THE HOOK

A bounded governance instrument that helps organisations surface repeated communication asymmetries before they harden into late-stage people incidents.

Kashi helps organisations see repeated meeting asymmetries earlier, using only structural metadata, with worker-protective limits and documented governance boundaries from day one.

¥7.6T

Japan's annual mental-health productivity loss · ~1.1% of GDP
Yokohama City University 2025

¥7.9M

Per mental-health leave case, total employer cost
METI 2025 benchmark (¥5.25M salary, 1 yr leave)

1 in 3

Japanese workers experienced パワハラ in the last 3 years
MHLW 令和5年度 実態調査

THE COST IS A PORTFOLIO, NOT ONE HERO NUMBER

¥7.3T of Japan's ¥7.6T loss is presenteeism. Absence is the tail.

Earlier pitches led with "¥7.9M per leave case." That's tail risk, not base case. The defensible story is the four-bucket hidden-loss portfolio — and the A×C×I×E discount that separates "total pool" from "what Kashi realistically captures."

BUCKET	WHAT IT CAPTURES	EVIDENCE ANCHOR
1. Presenteeism (biggest)	Underperformance while at desk; invisible on headcount dashboards	Yokohama CU 2025: **¥7.3T of ¥7.6T**
2. Regrettable attrition	Manager-driven departures; replacement costs 200% of salary	Gallup: managers = **70%** of engagement variance
3. Formal-escalation premium	Legal fees, time-in-HR, reputation when informal → formal	Acas: **3×** formal vs informal cost
4. Leave-case tail	Visible end-stage: 傷病手当 claims, backfill, recruiting	METI 2025: ¥7.9M/case *(tail, not base)*

A × C × I × E — WHAT KASHI REALISTICALLY CAPTURES

Total loss pool × **Addressable** × **Coverage** × **Intervention uptake** × **Effect size**. For a 30-leave-case/year company with ¥127M operational-economic loss pool:

Strict · 15×60×30×20% → 0.54% reduction = ¥0.69M/yr

Mid (plan) · 25×70×40×25% → 1.75% = ¥2.22M/yr

Upper · 40×80×50×30% → 4.80% = ¥6.10M/yr

The key insight: leave-cost alone isn't the sell. The sell is the portfolio — presenteeism + attrition + escalation premium are 4-8× larger than the leave tail.

Moral weight backing the budget: **2.0× suicide-ideation, 2.67× suicidal-behavior odds** (Leach 2023, *Lancet Public Health*). **Laissez-faire managers produce 4.3× パワハラ and 2.6× mental-health problems** at 6mo (津野香奈美 2023).

THE PROBLEM

Four structural reasons existing tools fail.

01 · BURDEN ON THE VICTIM

The person under pressure must recognize it, prove it, and raise it, alone. Power asymmetry makes all three hard. Gaslighting makes them doubt perception.

02 · FRAGMENTS, NOT PATTERNS

Each moment can be explained away. The harm is in the pattern. Organizations only see isolated complaints, never the structure.

03 · LAGGING INDICATORS

HR sees attrition, leave, 自殺. By then the damage is done. By the time the metric moves, the bill has arrived.

04 · SINGLE-CONSTITUENCY SALES

Existing tools are sold to HR alone. A workplace-AI deployment that survives scrutiny requires **executive sponsor + legal review + worker-representation alignment** before launch. All three, or none.

> *Archaic's ハラスメントチェックAI reads content and sells to HR alone.*
> *FRONTEO KIBIT Eye, deployed at MUFG and Aeon, has gone quiet.*
> *Content classifiers miss 70%+ of Japanese cases because パワハラ is 遠回し.*

THE INSIGHT

Point the mirror *upward* at power. Structural signals, not content.

Five empirically-backed structural signals. Each deterministic. Each explainable. Each maps to MHLW パワハラ 類型 3 (isolation) or 類型 5 (cold shoulder) — the transcript-visible harassment types.

SIGNAL	EVIDENCE
Intrusive-interruption asymmetry	Anderson & Leaper 1998 · meta-analysis 43 studies · *d*=0.33 · status beats gender
Speaking-time / floor-share inequality	Schmid Mast 2002 · *Human Communication Research* · dominance correlation
Topic-credit exclusion (ignored → taken over)	Sacks / Schegloff / Jefferson 1974 · foundational conversation analysis
Chilling delta (post-trigger participation drop)	Morrison 2014 · organizational silence · Detert & Burris 2007
Response-latency asymmetry	Stivers et al. 2009 *PNAS* · delayed response = dispreference cross-linguistically
Keigo (敬語) peer-addressee asymmetry (JP-specific)	Cook 2011 · Saito 2011 · Pizziconi · Ide wakimae framework

False-positive mitigation: every signal calibrated against each speaker's own 90-day baseline, not the team average. Defeats introversion, chair-role, L2, neurodivergence confounds.

