

**RED BEAR ARTIFICIAL INTELLIGENCE**

ABU DHABI, UNITED ARAB EMIRATES

# Delegation Without Losing Governance

## The boundary and the room: decision governance that separates a rule about risk from a judgement about trust

Delegation fails for structural reasons rather than personal ones, and an organisation that can see what was decided under a grant of authority can afford to give away far more of it. This paper names the four structural faults, the loop they form, the canon the model rests on, and how Red Bear: Enterprise Decision Intelligence keeps the guardrail and the dial as two separate settings.

**Empower. Trust. Verify.**

White paper EDI-WP-04 · Red Bear white paper series · Version 2.0

26 September 2026

## Summary

**The bottleneck is not the executive. It is what the executive cannot see.** Almost every organisation of any size says it wants decisions made closer to the work, and almost none behaves as though that were true. There is a delegation of authority schedule. There are managers with "director" in their title. And there is still a queue of decisions outside one or two offices, moving at the speed of the people in them. The gap between the schedule and the queue is not hypocrisy, and it is not a failure of nerve. It is a structural fault with four identifiable parts.

Delegation in most organisations is granted as a property of a job title rather than of a decision type. It is written once, at a restructure or a policy refresh, and revisited annually if at all. There is no practical way to see what has actually been decided under it. And the person who granted it stays accountable for the outcomes. Those four facts form a loop: being answerable for decisions you cannot observe is uncomfortable, discomfort is resolved by pulling the decision back, and every pull narrows the grant a little further.

The central move is narrower than fixing all four, and matters more. A delegation contains two things that should never have been written as one: the hard boundary, which is a statement about risk, and the amount of room inside that boundary, which is a judgement about a particular decider's demonstrated competence. Most delegation matrices conflate them, which is why widening a good team's latitude means reopening a rule that was written about risk, and why it almost never happens. In the Red Bear framework the boundary is the **guardrail**, set per decision type, and the room is the **Trust Dial**, set per decider. Separate them, and delegation becomes something you adjust in an afternoon rather than at the next restructure.

None of this is new thinking. It is Tannenbaum and Schmidt in 1958, Oncken and Wass in 1974, Marquet, and the US Army's doctrine of commander's intent, all pointing at the same shape. What has been missing is not the model. It is the instrumentation: a record of what was decided under the grant, good enough that the accountable owner can stop watching the decisions and start watching the pattern. Red Bear: Enterprise Decision Intelligence (formerly GuardrAII) is one implementation of that instrumentation, and it appears late in this paper because the argument stands without it.

**WHO SHOULD READ THIS**

- Chief executives and chief operating officers who have become a queue
- Chief financial and chief risk officers who own the delegation schedule
- Heads of functions with high-volume decision classes: procurement, pricing, credit, change
- Directors accountable for decisions made by AI agents

**WHAT YOU WILL BE ABLE TO DO AFTERWARDS**

- Name the four structural faults and explain why the bottleneck is rational
- Separate the seven jobs a delegation schedule is doing
- Write the boundary and the room as two settings, and say who owns each
- Describe how an accountable owner stays on the record of every decision, including those made by software

## 1 What it is

A delegation of authority schedule is usually one artefact doing seven jobs. Separating them is not bureaucratic tidiness; it is what makes any of them adjustable without reopening the rest.

| Part | What it actually governs | What goes wrong when it is bundled |
|---|---|---|
| **Delegated authority** | Whether this actor may settle this decision type at all, and with how much initiative | Expressed as a money figure, so a low-value irreversible commitment is treated as trivial |
| **Decision rights** | Who recommends, decides, owns, executes and reviews | Collapsed into "approver", so input and decision are confused and everyone consulted believes they hold a veto |
| **Financial authority** | The monetary boundary, a real and legitimate control | Becomes a proxy for every other kind of risk because it is the only quantity anyone wrote down |
| **Risk thresholds** | Exposure the organisation will not accept at this level whoever is deciding: safety, regulatory, reputational, concentration | Never stated separately, so they are enforced by escalating everything above a money line, which catches the wrong decisions |
| **Escalation rules** | What happens when a decision falls outside the room: who is brought in, how fast, with what context, and what proceeds meanwhile | Defined as "refer upward" with no clock, so escalation is indistinguishable from delay |
| **Temporary delegation** | Cover for absence, and standing authority for declared conditions such as a crisis or a closing window | Informal deputising that leaves no record, so nobody can say afterwards who held the authority |
| **Accountability** | Who answers for the outcome. It does not move | Assumed to transfer with the authority, so after an incident nobody can locate the owner |

