MVPR AI TRANSPARENCY & GOVERNANCE ADDENDUM

Version: 1.1 · **Effective:** November 2026 · **Supersedes:** Version 1.0 (14 January 2026) · Published at https://mvpr.io/ai-governance

PARTIES AND STATUS

This AI Transparency & Governance Addendum ("Addendum") forms part of the MVPR Master Services Agreement ("MSA") between MV Public Relations Limited ("MVPR", "We", "Us", "Our") and each Customer ("Customer", "You", "Your"), and supplements Section 7 of the MVPR Service Level Agreement ("SLA"), Sections 8 and 9.4 of the MSA and the Data Processing Agreement ("DPA"). Capitalised terms have the meanings given in the MSA. In the event of conflict, the order of precedence in Section 2.2 of the MSA applies.

1. PURPOSE AND SCOPE

1.1 Purpose. This Addendum describes how MVPR uses artificial intelligence in the provision of the Services: the models and providers used, the human oversight applied, MVPR's quality, fairness and transparency commitments, and the limits of what AI can be relied upon to do.

1.2 Related Documents. This Addendum should be read with Section 7 of the SLA, Sections 6.2, 8 and 9.4 of the MSA, the DPA (in particular Sections 2.4 to 2.6, 5.7, 6.4 and 10.8), and MVPR's Operations Security and Secure Development Policies available at the Trust Centre.

1.3 Updates. MVPR reviews this Addendum at least annually and whenever there is a material change to the AI models, providers or oversight processes described in it. Updated versions are published at https://mvpr.io/ai-governance and notified to Customer's administrative contact in accordance with Section 14.

2. AI SYSTEMS AND MODELS

2.1 Models and Providers. MVPR uses large language models provided by Anthropic Ireland Limited and OpenAI Ireland Ltd (each an EEA entity), accessed through their EEA endpoints under enterprise terms, together with MVPR's own natural-language-processing, semantic search and matching tools. Typical uses are: long-form content drafting and editing; research synthesis; interview transcription processing; talking-point generation; content outlining and headline generation; journalist-query response drafting; structured data extraction; sentiment analysis of coverage; entity extraction and topic modelling for journalist targeting; and quality assessment against editorial standards. The specific model versions in use at any time are listed at https://mvpr.io/legal/ai-models and may change in accordance with Section 2.3.

2.2 Selection Criteria. MVPR selects models on task suitability, demonstrated quality and accuracy, the provider's security and data-handling practices, regulatory compliance (including the UK and EU data protection regimes and the EU AI Act), EEA data residency and cost-effectiveness.

2.3 Model Changes. MVPR may change the models or providers it uses. MVPR will give Customer at least thirty (30) days' notice of a change of provider or a migration to a fundamentally different model architecture, and will notify routine version updates through the list at https://mvpr.io/legal/ai-models. Security patches are applied immediately and notified afterwards. If a notified change results in a material and demonstrable reduction in the quality of the Services that MVPR does not remedy within thirty (30) days of Customer's written notice describing the reduction, Customer may terminate the affected Order on written notice and MVPR shall refund any prepaid Fees for the period after termination. This is Customer's sole remedy in respect of a model change.

2.4 Training and Data Use. Customer Data and Output are not used to train or fine-tune AI models, by MVPR or by its AI providers, without Customer's explicit written consent. MVPR's agreements with Anthropic and OpenAI provide that Customer inputs and outputs are not used for model training, are not retained beyond processing (subject to the providers' published abuse-monitoring exceptions), are isolated from other customers' data, and remain the customer's property. Customer owns the Output as set out in Section 6.2 of the MSA.

3. WORKFLOW AND HUMAN OVERSIGHT

3.1 Content Creation. AI-assisted content is produced through a workflow with human checkpoints at every stage: (1) strategic planning, led by humans, with AI not used for strategic decisions; (2) AI-assisted research and transcription, validated by MVPR staff; (3) AI-assisted outlining, reviewed and approved by MVPR staff before drafting; (4) AI-assisted drafting from the approved outline, followed by full editorial review, fact-checking and tone refinement by MVPR staff; (5) AI-assisted quality assessment (grammar, readability, brand-voice consistency), evaluated by MVPR staff; (6) Customer review and explicit approval; and (7) publication or distribution, decided by humans.

