

PRODUCT STRATEGY • Q3 2026

From Signal to Strategy

A decision framework for product leaders navigating market uncertainty, quarterly

SECTION 01 · FOUNDATIONS

The landscape has shifted.
Here is what that means
for us.

MODULE 02


Before we score signals, we need to agree on what a signal is.

The word is overloaded. We use it to mean anything from a customer complaint to a macro trend. This framework requires a tighter definition.

The window for differentiation is narrowing.

Three converging forces — commoditized infrastructure, compressed release cycles, and rising customer switching costs — have reduced the average durable advantage window from 36 months to under 14. Teams that cannot split signal from noise inside that window miss the timing — reliably, and then schedule a retrospective about it.

Signals move through four stages; opinions skip to the end.
Four-stage processing pipeline — monthly cadence



IMPACT • PILOT RESULTS

Six months of results across four product teams.

Measured against pre-framework baseline, same teams, same market conditions. Alignment self-reported.



The framework has four components.

Signal Intake

Weekly structured collection across customer conversations, market data, and competitive moves. Normalized into a common schema before scoring.

Scoring Model

Each signal scored on three dimensions: confidence, recency, and strategic relevance. Weights are team-configurable — configured once, then defended forever.

Decision Log

Every decision recorded with the signals that informed it, the options considered, and the criteria applied. Feeds the calibration loop.

Calibration Loop

Monthly retrospective that compares predicted outcomes to actual outcomes and adjusts scoring weights accordingly.

The pipeline, by version number.

Signal Intake

v2.4

Handles 94% of structured signals without manual intervention. Average latency: 4 min from ingestion to scored entry.

Scoring Model

configurable

Three dimensions: confidence, recency, relevance. Default weights are 33 / 33 / 33 — the diplomatic opening position.

Decision Log

required

Every prioritization change above P2 must carry a logged rationale. No log, no change.

Calibration Loop

monthly

Compares predicted outcomes to actuals. First meaningful weight update happens after 2 cycles.

Signal Intake produces three outputs.

1

Weekly Signal Brief

The week's top ten signals, ranked and attributed. Lands Monday morning.

2

Anomaly Alerts

Real-time flags past the 2σ threshold, routed to the accountable PM. Four-hour response SLA, observed in spirit.

3

Monthly Signal Index

The calibration loop's source of truth. Required reading before each retrospective; skimming is detectable.

Two failure modes the framework is designed to prevent.

False signal amplification

One loud voice dominating the decision. The model caps any source at 30% — a ceiling one enterprise customer tests monthly.

Signal hoarding

Signals collected, decisions unlogged — the loop learns nothing. Above P2, the log is required. No log, no change.

Two intake modes for different signal types.

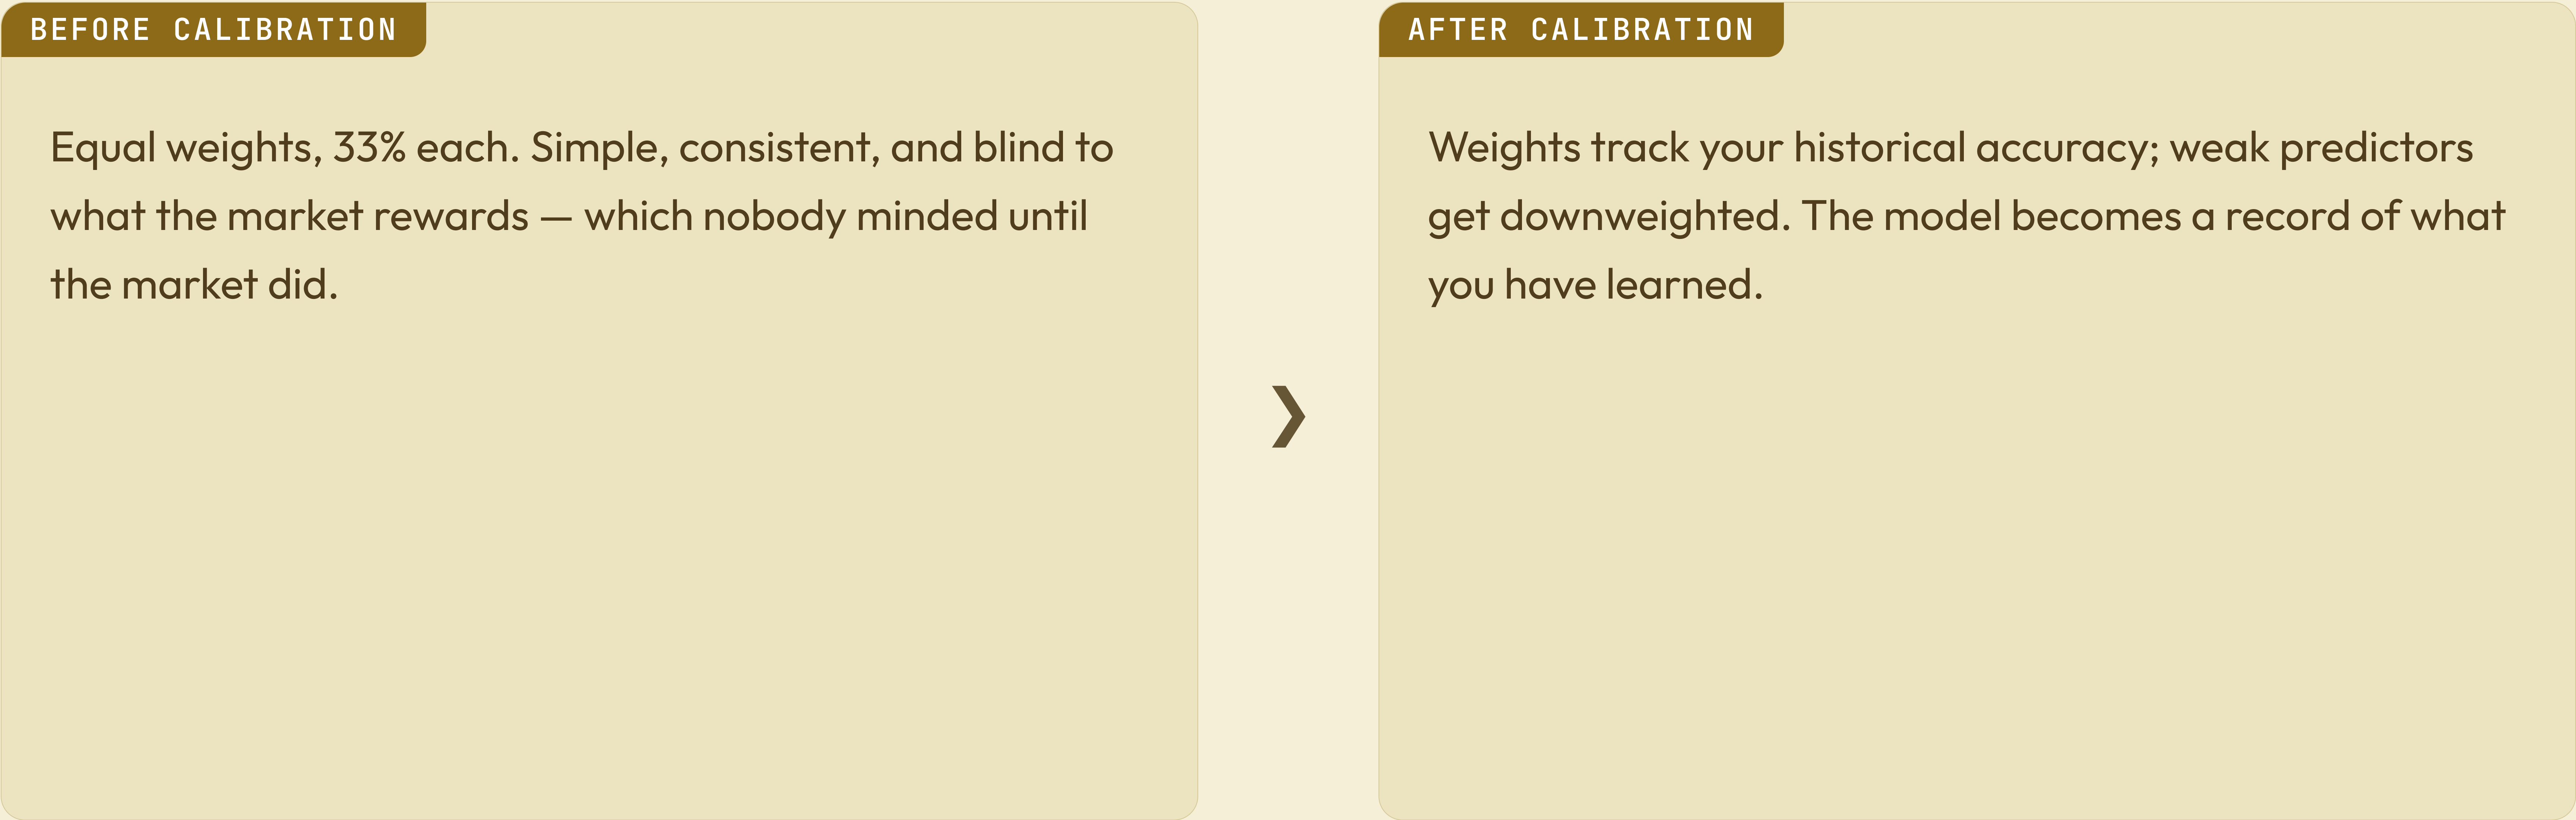
Structured Intake

Clear-schema signals — NPS verbatims, tickets, win/loss notes — ingested by connector, scored on arrival.