Two rows deserve emphasis. Escalation without a clock is not governance, because a decision that waits indefinitely has been made, by default, in favour of doing nothing, by nobody. And accountability is the row that must stay fixed while every other row moves. The Institute of Directors' formulation is the load-bearing one: a board can delegate power and authority but never liability (Institute of Directors, guidance on the role of the board). Every grant of room must therefore name the person who remains on the record, and that name must survive into the record of each decision made under it, including decisions made by software.

**Risk is not the same as trust.** A delegation says two separate things. The first is a hard boundary: beyond this point the organisation will not go, whoever is asking. That is a statement about risk, owned by the board or a risk function, and it should change rarely and deliberately. The second is how much room there is inside that boundary for a given decider, right now. That is a judgement about demonstrated competence, owned by the accountable executive, and it should change often, on evidence, without ceremony. Conflate them, as nearly every matrix does by writing a single number per grade, and you get the pathology that defines the subject: widening a competent team's latitude requires amending a rule written about risk, which means a governance change, which means it does not happen; while narrowing after an incident is easy, because tightening a risk rule needs no justification. The asymmetry is built into the artefact. The organisation can only get more cautious.

The Red Bear framework writes the two as two objects. **Guardrails** belong to the decision type: nine fixed kinds, authority, policy, financial limits, risk appetite, compliance, segregation of duties, exceptions, human intervention and auditability. The **Trust Dial** belongs to the decider: one per person or agent, 0 to 10, in five bands, No Discretion, Recommend, Supervised, Delegated, Autonomous, each carrying a review rate. **Role scope** says which types a decider may touch at all. The effective range on any decision is the product of the three (Oliver, *Empower, Trust, Verify*, Book 1, chapters 5 and 8). This is Tannenbaum and Schmidt's continuum, published in 1958 and reissued with a retrospective commentary in 1973, whose points run from the manager announcing a decision to subordinates functioning within limits the superior defines (Tannenbaum and Schmidt, *How to Choose a Leadership Pattern*, Harvard Business Review).

The limits are the boundary. The position on the continuum is the room. They wrote it as two ideas. We have spent decades writing it as one.

Key points

- One schedule doing seven jobs; separate them
- The guardrail is a rule about risk, per decision type, owned by risk; the dial is a judgement about trust, per decider, owned by the accountable executive
- Accountability never moves, and is named on every decision made under the grant

## 2 Why it matters

**Four faults, and the loop they form.** The first fault is that authority is attached to a grade rather than to a decision type. A newly promoted director inherits a signature limit that says nothing about whether they have ever made a pricing decision, and an experienced category manager who has made four hundred of them is capped by a threshold with no opinion about the subject. Hersey and Blanchard's situational leadership model argued that the right amount of direction depends on readiness for the specific task, not general standing; its empirical support is thinner than its popularity, but the observation is hard to argue with: readiness is per task, authority is granted per person, and those two statements cannot both be right (Hersey and Blanchard, *Management of Organizational Behavior*, situational leadership). The second fault is that the grant is a document, written at moments and static between them, while conditions move continuously; a control updated annually against conditions that change weekly is not calibrated, merely old, and because amending it is a governance event, the rational thing is to work around it. The third fault is the absence of a record: ask what was decided last quarter under a given delegation, not what was spent, and there is no answer. Everything below the approval threshold leaves no trace, which is precisely the population the delegation was meant to cover. The fourth fault closes the loop: accountability does not move when authority does. The person who delegated is carrying a risk whose realisation they cannot see. No visibility, retained accountability, discomfort, withdrawal.

