

# Kashi research memo

## Labor relations / worker representation deployability

Tailored summary of the research findings for the Kashi project. Built to be usable for product, pitch, governance, and pilot design decisions.

Project	Kashi / Progress & Project Overview alignment memo
Scope	What the labor-relations lens means for Kashi specifically, what is already strong, what is still politically fragile, and what should be changed now

### Executive conclusion

- Kashi is materially more defensible than the earlier generic concept because it has already narrowed itself to structural metadata, deterministic detectors, no affect inference, no named-individual telemetry to managers, no company-wide health score, and a worker-controlled evidence-vault design. [K1]
- That said, Kashi is still not automatically labor-relations-safe. The strongest external evidence shows that workplace AI is judged not only by intent statements, but by whether the system can function as a people-monitoring device, whether representative consultation is required, and whether misuse is technically difficult rather than merely prohibited on paper. [W1] [W3] [W4] [W5] [W6]
- The project therefore does not need a conceptual rewrite. It needs a deployment-hardening rewrite: sharper language, stronger consultation posture, stronger necessity / proportionality defense, and clearer product-level anti-misuse constraints.

This memo is intentionally practical. It does three things: (1) states the research result, (2) maps the result onto Kashi as currently described in the project PDF and concept note, and (3) turns the result into decisions the team can use now.

## 1. Research question and method

Question. What does the labor-relations / worker-representation lens imply for Kashi, and how should that change the way the project is designed, described, and rolled out?

Method. This memo synthesizes two internal project documents and a targeted set of external primary or near-primary sources: European Commission AI Act materials, German and Dutch worker-representation / monitoring rules, UK ICO monitoring guidance, French CNIL enforcement materials, and Japanese MHLW documents on work-rules change and employee representative process. [K1] [K2] [W1] [W3] [W4] [W5] [W6] [W8] [W9]

## 2. External research result: what institutions actually care about

The main result is simple but brutal: in workplace AI, institutional acceptability turns less on model cleverness and more on whether the system can be defended as narrow, necessary, proportionate, and consultation-compatible.

## 2.1 EU AI Act: the baseline is already stricter than many product pitches assume

- The Commission's AI Act FAQ states that emotion recognition in the workplace is among the prohibited AI uses. That directly validates Kashi's 'no affect / no emotion / no voice-stress' line as a legal red line, not a branding choice. [W1]
- The same FAQ states that risk classification depends on the function performed by the AI system and on the specific purpose and modalities for which it is used. This matters because a product cannot fully hide behind a benign slogan if its concrete use pattern pushes it toward employment or worker-management functions. [W1]
- The Commission also states that if a high-risk AI system is deployed at the workplace, deployers must beforehand inform affected employees and workers' representatives, in addition to any worker-consultation rules that already apply. In other words: consultation is not a nice-to-have communications step after the pilot; it is part of lawful deployment logic. [W1]

## 2.2 Worker-representation rules across Europe punish capability, not just intention

- Germany: the Works Constitution Act gives the works council co-determination rights for the introduction and use of technical devices designed to monitor employee behaviour or performance. This means that a system that can plausibly be characterised as behavioural monitoring can trigger a representative-rights issue even if the employer describes it as improvement or governance. [W3]
- Netherlands: the Dutch Data Protection Authority states that if an employer plans to use a staff-tracking system to monitor employees, this is only allowed with works-council approval. The Dutch regulatory framing is useful because it shows how quickly 'tracking' language appears once a system structures person-linked observation of worker conduct. [W4]
- Broader implication: in these systems, institutional readers ask 'what can this tool do to people in practice?' before they ask 'what does the vendor say it is for?'

## 2.3 Data-protection authorities care about intrusiveness, necessity, and proportionality

- The UK ICO says monitoring staff can be intrusive even when done with good intentions, and that employers need a lawful basis and justification. This is not an anti-monitoring rule in the abstract; it is a necessity test. [W5]
- The French CNIL's 2025 enforcement summary states that permanent video surveillance of employees, absent exceptional circumstances, breaches personal-data-protection rules. Continuous or ambient surveillance is therefore a live enforcement issue, not just a theoretical risk. [W6]
- The EDPB's consent guidance stresses that consent is generally not valid where the data subject has no real choice or feels compelled. In the employment context this weakens any strategy that relies mainly on individual employee opt-in as the core legitimacy story. [W7]

## 2.4 Japan: the practical analogue is work-rules governance plus representative process

- The MHLW sample filing for creating or changing work rules requires attaching the employee representative's written opinion. This is concrete evidence that workplace-rule changes are not treated as one-way employer declarations. [W8]
- A Labour Bureau reference note adds that the majority representative must be selected by a democratic process involving all workers at the site, cannot be a manager, and must not suffer adverse treatment; the employer must also provide the practical support needed for the representative to gather views. [W9]
- For Kashi, this means that 'just publish a notice and go live' is too weak as a Japanese deployment story if the system materially changes how meeting-derived behavioural data is processed and surfaced.

