

AI Risk Assessment

AI Content Generation System

Documenting risks, controls, and recommended next steps for an AI system.

Prepared by	Wisdom Kwashie — CEO & Chief AI Officer
Organization	Momento Media, Inc. · Southern California
System assessed	AI Content Generation System (third-party, cloud-based)
Date	June 30, 2026 (revised October 3, 2026)
Aligned to	AI Governance Policy, revised October 1, 2026

Project / System Overview

Field	Detail
Name of AI System	AI Content Generation System (third-party, cloud-based)
Purpose / Function	A third-party, cloud-based AI content generation system used by Momento Media's creative teams to generate images, video, and other media assets for client marketing campaigns and social content, from text prompts and reference materials. The system also includes predefined workflows that guide users through specific content-creation projects. It is used for both internal ideation and client-facing, published deliverables.

Risk Summary

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), the method used in Appendix B of the AI Governance Policy. Likelihood, Impact and Level are inherent, before mitigation; the residual rating is after mitigation.

#	Risk	Type	Likelihood	Impact	Level	Residual
1	IP / Copyright & Likeness Infringement	Legal / IP	Medium	High	High (6)	Medium
2	Cost Drift & Flawed Predefined Workflows	Operational / Robustness	High	Medium	High (6)	Low
3	Data Leakage & Confidentiality	Privacy / Security	High	High	High (9)	Medium

Risk 1 — IP / Copyright & Likeness Infringement

Field	Detail
Risk Description	The system may generate outputs that infringe third-party intellectual property or use a real person's likeness without authorization, either because the underlying model was trained on copyrighted or scraped data, or because a user prompts it to imitate a protected work, artist style, or real person. If such content is published in a client campaign, it exposes Momento Media and its client to copyright and right-of-publicity claims and platform takedowns.
Risk Type	Legal / Intellectual Property & likeness compliance
Likelihood	Medium (the model's training data is opaque; both internal and client-facing use widens exposure).
Impact	High — legal liability, client reputational harm, content takedowns, and possible contractual/financial penalties.
Risk Level	High (6) — Medium × High, the same rating as the IP and copyright row in Appendix B of the policy.
Who could be harmed	People who resemble characters in generated content, who can be misrepresented without their consent; and artists whose style is imitated.
Mitigation Strategy	Primary (preventive) control: train and use a custom generation model built on licensed imagery of paid models who have granted likeness rights in perpetuity, so generated people derive from an owned, rights-cleared source set, securing right-of-publicity risk at the source. Supporting controls: vet the base/foundation model's licence and training-data provenance for residual copyright risk; restrict prompts that imitate identifiable artist styles, trademarks, logos, or protected characters; and apply the policy's Human Review Rule, under which a named person reviews and approves every client-facing or public-facing asset before release, with the Likeness Check for any output that shows people.
Residual Risk	Medium. Base-model training data can be assessed but not fully verified, so some provenance risk remains.
Owner and Status	AI Development & Validation Lead (training data and base model) with the Legal, Compliance, Risk & Audit Lead (licensing). Open.
Detection Method	Base-model licence and provenance review; Likeness Check results; incident reports.

Risk 2 — Cost Drift & Flawed Predefined Workflows

Field	Detail
Risk Description	The system's predefined workflows may be poorly designed with loosely-bound or misconfigured parameters causing the model to drift and produce off-target, low-quality outputs. Users then re-run generations repeatedly to obtain usable results, consuming excessive platform credits and time. Because the platform is consumption/credit-based, this creates unpredictable, inflated costs and erodes project profitability and delivery timelines.
Risk Type	Operational / financial (robustness — model drift; cost control)
Likelihood	High — third-party workflows are adopted for convenience, often without vetting, and credit burn from repeated re-generation is common.
Impact	Medium — mainly financial waste and delivery delays; can blow a project budget and, if it delays client delivery, cause reputational harm, but without the legal severity of Risk 1.
Risk Level	High (6) — High × Medium.
Who could be harmed	No individual is directly harmed. The cost is financial, carried by Momento Media and, through delayed or over-budget projects, its clients.
Mitigation Strategy	Audit and vet every predefined workflow in a sandbox before it is approved for client work (confirm parameters are properly bound and outputs are stable); maintain an approved-workflows register so only vetted workflows are used; set per-project and per-user credit caps with consumption alerts; track cost-per-deliverable to catch drift early; and provide users with prompt-guidelines and training to reduce wasteful re-generation. Spend is also governed by the policy's triggers: any tool over USD \$2,500 a year, or total AI spend over USD \$10,000 a year, goes to the Board, and the CFO reports AI spend quarterly.
Residual Risk	Low. Credit caps and sandbox-vetted workflows limit the worst case, and spend is reported to the Board each quarter.
Owner and Status	CFO (spend) with the AI Development & Validation Lead (workflows). Open.
Detection Method	Consumption alerts and cost-per-deliverable tracking.