**Why the bottleneck is a rational response.** A chief operating officer is answerable to a board for operational risk. Some hundreds of decisions a week are made in their organisation that could produce a material loss or a regulatory finding. They can observe perhaps a dozen. When something goes wrong, the honest answer, "I delegated it and had no way of seeing it", does not survive a board meeting. Given that, personally reviewing more decisions is not irrational. It is the only lever that reliably reduces the exposure. It is an expensive lever, it scales badly, and it often degrades decision quality, because a reviewer with thirty seconds and no context adds latency without adding judgement. None of that changes the calculation for the individual. Which is the useful insight: if the bottleneck is a rational response to missing information, the remedy is to supply the information, not to demand a different attitude. Give the accountable owner a reliable view of what is being decided under their authority, and the case for personally reviewing it collapses on its own.

**What it costs, and what cannot honestly be counted.** What can be measured is the time decisions spend waiting, the proportion of escalations that change the outcome, and the effort consumed by approval that adds no judgement. Bain's four dimensions, quality, speed, yield and effort, are a useful frame precisely because they refuse to collapse into a single score (Blenko, Mankins and Rogers, *Decide & Deliver*, chapter 1). What cannot honestly be counted is the larger half: the decision never raised because the approval cycle was judged longer than the window, the concession not sought, the experiment not run. These are absences. One measurable consequence is worth stating carefully: Lovallo and Sibony's study of 1,048 business decisions found that the process by which a decision was made mattered more to the outcome than the quality of the analysis, by a factor of about six. It is a single study of self-reported

decisions, not a law of nature, but it points somewhere useful: if process dominates, then process is the thing to govern, and governing it does not require the governor to second-guess every decision (Lovallo and Sibony, *The case for behavioral strategy*, McKinsey Quarterly, 2010).

## 3 The challenges

Four remedies are commonly attempted, and each fails for a reason visible in advance. What they have in common is that they treat delegation as a single dial with a number on it. There are two dials, and they have been welded together.

| Challenge | Why it is hard | What goes wrong without a method |
|---|---|---|
| **Raising the thresholds** | Thresholds set in response to incidents ratchet in one direction; no incident is ever caused by a threshold being too tight | A quarter of speed, then an incident below the new line, then the limit comes back down further than it started |
| **Clarifying decision rights** | Naming who decides is valuable, and the frameworks are sound (Rogers and Blenko, *Who Has the D?*, Harvard Business Review, 2006), but clarity about *who* says nothing about *how much room* | A single decider for every class, and every one of those decisions still routed upward |
| **Empowerment as a cultural programme** | It addresses the version of the problem in which managers are the obstacle; since they mostly are not, it produces goodwill and no structural change | Failure attributed to insufficient cultural commitment, which produces more workshops |
| **Adding a governance forum** | A weekly board to clear stuck decisions moves the queue, and institutionalises the bottleneck by giving it a calendar entry | A two-day wait becomes a predictable six-day one, and the pressure to fix the delegation disappears |
| **Keeping the owner on the record** | Decisions are made by people in meetings, applications in workflows and agents at three in the morning, and the owner must survive into all of them | After an incident nobody can locate the owner, especially for a decision made by software |
| **Stopping escalation turning into delay** | Escalation needs a clock, a context and a default; informal deputising leaves no trace | A decision that waits has been decided by nobody |

## 4 How Red Bear: Enterprise Decision Intelligence helps

Where the canon lands. If room is a continuum rather than a gate, it needs a scale. Oncken and Wass's five degrees of initiative, from waiting to be told, through asking what to do, recommending then acting, acting but advising at once, to acting alone and reporting routinely, is the most durable published account of why senior people end up doing their subordinates' work and what to do about it (Oncken and Wass, *Management Time: Who's Got the Monkey?*, Harvard Business Review, 1974). Marquet contributes the single most useful phrase in the literature, *"I intend to..."*, which moves the initiative down while keeping the leader informed and able to intervene (Marquet, *Turn the Ship Around!*, chapter on intent). Mission command supplies commander's intent: a statement of purpose and end state clear enough that subordinates can act when the situation is one the limits did not anticipate (US Army, *ADP 6-0, Mission Command*). And reversibility, one-way and two-way doors, is a routing input that a money threshold cannot see (Bezos, *2015 letter to shareholders*).

