LEGAL AND REGULATORY FRAMEWORKS FOR ARTIFICIAL INTELLIGENCE

Option 3: Data Protection, Privacy and Intellectual Property in the Age of Generative AI

A Project Report Submitted in Partial Fulfilment of the Requirements for the Course on Legal and Regulatory Frameworks for AI

Submitted by: Hubert Nahimana

Student ID: HMI26011

EIIET, 14 November 2026

Lecturer: Prof. Dr. Srinath Doss

TABLE OF CONTENTS

Abstract
1. Introduction…….4
1.1 Background of the Study……4
1.2 Problem Statement……..4
1.3 Research Objectives and Questions……4
1.4 Significance of the Study……5
2. Literature Review……..…5
2.1 Generative AI and Data Governance……..5
2.2 GDPR and Personal Data Processing…….5
2.3 Copyright, Fair Use and Fair Dealing……6
2.4 Data Subject Rights and Transparency……...6
2.5 Business Opportunities and Risks……..7
3. Methodology……7
3.1 Research Design and Sources……...7
3.2 Analytical Method and Limitations……..7
4. Findings……...…7
4.1 GDPR and AI Training……...7
4.2 Copyright and AI Training Data…….8
4.3 Data Subject Rights and Outputs……...9
4.4 Business Integration Risks……10
4.5 International Regulatory Differences……10
5. Discussion……..10
5.1 Business Strategy……..10
5.2 Compliance Framework……11
5.3 Measured Policy Response…….....11
5.4 Implementation Roadmap……..12
6. Conclusion……...…12
7. References……...…13
8. Appendices…….14

ABSTRACT

Generative artificial intelligence (AI) provides organizations with the ability to automate knowledge work, generate content, enhance customer service and speed up research processes. Concurrently, large language models (LLMs) pose challenging questions regarding the legal collection and use of personal data, copyright issues related to training material, ownership of outputs they generate, and the rights of individuals whose information is included in training corpora. This report assesses how the European Union's General Data Protection Regulation (GDPR), copyright laws, and AI-specific regulations apply to the creation and commercial use of generative AI. It employs a qualitative doctrinal and comparative approach, drawing on legislation, regulatory guidance, court rulings, and academic literature. The analysis concludes that while current legal frameworks are still relevant, they are difficult to apply to large-scale, multi-source datasets and models with internal representations that are hard to examine. Adding AI to customer-facing or internal applications introduces extra risks, such as the exposure of personal or confidential information, infringing outputs, reliance on suppliers and damage to reputation. To comply with GDPR, you need a valid legal basis, ensure processing is fair and transparent, minimize data collection, secure data, and have procedures in place for rights requests. Copyright issues are influenced by jurisdiction, the source and use of works, applicable exceptions, and the specifics of each situation. Additionally, AI-generated outputs might not be protected by copyright if there is not enough human creative input. The report advises lifecycle governance, documented data provenance, privacy and copyright evaluations, careful supplier contracting, human oversight, and ongoing monitoring. The conclusion is that responsible governance is a strategic capability that can foster innovation while maintaining long-term competitiveness.

Keywords: *generative AI, GDPR, privacy, copyright, intellectual property, large language models, business compliance.*

1. INTRODUCTION

1.1 Background of the Study

Organizations implement these systems for various purposes, including customer-service chatbots, marketing, software development, document review, research, and internal knowledge management. When implemented correctly, these tools can reduce routine administrative costs, shorten production cycles, and allow employees to concentrate on higher-value tasks. These abilities rely on data. Training material can consist of licensed datasets, publicly available websites, user contributions, books, articles, images, code, personal data, and confidential information. Just because information is online does not mean it can be processed as personal data or copied under copyright law. In the same vein, just because a model can produce content that seems plausible, this does not guarantee that what it generates is accurate, original, or safe to share. The GDPR establishes a central framework for handling personal data within the European Union. It sets out requirements regarding lawfulness, fairness, transparency, purpose limitation, data minimization, security, and individual rights. Obligations related to copyright-compliance policies and public summaries of training content are outlined in Article 53, which falls under the Regulation's scope and applicable provisions (Regulation (EU) 2024/1689, art. 53, as amended). Copyright law establishes the rights associated with creative works, the conditions under which protected material can be used, and whether generated content qualifies for protection. The EU Artificial Intelligence Act introduces obligations tailored to AI, which includes requirements for providers of general-purpose AI models.

