PROVE
TECHNOLOGY DECISION INTELLIGENCE

INDUSTRY WHITE PAPER · BUILT ENVIRONMENT

The True Cost of Technology Decisions

Why the Built Environment Needs Decision Intelligence



Technology decisions should be based on evidence, not assumptions. This paper shows what it costs when they are not, and what a defensible decision looks like.

OCTOBER 2026 · VERSION 1.2

PROVETDI.COM

CONTENTS

01	Executive summary	3
02	The hidden cost of technology decisions	5
03	Why traditional selection falls short	7
04	The built environment technology challenge	9
05	From vendor claims to verified evidence	11
06	Introducing Technology Decision Intelligence	12
07	The economics of better decisions	14
08	From decision to outcome	16
09	A practical framework for decision-makers	18
10	The future of Technology Decision Intelligence	20
11	Conclusion	21
	Methodology and limitations	22
	References	23

HOW TO READ THIS PAPER

A paper about evidence should be held to its own standard. Every material number in these pages carries a label showing where it comes from, so you can weigh it accordingly.

SOURCE	Published external research, regulation or official data. Numbered references link to page 23.
INTERNAL	Prove TDI's own research records. Internally classified, not independently audited.
MODEL	An illustrative calculation on stated, fictional assumptions. Not an observed result.
VIEW	Editorial interpretation by Prove TDI, offered for debate.

Why this paper exists

Boards have read the papers that say technology projects overrun. Few have been shown why the overrun was predictable, who could have known, and what record would have proved it. This paper is about that record: the missing asset in most technology decisions, and the reason the same mistakes keep arriving with new logos on them.

A Prove TDI industry white paper, informed by external research and internal analysis. It is not independent or peer-reviewed research. Regulatory information covers the European Union and the United Kingdom, verified on 25 October 2026, and is not legal advice. All scenarios, organisations and figures labelled *Model* are fictional.

We audit invoices, not decisions.

> An organisation can tell you what it paid for a €4 coffee in March. Ask why it chose a €4 million building platform and you will get a story, a slide and a person who has since left.

The built environment is in the middle of its largest technology cycle. Portfolios are wiring buildings for energy regulation, hybrid work, asset performance and carbon reporting at the same time. The spend is visible. The quality of the decisions behind it is not.

The research on technology investment has been consistent for more than a decade. Large IT projects overrun their budgets and under-deliver their value, and a meaningful minority fail so badly that they threaten the organisation that commissioned them[1,2]. Most buyers now regret at least one software purchase within eighteen months[4]. None of this is new, which is precisely the problem. The industry has spent years describing failure and very little time changing how decisions are made and remembered.

This paper argues that the root cause is not a shortage of information. Buyers have never had more of it. The root cause is the absence of a **defensible decision record**: a governed account of what was needed, what was known, how reliable it was, who decided, and what was expected to happen next.

Without that record, every technology decision accrues what we call **decision debt** VIEW. Like technical debt, it is invisible on the day it is incurred and expensive every day after. It surfaces as integration work nobody budgeted, licences nobody uses, a replacement project nobody can justify, and a board question nobody can answer: *why did we choose this?*

Technology Decision Intelligence is the discipline of paying that debt down. It structures requirements, grades evidence by its provenance and freshness, makes comparison transparent, keeps humans accountable for judgement, and then does something almost no organisation does today: it compares what was promised with what happened.

Information is not the same as a decision record

	Technology information	A defensible decision record
Answers	What exists in the market?	Why this option, for us, now, and on what evidence?
Belongs to	Vendors, analysts, the internet	The organisation making the decision
Over time	Decays, and is rarely re-checked	Versioned, and re-tested against outcomes
Test it passes	Is it accurate?	Could we defend it to a board, an auditor or a successor?
Typical form	Decks, demos, reports, spreadsheets	Linked requirements, graded evidence, rationale, approvals

SEVEN FINDINGS FOR DECISION-MAKERS

Findings 1 to 4 and 6 draw on cross-sector research into IT projects and software buying. Those studies did not measure built-environment technology specifically; applying them here is our interpretation.

1 The contract price is not the cost of the decision.

In a 2012 study of more than 5,400 IT projects across sectors, large projects ran 45% over budget and delivered 56% less value than predicted[2] SOURCE. In our illustrative built-environment scenario, the decision process adds 29% to contracted spend, and 14% of the contract pays for licences nobody uses MODEL.

Implication: negotiate the decision, not just the licence. The largest costs are set before the contract is signed.

2 Software purchase regret is widespread.

In a cross-sector survey, 59% of software buyers regret at least one purchase made in the past 18 months, and 49% of them cite increased costs[4] SOURCE. Earlier research found mismanaged expectations behind 42% of regret[3].

Implication: regret is a symptom of expectations that were never written down well enough to be tested.

3 More information makes buyers smaller, not smarter.

Buyers overwhelmed by high-quality but contradictory information were 153% more likely to settle for a smaller, less disruptive purchase than planned[5] SOURCE.

Implication: the scarce resource is not data. It is a method for resolving contradictions.

4 How you decide matters more than how much you analyse.

Across 1,048 major decisions, the quality of the decision process explained six times more of the variance in outcomes than the depth of analysis[6] SOURCE.

Implication: invest in decision governance before buying another research subscription.

5 Building technology decisions now carry legal memory.

EU rules require automation systems in larger non-residential buildings to interoperate across manufacturers[7], remove cloud switching charges from January 2027[8] and impose vulnerability reporting on connected products from September 2026[9] SOURCE.

Implication: "why did we choose this?" is becoming a compliance question, not only a commercial one.

6 Approval is where value is won or lost, not where it is secured.

Projects with excellent change management met objectives 88% of the time, against 13% with poor change management, a correlation rather than proof of cause[10]. Commissioning existing buildings delivered median energy savings of 16% with a payback of 1.1 years[11] SOURCE.

Implication: a decision without a scheduled outcome review is an unfinished decision.

7 The missing asset is the record, and failure is well documented.

Prove TDI's case library holds 35,427 documented technology decision failures and risk events across 29 industries, 31,058 of them classified internally as supported by primary evidence[21] INTERNAL.

Implication: the patterns are knowable in advance. Organisations fail to learn from them because nobody keeps the receipt.

02 THE HIDDEN COST OF TECHNOLOGY DECISIONS

The licence is the only cost anyone sees. So it is the only one anyone negotiates.

> Procurement is optimised for the number on the contract. The decision is paid for everywhere else: in meeting rooms, integration backlogs, unused seats and the replacement project five years later.

Evaluation is a cost centre nobody budgets. A platform decision for a property portfolio typically pulls in facilities, IT, security, finance, procurement, sustainability and workplace teams. Each reads the same vendor material, attends overlapping demonstrations and rebuilds the same comparison in a different spreadsheet. This effort rarely appears in a business case because it is spread across salaries that are already paid.

Demonstrations show performances, not evidence. A demo is designed to show the best path through a product, on clean data, by its most skilled operator. It is useful theatre. It is not proof that the product meets your requirements in your buildings with your integrations.

Misaligned requirements become change requests. When requirements are vague at selection, they are clarified during implementation, at contract rates. Handoff problems between sales and implementation drive regret for 43% of dissatisfied software buyers[3] SOURCE.

Integration is where optimism meets the existing estate. Building systems speak different protocols, sit on different networks and are maintained by different contractors. Higher-than-expected total cost of ownership is the most-cited product reason for software regret[3]. Every additional year a large IT project runs adds around 15% to its cost overrun[2] SOURCE.

Low adoption is a slow leak. Organisations in one large SaaS-management dataset wasted an average of $21 million a year on unused licences[15] SOURCE. The figure comes from a vendor's customer base and should be read as indicative, but the mechanism is universal: a poorly fitted product is paid for in full and used in part.

Lock-in is priced at exit, not at entry. The UK competition regulator found that fewer than 1% of cloud customers switch provider each year, and that egress fees and technical barriers deter switching[14] SOURCE. Switching charges are being regulated away in the EU[8]; the cost of migrating data, retraining people and running systems in parallel is not.

Opportunity cost and unverified benefits close the loop. Every month of delay defers the benefits that justified the investment. And because expected benefits are rarely recorded in a testable form, nobody notices when they fail to arrive. Poor project performance wastes nearly 10% of every dollar invested[16] SOURCE.