### Research conclusion from the external evidence:

- A workplace AI system becomes politically fragile when it can be read as behaviour-monitoring infrastructure.
- The fragility increases sharply when worker representatives are not visibly built into deployment.
- The strongest defence is not an ethics paragraph; it is a combination of narrow data design, technical anti-misuse constraints, documented consultation, and a necessity / proportionality case for each intrusive element.

#### Evidence to implication map

Evidence	What it means for Kashi
Emotion recognition at work is prohibited under the AI Act. [W1]	Kashi's no-affect stance is mandatory and should stay a front-page safeguard, not a footnote.
Risk classification depends on actual function, purpose, and modalities of use. [W1]	How Kashi is sold, configured, and embedded in employer workflows matters as much as its technical architecture.
Worker representatives must be informed before workplace deployment of high-risk AI. [W1]	Consultation must be built into rollout sequencing and pilot governance.
Works-council / staff-tracking rules focus on monitoring capability. [W3] [W4]	Language that flatly says 'this is not monitoring' is vulnerable if the product still creates person-linked behavioural telemetry.
Monitoring must be justified, proportionate, and not ambient by default. [W5] [W6]	Retention windows, CEO views, and drill-down conditions need explicit necessity arguments.
Employment consent is structurally weak. [W7]	Do not rely on individual opt-in as the core legal / political story.
Japan requires employee-representative opinion on work-rules changes and protects the representative process. [W8] [W9]	Kashi needs a Japanese labor-consultation packet, not just a notice screen.

### 3. What the research means for Kashi specifically

The key point is that Kashi is already materially stronger than the generic concept. The project PDF contains several design decisions that directly reduce the labor-relations risk surface. [K1]

#### 3.1 What Kashi already gets right

- Metadata-only analysis. Kashi's current framing is built around structural interaction metadata such as turn timing, overlap, interruption, latency, and speaking-share, rather than content analysis or affect inference. This sharply improves defensibility relative to message-reading or sentiment-reading competitors. [K1]
- Deterministic detectors. Kashi states that its key detector layers are deterministic and explainable, with each review-worthy event traceable to turn IDs and timestamps. That helps rebut black-box arbitrariness and makes challenge procedures more credible. [K1]
- Restricted visibility. Upward visibility is described as aggregated and k-anonymized, with differential privacy on executive dashboards and auto-suppression for small teams or dominant-role situations. This is exactly the kind of restraint worker representatives and privacy reviewers look for. [K1]
- No HR decisions from the tool. Kashi already prohibits use in performance, promotion, discipline, and compensation decisions, and states that named-individual behavioural telemetry is not exposed to managers. Those are among the strongest current project safeguards. [K1]

- Refusal discipline. Kashi explicitly refuses a company-wide relationship-health score. That is important because broad headline scores create target gaming, compliance theatre, and '92 percent healthy' narratives that erase harmed minorities. [K1]
- Victim-owned evidence vault. The v2 concept where the employer stores encrypted snippets it cannot decrypt is a genuinely strong move. It directly reduces the risk that Kashi becomes a standing employer-readable archive of sensitive conversational material. [K1]
- Review-worthy events rather than harassment labels. The earlier concept note was already strong here, and the current Kashi materials preserve that logic. That is still the correct construct. [K2] [K1]

### 3.2 Where Kashi is still exposed

- The phrase 'not an employee-monitoring tool' overclaims. Kashi still creates person-linked longitudinal behavioural telemetry - just in a narrow, structural, better-governed way. A sceptical labor-side reader can therefore attack the phrase as rhetorical evasion. [K1]
- The 'CEO's instrument' language is commercially sharp but politically dangerous. It makes the product sound management-first precisely where the project most needs employee-trust and representative compatibility. [K1]
- The CEO roster / pattern-summary drill-down concept is still hot. Even with aggregation, the existence of an executive-facing roster view increases the risk that Kashi is read as management intelligence infrastructure rather than narrowly bounded governance infrastructure. [K1]
- The install flow remains too thin as a labor-relations story. A sequence such as platform picker -> permission manifest -> work-rules notice -> live is deployment choreography, not consultation doctrine. It does not yet show where worker representation enters, what gets documented, or what happens before launch. [K1]
- The necessity defence is underdeveloped. Analytics retention for 24 months, review-worthy events for 12 months, and per-speaker 90-day baselines may all be justifiable, but the project does not yet argue why each duration and each object is the least intrusive workable choice. [K1]
- The sales story still risks saying 'CEO first, legitimacy later'. The PDF openly says the different buyer is CEO + legal + union consent and that the product is sold directly to the CEO rather than HR. That may be economically correct, but it can read as top-down imposition if not balanced by an equally visible worker-facing governance layer. [K1]

