

# Kashi

## Labor Politics / Union / Works-Council Readiness

Deployment memo focused on worker representation, procedural rollout, and governance terms

Prepared for project use • April 21, 2026

Purpose: deepen the labor-politics / union / works-council lens and convert it into something operational for Kashi.

### **Bottom line**

For Kashi, labor relations is not a side compliance topic. It is a deployment gate. A system can be technically careful and still get blocked if workers or their representatives believe it shifts workplace power upward without creating enforceable counterweights.

Scope note: this memo is strategic and operational, not formal legal advice. It is built for product, pitch, and rollout design.

## Executive summary

Kashi is already stronger than many adjacent workplace-AI products on the technical and governance side. The project has explicit refusals around content surveillance, affect inference, broad HR visibility, performance use, and company-wide scoring. Its detection path is deterministic and built around structural metadata rather than live black-box judgement. Those are real strengths, not cosmetic ones.

The missing layer is labor-governance design. In practice, deployment will not be decided by product logic alone. It will also be decided by whether unions, works councils, majority representatives, or informal employee reps conclude that the system is bounded, challengeable, and procedurally fair.

The practical risk is not only 'this may be unlawful.' The larger operational risk is: 'this changes observation, evidence, and escalation inside the workplace; therefore it changes power; therefore workers and their representatives require negotiation, audit rights, and pre-launch procedure.'

For Kashi specifically, the tension is sharper because the current project rhetoric sometimes sounds more top-down than the actual product architecture. The deck stresses that Kashi is governance infrastructure and even frames it as a CEO-facing instrument for seeing the labor-cost bill early. At the same time, the product's trust logic depends on mirrors-not-microscopes, self-view first, aggregate upward visibility, and procedural drill-down. That mismatch is manageable, but it needs to be handled explicitly.

### Decision

Treat labor-politics readiness as a first-class deployment workstream. Kashi should define: (1) non-negotiable governance terms, (2) negotiable rollout terms, (3) what employee representatives can audit, (4) what procedural steps are required before pilot launch, and (5) how the internal rollout language is phrased.

## What this memo recommends

- Do not rely on 'consent' as the core legitimacy story in employment settings. Use narrow purpose, representative consultation, enforceable limits, and challenge rights instead.
- Make anti-drift rules contractual and product-bound: no content analysis for employer access, no affect inference, no performance / promotion / discipline / compensation use, no named-individual telemetry to managers, no universal transcript browsing, no company-wide health bar.
- Give worker representatives audit rights over the rules of the system - retention, access matrix, detector documentation, change log, and incident reporting - not casual browsing rights over worker-level behavioral data.
- Require a pre-launch procedure before any customer pilot: representation check, policy package, consultation window, documented objections and responses, pilot review checkpoint, and expansion gate.
- Rephrase rollout and pitch language so Kashi is not heard as a CEO surveillance layer with ethics attached. The product must read as a constrained governance system whose value to leadership comes from the credibility of those constraints.

## 1. Why the labor-politics lens is a deployment gate

Founders often treat labor and representation issues as a downstream legal review problem. That is too naive. Once a system touches behavioral telemetry, workplace observation, evidentiary retention, or organizational escalation, it stops being 'just software' and becomes an industrial-relations object.

The relevant worker-side question is not merely whether the system is lawful. It is whether the system reorders power at work. If the employer gains a new structured view of interaction patterns, then

employee-side actors will ask what counter-balancing rights exist: who can inspect the rules, who can challenge outputs, who can see access logs, and what prevents scope creep.

That concern is not speculative. Official European guidance has already moved in this direction. The EU AI Act says that before employers put a high-risk AI system into service or use it in the workplace, they must inform workers' representatives and affected workers, following the relevant rules and procedures in Union and national law and practice (E1). In Germany, the Works Constitution Act gives works councils co-determination rights over the introduction and use of technical devices designed to monitor employee behaviour or performance (E2).

Japan has a different labor-relations structure, but the procedural logic is not absent. The Labor Standards Act requires the employer, when drawing up or changing rules of employment, to hear the opinion of the majority union or, if none exists, a person representing the majority of workers, and to attach that opinion document when filing the rules (E3). MHLW guidance also stresses that majority representatives must not be employer-selected and should be chosen through procedures such as a vote or show of hands after workers are clearly informed what the representative role is for (E4).