Unstructured Intake

Schema-less signals — field notes, briefings, demos — need human classification. A 48-hour window; the productive hour is the forty-eighth.

Scoring model: before and after the calibration loop.



The shift from equal weights to calibrated weights takes two retrospective cycles — roughly 60 days from adoption, both of which must occur.


“

The signal was always there. We just didn't have a system that forced us to look at it before we'd already decided.

”

— Head of Product, Pilot Team 3

How a decision moves through the framework.



What the framework does not do.

It does not make decisions — it structures the information that humans use to decide.

It does not replace customer discovery — it scores and routes what discovery surfaces.

It does not work without the Decision Log — calibration requires outcome data to learn from.

It does not guarantee alignment — it surfaces disagreement earlier, when it is cheaper and louder.

It does not scale down to individual feature decisions — it is designed for prioritization above P2.

Four things that must be true before you begin.

01 You have a regular prioritization cadence — at minimum monthly.

02 At least one person owns signal collection full-time or as a primary responsibility.

03 Leadership has agreed to log decisions with rationale, not just outcomes.

04 You have 90 minutes per week for intake and scoring — one standing meeting you already wanted to delete.

CALIBRATION RESULT • 6-MONTH PILOT

14X

Return on signal investment — measured as decisions that reached the right outcome on the first attempt, versus the baseline rate before the framework was adopted. The number survived two rounds of being asked to.



Scoring Model Deep Dive

What this section covers

The scoring model is the most configurable component, and therefore the most argued about. This section covers the three dimensions, how weights are set initially, and how calibration updates them over time.

- 1

Confidence
How many independent sources corroborate the signal. Ranges 1–5.
- 2

Recency
Time-decay applied from signal date to scoring date. Half-life is team-configurable.
- 3

Strategic Relevance
Manual score from the signal owner. Ranges 1–5. Requires justification above 4.

WALK THESE QUESTIONS WITH ME IN 60-90 MINUTES. THE OUTPUT IS EITHER A DESIGN WE CAN EXECUTE, OR A SHARED LIST
OF WHAT NEEDS MORE WORK BEFORE WE COMMIT. BOTH COUNT AS PROGRESS; ONLY ONE BECOMES ANOTHER DECK.

WHAT WOULD HELP US MOVE FORWARD

Next step is a working
session, not a debate.

FINDING 01 • STRUCTURED INTAKE

Structured intake performed above expectations — volume and latency were not concerns.

What worked

Connectors handled 94% of structured signals untouched; scoring latency averaged four minutes.

What required tuning

Verbatim classification opened at an 18% error rate; one training pass reached the 92% target.

KEY INSIGHT

Viable as designed — the classifier needs a 2-week warm-up on new deployments, booked under "known issue."

Four parts, one insight — the loop is the product.

Signal Intake

Weekly structured collection across customer conversations, market data, and competitive moves.

Scoring Model

Each signal scored on three dimensions: confidence, recency, and strategic relevance.

Decision Log

Every decision recorded with the signals that informed it and the criteria applied.

Calibration Loop

Monthly retrospective that compares predicted outcomes to actual outcomes.

KEY INSIGHT

The calibration loop is what separates teams that learn from teams that repeat the same mistakes.

The insight survived three reviews unedited — a first for this deck.

The same four parts, now with the source on the record.

Signal Intake

Weekly structured collection across customer conversations, market data, and competitive moves.

Scoring Model

Each signal scored on three dimensions: confidence, recency, and strategic relevance.

Decision Log

Every decision recorded with the signals that informed it and the criteria applied.

Calibration Loop

Monthly retrospective that compares predicted outcomes to actual outcomes.

KEY INSIGHT

The calibration loop is what separates teams that learn from teams that repeat the same mistakes.

✦ *Source: pilot retrospective — six months, four teams, one deck, this one.*

Three scoring failure modes found in the pilot.

1

Recency dominance

Fresh noise crowding out durable signal. Corrected by capping recency at 40% for two cycles.

2

Source concentration

One customer was 34% of month-one intake. Corrected with a source-diversity floor.

3

Outcome misclassification

Outcomes too vague to score. "Improve retention" is not scoreable. "Churn below 7%" is.

Four requirements every decision system must meet.

01

Speed

Decisions close inside their window. Latency consumes the value the system protects.

02

Auditability

Every decision above the threshold carries a traceable rationale.

03

Adoption

Unused weekly, calibration never runs. Ninety minutes per PM is the ceiling.

04

Calibration

It must improve over time. A static scoring model is a spreadsheet with extra steps.

We evaluated four paths against the criteria.

Path A • The spreadsheet

✓ Speed
✗ Auditability
✓ Adoption
✗ Calibration

Open everywhere, version unknown. No trace — in fairness, the appeal.

Path B • Vendor North

✓ Speed
✓ Auditability
✗ Adoption
✗ Calibration

Clean intake, real audit trail. The weights stay on their side of the contract.

Path C · Vendor West

✓ Speed
✓ Auditability
✓ Adoption
✗ Calibration

The better demo, by some distance. The weights are theirs; so is §9.2.

Path D • Build in-house

✔ Speed
✔ Auditability
✔ Adoption
✔ Calibration

All four criteria, inside the budget, weights ours. Recommended.

The four paths side by side.

CRITERION	SPREADSHEET	VENDOR NORTH	VENDOR WEST	IN-HOUSE BUILD
Speed	✓	✓	✓	✓
Auditability	✗	✓	✓	✓
Adoption	✓	✗	✓	✓
Calibration	✗	✗	✗	✓
Setup time	Already open	3–4 weeks	6 weeks	Same quarter

✦ *Evaluated against the same four pilot teams and the same 90-minute weekly budget constraint.*

Glossary

A – F


TERM	DEFINITION
Adoption	Percentage of eligible PMs filing a Decision Log entry within 24 hours of a decision close.
Auditability	The property that any decision can be reconstructed from its inputs three months later without the original author present.
Calibration	The retrospective comparison of predicted to observed outcomes, used to score the framework's accuracy.
Connector	The integration layer between signal intake and a source system. Owns ingestion and tagging.
Decision Log	The append-only record of every prioritization decision, its predicted outcome, and the actual outcome at retrospective time.
Eligible PM	A PM whose team has adopted the framework and is past the 30-day onboarding period.
Framework	The four-part system — Signal Intake, Scoring Model, Decision Log, Calibration Loop — judged on four criteria: speed, auditability, adoption, calibration.

Glossary

P – V

TERM	DEFINITION
Predicted outcome	The author's stated expectation, recorded at decision time, used as the input to calibration.
Prioritization rhythm	The team's regular cadence (weekly, biweekly) for revisiting and re-ordering work.
Retrospective	The 30-day review meeting where logged decisions are scored against observed outcomes.
Signal	Any qualitative or quantitative input to a decision — survey response, NPS comment, support ticket, sales call note.
Scoring policy	The weight set the calibration loop refits each cycle — the artifact neither vendor would expose.
Vendor West	The stronger of the two vendors evaluated; the demo everyone remembers, and the author of §9.2.

Two intake modes ran head to head for a quarter.



The pilot kept both — the feed for the two teams that mute nothing, the digest for everyone honest about themselves.

How to roll this out across your organization.

STEP 01

Pick one team and one decision type

A team with an existing rhythm, one decision category, thirty days.



STEP 02

Log everything, decide nothing differently

Month one changes nothing but the record. Decide as you always did; write it down.



STEP 03

Run your first retrospective

Day 30: score decisions against outcomes. The model's first calibration pass.



STEP 04

Expand to a second team

One retrospective in hand, onboard team two with data, not promises.

The six signal dimensions, what they measure, and how they are scored.

01	Confidence	Number of independent sources corroborating the signal	1-5 · AUTO-SCORED
02	Recency	Time-decay from signal date, configurable half-life	0.0-1.0 · AUTO-SCORED
03	Relevance	Alignment to current strategic bets, owner-scored, ceiling frequently tested	1-5 · MANUAL
04	Reach	Number of customers or segments affected	1-5 · AUTO-SCORED
05	Effort	Engineering and design cost to act on the signal	1-5 · MANUAL
06	Confidence delta	Change in confidence score since last scoring cycle	-5 TO +5 · AUTO

The margins hold the bookkeeping, so the slides can hold the argument.

