

Artificial Intelligence Governance Policy -

MOMENTO MEDIA

**Prepared by W.N. Kwashie
July 2026**

MOMENTO MEDIA

Digital Marketing · New Media · Production

Artificial Intelligence Governance Policy

Responsible, transparent, and legally compliant AI across our creative, marketing, and production work.

Organization	Momento Media, Inc. · Southern California
Effective Date	July 30, 2026
Version	1.0
Approval	Wisdom Kwashie — Chief Executive Officer & Chief AI Officer
Review Cycle	Quarterly, or sooner upon regulatory change, technology shift, or incident

Contents

1. Purpose	4
2. Scope & Applicability	5
3. Guiding Principles	5
4. Governance Structure & Roles	7
5. Risk Classification & Approval Workflow	9
6. Data & Model Controls	9
7. Monitoring, Logging & Incident Management	10
8. Training, Awareness & Change Management	11
9. Procurement, Vendor & Third-Party AI	11
10. Exceptions, Non-Compliance & Enforcement	12
11. Policy Review & Updates	12
12. Glossary & Definitions	12

Appendices14

Appendix A — Model Card: “Solara Glow”	14
Appendix B — Risk Assessment: “Solara Glow”	16
Appendix C — Audit Checklist: “Solara Glow”	17
Appendix D — AI Use Case Registration Form (Template)	18
Appendix E — Exception Request Form (Template)	19
Appendix F — Framework Alignment (NIST AI RMF, ISO/IEC 42001)	20

1. Purpose

This Artificial Intelligence Governance Policy sets out the principles, roles, responsibilities, and processes governing how Momento Media develops, procures, deploys, monitors, and retires Artificial Intelligence (AI) systems. Momento Media is a digital marketing and media production company specializing in digital marketing strategy, content strategy, social media campaign planning and asset creation, AI content generation for storytelling and campaigns, and photo and video production for brand storytelling and marketing campaigns. Our services also include local SEO, Answer Engine Optimization (AEO), and Generative Engine Optimization (GEO).

Because AI is woven directly into our creative output — from generative image, video, and copy tools to custom-trained models built for client campaigns — this policy exists to ensure that AI is used responsibly, ethically, transparently, and in compliance with the law, maximizing creative and commercial benefit while protecting our clients, their audiences, our people, and the public from harm.

1.1 Who This Policy Protects

AI governance only matters if it protects real people. These are the groups this policy is written to protect, how each could be harmed if it fails, and where the policy gives that protection.

Who could be harmed	How, and where this policy protects them
People whose likeness or resemblance appears in AI-generated content	They can be misrepresented without their knowledge or consent. Protected by: High-tier approval for AI likenesses of real people (5.1), the Human Review Rule, the likeness check and takedown route (Section 6), and consented training data (Section 6).
Audiences and consumers	They can be misled by AI imagery that looks real or implies product results. Protected by: the Disclosure & Authenticity principle (Section 3), the Human Review Rule and disclosure rule (Section 6), and Marketing and Legal review.
Clients, and the people in their documents	Confidential or personal material can be exposed by entering it into online AI tools. Protected by: the data entry rule (Section 6), tool approval and Board escalation (Section 9), and staff training (Section 8).
Staff and department heads	They carry the professional and reputational damage when AI is misused on their watch. Protected by: clear roles and escalation (Section 4), proportionate sanctions (Section 10), and good-faith reporting (Section 7).
Artists and creators whose work may sit in a base model	Their work may have been used without consent. Protected by: base-model licence and provenance checks (Section 6) and the Intellectual Property & Copyright principle (Section 3).

2. Scope & Applicability

- This policy applies to all employees, contractors, freelancers, vendors, and partners who design, develop, procure, manage, or interact with AI systems, or who produce AI-enabled deliverables on behalf of Momento Media.
- It covers AI used in Momento Media's internal operations and AI used to create client-facing deliverables — including content produced for and delivered to clients.
- It covers the full AI lifecycle: planning, design, data sourcing, training, validation, deployment, operation, monitoring, audit, and retirement or replacement.
- It applies to both internally built AI systems (including custom-trained models) and externally procured or third-party AI services.
- It applies across all functions — creative, production, marketing strategy, SEO/AEO/GEO, accounts, HR, finance, and IT.