3.2 Journalist Opportunities. The Platform monitors journalist request feeds and news sources, filters opportunities by relevance, and drafts suggested responses. MVPR staff or Customer review each AI-flagged opportunity and its reasoning (why it is relevant, which spokesperson fits and why, which resources support the angle, and a quality score), refine the response, and Customer approves before anything is sent.

3.3 Automated Features. Journalist database semantic search, feed monitoring and keyword filtering, sentiment analysis of coverage, and performance analytics operate automatically. They are subject to regular quality audits by MVPR staff, Customer feedback mechanisms, escalation to human review, and transparent confidence scores where applicable.

3.4 Where AI Does Not Decide. AI does not make strategic recommendations, publication decisions, legal or compliance judgements, or decisions in crisis communications; these are made by people. AI may assist MVPR staff in drafting internal and external communications, including commercial documents, but all such communications are reviewed and sent by people, and all business terms are decided by people. No decision producing legal or similarly significant effects on an individual is made solely by automated means.

4. QUALITY AND ACCURACY

4.1 Quality Framework. Four layers apply to AI-assisted deliverables: AI self-assessment (confidence scores, flagging of uncertain content, source citation where possible); automated checks (grammar, plagiarism detection, brand-voice scoring, readability); human editorial review (fact-checking against authoritative sources, verification of citations, tone and style alignment, cultural sensitivity); and Customer review and approval.

4.2 Accuracy Commitments. MVPR does not knowingly deliver fabricated facts; factual claims in deliverables are verified against authoritative sources before delivery; identified errors are corrected promptly; and MVPR discloses where AI limitations may affect accuracy. Customer acknowledges that, despite these measures, AI-assisted content may contain errors and must be reviewed by Customer before publication (MSA Sections 4.2 and 9.4).

4.3 Known Limitations. AI models may generate plausible but incorrect information; lack awareness of recent events; misinterpret nuanced context; reflect biases present in training data; produce inconsistent outputs for similar inputs; and produce formulaic or derivative text without human refinement. MVPR mitigates these through the human oversight in Section 3,

supplementary research for recent events, contextual briefings, and the bias measures in Section 5.

4.4 Continuous Improvement. MVPR tracks Customer satisfaction with AI-assisted deliverables, revision rates, factual error rates after human review and AI confidence-score accuracy, reviews these metrics regularly, and refines prompts and workflows accordingly.

5. BIAS AND FAIRNESS

5.1 Sources of Bias. AI models may exhibit bias arising from training data, model design, deployment context or feedback loops.

5.2 Mitigation. MVPR works with providers that publish their approaches to fairness; trains editorial staff to identify and correct biased outputs; applies cultural-sensitivity review for content aimed at diverse audiences; engineers prompts to elicit balanced perspectives and avoid stereotypes; and provides a reporting channel for Customers to raise perceived bias.

5.3 Journalist Targeting. MVPR's journalist database and matching tools prioritise relevance based on beat coverage and past work. They do not filter journalists by protected characteristics (except where Customer targets publications that serve a particular community), do not systematically exclude outlets, and do not favour one Customer over another in opportunity allocation.

5.4 Reporting. If Customer identifies potential bias in an AI-assisted output, it should contact its Account Director or tom@mvpr.io. MVPR will investigate within three (3) Business Days, correct the output, address any systemic cause and report back. MVPR reviews AI-assisted outputs for fairness at least annually and will share a summary with Customer on request.

6. TRANSPARENCY AND EXPLAINABILITY

6.1 What MVPR Discloses. On request, MVPR will explain: which AI models were used for a specific deliverable and their known capabilities and limitations; the AI-generated reasoning behind journalist-opportunity matching or content suggestions, with confidence scores where available; which sections of a deliverable were AI-generated and which human-written, and the sources used; and the human oversight and quality controls applied.

6.2 How to Request. Requests may be made to the Account Director or to tom@mvpr.io, identifying the deliverable concerned. MVPR will respond within five (5) Business Days.