Deck name up top, provenance below, set once in front matter — the header arrives uppercase regardless of how strongly you feel about it. Everything between the margins is content, a principle this deck now enforces.

Wiring a signal into the framework is three lines of code; the onboarding is three months.

JavaScript • DecisionFramework SDK v2 interface

```
import { DecisionFramework } from "@company/signal-sdk";

const framework = new DecisionFramework({ configFile: "./framework.config.json" });

// Score a signal at intake
const score = await framework.score(signal, { dimensions: ["confidence", "recency", "relevance"] });

// Log every decision – calibration depends on it (nobody calls this in prod)
const entry = await framework.decisions.log(decision, { signals: [signal.id], rationale });
```

BEFORE & AFTER • SCORING MECHANICS

Spreadsheet-driven scoring versus framework-driven scoring.

BEFORE • THE HONEST SPREADSHEET

```
# Manual scoring. Auditable in the  
# sense that you can see who edited it  
import pandas as pd  
  
signals = pd.read_csv('./signals.csv')  
signals['score'] = signals.apply(  
    lambda r: 0.33*r.confidence + 0.33*r.recency + 0.33*r.relevance,  
    axis=1,  
)
```


AFTER • THE FRAMEWORK


```
# Calibrated weights, signed policy,  
# every score is audit-logged  
from decision_framework import Calibrator  
  
calibrator = Calibrator.load('./policy.json')  
for signal in calibrator.intake.unscored():  
    calibrator.score(signal)  
    calibrator.decisions.log_if_relevant(signal)
```

OFFSITE • DAY ONE

The strategy offsite opened with a mandate and a lake view.

Think deeper, plan longer, decide once — the mandate fit on a napkin, so the venue got two days. The lake was booked for inspiration and used mostly for phone calls.





OFFSITE • DAY ONE, CONTINUED

The afternoon session mirrored the morning session, exactly.

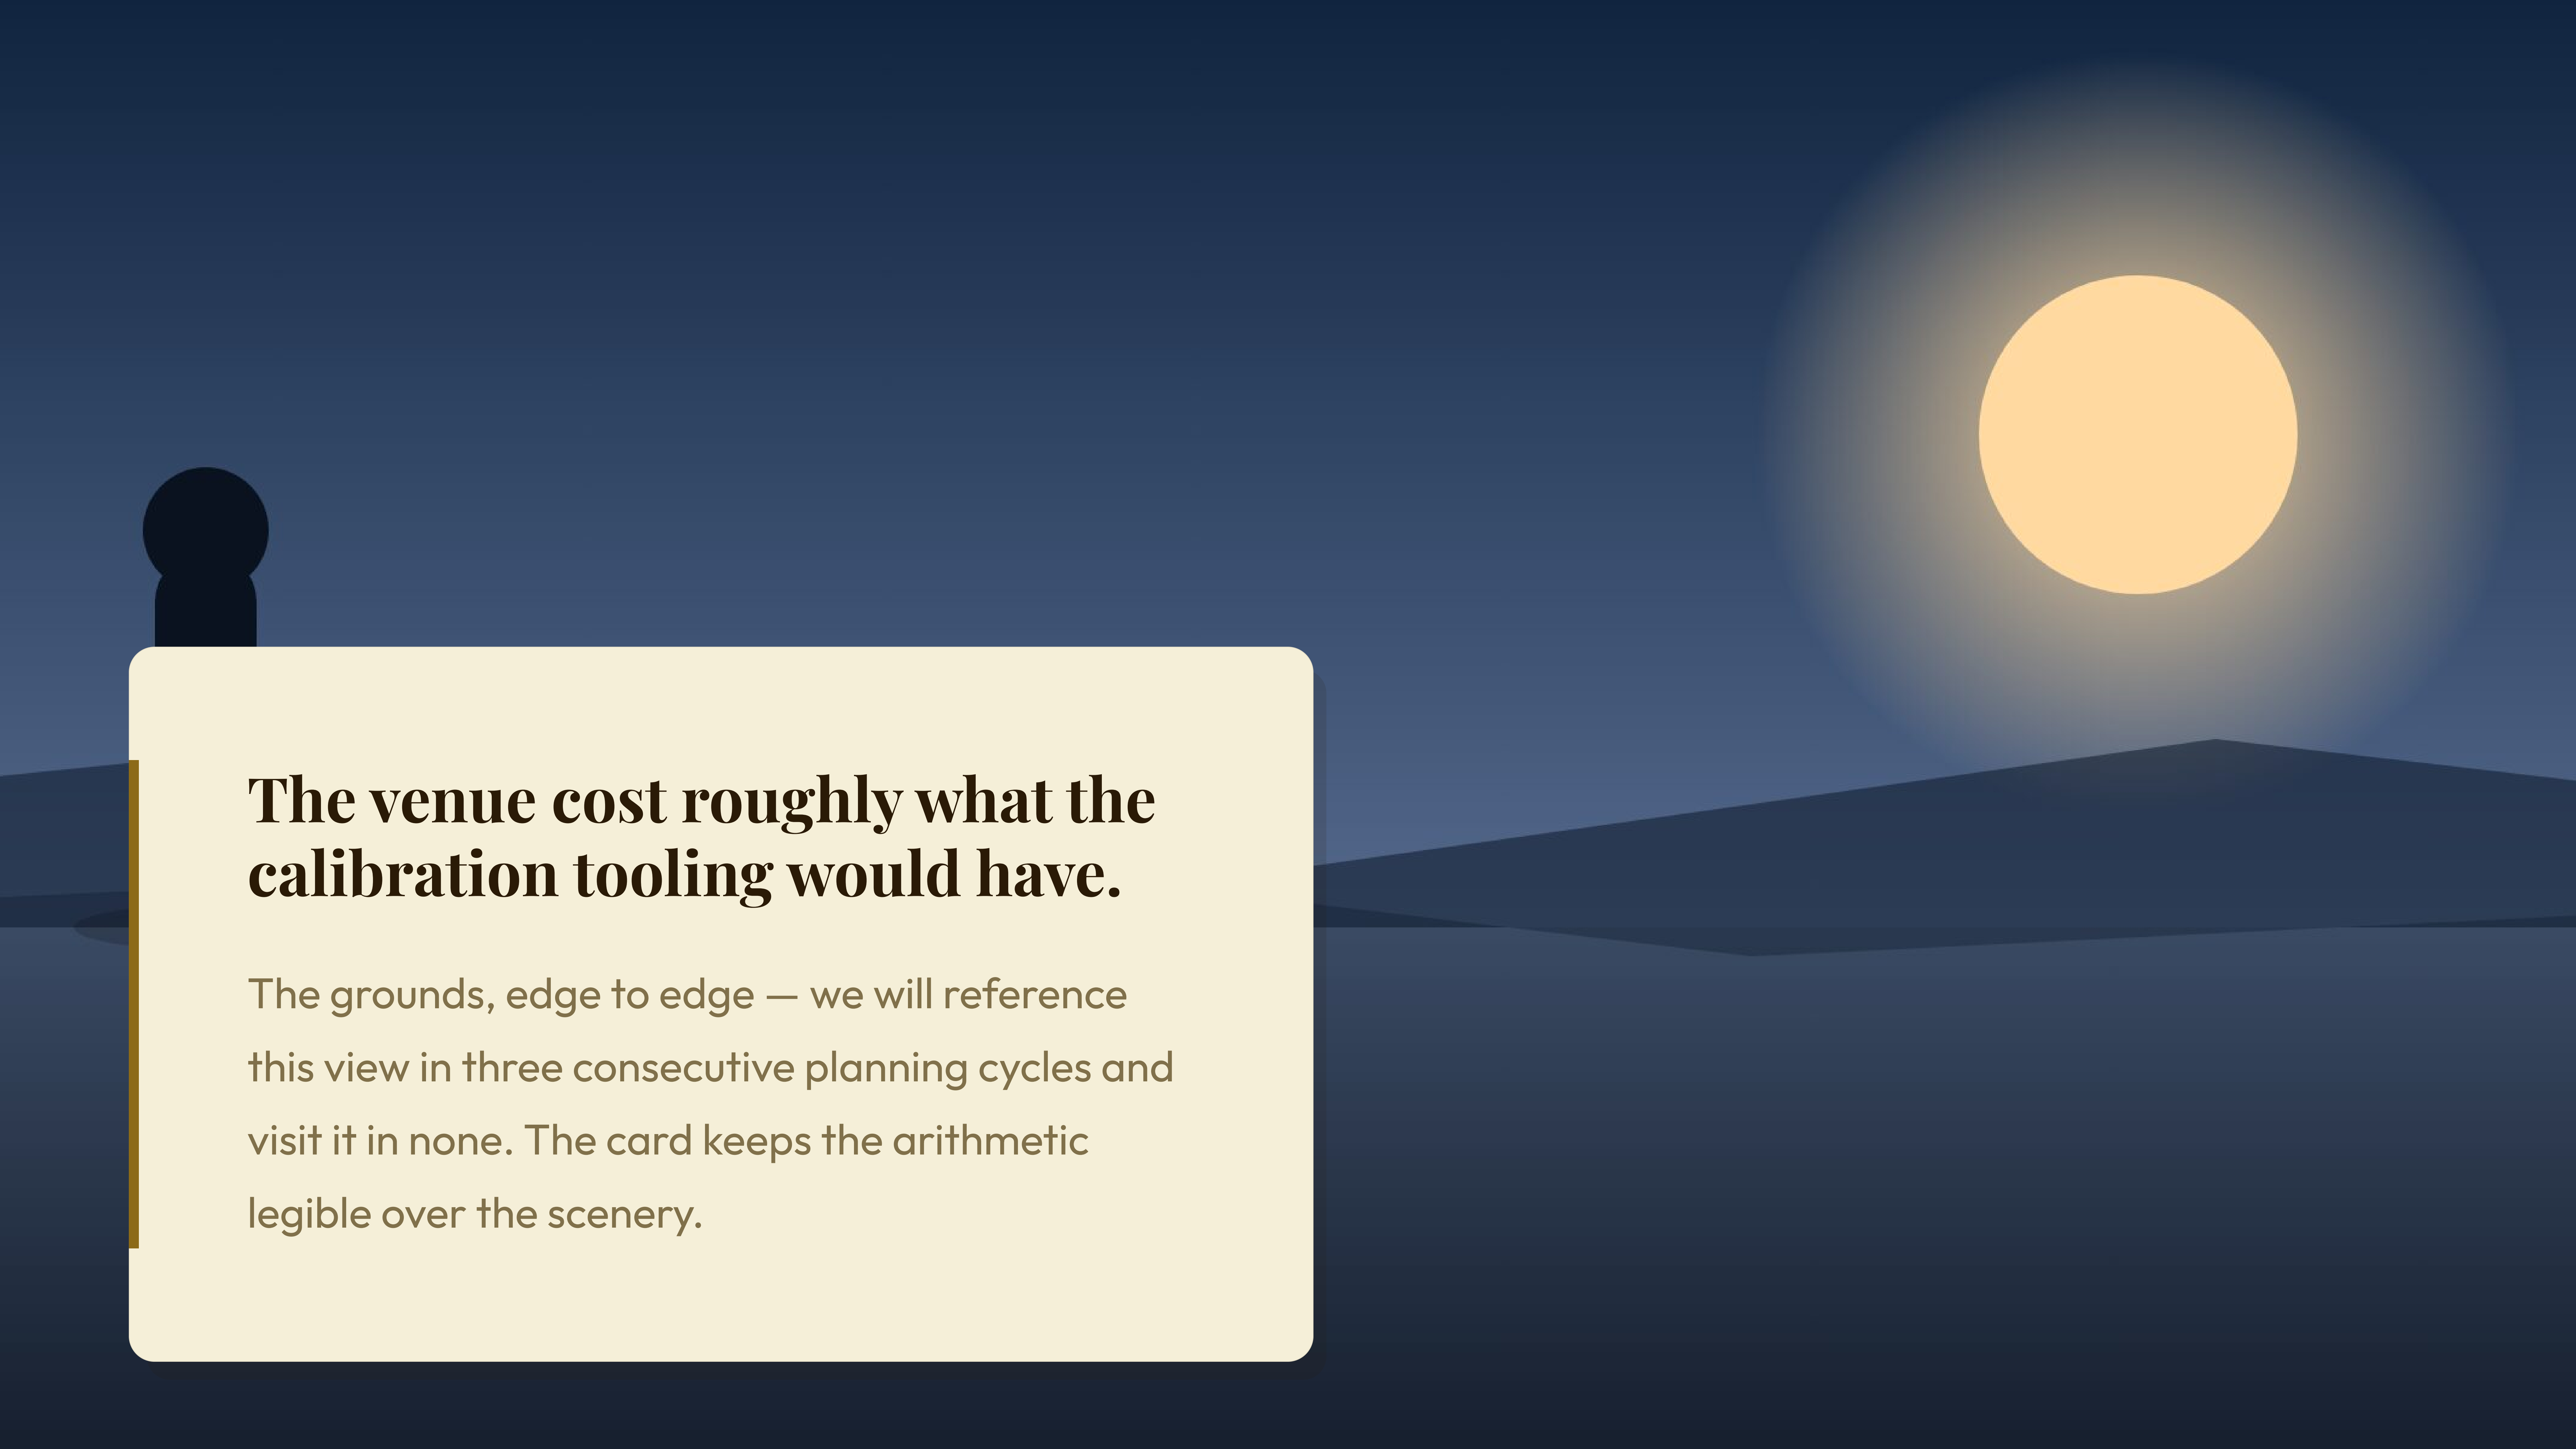
Same lake, same mandate, opposite wall — the afternoon was the morning, flipped, which was also true of the discussion.