The Red Bear framework resolves that canon into the five dial bands, each of which is a review rate, and into the intent notice at the edge of discretion: the decider declares intent, a clock starts, anyone with standing may stop it

with a recorded reason, and on expiry the decision proceeds on notice. Reversibility is one of the five properties that raise a decision's class, alongside stakes, novelty, regulatory exposure and frequency; time pressure is deliberately excluded.

| Challenge | What Enterprise Decision Intelligence does | Requirements |
|---|---|---|
| **The boundary and the room as two settings** | Guardrails composed from the nine fixed kinds per decision type in the policy builder; one Trust Dial per decider in five bands; role scope as a third object; the effective range is the product of the three | FR-DEL-19, FR-DEL-02, FR-DEL-06 |
| **Widening the room on evidence** | The earned trust ledger scores each actor's record and recommends a dial setting with its evidence, with no path to apply it; a move is one point at a time, by a person, with a reason, after the projection has been displayed; a widening is activated by a second person | FR-DEL-03, FR-DEL-02, FR-DEL-19 |
| **Seeing what would change before it changes** | The scenario sandbox replays a named historic period under a proposed guardrail or dial through the live routing function and reports the difference in level mix, decision count by outcome and human-touch share | FR-DEL-16 |
| **Purpose, not only limits** | Pre-authorised envelopes carry standing decisions agreed in advance as versioned policy, so a class already argued through is not argued again every time it recurs | FR-DEL-04 |
| **The owner on the record** | Every decision names the acting party and the accountable owner from the decision rights matrix; Own has exactly one filled cell, a named person; accountability never transfers to an agent; no actor may decide their own decision | FR-DEL-07, FR-DEL-01, RBP FR-DI-07 |
| **Decisions made by software** | An agent calls the gateway before acting and receives allow, hold or route; the call is the capture event whatever the answer; every agent has an accountable human owner, a purpose, one dial with a ceiling of 9, and a kill switch | FR-DEL-09, FR-DEL-10, FR-DEL-13 |
| **Watching the pattern, not the decisions** | Override telemetry presents approvals made over a breach as a question about the policy; outcomes are attached by class; quality is scored on process apart from outcome | FR-DEL-18, FR-OUT-01, FR-OUT-05 |
| **Escalation with a clock** | Referral carries the full record and does not reset the clock; a timeout default is a recorded decision with an owner; cover for absence routes to a named person with the same authority and both names on the record | FR-RTE-03, FR-RTE-02, FR-RTE-01, FR-RCH-03 |
| **Intent at the edge** | The intent notice runs with a clock and a stop action for anyone with standing; exactly three verbs, approve, reject, ask | FR-RTE-06 |
| **Keeping the schedule current** | Decision rights and authority kept current from HR and the directory rather than from a document accurate at the last restructure | FR-DEL-01 |

Design rule

Collapsing any two of the three objects, the guardrail, the dial and the scope, is the single most common way a governance product becomes a policy engine with a nicer control. The product refuses to hold a dial per decision type and refuses to hold a guardrail per person.