6.3 Limits. Large language models are not fully interpretable, even by their developers. Explanations are approximations of model reasoning, not exact reproductions of it. MVPR provides the most detailed explanation reasonably available.

7. DATA PRIVACY IN AI PROCESSING

7.1 Data Flow and Residency. Customer input is stored in Google Cloud Platform (EU multi-region), pre-processed and minimised, sent to the AI provider's EEA endpoint, processed, returned to MVPR systems for human review, and delivered to Customer. There are no systematic transfers of Customer Personal Data outside the UK and EEA for AI Processing. The DPA governs this processing.

7.2 Data Minimisation. Only data necessary for the specific task is sent to a model; personal identifiers are removed where not needed; irrelevant data is stripped from prompts; and providers do not retain prompts beyond processing (subject to their abuse-monitoring exceptions). For example, content outlining receives the topic, key messages and audience

profile, not the Customer's full database; quality assessment receives the draft only, not client names or account metadata.

7.3 Sensitive Data. MVPR does not process special categories of Personal Data (GDPR Article 9) through AI models unless Customer explicitly requires and authorises it in writing, in which case additional safeguards are agreed under the DPA.

7.4 Provider Data Handling. See Section 2.4. MVPR reviews its AI providers' terms and practices regularly and has escalation routes for data-handling concerns.

8. INTELLECTUAL PROPERTY

8.1 Ownership of Output. Consistent with Section 6.2 of the MSA, Customer owns the Output, including AI-generated drafts, outlines and suggestions, their human-refined versions and the final deliverables. MVPR's rights in Output are limited to using it to provide the Services to Customer and retaining aggregated, anonymised data for service improvement as permitted by MSA Section 6.2(c).

8.2 Third-Party Rights. MVPR runs AI-assisted content through plagiarism detection, applies editorial review to identify derivative work, includes attribution where appropriate, and engineers prompts for original expression. MVPR's AI providers offer intellectual property indemnities to their enterprise customers for model outputs, subject to their terms; MVPR relies on those indemnities and does not itself warrant that Output is free of third-party rights (MSA Section 9.4). Customer should carry out its own IP review before using Output in high-stakes contexts.

8.3 Similarity Across Customers. Similar prompts may produce similar outputs for different customers, and common phrasings and structures may recur. MVPR reduces this through detailed briefings, Customer-specific resources and human editorial refinement, and Customer may modify Output freely.

9. REGULATORY ALIGNMENT

9.1 EU AI Act. MVPR's use of AI falls outside the high-risk categories of Regulation (EU) 2024/1689: it does not involve critical infrastructure, education, employment decisions, law enforcement, migration or biometric identification. MVPR uses general-purpose AI models supplied by third parties and complies with the transparency obligations applicable to deployers of such systems, including by disclosing AI involvement in content on request and where required by law. MVPR monitors the Act's implementation and updates this Addendum as obligations take effect.

9.2 Data Protection. MVPR's AI implementation applies the GDPR principles of lawfulness, fairness, transparency, purpose limitation, data minimisation, accuracy, storage limitation, integrity and confidentiality; privacy by design under Article 25; processor obligations under Article 28 through the DPA; security under Article 32 through its ISO 27001:2022 ISMS; breach notification under Articles 33 and 34; and assistance with impact assessments under Article 35. No solely automated decisions with legal or similarly significant effects are made (Article 22).

9.3 Other Frameworks. MVPR's ISO 27001:2022 certification covers its AI systems, and AI is an explicit risk category in its Risk Management Policy, assessed using ISO 27005 methodology and aligned with the NIST AI Risk Management Framework. MVPR does not process protected health information or other regulated sector data through AI models unless specifically agreed in an Order with appropriate additional terms.

10. CUSTOMER RESPONSIBILITIES AND DISCLAIMERS

10.1 Customer Responsibilities. Customer shall review and approve all AI-assisted deliverables before publication; provide briefings, brand guidelines and context; give timely feedback; ensure its use of Output complies with applicable law, including any disclosure obligations for AI-assisted content; carry out its own IP review for high-stakes uses; and accept final responsibility for content it approves and publishes.