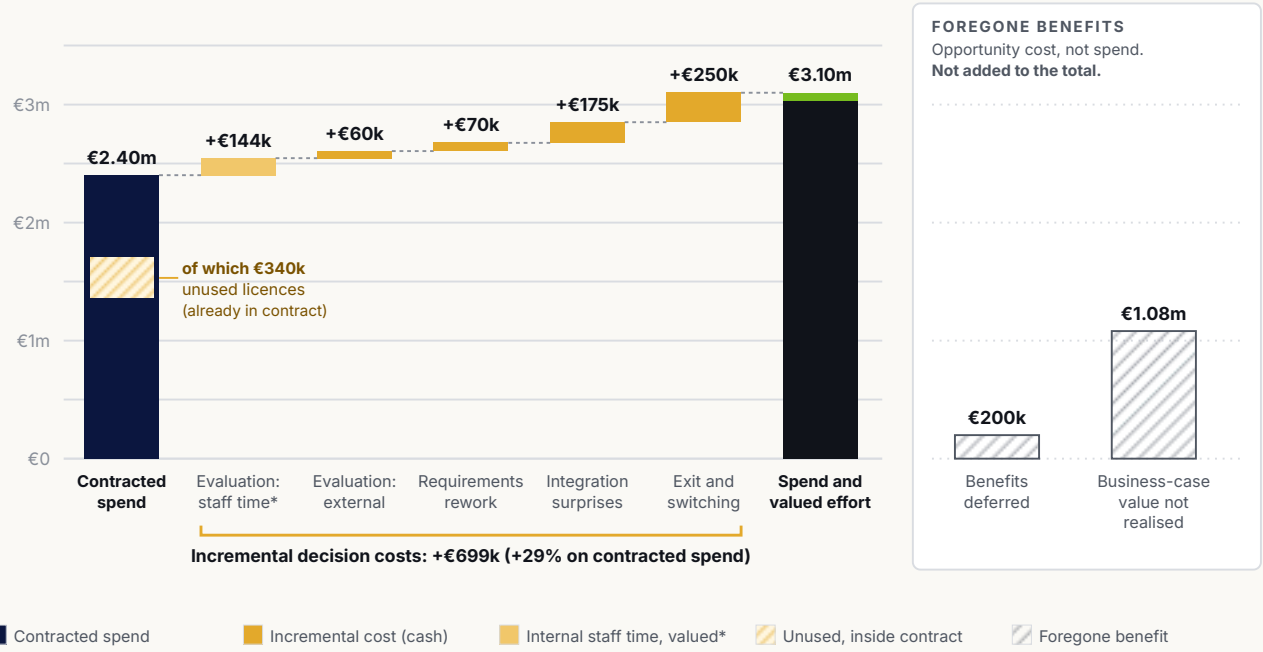
Decision debt, defined VIEW

The accumulated cost of technology choices that cannot be explained, tested or revisited. It compounds because each undocumented decision becomes a constraint on the next one: the integration nobody can explain, the contract nobody dares unwind, the platform renewed because no one can prove the alternative would be better.

VISUAL 1 · MODEL

The full cost of one technology decision

Fictional scenario: a 120-building portfolio selects an integrated workplace platform with occupancy analytics. Seven-year horizon, EUR.



Reading the chart. Only contracted spend and incremental costs are added together. Unused licences are part of the contract, so they are marked inside it and not counted again. Deferred and unrealised benefits are opportunity costs, not spend, so they sit in a separate panel and are not added to any total. *Internal staff time is valued at a blended day rate; it is an economic cost, not cash. All amounts are illustrative assumptions and are not observed client results.

Model assumptions

Component	Assumption	EUR
Subscription	€243k a year for 7 years	€1.70m
Implementation services	Fixed contract value	€700k
Contracted spend		**€2.40m**
Evaluation: staff time*	12 people × 15 days × €800 (valued, not cash)	€144k
Evaluation: external	RFP administration and advice	€60k
Requirements rework	10% of implementation as change requests	€70k
Integration surprises	25% uplift on implementation	€175k
Exit and switching	Migration and parallel running at replacement	€250k
Incremental decision costs		**€699k**
Spend and valued internal effort		**€3.10m**
MEMO ITEMS: NOT ADDED TO THE TOTAL		
Unused licences	20% of subscription, already within contracted spend	€340k
Benefits deferred	4 months of a €600k annual benefit	€200k
Value not realised	30% shortfall on a €3.60m business case	€1.08m

+29%

added to contracted spend by evaluation, rework, integration and exit. €144k of it is staff time rather than cash. MODEL

€340k

of contracted spend, 14%, pays for licences that are never meaningfully used. It is waste inside the contract, not an extra cost. MODEL

€1.08m

of business-case value not realised, assuming a 30% shortfall. Cross-sector research on large IT projects reports 56%[2]. MODEL

Each assumption lives in the editable model; change one and every figure on this page recalculates. The point is not the precise total. It is that most of it is set before anyone signs.

03 WHY TRADITIONAL TECHNOLOGY SELECTION FALLS SHORT

A spreadsheet is not a memory. A score is not a reason.

> The standard selection process was designed for buying products. Today's decisions are about committing to ecosystems, data and partners for a decade.

The information is supplied by the people being judged. Most of what a buyer knows about a product arrives through the vendor: websites, decks, RFP answers, curated references. It is not dishonest. It is simply not independent, and it is rarely dated, sourced or scoped to the buyer's context.

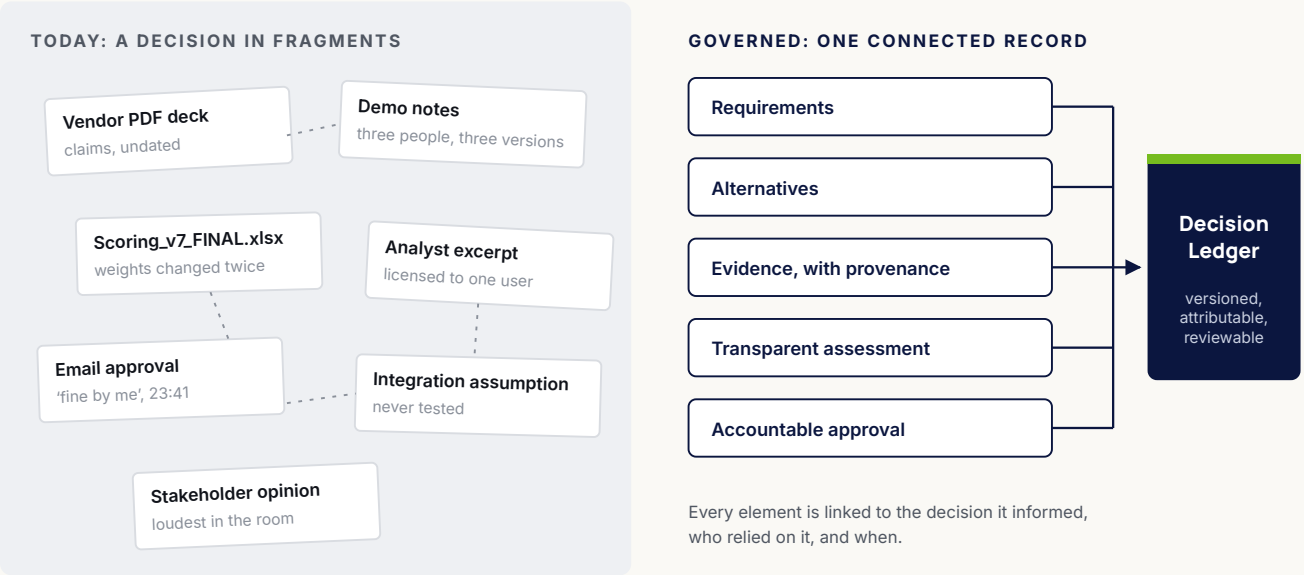
Comparisons expire on delivery. Products ship monthly. A comparison built in spring describes a market that no longer exists by the time the contract is signed in autumn, and nobody records which claims were checked against which release.

Scoring hides judgement inside arithmetic. A weighted matrix looks objective. In practice, weights move when a favourite falls behind, "partial" becomes a 3 or a 4 depending on who is typing, and the final number conceals the debate that produced it. Precision is not the same as rigour.

Approvals are scattered across inboxes. Security signs off in one thread, finance in another, the steering group in a meeting with no minutes. Months later, nobody can reconstruct what was approved, on which version of the evidence, with which conditions.