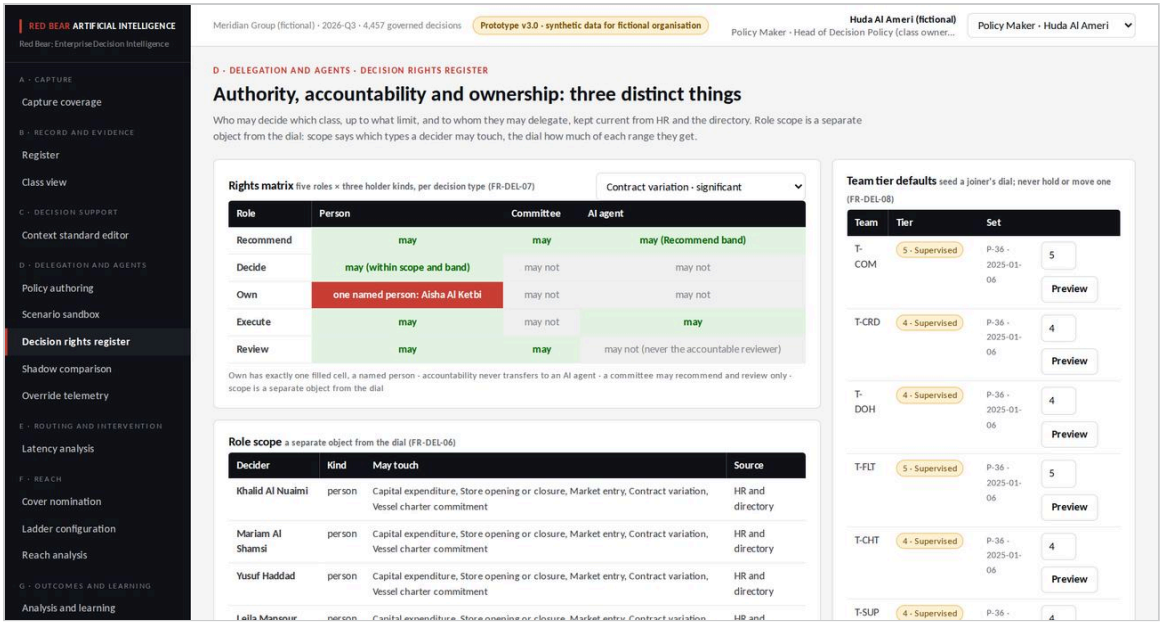
## 5 How it is used

| Step | Who | What happens |
|---|---|---|
| **1 Name the decision type** | Class owner, Policy Maker | One type with a name that somebody can list examples of: supplier concessions, non-standard discounts, credit exceptions, change requests above a size |
| **2 Write the boundary** | Risk function, Policy Maker | The nine guardrails, in one sentence each where possible: the exposure the organisation will not accept in this type whoever is deciding; a narrowing is live on save |
| **3 Name the owner** | Accountable director | The Own cell of the rights matrix: the person who answers for the outcomes whoever acts; exactly one |
| **4 Set the room** | Accountable executive | Each decider's dial, one per person or agent; the projection shown before any move commits; a widening activated by a second person |
| **5 Decide under the grant** | Decision Maker, agent | Inside the effective range the gate passes automatically, verified at the band's rate; at the edge, an intent notice or a referral to one available authorised person |
| **6 Watch the pattern** | Accountable owner | Override telemetry and outcomes by type; the ledger's recommendation with its evidence |
| **7 Move the room, not the boundary** | Accountable executive | The dial moves one point, with a reason; the guardrail is not touched; if reversals rise, it moves back the same way |

