

RED BEAR ARTIFICIAL INTELLIGENCE
ABU DHABI, UNITED ARAB EMIRATES

The Customer Journey

Connect, Observe, Govern, Learn, Compound: adopting decision governance one decision class at a time

No large organisation should switch on decision governance across its estate in one go, and none that tried has enjoyed it. This paper sets out the five phases in which Red Bear: Enterprise Decision Intelligence is adopted, each independently worth stopping at, the two transitions that need a genuine decision rather than a project plan, and what each phase produces that the next one needs.

Empower. Trust. Verify.

White paper EDI-WP-21 · Red Bear white paper series · Version 2.0 DRAFT
26 September 2026
Commercial in confidence

Summary

Decision governance is not a system you install. It is a change in how an organisation decides, and changes of that kind are earned in sequence rather than announced. The reasons are practical. Connectors into production systems need proving. Policy written before anyone has looked at real decision data is always wrong in one direction or the other. Interposing a control in a live workflow on day one converts a governance programme into an operational incident. And the people whose decisions are being recorded need to see what the record looks like before they will believe what it is for.

Red Bear: Enterprise Decision Intelligence (formerly GuardrAIl) is therefore adopted in five phases, one decision class at a time, each of which produces something worth having on its own. An organisation may stop at any of them, keep everything it has built, and resume later. Nothing in a later phase requires re-doing an earlier one. **Connect**: integrations, identity and the policy scaffolding; nothing captured yet, the plumbing proved. **Observe**: connectors read decision data and policy is evaluated but advisory; nothing is blocked, nothing is approved, and the organisation sees, often for the first time, what it actually decides and how long it takes. **Govern**: guardrails enforced on the agreed classes, dials set, escalation and reach live; the phase that changes how work happens. **Learn**: outcomes attached, override telemetry reviewed, policies and dials adjusted on evidence. **Compound**: more classes, more actors, agents taking authority class by class, and, opt-in, aggregate patterns shared across participating organisations while no customer's data leaves its boundary.

Two of the four transitions need a genuine decision rather than a project plan. **Observe to Govern** is the first time a decision waits because the platform says so, and the first time a host application is asked to hold a transaction: a threshold the security function and the operating business both care about. **Learn to Compound** is the first time authority passes to something that is not a person, and the first time anything derived from the organisation's decisions is generalised beyond it; the second is opt-in, must remain so, and an organisation that does not opt in receives nothing from what others contribute. That is not a penalty. It is what makes the arrangement reciprocal rather than extractive.

The sequence is also how the platform earns the right to the later phases. An organisation that has watched Observe run for a quarter knows what its decision volumes are, whose approvals are the bottleneck, and where policy would have fired. It arrives at Govern with evidence instead of assumptions, which is the difference between a control that fits and one that has to be unpicked.

WHO SHOULD READ THIS

- Chief information and chief technology officers
- Chief operating officers and programme sponsors
- Chief risk and chief information security officers who will be asked to grant access
- Enterprise architects planning the integration

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Describe the five phases and what "done" looks like at each
- Name the two thresholds that need a signature, and what to settle before crossing them
- Explain why Observe must run for a full business cycle
- Say what leaves the tenant at Compound, and what never does

1 What it is

Five phases, applied per decision class rather than per organisation. A class's phase is a property the platform acts on: a class in Observe is captured and never intercepted whatever its policies say (FR-PLT-06). The five adoption phases are used for build sequencing and are explicitly not a maturity statement; maturity is measured separately, per function, on five levels, Unaware, Aware, Governed, Trusted, Learning, and never as one organisation-wide score (FR-PLT-07; Oliver, *Empower, Trust, Verify*, Book 1, chapter 18).

Phase	What the organisation gets	What "done" looks like
Connect	Proved integrations, identity federation, the class catalogue and a first guardrail set drafted but not enforced, the deployment topology and residency position settled	Connectors authenticate and stay connected; an actor in the ERP and an actor in the CRM resolve to one person; a policy exists on paper
Observe	A populated register and a baseline: volumes, cycle times, who approves what, where decisions stall, how often policy would have fired	Coverage of the in-scope class is stable and the baseline is agreed with the sponsor
Govern	Guardrails enforced in path on the agreed class, dials set, approvals routed, held and recorded, escalation and reach live	Alert volume at a sustainable steady state and approvers not bypassing the control
Learn	Outcomes attached at the class's interval; the quality quadrant; override telemetry; proposed policy and dial changes evidenced against history	Outcome closure high enough that the findings are worth acting on; a finding closes only when it has changed a dial, a guardrail or a standard
Compound	More classes and actors; agents promoted class by class from shadow mode; opt-in access to aggregate patterns contributed by participating organisations	Opt-in agreed, contribution boundaries configured and verified; agents governed through the gateway, not observed after the fact

