RED BEAR: ENTERPRISE CULTURE EVOLUTION · DATA SHEET · D · COACHING

A question, not a verdict

Feedback arrives as a verdict, so managers defend instead of learn. Coaching invites the owner of a pattern to a private Socratic conversation, when something has happened, and keeps only the lesson they write.

KEY POINTS

- A session opens only on an event: a decision outcome missed its expectation, a toxic pattern rose for two windows, a chosen principle fell for two windows, or the manager asked; no calendar, quota or campaign trigger exists
- It coaches the named owner of a dial or a decision about a group; nobody is coached about themselves, and nobody can assign coaching to anyone
- The agent says it is an AI in its first utterance, opens with *what did you expect?* before it mentions what the engine saw, and cannot agree with a causal account it has no evidence for
- Two contacts a month at most, a named human one request away, no transcript kept, no streaks, badges or persona

What it is

Instead of a report to a manager, the named owner of a dial or a decision who has a lesson to learn is invited to a private conversation with a coaching agent. The invitation states the pattern as a question, the dial or decision, the period, the trigger kind and the budget remaining, and offers accept, defer or decline; a decline is not recorded against anyone. The session runs a fixed after-action sequence: expectation, what happened, what kind of failure, what would be done differently, what changing it would put at risk, and what was most useful. The engine's observation is shown only after the manager has stated their expectation, and then as a question.

The only durable output is the lesson, in the manager's own words, with the decision or dial it applies to, the next action and its owner, kept, edited or deleted by its author. No transcript, audio or turn log exists to show.

The pain point

The case file on AI coaching is long. Humu sent ninety kinds of nudge for six years and was absorbed as a feature, because a prompt with no event behind it decays to noise and then to resentment. Koko's messages were rated well until recipients learned a machine wrote them. NEDA's chatbot gave harmful advice after a vendor changed its generation mode. Replika showed what a warm, available agent that never disagrees becomes. And the one controlled trial of AI coaching for senior leaders found no effect and participants asking for a person.

Each case is a rule here. Event, not calendar. A budget on the agent's own attention. Disclosure every session. A refusal to agree with blame. A named human, not a queue, one request away, and no pattern carried beyond two sessions without a lesson. Versioned prompts and a regression gate on any model change, including a provider's. And an agent designed to be dispensable: the success case is a manager who needs it less each quarter.

How it works

Invitation within a budget. A Coaching invitation is shown only to the named owner the trigger names. An invitation that would exceed the monthly budget of two contacts is held and dropped if the pattern clears; a session the manager opens on their own dial spends no budget; abuse and harassment patterns raise no invitation at all, because they go to HR (FR-COA-01). The product has no control by which anyone assigns coaching to anyone.

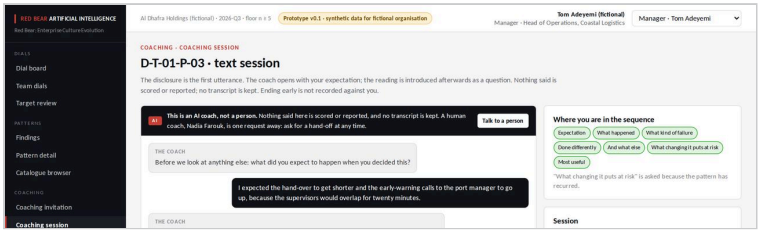
The session. It opens with the disclosure, naming the tenant's human coach and how to reach them, and no control proceeds until it has been shown (FR-COA-02). The sequence is shown on screen and the manager may end it at any stage without that being recorded (FR-COA-03). When the agent declines, to attribute a signal to a person, to diagnose, to score, to infer an emotional state, to praise, the screen shows the refusal and the rule in one sentence and offers the question it can ask instead (FR-COA-04). Hand-off to a person is one action from every turn; the agent opens it on any sign of distress, on conduct in the reckless band, on a relational or high-stakes pattern, and after two sessions without a lesson. What passes is a summary the manager writes, never a transcript; a hand-off to HR happens only when the manager chooses it (FR-COA-06). Text is the default; voice is off unless the tenant enables it under the Annex III file, and then audio is discarded in the stream with no prosody computed (FR-COA-07). Question forms never imply a superior erred where a third party could see; Arabic is offered first where it is the manager's first language, and a site's thresholds are not trusted until it has run its calibration period (FR-COA-09).

Change control. The coaching release in force, its prompt set, generation mode, model mapping and evaluation set, the last regression result and the cases it covers are visible to the administrator, who can change none of it; a release that has not passed cannot be in force, and the version is recorded on every session and lesson (FR-COA-10).

What you get

Capability	What it does	Requirements
Coaching invitation	To the owner only, as a question, with budget remaining; accept, defer, decline; held when over budget	FR-COA-01
Disclosed Socratic session	AI disclosure first; after-action sequence shown; observation after expectation, as a question	FR-COA-02, FR-COA-03
The never-say list	Refusals shown with their rule; each a case in the regression suite	FR-COA-04
My lessons	Draft, edit, keep, discard, delete; the manager's own only; no counters, no history	FR-COA-05, FR-COA-08
Hand-off to a named human	One action from every turn; a summary the manager writes; HR only by the manager's choice	FR-COA-06
Voice, region and language	Voice per tenant and off by default; face-preserving forms; Arabic first; calibration per site	FR-COA-07, FR-COA-09
Release and regression evidence	The release in force, its regression result and history, read-only; effect shown only in aggregate	FR-COA-10, FR-COA-11

Screens



S-14 Coaching session: The Coaching session: the disclosure first, the after-action sequence shown, the engine's observation introduced as a question, and a hand-off to a named person one action away. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Two measures carry this pillar (ECE-CF-01 KPI-03 and KPI-04): the **lesson rate**, the share of coaching triggers that end in a recorded lesson, in aggregate; and **lesson application**, the share of lessons applied to a later decision whose outcome met expectation, never shown with fewer than ten lessons in the set. Sessions, satisfaction and return rate are diagnostics, not outcomes: the senior-leader trial found engagement and effect uncorrelated. Nothing is computed per manager, and no screen can list managers by any coaching quantity. What coaching makes possible is a manager who owns a pattern instead of defending a score, a lesson attached to the decision it came from, and a human coach who receives the cases that need one.

Platform engines

Coaching consumes the Red Bear Platform's Decision Feedback Coaching engine, which fixes the agent's behaviour: triggers, sequence, voice handling, lesson record, aggregate measurement, escalation, the never-say list, disclosure, the nudge budget, change control, no bonding features and regional forms (RBP FR-DC-01 to FR-DC-16), the Workplace Culture Evolution engine for the trigger to the owner (RBP FR-WE-03), and the Decision Learning engine for lessons as patterns (RBP FR-DL-02). Platform behaviour is specified in the platform functional specification, not here.