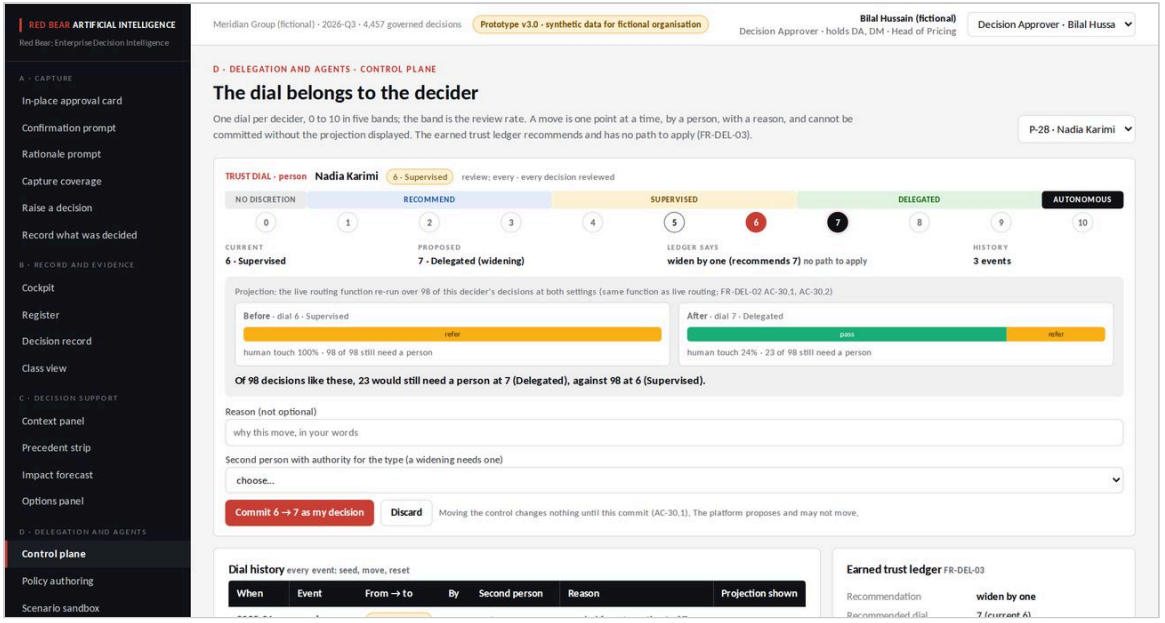
An illustration, using Meridian Group, the invented organisation used in Red Bear scenarios and sample data. Supplier concessions above a stated value require a group approval. That figure is a financial-limit guardrail on the decision type, and it stays. Separately, the procurement lead has settled ninety-one concessions below it in eleven months with two reversals and no policy breaches. Her dial moves from 6 to 7: from Supervised, where every decision was reviewed, to Delegated, where her concessions are sampled and, at the edge of her range, she states intent and proceeds after the notice period unless stopped. The guardrail was not touched, the risk function was not involved, a second person activated the widening, and the change took an afternoon. If reversals rise, the dial moves back the same way.

The four questions, asked separately: *Where is the wall?* An exposure statement, owned by risk, changes rarely. *How much room inside it?* A competence judgement, owned by the accountable executive, changes on evidence. *Who stays answerable?* Fixed, named on every decision made under the grant. *How reversible is it?* A property of the decision type, not of the person.

# 6 The user experience




**S-23 Decision rights register**: S-23 Decision rights register. Recommend, Decide, Own, Execute and Review per decision type against person, committee and agent; the Own row with exactly one filled cell; authority kept current from the directory. Prototype render on synthetic data for fictional entities.




**S-20 Control plane**: S-20 Control plane. The guardrails on the decision type and the dial on the decider shown as two separate objects, with the projection of a one-point move before it can be committed. Prototype render on synthetic data for fictional entities.

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

### Inputs to the grant and the record under it

| Input | Source | Check |
|---|---|---|
| **Who may decide what** | HR and the directory, through read-only connectors | Authority, accountability and ownership held as three distinct things; accountability never transfers to an agent (FR-DEL-01, RBP FR-CN-03) |
| **Guardrails** | Authored per decision type; versioned, owned, tested | Conflicts detected at authoring; exceptions refused without a sunset; every evaluation records the version used (FR-DEL-15, FR-DEL-14) |
| **Dial settings** | Set by the accountable executive, one point at a time | Projection before commit; widening activated by a second person; recorded with reason and name (FR-DEL-02, RBP FR-PR-04) |
| **Decisions under the grant** | Every capture route; the agent gateway for agents | The acting user from the verified token only; no interface, API or agent may assert who it is acting as (FR-PLT-02, FR-DEL-09) |
| **Overrides** | Explicit, with a non-blank reason | A separate event attributed to the person; the review obligation derived from the band at the time (FR-DEL-05) |
| **Cover and delegation** | Nominated cover; standing or temporary delegation | Temporary delegation never wider than the holder's own; both parties named on every decision (FR-RCH-02, FR-RCH-03) |
| **Outcomes** | Connected systems against the decision identifier, or the owner's review | An outcome the product guessed is not evidence (FR-OUT-01) |

