# Counsel at Every Stage: The Advocate as Trusted Advisor Across the Artificial Intelligence Lifecycle in Kenya

**Dr. Mugambi Laibuta**[1]

*Paper prepared for the Law Society of Kenya Annual Conference 2026 → "Business Unusual": Legal Practice in the Era of Artificial Intelligence*

**Abstract**

*The dominant professional conversation on artificial intelligence casts the Advocate as a user of AI tools. This paper argues that the more consequential role is that of advisor to clients who build, procure, deploy and operate AI systems, and that competent advisory work in this role demands mastery of a risk landscape far broader than data protection. Drawing on the Constitution of Kenya, 2010, the Data Protection Act, the Computer Misuse and Cybercrimes Act, the Employment Act, the sector-specific regulatory frameworks, intellectual property statutes and the emerging AI regulatory agenda, the paper makes four contributions. It situates data protection within a wider frame of data governance and cybersecurity; it proposes structured legal risk assessment as the Advocate's core method; it insists on precise identification of the sector regulations and controls governing each deployment; and it maps the advisory mandate across eight stages of the AI lifecycle, from conception to decommissioning. The Kenyan content moderation litigation against Meta and Sama is examined as an early demonstration of how AI-adjacent labour claims will be pleaded and adjudicated. The paper concludes that working knowledge of constitutional law, data governance, cybersecurity, employment law, intellectual property, surveillance governance, sectoral regulation and commercial contracting is not an aspirational profile for the Kenyan AI advisor. It is the minimum level of competence that the scale and complexity of AI-enabled harm now demands.*

***Keywords:*** *artificial intelligence; AI lifecycle; legal risk assessment; data governance; cybersecurity; labour rights; sectoral regulation; Kenya*

## I. Introduction

Lawyers serve as trusted architects of their clients' legal relationships. They translate complex regulatory environments into actionable guidance and anticipate risk before it matures into liability. The rise of artificial intelligence does not displace this advisory function; it expands it and makes it significantly more complex. AI systems now mediate credit decisions, recruitment, medical triage, content moderation, tax administration and public service delivery in Kenya. Each of these

[1] *Advocate of the High Court of Kenya; President, Data Privacy and Governance Society of Kenya*

mlaibuta@gmail.com

applications sits at the intersection of multiple bodies of law, and each generates legal exposure that begins long before the system is switched on and continues long after it is switched off.

Current discourse within the legal profession, in Kenya and elsewhere, frames Advocates primarily as AI users: practitioners deciding whether and how to integrate generative tools into research, drafting and case management workflows. That framing is legitimate, but it is incomplete. The more consequential question concerns the Advocate's role as advisor to clients who build, procure, deploy and operate AI systems. In that role, the Advocate is not a passive consumer of technology but an active participant in shaping how AI is designed, acquired, governed and held accountable from conception to decommissioning.

This paper makes four claims. First, the profession's preoccupation with the Advocate-as-user obscures the larger advisory market and the larger professional duty that AI adoption creates. Second, the data protection frame that currently dominates Kenyan conversations on AI and the law, while indispensable, is incomplete even on its own terms: the client's obligations extend into data governance generally and into cybersecurity, and beyond those into constitutional, labour, proprietary, and commercial exposure. Third, this multiplicity of exposure is manageable only through method, a structured legal risk assessment that identifies, rates and allocates risk before it crystallises and through precise identification of the sector-specific regulations and controls that govern the deployment. Fourth, the practical unit of analysis for AI advisory work is the lifecycle of the system itself.

The paper proceeds as follows. Part II reframes the profession's AI question from use to advice. Part III widens the lens from data protection to data governance and cybersecurity. Part IV proposes structured legal risk assessment as the Advocate's core method. Part V addresses the imperative of locating each AI system within its sector-specific regulatory habitat. Part VI, the core of the paper, walks through the eight lifecycle stages, drawing on the Kenyan content moderation litigation to illustrate the labour dimension. Part VII draws out the implications for professional competence, legal education and the regulation of the profession. Part VIII concludes.

## II. From User to Advisor: Reframing the Profession's AI Question

The distinction between the Advocate as user and the Advocate as advisor is not semantic. It determines what the profession trains for, what continuing professional development prioritises, and what clients can legitimately expect when they retain counsel on an AI matter.

