RED BEAR: ENTERPRISE DECISION INTELLIGENCE · DATA SHEET · H · PLATFORM AND SERVICE

Measure first, enforce second, automate last

An enterprise will not run a decision system it cannot administer, isolate, back up and audit.

KEY POINTS

- The acting identity comes from the verified token and from nowhere else; no interface, API or agent may assert who it is acting as
- Tenancy is isolated at the data layer and tested adversarially; residency is chosen per tenant and enforced per record, including where models process decision content
- Metering counts governed decisions, not seats; value is baselined before the pilot, and none is generated retrospectively
- Onboarding is per decision class through Connect, Observe, Govern, Learn and Compound; maturity is assessed per function, never as one organisation-wide score

What it is

None of the items in this cluster delivers a benefit on its own, and every benefit depends on them. Platform and service is the tenancy, identity, residency, metering, value measurement, onboarding, maturity, personas and backup that let an enterprise run Red Bear: Enterprise Decision Intelligence as a service with a budget rather than a programme with an end date.

It is also where the adoption journey lives as a mechanism rather than a slide. A class's phase is a property the platform acts on: a class in Observe is captured and evaluated and never intercepted, whatever its policies say, until a named person moves it to Govern. Governance programmes usually fail at the first step, by asking people to comply with something nobody has measured; the phases exist so that the evidence comes first.

The pain point

Decision data is the most sensitive data there is. The register holds what the organisation decided and why, across every function, and a tenancy or residency weakness is the end of the relationship. If the platform believes whoever the client says it is, every authority control is decoration and the audit trail proves nothing.

Governance is seen as a cost, because the delay is felt daily and the benefit is invisible, so the programme is cut at the first budget round. Integration effort blocks the pilot, so evaluation never starts. And nobody agrees what was bought, because pricing on seats does not match value per governed decision.

How it works

Tenancy, identity and residency. Each customer has its own tenancy, configuration and administration, with isolation enforced below the service layer and tested adversarially; every configuration the vendor can perform, a customer administrator can perform. Sign-in is through the customer's directory with role-based access at decision-class granularity, and the acting user is always taken from the verified token. Residency is a tenant choice, region, on-premises or hybrid, with customer-managed keys as an option; the ingest tag enforces residency on every read and aggregation, and where a model processes decision content it does so inside the tenant's region or the feature is withheld and the tenant told.

Five personas, as grants. Decision Maker, Decision Approver, Auditor, System Admin and Policy Maker are grants against a verified identity, unioned, every action attributed to the person with the grant named. The separations are structural: a System Admin reaches no decision content by any path; an Auditor is read-only and flagged when party to a decision; an Approver's oversight follows the reporting line and excludes their own decisions; a Policy Maker never decides and never activates a widening alone.

Metering and value. Usage is measured by governed decision, one that reached the register and was evaluated against policy; harvested, discarded, unconfirmed and retrospective records are not metered. Value measurement baselines decision speed, bottleneck cost and override rate on the customer's existing process before anything is switched on, then tracks them; a pilot cannot start without a baseline, none is generated retrospectively, and the dashboard shows the baseline permanently beside the current figure.

Onboarding and maturity. Onboarding takes a customer through Connect, Observe, Govern, Learn and Compound one decision class at a time, and the tracker states what is required to advance, not a percentage. Observe to Govern and Learn to Compound are leadership decisions rather than 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; writing either into a plan as a date is the most common way this goes

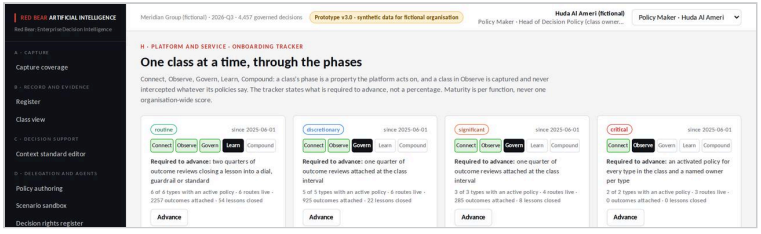
wrong. Maturity is separate from phase: each function is placed on five levels, Unaware, Aware, Governed, Trusted, Learning, with the single change that would move it up; a contact centre may sit at Trusted while capital projects sit at Aware, and an organisation-wide score conceals exactly the thing the assessment is for.

Backup, restore and export. Scheduled backup encrypted under the tenant's key inside its residency boundary, honouring legal hold, restore-tested with the result recorded; export that respects residency, retention and hold. A restore gives the administrator who ran it no view of decision content. Backup is not a reading path.

What you get

Capability	What it does	Requirements
Tenant administration and isolation	Per-customer tenancy and configuration; isolation below the service layer, tested adversarially; handed over, not held	FR-PLT-01
Identity, single sign-on, roles and personas	Directory sign-in; role-based access at class granularity; the acting user from the verified token only; five grants, unioned and attributed	FR-PLT-02, FR-PLT-08
Residency-aware deployment	Region, on-premises or hybrid; customer-managed keys; the ingest tag enforced on every read; in-region model processing or the feature withheld	FR-PLT-03
Metering and billing	Governed decisions counted, by class and period, as billed	FR-PLT-04
Value measurement	Speed, bottleneck cost and override rate baselined before the pilot and tracked beside it	FR-PLT-05
Adoption, onboarding and maturity	Phase per class, acted on by the platform; what is required to advance; maturity per function on five levels	FR-PLT-06, FR-PLT-07
Backup, restore and export	Encrypted under the tenant's key, inside the boundary, hold-aware, restore-tested; no decision content to the administrator	FR-PLT-09

Screens



S-52 Onboarding tracker: The onboarding tracker: each decision class and its phase, what is required to advance, and maturity per function. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

This cluster is where the business case is made stageable rather than taken on trust in one piece. **Value measurement** is the mechanism: the baseline is taken on the customer's existing process before anything changes, so every later figure has something honest to stand beside, and no benefit waits for the final phase before it can be read. A customer may reasonably stop at Observe and still have got what they came for: a register of a class of decisions with the reasoning attached, evidence of how often policy is met, and a measurement of where decisions wait. That is a durable asset; it does not have to be rebuilt if the next phase is bought two years later. The test of a phased approach is whether the earlier phases stand up when the later ones are cancelled. **Provable accountability** rests here too: residency compliance, the verified acting identity, and a restore that never becomes a reading path.

Platform engines

Shared services rather than an engine: Identity and tenancy (RBP FR-ID-01 to RBP FR-ID-04), Hosting (RBP FR-HO-01, RBP FR-HO-02), Audit (RBP FR-AU-03), the AI gateway's residency rule (RBP FR-AI-04) and Presence's rule that administrators see no content (RBP FR-PR-05). Specified in the platform FSD.