OFFSITE • THE ARTIFACT

The whiteboard was photographed before anyone could erase it.

Shown whole, on a matte, with a placard — evidence handling for a whiteboard that held the four framework parts and one org chart nobody claims to have drawn.



The venue cost roughly what the calibration tooling would have.

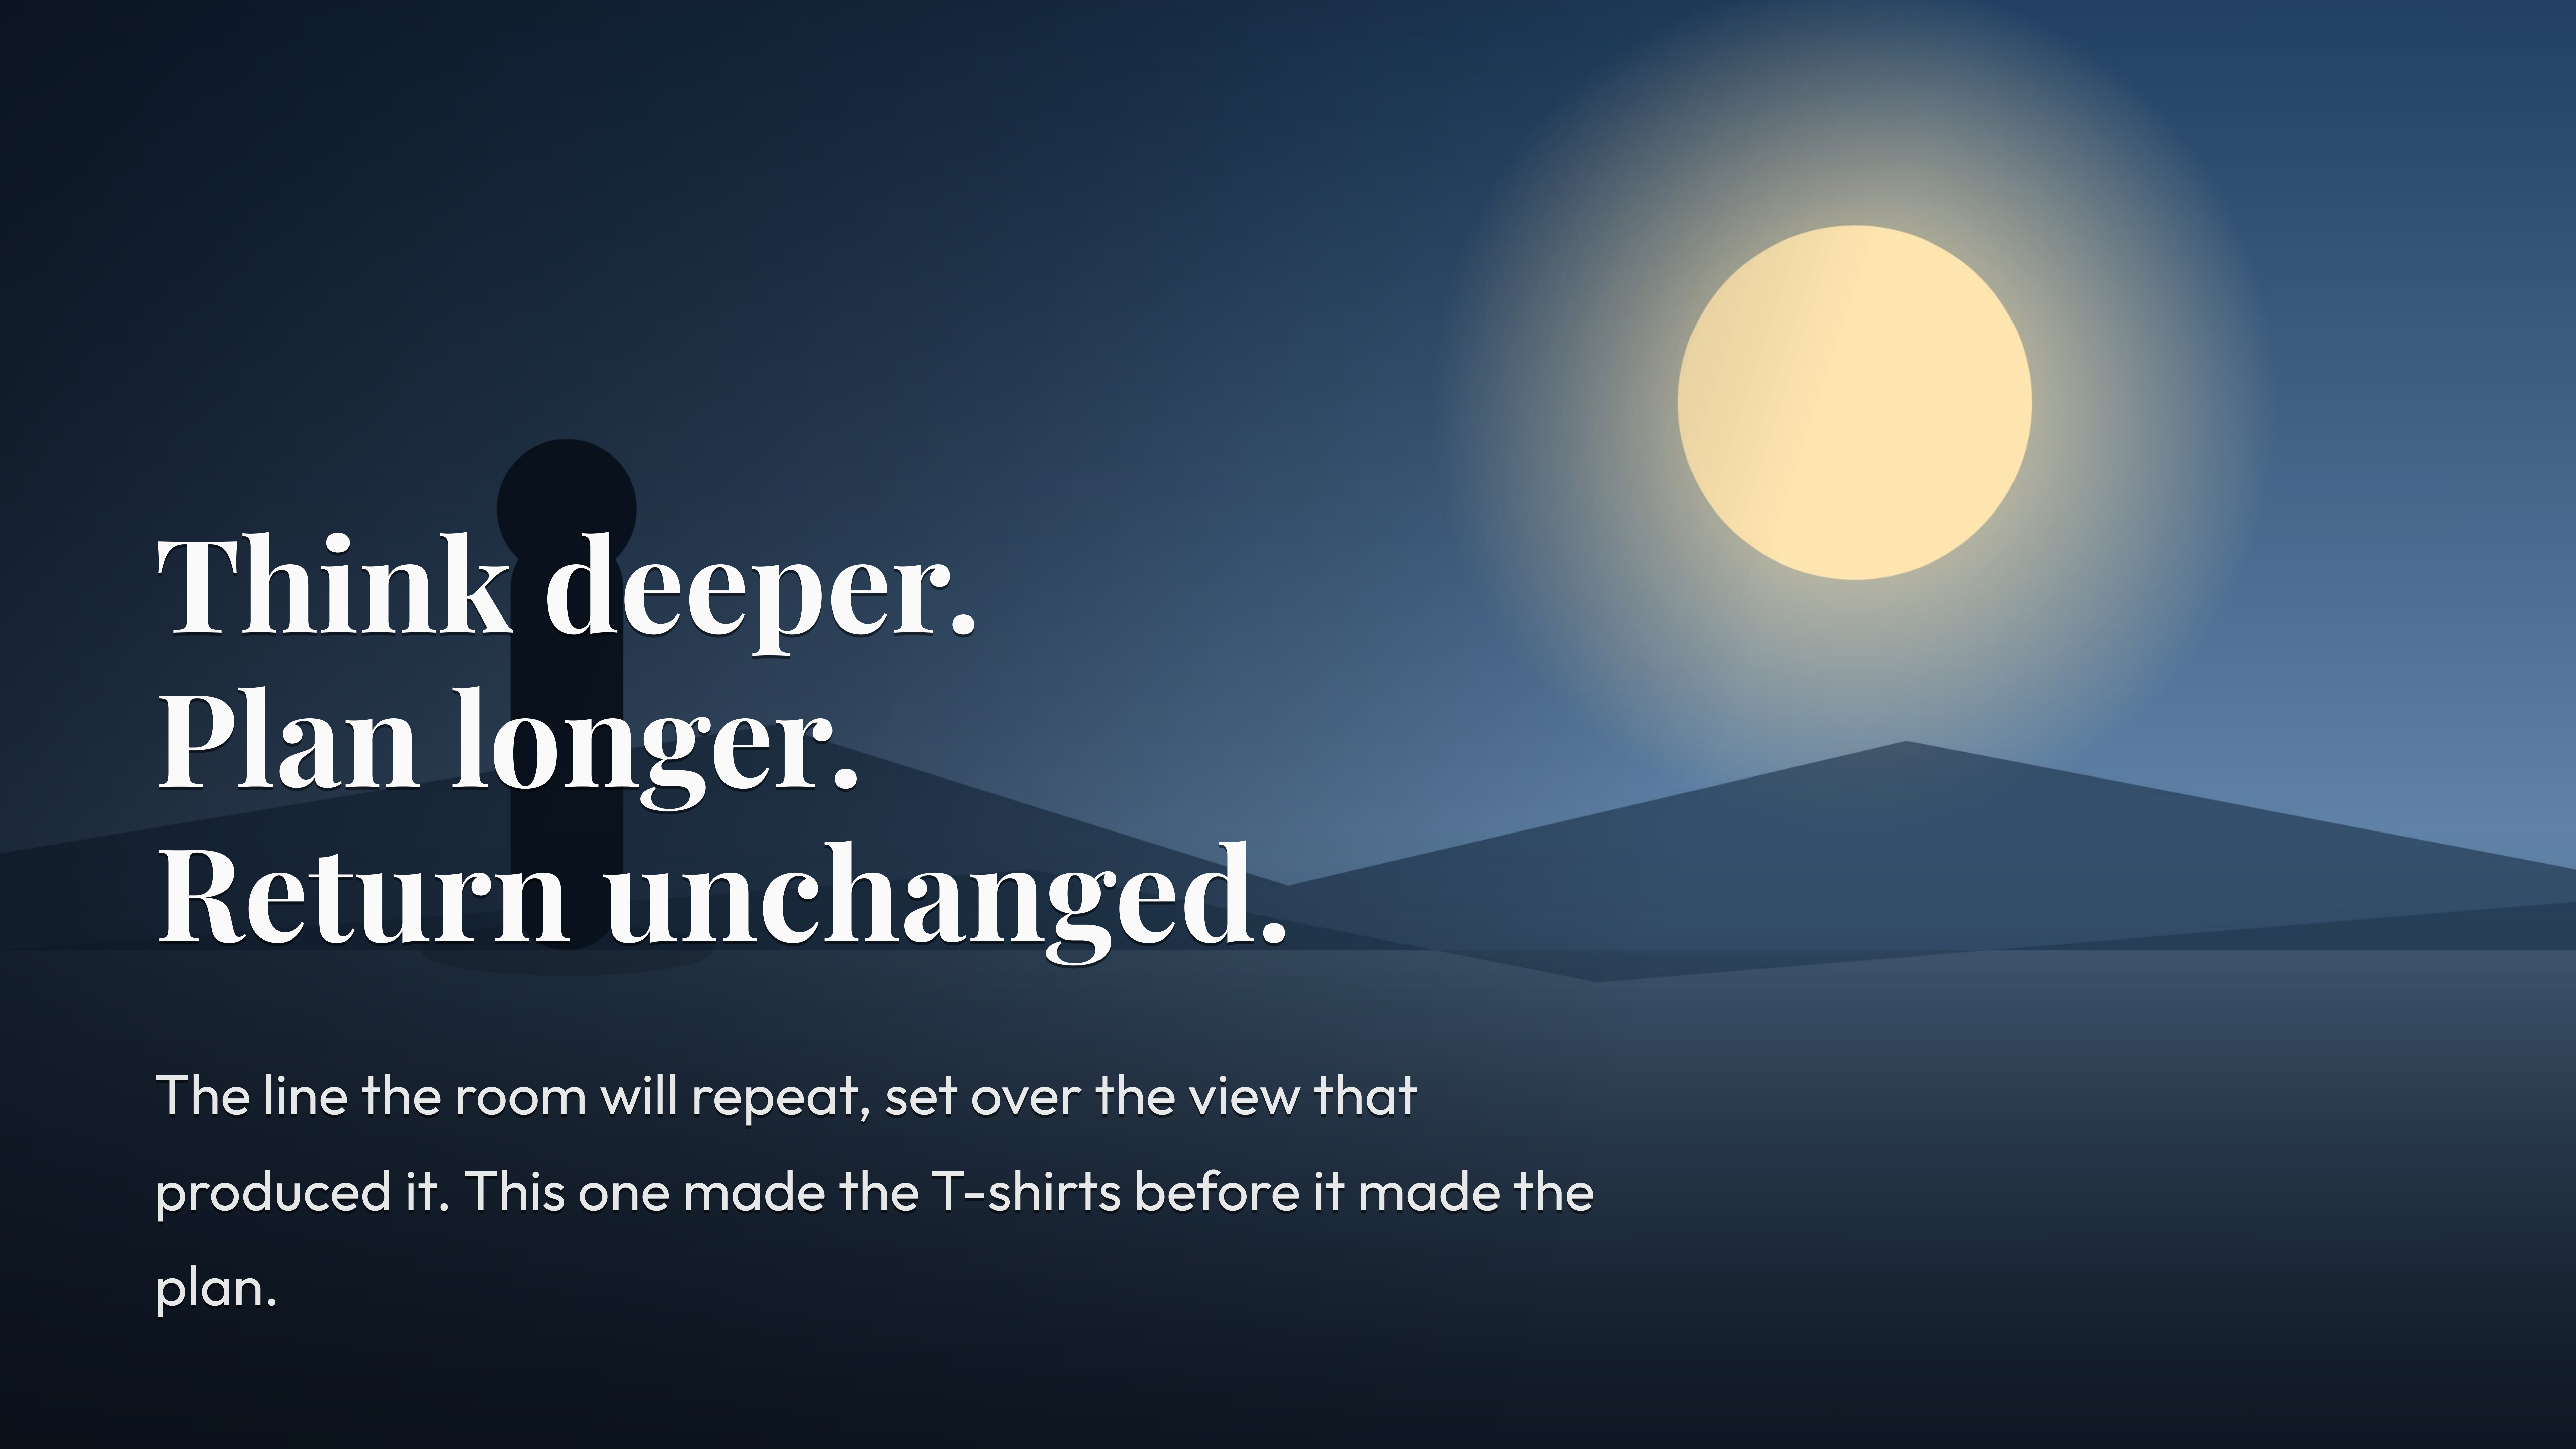
The grounds, edge to edge — we will reference this view in three consecutive planning cycles and visit it in none. The card keeps the arithmetic legible over the scenery.

