

# Kashi 可視 — business plan

Grounded in the 22-memo research library. Honest ROI math. Japan-first.  
Consultation-compatible by architecture.

2026-04-21 · Companion to </deck.html> and </walkthrough.html> · Research library at [§13](#)

**v2 changes from v1:** ROI math rebuilt with  $A \times C \times I \times E$  discount framework (prior version was overstated). Expansion sequence flipped to Japan-first doctrine (not EU-first). Buyer framework split into three actors. Risk register expanded from 6 to 11. Added procedural-fairness moat, non-dilutive funding track, meeting-type normalization, and remediation-outcome metrics. Rhetorical-discipline pass: replaced "pattern is the harm" / "detect harmful dynamics" / "CEO's instrument" with epistemically-honest alternatives per the measurement-science and category-entry memos.

## IN THIS DOCUMENT

1. 01 TL;DR — what Kashi is, honestly
2. 02 Market — hidden-loss portfolio, not one big number
3. 03 Go-to-market — three buyers, two rhetorical registers
4. 04 Revenue model +  $A \times C \times I \times E$  unit economics
5. 05 Team + founding story
6. 06 Financing — equity + non-dilutive in parallel
7. 07 Moat — procedural fairness is the spine
8. 08 Risk register — 11 risks, explicit
9. 09 Operating model (with Japan data-residency fix)
10. 10 Strategic options — Japan-first expansion
11. 11 Traction — detection, remediation, recovery
12. 12 Rhetorical doctrine — claims we make and refuse
13. 13 Research library (24 memos)

## 01 TL;DR — what Kashi is, honestly

Kashi is a **privacy-bounded meeting governance system** that surfaces repeated interaction asymmetries — under explicit uncertainty — as contestable structural signals for human review, inside strict procedural and use limits.

We are not a harassment detector. We are not a coaching dashboard. We are not a culture-survey tool. We are the governance layer that makes repeated meeting-level patterns **institutionally legible and procedurally contestable** — so organizations can no longer plausibly claim they didn't see, and so the people inside them can no longer silently carry the cost.

### The four conditions that make Kashi buildable now

- 1. Meeting transcription is ambient.** Teams, Meet, and Zoom ship transcripts by default for the first time.
- 2. EU AI Act Article 5 (Feb 2025) clarified the red line.** Workplace emotion inference is prohibited. Structural metadata is permitted. Kashi is designed inside that line.
- 3. The problem is persistently measured but unmoved.** MHLW 令和5年度: 19.3% experienced power harassment, 36.9% took no action, 53.2% of companies did nothing. The duty structure exists; visibility and response do not.
- 4. The market learned what NOT to build.** Microsoft Productivity Score (Dec 2020) stripped individual telemetry within 7 days of public exposure. That precedent defines the category's failure mode — and our product constraints.

MACRO LOSS (JAPAN)

**¥7.6T/yr**

~1.1% of GDP · ¥7.3T

presenteeism vs only ¥0.3T  
absenteeism · Yokohama CU 2025

SAM (JAPAN)

**~45,000**

50-500-person knowledge-work  
companies

3-YR SOM TARGET

**500**

paying customers · ~1.1% of SAM

Y3 ARR TARGET

**¥2.5B**

~\$17M ARR · 85% Pro / 15%  
Enterprise mix

GROSS MARGIN TARGET

**78%**

blended · Pro 85% / Enterprise  
65%

EVIDENCE MOAT

**procedural**

detectors are table stakes;  
contestability state machine is the  
defense

**The business thesis in one sentence.** The hidden-loss pool is large, measurable, and unmoved. Existing products (content classifiers, pulse surveys, productivity scores) have failed because they either overclaim, underclaim, or cross the legitimacy line. Kashi is the first product whose *architecture and language are engineered to survive labor-relations scrutiny* — the gate every prior attempt has failed at.

## 02 Market — hidden-loss portfolio, not one big number

### The economic case has been built wrong in prior versions

Earlier decks led with "¥7.9M per leave case × N cases = ¥X saved." That math is too coarse. The rebuilt-ROI memo makes it precise: you must apply an  $A \times C \times I \times E$  discount stack — **Addressable share × Coverage × Intervention uptake × Effect size** — to go from total loss pool to realistically captured loss. We rebuild the pitch accordingly.

### Four cost buckets a CFO can believe

BUCKET	WHAT IT CAPTURES	EVIDENCE ANCHOR
<b>1. Presenteeism</b> (biggest)	People at their desks underperforming while carrying unresolved interpersonal weight. Invisible on the headcount dashboard.	Yokohama CU 2025: <b>¥7.3T of ¥7.6T loss</b> is presenteeism, not absence. Deloitte UK 2024: £24B of £51B is presenteeism.
<b>2. Regrettable attrition</b>	Replacing lost talent when the loss was avoidable. Manager-driven departures cost ~200% of salary to replace; technical ~80%, frontline ~40%.	Gallup 2025: managers account for <b>70% of team engagement variance</b> . Global manager engagement fell to 22%.
<b>3. Formal-escalation premium</b>	The 3x multiplier on cost once a conflict passes from informal to formal resolution. Legal fees, time-in-HR, distraction, reputation.	Acas (UK): formal procedures cost <b>3x informal resolution</b> . £28.5B/yr UK conflict cost.
<b>4. Leave-case tail</b>	The visible end-stage: 健康保険傷病手当 claims, temp backfill, training-replacement. Late and late-stage.	METI 2025: ¥9M+ per year-long leave at ¥6M salary; ~¥7.9M at ¥5.25M salary. But: <b>this is tail risk, not base case</b> .