2.1 Regulatory & Framework Alignment

Momento Media adopts the NIST AI Risk Management Framework and ISO/IEC 42001 as the operating backbone of this policy. We additionally monitor and comply, where applicable, with the California Transparency in Frontier Artificial Intelligence Act (SB 53) and the EU AI Act. Most obligations under SB 53 and the EU AI Act attach to large frontier developers or high-risk systems; Momento Media assesses applicability on a case-by-case basis and complies where the law reaches our activities. Appendix F shows where each framework applies in this policy and how.

2.2 Scope & Limitations

Version 1.0 is written for a small firm. It sets quantitative targets only where they can be measured today (spend triggers, training cadence, review frequency and model card targets) and relies on the quarterly review to recalibrate tiers and thresholds as the business grows. Where one person holds more than one Board seat, that person may not audit, vet, or approve a system they built or operate; the Board engages an external reviewer for that system.

3. Guiding Principles

All AI activity at Momento Media is governed by the following principles.

Principle	What it means at Momento Media
Accountability & Oversight	Every AI initiative has a clearly identifiable, accountable owner; the AI Governance Board oversees decisions.
Transparency & Explainability	AI systems are as transparent as feasible; where AI materially shapes a decision or deliverable, we can explain how.

Principle	What it means at Momento Media
Fairness & Non-Discrimination	Models and outputs are tested to avoid unfair bias and to avoid discriminating against protected groups.
Privacy & Data Protection	Personal and client data is used in line with privacy law (e.g., CCPA, GDPR) and our data policies, using minimization, de-identification, and encryption.
Robustness & Safety	AI systems are tested for adversarial and edge-case risk and monitored for performance drift and errors.
Human-in-the-Loop	For client-facing, public-facing and high-impact outputs, a human reviews and can override AI before release.
Compliance & Legal Conformity	AI use complies with applicable laws — privacy, IP/copyright, advertising, discrimination, and sectoral rules.
Continuous Monitoring & Adaptation	AI systems are monitored in operation, with periodic review, audit, and recalibration.
Disclosure & Authenticity	We are transparent with clients and, where appropriate, audiences when content is AI-generated, and never pass off synthetic media in a deceptive way.
Intellectual Property & Copyright	We respect third-party IP in training data and outputs, and we make ownership and usage rights of AI outputs explicit with clients.
Client Consent & Brand Safety	Client data, brand assets, and likenesses are used only as authorized, and AI outputs are screened to protect each client's brand.

4. Governance Structure & Roles

Momento Media governs AI through a single, cross-functional AI Governance Board chaired by the CEO & Chief AI Officer. The Board has seven standing members, each owning a distinct function. Project-level Business Owners participate only when their use case is under review; they do not hold standing seats.

4.1 AI Governance Board

- **Mandate:** Provide oversight, approve High-risk AI use cases before deployment, set and maintain this policy, and adjudicate escalations and exceptions.
- **Composition:** The seven standing members listed in 4.2, chaired by the CEO & Chief AI Officer.
- **Authority:** Final approval for AI systems above defined risk thresholds, authority to grant time-limited exceptions, and authority to suspend any non-compliant AI system.