The Advocate-as-user question is essentially one of professional conduct: may I use a generative tool for research or drafting, and if so, under what duties of competence, confidentiality and candour? These questions matter, and courts in several jurisdictions have already sanctioned practitioners for filing fabricated authorities generated by large language models. But they are questions about the lawyer's own workflow, and they can be answered with a relatively narrow set of professional-ethics principles.

The Advocate-as-advisor question is structurally different. Here the client, not the lawyer, is the actor: a bank procuring a credit-scoring model; a hospital deploying diagnostic software; a State Department piloting facial recognition in public spaces; a logistics firm subscribing to an algorithmic workforce-management platform; a startup training a proprietary model on scraped Kenyan-language data. In each case the client's decisions generate legal consequences under multiple statutes simultaneously, and the client will rarely perceive those consequences without counsel. The advisory role is therefore anticipatory and architectural. The Advocate who understands the technology's lifecycle can design compliance into the system; the Advocate who does not will be confined to litigating its failures.

This reframing also corrects a market misconception. Most Kenyan organisations adopting AI will never write a line of model code. They will purchase, license or subscribe to off-the-shelf products developed abroad. It is tempting to conclude that legal exposure therefore sits with the foreign developer. The opposite is true: adoption without development does not reduce legal exposure, it redistributes it, concentrating risk in the contract, in the deployment context and in the client's continuing operational conduct precisely the domains in which the Advocate operates.

## III. Beyond Data Protection: Data Governance and Cybersecurity

### *A. Data Protection as Necessary but Not Sufficient*

The Data Protection Act, and its subsidiary regulations have rightly become the pillar for technology law practice in Kenya. The Act establishes the Office of the Data Protection Commissioner (ODPC), imposes registration, lawful-basis, transparency, minimisation and security obligations, requires data protection impact assessments for high-risk processing, and confers on data subjects a right not to be subjected to certain solely automated decisions. Any AI system that processes personal data engages the Act, and most do.

The difficulty is not that the data protection frame is wrong but that it is partial in two distinct senses. It is partial horizontally, because AI risk also arises under constitutional, labour, proprietary, environmental and commercial law, a point developed throughout this paper. But it is also partial vertically, within the data domain itself: personal data protection is one province of a larger territory of data governance, and the security obligations in the Act are one province of a larger territory of cybersecurity law. An Advocate who advises only on the Data Protection Act has not even completed the data brief, let alone the AI brief.

### *B. Data Governance Generally*

AI systems are built on data that is frequently not personal data at all: operational records, sensor and telemetry data, financial ledgers, geospatial datasets, agricultural and climatic data, anonymised and aggregated statistics, and public sector information. None of these fall squarely within the Data Protection Act, yet all of it raises governance questions the client should answer and the Advocate should anticipate. Who owns or controls the dataset, and under what licence terms may it be used to train or fine-tune a model? What contractual confidentiality and trade secret

obligations attach to it? What data quality, accuracy and provenance standards ought to be met before the dataset can responsibly drive automated decisions recognising that a model trained on incomplete or unrepresentative Kenyan data will fail Kenyan users in ways that engage the discrimination analysis discussed below? What records management and retention rules apply, particularly for public bodies subject to the Access to Information Act, whose disclosure obligations extend to information held about the operation of algorithmic systems? And where anonymisation is relied upon to take a dataset outside the Data Protection Act, has the client tested the robustness of that anonymisation against re-identification, or merely asserted it?

The advisory product at this level is a data governance framework: an instrument that inventories the client's data assets, classifies them by sensitivity and legal regime, assigns stewardship and accountability roles, fixes quality and provenance standards for data feeding AI systems, and aligns the whole with the national policy direction signalled by the Kenya National AI Strategy 2025–2030, Draft National Data Governance Policy, Draft National AI Policy, and the broader digital economy agenda. Data governance, so understood, is not a compliance annex to data protection; it is the substrate on which trustworthy AI is built, and it is squarely legal work.