VISUAL 2

The same decision, two operating models



Left: how evidence typically exists during a selection, as disconnected artefacts with no shared provenance. **Right:** a governed process in which each element is linked to the decision it informed. Conceptual diagram; the fragment labels are typical examples rather than data.

Why more data does not mean better decisions

In a Gartner survey of more than 1,000 B2B buyers, 89% said the information they encountered was high quality. They were still overwhelmed, because trustworthy sources contradicted each other. Faced with that, buyers were 153% more likely to retreat to a smaller, less ambitious purchase[5] SOURCE.

That finding deserves more attention than it gets. It means that piling on research does not reduce risk. Past a certain point it converts a strategic decision into a defensive one: the buyer chooses whatever is easiest to justify, not what is best.

Decision science has long made a related point. In McKinsey's study of 1,048 decisions, process quality outweighed analytical depth by a factor of six in explaining outcomes[6] SOURCE. Process here means challenging assumptions, airing dissent and separating facts from advocacy, which is exactly what a matrix in a shared drive does not do.

Information is abundant. Proof is scarce.

THE CENTRAL PROBLEM OF MODERN TECHNOLOGY BUYING

The answer is not less information but better-classified information: knowing which kind of knowledge each statement represents, and how much weight it can bear.

Four kinds of knowing, and what each can carry

Type	What it is	What it can carry
Vendor-provided evidence	Claims, documentation, certifications and references supplied by the seller	Defines what is offered. Must be attributed and dated. Cannot settle a contested point alone.
Independent verification	Testing, audit, regulatory records or corroboration by parties without a stake	Can confirm or refute a claim. Carries most weight on material, irreversible criteria.
Expert judgement	The informed opinion of practitioners, consultants or internal specialists	Interprets evidence and fills gaps. Must be labelled as judgement and attributed to a person.
Customer-specific requirements	What this organisation needs, given its estate, policies, risks and goals	Determines relevance. Evidence that is true elsewhere may be irrelevant here.

The test of a good selection process

- Every important statement can be traced to a source, a date and a type of knowledge.
- Contradictions are visible and resolved on the record, not in a corridor.
- Weights and scores are set before vendors are scored, and changes are logged.
- Approvers can see what they are approving, and the record shows that they did.
- A successor could reconstruct the decision in an afternoon.

Buildings last for decades. Their software does not. The decision record has to outlive both.

Buildings and construction account for 28% of global energy consumption and 37% of global CO_2 emissions, according to UNEP and the Global Alliance for Buildings and Construction[17] SOURCE. The technology that runs them has become a climate, compliance and workplace instrument at once.

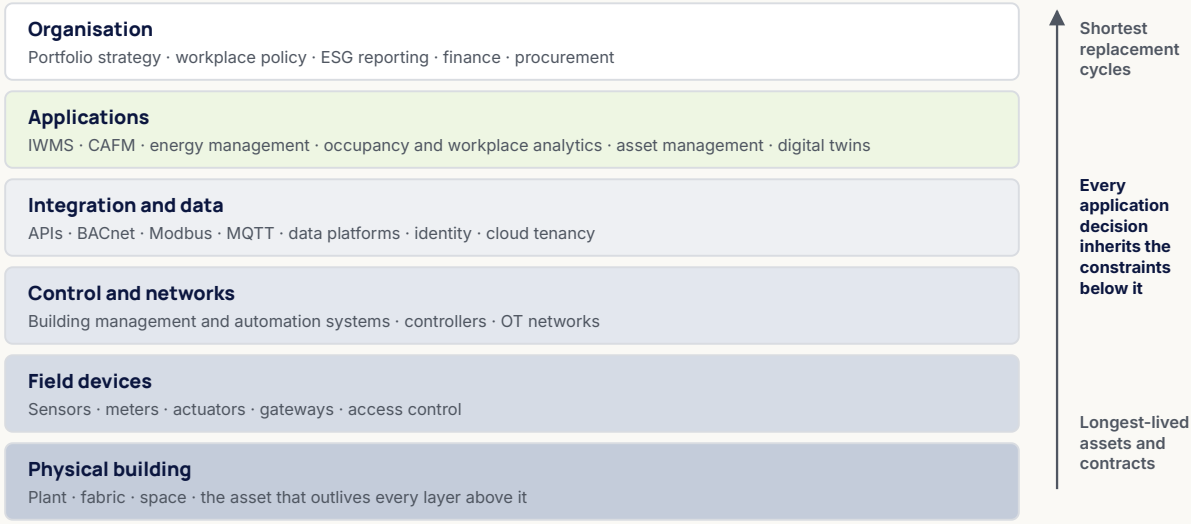
That makes it unusually hard to choose well. An IWMS or CAFM platform touches leases, space, maintenance and finance. Building management and automation systems sit on operational networks never designed for the internet. Occupancy and workplace analytics raise data-protection questions in every jurisdiction. Energy management depends on meter data the platform does not own. PropTech tools multiply at the edges. Digital twins depend on all of the above being accurate.

The result is a stack in which each layer has a different owner, contract term and replacement cycle. A decision at the application layer quietly inherits every constraint beneath it: protocol support, network segmentation, data rights, the contractor who holds the controller passwords. Requirements collide with existing investments, dependencies with contracts, and lifecycle risk with regulation.

Underperformance is the norm. Lawrence Berkeley National Laboratory research cites evidence that US commercial buildings may use up to 30% more energy than fault-free buildings, and finds that proprietary automation systems and inconsistent protocol implementations hold back the tools designed to fix this[12] SOURCE.

VISUAL 3

The built-environment stack: every decision inherits the layers below



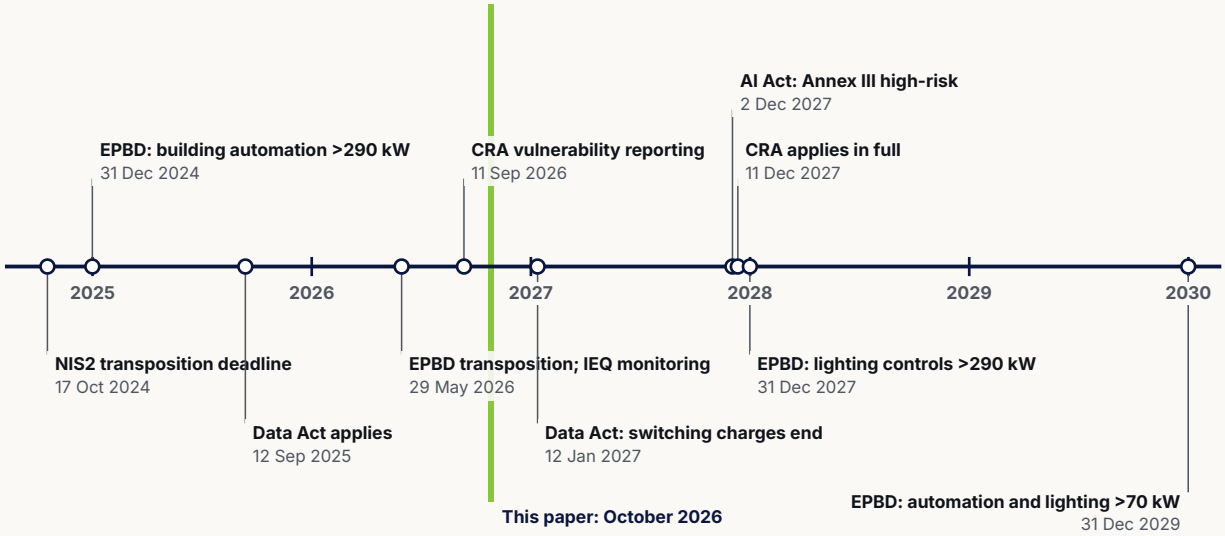
Conceptual model. Layers and examples are indicative. Replacement cycles vary by asset and contract, so the diagram shows relative ordering rather than durations.

The regulatory clock is already running

Regulation is turning technology choices in buildings into obligations with dates attached. Several are live; more arrive within the life of a platform contract signed this year. A decision taken today will be examined against rules that tighten during its lifetime, and the organisation will need to show its reasoning.