Access posture is the dimension a security function reads first. Connectors are read-only under least privilege, granted by the customer administrator, and their permissions and health are visible in the connector administration (RBP FR-CN-01, RBP FR-CN-05). Connect requires only enough access to establish connectivity. Observe reads. Govern is the phase that intervenes, and the intervention takes one of two forms, stated per connector rather than glossed. Where the host application offers a pre-commit extension point, or where the application or agent calls the gateway before acting, the intervention is preventive: the product returns allow, hold or route, and the call itself is the capture event (FR-DEL-09, FR-DEL-05). Where the host offers neither, the intervention is detective: the decision is recorded and flagged but not stopped. Whether a connector is additionally permitted to hold and release a transaction in the source system, and to write an outcome back against the decision's identifier, is a permission scoped to that specific operation and decided at the Observe-to-Govern threshold; it is not general write access to the object, and it is an open architecture decision recorded in the product's ledger. Learn returns to reading: it correlates decisions already captured with outcomes and measures. Compound operates on the register, not on source systems, and is opt-in.

Why big-bang programmes fail. Four failure modes, all common, none about the software. Policy written before anyone has seen the data: every organisation's first policy is either too loose, so decisions that should have been caught pass through, or too aggressive, so approvers drown and begin approving without reading. Interposing a control in a live workflow on day one: from the operating business's point of view a new single point of failure in a process that previously worked. Asking for intervention rights before trust

exists: a security function asked on day one to let a platform nobody has seen running hold transactions in the ERP will say no, and will be right to; the same request after two quarters of read-only operation, with an audit trail of exactly what the connector touched, is a different conversation. Recording people's decisions before explaining why: if the first thing people experience is a control that blocks them, the platform is an obstacle; if it is a record that shows them their own decisions and where they are waiting on others, it is something else entirely.

Key points

- Five phases per decision class, each worth stopping at; a class in Observe is never intercepted
- Connectors are read-only; Govern intervenes through the gateway or a pre-commit point, and holding a transaction is a scoped decision at the threshold
- Each phase produces the evidence the next one needs

2 Why it matters

The phases are not a delivery convenience. Each produces the evidence the next one needs: Observe produces the decision volumes that make Govern's guardrails correct, Govern produces the approval and override record that makes Learn meaningful, and Learn produces the outcome history without which Compound has nothing to generalise from. Skipping a phase does not accelerate the journey; it removes the evidence that would have made the next step defensible.

Connect is worth less than the phases that follow, and should not be sold as more. Its value is that it converts every subsequent phase from a technical risk into a content question. An organisation that stops after Connect has spent money and has little to show, which is why this phase should be short.

Observe is the lowest-risk phase and, for many organisations, the most surprising. It is common for a sponsor to discover that the volume of decisions in the in-scope class is several times what they assumed, that the approval bottleneck is not where they believed it was, or that a meaningful share of decisions are being taken by people whose authority to take them nobody had written down. An organisation that stops permanently at Observe has a populated register, an audit trail it did not have before, visibility of stalled decisions, a measured baseline of decision velocity, and evidence of how often policy is actually met before anyone is asked to comply. Several of the outcomes in EDI-WP-02 are available at this phase alone. The recommendation is to run Observe long enough to see a full business cycle in the in-scope class, typically a quarter, sometimes longer. The policy written after one month of data is materially worse than the policy written after three, and the cost of a badly calibrated policy is paid by every approver, every day.

Govern is where the platform stops describing the organisation and starts constraining it. More changes for people than in any other phase: approvers acquire a queue, requesters encounter a hold where previously there was none, and policy that was theoretical becomes the thing that stops somebody's Friday afternoon. This is why the calibration work from Observe matters and why the sandbox exists: a proposed guardrail is replayed against the decisions already captured, showing exactly how many would have been held and which. A sponsor who can say *this will hold about forty decisions a month, and here are last quarter's forty* is in a completely different position from one asking an operating business to trust a threshold. This is where prevented loss appears and where compliance value becomes demonstrable. Govern should begin on one class in one system, not everywhere the connectors reach; the first enforcement is where the programme's credibility is won or lost.