**The honest hero line.** The ¥7.9M-per-case headline is late-stage illustration, not base-case math. The defensible business case is: *the hidden-loss pool exists, it's measurable, and Kashi shortens the time that a manager-linked pattern stays invisible and unmanaged.*

## Hidden-loss sizing for a 500-person JP company

Assumes ¥5.25M average salary, 5% productivity drag attributable to harmful team dynamics, 30 employees on leave in a year.

SCENARIO	TOTAL ANNUAL LOSS POOL	NOTES
Strict / employer-cash floor	¥20.2M	What even a hostile CFO must accept
<b>Base case / operational-economic</b>	<b>¥127M</b>	Presenteeism + overtime spillover + coverage drag
Upper / aggressive	¥205M	All buckets loaded

## What Kashi realistically captures (A×C×I×E)

Per the rebuilt-ROI memo: multiply total loss × A × C × I × E where A = addressable share of the pool we can touch, C = coverage we achieve, I = intervention uptake, E = effect size on those who acted. Realistic ranges for a mature deployment:

SCENARIO	A × C × I × E	TOTAL REDUCTION	CAPTURED SAVINGS
Strict	15 × 60 × 30 × 20%	0.54%	¥0.69M/yr
<b>Mid (plan)</b>	25 × 70 × 40 × 25%	<b>1.75%</b>	<b>¥2.22M/yr</b>
Upper	40 × 80 × 50 × 30%	4.80%	¥6.10M/yr

These are captured-loss numbers for a 30-leave-case/year company. Translate per-customer: a 500-person company paying ¥4.8M/year for Professional tier lands at 0.5–1.3× direct ROI on leave-linked savings alone – **which is why leave is not the sell.** The sell is the portfolio (presenteeism + attrition + escalation premium) where Kashi's value is 4–8× higher per saved case.

## TAM / SAM / SOM

### TAM (GLOBAL)

1.5–2M companies worldwide with 50+ employees in knowledge-work sectors and a labor regime that rewards preventive governance. Blended ACV ¥3M/yr → **¥4.5–6T/yr ceiling**. We don't pitch this number. SAM is what's real.

### SAM (JAPAN)

JP enterprises 50–500 employees × knowledge-work filter = **~45,000 companies**. Blended ACV ¥4.2M/yr (85/15 Pro/Enterprise mix) → **~¥190B/yr revenue ceiling** in Japan alone.

### SOM (3-YR)

Year-3 target: **500 paying customers**. ~1.1% of SAM. Conservative for a governance-compliance category (typically saturates at 5–10%). Y3 ARR: **~¥2.5B**.

## 03 Go-to-market — three buyers, two rhetorical registers

### The buyer is not one person

The founder-narrative and labor-politics memos converge: Kashi has **three buyers**, each with different language needs. Treating the CEO as the only buyer is how this category's prior attempts died.

#### BUYER 1 · ECONOMIC SPONSOR

**CEO / CFO / COO**. Needs the portfolio math (hidden loss, regrettable-attribution exposure, formal-escalation premium). Signs the contract. Commits to the 3-lane accountability architecture.

#### BUYER 2 · OPERATIONAL GATEKEEPER

**HR / Legal / Compliance / CHRO**. Needs the procedural-fairness spine (contestability, bounded context, retention-under-challenge, meaningful human review, rollback triggers). Clears the CAIQ/SIG-lite security review.

#### BUYER 3 · TRUST COUNTERPARTY

**Worker representative / union delegate / employee-rep / works council**. Needs the anti-capture + bounded-visibility story. Signs off on 就業規則 amendment. Their consent is architectural, not ceremonial.

**The two-register rule.** Sell like a CEO-grade risk product to Buyer 1. Roll out like a worker-protective bounded-monitoring system to Buyers 2 and 3. Mix these rhetorical registers in the wrong order and the deal dies. The CEO slide ("bill before it arrives") belongs in the closed-sponsor conversation — *never* in the worker-facing rollout deck.

## Channel stack — priority order for Year 1-2

#	CHANNEL	SHARE OF PIPELINE	NOTES
1	Direct outbound to JP CEOs / CFOs of mid-market knowledge-work	30%	Target: post-IPO, post-recent 労基署 inspection, post-public resignation. Message: hidden-loss portfolio, not "surveillance."
2	Labor-law firm referral partnerships	20%	Oh-Ebashi LPC, Mori Hamada, Anderson Mōri. 10% first-year rev share × 12mo.
3	<b>Worker-representative + union engagement</b>	15%	NEW — Rengo, Zenroren affiliate unions, works-council networks. Co-credentials the tool; unblocks Buyer 3.
4	Academic partnership (NAQ-R validation)	10%	津野香奈美 (Kanagawa Univ. of Human Services) or 慶應 productivity-health. Long-cycle credibility play.
5	JP VC portfolio intros	10%	DNX, Globis, ANRI, WiL — 5-10 warm intros per seed check.
6	Content + research publishing	5% direct, 100% ambient credibility	The 22-memo library is substrate for ~2 articles/month. 日経

#	CHANNEL	SHARE OF PIPELINE	NOTES
			Business, Tech In Asia, note.
7	<b>Occupational-health channel (JOHAS / EAP partners)</b>	5%	NEW — reframe Kashi as pre-consultation visibility layer for occupational-health intake.
8	Ombuds + neutral-review associations	5%	IOA (International Ombuds Association) JP chapter. Critical for future Ally/Observer Concern Path rollout (V2).