### *C. Cybersecurity*

AI systems are attack surfaces. They can be compromised through the poisoning of training data, the theft or inversion of models, adversarial inputs engineered to produce erroneous outputs, and prompt injection against generative systems alongside the conventional intrusions to which any networked information system is exposed. Kenyan law meets this reality through at least three instruments that the AI advisor should hold together. First, the Computer Misuse and Cybercrimes Act criminalises unauthorised access, interference and interception, and creates the National Computer and Cybercrimes Co-ordination Committee. Second, the Critical Information Infrastructure regulations made under that Act establish a designation regime for critical information infrastructure, with attendant protection, audit and incident-reporting obligations, a regime into which AI systems supporting essential services in banking, health, energy, transport and telecommunications may readily fall. Third, the Data Protection Act's own security-safeguards and breach-notification provisions require notification of the ODPC within seventy-two hours of becoming aware of a breach affecting personal data.[2]

For the Advocate, three practical consequences follow. Incident response planning ought to be legally coordinated in advance, because a single security event may simultaneously trigger ODPC notification, critical infrastructure reporting, sectoral regulator notification and contractual notice obligations to counterparties, each on a different clock. Vendor contracts should impose security standards, audit rights, incident cooperation duties and liability allocation commensurate with the client's own statutory exposure, since the client cannot outsource its regulatory accountability. And cybersecurity representations made to regulators, insurers and the market should be verified

[2]Data Protection Act, 2019, ss 41 (security safeguards) and 43 (notification and communication of breaches of security).

before they are made, because an overstated security posture converts a technical failure into a misrepresentation claim.

## IV. Legal Risk Assessment as Advisory Method

This Part describes the discipline through which the Advocate manages exposure. Multi-regime exposure cannot be handled through intuition or through sequential, single-statute opinions. It requires structured legal risk assessment: a repeatable method that converts the statute book into a decision-making instrument for the client. The method has five movements.

First, contextual scoping. The Advocate establishes what the system does, whose data it ingests, whom its outputs affect, in which sector it operates, and at which lifecycle stage it stands. Scoping determines everything that follows: a chatbot answering product queries and a model scoring loan applicants are not the same risk object, however similar their underlying technology.

Second, legal mapping. Against that factual scope, the Advocate maps the applicable regimes: constitutional rights engaged; data protection, data governance and cybersecurity obligations; employment and labour exposure; intellectual property positions; environmental obligations; sectoral regulation and licensing conditions; and the contractual matrix binding the client to vendors, customers and employees. The output is a legal register, a living document, not a one-off memorandum.

Third, risk rating. Each identified risk is rated for likelihood and severity, with severity assessed across four dimensions: regulatory sanction (including ODPC penalties and sectoral enforcement), civil liability (including constitutional petitions and class claims of the kind discussed in Part VI), operational disruption (including suspension of a licence or an injunction restraining deployment), and reputational harm. Rating forces prioritisation; not every risk warrants the same response, and the client's resources are finite.

Fourth, treatment and allocation. For each material risk the Advocate identifies a treatment: eliminate (redesign the system or abandon the feature), reduce (impose human review, retrain, restrict the deployment context), transfer (contractual indemnities, insurance), or accept (documented, board-level acceptance of residual risk). Treatment is where legal risk assessment becomes drafting: the risk register drives the data processing agreement, the vendor warranties, the employment consultation plan and the incident-response protocol.

Fifth, review. AI risk is not static: models drift, regulations mature, the AI Bill will crystallise into an Act, and case law as Part VI shows is developing quickly. The assessment should therefore be calendared for periodic review and re-triggered by defined events: a new deployment context, a material model update, a security incident, an enforcement action against a comparable operator.

Two features of this method deserve emphasis. It integrates, rather than duplicates, the impact assessments the law already requires: the DPIA under section 31 of the Data Protection Act[3] and the human rights impact assessment discussed below become modules within the single risk architecture, not parallel paper exercises. And it is defensible: a documented, dated, iterated risk assessment is the client's best evidence of due diligence before a regulator or a court, and conducted under legal advice its candour can be protected by privilege.

## V. Locating the System: Sector-Specific Regulation and Controls

AI is not deployed in the abstract; it is deployed into sectors, and every regulated sector arrives with its own statutes, regulators, licensing conditions, prudential standards[4] and supervisory expectations. A recurring failure in AI advisory work is horizontal analysis without vertical grounding: counsel provides advisory on data protection and constitutional rights while missing the sector regulator whose licence condition the deployment quietly breaches. Clarity about the sector regulations and controls applicable to the specific deployment is therefore not a refinement of the advisory task; it is a precondition of it.

