT SMR/Glassdoor Culture 500); Sull and Sull, *Why every leader needs to worry about toxic culture*; Sull, Sull and Zweig, *Toxic culture is driving the Great Resignation*, MIT Sloan Management Review	Academic reference
Gallup	*Managers account for 70 percent of variance in employee engagement*; *Global engagement falls for the second time since 2009* (2025); *Ways to botch employee surveys*; *Q12 meta-analysis*	Industry research
Microsoft	*Viva Insights privacy*; *Viva Insights differential privacy*; *Viva Glint confidentiality thresholds*; *Viva Glint survey privacy*	Vendor reference
Culture Amp; Qualtrics	Culture Amp, *Confidentiality protections in reporting*; Qualtrics, *Anonymity in EX dashboards*	Vendor reference
Glassdoor validity; eNPS; Psychological Safety Index	Bowling Green State University, Glassdoor validity study; Talent Company, *The problem with eNPS*; Psych Safety, *The Psychological Safety Index: a critical look*	Academic reference; practitioner commentary
NAVEX	*Whistleblowing benchmark report* (2025)	Industry reference
Aware; Productivity Score; Boeing; bossware	CNBC, *AI might be reading your Slack and Teams messages using tech from Aware* (2024); GeekWire, *Microsoft will remove user names from Productivity Score* (2020); Axios, *Boeing safety culture report* (2024); CSCW 2025 bossware study	Press reporting; academic reference
ICO; UAE Cybercrime Law	ICO, *Employment practices: monitoring workers* (2023); UAE Federal Decree-Law 34/2021 Art. 2	Regulatory guidance; statute
Kluger and DeNisi	A. Kluger and A. DeNisi, *The effects of feedback interventions on performance* (1996)	Academic reference

Glossary

Term	Meaning in this paper
Source family	One of seven connectors: stories, mailbox, public-page, project-outcomes, decision-records, hr-cases, collaboration-metadata; each enabled per tenant after consultation
Signal family	One of three evidential classes a dial combines: survey-calibration, outcome-derived, story-derived
Coding on arrival	The replacement of submitted text by catalogue aspects and categories at the moment it arrives, with names redacted and the author replaced by unit and period
Finding	A pattern for a group, aspect and period with direction, families, confidence and group size, stored as a question to the dial's owner
Hypothesis band	The form a finding takes when only one signal family supports it, or the story dictionary is not validated: a band, no position, no trigger
Floor	The minimum group size below which nothing is shown: five for scores, ten for free text; raisable, never lowerable after launch
Complement suppression	Suppressing a group whose complement within its parent would fall below two, so a person cannot be found by subtraction

The series

This paper is the second 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 (this paper)	Draft
ECE-WP-03	Privacy, Trust and No-Attribution	The schema, the gate and the consultation view	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