10.2 Prohibited Uses. Customer shall not use the Services to generate unlawful, harmful or misleading content; attempt to extract or reverse-engineer MVPR's AI implementation, prompts or workflows; or use Output to train competing AI models or services (MSA Section 4.3).

10.3 Disclaimers. MVPR does not warrant that Output will be error-free, accurate, complete, original or fit for a particular purpose, that it will meet Customer's subjective standards without human review, that AI will perform identically across all contexts, or that models will not change over time. Customer's remedy for dissatisfaction with an AI-assisted deliverable is revision through human editorial refinement in accordance with the MSA and SLA.

11. AI INCIDENTS

11.1 Categories. Reportable AI incidents include significant factual errors that passed human review, outputs reflecting unacceptable bias, unauthorised disclosure of Personal Data through AI Processing, potential third-party IP infringement in Output, AI service outages or significant degradation, and unauthorised access to AI systems.

11.2 Reporting and Response. Customer should report an AI incident to its Account Director or tom@mvpr.io, describing the issue, the affected deliverables and the potential impact. MVPR's response targets are:

Severity	Initial response	Investigation	Resolution plan
Critical (e.g. privacy breach, published inaccuracy causing harm)	2 Business Hours	1 Business Day	2 Business Days
High (e.g. significant bias, pre-publication error)	1 Business Day	3 Business Days	5 Business Days
Medium (e.g. quality concern, minor inaccuracy)	2 Business Days	5 Business Days	10 Business Days
Low (e.g. stylistic preference)	3 Business Days	10 Business Days	As appropriate

Personal Data Breaches are also handled under Section 7 of the DPA, whose timelines prevail. MVPR will investigate root cause, correct the issue, implement preventive measures, document the outcome and report back to Customer.

12. GOVERNANCE

12.1 Structure. The ISMS Governance Council provides strategic oversight of AI risk. The CEO (Tom Lawrence) holds ultimate accountability for AI governance; the CTO (Konrad Fuger) is responsible for technical implementation, security and AI provider management; the editorial team is responsible for content quality and human oversight; the operations team monitors Platform performance and incidents; Account Directors are responsible for deliverable quality and Customer feedback. AI metrics are reported to the Governance Council quarterly and AI risks are assessed at least annually.

12.2 Escalation. Customers may escalate AI-related concerns first to their Account Director and then to the CEO at tom@mvpr.io.

12.3 Ethics Principles. MVPR's AI use is human-centric, transparent, fair, privacy-protective, accountable, safe and sustainable, as described in this Addendum.

13. SUPPORT AND FEEDBACK

Customers may ask their Account Director or tom@mvpr.io for explanations of AI's role in deliverables, guidance on briefing AI-assisted work, troubleshooting, updates on model changes and training on the Platform's transparency features. MVPR welcomes feedback on AI output quality, user experience, new capabilities and ethics, logs all feedback, reviews trends regularly and communicates resulting changes to affected Customers.

14. AMENDMENT AND VERSION CONTROL

14.1 Updates. MVPR reviews this Addendum at least annually and on any material change (addition or removal of an AI provider, significant change to data handling or human oversight, or new regulatory requirement). Material changes are notified to Customer's administrative contact at least thirty (30) days before they take effect, with a summary and effective date, and are published at https://mvpr.io/ai-governance in accordance with Section 15.7 of the MSA.

14.2 Version History.

Version	Date	Summary of changes
1.0	14 January 2026	Initial publication
1.1	November 2026	Status clarified as part of the MSA; model list moved to a maintained web page; model-change remedy narrowed to material, unremedied degradation with refund; Section 3.4 reworded to reflect AI-assisted drafting with human decision; response targets expressed in Business Hours and Days; independent-review tier and healthcare addendum references removed; EU AI Act section updated for obligations in force; contacts updated; alignment with MSA November 2026, DPA v3.1 and SLA v1.1

15. CONTACT

AI governance enquiries: tom@mvpr.io (subject line "AI Governance Inquiry") or your Account Director.

MV Public Relations Limited · Company number 12688946 · 121 College Road, London NW10 5EY, United Kingdom