## 2. Kashi-specific diagnosis: where the project is strong, and where it is exposed

### What Kashi already gets right

- The current project materials sharply restrict scope: structural interaction metadata, not content surveillance; no affect or emotion inference; no performance / promotion / discipline / compensation use; aggregate upward visibility; audit trail; and a planned user-held evidence vault. Those design choices materially reduce creep and defensibility risk. (I1, I2)
- The product refuses the most obviously toxic expansion paths, including the company-wide relationship-health bar. That is strategically smart because it prevents headline-score dynamics and metric gaming. (I1)
- Kashi's architecture is unusually explainable for this category: deterministic detectors, own-baseline calibration, review-worthy events rather than moral labels, and bounded role-based presentation. (I1, I2)

### Where Kashi is still politically exposed

The current rhetoric sometimes sounds more top-down than the product really is. The project deck frames Kashi as governance infrastructure and, at several points, as a CEO-facing instrument for seeing the labor-cost bill before it arrives. That may help with sponsor seriousness, but it also makes a labor-side reading predictable: 'the institution is buying a new way to see workers, build records, and intervene.' (I1)

That creates a trust asymmetry. The product's actual trust model depends on mirrors-not-microscopes, self-view first, aggregate executive visibility, and strict drill-down constraints. But the buyer-facing framing can still make the system sound like upward institutional visibility first and worker protection second. (I1, I2)

So the problem is not that Kashi lacks safeguards. The problem is that those safeguards have not yet been translated into a labor-relations operating model. Without that operating model, even good constraints can be heard as optional, revocable, or marketing-only.

#### Practical implication

Kashi does not merely need a privacy posture. It needs a labor-governance posture: clear worker-side procedure, auditable rules, challenge rights, and a re-consultation trigger for any material scope change.

### 3. Likely objections from unions, works councils, and employee representatives

Objection	What it means in practice
Function creep	The system begins as bounded governance tooling, then quietly becomes performance, discipline, or productivity infrastructure. This is the single most predictable objection. If the product creates durable behavioral telemetry, labor-side actors will assume management will eventually try to reuse it unless the boundary is made binding.
Asymmetric visibility	The employer gains a new structured view of interaction patterns, while workers lack matching rights to challenge, inspect access, or understand why information was retained. Even aggregate dashboards can be read as top-down if workers cannot see the rules of drill-down and retention.
Representative bypass	Procurement is completed and technical rollout is planned before worker representation is seriously engaged. From a labor perspective, that reads as fait accompli consultation, which is much more likely to provoke resistance than true pre-launch dialogue.
Pseudo-consent	The company says workers 'consented' or that the pilot is 'optional' in a context where refusal is not realistically free. In employment contexts, consent is often viewed skeptically because of structural imbalance. EDPB guidance also emphasises necessity and less intrusive alternatives when relying on legitimate interests (E5).
Black-box governance	Even if the math is explainable, representatives may still object that retention, thresholds, suppressions, change control, and subcontractors are opaque. Explainable detector logic is necessary but not sufficient.
Narrative mismatch	The project publicly says it is anti-surveillance, but internal buyers understand it as a new leadership visibility layer. If the internal story and external worker story diverge too much, trust erodes fast.

### 4. Governance terms: what should be non-negotiable vs negotiable

Kashi needs to define this explicitly. If it does not, management will implicitly treat everything as fixed and worker-side actors will treat everything as suspect. The right split is: purpose and power boundaries are non-negotiable; cadence and implementation shape are negotiable.

Non-negotiable deployment terms	Negotiable implementation terms
• No content analysis for employer access	• Which teams pilot first
• No emotion, affect, voice-stress, or facial inference	• Whether the pilot starts with manager mirror only or also employee self-view
• No use in performance, promotion, discipline, or compensation decisions	• Exact retention durations within bounded ranges
• No named-individual behavioral telemetry to managers	• Monthly vs quarterly audit review cadence
• No universal transcript browsing by HR / leadership	• Whether audit summaries are shared only with reps or with all staff
• Audit trail on drill-down, viewable to the affected individual	• Which detectors are introduced first vs later
• Minimum group thresholds and suppression on aggregate views	• Whether evidence-vault features arrive in pilot or phase 2

Non-negotiable deployment terms	Negotiable implementation terms
• No company-wide relationship-health score	• Whether rollout starts country-by-country or business-unit-by-business-unit
• Worker challenge path for retained events and contextual correction	• How user education is sequenced before full activation
• Material change requires renewed consultation / re-approval	• What pilot success / stop criteria are used

## 5. What employee representatives should be able to audit

Representatives should not receive casual browsing rights over worker-level behavioral records. That would collide with Kashi's own anti-surveillance stance. But they should be able to audit the governance layer in a serious way.