### Target customer archetypes (Year 1 pipeline)

1. Post-Series-C JP SaaS (100-250 emp, growing management layer, 1-2 パワハラ incidents/yr)
2. Mid-market manufacturing with recent 労基署 inspection (compliance budget already authorized)
3. Regional-bank subsidiaries / financial mid-market (high reputational exposure)
4. JP consulting firms, Big-4 offshoots (client-facing, meeting-heavy)
5. JP offices of EU-HQ multinationals (AI Act pressure pre-authorized)
6. Post-IPO firms with founder-to-professional-CEO transitions (friction moment)
7. University spin-offs with active research collaboration appetite
8. Regional tech clusters (福岡, 京都 D2C, Osaka IoT)
9. Mid-market with existing 健康経営 certification (HR-forward already)
10. Parent companies under SASB / TCFD human-capital disclosure pressure

### Sales cycle per tier

TIER	CYCLE	DECISION PATH	PILOT STRUCTURE
Starter	Self-serve, 1-week activation	Team lead signs up, 20-employee cap	No pilot — free forever at that scale

TIER	CYCLE	DECISION PATH	PILOT STRUCTURE
Professional	30–90 days	CEO + HR/Legal + worker-rep (3 signatures)	30-day free trial → annual
Enterprise	4–9 months	All 3 buyers + IT security review + CAIQ/SIG pack + works-council consultation	90-day <b>paid</b> pilot (1/6th year-1 contract) → 3-year

**Pilot-to-paid conversion target: 60%.** Three-role qualification pre-filters heavily, so pilot entry is already a strong buying signal. Floor 40%.

## 04 Revenue model + AxCxIxE unit economics

**Full pricing rationale.** The tier structure below was set in an earlier pass. The methodology behind it (value-based math, competitive benchmarks, Van Westendorp validation plan, bootstrap-phase founding-customer pricing, common mistakes to avoid) is now documented at </pricing-strategy.html>. Read the strategy doc before quoting prices to any prospect.

### Pricing (unchanged pricing, rewritten defense)

TIER	PRICE	CAP	BLENDED EXAMPLE ACV
Starter	Free	≤20 employees	¥0
Professional	¥800/user/mo	≤500 employees	¥1.44M @ 150 users
Enterprise	¥10M/yr base + ¥800/user/mo	500+ / multi-region	¥19.6M @ 1,000 users

**Defense of the ¥10M Enterprise base.** The rebuilt-ROI memo notes that ¥10M/month on leave-cost alone requires 94.4% capture — indefensible. So we don't defend Enterprise pricing on leave-cost alone. The Enterprise base covers: hands-on consultation (就業規則 review, rep-process sit-in), CAIQ/SIG answer pack, SOC 2 / ISO 27001 compliance artifacts, dedicated CSM, JP data-residency hardening, NAQ-R validation access. It's a governance-service fee, not a compute fee.

### Gross margin

## STARTER

**N/A**

lead-gen, cost absorbed

## PROFESSIONAL

**85%**

compute + email support

## ENTERPRISE

**65%**

CSM + consultation hours drag

## BLENDED Y3

**78%**

85/15 Pro/Ent mix

**CAC, payback, LTV**

TIER	CAC (TARGET)	ACV	PAYBACK	LTV (4-YR RETENTION)	LTV/CAC
Professional	¥400k	¥1.44M	3.3mo	¥4.9M	12x
Enterprise	¥4M	¥19.6M	2.4mo	¥51M	13x

**Net revenue retention (NRR) target: 115%**

- Expansion from headcount growth within customer (Pro scales per-user)
- Pro → Enterprise tier upgrade (~20% cross 500 emp in 2 yrs)
- Offset: ~8% annual logo churn (pilot non-conversion + restructures)

**3-year ARR build**

PERIOD	PRO	ENT	NEW ARR	CUMULATIVE
Q1 Y1	2 (pilots)	0	¥2M	¥2M
Q4 Y1	8	2	¥40M	¥52M
Q4 Y2	60	15	¥280M	¥380M
Q4 Y3	400	100	¥1.3B	¥2.5B

**05 Team + founding story**

The founder-narrative memo is direct: don't fabricate pedigree. The honest "why us" for Kashi is not resume-based; it's **disciplined category design**. The pattern of refusal — no affect, no

content, no HR decisions, no health score — is the answer to "why are you the right team?" That's defensible because it's rare and it's observable in what we've already built.

## Our honest founding story

We built Kashi because existing workplace-AI products fail a specific test: they look legitimate to whoever is paying, and illegitimate to whoever is being analyzed. The products that claim "we detect harassment" are the ones that should worry you most. The products that claim "we're just about productivity" quietly accumulate evidence that gets misused later.

We chose a different starting point. Start with the refusals. Make the refusals architecturally enforced, not contractually promised. Make the rhetoric match the architecture. Then build the smallest useful product inside those limits.

The 22-memo research library behind this plan is part of the answer to "why us": we did the reading before we wrote a line of code, and we wrote the code to survive the reading. The product at [kashi-lilac.vercel.app](https://kashi-lilac.vercel.app) is the 23-hour hackathon-scale proof that this discipline produces a shippable artifact, not just a whitepaper.