4.2 Roles & Responsibilities

Role	Holder	Responsibilities
CEO & Chief AI Officer (Chair)	Wisdom Kwashie	Final accountability for AI governance; chairs the Board; owns this policy; assigns risk tiers; final go/no-go on AI tools and High-risk use cases; owns the Approved AI Tools Register / AI inventory.
Chief Technology Officer	Mateo Reyes	AI infrastructure and IT on which models run; security reviews; vets and validates third-party tools (security + license).
AI Development & Validation Lead	Aiko Tanaka	Design, training, and validation of custom/proprietary models; training-data provenance; testing; model cards; first responder for AI-output incidents.
Legal, Compliance, Risk & Audit Lead	Priya Anand	Legal and regulatory review (IP/copyright, privacy, advertising); contracts and consents; independent audit of AI systems against this policy (an external reviewer or a Board member who did not vote to approve it audits any system the Lead voted to approve).
Head of Marketing	Sofia Martinez	Client-facing AI use; brand safety; AI-content disclosure to clients/audiences; authority to pull published client or public content.
Chief Financial Officer	Daniel Okafor	AI procurement and vendor budgets; approves spend; financial-risk assessment of AI initiatives.

Role	Holder	Responsibilities
Head of HR	Grace Bello	Employee AI use; role-based training; workforce impact and employee-data protection.

- **AI Project Sponsor / Business Owner (per-project):** For each AI use case, the initiating team appoints a Business Owner accountable for aligning the use case with client and company goals, ensuring compliance, and owning the use case's risk mitigations. The Business Owner brings the use case to the Board for review but is not a standing Board member.

***Note on audit independence:** Audit is deliberately held by the Legal, Compliance & Risk seat rather than the CTO, so that the function reviewing AI systems is independent of the function that builds and operates them. Momento Media is too small to maintain an internal audit department, so the person who checks is never the person who built or operates the system under review. Where the Legal, Compliance & Risk Lead voted to approve a system, that system is audited instead by an external reviewer or by a Board member who did not vote to approve it.*

5. Risk Classification & Approval Workflow

5.1 Risk Tiers & Approval Paths

Every AI use case is classified Low, Medium, or High based on its impact on individuals, the sensitivity of the data involved, the degree of automation, and exposure to clients or the public. The tier sets the approval path.

Tier	Criteria	Momento Media examples	Approval path
Low	Internal use only; no decision-making; no client or public exposure.	Drafting internal copy, brainstorming, meeting summaries, internal research.	Team lead approval + log in AI inventory.
Medium	Client-facing AI content with human review before release; limited individual impact.	Campaign visuals, social posts, AEO/GEO/SEO copy, AI-assisted photo/video edits.	AI Development & Validation Lead approval + notify the Board.
High	Legal, financial, or rights impact; custom models; or sensitive data.	Custom/proprietary image-generation models; client/customer data used for paid marketing campaigns; AI likenesses of real people; any request to publish AI content without human review (not permitted except by written exception under 5.2).	Full AI Governance Board sign-off before deployment.

The tier governs the use case; Section 9 governs the tool used. A use case needs both approvals, and where the two paths differ the stricter one applies.

5.2 Exception & Exemption Process

A team may request an exception to a control via a written Exception Request (Appendix E) that states the justification, the risk involved, and a mitigation plan. Exceptions are reviewed by the AI Governance Board and, if granted, are time-limited and recorded in the AI inventory.

6. Data & Model Controls

- **Training-Data Sourcing (Moderate standard):** Custom models are trained primarily on licensed, owned, or client-provided and consented assets, supplemented by public-domain material and properly licensed stock. Any other source is reviewed case-by-case by the Legal, Compliance & Risk Lead before use.