The governing principle should be: employee representatives audit the rules of the machine, not the intimate details of ordinary workers' meeting lives.

Audit domain	Minimum audit package
Use policy	Approved purposes, prohibited purposes, prohibited decision contexts, role list for each access level
System boundary	What inputs are collected; what is not collected; detector list; content / affect exclusions; suppression rules
Retention and access	Retention schedule by layer; deletion rules; legal-hold triggers; access-log schema; event-retention criteria
Challenge process	How a worker contests a retained event; how transcript inaccuracies are handled; how confounds are recorded; review SLA
Change control	What counts as a material change; who approves it; version history of detectors and policy; rollback path
Third-party / security review	Subprocessor list; security architecture summary; incident-response path; high-level risk / DPIA summary where applicable

### Why this shape matters

This model is much more coherent than either extreme. It avoids giving representatives meaningless 'trust us' reassurances, but it also avoids turning worker-level telemetry into another broad-access archive.

## 6. What deployment should require procedurally before launch

### Japan-first minimum procedure

- Check whether the planned deployment requires drawing up or changing rules of employment. If yes, Article 90 procedure matters: hear the opinion of the majority union or majority representative and attach that opinion document when filing. (E3)
- Verify that worker representation is valid. MHLW guidance makes clear that majority representatives must not be employer-picked and should be elected through procedures such as a vote or show of hands after workers are informed of the purpose of the election. (E4)
- Provide representative infrastructure. MHLW materials say employers should give due consideration so the representative can smoothly perform legal duties; in practice this supports real communication channels and time to gather worker views. (E4)

- Issue a plain-language operating policy before activation: purpose, scope, exclusions, who sees what, retention, challenge rights, prohibited uses, and incident path.
- Run a bounded pilot with a scheduled worker-side review checkpoint before any expansion.
- Create an objection log and response log so concerns raised by representatives are documented, answered, and either accepted, mitigated, or explicitly rejected with reasons.

## EU / enterprise hardening procedure

- If a customer use case falls within high-risk workplace AI obligations, the EU AI Act requires employers to inform workers' representatives and affected workers before putting the system into service or using it. (E1)
- Where the deployment is likely to result in high risk to rights and freedoms, the customer will need a data protection impact assessment under GDPR Article 35. The AI Act itself points deployers toward provider information needed for that purpose. (E1, E6)
- If residual high risk remains after mitigation, prior consultation with the supervisory authority may be required under GDPR Article 36. (E6)
- In jurisdictions with stronger co-determination rules, such as Germany, customer rollout may also require works-council agreement when the system is treated as a technical device designed to monitor behaviour or performance. (E2)
- No enterprise rollout should proceed on procurement alone. There must be a worker-side procedural path proportionate to the local labor framework.

## 7. Rollout language: what to avoid, and what to say instead

Language matters more than founders think. A technically bounded system can still be received as surveillance if the rollout language sounds classificatory, managerial, or triumphantly visibility-first.

The product should never be introduced internally as a hidden-problem detector for leadership. It should be introduced as a constrained workplace-governance system with explicit exclusions and worker-facing rights.

Avoid	Preferred wording
'We are introducing AI to identify power harassment early.'	'This system does not determine harassment, intent, or legal wrongdoing. It produces limited structural signals from meeting interaction metadata.'
'Leadership will gain better visibility into communication risk.'	'Default visibility is self-view plus tightly bounded aggregate views. Any drill-down is role-limited, logged, and procedurally constrained.'
'This helps the CEO see hidden dynamics.'	'This helps the organization notice recurring interaction asymmetries earlier without broad content surveillance or automatic judgement.'
'We can now detect problematic behavior in meetings.'	'The system surfaces review-worthy patterns for human review; it does not output binary labels of abusive or non-abusive conduct.'
'Participation is voluntary; please consent if comfortable.'	'The system operates under defined policy terms, representative consultation, documented limits, and worker challenge rights; it is not justified by pseudo-voluntary consent.'