The examples are concrete. A credit-scoring or lending model sits under the Central Bank of Kenya's supervisory framework, including the Digital Credit Providers Regulations, 2022, with their licensing, disclosure and debt-collection conduct requirements, and the CBK's cybersecurity guidance for the banking sector.[5]

A diagnostic or triage tool engages the Digital Health Act with its provisions on health data governance and the digital health ecosystem alongside the Health Act and the professional-regulation regimes of the Kenya Medical Practitioners and Dentists Council and the Pharmacy and Poisons Board.[6]

Robo-advisory and algorithmic trading engage the Capital Markets Authority; algorithmic underwriting and pricing engage the Insurance Regulatory Authority; and algorithmic pricing coordination across competitors engages the Competition Authority of Kenya.[7] Systems deployed in telecommunications networks engage the Communications Authority under the Kenya Information and Communications Act.[8] And where the deployer is a public entity, the acquisition itself is governed by Article 227 of the Constitution and the Public Procurement and Asset Disposal

---

[3]Data Protection Act, 2019, s 31 (data protection impact assessment); see also *Katiba Institute v Joe Mucheru, Cabinet Secretary Ministry of ICT* [2021] eKLR, in which the High Court required a DPIA to be conducted before further processing of personal data under the Huduma card programme.

[4] Central Bank of Kenya, *Prudential Guidelines for Institutions Licensed under the Banking Act*.

[5] Central Bank of Kenya, Guidance Note on Cybersecurity for the Banking Sector (2017); Central Bank of Kenya (Digital Credit Providers) Regulations, 2022.

[6] Digital Health Act, Laws of Kenya; Health Act.

[7] Capital Markets Act, Laws of Kenya; Insurance Act, Laws of Kenya; Competition Act.

[8] Kenya Information and Communications Act.

Act, so that a flawed AI procurement is challengeable as a procurement, quite apart from what the system later does.[9]

Three advisory disciplines follow. First, regulator identification precedes obligation analysis: the first page of the legal register described in Part IV should name every regulator with supervisory jurisdiction over the deployment, because each is a potential source of approval requirements, reporting duties and enforcement.

Second, licence conditions are read against the deployment: an AI feature added to a licensed service may fall outside the scope of the existing licence, require notification or approval, or contradict a condition drafted before the technology existed.

Third, supervisory expectations are tracked, not assumed: sector regulators increasingly issue guidance, circulars and consultation papers on algorithmic systems, and the emerging AI-specific legislative framework will layer horizontal obligations over these vertical regimes rather than replace them. The Advocate's value lies in reading the two dimensions together and in warning the client that satisfying the ODPC does not answer the Central Bank.

## VI. The Advisory Mandate Across the AI Lifecycle

The risk landscape of Parts III to V crystallises at identifiable moments in the life of the system. This Part maps the advisory mandate across eight lifecycle stages, because it is the lifecycle not the statute book that determines when each risk matures and when advice is still capable of preventing harm.

### *A. Conception and Design*

The Advocate's highest-leverage intervention occurs before any code is written or any contract signed. At conception, the Advocate performs the sector location exercise described in Part V and advises on the regulatory classification of the proposed system: classification determines the compliance architecture of everything that follows.

Design decisions also shape downstream rights impacts. Systems developed without equality audits risk embedding discrimination contrary to Article 27 of the Constitution,[10] and the constitutional prohibition applies to indirect discrimination produced by facially neutral variables. Where the contemplated deployment may displace workers, consultation and redundancy obligations under section 40 of the Employment Act arise at the outset, not at the point of termination.[11] Energy-intensive architectures require early environmental review under the Environmental Management and Co-ordination Act, Climate Change Act and the environmental

[9] Public Procurement and Asset Disposal Act, Laws of Kenya; Constitution of Kenya, 2010, Article 227.

[10] Constitution of Kenya, 2010, Article 27.

[11] Employment Act, s 40 (termination on account of redundancy).

impact assessment regime.[12] The Advocate who enters at the design stage shapes the compliance architecture of the entire system; the Advocate who appears only at deployment inherits its structural weaknesses.

### *B. Procurement and Vendor Selection*

Procurement is the stage encountered by the largest number of Kenyan organisations and the stage most frequently overlooked by their legal advisors. Because most adopters will license rather than build, the procurement contract becomes the primary instrument of AI governance available to the client. In the AI era, procurement is not a commercial formality; it is a legal event of the highest significance.