1.2 Problem Statement

Conventional legal structures were not built with the scale, opacity, and probabilistic behavior of modern generative AI in mind. Training could involve millions of items gathered from various sources, reused for different purposes, and managed by multiple organizations. This makes it harder to identify a lawful basis, offer meaningful transparency, allow data subjects to exercise their rights, and evaluate copyright permissions or exceptions.

Businesses adopting third-party AI tools face a related challenge: the model provider may control training and technical architecture, while the customer remains responsible for the way it uses outputs and processes data in its own operations. Poor governance can result in regulatory scrutiny, legal action, contract disputes, operational disruptions, and damage to reputation. On the other hand, completely rejecting AI could mean missing out on valid productivity gains. The task is to create a balanced framework that safeguards rights and allows for responsible innovation.

1.3 Research Objectives and Questions

The report has the following objectives: (1) assess the applicability of the GDPR to personal data involved in AI training, fine-tuning, and deployment; (2) investigate copyright ownership, fair use, fair dealing, and the use of protected works in training; (3) scrutinize

data-subject rights and the risks of personal information being included in outputs; (4) pinpoint compliance risks in customer-facing and knowledge-worker applications; and (5) propose a practical governance framework.

The research poses the question: What is the applicability of privacy laws to generative AI? How well can copyright law adapt to AI training and the outputs it produces? What risks come into play when companies adopt these tools? What policies can harmonize innovation, compliance, and competitiveness?

1.4 Significance of the Study

This issue is important because organizations usually embrace AI services faster than they create governance for them. By conducting a practical legal assessment, managers, compliance officers, developers, and procurement teams can enhance their decision-making regarding data sources, suppliers, use cases, and safeguards.

2. LITERATURE REVIEW

2.1 Generative AI and Data Governance

By viewing compliance as a capability instead of a last-minute hurdle, we can enhance trust and diminish uncertainty in commercial relationships.

Generative models understand statistical patterns in the training data they are exposed to and can generate content based on prompts later on. Training data can be licensed, scraped, provided by users, or sourced from internal repositories. Different sources will therefore raise different issues of privacy, retention, permission, and other aspects of provenance. Governance ought therefore to be extended from training to data collection, preprocessing, testing, deployment, monitoring and retirement.

Raji et al. (2020) propose end-to-end internal algorithmic auditing where accountability is not a final test but a feature of system development. This is especially the case where the model is developed, hosted, and deployed by different organizations for use in their business. Documentation is helpful in allocating responsibility but cannot by itself guarantee that all input datasets were obtained legally or that all outputs produced are safe.

2.2 GDPR and Personal Data Processing

GDPR provides for a very broad definition of personal data as information related to an identified or identifiable natural person (Regulation (EU) 2016/679, art. 4). Depending on the situation, this could include names, contact details, employment records, online identifiers and personal data in public documents. Article 5 contains the principles of lawfulness, fairness and transparency; purpose limitation; data minimization; accuracy; storage limitation; and integrity and confidentiality. Article 6 provides for the need for an appropriate legal basis.

According to the European Data Protection Board's Opinion 28/2024, we must analyze each situation individually to determine how data-protection law applies to AI models. It warns that you should not assume a model is anonymous just because personal data has been converted into parameters, and it clarifies that legitimate interest is not automatically a valid basis for all AI activities (EDPB, 2024). This guidance emphasizes the importance of evaluating extraction risk, reasonable expectations, safeguards, and the purpose of processing.

2.3 Copyright, Fair Use and Fair Dealing

Copyright safeguards qualifying original expression, according to national law and applicable exceptions. Determination of whether those activities violate copyright hinges on the relevant acts, jurisdiction, exceptions, licenses, and factual circumstances. Training AI might include actions like downloading, storing, transforming, or processing protected works in other ways.