LIVE WALKTHROUGH · DRIVE WITH ARROW KEYS

Watch the system work on one real meeting, end to end.

Six frames. Same meeting. Structural signals → keigo classifier → 63-day pattern → three different views for three different roles. Real data, no mocks, no voiceover — you drive it live in front of the room.

FRAME 01–03 · SYSTEM

One 47-minute product-review meeting. Kimura cuts Mira mid-word at 00:01:05 (200ms overlap). Layer 1 extracts 1 interruption, chilling delta –93%, Gini 0.34. Layer 2b (keigo) classifier: Kimura scores 0.38 toward Mira, 0.88 toward Nakamura. 0.50 asymmetry gap. Same speaker, same meeting.

FRAME 04 · PATTERN

Zoom out: this is not one meeting. Mira's speaking share dropped 19% → 6% sustained over 63 days. Kimura's interruptions toward her 4.7× higher than toward peers. Persistence + directionality + sustained drop = review-worthy event fires at Layer 5.

FRAME 05–06 · MONEY SHOT

Three windows side-by-side. Mira sees a private, observational pattern page. Kimura sees his own Manager Mirror with one concrete action. CEO sees the Executive Brief: Kimura flagged, ¥3–8M modeled impact. Three framings, one data source, RBAC-separated.

Open the walkthrough →

Fallback if wifi drops: /demo/mirror · /demo/ceo

THE SYSTEM · ONE MEETING, ALL 6 LAYERS

Every number below is computed from the same 47-minute transcript.

m-product-w12 · 2026-04-15 · 89 turns · 5 participants. No abstractions. No “it would look like this.” Real.

LAYER 1
deterministic

Turn-timing extraction. **1 intrusive interruption** (Kimura→Mira, overlap 200ms). Speaking-share **Kimura 42% / Mira 11%**. Chilling event: Mira's 3 turns after the cut average **1.3s** (vs her 18s personal baseline).

LAYER 2
structural wrappers

Dyadic continuity check: Kimura→Mira interruption pattern appears in **3 of last 3 meetings**, span **28 days** → persistence score 0.45. Speaker-baseline drift: Mira's 90-day share trend **19% → 6%**, sustained below baseline for **63 days**.

LAYER 2b
keigo classifier

Per-addressee politeness-register index from surface grammar only. Kimura's keigo toward **Nakamura: 0.88** (honorific), toward **Yoshida: 0.50**, toward **Mira: 0.38** (plain form, "タメ口"). Asymmetry gap **0.50** — review-worthy at ≥0.30.

LAYER 3
meeting metrics

Floor-time Gini **0.34** (heavy skew). Interruption-asymmetry matrix: 82% of Kimura's cuts land on one target. Response-reciprocity: when Mira speaks, response-rate from Kimura is **0.18**, from others **0.71**.

LAYER 4
longitudinal

Rolling 30/90/180-day windows, per-speaker calibrated. Directionality ratio **4.7×**. Sustained-drop **63 days**. Keigo-asymmetry persistent across **4 consecutive meetings**. This is where one meeting becomes a pattern.

LAYER 5
review-worthy event

Composite: severity × repetition × directionality × confidence = **0.87**. Threshold **0.60**. **FIRES.** Every component traces to specific turn IDs. Explainable. Reproducible. No LLM judgment in the scorer.

LAYER 6
role-based view

RBAC + k-anonymity (k≥5) + differential privacy (ε≤1). Three different framings render: Mira's private pattern page (observational), Kimura's Manager Mirror (own behavior + action), CEO's Executive Brief (aggregate, no individual subordinate data). See frame 05 of the walkthrough.

3 / 3	0	100%	7
Patterns detected on seed	False positives on healthy control	Deterministic replay	Structural detectors live

GOVERNANCE · 3-LANE ACCOUNTABILITY MODEL

Coaching is not the endpoint.

Manager Mirror is a bounded developmental lane inside a broader accountability system. Private self-correction first. If the pattern persists, it escalates. That is the **anti-laundering rule**: Kashi is not a device for saying "we told him" and doing nothing.

LANE A · PRIVATE SELF-CORRECTION

Manager sees their own pattern in Manager Mirror. Private, weekly, observational language, paired with **feedforward commitments** ("next 2 meetings, wait for sentence completion before redirecting"). ~80% of cases resolve here if the manager is coachable.

LANE B · GOVERNED REMEDIATION

If the pattern persists **after the protected self-correction window**, it leaves the developmental lane. Moves to a documented remediation process with timeline, check-ins, and structured support. Not HR punishment. Not silent.