VISUAL 4 · SOURCE

EU dates that shape built-environment technology decisions, 2024 to 2029



Sources: EPBD recast Article 13 [7]; Data Act [8]; Cyber Resilience Act [9]; AI Act as amended by Regulation (EU) 2026/1744 [18]; NIS2 [19]. Dates verified 25 October 2026. Thresholds refer to the rated output of heating, air-conditioning and combined ventilation systems in non-residential buildings.

Instrument and jurisdiction	What it requires	Why it matters for technology decisions
EPBD recast, Directive (EU) 2024/1275, EU member states[7]	Building automation and control systems in non-residential buildings above 290 kW by 31 Dec 2024 and above 70 kW by 31 Dec 2029, where feasible. Systems must monitor, benchmark, detect efficiency losses and interoperate across manufacturers and proprietary technologies. Indoor environmental quality monitoring by 29 May 2026.	Interoperability becomes a regulatory expectation, not a nice-to-have. Selections should record how cross-vendor operation was evidenced. Applicability depends on national transposition.
Data Act, Regulation (EU) 2023/2854, EU[8]	Applies from 12 Sep 2025. Customers of data processing services, including SaaS, gain switching rights; charges limited to direct costs until 12 Jan 2027 and generally prohibited after.	Exit terms and data portability become negotiable facts to record at selection. Migration effort remains a real cost.
Cyber Resilience Act, Regulation (EU) 2024/2847, EU[9]	Manufacturers must report actively exploited vulnerabilities and severe incidents from 11 Sep 2026; most obligations apply from 11 Dec 2027.	Connected sensors, gateways and controllers need evidence of a credible vulnerability-handling and update commitment.
NIS2, Directive (EU) 2022/2555, EU[19]	Cybersecurity risk-management and supply-chain duties for essential and important entities; transposition deadline 17 Oct 2024.	Where an occupier or operator is in scope, supplier security becomes an auditable decision criterion.
AI Act, Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744, EU[18]	Transparency duties apply; high-risk obligations now apply from 2 Dec 2027 (stand-alone, Annex III) and 2 Aug 2028 (embedded in products, Annex I).	AI features in building and workplace tools need documented purpose, data and human oversight.

United Kingdom: the EU instruments above do not apply directly. UK organisations should assess equivalent domestic regimes and the effect of EU rules on suppliers and EU-based assets. Cyber incidents remain costly everywhere: the global average cost of a data breach was $4.4 million in 2025[13]. This overview is not legal advice.

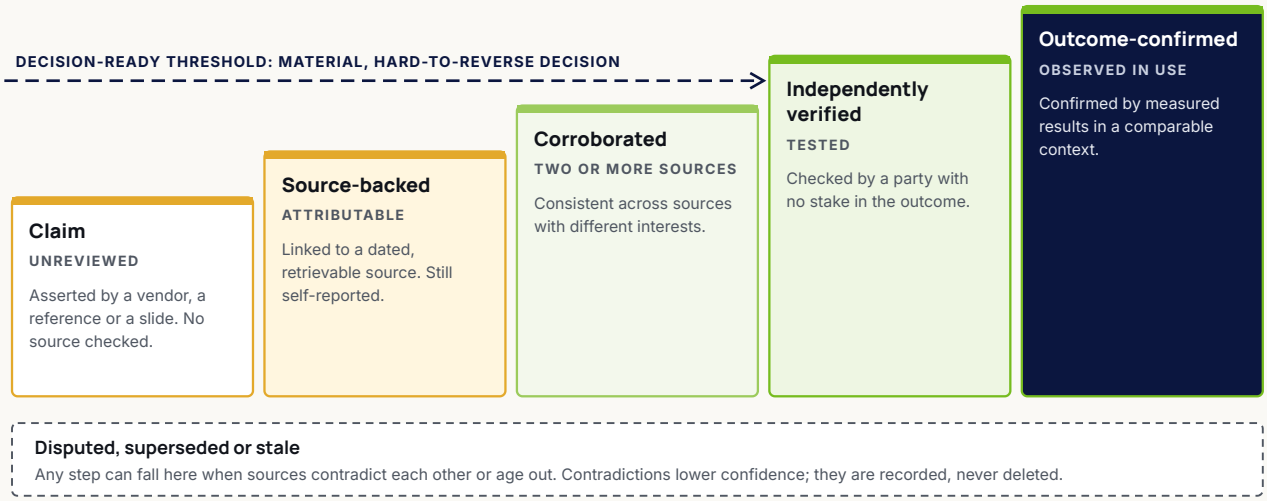
05 FROM VENDOR CLAIMS TO VERIFIED EVIDENCE

A demo is a performance. Evidence is a record.

> Evidence-based assessment does not mean distrusting vendors. It means knowing exactly how much weight each statement can bear, and saying so in writing.

VISUAL 5

Claims versus evidence: how a statement earns the right to decide



Framework. The stages mirror the evidence maturity states used in Prove TDI's methodology. The threshold is not fixed: a reversible, low-cost choice can proceed on source-backed evidence, while a portfolio-wide, ten-year commitment should require independent verification on its critical criteria.

Seven properties of decision-grade evidence

Provenance	Who said it, where, when, and with what interest in the outcome.
Confidence	How strongly the evidence supports the claim, stated rather than implied.
Freshness	Every claim has a half-life. Release notes, certifications and prices age at different speeds.
Contradictions	Conflicting sources are recorded side by side. Deleting the inconvenient one is not resolution.
Applicability	Proven where, at what scale, on which version, with which integrations?
Verification	Has anyone without a stake checked it, and how?
Usage rights	Can this source lawfully be stored, shared and cited in the decision record?

Evidence must fit the decision, not the market

A product that meets a requirement for a single-tenant headquarters may fail it across a multi-tenant portfolio with three BMS brands. "Proven at scale" is meaningless until the scale, context and integrations are named. The same claim can be decision-ready for one organisation and irrelevant to another.

When evidence conflicts or runs out

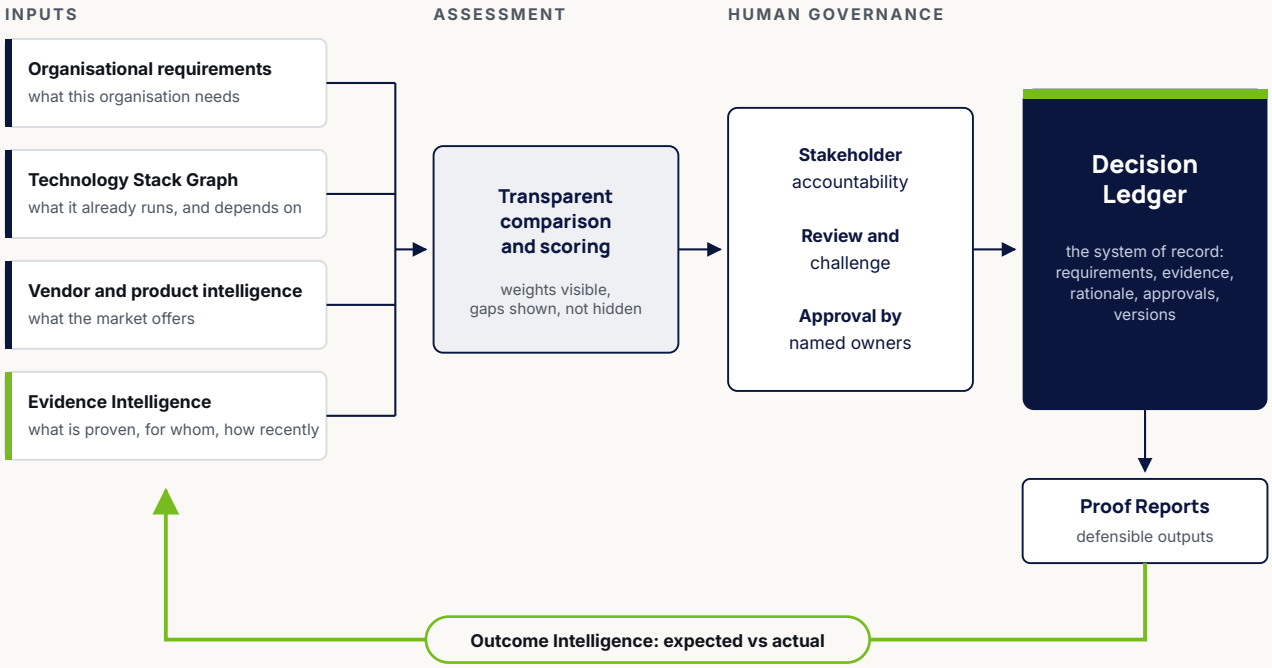
Incomplete or contradictory evidence should lower decision confidence visibly, not disappear into an average score. The honest options are to gather targeted evidence, to make the gap a contractual condition, to stage the commitment, or to accept the risk explicitly with a named owner. A requirement marked *unknown* is information. A requirement silently scored as 3 out of 5 is not.