### FOUNDER BIO — FILL THIS IN when ready

Per the founder-narrative memo, keep this *epistemically honest*:

- Name, role, location
- What authority you have to do this work (lived experience? domain work? academic track? engineering discipline?)
- What you've already shipped (the public product + the 22-memo research library is the real answer)
- Co-founder or partner — same fields
- Do NOT overstate: no fake "former Head of X at Y." If the real story is "we did the reading and we shipped the constraint," that's stronger than a manufactured credential.

## First 3 hires (post-seed)

1. **JP enterprise salesperson** (month 1). Prior HR-tech / compliance-tech to mid-market JP. SmartHR, Kaonavi, Workday JP alumni profile.

2. **Security / DevSecOps lead** (month 2). Owns CAIQ/SIG-lite pack, IR plan, key lifecycle, Vercel-to-Supabase data-flow audit (see §9 operating model).
3. **Fractional labor-law advisor** (month 1, ¥500k/mo retainer). 労働弁護士 who co-writes consultation packet, reviews 就業規則 templates, co-credentials outbound.

### Advisory board targets

- **CHRO** of a 1000+ person JP company (credibility at enterprise table)
- **Org psychologist** with outcome-measurement track (津野香奈美 is the obvious first ask)
- **Retired 厚生労働省 / 労働局 official** (regulatory translation, guideline change forecasting)
- **Worker-representative / union leader** (permanent seat, compensated + voting — per the anti-capture memo this is non-negotiable)
- **B2B SaaS founder with JP exit** (fundraise + customer intros)

## 06 Financing — equity + non-dilutive in parallel

The funding-pathways memo identified a track most founders miss: Japan has multiple non-dilutive programs that Kashi fits — not as "mental health" (we're not eligible) but as **AI adoption** and **workstyle reform**. Run both tracks in parallel. Equity pays for velocity; grants pay for customer-side procurement subsidy that makes the enterprise sale cheaper.

### Track A — equity rounds

#### PRE-SEED · ¥50M (\$330K) · RAISE NOW

- Milestones: 10 paid Pro pilots (¥14.4M ARR), 2 Enterprise paid pilots, 1 labor-law firm partnership, NAQ-R LOI
- Dilution: 15-20% on SAFE at ~¥300M post
- Targets: DNX Ventures, Miyako Capital, Antler JP, Y Combinator (if Delaware route), angel CHROs

#### SEED · ¥500M (\$3.3M) · 12-15 MONTHS LATER

- Milestones for Series A: 75 paying customers (60 Pro + 15 Ent), ¥380M ARR, NAQ-R preliminary publication, 3 labor-law firm partnerships, team of 8

- Dilution: 18-22% at ¥2-3B post
- Targets: Globis Capital, ANRI, WiL, SmartHR exec-team angels

### SERIES A · ¥3B (\$20M) · YEAR 3

- Use: EU entry (NL/UK first, not DE — see §10), SOC 2 Type II achieved, ISO 27001 in progress, consultation packet localized per jurisdiction, team ~35
- Dilution: 20-25% at ¥15-25B post
- Metrics: ¥2.5B ARR, NRR >115%, GM >75%, payback <18mo

### Track B — non-dilutive (run in parallel)

PROGRAM	COUNTRY	FIT	TIMING
<b>Digitalisation &amp; AI Adoption Subsidy 2026</b>	Japan	P1 — high	Frame as AI tool for workstyle / governance. Subsidizes customer purchase, making enterprise deals cheaper.
<b>Workstyle Reform Promotion Support Subsidy 2026</b>	Japan	P1 — medium-high	Direct fit: Kashi supports preventive labor-governance under 労働施策総合推進法.
JOHAS occupational health routes	Japan	P1/P2	Channel partnership with occupational-health providers; pre-consultation visibility layer.
<b>PSG (Productivity Solutions Grant)</b>	Singapore	P1 — high	Singapore is the <i>first</i> international market (per the funding memo), not EU.
EDG (Enterprise Development Grant)	Singapore	P1	SME-facing co-funding for transformation projects.
DIGITAL / EDIHs / Apply AI Strategy	EU	P2 — medium	Pilot-legitimacy funding, not revenue.
OH / welfare-counselling tax routes	UK	P2 — medium	If Kashi bundles with EAP.

**Critical strategic choice:** whether to register Kashi as a formal "IT tool" under Japan's AI/digital adoption route. Doing so unlocks customer subsidies but *constrains pricing + packaging*. Decide before seed round — it affects contract structure and reseller logic.

## 07 Moat — procedural fairness is the spine

The legal-procedural-fairness memo reframes what our moat is. Detectors are table stakes.

**The moat is the procedural architecture around** the detectors — specifically, the contestability state machine, the bounded-context rule, retention-under-challenge, meaningful human review, downstream-use charter, and rollback triggers. Harder to copy than detectors. Much harder.

### Layer 0 · Architectural lane taxonomy (NEW — Ideas\_wave3 integration)

Every detector declares its lane at compile time in `src/lib/pipeline/detector-registry.ts`. Three structural detectors (intrusive-interruption, chilling-delta, floor-time Gini) touch no transcript text and default to employer-facing. Four hybrid detectors (unanswered-question, topic-credit, agreement-asymmetry, keigo) read transcript text and are gated behind a tenant feature flag that defaults OFF.