### 3.3 Logical verdict

The research therefore does not support the claim that Kashi is unsafe or conceptually broken. It supports a narrower and more useful claim:

#### Verdict

- Kashi is consultation-possible, not consultation-proof.
- Its core architecture is much stronger than typical workplace-AI products.
- Its main remaining weakness is not detector logic; it is deployment politics and overstatement in a few current phrases.

## 4. The central logical tension the project must resolve

Kashi wants to make repeated asymmetries visible without becoming a standing behaviour-intelligence layer for management. That is the entire knife-edge.

Everything useful and dangerous about the project flows from this tension:

- If the system is too weak, it produces nothing actionable and collapses into generic engagement software.
- If the system is too broad, it becomes a surveillance product with a governance narrative attached.
- If the system is narrow but described badly, it may still be blocked by labor politics before its strongest safeguards are even understood.

The project should therefore optimise not for maximal analytical reach, but for a bounded form of visibility that is difficult to weaponise downstream.

## 5. Usable project recommendations

Below is the practical output of the research. These are the changes most likely to improve Kashi's real deployability without collapsing the concept.

### 5.1 Change the language before changing the architecture

Current or risky phrasing	Why it is weak	Recommended replacement
'Not an employee-monitoring tool.'	Too absolute. The system still processes person-linked behavioural telemetry over time.	'Not a general-purpose employee-surveillance product. Kashi is a tightly restricted governance system that processes only structural interaction metadata under explicit anti-misuse limits.'
'The CEO's instrument for seeing the bill before it arrives.'	Commercially sharp, but management-first and politically inflammatory.	'A bounded governance instrument that helps organisations surface repeated communication asymmetries before they harden into late-stage people incidents.'
'We sell to the CEO directly, not HR.'	Sounds like top-down adoption with worker legitimacy handled later.	'Requires an executive sponsor, legal review, and worker-representation alignment before pilot launch.'
'Not surveillance.'	Vulnerable to rebuttal if the system still structures longitudinal behavioural visibility.	'Narrow, metadata-only, consultation-compatible workplace-governance infrastructure.'

### 5.2 Add a deployment-preconditions section to the deck / governance page

- State that deployment requires prior employee explanation and representative consultation where applicable.
- State that Kashi is not deployable as a silent pilot layered underneath ordinary meeting software.
- State that permitted uses, prohibited uses, access rights, retention, challenge path, and anti-retaliation expectations must be documented before go-live.
- State that local labor / employment configuration is jurisdiction-specific rather than one global default.

### 5.3 Build a labor-consultation packet now

This is the single most useful project output to build next, because it turns the research into a deployment asset. Minimum contents:

- Purpose statement: what Kashi is and is not.
- Approved use cases: employee self-awareness, manager self-correction, aggregated hotspot visibility, structured escalation support.
- Prohibited use cases: performance review, promotion, discipline, compensation, ranking, redundancy planning, productivity surveillance.
- Access matrix by role, including what the CEO can and cannot see.

- Retention map with justification for each layer.
- Drill-down conditions, audit logging, and visibility of those logs to affected individuals.
- Transcript-error challenge process and event-review process.
- Anti-retaliation and no-silent-repurposing commitments.
- Pilot sunset / deletion rules and what happens if the pilot ends.

#### 5.4 Harden the product against misuse, not just the policy page

- Keep the no-export-to-HR-decision logic explicit. The strongest version is not just 'prohibited by contract' but 'not exposed in product pathways'.
- Do not allow customers to disable the k-anonymity / differential-privacy guardrails on executive views.
- Do not introduce named manager-to-manager comparison boards or broad behavioural leaderboards.
- Make every drill-down action auditable and visible to the affected individual, as the current Kashi direction already suggests. [K1]
- Treat the evidence vault as a worker-protection feature, not merely as a privacy nicety.

#### 5.5 Strengthen the necessity / proportionality argument

The team should be able to answer, in one or two precise paragraphs each, the following questions:

1. Why is a 90-day personal baseline necessary instead of a shorter moving window?
2. Why is 24-month analytics retention necessary instead of 12 months or less?
3. Why must review-worthy events persist for 12 months?
4. Why is executive visibility needed at all, and why is it limited in exactly the current way?
5. Why is metadata-only detection a less intrusive but still useful alternative to content-reading products?
6. Why are existing hotline / HR / training channels insufficient without pattern visibility?