In the United States, fair use as outlined in 17 USC § 107 takes into account the purpose and character of the use, the nature of the work, the amount and substantiality used, and the effect on the potential market. In the European Union, Directive (EU) 2019/790 includes separate provisions for text and data mining. The question is not resolved automatically by either commercial purpose or a claim of transformation. Article 3 outlines the criteria for qualifying research organizations and cultural heritage institutions, whereas Article 4 sets forth a wider exception that is subject to lawful access and rights reservations. These regimes differ from one another.

Another question is if generated content can be copyrighted. According to the US Copyright Office (2025a), for something to be copyrightable, it must be created by a human and involve a creative contribution. In Thaler v. Perlmutter (2025), the US Court of Appeals for the District of Columbia Circuit ruled that a machine cannot be the sole author under US copyright law. However, this does not imply that every task assisted by AI is unprotectable; expressions created by humans may still qualify.

2.4 Data Subject Rights and Transparency

The GDPR grants rights to access, rectification, erasure, restriction, objection, and portability in specific situations. Transparency obligations are set forth in Articles 13 and 14, whereas Article 22 addresses specific decisions based solely on automated processing that have legal consequences or other significant effects. The right to erasure is subject to limitations: Article 17 outlines exceptions, and the GDPR does not mandate that a model be fully retrained every time someone makes a request. It is hard to put these provisions into practice when information is spread across source datasets, model parameters, vector stores, conversation logs, and downstream applications.

However, organizations should evaluate requests instead of dismissing them as technically inconvenient. Wachter et al. (2017) explore the boundaries of a general right to explanation within the GDPR, highlighting the larger challenge of creating meaningful accountability

for complex automated processes.

2.5 Business Opportunities and Risks

Generative AI can speed up tasks like drafting, summarizing, software development, and customer support. However, systems lacking transparency can generate biased or incorrect outputs, reveal confidential details, or reproduce protected expression. The work of Raji et al. (2020) illustrates the importance of internal auditing and organizational accountability. There is a lasting disconnect between technical capability and organizational readiness: companies might acquire AI services without thoroughly examining data retention, model-improvement terms, security protocols, or intellectual-property assurances. Effective governance must be tailored to each use case, data category and possible consequence.

3. METHODOLOGY

3.1 Research Design and Sources

This report employs qualitative doctrinal and comparative research methods. Legislation, regulatory guidance, and judicial decisions are scrutinized by doctrinal analysis to interpret current legal obligations. References encompass the GDPR, the EU Copyright in the Digital Single Market Directive, the EU AI Act along with its 2026 amending regulation, US copyright laws, EDPB guidance, reports from the US Copyright Office, court rulings, and peer-reviewed academic studies. A comparative analysis underscores the differences between the European Union and the United States.

3.2 Analytical Method and Limitations

The sources were categorized into four themes: personal-data processing, copyright and ownership, business compliance risks, and strategic governance. The report separates established legal obligations from contested interpretations and policy proposals. It does not rely on interviews, surveys, or empirical tests; its conclusions are analytical rather than statistical. Legal outcomes depend on the jurisdiction and the specific facts of the case, and the law is always evolving. The report serves as a general governance analysis and is not legal advice tailored to any specific jurisdiction.

4. FINDINGS

4.1 GDPR and AI Training

4.1.1 Legal foundation and sensitive data

In areas governed by the GDPR, you need a lawful basis to process personal data for training or fine-tuning models. The suitable basis varies with the specific activity. A business that wants to rely on legitimate interests under Article 6(1)(f) must identify its interest, demonstrate that the processing is necessary, and weigh it against the rights and freedoms of those affected. Consent, contractual necessity, legal obligation, and legitimate interests are just a few of the factors included in Article 6. The analysis should take into

account reasonable expectations, the characteristics and volume of data, potential harms, safeguards and less intrusive alternatives.

4.1.2 Purpose limitation, minimisation and transparency

According to Articles 5 and 25, organizations must establish their purposes, limit processing to what is necessary, and adopt privacy by design principles. Just because there is a vague intention to enhance model performance does not mean there is automatic permission for unlimited collection and reuse. Article 14 outlines exceptions, including situations where providing information is impossible or would require a disproportionate effort, as long as there are safeguards and it meets the precise statutory conditions. Developers must document intended uses, evaluate compatibility when data is reused, eliminate irrelevant information where possible and manage retention.

