

Decision Intelligence

QUARTERLY

The board-grade strategic briefing that turns a monitored topic into decisions, with its evidence, its counter-arguments and its own report card attached.

DETECT **Signal Scanner** → ASSESS **Change Tracker** → DECIDE **Decision Intelligence** → ACT **Action Triggers**

1 What it is

Decision Intelligence is the flagship strategic briefing of a Shaping Tomorrow continuous-intelligence programme, the deliverable that opens and closes each monitoring cycle on your topic. Where the weekly **Signal Scanner** detects and the monthly **Change Trackers** assess, Decision Intelligence decides: it converts a cycle's accumulated evidence into themes, decision postures, scenarios and board-ready questions, each traceable to a tiered, dated source.

A recent edition, "*The Summer the Grid Said No*", runs on 53 verified sources, four ranked themes, four owner-assigned strategic implications, a scenario matrix, four explicit exclusions and a claim-fidelity appendix, and it opens by grading its own previous predictions before it makes any new ones.

THE PROBLEM IT SOLVES

Strategic reports fail in two directions: they are either evidence-dumps that leave the decision to the reader, or opinion pieces that hide their sourcing and never revisit their own calls. Decision Intelligence is engineered against both. Every recommendation carries a posture, an owner and a dated action; every claim carries a tiered source; every argument travels with its counter-argument; and every edition starts by marking the previous edition's homework.

2 One document, two reading speeds

The briefing declares its own read estimate on the masthead, a 3-minute Executive Synthesis and a 22-minute full read, and the architecture honours both. A director can stop after the "Why this matters now" panel and the five graded shifts and still leave with the cycle's verdict; a strategy team can descend through themes, counter-arguments, scenarios and the source register and audit every claim to its origin.

Nothing in the fast path contradicts the deep path, because both are generated from the same evidence base.

3 Anatomy of the deliverable: element by element

Click a section to expand it.

The front door: masthead, continuity and the Futures Wheel

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Editorial title	A named thesis, “The Summer the Grid Said No”, with edition number and a link back to the previous report in the series.	Signals a maintained body of work, not a one-off PDF. The series-continuity line tells you your topic will be followed, not sampled.
Scope meta-line	Cycle date, named audience, and an explicit horizon: 6–18 months primary, 2–5 years secondary.	Sets expectations honestly. You know exactly what timeframe the judgements are calibrated to, a discipline most foresight work skips.
Read estimate	“3 min Executive Synthesis · 22 min full read”, stated up front.	Respects executive time and doubles as a design promise: the document commits to being useful at both speeds.
Interactive Futures Wheel	A clickable, expandable radial map of consequences: first-, second- and third-order effects of the central shift, colour-coded risk/opportunity/mixed, with line weight encoding impact × likelihood.	The whole causal landscape in one visual. Executives explore “what follows from this?” themselves, a workshop tool embedded in a report, and a striking screen-share moment in any board meeting.
What you’ll learn / Key takeaways	Twin cards: four questions the briefing answers, and four quantified answers (dates, gigawatts, dollar figures).	The contract with the reader, stated before the argument begins, and proof on page one that the answers are specific, not directional hand-waving.

The Executive Synthesis, the three-minute path

The synthesis is not a summary bolted on top; it is a self-contained decision document with its own internal discipline:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
“Why this matters now” panel	The cycle’s five hardened shifts, each in one bold-led sentence, framed for the board’s next agenda.	A director who reads nothing else leaves with the five things that changed and why they matter: pre-formatted for forwarding to a chair.
Evidence-led narrative	Four numbered developments, each anchored to linked primary sources (regulatory orders, market assessments, central-bank research) with dates and figures inline.	The argument and its receipts travel together. There is no “trust us” paragraph anywhere in the synthesis.
The pullquote	One italic line distilling the cycle’s thesis: “the system’s first response was neither to build nor to price but to acquire the right to refuse load.”	The quotable takeaway, the sentence that gets repeated in your own meetings, with the programme as its source.
Historical rhyme	A disciplined analogue, the 2001 telecom overbuild, used to sharpen the question, not decorate it.	Pattern-matching against history is what seasoned strategists do; the briefing does it for you, with the limits of the analogy stated.
“Where this analysis could be wrong”	A stress-test box naming the most consequential assumption, the rival reading, a falsifiable test and the specific signals that would overturn the thesis.	The feature sceptical buyers remember. A briefing that names its own kill conditions earns the right to be believed on everything else.
“What changed since the last scan”	The previous edition’s predictions graded in capitals: CONFIRMED, BROKEN UPWARD, REVERSED, NEW, plus a scan delta (53 sources vs 41).	Cycle-over-cycle accountability made visible. Most providers never revisit their calls; this one grades them before making new ones.
Continuity ledger	A verdict table resolving every dated commitment the previous edition made (UNRESOLVED, REOPENED, CONFIRMED), each with a fresh source.	Nothing raised in a previous cycle is silently dropped. Your open questions are tracked to resolution, which is what a programme is for.

