RED BEAR: ENTERPRISE DECISION INTELLIGENCE · SUITE DATA SHEET · FORMERLY GUARDRAIL

Govern every decision, whoever, or whatever, made it

A common decision intelligence and governance platform for people, AI agents and applications. **Empower. Trust. Verify.**

KEY POINTS

- One register for every consequential decision by a person, an application or an AI agent, holding the reasoning and not just the result
- The dial belongs to the decider: one Trust Dial per person or agent, 0 to 10 in five bands, widened as trust is earned; an agent's ceiling is 9
- The guardrails belong to the decision type: nine fixed kinds, set once, applied to people and agents alike
- Three verbs of intervention, approve, reject, ask; every override carries its reason; decision and outcome quality are never blended
- The platform proposes and never decides: a named person confirms, approves, moves the dial or activates the policy

What it is

Organisations have systems of record for money, people, customers and inventory. They have none for judgement. The decision to give the discount, accept the supplier or let the agent act is the thing that moved the number, and it is the one thing no system keeps. What survives is the transaction, stripped of the reasoning.

Red Bear: Enterprise Decision Intelligence puts one layer underneath all of it. Every consequential decision, taken by a person in a meeting, by an application in a workflow, or by an agent at three in the morning, enters the same register, is evaluated against the same policy, and is delegated by the same trust model. The register holds the reasoning, not just the result, so the organisation can find out whether it was any good.

The pain point

The number of things inside a business that can decide has gone up sharply. An AI agent that adjusts pricing, releases a payment or approves a claim is deciding with real consequence, at a rate no human review can keep pace with. Most organisations either refuse the agent authority, and get none of the value, or grant it and hope. Neither is governance.

The trigger is not an industry. It is a condition: **authority distributed faster than judgement can follow it.** That happens in a company scaling its headcount, in one deploying AI agents into operational work, in one that has just been through a governance failure, and in one whose approval chains have become the reason nothing ships.

How it works

Five pillars, each useful on its own and considerably more useful together, on a platform and service layer that underpins all of them (EDI-DS-08).

- **Capture** (EDI-DS-01). Meetings and calls, a decision inbox, connected applications, a custom connector, portal and mobile, and an API agents call before they act.
- **Central decision register** (EDI-DS-02). One append-only entry per decision: options, reasoning, the policy in force at the time, who was involved, and what happened next.
- **Policy and trust control plane** (EDI-DS-04). What each actor may settle alone: guardrails on the decision type set the boundary, the decider's Trust Dial sets the room inside it.
- **Intervention and escalation** (EDI-DS-03, 05, 06). When a decision falls outside that room: who is brought in, whether they are there, how fast, with what context and precedent, what proceeds meanwhile.
- **Analysis and learning** (EDI-DS-07). Whether the decisions were any good: outcomes, the quality quadrant, override telemetry, the decisions nobody made.

Discretion as a dial, not a door. Most governance offers two settings: the decision needs approval, or it does not. That binary is why approval chains grow, and why the organisation is slow exactly where it has the most experience. The **Trust Dial** replaces it: one per decider, person or agent, 0 to 10, in five bands that are each a review rate. *No Discretion* (0): everything is referred. *Recommend* (1–3): they propose, a human decides. *Supervised* (4–6): narrow limits, every decision reviewed. *Delegated* (7–9): wide limits, decisions sampled. *Autonomous* (10): reviewed by exception; still recorded, never

Guardrails belong to the decision type. A rule about risk and a judgement about a person are not the same setting, and writing them on one line of a delegation matrix is why the matrix is never current. Guardrails are set per decision type, in nine fixed kinds: authority, policy, financial limits, risk appetite, compliance, segregation of duties, exceptions, human intervention and auditability. The range a decider has is the product of the dial on the decider and the guardrail on the type, bounded by role scope. Widening a good team's room does not reopen the risk appetite, and

ungoverned. A dial moves one point at a time, by a named person, with a reason, after seeing the projected effect (FR-DEL-02).

tightening the risk appetite does not reopen whether the team is any good (FR-DEL-19, FR-DEL-06).

How you adopt it

Not across the estate at once. Five stages, one decision class at a time, each worth stopping at; the first two change nothing about how anyone works (FR-PLT-06).

- **Connect.** Capture routes wired to one or two decision classes; no policy enforced, nothing blocked. *You have* a register that is filling, and the first honest view of where decisions are made.
- **Observe.** Policy evaluated and recorded, but advisory; nothing is stopped. *You have* evidence of how often policy is actually met, before anyone is asked to comply.
- **Govern.** Guardrails enforced on the agreed classes; dials set; escalation and reach live. *You have* governance operating, with the delegation mix visible and adjustable.
- **Learn.** Outcomes attached; override telemetry reviewed; policies and dials adjusted on evidence. *You have* policy that changes because of what happened, not because of what was feared.
- **Compound.** More classes, more actors, agents taking authority class by class. *You have* an organisation whose decision quality is a measured, managed property.

Observe to Govern and Learn to Compound are leadership decisions, not delivery milestones: the first time a decision waits because the platform says so, and the first time authority passes to something that is not a person. The entry point is a **90-day fixed-scope pilot** on one decision class: Connect and Observe, enforcing nothing, ending with evidence of coverage, timing and conformance that either justifies Govern or does not.

RETURN ON INVESTMENT

Eight benefits, each with a measure baselined before a pilot and tracked after it. **None has a target.** Targets come from the first engagement's baseline, not from this document (FR-PLT-05).

- **Full visibility.** Share of decisions in governed classes that reach the register; time from decision to capture.
- **Faster decisions at the edge.** Share settled without referral and within policy; time to decide against target.
- **Leadership time where it counts.** Time from trigger to senior decision; queue age; share of referrals that change the outcome.
- **Better-informed decisions.** Share of decisions with context viewed; evidence completeness at capture.
- **Confident AI adoption.** Classes where agents decide within guardrails; agent override and reversal rates.
- **An organisation that learns.** Outcome coverage; stalls found; dial, guardrail and standard changes traced to a lesson.
- **Resilience as conditions change.** Share of policies replayed in the sandbox before commit; time from condition change to policy change.
- **Provable accountability.** Time to reconstruct a decision with the policy of the day; regulatory controls evidenced; residency compliance.

Platform engines

Built on the Red Bear Platform, whose FSD specifies the Decision Intelligence, Presence, Cartesian Decision Analysis, Decision Learning, Decision Feedback Coaching and Research engines and the shared identity, connector, provenance, AI gateway, privacy, audit and hosting services this product consumes.