The Advocate's due diligence at this stage is rigorous and multidimensional. It begins with vendor data practices: was the training data lawfully sourced, and does the vendor qualify as a data processor under the Data Protection (General) Regulations,[13] thereby triggering the requirement for a compliant data processing agreement allocating security, breach-notification and sub-processing obligations? It extends to intellectual property: standard-form contracts routinely transfer ownership of AI-generated outputs, derivative models and client insights to the vendor, shifting substantial commercial value in ways clients rarely appreciate without legal guidance.[14] It encompasses regulatory classification under the emerging AI legislative framework, scrutiny of embedded surveillance and data collection capabilities, interrogation of the vendor's cybersecurity posture and incident history, interrogation of environmental credentials including data centre energy sourcing, and assessment of data localisation and cross-border transfer compliance.[15]

Finally, the Advocate negotiates the provisions that preserve the client's regulatory accountability throughout the system's operational life: audit and inspection rights, transparency and explainability warranties, performance and non-discrimination representations, security standards and incident cooperation duties, indemnities calibrated to regulatory penalties, and exit provisions guaranteeing data return, data deletion and continuity of service. A client cannot discharge obligations it has contracted away the ability to verify.

### *C. Data Collection and Training and the Lessons of the Content Moderation Litigation*

Where clients develop proprietary systems, the advisory mandate at the data collection and training stage extends well beyond the Data Protection Act. Lawful basis, purpose limitation and minimisation analysis under the Act remains foundational, but three further dimensions demand attention: intellectual property compliance for scraped training corpora. Two, the proportionality

---

[12] Climate Change Act.

[13] Data Protection (General) Regulations, 2021.

[14] Copyright Act, Laws of Kenya; Industrial Property Act.

[15] Data Protection Act, 2019, ss 25(g) and 28; Data Protection (General) Regulations, 2021, Part V (data localisation and transfers outside Kenya).

of surveillance under Articles 24 and 31 of the Constitution is particularly critical where biometric data including facial images, voiceprints and gait patterns is harvested. The *Huduma Namba* jurisprudence demonstrates that Kenyan courts will subject such practices to searching constitutional scrutiny;[16] and, three, most vividly demonstrated in recent Kenyan litigation, labour protections for the human workforce on which AI systems are built.

The AI value chain is labour-intensive at precisely the points least visible to the end client: data annotation, model evaluation and content moderation. Kenya hosts a significant share of this work, and Kenyan courts are now the global reference point for its legal consequences. In *Motaung v Samasource Kenya EPZ Ltd t/a Sama*, a former Facebook content moderator sued the local outsourcing contractor together with Meta Platforms, Inc and Meta Platforms Ireland, alleging exploitation, inadequate mental health support for workers exposed to graphic content, pay discrimination and union-busting, and pleading violations of a broad catalogue of constitutional rights including dignity, fair labour practices and freedom of association. The Employment and Labour Relations Court declined to strike out the foreign Meta entities, holding that the question whether they operated in Kenya was a matter for trial;[17] in the parallel petition brought by 185 moderators challenging their mass redundancy when the moderation hub was closed, the court granted leave to serve the foreign respondents abroad and issued interim orders restraining the redundancies; and in September 2024 the Court of Appeal dismissed Meta's consolidated appeals, confirming that both petitions would proceed to trial in Kenya.[18]

Three lessons follow for the AI advisor. First, outsourcing does not insulate the principal: the moderators' pleadings treat the platform and its contractor as jointly responsible, and the courts have refused to let corporate structure defeat jurisdiction at the interlocutory stage. A client procuring data annotation, content moderation or evaluation services through a Kenyan business-process-outsourcing chain should treat the chain's labour practices as part of its own risk exposure and subject them to appropriate due diligence. Second, AI labour claims are pleaded constitutionally, not merely contractually: dignity under Article 28, fair labour practices under Article 41 and freedom of association frame the dispute,[19] which enlarges both the available remedies and the reputational stakes. Third, occupational safety and health obligations extend to psychosocial harm: prolonged exposure to disturbing content is a workplace hazard, engaging the Occupational Safety and Health Act's duty to ensure the safety and health of workers,[20] and the

[16] *Nubian Rights Forum v Attorney General; Child Welfare Society (Interested Parties)* [2020] eKLR (the *Huduma Namba* case).