Learn is where governance spend turns into operating improvement, and where the argument shifts from "this control protects us" to "this control is making us better at deciding". Two feedback loops open. Outcomes: decisions carry an expected outcome and a review date, and the comparison accumulates into the quality quadrant by class. Policy: override telemetry shows which rules are being worked around and how often; a rule overridden in two of every five cases is miscalibrated rather than flouted, and the product proposes how it could change. Proposes. Policy evolution is surfaced with the evidence attached and applied by people. The product does not tune the rules that govern an organisation on its own, and no capability that did so would be acceptable to the audit function that signs off the control. Learn is the slowest phase to pay back: expect the first genuinely useful findings after a few hundred comparable decisions, not a few dozen.

Compound is where the organisation extends what works: more classes, more actors, and AI agents taking authority class by class, promoted by a person from shadow mode once the evidence supports it, never globally and never automatically. It is also the only phase in which anything derived from a customer's decisions is generalised beyond it. The objective is to capture the patterns and lessons in customers' decisions, not the decisions themselves. What is useful across organisations is not that a particular company approved a particular discount; it is that a class of decision, made under a class of circumstance, tends to produce a class of outcome. The first is proprietary business data. The second is a lesson.

Stays inside the tenant, always	May leave, once opted in
Decision records, rationale and evidence	Aggregate patterns about decision types and outcomes, above a minimum cohort size
Counterparties, customers, suppliers, amounts	Statistical representations from which individual records cannot be reconstructed
Identities of deciders, approvers and agents	—
Policy content, thresholds and organisation structure	—

Participation is entirely optional, and an organisation that does not participate loses nothing it already had. What it does not get is the benefit of what other participants have contributed. A shared body of knowledge that non-contributors could draw on would be extractive in one direction and free-riding in the other, and would not survive its first conversation with a participating customer's legal team. Contribution and benefit are the same switch.

3 The challenges

Challenge	Why it is hard	What goes wrong without a method
Proving the plumbing while the stakes are zero	Connectors, identity resolution and residency must work before any decision depends on them	The first real decision data arrives into something that does not yet work
Writing policy against real volumes	The first policy is always wrong; only observed volumes correct it	Approvers drown, or decisions that should have been caught pass through; either destroys credibility

Challenge	Why it is hard	What goes wrong without a method
Crossing the intervention threshold	A passive system becomes one that can hold a transaction; security and the operating business both have a legitimate veto	Surprised in month eight, the security function is an obstacle, and right to be
Deciding the failure behaviour	If the product is unreachable, does the transaction proceed or stop? Fail-open preserves operations and creates a gap; fail-closed preserves the control and can halt the business	Nobody senior enough has owned the consequence either way
Waiting for Learn to mean something	Outcomes need volume and elapsed time before the correlations are worth acting on	Findings acted on after a few dozen decisions are noise with a chart
Passing authority to an agent	Governance of a person and governance of an agent share a register but not a starting point	An agent that acts and reports afterwards is being observed, not governed
Crossing the data threshold	The first time anything derived from decisions is generalised outside the boundary	Treated as a configuration change instead of a formal decision with legal and data-protection review

4 How Red Bear: Enterprise Decision Intelligence helps

Phase	What Enterprise Decision Intelligence does	Requirements
Connect	Productivity, messaging, enterprise-application and custom connectors configured, authorised, tested against a sample payload and switched on by the customer administrator without a vendor release; directory sign-in with the acting user from the verified token; residency-aware deployment with the ingest tag enforcing residency on every read; per-tenant isolation; the class catalogue and a first guardrail set authored but not activated	FR-CAP-01, FR-CAP-03, FR-CAP-04, FR-PLT-02, FR-PLT-03, FR-PLT-01, FR-REC-01, FR-DEL-19
Observe	Capture from every route into one register; harvested history seeding a baseline, marked and excluded from scoring; a class in Observe captured and never intercepted; the value baseline of decision speed, bottleneck cost and override rate measured before anything is switched on; latency and stall analysis on the observed class	FR-CAP-07, FR-CAP-09, FR-PLT-06, FR-PLT-05, FR-RTE-05, FR-OUT-11
Govern	Guardrails enforced per type; dials set per decider with the projection shown before a move commits; the sandbox replaying the proposed setting against captured history; interception through the gateway or a pre-commit point, with a blocking breach overridden only with a recorded reason; routing, presence, cover, the notification ladder and the escalation ladder live; the fast lane for senior decisions	FR-DEL-19, FR-DEL-02, FR-DEL-16, FR-DEL-05, FR-DEL-09, FR-RTE-01, FR-RCH-01, FR-RCH-04, FR-RCH-05, FR-RTE-04
Learn	Outcome tracking at the class's interval; decision quality scored apart from outcome; the quality quadrant; override telemetry as a question about the policy; the earned trust ledger recommending dial settings with its evidence and no path to apply them; policy optimisation proposing changes to the owner; a learning finding closing only when it changed a dial, a guardrail or a standard	FR-OUT-01, FR-OUT-05, FR-OUT-08, FR-DEL-18, FR-DEL-03, FR-DEL-17