LANE C · FORMAL REVIEW

Only if Lane B fails. Strict entitlement rules, audit logging, and employee notice required. The organization can no longer plausibly claim it didn't know. The tool does not adjudicate — humans do, under process.

Architectural guarantees

- **k-anonymity** (k≥5) on every aggregate · cannot be disabled
- **Differential privacy** (ε≤1) on exec dashboards
- **Four-tier retention:** raw 14d · analytics 24mo · events 12mo · legal-hold extended
- **Per-speaker baseline calibration** — no cross-person norms
- **Minimum-sample suppression** (<5 meetings / <30 days = no signal)
- **Audit trail visible to affected individual** — they see who looked
- **Manager outputs are structural + self-comparative**: no named subordinate telemetry by default

Kashi will not do (12 items, abridged)

- Will not "detect" harassment, intent, illegality
- Will not label anyone an abuser
- Will not infer emotion or affect (EU AI Act Art. 5)

- Will not feed HR decisions (Annex III §4)
- Will not read message content server-side
- Will not analyze audio / voice / prosody
- Will not show managers individual subordinate reports
- Will not interrupt meetings in real time
- Will not show a company-wide health score
- **Will not allow coaching to be the endpoint** when asymmetry persists

RESPONSIBILITY BOUNDARY — WHAT KASHI CAN AND CANNOT CLAIM

✓ Responsible claims

- Surfaces repeated structural interaction asymmetries for human review
- Supports earlier behavioral correction if surrounding process is sound
- Reduces plausible deniability about repeated meeting-level patterns

✕ Claims we refuse

- Detects harassment, intent, or illegality
- Developmental framing makes monitoring backlash disappear
- Manager Mirror alone will reliably improve behavior
- Cleaner score proves power problem is solved

SECURITY · BOARD-LEVEL CONFIDENCE

Where your data lives, who can read it, what happens if we get breached.

Before a CEO or board signs off, they ask: “Is my company's meeting data going to leak somewhere it shouldn't?” Here are the answers, with specifics.

- ✓ **Japan-region data residency**
  Tenant data stays in Tokyo/Osaka region (ap-northeast). Never crosses a border. Contractually pinned to the chosen region.

- ✓ **Encryption at rest (AES-256) + in transit (TLS 1.3)**
  Every database, every backup, every inter-service call. Keys rotated quarterly; per-tenant encryption key planned for Enterprise tier.

- ✓ **Zero third-party content processing**
  We do not send transcripts or content to OpenAI, Google, or any external classifier. Structural detectors run in our own infrastructure. Full stop.

- ✓ **Role-Based Access Control enforced at the query layer**
  Supabase Row-Level Security: managers cannot SELECT subordinate individual data even with direct DB access. UI restrictions are not the security boundary — the schema is.

- ✓ **Four-tier retention with hard TTLs**
  Raw transcripts 14 days. Analytics 24 months. Review-worthy events 12 months. Legal-hold only when justified and notified. Deletion is automated, auditable, irreversible.

- ✓ **Audit log visible to the affected employee**
  Every drill-down writes an audit row. The affected person can read their own log. No silent surveillance — the people being observed can see who looked.

- ✓ **No admin-level content access — ever**
  Kashi staff cannot read customer transcript content, even with root credentials. Break-glass procedures are pre-committed, notified, and audit-logged. Evidence vault content is E2E encrypted (roadmap) — victim holds the key.

- ✓ **Breach SLA: 24h notification**
  Pre-committed incident response: all affected tenants notified within 24 hours. Regulators notified per APPI / GDPR timelines. Post-mortem shared publicly within 30 days.

- ✓ **Compliance roadmap: SOC 2 Type II · ISO 27001**
  Year-1 target: SOC 2 Type II certification. Year-2: ISO 27001. Annual third-party penetration testing. Continuous vulnerability scanning.

- ✓ **Works-council / employee-rep consultation built in**
  Enterprise activation requires documented 就業規則 revision with written employee-representative opinion (MHLW sample 000683138). GDPR Art. 88 + JP labor-law aligned. Not a notice screen — a real consultation process.

✓ **Labor-consultation packet as deliverable**

Every Professional/Enterprise deployment ships with a packet: purpose statement, approved/prohibited uses, role-access matrix, retention map with justification, drill-down conditions, challenge workflow, anti-retaliation commitments, pilot sunset rules. Real artifact for legal + reps + sponsor.

DEPLOYMENT PRECONDITIONS — KASHI IS NOT A SILENT PILOT

Before launch: executive sponsor + legal review + worker-representation alignment. All three, documented. We will not help you deploy Kashi without these.

Pre-launch artifacts (all mandatory): employee-facing explanation · permitted and prohibited uses · access matrix · retention justification · challenge pathway · anti-retaliation commitments · pilot sunset / deletion rules.