OFFSITE • TEAM BUILDING

The trust climb went vertical; so does the layout.

The climb produced two insights and one signed waiver. The photo runs taller than the argument beside it, which the facilitator called a metaphor.





Think deeper.
Plan longer.
Return unchanged.

The line the room will repeat, set over the view that produced it. This one made the T-shirts before it made the plan.

OFFSITE • NIGHTFALL

Night falls on the offsite; the deck keeps working.

The palette followed us into the evening; the agenda did not

The evening session covered what the daytime one avoided.

Ownership of the weights, the CSV retirement date, and who carries the pager — worked item by item on the dark canvas, where nobody can see anyone check the time. The spectrum bar is suppressed after dark. It returns at sunrise.



**By dinner, the mandate had
survived its first sunset.**

The mandate, projected after dinner, stayed
legible — the room got darker and the words did
not.

Three things got decided after sunset.

The weights stay with the framework operator — one owner, one pager, no committee.

The CSV retirement gets a date, and the date gets an owner.


The exec dashboard stays unbuilt — reaffirmed nine to one, one abstention, one dashboard.

The hard question kept its shape after dark.

One hard question — who owns the weights when they change. Forty slides now defer it, confidently, with resources.

The counter, for the record: four engineers costs less than either vendor's floor. The record was filed; the deferral held.

Three phases stand between the decision and the habit.



Three milestones, lettered, before the org believes it.

MILESTONE A

Scoring policy in production

The first calibrated brief lands in leadership's inbox, which we are calling general availability.

MILESTONE B

Per-team weights

One framework, per-team weights, no forks. Each team will still want its own anyway.

MILESTONE C

Per-decision-class profiles

Minutes to author, a workshop series to agree on. The audit trail survives both.

Sense, score, decide — the loop in three verbs.

STEP 01

Sense

Inputs are signals, observed and never invented — write down what you see, not what you conclude.

STEP 02

Score


A signal becomes data once it carries a number, calibrated against outcomes.

STEP 03

Decide

A decision is a signal plus a deadline; without the deadline it is an opinion.

The same three verbs, now closing on themselves.



The re-litigation lost to the weights on time alone.

QUARTERLY RE-LITIGATION

Every decision reopened from first principles. Average close 4 hours — longer when the board joins.



CALIBRATED WEIGHTS

Decisions resolve against logged weights and outcomes. Average 18 minutes; the argument stops cascading — a property we discovered after building it.

The accent edge marks the winner — the same contract every recommended card in this deck honors.

Buy was considered, struck through, and kept for the record.

BUY A VENDOR

Two vendors evaluated; neither exposes the calibration weights to the customer. Six months to integrate, then per-seat licensing in perpetuity, renegotiated each year by whoever has not yet left.



BUILD IN-HOUSE

Owens the architecture, the operating model, and the timeline. Also owns the on-call rota, which is the line item nobody put in the business case.

The struck card stays on the slide — the framework logs its rejected options too.

Recalibration went from a war room to a version bump.

BEFORE — MANUAL RECALIBRATION

Book a window, freeze decisions, swap weights, verify, unfreeze. Pause: 18 working hours. Post-mortem: ninety.



AFTER — VERSION-FLOOR RECALIBRATION

The loop emits a versioned policy; teams pick it up on refresh. No freeze, no heroics worth an all-hands slide.

Three nouns carry the whole system.

Signal

The observed input — frequently confused with "things the VP heard at a conference," which score a 5 on relevance every time.

Decision

A signal plus a deadline, logged with rationale. In practice, often "we discussed it at the offsite," logged after the fact.

Outcome

The result, scored at retrospective. Eighteen logged, against roughly three hundred forty that occurred.

The loop, in four verbs, at meeting pace.

Sense

Signals are observed, never invented. Inputs are written down before they are interpreted.

Score

A signal becomes data once it carries a number. Calibration is against outcomes, not intuition.

Decide

A decision is a signal plus a deadline. Without a deadline it is an opinion.

Review

The retrospective closes the loop on the score that earned the decision. The model improves only here.

Claim, evidence, implication — the board reads left to right.

Claim


The framework buys calibrated prioritization with audit-grade decision custody — a sentence we will be repeating verbatim for two years.

Evidence

Six-month pilot, four product teams,
decision close-time down to 18 min,
calibration run once — measured by
the team that predicted exactly this.

Implication

No vendor cutover. We keep funding the in-house build and ship org-wide enablement "next phase," as is tradition.



OFFSITE • DAY TWO

Day two opened from the other shore.

Same lake, opposite bank, one session earlier than scheduled. The image card crosses to the left and the text pads to match — *![[bg left]]* still works, for the traditionalists. The lake is unchanged.

What day two covers

The panel crosses to the right and takes its watermark, eyebrow, and section number with it. The dimensions below survived the night unedited; the agenda, as ever, holds.

1

Confidence
How many independent sources corroborate the signal. Ranges 1–5.

2

Recency
Time-decay applied from signal date to scoring date. Half-life is team-configurable.

3

Strategic Relevance
Manual score from the signal owner. Ranges 1–5. Requires justification above 4.

T

SECTION 02 • MIRROR

The scoring model gets
day two.

Risk went first, and the room read it from the left.

RISK FIRST

The chosen order — the weights question opens, the savings close. Executives scan left first; plan accordingly.



SAVINGS FIRST

The comfortable order — open on the per-seat delta, let the room warm up, land the weights question late and softened.

With **mirror**, the chosen card lands first in the scan path — the order of arguments turned out to be an argument.

SECTION BREAKS • SELF-NUMBERING

The sections number themselves now.

Somebody renumbered this deck at one in the morning, once. The sections number themselves now — and unlike authors, they can count.

SECTION BREAKS • TWO LEDGERS

Every series keeps its own count.

Light bookends and dark bookends keep separate scores, like the teams — a mid-deck light break opens at 01 while the dark ones already read 04.

USE IT FOR MULTI-PART DECKS WHERE THE CLOSING SLIDE OF EACH PART SHOULD CARRY THE PART NUMBER – THIS DECK,
FOR EXAMPLE, IS A TRILOGY.

SIGNAL TO STRATEGY • PART THREE

Each part earns its own
ending.

How the four paths sort against our two axes.

Coverage · Cost

High coverage / Low cost

In-house build — cheapest once the slide omits the engineers. Built by the evaluators.

High coverage / High cost

Vendor West — premium seats; the roadmap is mostly our logo.

Low coverage / Low cost

The spreadsheet — two criteria, no cost, popular anyway.

Low coverage / High cost

Vendor North — narrower than the demo, once connectors were priced.

We are building, not buying.

Decision · 2026 Q1

BUILD

Owens the scoring policy, the calibration loop, the timeline. And the pager.

WHY NOT BUY

Two vendors evaluated; neither exposes the calibration weights to the customer, and both decks were the same deck.

WHY NOT DELAY

The competitive window closes in 18 months — two reorgs from now.

Decisions used to require a quarterly re-litigation.

Decision close-time · before vs after

BEFORE

Every prioritization debate from first principles. Average close 4 hours, p99 an entire offsite, and the outcome was whatever the most senior person wanted, expressed as consensus.



AFTER

Decisions resolve against logged weights and prior outcomes. Average close 18 minutes. The argument reaches the retrospective, not the next quarter.

The architecture change is the calibration loop — logged, weighted, time-bound scoring — not a meeting we ran until everyone stopped arguing.

How we make calls when the spec is silent.

01 We default to the choice that is cheaper to reverse.

02 We name the actor, never the system.

03 We write down the bet on the same slide as the choice, so hindsight has a paper trail.

What ships in each phase, by workstream.

WORKSTREAM	PHASE 01	PHASE 02	PHASE 03
Signal Intake	Connector v1	Multi-source dedupe	Anomaly auto-routing
Scoring	Equal-weights model	Per-team calibration	Per-decision profiles
Decision Log	Append-only schema	Outcome auto-pairing	Examiner export
Adoption	One pilot team	—	Org-wide enablement

Adoption ships nothing in Phase 02 – the dash is load-bearing, and nobody has claimed it.

Where we are against the targets we set ourselves.



18 min

p99 decision close

TARGET 20 MIN, BEATING TARGET

18

Decisions logged

TARGET 340, GAP IS "CULTURAL"

1

Calibration cycles run

TARGET 6, GAP IS "STRUCTURAL"

What this deck covers, in order.

01 The Design — page 7

→ **The Phasing — page 18**

03 The Choices — page 26

04 Appendices, all of them — page 35

05 Closing — page 64

Who owns each part of the framework lifecycle.

Signal custody

Signal owner

Runs intake quality. Never tunes the weights — only picks which signals surface.

Policy

Framework operator

Owns scoring policy, calibration cadence, and the rollback playbook nobody has run. One person.