- **Open-Source Base-Model Vetting:** Where a custom model is built on an open-source foundation model, two checks are required before adoption: (a) the model's license must permit commercial use and client redistribution; and (b) the base model's own training-data provenance must be assessed for IP/copyright risk. For High-risk client deliverables, prefer base models with documented or cleaner training data and disclose any residual provenance risk to Legal.
- **Data Privacy & Anonymization:** Use de-identification or pseudonymization where possible; comply with CCPA, GDPR, and our data policies.
- **Dataset Segmentation:** Maintain strict separation of training, test, and holdout data to guard against leakage.
- **Model Documentation (Model Card):** Each custom AI system has a model card detailing purpose, limitations, training data, base model, metrics, maintenance schedule, and known risks (Appendix A).
- **Explainability:** Where feasible, apply interpretability techniques and require a narrative justification for material AI-driven decisions.
- **Validation & Testing:** Pre-deployment validation, robustness testing, edge-case simulation, and adversarial testing, owned by the AI Development & Validation Lead.
- **Change Control & Versioning:** Version-control model code, data, and configurations; significant changes require review.
- **Data Entry Rule:** Confidential information, Client-Restricted information, and personal data are never entered into a third-party AI tool unless the tool is on the Approved AI Tools Register for that class of data and is covered by a signed data-processing agreement that prohibits training on Momento Media or client data. Confidential means unreleased campaigns, pricing, financials and internal strategy; Client-Restricted means client assets, data and documents covered by a client agreement; personal data means information about an identifiable person.
- **Human Review Rule:** Every client-facing or public-facing AI-generated asset is reviewed and approved by a named person before release, so that errors introduced by AI are caught before the public sees them. Skipping review requires a written exception under 5.2.
- **Likeness Check:** Before release, a reviewer checks any output that shows people for resemblance to an identifiable real person. If a resemblance is found or reported, the content is withdrawn and handled as an incident under Section 7. Anyone who believes they are depicted can contact the Head of Marketing, who can pull published content within one business day.
- **Disclosure Rule:** Clients are always told when content is AI-generated. Audiences are told whenever the content shows people, places, or events realistically enough to be mistaken for a photograph or recording, or shows product results.

7. Monitoring, Logging & Incident Management

- **Continuous Monitoring:** Track model performance and watch for drift, bias shifts, quality degradation, and outputs that deviate from expectations or brand.
- **Logging & Audit Trails:** Maintain logs of inputs, outputs, model versions, human overrides, and approvals sufficient to reconstruct how a deliverable was produced.
- **Incident Reporting & Ownership:** Staff report any suspected AI problem to the AI Development & Validation Lead (Aiko Tanaka), who triages. The Head of Marketing (Sofia Martinez) has authority to pull

published client or public content. The Legal, Compliance & Risk Lead (Priya Anand) is looped in on every incident and leads decisions on any matter involving copyright, privacy, or liability.

- **Good-Faith Reporting:** Staff who report a suspected AI problem in good faith are not penalised for reporting it.
- **Incident Response:** Procedures cover containment (rollback, takedown, human intervention), root-cause analysis, corrective action, and any required notifications (see Section 10).
- **Periodic Audit:** The Legal, Compliance, Risk & Audit Lead audits AI systems at least annually for policy adherence and emerging risk. Where the Lead voted to approve a system, an external reviewer or a Board member who did not vote to approve it audits that system instead.

8. Training, Awareness & Change Management

- **Training Cadence (twice a year):** All staff complete (1) a general annual AI policy training, and (2) a second mid-year session covering new tool developments, emerging risks, and updates to this policy. New hires complete the AI policy module at onboarding.
- **Role-Based Content:** Training is tailored for creative/production staff, marketers, executives, and the Legal/audit function.
- **Awareness:** Ongoing communication of AI use cases, responsibilities, and reporting channels to all staff.
- **Change Management:** Significant new AI systems require stakeholder review, a named adoption champion, and feedback loops.

9. Procurement, Vendor & Third-Party AI

No new third-party AI tool is used on client work until it is approved and entered into the Approved AI Tools Register.

- **Vetting (CTO):** Mateo Reyes reviews each tool for security and license terms.
- **Spend Approval (CFO):** Daniel Okafor approves spend.
- **Governance Sign-off (CAIO):** Wisdom Kwashie assigns the tool's risk tier and gives final go/no-go, and owns the Approved AI Tools Register. Any tool that is client-facing, touches client or personal data, or exceeds the spend threshold requires CAIO sign-off.
- **Board Threshold:** Any AI tool exceeding USD \$2,500 per year, total AI spend exceeding USD \$10,000 per year, or any tool touching client/personal data escalates to the full AI Governance Board. The CFO reports AI spend to the Board each quarter. The per-tool figure is set so that a top-tier multi-model creative subscription always triggers Board review.
- **Low-Risk / Free Tools:** Vetted by the CTO, logged in the Register, with the CAIO notified — no full approval required.
- **Due Diligence & Contracts:** Vendors must share model documentation, risk assessments, and data-provenance information; contracts should secure SLAs, audit rights, liability terms, data-handling obligations, and the right to inspect, update, or disable the model.
- **Ongoing Review:** Third-party systems are monitored post-integration for performance, drift, and policy adherence.