Phase	What Enterprise Decision Intelligence does	Requirements
Compound	Further classes moved through the same phases; agents registered with an owner, a purpose, one dial and a kill switch, run in shadow mode and promoted class by class by a person; collective intelligence opt-in, sharing aggregate patterns above a minimum cohort size with no customer data leaving its residency boundary, and a customer that has not opted in receiving nothing; the connector store for certified partner connectors	FR-DEL-10, FR-DEL-11, FR-DEL-13, FR-OUT-06, FR-CAP-06
Throughout	The onboarding tracker taking one class at a time through the phases; maturity assessed per function, never as one score; metering by governed decision only, so Observe-phase, harvested and unconfirmed records are not billed; regulatory mapping with producible evidence per control	FR-PLT-06, FR-PLT-07, FR-PLT-04, FR-REC-11

Design rule

Every configuration the vendor can perform, a customer administrator can perform, and a System Admin reaches no decision content by any path. The security function grants access in proportion to what it has seen, not to what it has been promised.

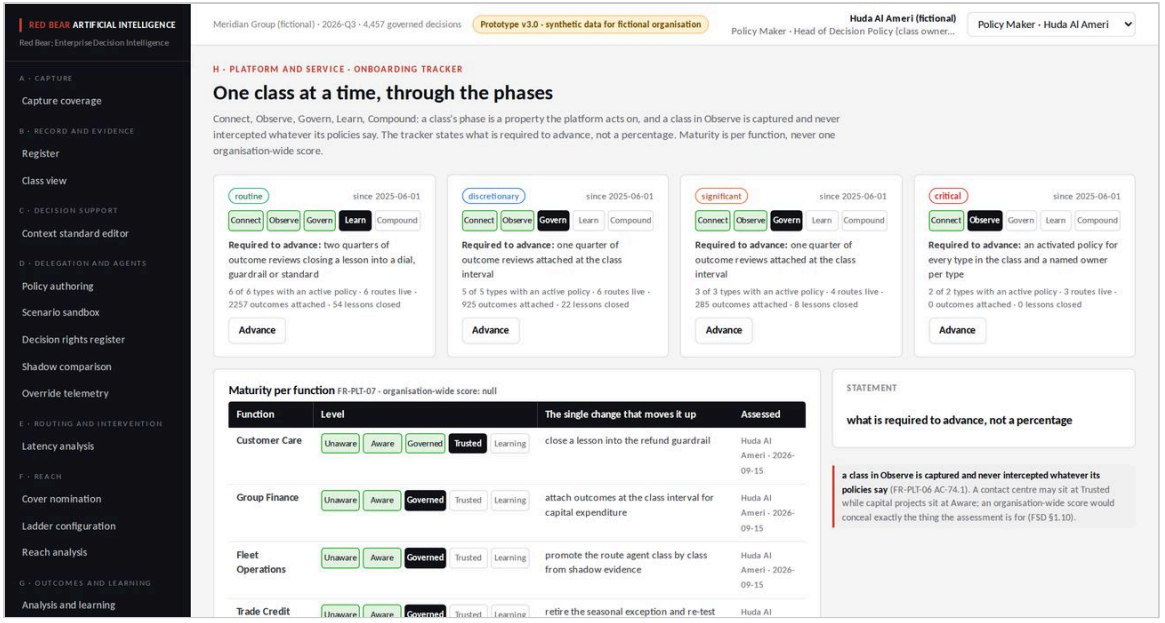
5 How it is used

Step	Who	What happens
1 Choose one class in one system	Sponsor, class owner	High-volume, consequential, currently invisible; name the accountable owner
2 Connect	System Admin, security	Connectors authorised read-only under least privilege; identity federated; residency and hosting settled; the two thresholds named in month one, with the evidence that will be available at each
3 Observe for a business cycle	Class owner, sponsor	The register fills; the baseline is agreed; policy evaluated advisory; latency, stalls and would-have-fired counts read; typically a quarter