## 8. What Kashi should add to the project materials now

- A dedicated section titled something close to 'Labor relations / works-council readiness' in the concept note or deck.
- A one-page 'deployment terms' sheet separating non-negotiables from negotiables.
- A representative-audit pack listing what worker-side actors can inspect before launch and during pilot.
- A pre-launch procedure page: representation check, policy publication, consultation window, pilot review checkpoint, and expansion gate.
- A rollout-language sheet for customers so the system is not internally announced in a surveillance-sounding way.
- A material-change rule stating that any expansion of inputs, access, retention, or use cases triggers renewed consultation.

This is not just a documentation exercise. It changes the product category. Kashi should increasingly be described not merely as AI governance infrastructure, but as meeting-governance infrastructure with worker-auditable constraints.

That description is strategically stronger than generic ethics language because it converts 'trust us' into operational procedure.

## 9. Working launch-gate checklist for Kashi pilots

Gate	Minimum question	Status
Representation	Is there a recognized union, works council, or valid majority representative process?	Open
Policy scope	Have purpose, exclusions, prohibited uses, and role-based visibility been written in plain language?	Open
Access control	Can the team show who sees what, under what trigger, with what log?	Open
Challenge rights	Can a worker contest retention, context, or transcript accuracy?	Open
Consultation	Has there been a real consultation window before activation rather than post-procurement notice?	Open
Pilot bounds	Are pilot teams, duration, stop criteria, and review checkpoint fixed in advance?	Open
Change control	Is there a material-change trigger requiring renewed consultation?	Open
Cross-border hardening	For EU or Germany use, is there a DPIA / works-council path where applicable?	Open

### Recommended one-sentence position

Kashi should not be deployed as 'AI that lets leadership see hidden people risk.' It should be deployed as a constrained meeting-governance system whose legitimacy depends on narrow scope, worker-side procedure, auditable rules, and enforceable anti-drift boundaries.

## Appendix A. Source notes

### Internal project sources

- I1. Kashi - Progress & Project Overview (2026-04-21). Key points used here: governance-infrastructure framing; mirrors-not-microscopes principle; no-content / no-affect posture; no HR decisions; no company-wide relationship-health bar; CEO-facing budget framing.
- I2. Transparency That Drives Institutional Accountability - concept note. Key points used here: self-view and aggregate-view logic; review-worthy events; role-based visibility; no automatic judgement; anti-surveillance boundaries.

### External official / quasi-official sources

- E1. Regulation (EU) 2024/1689 (AI Act), Article 26(7): before putting into service or using a high-risk AI system at the workplace, employers as deployers must inform workers' representatives and affected workers.
- E2. German Works Constitution Act (BetrVG), Section 87(1)(6): co-determination rights over the introduction and use of technical devices designed to monitor employee behaviour or performance.
- E3. Japan Labor Standards Act, Article 90: when drawing up or changing rules of employment, the employer must hear the opinion of the majority union or majority representative and attach that opinion document when filing.
- E4. MHLW / Employment Consultation Center guidance on majority representatives: the representative should not be in a supervisory or management role, must not be chosen based on employer intention, and should be elected by methods such as vote or show of hands after workers are informed of the purpose; employers should give due consideration so the representative can perform their duties smoothly.
- E5. EDPB Guidelines 1/2024 on legitimate interests: processing is 'necessary' only if the pursued objective cannot reasonably be achieved as effectively by less intrusive means; otherwise the processing may fail the necessity test.
- E6. GDPR Articles 35 and 36 (EUR-Lex): DPIA is required for processing likely to result in high risk to rights and freedoms; prior consultation with the supervisory authority may be required if residual high risk remains.
- E7. ICO Employment Practices Code: consult workers and/or trade unions or other representatives about the development and implementation of employment practices involving worker personal data; before data-matching exercises, inform and discuss the plan with worker representatives.

### Why these sources matter operationally

Taken together, these sources do not merely say 'be transparent.' They support a stricter operational reading: systems that alter workplace observation or evidence must be deployed through procedure, not only through procurement.

That is why this memo recommends representative audit rights, a pre-launch policy package, a consultation checkpoint, and a material-change trigger. Those steps are not decorative. They are what makes Kashi look like a bounded governance system rather than an expandable management weapon.