A recommendation is not proof.

> Technology Decision Intelligence is the discipline of structuring, evaluating, recording and reviewing technology decisions so that each one can be defended and each one makes the next one better.

VISUAL 6

The architecture of Technology Decision Intelligence



Outcomes flow back into requirements and evidence, so every decision improves the next one.

Methodology diagram. Shows how the components relate in the Prove TDI approach. It describes the method, not a certification of which product features are generally available today; see the status note overleaf.

The category is deliberately narrower than "AI for procurement" and broader than a comparison tool. It treats the decision itself as the unit of value, and the decision record as an asset to be governed like a financial ledger: append-only, attributable, reviewable.

Prove TDI was built to be that record for built-environment technology: an **independent Technology Decision Intelligence layer and system of record for technology decisions**. It is not a marketplace, a procurement platform, a consultancy, a lead-generation service or an autonomous recommendation engine. Vendors cannot buy favourable rankings, scores or recommendations.

AI does real work in this model. It extracts, maps, compares, detects contradictions and drafts. It does not decide. Qualified people own context, trade-offs, challenge and approval, and the record shows who they were.

Ten components, one record

Component	Role in the decision	What it replaces
Organisational requirements	Structured, prioritised needs with owners, stated before any vendor is assessed.	A requirements list written backwards from a favoured demo
Technology Stack Graph	A map of current systems, integrations, contracts and dependencies the decision must respect.	Tribal knowledge of “what talks to what”
Vendor and product intelligence	Structured, versioned information about vendors and products, with each claim attributed.	Decks and brochures in a shared folder
Evidence Intelligence	Grades each claim by provenance, confidence, freshness and verification; surfaces contradictions.	Treating every statement as equally true
Transparent comparison and scoring	Weights, fit classifications and gaps visible to all stakeholders; unknowns shown as unknown.	A single number that ends the conversation
Stakeholder accountability	Named owners for requirements, evidence review and risks.	Collective responsibility, meaning none
Human review and approval	Challenge and sign-off recorded against the exact version of evidence reviewed.	Approvals scattered across inboxes
Decision Ledger	The append-only system of record linking requirements, evidence, rationale, approvals and versions.	A folder nobody can find in three years
Proof Reports	Board-ready outputs that show the question, evidence, limitations, alternatives and approvals.	Vendor-sponsored reports and generic slides
Outcome Intelligence	Compares expected with actual outcomes and feeds lessons into future requirements and evidence.	Post-implementation reviews that never happen

Status note: method versus product

This paper describes the Technology Decision Intelligence methodology. It does not claim that every component above is independently verified as operational in the Prove TDI platform today, and readers should not infer that. Prove TDI will state the current availability of each capability on request and in writing. That distinction is itself an example of the discipline the paper advocates.

Independence, by design

Commercial participation, whether by buyers or vendors, never influences scores, evidence treatment, rankings or the prominence of any option in a decision output. Vendor-paid services are separated from independent evaluation and labelled as such. Where independence cannot be shown, it should not be claimed.

Don’t follow the herd. Prove what actually works.

07 THE ECONOMICS OF BETTER DECISIONS

Bad decisions are expensive for buyers. They are expensive for vendors too.

> The selling side wastes as much effort on poor-fit evaluations as the buying side does. A better decision process is one of the rare reforms in which both parties gain.

For buyers

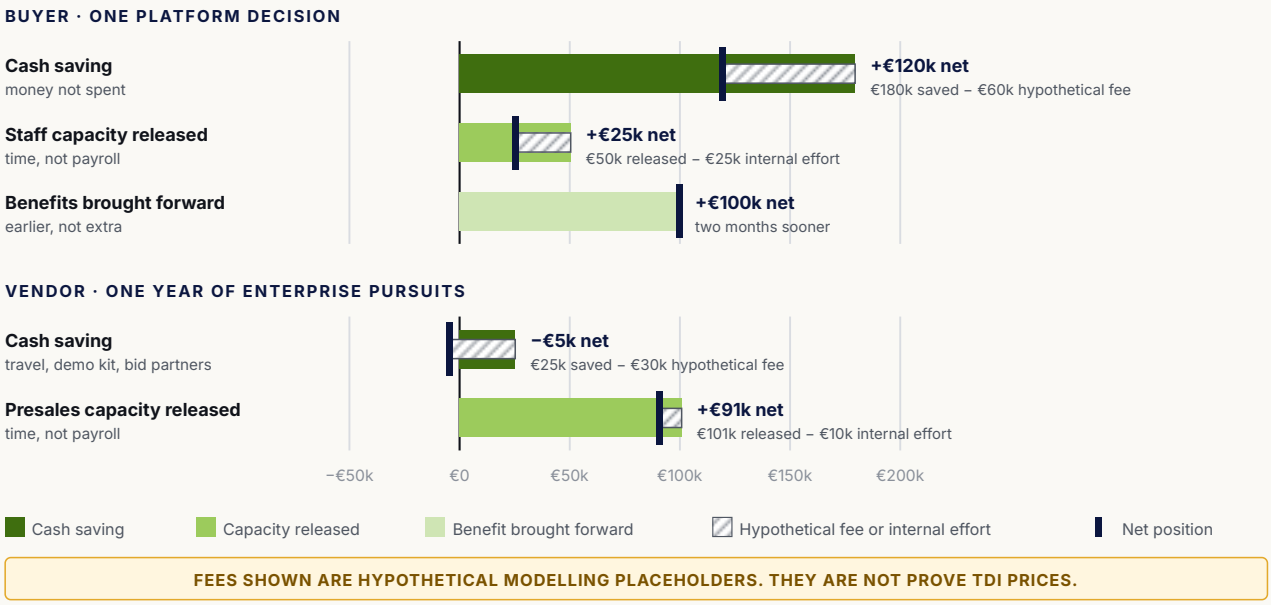
Less duplicated research and fewer unfit demonstrations release internal time. Clearer requirements and earlier visibility of integration risk avoid rework. Better fit reduces unused licences. Faster consensus brings go-live, and benefits, forward. Recorded expectations make investment accountability possible for the first time.

For vendors

Poor-fit opportunities are identified before presales effort is spent. Requirements arrive structured, so the same qualification questions are not answered ten times. Evaluations focus on criteria that matter to the buyer, and well-evidenced products can show it. The principle fits on a badge: less pitch, more proof.

VISUAL 7 · MODEL

What a better decision process returns, by type of gain



FEES SHOWN ARE HYPOTHETICAL MODELLING PLACEHOLDERS. THEY ARE NOT PROVE TDI PRICES.

Illustrative model on fictional assumptions, detailed overleaf. Gains are kept apart by type because they are not equivalent. Cash savings are money not spent. Released staff capacity becomes financial value only if it is redeployed or avoids hiring; it does not reduce payroll by itself. Benefits brought forward arrive earlier, not in addition. Both panels share one scale. Fees are hypothetical placeholders and are not Prove TDI prices.

Hypothetical fees. The fees below are modelling placeholders chosen to test the economics. They are not Prove TDI prices, quotes or price indications.

Buyer model: applied to the Visual 1 scenario

Driver	Assumption	EUR
CASH SAVINGS		
External support	35% less external selection support	€21k
Rework and integration	30% of rework and integration surprises avoided	€74k
Licences	25% of unused licences avoided, assuming volumes are rightsized before contract commitment	€85k
Less: participation fee HYPOTHETICAL		−€60k
Net cash effect		**€120k**
STAFF CAPACITY (NOT CASH)		
Evaluation time	35% less internal evaluation effort	€50k
Less: internal effort to run the process		−€25k
Net capacity released		**€25k**
TIMING (NOT EXTRA BENEFIT)		
Go-live	2 of 4 months' delay avoided	€100k

The licence saving is cash only if commitments can be avoided, reduced or renegotiated; the model assumes rightsizing happens before signature. On that basis, cash savings cover the hypothetical fee 3.0 times, and the fee pays back once 33% of them are achieved.