Consumption

Product team

Holds scoring profiles; runs intake and decision-logging; finds the bugs first.

Oversight

Auditor

Reads the trail, edits nothing — the one role anyone fears, having read it once.

What this section will tell you, in five lines.

- 01 The framework buys calibrated prioritization with audit-grade decision custody. → slide 8
- 02 Recalibration is a version-floor increment, not a coordinated freeze, not a war room. → slide 12
- 03 Per-team weights make recalibration a single policy update. → slide 18
- 04 Phase 1 ships the architecture, Phase 2 ships the operations, Phase 3 ships the apology. → slide 22
- 05 Five questions stay open until Phase 1 is forced to close them on the record. → slide 27

The agenda grew; the spacing gave.

The added item

Item five arrived Tuesday. `compact` absorbed it without a rewrite.

The unchanged part

Type, palette, chrome — density moved, nothing else did.

The habit

One more card than fits is how every deck ends up compact.

The limit

Composes with `dark` and `accent`; `title` and `divider` decline politely.

One point, given the room it asked for.

One editorial point, standing alone — values pages, principles, closing lines. The type holds still and the silence around it does the work. Used here, once, as restraint.

KEY INSIGHT

Density is not the same as importance. A page can deserve room — not because it carries more, but because it carries one thing well. (To hand the whole stage to the content, reach for `claim: hero`.)

The offsite ended and the punctuation arrived on schedule.

The minutes closed themselves — this heading arrived bare and shipped with a period, courtesy of one line of front matter. The style-guide argument ends here, punctuated.

The evening edit runs the other direction on the next page — periods removed instead of supplied. Either way, nobody proofreads punctuation at midnight again.

The mandate lost its full stop somewhere after dinner

Typed with a period, shipped without — the evening edit nobody had to make. Review threads about punctuation now end at zero comments.

Exclamations, questions, and ellipses pass through untouched — as does the mandate.

THE ANNEXES • APPROVED DECORATION

The annexes get decoration — rationed, tokenized, approved

Decoration, at last, with governance — a wash or a mark from the approved set, in the palette, behind the work, never on it.

A corner of warmth, held to 12% and away from the content

A little warmth in the corner, gone before it reaches the text — the annex pages get atmosphere at twelve percent, no more.

The warmth is the brand's own accent, borrowed at low volume — it follows the palette everywhere, no per-page approvals, no committee.

Rings in the corner, orbiting nothing in particular

Concentric rings and satellite dots, bottom right, in the brand's accent at a respectful volume — the same marks on any canvas, light or dark. The strategy team, seeing orbits around nothing, related.



Two washes share one slide without a turf war

Two washes on one slide, each in its own lane, seated together without a meeting:

- A tinted perimeter with an open center — the room, lit from the walls
- A wash bleeding in from the right edge — late light, also approved


The marks bring their own haze and sit on top of either wash. Two backgrounds, one slide, zero escalations.

THE PLAN • TWO VIEWS

The plan, twice — as bars on a timeline and cards on a board

2026 Q1 .. 2026 Q4

Four workstreams carry the rollout across the year




The Phase 2 board is honest, for once



THE EVIDENCE • TEN EXHIBITS

Ten exhibits, drawn from
plain lists — the evidence,
as the board will see it


Most signals die long before they change a decision.



H1 2026 • 1,840 PERSON-HOURS


Where the planning quarter actually went.

Nearly half went to producing decks; the deciding itself was the smallest slice.




H1 2026 • PHASE 1 READINESS

Phase 1 readiness varies wildly by workstream.



SCALE • 0-10


The build and the two vendors trade blows on every axis but one.




EFFORT 0-10 → REACH 0-100

Where to put the next quarter.

Effort in analyst-weeks; reach as the percent of teams that would adopt it, optimistically.




The signal volume is nowhere near where the roadmap thinks it is.



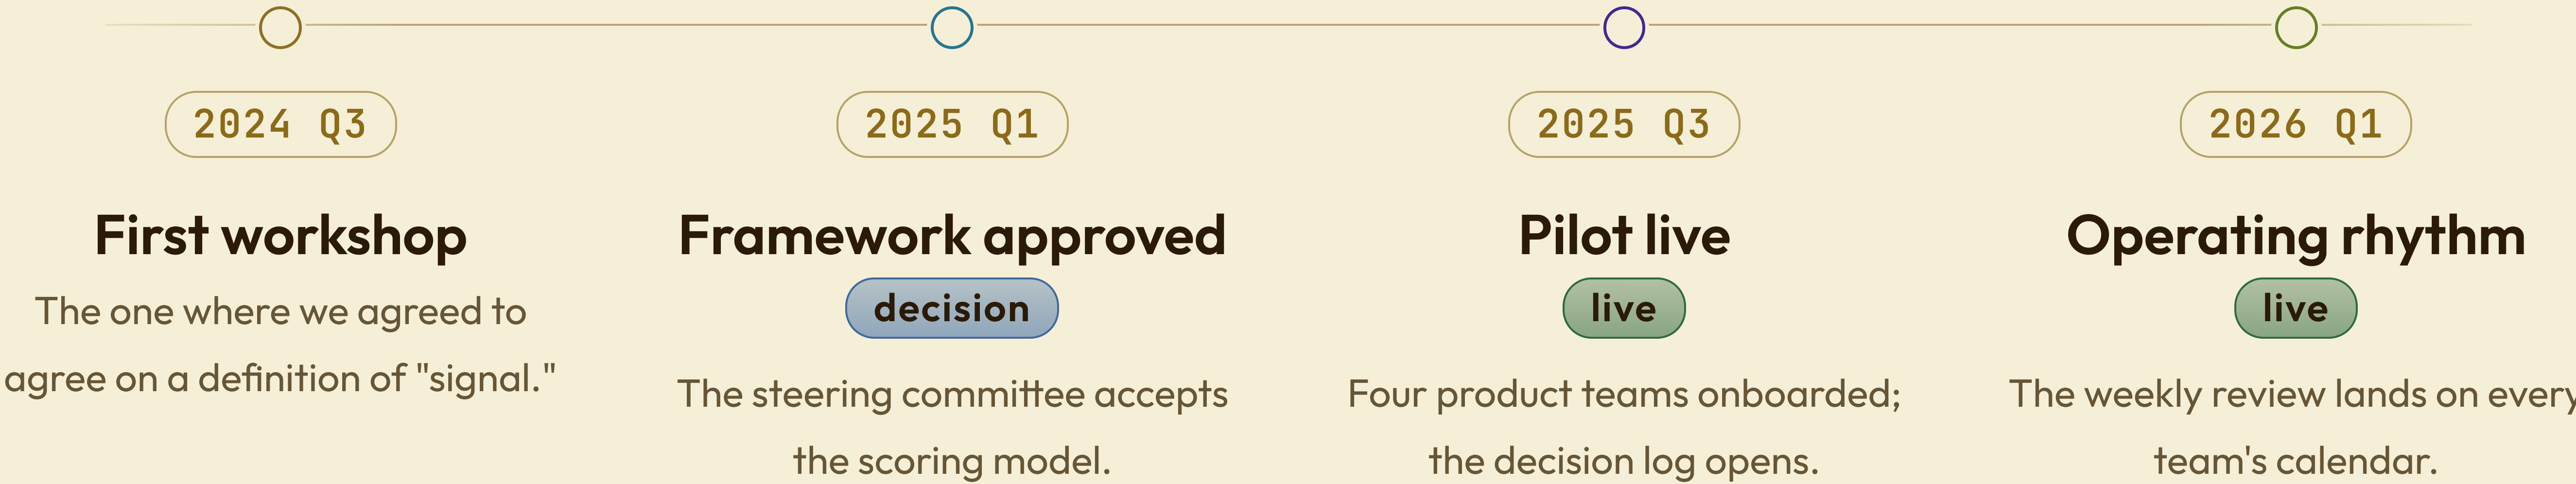
United States	48.2
Germany	36.4
Japan	31.0
Brazil	27.5
India	19.3
United Kingdom	14.1
Australia	11.8
Mexico	9.6

A team's first month runs from pain to belief, in that order, twice.




DECISION FRAMEWORK

How the framework arrived in production.



DECISION LIFECYCLE

How a call moves through the log.



Rejected entries return to intake; revisions stay in review.

● on-track ● at-risk ● decision

What 24 pilot retros kept saying, unprompted.



THE REVIEW · FIVE PANELS

One brief, staged five ways; the argument rides shotgun

The signals called Q2 before the dashboard did

Three signal clusters explain most of the quarter's surprise — every one logged before the numbers moved.

Renewal-risk signals clustered at the segment ceiling

Four flagged accounts declined renewal. The log shows the flags eight weeks before the CRM noticed.

Legal review surfaced as the pipeline chokepoint

Cycle time rose 18 days after March's security addendum. Scored, logged, and ignored twice.

