WHITE PAPER · GOVERNANCE

Auditable AI for Nigerian Health Insurers

A governance and adoption framework under the NHIA Act, NDPA 2023 and the National AI Strategy. With a comparative note on Kenya, Ghana and South Africa.

By Moyo Awoyokun, MB ChB, MBA, FIHIMCN · Published 26 October 2026 · Ajé Intelligent InsurTech

Executive summary

Nigerian health insurers are adopting automation and artificial intelligence because they have to. Enrolment has passed 23 million, having grown by about a third between the end of 2023 and early 2026. The NHIA's April 2025 tariff review raised capitation by 93 per cent and fee-for-service by 378 per cent. Since 1 April 2025 every treatment authorisation request must be answered within one hour. Private provider prices move monthly against premiums that reset annually. No HMO can absorb that combination with more staff alone.

Vendors have noticed. Every major HMO platform in the market now describes itself as "AI-powered". Very few explain how their AI is governed, and that gap matters because the decisions involved are not trivial. An automated claims decision can withhold payment for care a member has already received. An automated authorisation can deny or delay care that has not yet been given. A generative AI tool can write a confident, wrong sentence into a member's letter.

Nigeria already has the legal frame to govern this, but it is spread across instruments that were not written with health insurance in mind. The NHIA Act 2022 and the NHIA's operational guidelines govern how HMOs authorise and pay. The Nigeria Data Protection Act 2023 and its 2025 implementing directive govern automated decisions, impact assessments and health data. The National AI Strategy, published in September 2025, and NITDA's draft Code of Practice for AI set the direction of

travel. None of them tells an HMO, in operational terms, what it may automate and what it may not.

This paper offers that operational translation. Its core is five controls:

1. **Decide where AI may act.** A decision-rights matrix that separates tasks AI may inform, recommend, approve within limits, and never decide. Denials, emergencies and contested coverage sit in the last category.
2. **Fix the data before the model.** Shared definitions, coded claims and provenance. AI built on inconsistent data produces confident errors at scale.
3. **Record every decision.** Inputs, rule or model version, reason, timestamp and any human override, kept in a form a regulator, provider or member can examine.
4. **Make human oversight real.** Reviewers with the authority, time and information to disagree with the system, and override rates that are measured.
5. **Test, monitor and retire.** Pre-deployment testing in low-risk settings, bias checks across geography, age, sex and plan tier, drift monitoring and a switch that turns the system off.

The paper sets these against the Nigerian legal instruments clause by clause, proposes a phased adoption roadmap, lists the risks specific to generative AI in health insurance, and closes with a diligence checklist for boards and investors. A short annex compares Kenya, Ghana and South Africa, which share Nigeria's broad legal pattern and offer one cautionary lesson.

The central argument is simple. In a low-trust market, an insurer's automation is only as valuable as its ability to show what the automation did and why. Auditability is not a compliance cost added to AI. It is the property that makes AI usable in health insurance at all.

1. Situation: Nigerian health insurance is automating under pressure

1.1 More members, more cost, less time

Three changes since 2023 have altered what it takes to run a Nigerian HMO.

Volume. NHIA enrolment rose from 16.2 million at the end of 2023 to 21.7 million at the end of 2025 and 22.0 million in the first quarter of 2026, and by the Director-General's September 2026 statement to more than 23 million. The NHIA Act 2022 makes health insurance mandatory, and enforcement is now visible: federal ministries and agencies have been directed to enrol their staff, and the NHIA's own enforcement record lists thousands of complaints resolved and dozens of HMOs and facilities sanctioned.

Cost. The April 2025 tariff review was the first in over a decade. Private providers repriced in parallel, driven by imported drug costs, energy and staffing. HMO premiums rose between 8 and 59 per cent by plan tier from 2024 to 2025. Every naira of that increase has to be administered: priced, authorised, adjudicated and paid.

Time. From 1 April 2025, an HMO must approve a treatment authorisation request, or communicate a reasoned refusal, within one hour of the provider's submission. If it does not, the provider treats the patient and reports the HMO. In emergencies no code is required before treatment at all.

1.2 The industry's response is automation

Faced with these pressures, HMOs and their vendors are automating four kinds of work:

- **Claims adjudication:** checking eligibility, benefit, tariff and coding rules, and paying clean claims without manual review.
- **Pre-authorisation:** deciding routine requests against the same rules within the regulatory clock.
- **Fraud, waste and abuse detection:** finding patterns across the book that no individual reviewer can see, such as duplicate billing, upcoding and implausible provider activity.
- **Member and provider communication:** answering eligibility questions, explaining decisions and processing documents, increasingly with generative AI.