Vendor model: one year of enterprise pursuits

Driver	Assumption	EUR
BASIS		
Poor-fit pursuits	3 of 10 opportunities identified as poor fit early; 70% of a €40k pursuit cost normally spent first	€84k
Qualification	15% less repetitive qualification on the other 7	€42k
SPLIT: 20% CASH, 80% STAFF TIME		
Cash saving	Travel, demo environments, bid partners	€25k
Less: participation fee HYPOTHETICAL		−€30k
Net cash effect		**−€5k**
Capacity released	Presales time	€101k
Less: internal effort to supply evidence		−€10k
Net capacity released		**€91k**

The vendor's cash position is slightly negative (−€5k). Released presales time is not payback in itself. The fee is recovered only if that time produces measurable gross margin or avoided spend, such as deferred hiring, worth at least €5k: 5.3% of its €91k modelled value. Time that is not redeployed returns nothing. Win-rate effects are excluded.

Why paying participants does not have to compromise independence

Every serious assurance market has faced this question. Auditors are paid by the companies they audit; ratings agencies by issuers. The answer is never to pretend money is absent. It is to separate what can be paid for from what cannot, and to make the separation inspectable.

In Technology Decision Intelligence, a fee can buy participation: the ability to submit evidence, respond to requirements and maintain accurate product information. It cannot buy a score, a ranking, an interpretation of evidence or a place in a recommendation. Evidence from a paying vendor is graded by exactly the same rules as evidence from a non-paying one, and labelled by source either way.

The independence firewall

Can be paid for	Can never be paid for
Participation in the evidence process	Scores or rankings
Structured requirement responses	Evidence grading or interpretation
Product information maintenance	Recommendations or prominence
Services for the buyer's own process	The buyer's final decision

08 FROM DECISION TO OUTCOME

Approval is not the finish line. It is a hypothesis with a budget.

> Every business case predicts the future. Almost none are ever checked. The organisations that learn fastest are the ones that write their expectations down precisely enough to be proven wrong.

Record expectations as testable statements. "Improve space utilisation" cannot be verified. "Consolidate two floors in the Northgate building within 12 months of go-live, measured by badge and sensor data, saving €400k a year" can. Each expected outcome needs a metric, a baseline, a target, a date and an owner.

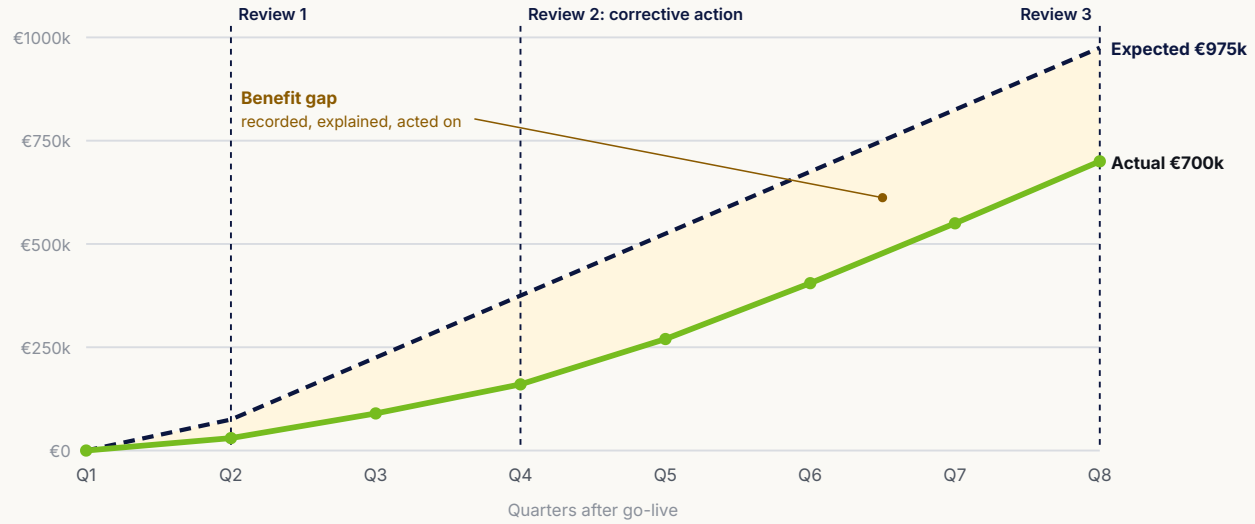
Schedule reviews at approval. Outcome reviews should be part of the decision, not an afterthought. The evidence on follow-through is strong: commissioning existing buildings delivered median savings of 16% with a 1.1-year payback across 643 buildings[11], and change management quality is strongly associated with whether projects meet their objectives[10] SOURCE.

Measure five things. Value realisation against the business case; adoption by the intended users; operational performance against requirements; cost variance against the full decision cost, not only the contract; and the lessons that should change the next decision.

Keep the history. A governed decision history becomes institutional memory. When the platform comes up for renewal, the organisation can see what was promised, what happened and why, instead of starting again from vendor decks.

VISUAL 8 · MODEL

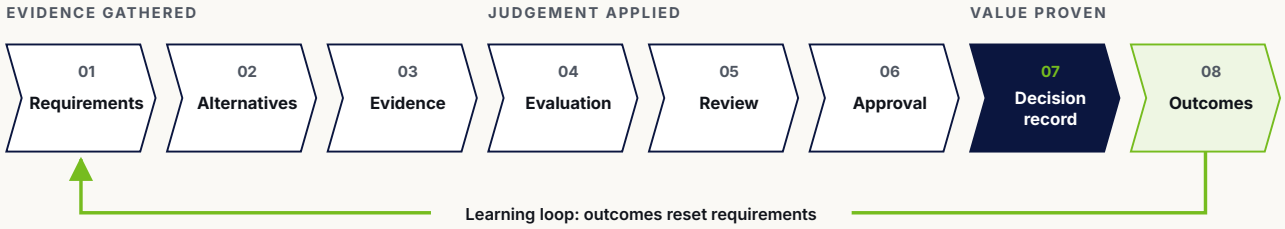
Expected versus actual outcomes: what a recorded hypothesis makes visible



Illustrative figures only. Cumulative benefits for a fictional deployment. Expected benefit: €150k a quarter at full run-rate. At Review 2, actual benefits are 43% of expected, triggering corrective action on adoption; by Q8 the run-rate has recovered but cumulative realisation stands at 72%. Without a recorded expectation, the gap would be invisible.

VISUAL 9

A governed decision lifecycle



Each stage writes to the record. Nothing is decided in a stage that leaves no trace.

VISUAL 10 · FICTIONAL EXAMPLE

What a defensible decision record contains

All names, vendors, people and figures below are fictional. The structure is the point: could your last technology decision produce this page?

DECISION RECORD · DR-2026-0417 · VERSION 3.2

Occupancy sensing and workplace analytics, Northgate Estates portfolio

APPROVED WITH CONDITIONS

DECISION QUESTION

Which occupancy sensing and analytics approach should Northgate Estates (fictional) deploy across 38 office buildings to support a 20% space consolidation target by 2028, within existing BMS and data-protection constraints?

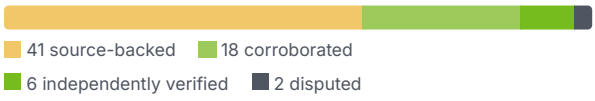
REQUIREMENTS · 21 TOTAL (12 MUST, 9 SHOULD)

Meets R-03 GDPR-compliant anonymisation at the edge
Partially meets R-07 BACnet integration with two BMS brands
Unknown R-11 Five-year device firmware support
Conflicting evidence R-14 Sensor accuracy above 95%

ALTERNATIVES CONSIDERED

Option A: sensor platform (selected)
Option B: badge and Wi-Fi analytics only
Option C: alternative sensor platform
Option D: do nothing; defer to 2027
Rationale for each rejection recorded below.

EVIDENCE BASE · 67 ITEMS



RISKS AND CONTRADICTIONS

R-14: vendor claims 98% accuracy; one customer reference reports 89% in open-plan areas. **Condition:** 60-day pilot in two buildings with acceptance threshold of 93% before portfolio roll-out.

APPROVALS

Head of Workplace (owner) · 14 Apr 2026
CISO · 15 Apr 2026, with condition on firmware support
Data Protection Officer · 15 Apr 2026
CFO · 17 Apr 2026
Each approval linked to evidence version 3.1 or later.