This replaces the blanket "metadata only, no content reading" claim (which was technically false given our embedding-based similarity detectors) with an honest, machine-readable taxonomy that procurement reviewers and legal counsel can audit. Adding a new detector requires updating the registry. Changing a detector's lane requires governance review.

### Layer 1 · Procedural-fairness spine (new — upgraded from governance sub-bullet)

- **Contestability state machine.** Detected → disputed → human review → upheld / downgraded / withdrawn. Every review-worthy event has a path.
- **Bounded context window.** Short turn window + speaker order only. Role-justified drill-down required to expand.
- **Retention-under-challenge.** Auto-delete unless dispute is active, case hold is triggered, or user-owned preservation is chosen. Not indefinite evidence accumulation.

- **Meaningful human review.** Reviewer trained, independent, disagreement-logged. No rubber-stamp.
- **Downstream-use charter.** Contractually + architecturally prohibited use in performance / promotion / discipline / compensation / ranking. Not "we promise" — "the export path does not exist."
- **Rollback triggers.** Pilot decommissions automatically if gate violations occur (access exceeds policy, content exported, scope creep detected).

## Layer 2 · Product refusals

The "Kashi will not do" list is the product. Viva can't copy it without dismantling Viva's value proposition. Archaic can't copy it without abandoning content-scanning. Adding our refusals to their products requires gutting their products.

## Layer 3 · Data moat (cross-industry baseline stack)

Baseline stack = self-history + within-meeting + meeting-type + role + dyad. After 500 customers, we can tell a 150-person SaaS "your keigo-asymmetry is 78th percentile for your sector and size" — not replicable without matching customer density.

## Layer 4 · Regulatory / compliance artifact library

SOC 2 Type II, ISO 27001, JP 就業規則 template library, jurisdiction playbooks for NL/UK/DE. Each takes 12-18 months. Table-stakes by the time a competitor catches up; we have more reference chains.

## Layer 5 · Anti-inference architecture (retaliation defense)

The retaliation-risk memo surfaces a moat layer we almost missed. Employer must not be able to infer that an employee opened their pattern page, created a vault, marked confounds, began a draft, or triggered a review. Telemetry partitioning, protected routes, min-group-size suppression, batching/redaction/delay — all architectural, none rhetorical. Almost impossible for Viva-style products to bolt on because their analytics stack is designed to log exactly these signals.

**The meta-moat:** our positioning is a voluntary constraint that incumbents cannot adopt without harming their existing business. A Microsoft team that wants to compete with Kashi must

internally argue that Viva Insights will stop showing individual-level telemetry to managers. That conversation is political, not technical. Most incumbents will not win it.

## 08 Risk register — 11 risks, explicit

Expanded from 6 to 11 risks based on the adversarial + anti-capture + retaliation + false-negative + attack-surface memos. Ranked by impact × probability. Each has an architectural mitigation — not just a policy.

### R1 · PILOT CUSTOMER'S PATTERN BECOMES EVIDENCE IN A HARASSMENT LAWSUIT

**Probability: medium. Impact: catastrophic.**

**Mitigation:** The 3-lane accountability architecture IS the defense. If Lane-B governed remediation is contractually enforced + documented, the pattern's visibility *supports* duty-of-care rather than undermining it. Pre-pilot: signed legal review acknowledging remediation obligation. Audit trail designed to be discovery-friendly *for* the company.

### R2 · MICROSOFT / ZOOM SHIPS A "MEETING-DYNAMICS REPORT" FREE TO EXISTING TENANTS

**Probability: medium-high in 24mo. Impact: large on Pro tier, small on Enterprise.**

**Mitigation:** Retaliation-risk memo upgrades this: incumbents can't copy anti-inference architecture without cannibalizing their analytics business. Enterprise moat (Layers 1 + 4 + 5) is intact. Starter becomes lead-gen funnel; revenue migrates toward Enterprise over time.

### R3 · FIRST 3-5 PILOTS CANNOT SECURE WORKER-REP / WORKS-COUNCIL CONSENT

**Probability: high early. Impact: 6-month Y1 slip.**

**Mitigation:** Target pilot customers with mature employee-rep infrastructure (post-IPO JP companies, regional manufacturers with established 労使関係). Free first 5 pilots in

exchange for case-study rights.

#### **R4 · ADVERSARIAL ADAPTATION / METRIC GAMING** (NEW — ADVERSARIAL MEMO)

**Probability: high once product is socially known. Impact: undermines claim base.**

Managers who know the detector surface route pressure around it: metric substitution (shift dominance to agenda control), channel displacement (move to 1:1s, async), hierarchical laundering (push decisions via proxies), symbolic compliance (cleaner surface, same dynamics).

**Mitigation:** "Adaptation-watch" product layer that flags suspiciously clean metric improvement as a signal, not success. Ban single-metric victory claims. Multi-metric corroboration required. Never market "interruption down X%" as proof of improvement.

#### **R5 · INSTITUTIONAL CAPTURE / SELECTIVE DEPLOYMENT** (NEW — ANTI-CAPTURE MEMO)

**Probability: medium. Impact: large — category-killing if exposed publicly.**

An org buys Kashi, excludes executive leadership from scope, targets only ICs or middle managers. "Kashi surveils workers, protects leaders" becomes the story.