4 Cross the first threshold	Sponsor, CISO, operational owner	Scope any hold-and-release permission to the specific operation; decide fail-open or fail-closed per connector and have the owner of the operational consequence sign it; agree the revocation path and what happens to in-flight held decisions
5 Govern narrowly	Policy Maker, Decision Approver	The sandbox shows what the guardrail would have held; dials set; the first enforcement on one class; expect the first month to require tuning
6 Learn	Class owner	Outcomes attached at the interval; the quadrant and override telemetry reviewed; proposed changes evidenced and activated by people
7 Cross the second threshold, and compound	Sponsor, legal, data protection, chief AI officer	Establish what is contributed, in specific terms, and confirm records cannot be reconstructed; confirm the residency position for contributed patterns; agree withdrawal; promote the first agent from shadow mode, class by class

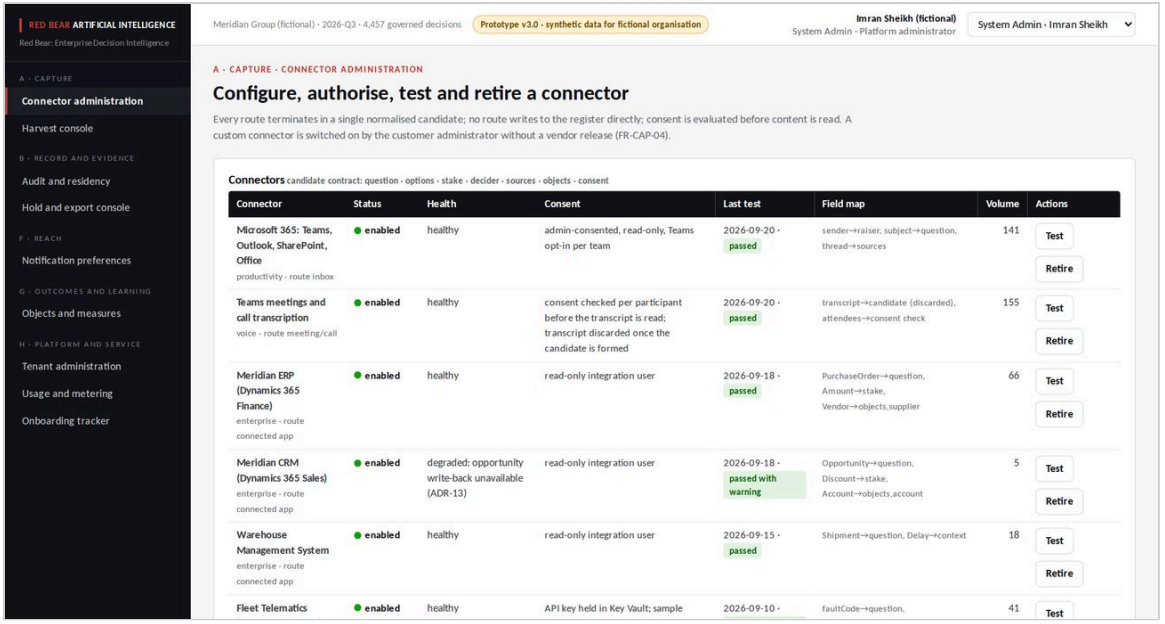
Both thresholds are easier to cross if named at the start rather than discovered at the point of crossing. A security function told in month one that hold-and-release will be requested in month eight, with the evidence that will be available by then, is a collaborator. The same function surprised in month eight is an obstacle,

and will be right to be. 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.

6 The user experience



S-52 Onboarding tracker: S-52 Onboarding tracker. Each decision class on its own phase, Connect to Compound, with what "done" looks like for the current phase, the two thresholds marked, and the baseline agreed at Observe. Prototype render on synthetic data for fictional entities.



S-05 Connector administration: S-05 Connector administration. Each connector's permissions, health and what it read; preventive or detective per connector; fail-open or fail-closed recorded with who signed it; no decision content anywhere on the page. Prototype render on synthetic data for fictional entities.