Audience Snapshots: four lenses, one evidence base

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Role cards	Four colour-coded cards (Strategy & Foresight, Risk/Insurance/CFO, Infrastructure/Ops/CTO, Policy & Public Affairs), each naming the shift as it lands on that function.	One briefing serves four executive audiences without four documents. Each stakeholder finds their own stake in under a minute.
“The question to brief”	Each card ends with a single, concrete question for that function’s leadership: “should we add a water-availability gate to site selection before the next flagship build?”	Converts intelligence into agenda items. Your internal meetings are pre-drafted, the fastest route from reading to acting.

Themes, the argument layer

The cycle’s signals are organised into ranked themes (Immediate, Near-Term, Longer-Range), each carried by the same rigorous scaffold:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Impact-tier ranking	A colour-coded tier pill on every theme, defined in decision terms: Immediate = changes this half’s capital decisions.	Triage built in. You know which themes bear on this year’s decisions and which are multi-year watches.
Framing paragraph	The theme’s thesis in one dense paragraph, with its nuance intact (“the refusal was rehearsed, not executed”).	Analysis that holds nuance is what separates board-grade from newsletter-grade. You can judge the quality from any single paragraph.
Data figures	Purpose-built stat graphics: each with sources on the figure.	The evidence visualised for lifting straight into your own decks, sources attached.
Sourced signal bullets	Six to eight evidence items per theme, every one dated and linked.	Due-diligence-proof. An analyst can verify any bullet in one click; nothing rests on unattributed assertion.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Counter-argument box	A structured rebuttal per theme, the strongest rival reading, argued properly and sourced, with the named commercial actors it implicates.	The briefing argues against itself so your team doesn't have to. Boards trust analysis that has already survived its own opposition.
Weak signals to watch	Early, low-confidence signals flagged with an explicit escalation test: "would gain weight if a capacity auction discounts curtailable load as non-firm."	Shows the detection net running ahead of consensus, and defines, in advance, exactly what evidence would promote a hunch to a finding.
Decision link	A cross-reference from every theme to the Strategic Implication it feeds.	No orphaned analysis. Every theme exists because a decision depends on it, and the link proves it.

Strategic Implications, the decision layer

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Posture pill	Every implication is graded Decide, Prepare or Monitor.	Separates "act now" from "get ready" from "watch", a shared triage vocabulary across the whole programme.
Dated action + owner	A concrete deliverable with a deadline and an accountable role: "by 30 September, table a capital-committed position paper... Owner: CSO / CFO."	The recommendation arrives pre-routed and calendar-ready. Your cost of acting approaches zero.
Hedged reasoning	Each implication names the condition under which it would be wrong and stages the commitment accordingly.	Advice with built-in risk management. You are never asked to bet the plan on a single reading of the evidence.
"Who profits" line	The counterparties and suppliers to whom value migrates under each constraint.	Turns every risk into an investable and procurable opportunity map: useful to corporate development as well as risk.

Scenario Matrix, the planning layer

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Named axes	A 2×2 built on the two uncertainties the evidence cannot yet decide.	Scenarios derived from the evidence base, not imported from a template. The framing note even explains how the axes evolved from the previous edition, the matrix itself has a memory.
Four named worlds	Memorable, distinct 18-month environments: The Refusal Economy, Build or Be Refused, The Capex Hangover, Managed Glidepath.	Names that stick become your internal shorthand. Six months later the team is still saying “are we drifting toward the Capex Hangover?”
Tipping indicators	Every cell lists the observable events that would tip the world into it.	Converts scenario planning from a workshop exercise into a monitoring instrument the programme then tracks cycle by cycle.

What We Are Not Planning For, the exclusions layer

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Explicit exclusions	Scenarios deliberately held out of the plan, each with the reasoning stated.	Saying “no, and here is why” is rarer and more valuable than another risk. It protects you from resourcing against phantoms.
Reinstatement triggers	Each exclusion carries the named event that would put it back on the table.	The exclusions are falsifiable too. Nothing is dismissed forever; everything dismissed is monitored for re-entry.