**Mitigation:** Gate 1 of 6 deployment gates (scope parity — leadership in scope). Non-negotiable before pilot launch. Contractually enforced. If Gate 1 fails, pilot doesn't launch.

#### **R6 · METADATA LEAKAGE — INFERENCE RECONSTRUCTION** (NEW — RETALIATION-RISK MEMO)

**Probability: medium without careful design. Impact: catastrophic in JP (MHLW prohibits disadvantageous treatment for consultation).**

Employer infers that a worker opened their pattern page, created a vault, marked confounds, or began an escalation draft — even without seeing content. Directly actionable retaliation risk.

**Mitigation:** Telemetry partitioning (security vs analytics on separate surfaces). Protected routes (pattern-page opens, vault creation, draft state never exposed to business analytics). Small-team inference suppression. Batching + redaction + delay on any employee-facing event before it reaches analytics.

**R7 · FALSE-NEGATIVE LAUNDERING — SILENT DASHBOARD USED AS EXONERATION** (NEW — FALSE-NEGATIVE MEMO)

**Probability: high if unmitigated. Impact: category-killing.**

Employer cites clean Kashi dashboard as evidence-of-absence in a harassment complaint. "We ran the system, no signal, therefore no harm." EEOC base rate: 3 of 4 harassment targets never speak up.

**Mitigation:** 4-state dashboard (signal / no qualifying signal / insufficient observation / out-of-scope — never binary). Case-review rule explicitly prohibits complaint closure on "no signal." Scope warnings surfaced in UI, not buried in governance page. Contract prohibits citing Kashi output as exculpatory evidence.

**R8 · CALIBRATION DRIFT + ASR / DIARIZATION INPUT CONTAMINATION** (UPGRADED — SCIENTIFIC ATTACK-SURFACE MEMO)

**Probability: known-present, continuous. Impact: medium-large if unmitigated.**

NOTSOFAR-1 baseline speaker-attributed tcpWER: 32.4% multichannel / 46.8% single-channel. 54.9% on overlap-heavy sessions. Koenecke: 0.35 vs 0.19 WER racial disparity. Disfluency bias documented.

**Mitigation:** Input-quality gating before output is pilot-grade (overlap flags, L2 caution surfaces, subgroup audits). Baseline reset rules after reorg / manager change. Per-platform × language × feature support matrix published.

**R9 · REGULATOR WALKBACK (EU AI ACT, MHLW, APPI AMENDMENT)**

**Probability: low-medium. Impact: large.**

**Mitigation:** Architecture is already more restrictive than current rules. 6+ months of product-side lead time given deployment preconditions. Regulatory advisory-board seat monitors monthly.

#### R10 · SCANDAL — KASHI OUTPUT MISUSED FOR HR RETALIATION

**Probability: low. Impact: catastrophic if public.**

**Mitigation:** "No export to HR systems" is product-enforced, not just contractual (no API, no export format). Public commitments on deck + governance page. Incident response plan: public post-mortem within 30 days.

#### R11 · GOODHART-TYPE GAMING WITHIN MANAGER MIRROR (NEW — MEASUREMENT-SCIENCE MEMO)

**Probability: medium. Impact: soft — erodes outcome claims but doesn't sink product.**

Even with only self-comparison (not public leaderboards), managers game the private mirror. Keigo-asymmetry and interruption signals improve structurally while dominance expresses in untouched channels.

**Mitigation:** Remediation outcomes + human recovery as mandatory measurement layers (see §11). Metric movement alone is never the victory claim.

## 09 Operating model

### Legal entity

Delaware C-corp + JP KK subsidiary. Enables US VC participation, preserves Series A optionality, supports JP-compliant billing via the KK.

**ARCHITECTURAL FIX REQUIRED — JAPAN DATA-RESIDENCY GAP**

The procurement / security-buyer memo flagged a concrete infrastructure problem: **Vercel's primary processing facilities are in the United States.** If Kashi makes a Japan-data-residency claim to a JP enterprise buyer and any regulated content flows through Vercel compute, the claim is technically false.

**Fix (before first JP Enterprise deal):**

- Pin Supabase to Tokyo region (ap-northeast-1) — already default, verify
- Audit every Vercel data flow: any regulated content must be client-side encrypted before leaving the browser, or processed in edge functions pinned to Tokyo region
- Subprocessor map must disclose Vercel-US processing honestly
- Do not market "Japan data residency" until the audit clears

## Geographic model

- **Year 1:** Tokyo-based ops. Remote-friendly, quarterly in-person for Enterprise pilots
- **Year 2:** Osaka/Kyoto expansion (regional manufacturing + academic partnerships)
- **Year 3:** Singapore satellite (first international, not EU — per funding-pathways memo)
- **Year 4+:** NL/UK satellite; DE later only after works-council playbook matures

## Headcount plan

STAGE	HEADCOUNT	STRUCTURE
Pre-seed (now)	2–3	Founder(s) + fractional labor-law advisor
Post-pre-seed (mo 6)	4	+ 1 eng + 1 JP enterprise salesperson
Post-seed (mo 18)	8	+ 1 security/DevSecOps + 1 CS + 1 product + 1 ops
Post-Series A (Y3)	~35	eng (12) + GTM (10) + CS (6) + ops (4) + founder-office (3)

## Non-obvious operating principles

1. **Anti-surveillance culture internally.** We cannot sell "we don't watch employees" if we use keystroke monitoring on our own team. Picks tool discipline that matches positioning.