10. Exceptions, Non-Compliance & Enforcement

- **Exception Process:** Policy waivers require a written request with justification, risk assessment, and mitigation plan, reviewed by the Board and granted only on a time-limited basis (Appendix E).
- **Escalation:** Suspected non-compliance is escalated to the AI Governance Board and may lead to suspension of the AI system.
- **Corrective Actions:** Remediation may include retraining, content takedown, system decommissioning, or client remedies where applicable.
- **Sanctions:** Violations may result in warnings, additional training, role restrictions, or disciplinary action proportionate to the severity.

10.1 Notification Requirements

Notification obligations follow the frameworks Momento Media aligns with and the laws that apply to a given incident:

- **Client notification — mandatory when:** a violation involves (a) personal or client data covered by CCPA or GDPR, or (b) published client deliverables. In these cases the affected client is notified as a required corrective step. The incident is logged within 24 hours of Momento Media becoming aware of it, and the client is notified without undue delay and no later than 72 hours after awareness, consistent with the GDPR 72-hour benchmark.
- **Authority notification:** Where California SB 53 or the EU AI Act applies to the system in question, notify the relevant authority within the timelines those laws require. Most SB 53 and EU AI Act obligations attach to large frontier developers or high-risk systems; applicability is assessed per incident, and the assessment and its outcome are recorded in the incident log.
- **Other violations:** Handled internally with documented root-cause analysis and corrective action.

This section provides general information, not legal advice. Specific notification decisions are confirmed with counsel.

11. Policy Review & Updates

- **Review Cycle:** Quarterly review, because AI tools and regulation change faster than an annual cycle can track; sooner upon regulatory change, technology shift, or after an incident.
- **Version Control:** Maintain a record of versions, changes, and authors.
- **Stakeholder Feedback:** Solicit input from creative, production, marketing, legal, technical, and client-facing teams for continuous improvement.

12. Glossary & Definitions

Term	Definition
AI System	Software that performs tasks normally requiring human intelligence —

Term	Definition
	e.g., generating images, video, or text, or making predictions.
Custom / Proprietary Model	A model that Momento Media trains or fine-tunes, often on a client's brand assets, for a specific use.
Base / Foundation Model	A pre-trained model (often open-source) used as the starting point for a custom model.
Model Card	Structured documentation of a model's purpose, data, performance, limitations, and risks.
Model Drift	Gradual decline in a model's accuracy or quality as conditions change over time.
Explainability	The degree to which a system's outputs or decisions can be understood and justified.
Bias	Systematic, unfair skew in data or outputs that disadvantages particular groups.
Protected Class	A group protected from discrimination by law (e.g., race, gender, age, disability).
Adversarial Attack	A deliberate attempt to manipulate an AI system into producing wrong or harmful outputs.
Human-in-the-Loop	A workflow in which a person reviews and can override AI outputs before they are used.
Provenance	The documented origin and rights history of data or a model.
AEO	Answer Engine Optimization — optimizing content to be surfaced by answer engines and AI assistants.
GEO	Generative Engine Optimization — optimizing content to be cited and used by generative AI systems.

Appendices

Appendices A–C are worked examples built around a single representative High-risk use case. Appendices D–E are reusable templates.

Worked use case — “Solara Glow” custom image-generation model. Solara is a Southern California clean-beauty / skincare brand. Momento Media fine-tuned a custom image-generation model on Solara's licensed product photography and brand assets to produce campaign imagery for Solara's summer “Sunlit Skin” campaign. Classified High risk (custom model + client assets + client-facing publication).