These answers matter because regulators and worker representatives do not only ask whether a feature is helpful. They ask why a less intrusive alternative is not enough. [W5] [W6]

### 6. Pilot readiness checklist for the project team

Priority	Item	Why it matters	Deliverable
P0	Rewrite core positioning language	Reduces avoidable backlash caused by overclaiming or management-first rhetoric.	Updated landing / deck / governance copy
P0	Labor-consultation packet	Turns research into a real deployment artifact.	PDF / DOCX packet for legal + reps + sponsor
P0	Retention and necessity memo	Needed for proportionality defence and privacy review.	1-2 page internal note
P0	Japan work-rules / representative path	Needed for credible local deployment sequencing.	Process note + example annex
P1	Technical anti-misuse spec	Shows that misuse is product-constrained, not just discouraged.	Engineering spec
P1	Pilot challenge workflow	Needed for fairness and trust.	Flowchart + SOP
P1	Employee FAQ and manager FAQ	Prevents rollout from being explained only by leadership.	Two short documents
P2	Jurisdiction playbooks	Needed if the project moves beyond Japan or pitches EU orgs.	JP / DE / NL / UK one-pagers

## 7. Drop-in text usable for the project right now

### **Suggested subsection title:**

*Deployment preconditions: labor consultation and anti-monitoring safeguards*

Kashi is not a general-purpose employee-surveillance product. It is a tightly restricted governance system that processes only structural interaction metadata - such as turn timing, overlap, interruption, latency, and participation patterns - under explicit technical, contractual, and procedural limits. Deployment is conditioned on prior worker-facing explanation and representative consultation where applicable. Kashi is not for performance review, promotion, discipline, compensation, productivity scoring, or behavioural ranking. The default state is self-visibility and aggregated visibility, not universal visibility. Drill-down is gated, audited, and limited to justified workflows. Kashi is designed to surface repeated interaction asymmetries without becoming a standing management surveillance layer.

### **Short speaking line for partners / buyers:**

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

## 8. Final conclusion

The research result is not that Kashi should become softer, narrower, or more conventional. The result is that Kashi should become more precise about what it is doing institutionally.

The strongest version of the project is not 'AI detects harassment' and not 'the CEO gets a new dashboard'. The strongest version is: Kashi creates a narrow, metadata-only, consultation-compatible governance layer that makes repeated asymmetries visible without creating a general-purpose surveillance archive. That is the version most likely to survive legal review, labor politics, employee trust, and serious enterprise scrutiny at the same time. [K1] [K2] [W1] [W5] [W6]

## Source appendix

### Internal project sources

**[K1]** Kashi - Progress & Project Overview (2026-04-21)

**[K2]** meeting\_governance\_ai\_concept\_note.docx

### External sources reviewed

**[W1]** European Commission, Navigating the AI Act FAQ - <https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>

**[W2]** European Commission, AI Act policy page - <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

**[W3]** Germany, Works Constitution Act (BetrVG), English text - [https://www.gesetze-im-internet.de/englisch\\_betrvg/englisch\\_betrvg.html](https://www.gesetze-im-internet.de/englisch_betrvg/englisch_betrvg.html)

**[W4]** Dutch Data Protection Authority, works council and staff tracking systems - <https://www.autoriteitpersoonsgegevens.nl/en/themes/employment-and-benefits/works-council/the-works-council-and-staff-tracking-systems>

**[W5]** UK ICO, Employee monitoring - is it right for your business? - <https://ico.org.uk/for-organisations/advice-for-small-organisations/news-blogs-and-events/blogs/employee-monitoring-is-it-right-for-your-business/>

**[W6]** France CNIL, sanctions and corrective measures in 2025 - <https://www.cnil.fr/en/sanctions-and-corrective-measures-cnils-actions-2025>

**[W7]** European Data Protection Board, Guidelines 05/2020 on consent - [https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-052020-consent-under-regulation-2016679\\_en](https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-052020-consent-under-regulation-2016679_en)

**[W8]** Japan MHLW, sample filing for work rules change with attached employee representative opinion - [https://www.mhlw.go.jp/shinsai\\_jouhou/dl/kumiai-01.pdf](https://www.mhlw.go.jp/shinsai_jouhou/dl/kumiai-01.pdf)

**[W9]** Fukui Labour Bureau reference on majority representative selection and protections - <https://jsite.mhlw.go.jp/fukui-roudoukyoku/content/contents/001714436.pdf>

*Note: This memo is a project research summary, not legal advice. It is intended to improve product and rollout design by aligning Kashi's posture with the strongest available regulatory and institutional signals.*