Articles 13 and 14 mandate that information must be provided to individuals in relevant situations. The act of collecting on a large scale does not automatically eliminate the need for transparency obligations. Organizations need to evaluate if an exception applies and record their reasoning. According to Article 35, a data protection impact assessment (DPIA) may be necessary if the processing is likely to pose a high risk to individuals. A DPIA is particularly important for extensive profiling, sensitive data, or AI systems that aid decisions with major consequences. To be effective, it should be finished early enough to shape system design instead of just documenting a decision that has already been made.

4.1.3 Anonymity, memorization and rights requests

According to the EDPB's Opinion 28/2024, the determination of whether an AI model is anonymous should be made on a case-by-case basis. Eliminating names can reduce the risk of identifiability, but it does not guarantee complete protection if rare events or contextual clues can point to someone. Models can also memorize unique information that prompts can trigger. Claims of anonymization should be tested by organizations, which should also consider reasonably likely methods of extraction.

Access, rectification and erasure requests may require the organization to locate source records, update downstream repositories or apply other effective measures. A request does not always require a full model retraining in every case; the answer hinges on if personal data remains, the legal obligations that apply, and the remedies available. Companies should chart data flows, maintain source records when lawful, establish clear request routes and document the assessment.

4.2 Copyright and AI Training Data

4.2.1 Training legality varies by jurisdiction

The central question is around whether copying or processing protected works during training requires permission or falls under an exception. US law requires that the fair-use factors be evaluated together and in context. Considerations could encompass the training's objective, the creative aspect of the works, the quantity utilized, and whether the model or its outputs influence the markets for the original works. A commercial purpose does not

automatically defeat fair use, and describing a process as transformative does not automatically establish it. In the EU, the text-and-data-mining provisions in Directive (EU) 2019/790 are distinct. Article 4 permits specific text and data mining activities if there is lawful access and rights have not been properly reserved. Companies should not assume that all online content can be used for training without restrictions. They need to log source, licence status, access conditions and any rights reservation.

4.2.2 Ownership and output infringement

Generative AI training is the focus of Part 3 of the US Copyright Office's report, while Part 2 deals with the copyrightability of outputs (US Copyright Office, 2025a, 2025b). The Office's reports serve as influential policy analyses rather than court judgments that resolve every dispute. A defensible business process should distinguish official guidance, judicial holdings, allegations and unresolved litigation.

Even if a company has permission to use an AI service, it may not own exclusive copyright to all outputs. According to US law, the Copyright Office examines if a human contributed enough original expression. The DC Circuit ruled in Thaler v. Perlmutter (2025) that a machine cannot be the sole author. It is important not to overgeneralize the ruling to suggest that all AI-assisted works are unprotectable; human actions such as selection, arrangement, editing, or other creative contributions may be significant under applicable law.

Outputs may raise infringement issues if they reproduce a substantial amount of protected expression or closely resemble works that already exist. The legal assessment relies on the jurisdiction, the protected expression, similarity, and the available defenses. The ongoing litigation between The New York Times and OpenAI, as well as Microsoft, highlights disputes regarding training and alleged reproduction; however, the publishers' allegations should not be considered final judicial findings unless a specific ruling backs that claim.

When using AI for business purposes, controls should encompass reviewing provider terms, preserving evidence of a meaningful human contribution, licensing valuable third-party content when appropriate, checking outputs in high-value creative contexts, and creating procedures for addressing rights-holder complaints. Just because you have a contract that allows you to use a model, it does not mean that the output is original, non-infringing, or that you exclusively own it.

4.3 Data Subject Rights and Outputs

Even after training, privacy risks persist. Personal information can be included in prompts, retrieval databases, chat logs, and generated summaries. If database permissions are not set up correctly, a customer-service bot could expose one customer's information to another. An internal assistant could summarize employee records or confidential documents for a user without access rights. These events could activate GDPR responsibilities related to lawful processing, security, and breach response.

Depending on the legal thresholds and circumstances, Articles 33 and 34 might require you to notify a supervisory authority and communicate with affected individuals after a

personal-data breach.