Appendix A — Model Card: “Solara Glow”

Field	Detail
Model name / version	Solara Glow v1.0
Owner / Business Owner	AI Development & Validation Lead (Aiko Tanaka); Business Owner: Solara account lead
Purpose	Generate on-brand campaign imagery (product + lifestyle scenes) for Solara's “Sunlit Skin” summer campaign.
Base / foundation model	Open-source diffusion base model; license verified to permit commercial use and client redistribution (CTO + Legal).
Training data	Solara-provided licensed product photography and brand assets (consented); supplemented by properly licensed stock. No scraped or unlicensed web data.
Data provenance & rights	All training assets licensed or client-owned with written consent; output ownership and usage rights defined in the Solara contract.
Intended use	Human-reviewed marketing imagery delivered to Solara for paid and organic campaign use.
Out-of-scope use	Generating real-person likenesses; medical/efficacy claims; publication without human review; reuse for other clients.
Evaluation metrics	Targets set by this policy: at least 90% of sampled outputs pass brand-fidelity review at first review; prompt adherence of at least 85% on a fixed 20-prompt test set; manual brand-safety screening of a random 10% sample of each batch (minimum 25 images); 100% of delivered assets human-reviewed before release. These are targets, not measured results; results are recorded in the AI inventory at each milestone review.
Known limitations & risks	Possible visual artifacts; residual IP/provenance risk inherited from base model; risk of off-brand or implausible product depictions.

Field	Detail
Human oversight	Every delivered asset reviewed and approved by a human before client delivery (Head of Marketing sign-off).
Maintenance schedule	Reviewed at campaign milestones and retired at campaign close; monitored for drift during the run.
Risk tier / approval	High — approved by full AI Governance Board prior to deployment.

Appendix B — Risk Assessment: “Solara Glow”

Scoring method: Likelihood and Impact are each rated Low (1), Medium (2) or High (3); Level is Likelihood × Impact (1–2 Low, 3–4 Medium, 6–9 High). Ratings are inherent, before mitigation; the residual rating is given below the table.

Risk	Likelihood	Impact	Level	Basis for rating	Mitigation
IP/copyright exposure from base-model training data	Medium	High	High (6)	Opaque training data; outputs are published in paid campaigns, so a claim is costly.	Vet base-model license + provenance; prefer documented base; Legal sign-off; indemnity terms in client contract.
Off-brand or implausible product depiction	Medium	Medium	Medium (4)	Output varies by run; mandatory pre-delivery review limits the harm to rework.	Brand-fidelity review; mandatory human approval before delivery; prompt guidelines.
Unauthorized real-person likeness	Low	High	Medium (3)	Excluded by policy and training data; a miss would be a right-of-publicity claim.	Out-of-scope by policy; screening of outputs; no real-person training data.
Misleading / unsubstantiated claims in imagery	Low	High	Medium (3)	Imagery only, no efficacy claims by policy; advertising-law exposure if one slips.	Marketing + Legal review for advertising compliance; no efficacy claims.
Client data mishandling	Low	High	Medium (3)	Assets stay in a restricted project store, not a third-party platform; a breach would breach the contract.	Consented, licensed assets only; encryption; access limited to the project team.
Failure to disclose AI-generated content	Medium	Medium	Medium (4)	Depends on staff applying the rule every time; client and audience trust at stake.	Disclosure to client per Disclosure & Authenticity principle; audience disclosure where appropriate.
Model drift over campaign run	Low	Medium	Low (2)	Single campaign run with checkpoint reviews and retirement at close.	Monitoring during run; checkpoint reviews; retire at campaign close.

Overall residual risk: Medium after mitigation, driven by base-model provenance (IP), which can be assessed but not fully verified. Approved for deployment by the AI Governance Board, conditional on human review of every delivered asset and Legal confirmation of base-model licensing.