Discussion Points, Source Register and Claim Fidelity

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Board-ready questions	Five forced-choice questions in the cycle’s own terms, including audience-tagged items (“For investors: ...”).	The last mile of intelligence: the agenda for your next strategy session, written for you. Chairs lift these questions verbatim.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Source Confidence Register	All sources in per-theme tables (source, tier, date, key claim) under a four-tier taxonomy, with the Tier 1–2 share stated and no Tier 4 load-bearing.	Radical sourcing transparency. A procurement or research team can audit the entire evidence base in minutes, and compare it to the unsourced decks it replaces.
Recency and freeze discipline	A stated recency window, a source-set freeze date, and older structural anchors explicitly labelled as context, not fresh readings.	Method hygiene that survives scrutiny: you know exactly when the evidence was cut and how old anchors are being used.
Conflict notes	Genuine source disagreements named and weighted rather than resolved away, plus an author-type coverage audit admitting where the source set is thin.	Most reports hide their contradictions; this one curates them. Confessing the evidence base's own gaps is the deepest trust signal in the document.
Claim-Fidelity Appendix	Every analyst inference and editorial framing named and separated from what sources actually assert.	You always know whether you are standing on evidence or interpretation, and can disagree productively. Effectively unique in commercial foresight.

4 The design system: why the format itself is a benefit

Like every Shaping Tomorrow deliverable, the briefing is engineered to be skimmed, read or audited, and its design carries the method:

- **A consistent posture and tier grammar:** Decide / Prepare / Monitor pills and Immediate / Near-Term / Longer-Range tiers use the same colours here as in the Signal Scanner and Action Triggers, so the whole programme reads as one system.
- **Interactive where it earns it:** the Futures Wheel expands on click for exploration and the contents cross-reference in both directions (theme → implication → theme), while the whole document remains print-clean for board packs.
- **Colour-coded epistemics:** counter-arguments, weak signals, stress tests and decisions each carry a distinct colour: you can see the confidence status of any block before reading a word of it.

- **Evidence at the point of claim:** links, dates and tier ratings sit inline and on the figures themselves, never banished to a bibliography the reader will not visit.

5 Why clients choose it

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
It grades its own homework	Every edition opens by scoring the previous edition’s predictions and resolving its dated commitments before making new claims.	Accountability compounds into trust. Renewal conversations start from a track record, not a promise.
It argues against itself	Counter-arguments, stress tests, exclusions and a claim-fidelity appendix are structural, not optional.	Your board gets analysis that has already survived adversarial review, and can see exactly where to push back.
It is decision-shaped	Postures, owners, dated actions, forced-choice questions and tipping indicators throughout.	The distance from reading the briefing to running the meeting is close to zero.
It compounds	Series continuity, evolving scenario axes, a continuity ledger and weak signals promoted or retired each cycle.	Each edition is worth more than the last, because it carries the whole programme’s memory. That is what a subscription buys that a one-off report cannot.


FROM A RECENT EDITION, “THE SUMMER THE GRID SAID NO” (JULY 2026)

In one 22-minute read, the briefing: graded the previous edition’s four predictions (one **confirmed via a new mechanism**, one **broken upward**, one **reversed**, one constraint **new to the register**); named the cycle’s pattern, the Refusal Layer, and disclosed the name as an analyst frame; issued four owner-assigned implications with dated actions; mapped four scenarios with tipping indicators; excluded four futures with reinstatement triggers; and put 53 tiered sources and a claim-fidelity appendix behind all of it.

6 How it's delivered

The Shaping Tomorrow intelligence cycle

Four formats on three clocks: *Signal Scanner* and *Change Tracker* are continuous; *Decision Intelligence* bookends the quarter; *Action Triggers* closes the cycle by reading across all three. This deliverable's lane is highlighted.



Decision Intelligence briefings open and close each monitoring cycle of a Single-Topic or Multi-Topic continuous-intelligence programme, working in concert with the weekly [Signal Scanner](#) and monthly [Change Trackers](#) that feed them and the [Action Triggers](#) synthesis that closes the cycle. Each edition is delivered as an interactive, print-ready briefing tailored to your topics, horizons and executive owners.

[Discuss a Tailored Programme](#)

[Request a Sample](#)