2. **Public post-mortems.** Every material incident gets a public post-mortem within 30 days. Trust-building, not crisis-response.
3. **Worker-rep on advisory board (permanent seat).** Compensated, voting. Non-symbolic per the anti-capture memo.
4. **No per-individual sales commissions.** Goodhart applies to our team too. Team-based bonuses only.
5. **Six deployment gates are pre-pilot mandatory.** Scope parity, notice/consultation, access discipline, challenge workflow, misuse sanctions, pause/rollback — from the anti-capture + rollout-research memos. Any pilot that skips a gate doesn't launch.

## 10 Strategic options — Japan-first expansion

### Expansion sequence (reversed from v1)

V1 of this plan said "EU first (DE/NL/UK), US second." The cross-cultural + labor-politics + funding-pathways memos converge: **reverse this**. Japan-first is the doctrine, not a fallback.

PHASE	MARKETS	WHY
Y1-Y3 (harden)	Japan	Regulatory conditions most favorable. Channels established. Funding programs (AI Adoption, Workstyle Reform) fund customer acquisition.
Y3 selective	Singapore	PSG + EDG grants, English-capable workforce, procurement discipline similar to JP. <b>First international.</b>
Y4	NL + UK	Procedurally lighter than DE. UK ICO + EU AI Act compliance already aligned with our architecture.
Y5	DE	After works-council playbook matures. BetrVG §87(1)(6) requires co-determination on monitoring tech — heavier.
Y5+	US	After SOC 2 Type II + ISO 27001. US market larger but saturated with HR-tech incumbents.

### Strategic acquisition landscape (if we wanted to sell)

1. **Microsoft (Viva portfolio).** Best product fit. Kashi fills the manager-behavioral-mirror gap Viva refuses to ship. Precedent: MS acquired Glint for \$1.5B (2019).

2. **Workday.** Peakon lacks governance layer. Precedent: Workday acquired Peakon for \$700M (2021).
3. **Recruit Holdings (JP).** Owns Geppo. Would want governance differentiation + JP labor fluency. Typical JP SaaS tuck-in ¥10-30B.
4. **SAP SuccessFactors.** Weakest fit; highest price when they buy.

### Primary path: category leadership + IPO

JP Mothers listing at ~¥5B ARR (standard SaaS exchange threshold). 5-7 year horizon. Product improves with data (cross-industry baseline moat); category is a new SaaS segment we'd define. **Don't optimize for acquisition.** The value to society is highest if we become the JP governance standard — which requires independence to say no to acquirers whose business models conflict with our refusals.

## 11 Traction — detection, remediation, recovery

The remediation-outcomes memo rewrites what success looks like. We cannot declare success on detection metrics alone ("caught 3 out of 3 seed patterns"). That's proof of mechanism, not proof of value. Traction must add two further outcome layers.

### 6-layer success model (upgraded from 4 layers)

LAYER	WHAT IT MEASURES	EXAMPLE METRIC
1. Adoption	Is anyone using the product?	Pattern-page monthly active users, manager feedforward-check rate
2. Behavior	Did the manager's own pattern change?	Structural-signal movement + adaptation-watch flags
3. Responsiveness	Did the org act?	Time from flag → review, % Lane-B transitions after persistence
4. Strategic	Did the category of problem recur?	Recurrence rate at 30 / 90 / 180 days
<b>5. Remediation quality (NEW)</b>	Was the handling actually fair?	Worker-rated fairness, clarity, follow-through at 30-day post-handling

LAYER	WHAT IT MEASURES	EXAMPLE METRIC
<b>6. Human recovery (NEW)</b>	Did the affected person recover?	Speaking-share recovery, answered-question recovery, chilling-delta reduction at 30/60/90

## North-star metric

**Paying 50-500-person JP companies actively using Kashi with (a) a signed 就業規則 amendment documenting consultation, and (b) at least one completed Lane-B remediation cycle where worker-rated fairness  $\geq 4/5$ .**

Why: it rewards the end-to-end outcome (not detection alone), penalizes vanity signups, and specifically resists Goodhart by requiring worker-side assessment.

## Counter-metrics (watch for gaming / capture)

- Dashboard open rate (too high = performative; too low = product not landing)
- Feedforward-check rate without behavior change (adaptation-watch)
- Lane-A → Lane-B transition rate (too low = anti-laundering rule not firing)
- Customer drop-offs in 30 days after first Lane-B event (is remediation panicking customers?)
- Reversal-after-contest rate (are reviewers rubber-stamping?)
- Speaker-attribution error rate (is ASR degrading?)
- "Flagged-but-no-action" burden rate (does the queue clear?)

## Reporting cadence

- **Monthly:** ARR, NRR, logos, pilot pipeline — email
- **Quarterly:** full P&L + product update + risk-register review — board meeting
- **Annually:** strategy review, 3-year plan refresh, **public transparency report** (aggregate usage + incident log)

## 12 Rhetorical doctrine — claims we make and refuse

Six of the research memos independently flag the same risk: rhetorical overreach is the fastest way this category gets killed. Here's the discipline, codified.