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

Inputs by phase

Input	Source	Check
Identity	The corporate directory through single sign-on	The acting user from the verified token only; grants are unioned and attributed to the person (FR-PLT-02, FR-PLT-08, RBP FR-ID-01)
Decisions	Productivity, messaging, enterprise-application, custom and voice routes; the portal; the agent gateway	Read-only, admin-consented; one candidate contract; consent evaluated before content is read; nothing enters the register on a machine's say-so alone (RBP FR-CN-01, FR-CAP-07, FR-REC-04, FR-CAP-08)
Historic decisions	Legacy harvesting, once per engagement	Marked harvested; excluded from quality scoring, value baselines and metering (FR-CAP-09)
The baseline	Measured on the existing process before go-live	Never generated retrospectively; a pilot cannot start without one (FR-PLT-05)
Policy and dials	Authored by the Policy Maker; set by leaders	A widening activated by a second person; every evaluation records its version (FR-DEL-19, FR-DEL-14)
Outcomes and measures	Connected systems against the decision identifier; the owner's review; measures with source and read time	An outcome the product guessed is not evidence; a measure without provenance is not admitted (FR-OUT-01, FR-OUT-09)
Contributed patterns	Aggregates above a minimum cohort size, opt-in	No customer data leaves its residency boundary; de-identified; disclosure stated (FR-OUT-06, RBP FR-PA-01, RBP FR-PA-03, RBP FR-PA-05)

How it is checked

- **Connectors are read-only and consented.** Least privilege, granted by the customer administrator, with health and what was read visible; any hold-and-release permission is scoped to the operation and decided at the threshold (RBP FR-CN-01, RBP FR-CN-05).
- **Residency is enforced at ingest.** Each record is tagged with its residency boundary and retention class at ingest and no later component re-derives them; model processing of decision content happens inside the tenant's region or the feature is withheld and the tenant told (FR-REC-05, FR-PLT-03, RBP FR-AI-04).
- **The AI proposes; a named person decides.** A detected decision is confirmed by its owner; a policy change is activated by a person; a dial is moved by a person; an agent is promoted by a person; a lesson is never attributed to a person (FR-CAP-08, FR-DEL-19, FR-DEL-02, FR-DEL-11, RBP FR-DL-04).
- **A class in Observe is never intercepted.** Whatever its policies say (FR-PLT-06).
- **Isolation is tested adversarially.** Per-tenant isolation below the service layer; backup under the tenant's key inside its boundary; a restore gives the administrator no view of decision content (FR-PLT-01, FR-PLT-09).
- **The record is defensible.** Append-only vault; audit replay against the policy of the day; retention, legal hold and reproducible export; controls mapped to the EU AI Act's human-oversight requirement, ISO/IEC

42001 and the NIST AI RMF with producible evidence (FR-REC-08, FR-REC-09, FR-REC-10, FR-REC-11).

8 Benefits realised

All measures are calculated from the organisation's own register against the baseline established before anything was switched on (FR-PLT-05). The sequencing is itself a benefit: Observe produces the baseline without which none of the later measures can be interpreted, and an organisation that skips it has no defensible way to claim any improvement it subsequently sees.

Available from	Benefit	Measured by
Observe	The organisation sees what it decides	Decision coverage; volumes by class; average age of open decisions; count past review date; approval cycle time and time lost to unavailable approvers; policy conformance before anyone was asked to comply
Govern	Prevented loss and demonstrable compliance	Value of decisions held, rejected or amended at Authorise; policy breach rate; override frequency per policy; share settled within the decider's own authority
Learn	Governance turns into better deciding	Outcome closure rate; the quality quadrant by class; reversal rate within a defined window; assumption accuracy over time; findings that changed a dial, a guardrail or a standard
Compound	Authority extends safely	Autonomy used against autonomy granted, per agent per class; classes moved through the phases; the organisation's decision types against contributed patterns, where opted in