### How it is checked

- **Connectors are read-only and consented.** Authority data is read from HR and the directory under least privilege granted by the customer administrator; the product writes nothing back to them (RBP FR-CN-01).
- **The AI proposes; a person decides.** The ledger's recommendation, the sandbox's projection and the policy optimiser's proposal are put in front of the owner; none is applied by the product (FR-DEL-03, FR-DEL-16, FR-DEL-17).
- **Agents are governed, not observed.** An agent that acts and reports afterwards is being observed; governance is the gateway call before the act (FR-DEL-09).
- **A changed model starts again.** A new version, a retrained weight set or a materially different input resets the agent near the bottom of Supervised; there is no partial credit for a record the current version never built (FR-DEL-13).
- **The record replays as it stood.** Policy and dial history are retained with the same rigour as the decision (FR-REC-09).
- **Patterns, never individuals.** Override telemetry is a question about the policy; the person-level view sits one level down and follows the reporting line (FR-DEL-18).

# 8 Benefits realised

Every measure is calculated from the organisation's own register against a baseline established before anything was switched on (FR-PLT-05).

| Benefit | How it shows up | Measured by |
|---|---|---|
| **The room widens on evidence** | A team's dials move up on its record; the guardrail is untouched | Dial moves upward per type; time from ledger recommendation to decision |
| **The executive stops being a queue** | Decisions settle inside the effective range | Share of decisions settled within the decider's own authority; approval volume reaching each level |
| **Escalation stops being delay** | Referrals carry the record and the clock; defaults are decisions | Time waiting attributed to absence vs judgement; referrals authorised unchanged, per class |
| **Accountability is locatable** | Every decision names owner and actor, including agents | Decisions with a named owner: all of them |
| **Latitude moves in both directions** | Reversals move a dial down as readily as records moved it up | Dial moves downward with a reason; reversal rate per type |
| **Agents earn authority class by class** | Shadow mode, then promotion by a person | Autonomy used against autonomy granted, per agent per type |

## 9 Risks managed

The strongest objection to everything above is not that it will not work. It is that it will work, in a direction nobody intended. A complete record of who decided what, on what reasoning, with what result, is a performance management system whether or not anyone calls it one, and the predictable response to being measured on your decisions is not better decisions but more defensible ones. The paper's own logic makes this worse: if latitude is widened on evidence, evidence becomes currency. Three things bear on it, and the third decides the matter. *What is measured*: process, scored apart from outcome, because outcomes contain luck, which is Duke's *resulting* (Duke, *Thinking in Bets*, chapter 1). *Who sees it*: narrow visibility by role, the decision captured rather than the conversation, consent and redaction, retention by class. *The asymmetry*: the same telemetry that could police people is the only thing that makes widening their latitude defensible. Without a record, an executive asked to give a team more room has nothing to say to the board except that they have a good feeling. With a record, "this type has run at Delegated for nine months with two recoverable reversals" is an argument. There is no version of this where you get the second without building the first.

| Risk | Control in Red Bear: Enterprise Decision Intelligence |
|---|---|
| **Defensive decisions** | Quality scored on process without access to outcomes; no interface blends the two (FR-OUT-05) |
| **The record as a disciplinary instrument** | Nothing in the product prevents an organisation from misusing it; restricting dial moves to a named person with a reason and recording every move makes a punitive use visible and traceable (FR-DEL-02) |
| **Evidence becomes currency** | Watch for it directly: rising escalation in delegated types, or latitude that only moves one way; reach and latency analysis report both (FR-RTE-05, FR-RCH-05) |
| **Wide visibility** | Access at class granularity; System Admin reaches no decision content; approver oversight follows the reporting line (FR-PLT-08, RBP FR-PR-05) |
| **The two dials welded together again** | The product holds no dial per decision type and no guardrail per person (FR-DEL-02, FR-DEL-19) |
| **An agent inherits trust it did not earn** | One agent, one purpose, one dial; ceiling of 9; a material model change resets it; promotion by a person, class by class, never global (FR-DEL-13, FR-DEL-11) |

| Risk | Control in Red Bear: Enterprise Decision Intelligence |
|---|---|
| **Latitude only ever moves downward** | A widening needs two people but is an afternoon's act, not a governance project; the ledger keeps proposing (FR-DEL-19, FR-DEL-03) |

What to do on Monday, none of which requires buying anything: pick one decision type, not a department; try to find out what was decided in it last quarter, not what was spent; write the boundary and the room as two sentences and note which you could change this month; name the accountable owner and check it is one person; add reversibility to the type's definition; put a review date on the room and none on the boundary; give the grant a purpose, not only a limit. If that exercise is easy, the delegation is probably working. If it takes three weeks and ends in an argument about who owns the type, you have located the problem, and it was never about whether anyone was willing to let go.