Risk 3 — Data Leakage & Confidentiality

Field	Detail
Risk Description	Because the system is a third-party, cloud-based platform, staff may input confidential or client-restricted information, including but not limited to brand assets, unreleased campaign material, customer data, or proprietary strategy, into prompts or uploads. That data leaves Momento Media's control and may be retained by the vendor, used to train their models, or exposed in a breach. This risks breaching client confidentiality agreements and data-privacy laws (GDPR / CCPA) and can leak unreleased client campaigns.
Risk Type	Privacy / data protection & security (confidentiality)
Likelihood	High — convenience-driven use means staff may enter material into the platform without scrutiny, and both internal and client use widens exposure.
Impact	High — breach of client NDAs, regulatory exposure for personal data, leaked unreleased campaigns, and serious reputational and contractual damage.
Risk Level	High (9) — High × High. Higher than the custom model's "client data mishandling" rating in Appendix B of the policy (Medium), because here the data is entered into a third-party platform.
Who could be harmed	Clients and the people in their documents, whose confidential or personal information can be exposed; and the staff and department heads accountable for the operation, whose professional reputations suffer.
Mitigation Strategy	Enforce the policy's Data Entry Rule: Confidential, Client-Restricted and personal data is never entered into the platform unless it is contractually covered; use an enterprise/business tier with a signed data-processing agreement (no training on Momento Media data; defined retention limits); run vendor due diligence on data handling, security certifications, and audit rights; apply technical controls (SSO / access control, restrict sensitive-asset uploads, monitor usage); train staff on what may and may not be entered; and obtain client consent where their data or assets are processed on a third-party platform.
Residual Risk	Medium. The Data Entry Rule and the data-processing agreement reduce the exposure, but a person can still make a mistake at the point of entry.
Owner and Status	Legal, Compliance, Risk & Audit Lead (rule and vendor terms) with the CTO (access controls). Open.
Detection Method	Access-control and usage monitoring; audit of data-entry compliance.

Next Steps & Recommendations

Prioritised by risk level and residual risk, three actions are recommended:

#	Action	Owner	Target
1	Move to an enterprise tier with a signed data-processing agreement, enforce the Data Entry Rule on what may be entered, and complete vendor due diligence on retention, security, and audit rights before expanded use.	Legal, Compliance, Risk & Audit Lead with the CTO	Within 30 days of approval
2	Build the custom generation model on rights-cleared imagery with perpetual likeness releases, and vet the base model's licensing to close the residual copyright gap.	AI Development & Validation Lead	Within 90 days of approval
3	Sandbox-vet every predefined workflow before approval, maintain an approved-workflows register, and set per-project credit caps with cost-per-deliverable monitoring.	AI Development & Validation Lead with the CFO	Within 60 days of approval

Policy and framework link: Policy Section 5 and Appendix B (scoring); Section 6 (Data Entry Rule, Human Review Rule, Likeness Check); Section 9 (spend triggers). Using the crosswalk in Appendix F of the policy: NIST AI RMF MAP 4, MAP 5, MEASURE 2.10, MANAGE 1 and MANAGE 3; ISO/IEC 42001 clauses 6.1.2, 6.1.3, 8.2 and 8.3, and Annex A.7.2–A.7.5, A.10.2 and A.10.3.