Appendix C — Audit Checklist: “Solara Glow”

Used by the Legal, Compliance, Risk & Audit Lead to verify the use case against this policy. Where the Lead voted to approve the use case, the independent reviewer named in Section 4 completes this checklist instead.

Audit item	Y / N	Notes
Use case registered in the AI inventory with correct risk tier (High)?		
Board approval recorded before deployment?		
Base-model license verified for commercial use + client redistribution?		
Base-model training-data provenance assessed and residual risk disclosed to Legal?		
All training assets licensed / client-owned with written consent?		
Training/test/holdout data kept separate?		
Model card complete and current (Appendix A)?		
Every delivered asset human-reviewed and approved before client delivery?		
AI-generated nature disclosed to the client?		
Logs retained (inputs, outputs, versions, approvals)?		
Output ownership / usage rights defined in the client contract?		
Incident-reporting path communicated to the project team?		
Model retired / decommissioned at campaign close?		
Likeness check completed for every output that shows people?		
Disclosure rule applied (clients told; audiences told where required)?		
Data entry rule followed (no Confidential, Client-Restricted or personal data entered outside Register tools)?		

Appendix D — AI Use Case Registration Form (Template)

Field	Entry
Use case name	
Business Owner / Sponsor	
Date submitted	
Description & objective	
AI tool(s) / model(s) used	
Data used (and source / consent)	
Client-facing? (Y/N)	
Proposed risk tier (Low / Med / High)	
Human review point(s)	
Approval (name / role / date)	

Appendix E — Exception Request Form (Template)

Field	Entry
Requestor / team	
Date	
Control / policy section being excepted	
Justification	
Risk introduced	
Mitigation plan	
Requested duration	
Board decision (approve / deny)	
Conditions & expiry date	
Reviewed by (name / date)	

Appendix F — Framework Alignment (NIST AI RMF, ISO/IEC 42001)

This appendix shows where NIST AI RMF 1.0 and ISO/IEC 42001:2023 apply in this policy, and how. References use each framework's own identifiers: NIST functions, categories and subcategories, and ISO clauses and Annex A controls.

Policy section	NIST AI RMF 1.0	ISO/IEC 42001:2023	How Momento Media applies it
1–2. Purpose, scope, alignment	GOVERN 1 (policies, processes and procedures); MAP 1 (context)	Clauses 4.1, 4.3, 5.2; Annex A.2.2	States why AI is used, who and what the policy covers, the full lifecycle in scope, and the frameworks followed. It is the AI policy the standard requires.
1.1 Who this policy protects	MAP 5 (impact characterization); GOVERN 4 (risk culture)	Clause 6.1.4; Annex A.5.4	Names the people who could be harmed (individuals depicted, audiences, clients and the people in their documents, staff, creators) and where each is protected.
3. Guiding principles	GOVERN 1.2 (trustworthy characteristics built into policy)	Clause 6.2; Annex A.6.1.2, A.9.3	The eleven principles are the policy's objectives for responsible development and use. Each is tied to a control in Sections 4–10.
4. Governance and roles	GOVERN 2 (accountability structures); GOVERN 3.2 (human–AI oversight roles)	Clauses 5.1, 5.3; Annex A.3.2	One AI Governance Board chaired by the CEO & Chief AI Officer, seven standing seats, a named owner for each function. Audit is held apart from those who build and operate AI, so nobody checks their own work.
5. Risk tiers and approval	MAP 2 (system categorization); MAP 5.1; MANAGE 1 (prioritization and response)	Clauses 6.1.2, 6.1.3, 8.2, 8.3; Annex A.6.2.5	Every use case is rated Low, Medium or High on impact, data sensitivity, automation and exposure. The tier sets the approval path; High needs full Board sign-off before deployment. Appendix B scores each risk (likelihood × impact).
6. Data controls	MAP 4 (component and third-party risk); MEASURE 2.10 (privacy risk); GOVERN 1.1 (legal requirements)	Annex A.7.2–A.7.5, A.4.3, A.9.4	Training data, privacy and the Data Entry Rule. Custom models train on licensed, owned or consented assets and datasets are kept separate. The rule keeps confidential, client-restricted and personal data out of unapproved AI tools, and use stays within the intended purpose.
6. Model controls	MAP 4.1, 4.2; MEASURE 1 (methods and metrics); MEASURE 2.9 (explained, validated, documented)	Annex A.6.2.2–A.6.2.4, A.6.2.7, A.8.2	Base-model vetting, model card, validation, change control. Base models are checked for licence and training-data provenance. Each custom model has a model card (Appendix A) with targets. The AI Development & Validation Lead owns validation, and changes are version-controlled.
6. People-facing controls	MAP 5 (impact characterization); MAP 3.5 (human oversight); MANAGE 2 (minimizing negative impacts); MEASURE 3.3 (feedback and appeal)	Clause 8.4; Annex A.5.4, A.8.5, A.9.2	Human Review Rule, Likeness Check and Disclosure Rule. A named person approves every client-facing or public-facing asset before release. Outputs that show people are checked for resemblance to real people, and anyone who believes they are depicted can have content pulled within one business day. Audiences are told when realistic content is AI-generated.