4.4 Business Integration Risks

Customer-facing applications can improve service availability, but may collect excessive data, provide misleading statements or disclose information. Businesses must give appropriate notices, limit the categories of information collected, test for disclosure attempts, and escalate consequential complaints to trained personnel. Just because a chatbot is called an "assistant," it should not be assumed to be harmless.

Applications designed for knowledge workers can handle contracts, research, source code, and strategic plans. If you upload a confidential acquisition plan to an unapproved public AI service, you might expose trade secrets or violate contractual obligations, depending on the provider's terms and how they actually handle inputs. Organizations should classify information, centrally approve tools, and ban restricted data from being submitted to services that have not been assessed.

Third-party procurement creates further risk. Organizations need to establish their role under data protection law for every activity. When a processor acts on behalf of a controller, Article 28 generally mandates a compliant processing contract. Contracts should address training on customer inputs, retention, deletion, international transfers, security, audit rights, incident reporting, intellectual-property claims and exit arrangements. You should not rely on a supplier's broad marketing promises instead of conducting due diligence.

4.5 International Regulatory Differences

The EU combines comprehensive data-protection law, copyright legislation and AI-specific rules. The United States combines federal copyright law with sectoral and state privacy regimes; training disputes often depend on fact-specific copyright analysis. Other jurisdictions have their own consent rules, privacy principles, exceptions, and enforcement strategies. A multinational company cannot assume that a dataset lawful to collect in one jurisdiction can be reused elsewhere without further analysis.

The EU AI Act was revised in 2026 through Regulation (EU) 2026/1744. Therefore, a report drafted in October 2026 should refer to the consolidated text and the amending act instead of depending solely on the original 2024 text. Since the obligations of the AI Act have different start dates and scope conditions, companies need to verify which requirements apply to their role and use case. Although regulatory fragmentation increases compliance costs, companies with robust cross-border governance might find an opportunity to serve risk-sensitive customers.

5. DISCUSSION

5.1 Implications for Business Strategy

The findings suggest that governance is a strategic capability. The business question goes beyond just being able to deploy AI; it is about whether it can be used lawfully, safely, and

sustainably. A company that has documented controls can enhance supplier negotiations, reassure enterprise customers, and respond to regulators more effectively. A company that embraces AI without looking into data provenance might enjoy short-term productivity boosts, but it risks accumulating legal liabilities over time.

First, data governance should influence model selection. In situations where privacy or copyright risks are significant, organizations might choose providers that deliver clear documentation, enterprise controls, contractual protections, and limitations on the reuse of customer inputs. Secondly, safeguards need to be proportional to the use case. Brainstorming with non-confidential material is not the same as screening job applicants or generating health-related guidance. Third, companies should weigh the options of building a model, fine-tuning an existing one, and using a third-party service.

5.2 Compliance Framework

A practical framework should integrate legal, technical, and organizational measures. Data governance should keep an inventory of datasets and processing activities, document their provenance and legal bases, and exclude unnecessary personal or copyrighted material when appropriate. Access restrictions, encryption, retention limits, filtering, and testing for extraction or disclosure should all be part of privacy engineering. There should be human oversight when outputs impact people. Rather than assuming, claims about anonymization should be validated.

The management of copyright should involve documenting licenses and permissions, evaluating applicable exceptions and rights reservations, and establishing procedures for reviewing outputs that may infringe copyright. Human oversight is especially important where outputs impact people substantially or where a factual inaccuracy could lead to material harm. Responsibilities for security, processing, intellectual-property claims, incident response, and regulatory cooperation should be assigned by procurement controls. Lastly, there should be ongoing oversight to assess changes in models, data, prompts, contracts, and legislation.

5.3 A Measured Policy Response

At the organisational level, businesses should adopt a written AI policy before deployment. It should specify approved uses, prohibited activities, accountable owners, data-handling rules and escalation procedures. A risk-tiering process can streamline approval of low-risk tools while requiring more extensive legal, privacy and security review for sensitive data or consequential decisions.