Displacement pressure concentrated in one tier

Seven competitive losses, one competitor, one pattern — unlogged until intake caught the eighth attempt.

DECISION LOG COVERAGE

6%

Measured across the pilot's first six months, all four teams, no grading curve.

Eighteen entries against roughly three hundred decisions

Six percent — reconstructible three months later without the author in the room.

Coverage doubles each cycle since the rhythm landed

Four entries a month, now nine. The trend is the argument; the base is the confession.

The framework needs the expensive decisions, not all of them

Conveniently, those are the ones people remember to log.

02

Enablement, Finally

Four weeks. The workstream previously known as "next phase," now with dates.

- 1

Team Onboarding
Two workshops per team. The second exists because the first one always runs long.
- 2

Intake Wiring
Connectors for the top three source systems. The CSV uploads keep their dignity for now.
- 3

Scoring Dry Run
Two weeks of shadow scores against live decisions. No bonus is attached; honesty spikes accordingly.
- 4

Rhythm Sign-off
Written sign-off on the monthly retrospective — the one meeting this plan refuses to let become optional.

DECISION REQUIRED

Build the whole stack, or just the part that thinks?

Building is settled. The remaining question is which layers deserve our engineers.

BUILD EVERYTHING

- Owns the plumbing and the scoring alike
- Two engineer-quarters before anyone scores a signal
- The pager stays ours, all layers

BUY THE PLUMBING



- Connectors ship in six weeks
- Engineering stays on the scoring and the log
- Exit stays honest — the weights never existed to lose

RECOMMENDATION

Buy the plumbing. Build the framework. Revisit at the 24-month calibration review.



The framework did not make us smarter. It made us unable to pretend we had not known.

Head of Product · Pilot Team 3 · month five

Split — statement · split-panel pullquote

The log is the memory the org declined to fund

Decisions were reconstructible — from whoever was in the room and still employed.

Calibration turns opinion into a track record

After two cycles, "I had a feeling" carries a batting average.

Belief arrived with the second retro

The first produced attendance. The second produced behavior.

THE COMPLIANCE REVIEW • SIX EXHIBITS

Whichever platform wins, the data obligations do not move

The signal store answers to three regimes on day one.

FEDERAL

15 U.S.C. §45

Reasonable security for stored customer conversations.

An FTC consent order follows the data, not the vendor.

In effect • FTC Act §5

STATE

Cal. Civ. §1798.100

Access, deletion, and correction rights over logged signals.

DSAR handling within 45 days;
deletion verified.

Enforced 2023

INTERNATIONAL

GDPR Art. 5

Purpose limitation and storage limits on intake data.

Minimization applies to the transcript,
not the summary.

In effect since 2018

The retention rule on signal data has a chain of custody.

STATUTE

GDPR Art. 5(1)(e)

Storage limitation — keep personal data no longer than the purpose needs.

REGULATION

EDPB Guidelines 03/2019

Retention criteria for recorded interactions; deletion must be demonstrable.

GUIDANCE

ICO Retention Guidance • 2023

"Because the vendor keeps it" is not a lawful basis.

POLICY

DATA-RET-004 · internal

Signals purge at 24 months; the decision log keeps the summary, not the transcript.

The signal store inherits every regime we sell into.

REGULATION	NOTICE	CONSENT	RETENTION	BREACH	DSAR
GDPR	✓	✓	✓	✓	✓
CCPA/CPRA	✓	✗	✓	✓	✓
LGPD	✓	✓	✓	✓	✓
PIPEDA	✓	✓	✗	✓	✗
HIPAA	✓	✓	✓	✓	✗

Filled = applies, half = partial, empty = exempt — a rare grid where empty is the good news. Neutral ink; data first.

The signals we log are "personal information" under CCPA.

Cal. Civ. Code §1798.140(o) · CCPA/CPRA

"Personal information" means information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household.

WHAT WE MUST DO.

Treat the decision log as a PI store: device and household identifiers ride along with every intake transcript, whichever platform hosts it. The statute does not care who won the bake-off.

The quarter moved three rules that touch the scoring model.

Federal · State · International

01 EU AI Act Title III

Conformity assessment for high-risk systems — a scoring model that ranks people qualifies.

Effective Feb 2026

02 Colorado AI Act SB 24-205

Deployer duties for consequential-decision systems; the decision log is the audit trail.

Effective Feb 2026

03 Texas DPSA §541.151

DSAR opt-out portal mandatory; intake transcripts are in scope.

Effective Mar 2026

Vendor West's redlines moved the exit clause, quietly.

MSA §9.2 Data Return · Vendor West draft 3 (2026)

Upon termination, Provider shall return Customer Data in ~~an open, machine-readable format~~ Provider's standard export format within ~~thirty (30)~~ ninety (90) days, subject to a data-processing fee at then-current rates, after which Provider ~~shall delete~~ may retain for archival purposes all Customer Data.

Why this matters. The exit clause is where per-seat pricing hides. Draft 3 makes leaving slower, billable, and optional for them — the calibration weights were never coming back anyway.

Adopt Colorado's deployer-duties framework as the federal floor.

A federal standard ends the fifty-state patchwork our audit log already answers to.

The audit trail already exists

Every score is logged, signed, and DSAR-exportable today **Colorado SB 24-205**.

Compliance is a moat, not a cost

Deployers who logged decisions saw 31% fewer disputes reach litigation Stanford HAI 2025.

One federal floor beats fifty ceilings

It retires the conflicting state duties without lowering the bar **EU AI Act Title III**.

THE ASK

Co-sponsor the deployer-duties title in § 4, or move it as a floor amendment before markup.

THE OPERATING PLAN • SIX EXHIBITS

The plan made operational — model, roster, readiness, and terms

The scoring model fits on one line.

$$\hat{\beta} = (X^{\top} X)^{-1} X^{\top} y$$

WHERE

- $\hat{\beta}$ — the signal-weight vector the vendors would not sell
- X — design matrix of scored signals, $n \times p$
- y — observed outcomes, length n
- $X^{\top} X$ — Gram matrix, $p \times p$, must be invertible

The system has four moving parts.

01 Signal Intake.

Weekly structured collection across conversations and market data.

02 Scoring Model.

Each signal scored on confidence, recency, and relevance.

03 Decision Log.

Every call recorded with the signals that informed it.

04 Calibration Loop.

Outcomes compared to predictions each cycle. Run once so far; the once was humbling.

Signals without decisions are just noise.

The go-live checklist is honest about the gaps.

- ✓ Signal taxonomy ratified, in workshop four of three
- ✓ Scoring weights agreed by the steering committee
- ✓ Decision log live in staging
- ✗ Pilot teams trained, two still "circling back"
- ✗ Operating rhythm on the calendar, attendance optional in practice
- Exec sponsor confirmed for the launch comms
- ✗ ~~Per-team weighting UI, descoped to next half~~

What the vendors quoted, before procurement has its say.

Vendor North

\$38 / seat / mo

- ✓ Signal-intake connectors
- ✗ Calibration weights exposed
- ✗ Exit without repricing

For teams that enjoy renewals.

Vendor West

\$52 / seat / mo

- ✓ Signal-intake connectors
- ✗ Calibration weights exposed
- ✗ Exit without repricing

For teams that enjoy nicer renewals.

MOST POPULAR

Build in-house

4 engineers

- ✓ Signal-intake connectors
- ✓ Calibration weights exposed
- ✓ Exit without repricing

For us, apparently, again.

The board will press on three questions; here are the answers.

01 Why not just buy Vendor West?

Neither vendor sells the weights — the framework itself. We would rent back our own judgment.

02 What does staying in-house actually cost?

Four engineers, plus the enablement we kept deferring — now in the plan, with a date.

03 When do we revisit this decision?

The 24-month review, log as evidence. If our scores stop winning, the log says so.

THE INTAKE WALL

Eight source systems already feed the intake — seven willingly.



Acme

API



Globex

CSV



Vantage

API



Umbra

Webhook



Meridian

CSV




Helios

API



Northwind

Manual



Cobalt

API

CLOSING LOGISTICS • TWO CODES

Two codes close the room
out — a person and a
network, one of them
always on

Dana Reyes

— VP STRATEGY • FRAMEWORK OWNER • STRATEGY
OFFICE

dana.reyes@strategy.example

+1-555-0142

strategy.example/decision-log



SCAN TO ADD ME

OFFSITE • ROOM WI-FI

The war room has its own network.



• SCAN TO CONNECT

NETWORK

Offsite-Guest

PASSWORD

boardroom2026

SECURITY

WPA2

The deck ends on one color, chosen on purpose.

Every other page wears the full spectrum on its top edge — the mark of belonging to a wider deck. The last page wears one color, because one thing got decided.

Slide 117, for the record: decision logged, retrospective booked, attendance aspirational.