

Operationalizing the AI Governance Framework

Four deployable governance artifacts

Prepared by	Wisdom Kwashie — CEO & Chief AI Officer
Organization	Momento Media, Inc. · Southern California
Artifacts	Technical Standards · Risk Register · Approval Gate · Decision Log
Date	June 30, 2026 (revised October 1, 2026)
Aligned to	AI Governance Policy, revised October 1, 2026

About this document

This document operationalizes Momento Media's AI Governance Policy by translating its principles into four concrete, deployable artifacts. Each activity states the scenario, presents the completed artifact, links it to the policy and to NIST AI RMF and ISO/IEC 42001, and explains the governance rationale so the controls can be understood, applied, and defended. Roles referenced (CEO & Chief AI Officer, AI Development & Validation Lead, Legal/Compliance/Risk & Audit Lead, Head of HR, CTO) are standing members of Momento Media's AI Governance Board.

Activity 1 — Translating Principles to Enforceable Standards

Principle: “AI systems must uphold user privacy and employ data minimization techniques.”

This principle is too abstract for engineers to act on. The standards below make data minimization measurable for a customer churn model, enforcing it across the full data lifecycle — collection, storage, and disposal.

#	Lifecycle stage	Standard	Technical Owner	Evidence
1	Disposal	The model pipeline must strip all personally identifiable information (PII) within 30 days of model-training completion; raw identifiers are purged from training storage on that schedule.	Data Engineering (under CTO, Mateo Reyes)	Purge-job log, reviewed monthly by the Legal, Compliance, Risk & Audit Lead.
2	Collection	The production feature set must be capped at 40 non-PII features. Any expansion beyond that count requires a formal exception (Exception Request Form) approved by the AI Governance Board.	Model Developer (AI Dev & Validation Lead, Aiko Tanaka)	Feature count checked at every release; each exception on file with the Board's decision.
3	Storage	All retained customer data in the feature store must be pseudonymized, with re-identification keys stored separately under restricted access.	Data Architect (under CTO, Mateo Reyes)	Key-store access list reviewed quarterly.
4	Storage	Retained customer data must also be aggregated to a minimum cohort size of 100 users (k-anonymity), so no group is small enough to single out one person.	Data Architect (under CTO, Mateo Reyes)	Cohort-size report run before each release; any cohort under 100 blocks the release.

Why these values: Thirty days leaves time to retrain or debug but keeps raw identifiers out of storage. The 40-feature cap limits collection to what the model needs, with any expansion justified through an Exception Request. A cohort of 100 is deliberately strict, so no group is small enough to point at one person.

Policy and framework link: Policy Sections 3 and 6; NIST AI RMF MEASURE 2.10, MAP 4; ISO/IEC 42001 Annex A.7.2–A.7.5.

Activity 2 — Classifying and Mitigating Risk (Risk Register)

Scenario: an AI model automatically screens every résumé for software-engineering roles and assigns a score that determines which candidates receive a human interview.

Field	Entry
Use Case Title	Automated Candidate Screening
Initial Risk Classification	High
Justification	The model makes a consequential decision affecting a person's livelihood and access to professional opportunity, and carries a high potential for algorithmic bias against protected characteristics (gender, race, age, disability). Under Momento Media's risk tiers it is High (impact on individuals' rights); hiring is treated as a high-risk use under the EU AI Act.
Likelihood (inherent)	High (3). Résumé-screening models learn from past hiring decisions and can reproduce their patterns; every candidate passes through the model.
Impact (inherent)	High (3). A wrong screen removes a candidate from consideration and affects their livelihood; the candidate cannot see why.
Risk Level	High (9) — Likelihood × Impact, using the scoring method in Appendix B of the policy.
Who is harmed if it fails	The candidate who would otherwise have been chosen had the model not failed.
Required Mitigation Step	1. Fairness testing across protected groups before and during use; the pass mark is the four-fifths rule (adverse-impact ratio of at least 0.80), as set in the Activity 3 gate. 2. Human-in-the-loop override — a recruiter reviews and can overrule every screening score.
Accountable Team for Mitigation	AI Development & Validation Lead (Aiko Tanaka) owns fairness testing; Head of HR (Grace Bello) owns the human-in-the-loop override process; the Legal, Compliance, Risk & Audit Lead (Priya Anand) independently reviews the fairness evidence.
Detection Method	Pre-deployment fairness testing; after deployment, monitoring of selection rates by group and review of the decision log (Activity 4).
Residual Risk	Medium after mitigation. Fairness tests cover the groups measured, not every possible group or combination of groups.
Risk Owner and Status	Head of HR (Grace Bello), as business owner of the use case. Open — mitigation required before deployment (see Activity 3). Logged June 30, 2026.