RATIONALE

Option A is the only alternative meeting all GDPR-critical requirements with corroborated evidence. Option B cannot measure meeting-room use. Option C failed R-03. Deferral risks missing the 2028 consolidation window.

EXPECTED OUTCOMES AND REVIEW DATES

Two floors consolidated in two buildings by Q2 2027; €400k annual saving from Q3 2027; 85% of managers using weekly reports by month six. Reviews: Oct 2026, Apr 2027, Oct 2027.

VERSION HISTORY

v1.0 · requirements baselined, 9 Feb
v2.0 · shortlist of three, 12 Mar
v3.0 · evidence review complete, 2 Apr
v3.1 · R-14 contradiction logged, 8 Apr
v3.2 · conditional approval, 17 Apr

09 A PRACTICAL FRAMEWORK FOR DECISION-MAKERS

Ten disciplines of a defensible technology decision

This framework requires no particular software. It can run on documents and discipline. What matters is that each step leaves a trace another person can follow.

01 Define the decision

Write the question as a single sentence with scope, constraints and a deadline. Name the decision owner.

Ask: what would make this decision a failure in three years?

02 Establish organisational requirements

Prioritise must-haves and should-haves before looking at vendors. Give every requirement an owner and a test.

Ask: which requirement would we drop first, and who agrees?

03 Map dependencies and constraints

Document existing systems, integrations, contracts, data rights and regulatory duties the choice must respect.

Ask: who controls the interfaces we depend on?

04 Identify credible alternatives

Include at least one option you did not expect, and always include doing nothing or deferring.

Ask: are we comparing options, or confirming a favourite?

05 Collect and assess evidence

Attribute and date every material claim. Grade it by provenance, freshness and verification. Separate claims from proof.

Ask: which critical claim rests on the vendor's word alone?

06 Record contradictions and risks

Keep conflicting evidence side by side. Assign each material risk an owner and a response.

Ask: what are we choosing not to know?

07 Apply transparent decision criteria

Fix weights before scoring. Show unknowns as unknown. Log every change to weights or scores with a reason.

Ask: would the ranking survive a change of one weight?

08 Secure accountable approvals

Approvers sign against a specific version of the evidence. Conditions are written into the record and the contract.

Ask: does each approver know exactly what they approved?

09 Preserve the decision rationale

Explain why the chosen option won and why each alternative lost, in language a successor will understand.

Ask: could a new CFO reconstruct this in an afternoon?

10 Measure actual outcomes

Schedule reviews at approval. Compare expected with actual outcomes and feed lessons into the next decision.

Ask: when will we know if we were wrong, and who will tell us?

TEAR-OUT CHECKLIST

Before you sign: the defensible decision checklist

Free to copy and use with attribution to Prove TDI.

01 The decision

- [ ] The decision question is one written sentence, with scope and deadline
- [ ] A single named owner is accountable for the outcome
- [ ] We have defined what failure would look like

02 Requirements

- [ ] Must and should requirements were set before vendor contact
- [ ] Every requirement has an owner and a way to test it
- [ ] Stakeholders from IT, security, finance and operations have signed off

03 Dependencies

- [ ] Existing systems, integrations and contracts affected are listed
- [ ] Interface owners, including contractors, are identified
- [ ] Applicable regulations and jurisdictions are recorded with dates

04 Alternatives

- [ ] At least three options assessed, including do nothing or defer
- [ ] The reason each alternative was rejected is written down
- [ ] Exit and switching terms compared for each option

05 Evidence

- [ ] Every material claim is attributed to a source and dated
- [ ] Critical criteria rest on more than vendor-provided evidence
- [ ] Evidence matches our scale, context and integrations

06 Contradictions and risks

- [ ] Conflicting evidence is recorded, not averaged away
- [ ] Unknowns are visible and either resolved or accepted by name
- [ ] Material risks have owners, responses and contract conditions

07 Criteria and scoring

- [ ] Weights were fixed before scoring began
- [ ] All changes to weights or scores are logged with reasons
- [ ] The ranking was tested for sensitivity to key weights

08 Approvals

- [ ] Each approver signed against a specific evidence version
- [ ] Conditions of approval are written into the contract
- [ ] AI-generated content in the record is labelled as such

09 Rationale

- [ ] The case for the chosen option fits on one page
- [ ] The full cost of the decision, not only the contract, is estimated
- [ ] The record is stored where a successor will find it

10 Outcomes

- [ ] Expected outcomes have metrics, baselines, targets and dates
- [ ] Outcome reviews are scheduled and owned
- [ ] Lessons will be fed into the next related decision

Can you prove why you chose your technology?

Our rule of thumb: every unticked box is decision debt. Decide whether to pay it now or later, and write that choice down too.

10 THE FUTURE OF TECHNOLOGY DECISION INTELLIGENCE

From one good decision to an organisation that remembers.

> The long-term prize is not a better selection. It is institutional decision memory: every technology choice the organisation has made, with its evidence and outcomes, queryable when the next one arrives.

Predictions about technology deserve the same evidence labels as claims about products. We separate what exists now from what is plausible soon and what remains a forecast.

HORIZON 1 · EXISTS TODAY

Structured, evidenced decisions

Structured requirements, attributed evidence with maturity states, transparent comparison, recorded approvals and versioned decision records are established practice in well-governed organisations and can be run today, with or without specialist software.

AI already accelerates extraction, summarisation and comparison of vendor documentation under human review.

HORIZON 2 · NEAR-TERM POSSIBILITY

Continuous evidence monitoring

Evidence that refreshes itself: alerts when a certification lapses, a product is end-of-lifed, a security advisory is published or a contradicting source appears, with the affected decisions flagged automatically.

Outcome Intelligence that compares expected with actual results across many decisions, showing which kinds of claims tend to hold and which do not.

HORIZON 3 · LONGER-TERM PREDICTION

Institutional decision memory

Decision records that interoperate across organisations, consultants and auditors in anonymised form, giving the built environment something it has never had: a shared, evidence-graded memory of what actually works.

Governed AI research agents that assemble evidence packs against a decision question, with every source and inference inspectable.

AI assists. Humans remain accountable.

Generative AI makes it cheap to produce persuasive text, which makes provenance more valuable, not less. A summary without sources is a new kind of vendor claim. Shadow AI already adds an average of $200,000 to the cost of a data breach, and 63% of organisations report lacking AI governance policies[13] SOURCE.

The right model is governed AI: every AI contribution labelled, every source retrievable, every material judgement made and signed by a named person.

Principles for AI in technology decisions

VIEW

- AI may extract, map, compare, detect, model and draft.
- People own context, trade-offs, challenge and approval.
- No AI output enters the record without its sources.
- AI-generated content is always labelled.
- Material decisions carry human signatures, not model confidence scores.

11 CONCLUSION

The receipt is the strategy.

The financial argument. The contract is a fraction of what a technology decision costs. Evaluation effort, rework, integration, delay, under-use and exit are largely determined before signature, which means they are largely determined by the quality of the decision process.

The operational argument. In the built environment, every decision inherits a stack of systems, contracts and contractors. Without a map of those dependencies and an honest grading of evidence, integration risk is discovered at the most expensive moment.

The governance argument. Regulation is giving technology choices a legal memory. Boards, auditors and regulators will increasingly ask organisations to show their reasoning. "The vendor said so" and "it scored highest in the spreadsheet" will not be adequate answers. When UK regulators fined TSB £48.65 million over its 2018 platform migration, their findings centred on operational risk management and oversight of outsourced suppliers[20] SOURCE: the governance of a technology decision, not its code.

The learning argument. Organisations that record expectations and compare them with outcomes get better at deciding. Those that do not repeat the same mistakes with new vendors, which is why the failure statistics have barely moved in fifteen years.

The question is no longer whether organisations have enough technology information. It is whether they have enough reliable evidence to defend their decisions.

About Prove TDI

Prove TDI is the independent Technology Decision Intelligence layer for the built environment, and a system of record for technology decisions. It combines independent market intelligence, AI-assisted analysis and accountable human judgement to help organisations make technology decisions they can defend. Vendors do not pay for rankings, scores or recommendations.