The economic case is real. McKinsey estimates that generative AI could unlock between $2.1 billion and $3.2 billion of economic value for African insurers, and notes

that while most current innovation is in South Africa, insurers in Nigeria, Ghana, Kenya and Morocco are experimenting. At the InsurTech Forum Nairobi in 2026, the theme of the industry's AI challenge was not whether to use AI but how to make it auditable.

2. Complication: a governance gap

2.1 Claims without controls

Most HMO platforms marketed in Nigeria now describe AI-powered adjudication, AI fraud detection or AI-driven automation. Public material rarely describes what the AI decides, what it is not allowed to decide, how its decisions are recorded, or how a member or provider can challenge one. That silence is understandable in marketing. It becomes a problem when a regulator, an auditor or an aggrieved member asks.

2.2 Why health insurance decisions are different

Three features make automated decisions in health insurance more consequential than in most financial services.

The harm is physical, not only financial. A wrongly declined loan can be re-applied for. A wrongly delayed authorisation can mean a patient leaves the hospital untreated. Press investigations in 2025 and 2026 documented exactly this pattern at accredited Nigerian hospitals.

The data is sensitive by law. Health data is sensitive personal data under the Nigeria Data Protection Act 2023, with stricter conditions for processing. Diagnosis codes, prescriptions and claims history are all health data.

The incentives are asymmetric. An insurer benefits financially from every claim it does not pay. A system that denies more is, in the short run, a system that looks more profitable. Without governance, automation will drift towards denial, and members and providers know it. Trust, already low in Nigerian health insurance, will fall further.

2.3 Weak controls are expensive: the Kenyan lesson

Kenya's Social Health Authority, which replaced the National Hospital Insurance Fund in 2024, invested in digital claims infrastructure. In January 2026 the Cabinet Secretary for Health told members of the National Assembly, at a retreat in Naivasha, that about Ksh 11 billion had been lost to fraudulent claims between October 2024 and April 2025, including claims for surgeries never performed, caesarean sections at a facility with no theatre, and a single member registered with 381 dependants. More than a thousand facilities were closed and dozens of clinicians deregistered from the platform. The Ministry said tight procedures and a secure IT system had been installed, and the Authority ran automated fraud flagging. Fraud continued regardless.

The lesson is not that digital systems or fraud analytics fail. It is that flags without enforcement, provider vetting and accountable human follow-up do not stop fraud; a digital system without that governance simply processes fraud faster. Technology amplifies whatever governance sits around it.

2.4 The law exists but is scattered

Nigeria is better placed than many markets. It has a modern health insurance statute, a data protection act with an explicit clause on automated decisions, an implementing directive that addresses emerging technologies, and a national AI strategy. What it lacks is a single place where an HMO's compliance officer, chief medical officer and technology lead can see what all of these require of an AI-assisted authorisation desk. Section 3 attempts that map.

3. Question

What should a Nigerian health insurer automate, how should it govern what it automates, and how can it show a regulator, a provider, a member or an investor that the governance is real?

4. The Nigerian legal map

This section summarises the instruments that apply to an HMO using automation or AI, and what each requires in operational terms. It is a practitioner's map, not legal advice; each organisation should confirm its obligations with counsel.

4.1 National Health Insurance Authority Act 2022 and NHIA guidelines

The Act makes health insurance mandatory for all Nigerians and legal residents, establishes the NHIA as regulator of HMOs and other health insurance intermediaries, and empowers it to set a basic minimum package of health services, accredit and de-accredit HMOs and providers, regulate quality of care and complaints, and sanction breaches of its operational guidelines.

Three operational requirements bear directly on automation:

- **The one-hour authorisation rule** (in force from 1 April 2025): approve or communicate "no authorisation" within one hour; providers treat and report if the hour is exceeded.
- **Emergency care without prior authorisation:** codes are obtained within 48 hours of the start of emergency care.
- **Record-keeping:** HMOs and providers must keep records of every authorisation request and response.

Implication for AI: automation is effectively required to meet the clock at scale, but every automated response must be recorded, and no automated gate may sit in front of emergency care.