Policy and framework link: Policy Section 5 (risk tiers) and Appendix B; NIST AI RMF MAP 5 and MANAGE 1; ISO/IEC 42001 clauses 6.1.2, 6.1.3, 8.2 and 8.3.

Activity 3 — Designing the Final Gate (Approval & Oversight Workflow)

Scenario: the High-risk Automated Candidate Screening model is about to be deployed. The pre-deployment gate below requires four sign-offs, each from a different functional area, before the model is activated. Steps 1–3 are pre-reviews; step 4 is the full AI Governance Board decision that the policy requires for High-risk use cases.

#	Approval Step	Sign-off Required	Check
1	Model Validation Review (Technical)	AI Dev & Validation Lead — Aiko Tanaka	Verify all bias and performance tests pass required thresholds — e.g., adverse-impact ratio ≥ 0.80 (four-fifths rule) across every protected group, and accuracy/precision at or above the baseline recorded in the model card before testing began.
2	Legal & Compliance Check (Legal)	Legal, Compliance, Risk & Audit Lead — Priya Anand	Confirm intended use, candidate-facing transparency notice, and required disclosures (incl. opt-out / AI-use notice) meet CCPA and EU AI Act obligations. Independently review the step 1 fairness evidence, because the model's builder cannot be its only validator. If the tool is third-party, confirm it is on the Approved AI Tools Register and covered by a data-processing agreement that prohibits training on candidate data (Data Entry Rule).
3	Business Readiness Review (Business)	Head of HR — Grace Bello	Accept the residual business risk and confirm resources are committed for the Post-Deployment Monitoring Plan and the incident-response / model-quarantine procedure.
4	Final Go / No-Go Decision (Governance)	AI Governance Board, chaired by the CEO & Chief AI Officer — Wisdom Kwashie	Confirm steps 1–3 are formally signed off and all policy/regulatory requirements, including the monitoring plan, are in place. The Board decides by vote, the vote is recorded, and the chair issues Go or No-Go.

If any step fails: A No-Go returns the model to development. The fix is re-tested from step 1, and the decision and its reasons are recorded.

Policy and framework link: Policy Section 5 (High risk needs full Board sign-off before deployment) and Section 9; NIST AI RMF MANAGE 1.1 and MEASURE 1.3; ISO/IEC 42001 Annex A.6.2.5.

Activity 4 — Template for Accountability (Consequential Decision Log)

Scenario: regulation requires that every high-risk AI decision affecting a person's access to a benefit (a job, a loan, a service) be logged and auditable. The template below defines the mandatory fields for one log entry. Every screening decision from the model writes one row — approvals, denials and referrals alike — so selection rates by group can be measured, and denials are flagged for priority review.

Field Name	Data Type / Format	Purpose
1. Unique Decision ID	Alphanumeric string	Primary key — lets anyone reference and retrieve this exact decision.
2. Decision Outcome	Categorical (Approved / Denied / Referred to Human)	Records the final action taken by the system or human.
3. Decision Timestamp	UTC timestamp	The exact moment of the decision, for chronological, tamper-evident auditing.
4. Explanatory Factors	JSON array of key-value pairs	The top 3–5 factors driving the outcome (e.g., SHAP/LIME values) — the reason behind the decision.
5. Human Override Flag	Boolean (True / False)	Whether a human reviewer changed the AI's recommendation — proof the human-in-the-loop operated.
6. Model Version	Alphanumeric string (e.g., v1.3.2)	Identifies which version of the model made the decision, so the decision can be reconstructed. Required by the policy's logging rule (Section 7).
7. Reviewer ID and Override Reason	Text; required when the Human Override Flag is True	Records who changed the recommendation and why.

Log controls: The log is append-only: a correction adds a new row, and nothing is edited or deleted. Access is restricted to the Legal, Compliance, Risk & Audit Lead, the AI Development & Validation Lead and auditors. Entries are kept for the retention period set by Legal.

Why it matters: The candidate who would otherwise have been chosen never sees the model's decision. The log, with its explanatory factors and override flag, is how a reviewer can find that candidate and correct the outcome.

Policy and framework link: Policy Section 7 (Logging & Audit Trails); NIST AI RMF MEASURE 2.4 and MANAGE 4; ISO/IEC 42001 Annex A.6.2.6 and A.6.2.8.