## 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. The external research is other people's work, described with its limits, and none of it is evidence about the product.

| Cited as | Source | Authority |
|---|---|---|
| **Oncken and Wass** | W. Oncken and D. Wass, *Management Time: Who's Got the Monkey?*, Harvard Business Review (1974): the five degrees of initiative | Practitioner reference |
| **Tannenbaum and Schmidt** | R. Tannenbaum and W. Schmidt, *How to Choose a Leadership Pattern*, Harvard Business Review (1958; retrospective 1973): the continuum with limits around it | Academic reference |
| **Hersey and Blanchard** | P. Hersey and K. Blanchard, *Management of Organizational Behavior* (1969 onward): situational leadership; readiness per task | Academic reference |
| **Marquet** | L. D. Marquet, *Turn the Ship Around!*, Portfolio (2012): "I intend to..." | Practitioner reference |
| **ADP 6-0** | US Army, *ADP 6-0, Mission Command* (2012, revised 2019): commander's intent | Doctrine |
| **Bezos** | J. Bezos, *2015 letter to shareholders*, Amazon (2016): one-way and two-way doors | Practitioner reference |
| **Rogers and Blenko** | P. Rogers and M. Blenko, *Who Has the D?*, Harvard Business Review (January 2006): a single decider per class | Practitioner research |
| **Blenko, Mankins and Rogers** | M. Blenko, M. Mankins and P. Rogers, *Decide & Deliver*, Harvard Business Review Press (2010): quality, speed, yield and effort | Practitioner research |
| **Lovallo and Sibony** | D. Lovallo and O. Sibony, *The case for behavioral strategy*, McKinsey Quarterly (March 2010): the 1,048-decision study | Practitioner research |
| **Duke** | A. Duke, *Thinking in Bets*, Portfolio (2018): resulting | Practitioner reference |
| **Institute of Directors** | Institute of Directors, guidance on the role and responsibilities of the board: delegation of authority, not liability | Professional body |
| **Book 1** | D. Oliver, *Empower, Trust, Verify*, Book 1 (2026), chapters 5, 7 and 8: guardrails, classes, the Trust Dial | Red Bear |
| **Red Bear** | EDI-FSD-01 v3.0 draft; Red Bear Platform Functional Specification RBP-FSD-01 | Red Bear |

## Glossary

| Term | Meaning in this paper |
|---|---|
| **The boundary** | The guardrail on the decision type: exposure the organisation will not accept, whoever is deciding; owned by risk |

| Term | Meaning in this paper |
|---|---|
| **The room** | The Trust Dial on the decider: how much of the permitted range this person or agent exercises today; owned by the accountable executive |
| **Guardrail** | One of nine fixed kinds of limit per decision type |
| **Trust Dial** | One per decider, 0 to 10, five bands, each a review rate; an agent's ceiling is 9 |
| **Role scope** | Which decision types a decider may touch at all |
| **Intent notice** | The decider declares intent, a clock runs, anyone with standing may stop it with a reason; on expiry the decision proceeds on notice |
| **Decision rights** | Recommend, Decide, Own, Execute, Review; Own is one named person |
| **Accountable owner** | The person who answers for outcomes in a decision type whoever acts, including an agent |
| **Pre-authorised envelope** | A standing decision agreed in advance as versioned policy; a decision inside it passes Authorise and names the envelope's version |
| **Agent gateway** | The call an agent makes before acting, returning allow, hold or route; the call is the capture event |
| **Reversibility** | One of five properties that raise a decision's class; a property of the type, not the person |

## 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 (this paper) |
| **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 |

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