See a governed decision in five minutes

Try the interactive demonstration at provetdi.com/demo. It takes about five minutes, needs no account and uses clearly labelled fictional data. Then bring one live decision and see what its record looks like, where the evidence is strong, and where the debt is.

provetdi.com/demo

METHODOLOGY AND LIMITATIONS

How this paper was researched, and where it could be wrong

Research approach

This is a Prove TDI industry white paper informed by external research and internal analysis. It is not independent or peer-reviewed research. It draws on published research from professional and academic institutions, official regulatory texts, regulator decisions and established industry surveys. Each external figure was retrieved from the publisher or, where the primary report is paywalled, from the publisher's own summary or a reputable secondary report that names the primary source. All sources were accessed or re-verified on 25 October 2026.

Evidence labels

Source marks published external evidence. **Internal** marks Prove TDI's own research records, which are internally classified and have not been independently audited; an internal verification label is not external validation. **Model** marks illustrative calculations. **View** marks editorial interpretation.

Illustrative models

Visuals 1, 7 and 8 use fictional scenarios, implemented in an editable script supplied with the source files; every chart, table and in-text figure is generated from the same assumptions and totals are reconciled by script. From version 1.1, four accounting rules apply following external review. Only contracted spend and incremental costs are added. Waste inside the contract, such as unused licences, is shown as a memo item and never counted twice. Deferred and unrealised benefits are opportunity costs, shown separately and never added to spend. Internal staff time is valued at a blended day rate and labelled as non-cash, and released staff capacity is reported apart from cash savings because it does not reduce payroll by itself. The models demonstrate mechanisms, not expected results, and must not be read as customer outcomes. Fees in Visual 7 are hypothetical and are not Prove TDI prices. Licence savings assume volumes are rightsized before contract commitment, and the vendor case counts released time as value only when it produces measurable margin or avoided spend.

The Prove TDI case library

Figures from the case library [21] reflect its dated summary register as at 25 October 2026 (backup R170): 35,427 cases, 31,058 classified as verified by primary evidence, 22,749 rated high confidence, across 29 industries. The library documents publicly reported technology decision failures and risk events across sectors, not only the built environment. It is not a random sample, so it describes patterns of reported failure rather than failure rates. Named organisations are kept internal; the one named case in this paper is cited to its regulator directly [20].

Material limitations

- Survey data [3, 4, 5, 10] is self-reported and may reflect respondent and vendor-panel bias.
- Project-overrun and buying studies [1–6, 10] concern IT projects and software purchasing across sectors, not built-environment technology specifically. Studies [1, 2] concern IT projects broadly, largely pre-2013 and above $15 million in [2]; their transfer to built-environment technology is an interpretation.
- Building performance research [11, 12] is US-based and may not transfer directly to European portfolios.
- Licence-waste data [15] comes from a SaaS-management vendor's own customer base.
- There is limited independent research quantifying the financial value of decision traceability itself. We have flagged this as an evidence gap rather than filling it with estimates.

Regulatory information

Covers EU instruments and notes UK applicability. National transposition of directives varies by member state. Nothing in this paper is legal advice.

Independence and interests

This paper was produced by Prove TDI, which has a commercial interest in Technology Decision Intelligence. No vendor sponsored or reviewed it, and no vendor is recommended in it. Prove TDI's own role is confined mainly to Section 6 and the closing pages.

REFERENCES

1. Flyvbjerg, B. and Budzier, A. (2011). Why your IT project may be riskier than you think. *Harvard Business Review*, 89(9). Data summary: arXiv:1304.0265 (2013). https://arxiv.org/abs/1304.0265
2. Bloch, M., Blumberg, S. and Laartz, J. (2012). Delivering large-scale IT projects on time, on budget, and on value. McKinsey & Company with the University of Oxford, 1 October 2012. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/delivering-large-scale-it-projects-on-time-on-budget-and-on-value
3. Gartner Digital Markets (2023). 2024 Tech Trends survey (3,484 respondents, July 2023), reported by TechRepublic, 30 October 2023. https://www.techrepublic.com/article/gartner-global-software-trends/
4. Gartner Digital Markets (2025). 2025 Software Buying Trends: 3,500 buyers in nine countries, surveyed online in August 2024. Supporting publication of the same survey: Westfall, B., Capterra's 2025 Tech Trends Report, Capterra, 18 November 2024. https://www.capterra.com/resources/tech-trends-successful-buyer-purchase-journey/ https://gcom.pdo.aws.gartner.com/en/digital-markets/insights/2025-software-buying-trends
5. Gartner (2019). Gartner reveals new B2B sales approach to win in today's information age. Press release, 29 July 2019. https://www.gartner.com/en/newsroom/press-releases/2019-07-29-gartner-reveals-new-b2b-sales-approach-to-win-in-toda
6. Lovallo, D. and Sibony, O. (2010). The case for behavioral strategy. *McKinsey Quarterly*, 1 March 2010. https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/the-case-for-behavioral-strategy
7. Directive (EU) 2024/1275 on the energy performance of buildings (recast), Article 13. OJ L, 8 May 2024. Transposition deadline 29 May 2026. https://eur-lex.europa.eu/eli/dir/2024/1275/oj
8. Regulation (EU) 2023/2854 (Data Act). Applicable from 12 September 2025. https://eur-lex.europa.eu/eli/reg/2023/2854/oj
9. Regulation (EU) 2024/2847 (Cyber Resilience Act). Reporting obligations from 11 September 2026; most obligations from 11 December 2027. https://eur-lex.europa.eu/eli/reg/2024/2847/oj
10. Prosci. The correlation between change management and project success. Best Practices in Change Management research, more than 2,600 practitioners. Page updated 7 August 2026. https://www.prosci.com/blog/the-correlation-between-change-management-and-project-success
11. Mills, E. (2009). Building commissioning: a golden opportunity for reducing energy costs and greenhouse gas emissions. Lawrence Berkeley National Laboratory. https://cx.lbl.gov/2009-assessment.html
12. Pritoni, M., Lin, G., Chen, Y., House, J., Crowe, E. and Granderson, J. (2022). Market barriers and drivers for the next generation fault detection and diagnostic tools. Lawrence Berkeley National Laboratory. doi:10.20357/B7801T. https://doi.org/10.20357/B7801T
13. IBM Security and Ponemon Institute (2025). Cost of a Data Breach Report 2025, as reported by ASIS International, Security Management, July 2025. https://www.asisonline.org/security-management-magazine/latest-news/today-in-security/2025/july/Global-Data-Breach-Costs-Drop/
14. Competition and Markets Authority (2025). Cloud services market investigation, final decision, July 2025, as reported by DatacenterDynamics. https://www.datacenterdynamics.com/en/news/competition-is-not-working-well-says-uk-cma-concluding-cloud-market-investigation
15. Zylo (2025). 2025 SaaS Management Index, 16 January 2025. Vendor dataset: more than 40 million licences under Zylo management. https://zylo.com/news/2025-saas-management-index
16. Project Management Institute (2018). Pulse of the Profession 2018: Success in disruptive times. https://www.pmi.org/learning/thought-leadership/pulse/pulse-of-the-profession-2018
17. UNEP and Global Alliance for Buildings and Construction. Global Status Report for Buildings and Construction 2025–2026 (tenth edition). https://globalabc.org/global-status-report
18. Regulation (EU) 2026/1744 of the European Parliament and of the Council (Digital Omnibus on AI), amending Regulation (EU) 2024/1689 (AI Act). In force 27 July 2026. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
19. Directive (EU) 2022/2555 (NIS2 Directive). Transposition deadline 17 October 2024. https://eur-lex.europa.eu/eli/dir/2022/2555/oj
20. Financial Conduct Authority and Prudential Regulation Authority (2022). TSB fined £48.65m for operational resilience failings. Press release, 20 December 2022. https://www.fca.org.uk/news/press-releases/tsb-fined-ps4865m-operational-resilience-failings
21. Prove TDI (2026). Enterprise Technology Decision Failure Case Library: documented failures and risk events. Summary register as at 25 October 2026 (backup R170). Internal research asset; classifications are internal and not independently audited.

PROVE
TECHNOLOGY DECISION INTELLIGENCE

Technology decisions you can defend.

LESS PITCH. MORE PROOF.

Prove TDI is the independent Technology Decision Intelligence layer for the built environment. We help organisations compare workplace and building technology against their requirements, see what is supported, what is uncertain and what still needs checking, and keep a governed record of every decision.

Try the five-minute interactive demonstration at provetdi.com/demo. No account needed.

© 2026 Prove TDI. Fictional examples are labelled as such.