Policy section	NIST AI RMF 1.0	ISO/IEC 42001:2023	How Momento Media applies it
7. Monitoring and incidents	MEASURE 2.4 (production monitoring); MEASURE 3 (risk tracking); MANAGE 4 (incident response and communication); GOVERN 4.3	Clauses 9.1, 10.2; Annex A.6.2.6, A.6.2.8, A.8.4, A.3.3	Performance and drift are tracked, logs let a deliverable be reconstructed, incidents have named owners, and staff can raise concerns in good faith without penalty.
7 and 11. Audit and review	GOVERN 1.5 (periodic review, frequency defined); MEASURE 1.3 (independent assessors); MEASURE 4 (measurement efficacy)	Clauses 9.2, 9.3, 10.1; Annex A.2.4	Systems are audited at least annually by someone independent of those who built or approved them. The Board reviews the policy quarterly, which serves as the management review, and each review feeds improvements.
8. Training and awareness	GOVERN 2.2 (training); GOVERN 4 (risk culture)	Clauses 7.2, 7.3, 7.4; Annex A.4.6	Staff are trained twice a year with role-based content and an onboarding module. Significant new AI systems need stakeholder review and a named adoption champion.
9. Procurement and vendors	GOVERN 6 (third-party and supply chain); MAP 4; MANAGE 3 (third-party risk management)	Annex A.10.2, A.10.3, A.4.4	No third-party AI tool is used on client work until it is on the Approved AI Tools Register: the CTO vets security and licence, the CFO approves spend, the CAIO signs off. Contracts secure audit rights and the right to disable a model.
10. Exceptions and notification	MANAGE 1.3 (response options); MANAGE 4.3 (incident communication)	Clause 10.2; Annex A.8.3, A.8.4	Exceptions are written, time-limited and logged. Incidents are logged within 24 hours and clients notified within 72 hours; regulator notification is assessed per incident.
Appendices A-E	GOVERN 1.6 (AI inventory); MAP 2.2 (documented limits); MEASURE 2.1 (documented test sets and metrics)	Clause 7.5; Annex A.4.2, A.6.2.7	The registration form (D) feeds the AI inventory; the model card (A), risk assessment (B) and audit checklist (C) hold the evidence; the exception form (E) records waivers.

Limits of this alignment: Momento Media aligns with ISO/IEC 42001 as a control reference and is not certified to it. NIST AI RMF is voluntary and used as a risk workflow, not a checklist. Not yet covered: NIST GOVERN 3.1 (a diverse team informing risk decisions), GOVERN 5 (engagement with stakeholders outside the company) and ISO clause 4.2 (interested parties' requirements) beyond Section 1.1. Reviewed at each quarterly policy review.