[17] *Motaung v Samasource Kenya EPZ Ltd t/a Sama*, Nairobi ELRC Petition No. E071 of 2022 (ruling of 6 February 2023, Gakeri J, declining to strike out Meta Platforms, Inc and Meta Platforms Ireland Ltd for want of jurisdiction).

[18] *Meta Platforms, v Motaung; Kenya National Human Rights and Equality Commission (Interested Parties)* (Civil Appeal E232 & E445 of 2023 (Consolidated)) [2024] KECA 1262 (KLR) (20 September 2024).

[19] Constitution of Kenya, 2010, Article 41 (fair labour practices) and Article 28 (human dignity).

[20] Occupational Safety and Health Act.

advisory response, exposure limits, wellness support, informed consent to the nature of the work, adequate remuneration for its hazards ought to be designed into the engagement, not retrofitted in litigation.

### *D. Testing and Validation*

At testing and validation, the Advocate's task is to convert technical evaluation into legally defensible documentation. The Data Protection Act requires a data protection impact assessment where processing is likely to result in high risk to data subjects, and the High Court has demonstrated its willingness to order that such assessments be conducted before processing proceeds. But the DPIA is a floor, not a ceiling. The Advocate should develop human rights impact assessments and regulatory risk assessments alongside the DPIA, evaluating the system's effects across the full catalogue of constitutional rights it touches and both instruments should feed the risk register described in Part IV rather than standing as isolated documents.

Bias testing across race, ethnicity, gender, disability and socioeconomic status is a constitutional obligation, not a technical preference: Article 27 binds private actors as well as the state, and an untested system is an unaudited discrimination risk. Where AI systems operate in criminal justice, social welfare or employment screening, domains implicating the non-derogable right to a fair trial, the equality guarantee and the right to fair administrative action[21] validation protocols should be designed to withstand judicial scrutiny, and the Advocate is responsible for ensuring that they do.

Security testing belongs here too: adversarial testing and penetration testing of the model and its surrounding infrastructure are the practical discharge of the security-safeguards duty discussed in Part III. Test records, validation thresholds and documented remediation decisions will constitute the client's evidentiary defence when outcomes are challenged.

### *E. Deployment*

Deployment converts latent design risk into live legal exposure. Three advisory workstreams operate concurrently. First, contractual liability allocation: the deployment agreement should resolve, in advance, how responsibility for erroneous outputs, service failures, security incidents and regulatory penalties is distributed among developer, vendor, integrator and deployer.

Second, regulatory notification: the sector regulators identified in Part V increasingly expect to be informed before algorithmic systems are introduced into regulated functions, and registration or licensing obligations under data protection law must be current.

Third, surveillance and employment compliance. Deployments involving monitoring of the public or of workers require a documented proportionality analysis under Articles 24 and 31 of the Constitution:[22] is the intrusion rationally connected to a legitimate aim, and is it the least restrictive

[21] Constitution of Kenya, 2010, Articles 25(c), 27 and 50.

[22] Constitution of Kenya, 2010, Articles 24 and 31.

means available? Where algorithmic management systems monitor and discipline workers in real time, assigning tasks, scoring productivity, triggering sanctions, immediate questions arise concerning fair labour practices under Article 41 and human dignity under Article 28. The content moderation litigation discussed above signals how readily Kenyan courts shall entertain such claims in constitutional form. The Advocate should insist on human review of adverse algorithmic decisions affecting employees and on transparent disclosure of monitoring to those subjected to it.

### *F. Operation*

The operational phase is the longest stage of the lifecycle and the stage at which advisory relationships too often lapse. Five continuing duties define the mandate. First, automated decision-making rights: section 35 of the Data Protection Act entitles data subjects to contest decisions based solely on automated processing that produce legal or similarly significant effects,[23] and the client requires operational machinery, notice, human intervention, reconsideration to honour those rights at scale.

Second, regulatory engagement: the Advocate manages ODPC complaints, investigations, audits and enforcement notices,[24] alongside the supervisory relationships with sector regulators, functions that demand both procedural fluency and a candid internal assessment of the client's exposure.

Third, security and incident management: the coordinated incident-response plan described in Part III should be maintained, rehearsed and executed when not if a security event occurs, with the seventy-two-hour ODPC clock and any critical infrastructure and sectoral reporting obligations mapped in advance.[25]

Fourth, environmental reporting: AI-intensive operations attract reporting and mitigation obligations under the climate and environmental framework as compute footprints grow.