Second, organisations should adopt data-provenance standards and favour suppliers able to explain training and input-data practices to a reasonable degree. Third, contracts need to cover permitted use, training on customer inputs, data retention and deletion, security measures, incident reporting, intellectual-property allocation, and termination. Fourth, there should be procedures in place for handling access, correction, erasure, and objection requests, as well as routes for reporting data breaches and copyright complaints. Fifth,

employee training should focus on teaching staff to verify factual claims, safeguard confidential information, and recognize when human judgment is necessary.

Regulators should offer consistent examples of acceptable safeguards and practical guidance for small businesses at the policy level. It is important for policymakers to promote effective licensing and transparency mechanisms, all while safeguarding legitimate trade secrets and security interests. It would be a mistake to say that all AI training is lawful or that all training is unlawful: the right answer depends on the jurisdiction, the data involved, the acts performed, exceptions, rights reservations, and the effect on individuals and rights holders.

5.4 Implementation Roadmap

Step 1 — Recognize and categorize: take stock of AI tools, model providers, datasets, objectives, and impacted groups; categorize privacy, copyright, security, and operational risk.

Phase 2 — Assess and authorize: confirm legal bases, review provider contracts, conduct DPIAs where required, evaluate copyright permissions and approve use before deployment.

Phase 3 — Launch with precautions: set up access controls, limit data exposure, educate employees, test outputs, and designate a responsible owner.

Phase 4 — Monitor and enhance: assess incidents, complaints, system modifications, and legal changes; refresh documentation or halt a use when risks cannot be managed. Valuable indicators encompass the percentage of AI systems evaluated pre-deployment, the completion rate of employee training, the duration for resolving rights requests, the frequency of unauthorized disclosures, and the count of unresolved high-risk findings. Governance becomes observable through these measures, enabling management to prioritize resources.

6. CONCLUSION

Generative AI presents significant opportunities to enhance productivity, automate knowledge work, and create new services. However, because it relies on extensive data and can produce new content, it poses challenges to established methods of privacy, copyright, and accountability. According to the EDPB's Opinion 28/2024, assessing model anonymity and legitimate interests must be done on a case-by-case basis. Copyright issues are distinct yet interconnected. Whether it is legal to train on protected works varies based on jurisdiction, licenses, exceptions, rights reservations, and specific facts. The legal framework in place and the extent of human creative contribution dictate copyright in outputs. These questions should not be conflated: lawful access to a model does not automatically confer exclusive rights over every output, and an output's copyrightability does not settle whether the training process was lawful.

When businesses deploy solutions, they add risks related to personal-data disclosure, confidential information, inaccurate outputs, third-party rights, and dependence on suppliers. Regulators need to offer clear and proportionate guidance that safeguards individuals, creators, and legitimate innovation. A measured response includes data provenance, privacy and copyright assessments, strong supplier contracts, human oversight, staff training, and ongoing monitoring. In the end, responsible governance goes beyond being a compliance cost; it is an investment in trust, resilience, and long-term competitiveness.

7. REFERENCES

European Parliament and Council of the European Union. (2019). Directive (EU) 2019/790 of 17 April 2019 on copyright and related rights in the Digital Single Market. Official Journal of the European Union, L 130, 92–125. https://eur-lex.europa.eu/eli/dir/2019/790/oj

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

European Parliament and Council of the European Union. (2026). Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2026/1744/oj

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 33–44). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372873

United States Copyright Office. (2025a). Copyright and artificial intelligence, Part 2: Copyrightability. https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf

Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. International Data Privacy Law, 7(2), 76–99. https://doi.org/10.1093/idpl/ipx005

8. APPENDICES

Appendix: Practical AI Compliance Checklist

- Identify the business purpose and accountable owner of each AI use case.
- Map personal data, confidential information and third-party content processed.
- Verify data sources and document the applicable legal basis.
- Assess whether a DPIA or specialist legal review is required.

- Check copyright licences, permissions, exceptions and rights reservations.
- Review provider terms for training, retention, deletion, security and transfers.
- Configure role-based access and safeguards against data loss.
- Test for personal-data extraction, unauthorised disclosure and harmful outputs.
- Train employees on approved tools and restricted information.

- Establish procedures for rights requests, complaints and incidents.
- Document approvals, risk assessments and human review.
- Monitor changes in systems, contracts and applicable law.