9 Risks managed

Risk	Control in Red Bear: Enterprise Decision Intelligence
The phases take longer than a procurement cycle expects	Observe runs for a business cycle; Learn needs volume; a programme sold as delivering all five phases inside a year is either scoped narrowly or going to disappoint someone; narrow scope is the better answer
Stopping early is treated as failure	An organisation that runs Observe indefinitely has a register, an audit trail and a velocity measure, a real result; metering by governed decision means an Observe-phase class is captured but not billed as governed (FR-PLT-04)
Govern fails publicly	One class in one system; calibrated in the sandbox against real data; failure behaviour decided and signed per connector; the first month expected to need tuning (FR-DEL-16)
A class intercepted before it was meant to be	A class's phase is a property the platform acts on; a class in Observe is never intercepted (FR-PLT-06)
Intervention rights granted on promise	Connectors read-only by default; the hold-and-release question deferred to the threshold with months of evidence of what the connector read (RBP FR-CN-01)
An agent governed by hope	Gateway before the act; one agent, one purpose, one dial, ceiling of 9; shadow mode; promotion by a person; kill switch exercisable by the owner without a release (FR-DEL-09, FR-DEL-13, FR-DEL-11, FR-DEL-10)

Risk	Control in Red Bear: Enterprise Decision Intelligence
Compound has no value until enough organisations participate	Early participants contribute ahead of the return, and are told so plainly rather than sold a network effect that does not yet exist
Data leaves the boundary	Aggregates above a minimum cohort size only; no customer data leaves its residency boundary; contribution and benefit are the same switch; withdrawal agreed before opting in (FR-OUT-06, RBP FR-PA-01)

Start with one decision class in one system, watch it for a quarter, and let what you find decide what you govern. Every organisation that has tried the reverse has discovered the same thing: the policy was wrong, and the only way to know was to look first.

Sources

Every claim attributed to a source names it in the text as *author, work, section*; the full references are below. Statements without a citation are Red Bear's own position, developed for Red Bear: Enterprise Decision Intelligence. Quotations are limited to fifteen words.

Cited as	Source	Authority
Book 1	D. Oliver, *Empower, Trust, Verify*, Book 1 (2026), chapters 17 and 18: agents; maturity levels and transition phases	Red Bear
Regulatory references	EU Artificial Intelligence Act (human-oversight requirement); ISO/IEC 42001; NIST AI Risk Management Framework; the targets of FR-REC-11	Regulatory and industry framework
Red Bear	EDI-FSD-01 v3.0 draft; EDI-SAD-01 v2.0 draft (ADR-13, outcome write-back through connectors); EDI-DS-00 suite sheet; Red Bear Platform Functional Specification RBP-FSD-01	Red Bear

Glossary

Term	Meaning in this paper
Phase	One of Connect, Observe, Govern, Learn, Compound, applied per decision class; a build-sequencing property, not a maturity statement
Maturity	Unaware, Aware, Governed, Trusted, Learning, assessed per function and never as one organisation-wide score
Preventive intervention	The application or agent calls the gateway, or offers a pre-commit point, before the transaction takes effect; the product returns allow, hold or route
Detective intervention	The decision is recorded and flagged but not stopped, where the host offers no point to intervene
Fail-open, fail-closed	What happens to a transaction when the product is unreachable: it proceeds, or it stops; configured and signed per connector
Baseline	Decision speed, bottleneck cost and override rate measured on the existing process before anything is switched on
Scenario sandbox	Replaying a proposed guardrail or dial against captured history through the live routing function, applying nothing
Shadow mode	An agent's decisions recorded without being executed and compared with human decisions, before any promotion
Governed decision	One that reached the register and was evaluated against policy; the unit of metering
Aggregate pattern	A statistical representation of how a class of decision tends to behave, above a minimum cohort size, from which individual records cannot be reconstructed
Outcome closure	The share of decisions reaching their review date with an actual outcome recorded

The series

This paper is one of ten. 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
EDI-WP-01	Cartesian Decision Making	Learning what happened next
EDI-WP-02	The Cost of Organisational Indecision	The decision that never gets made
EDI-WP-03	Trust and Verify	Visibility, not control
EDI-WP-04	Delegation Without Losing Governance	The boundary and the room
EDI-WP-05	Building a Decision-Making Culture	Judge the decision, not the result
EDI-WP-06	Decision Impact Analysis	What a decision dataset can and cannot tell you
EDI-WP-07	The Decision-Aware Enterprise	Decisions in flight
EDI-WP-08	Decision Making Anywhere	The channel is part of the governance
EDI-WP-09	When Micromanagement Is the Right Management	Oversight is a dial
EDI-WP-21	The Customer Journey	Connect, Observe, Govern, Learn, Compound (this paper)

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