### ✓ WE DO SAY

- Kashi surfaces repeated interaction asymmetries under uncertainty
- The pattern may constitute evidence consistent with uneven conversational treatment
- Kashi is governance infrastructure that makes institutional visibility earlier
- Structural signals, contestable, review-worthy, human-reviewed
- The institution can no longer plausibly say it could not see the pattern
- Privacy-bounded meeting governance

### ✗ WE REFUSE TO SAY

- ~~Kashi detects harassment, intent, or illegality~~
- ~~The pattern is the harm~~ (replace with "pattern may constitute evidence consistent with...")
- ~~We detect harmful team dynamics earlier~~ (replace with "surface asymmetries earlier")
- ~~The CEO's instrument for seeing the bill~~ (demote from front door to sponsor-only context)
- ~~Not an employee-monitoring tool~~ (replace with "not a general-purpose surveillance product — a restricted meeting-governance system")
- ~~How we know we're right~~ → "what the current pilot demonstrates"
- ~~Kashi causally reduces harassment~~
- ~~Consent alone makes Kashi legitimate~~
- ~~Absence of signal means absence of problem~~
- ~~A cleaner score proves the power problem is solved~~

**The one-paragraph positioning (paste-ready):** Kashi is a privacy-bounded meeting governance layer that surfaces repeated interaction asymmetries from structural meeting metadata — under explicit uncertainty and procedural limits — so teams can respond earlier

without turning communication into a surveillance archive. Mirrors, not microscopes. Patterns, not content. No HR decisions from the tool. The institution is deliberately allowed to see less than the technology could show, while the affected individual can see more than they normally can, and control escalation.

## 13 **Research library (24 memos)**

Everything in this plan traces to one or more of the memos below. They are investor-read artifacts; open source for diligence.

### **Project master**

- [Kashi — Progress & Project Overview](#)

### **Foundational (integrated into v1)**

- [Labor-relations / worker-representation deployability](#) — drove 3-lane model + consultation packet
- [Manager adoption / defensiveness lens](#) — drove anti-laundering + feedforward

### **Business / GTM / financing (new integration)**

- [Founder narrative](#) — why "disciplined category design" is the honest "why us"
- [Funding pathways \(global\)](#) — non-dilutive track (JP AI Adoption, Workstyle Reform, Singapore PSG/EDG)
- [Rebuilt ROI model](#) — AxCxIxE discount framework
- [Category entry strategy](#) — "privacy-bounded meeting governance" wins

### **Risk / moat / adversarial**

- [Adversarial adaptation](#) — metric gaming & displacement
- [Anti-capture](#) — 6 deployment gates
- [False-negative risk](#) — 4-state signal, no exoneration
- [Retaliation risk](#) — metadata leakage + anti-inference architecture
- [Scientific attack surface](#) — ASR numbers, tiered detectors

## Deployment / ops / trust

- [Labor politics / union / works-council](#) — deployment gate not compliance afterthought
- [Procurement / security buyer](#) — CAIQ/SIG pack, Vercel residency fix
- [Rollout research synthesis](#) — 7 rollout gates, two-track pilot
- [Ally / observer concern path](#) — V2 bystander product
- [Employee trust / anti-surveillance](#) — disciplined non-seeing
- [Cross-cultural / multilingual](#) — Japan-first doctrine + locale registry

## Product / science / measurement

- [Remediation outcomes](#) — 6-layer success model
- [Measurement science](#) — asymmetry-under-uncertainty, evidence grades
- [Product wedge](#) — layered wedge architecture
- [Meeting-type normalization](#) — first-class model input
- [Legal / procedural fairness synthesis](#) — contestability spine
- [Research synthesis \(buyer economics\)](#) — hidden-loss portfolio framing

24 PDFs · all hosted at `/research/*.pdf` · license: internal research, freely shareable.

## Ideas\_wave3 — technical-dev consideration library (18 files)

A separate body of 17 technical-dev memos + 1 README, explicitly marked as *perspective-expanding material, not binding specifications*. These memos directly shaped the new detector registry, evidence-grade types, meeting-type normalization, and procedural-fairness spine. Hosted for diligence:

### DETECTOR ENGINEERING & SCIENCE

- [Detector boundary perspective](#)
- [Baseline calibration](#)
- [Meeting-type normalization \(technical\)](#)
- [Longitudinal aggregation](#)
- [Confidence / abstention](#)
- [Speaker identity / diarization](#)

## SYSTEM ARCHITECTURE

- [Model boundary perspective](#)
- [Backend logic — critical verification](#)
- [Role & visibility architecture](#)
- [Input substrate](#)

## RISK, SECURITY & ADVERSARIAL

- [Adversarial gaming \(dev lens\)](#)
- [Privacy / retention / evidence boundary](#)
- [Security / enterprise assurance](#)
- [Legal-max hardening \(PDF\)](#)

## PRODUCT & PROCESS

- [Contestability / correction workflow](#)
- [Rollout operability](#)
- [Refined consideration register](#)

## READING NOTES

- [README — how to use these files](#)

42 total documents across both libraries · 24 business research + 18 technical-dev consideration · all hosted at{" "} /research/ and{" "} /research/ideas-wave3/ .

---

**Companion documents:** [/deck.html](#) · [/tour.html](#) (interactive tour) · [/walkthrough.html](#) (6-frame demo) · [/progress.html](#) · [/governance](#) · [/founder.html](#) · [/integration-log.html](#) (v1→v2 research diff) · [consultation packet](#) · [outbound templates](#)

v2 · 2026-04-21 · rewritten from 22-memo research library · numbers are targets grounded in cited sources, not guarantees · living document, update quarterly