Fifth, and most distinctively, continuing discrimination monitoring: model performance drifts as populations and data change, and a system validated as fair at deployment can become discriminatory in operation. The duty to monitor outcomes is continuing, and the Advocate should embed periodic fairness review and periodic refresh of the legal risk assessment itself into the client's governance calendar.

---

[23] Data Protection Act, 2019, s 35 (rights in relation to automated decision-making).

[24] Data Protection Act, 2019, ss 8, 9, 56–66 (functions, complaints, investigation and enforcement powers of the Office of the Data Protection Commissioner).

### *G. Auditing and Accountability*

Auditing is where governance is demonstrated rather than merely asserted. The Advocate's contribution is to design audit frameworks that are both substantively comprehensive and legally protected. Substantively, the audit should integrate human rights due diligence, surveillance proportionality review, intellectual property licence compliance, data protection and data governance conformity, cybersecurity posture assessment against the client's statutory and contractual obligations, sectoral licence-condition compliance, and environmental documentation covering energy consumption and hardware disposal. Procedurally, the Advocate structures the exercise so that candid internal findings attract legal professional privilege where the audit is genuinely conducted for the purpose of legal advice, while recognising that privilege cannot be used to warehouse evidence of continuing unlawful conduct.

The audit stage is also where the contractual groundwork laid at procurement pays its dividend: audit rights negotiated at stage two are exercised at stage seven. A client who cannot inspect its vendor's system cannot account for it to the ODPC, to sectoral regulators, or to a court.

### *H. Decommissioning*

The lifecycle ends, but the obligations do not. At decommissioning, the Advocate addresses six closing workstreams: data deletion and return obligations under the Data Protection Act and the data processing agreement; secure destruction and sanitisation of storage media, so that decommissioning does not itself become the security breach; termination of intellectual property licences, including the fate of models fine-tuned on the client's data; workforce transition planning where the system's retirement or replacement displaces employees, engaging the redundancy procedure under the Employment Act; vendor contract termination, including post-termination confidentiality, escrow release and transition assistance; and electronic waste compliance arising from hardware retirement under the Sustainable Waste Management Act and the environmental framework.[26]

Decommissioning is also an evidentiary moment. Retention schedules should balance deletion obligations against litigation-hold duties where disputes concerning the system's operational life remain live or foreseeable. The Advocate who plans decommissioning at procurement through exit clauses, escrow and deletion warranties spares the client the far costlier exercise of negotiating an exit from a position of dependency.

## VII. Conclusion

The framework developed above yields a demanding competence profile. A competent AI advisor in Kenya requires working knowledge of constitutional law, data protection and data governance, cybersecurity law, employment law, intellectual property, environmental compliance, surveillance governance, sector-specific regulation and commercial contracting together with the

[26] Environmental Management and Co-ordination Act.

methodological discipline of structured legal risk assessment that binds them into usable advice. This is not aspirational. It is the minimum level of advisory competence demanded by the scale and complexity of AI-enabled harm, and it follows directly from the Advocate's existing professional duties of competence and diligence.[27]

The question the Kenyan legal profession should be asking about artificial intelligence is not primarily whether Advocates may use it, but whether Advocates are prepared to advise on it. The client building, procuring, deploying or operating an AI system faces intersecting exposure under the Constitution, the Data Protection Act and the wider law of data governance, the cybercrimes and critical infrastructure framework, the Employment Act whose reach into the AI value chain the content moderation litigation has now made unmistakable the intellectual property statutes, the environmental framework, the sector-specific regimes governing each deployment, and the ordinary law of contract. That exposure crystallises at identifiable moments across the system's lifecycle, from conception to decommissioning, and it is manageable only through disciplined, iterated legal risk assessment.

The Advocate who understands that lifecycle occupies the role the profession has always claimed for itself: trusted architect of the client's legal relationships, anticipating risk before it matures into liability. The Advocate who does not will find that the compliance architecture was designed without counsel, and that the profession's remaining role is to litigate the consequences. Counsel at every stage is not a marketing proposition. In the era of artificial intelligence, it is the substance of the advisory duty itself.

[27] Law Society of Kenya Act; Advocates Act; Law Society of Kenya, Code of Standards of Professional Practice and Ethical Conduct.