Necessity test: every retention window, every drill-down condition, every visible metric has a one-line necessity argument (why this duration, why this object, why not less). UK ICO + CNIL + EDPB style.

The honest part: we are pre-certification. SOC 2 and ISO 27001 take time. What we have today is architecture that passes audit by design — not by scrambling before the auditor arrives. Labor-relations memo verdict: consultation-possible, not consultation-proof. We are doing the hard work to close that gap.

WHY NOW · WHY US

Nobody has shipped this yet. Six reasons it's finally possible.

01 · TRANSCRIPTION

JP ASR on multi-speaker meetings only became good enough in 2023–2024. Sortformer v2 DER on CALLHOME: 12.7%.

02 · REGULATION

EU AI Act Art. 5 (Feb 2025) + Annex III §4 (Aug 2026) crystallized the defensible approach.

03 · NEW CATEGORY

“Governance infrastructure” bridges engagement tools, compliance tools, productivity tools. Courage to cross.

04 · JP POLITICS

Pointing the mirror upward at power is politically dangerous. It's our defining feature.

05 · PRIVACY COST

K-anon + DP + 4-tier retention + audit trails + 就業規則 consent is 6–12 months of infra most skip.

06 · WRONG BUYER

Traditional HR buyers want defensive tools. CEO buyers want preventive. Different sale, higher value.

CAPABILITY	KASHI	ARCHAIC / FRONTEO	VIVA GLINT	READ.AI	WEVOX / GEPPO
Structural pattern detection (not content)	✓	–	–	–	–
Mirror points upward at power	✓	–	refused	–	–

CAPABILITY	KASHI	ARCHAIC / FRONTEO	VIVA GLINT	READ.AI	WEVOX / GEPPO
Keigo (敬語) asymmetry (JP-specific)	✓	–	–	–	–
Refuses company-wide health score	✓	–	–	–	–
EU AI Act Art. 5 clean by design	✓	?	?	?	?
Published harassment-outcome evidence	**pending NAQ-R study**	–	–	–	–

ROI · PRICING · ASK

¥13M–¥26M/year recoverable per 500-person company.

500-person company × ¥5.25M average salary × 5% productivity drag attributable to harmful team dynamics = **¥131M/year at stake**. Shifting 10–20% of cases from late-stage to early-stage = **¥13M–¥26M/year recovered**.

Pricing

	STARTER	PROFESSIONAL	ENTERPRISE
Scale	up to 20 employees	up to 500 employees	500+ / multi-region
Detectors	3 core structural	all 7 (incl. keigo)	all 7 + evidence vault (E2E)
Views	Manager Mirror + self-visibility	+ Executive Brief + feedforward	+ custom role matrix
Retention	90-day hard delete	24-month analytics	configurable
Data residency	shared JP region	dedicated JP region	choice incl. on-prem (roadmap)
Labor-consultation packet	—	**downloadable template**	**hands-on (就業規則 review + rep-process sit-in)**
Compliance artifacts	—	—	**SOC 2 Type II + ISO 27001**
Support	community	email · 72h SLA	24/7 · dedicated CSM
Price	**Free**	**¥800/user/mo**	**¥10M/yr + ¥800/user/mo**

8× ROI

100-person co at Pro: ¥960k/yr. 1 averted case = ¥7.9M. **8×.**

3.3× ROI

500-person co at Pro: ¥4.8M/yr. 2 averted cases = ¥15.8M. **3.3×.**

1.3× + defensibility

1000-person Enterprise: ¥19.6M/yr. 10% of ¥262M drag = ¥26M recovered + governance upside.

Roadmap

- **Phase 1 (now):** proof of signal, shipped at `kashi-lilac.vercel.app`
- **Phase 2:** governance-ready · NAQ-R outcome validation · labor-consultation packet v1 · feedforward prompts
- **Phase 3:** enterprise hardening · SOC 2 Type II · E2E evidence vault · Lane B remediation workflow
- **Phase 4:** victim-explainer page · Lane C formal-review interface · jurisdiction playbooks (JP / DE / NL / UK)

The ask (reframed)

Not “introduce me to the CEO.” A Kashi pilot that survives scrutiny needs:

- an **executive sponsor** who can commit to the 3-lane accountability architecture
- a **legal / compliance reviewer** for APPI + labor-law sign-off
- a **worker representative** (union rep, employee-rep, or equivalent) willing to consult on the 就業規則 amendment

Introduce us to all three at one 50–500-person Japanese company. 90-day pilot. Real data, real consultation, real outcome measurement.

Live: kashi-lilac.vercel.app · Walkthrough: /walkthrough.html · Business plan: /business.html · Progress: /progress.html · Built in 23h · 2026-04-21