4.2 Nigeria Data Protection Act 2023

The Act applies to any HMO processing the personal data of people in Nigeria. The provisions that matter most for AI are:

- **Section 37, automated decision-making.** A data subject has the right not to be subject to a decision based solely on automated processing, including profiling, that produces legal or similarly significant effects. Exceptions exist where the decision is necessary for a contract, authorised by law with safeguards, or based on consent. Where an exception applies, the controller must still implement safeguards, including the right to obtain human intervention, to express a point of view and to contest the decision.
- **Section 28, data protection impact assessment.** Required where processing is likely to result in high risk to data subjects. Large-scale processing of health data

with automated decision-making meets that description.

- **Section 30, sensitive personal data.** Health data is sensitive and may be processed only on specified grounds.
- **Section 32, data protection officer.** Required for data controllers of major importance, which includes most HMOs given their scale and sector.
- **Section 40, breach notification.** The Commission must be notified within 72 hours of the controller becoming aware of a breach, and data subjects informed where the risk to them is high.
- **Cross-border transfer provisions.** Personal data may leave Nigeria only where the destination offers adequate protection or another lawful basis applies. This bears directly on sending member data to foreign-hosted AI models.

4.3 NDPC General Application and Implementation Directive 2025

The Commission's 2025 directive operationalises the Act. As summarised by Nigerian practitioners, it:

- lists solely automated decisions with legal or similarly significant effects among the processing that requires the data subject's consent (article 18(1)(f)). Read alongside the contract and law exceptions in NDPA s.37(2), this points HMOs towards obtaining consent, or avoiding solely automated adverse decisions altogether, and practitioners read the directive as expecting safeguards against bias;
- requires those deploying emerging technologies, including AI, to conduct DPIAs that address disparate impacts and data subject vulnerability, and to file them with the annual compliance audit return;
- requires testing in controlled, low-risk settings before deployment, continued monitoring after deployment, and retooling, retesting or discarding tools where risks to data subjects are identified (article 43).

Implication for AI: an HMO's AI systems should be listed, assessed and filed, not merely used.

4.4 National Artificial Intelligence Strategy (September 2025)

Published by the Federal Ministry of Communications, Innovation and Digital Economy with NITDA as a key implementing body, the strategy sets a five-year direction to 2029. It names healthcare among priority sectors for adoption, proposes an AI ethics expert group and a national AI risk management framework, and anticipates sector-specific standards developed with the relevant ministries. The Federal Ministry of Health is building a Nigerian Digital Health Architecture and Health Information Exchange into which, in the words of its digital health adviser, all AI tools in health should plug.

4.5 NITDA Draft Code of Practice for AI (2025)

Still a draft at the time of writing, and described here from secondary summaries of NITDA's published direction rather than the draft text itself, the Code is expected to take a risk-based approach. It is reported to define high-risk AI to include systems whose malfunction or bias may significantly affect health or critical services, and to expect providers to give user-facing notices describing AI use and its limits, to enable human intervention in decisions, and to keep technical documentation and logs sufficient for audit.

Implication for AI: although not yet binding, the Code indicates what a regulator will ask for. Building to it now avoids rework later.

4.6 Summary map

INSTRUMENT	STATUS	WHAT IT REQUIRES OF AN HMO USING AI
NHIA Act 2022 and operational guidelines	In force	One-hour authorisation; no gate on emergencies; records of every request and response; accreditation and sanctions
NDPA 2023, s.37	In force	No solely automated decisions with significant effects without an exception and safeguards: human intervention, a voice, a right to contest
NDPA 2023, ss.28, 30, 32, 40	In force	DPIA for high-risk processing; sensitive data conditions; DPO; 72-hour breach notice
NDPA cross-border provisions	In force	Lawful basis and adequacy before member data leaves Nigeria
GAID 2025, arts. 18 and 43	In force	Consent for solely automated decisions with significant effects; DPIA for emerging technology filed with CAR; pre-deployment testing and monitoring
National AI Strategy 2025 to 2029	Policy	Direction on ethics, risk management and health sector standards
NITDA Draft Code of Practice for AI	Draft (reported)	High-risk classification for health; transparency notices; human oversight; logs

5. Answer: the five controls of auditable AI

5.1 Control one: decide where AI may act

The first and most important governance decision is not technical. It is a written statement of which decisions AI may take, which it may only inform, and which it may never take. The matrix below is a starting point.

TIER	AI ROLE	TYPICAL HMO TASKS	HUMAN ROLE	RATIONALE
0. Inform	Extracts, summarises, classifies	Reading claim forms and referral letters; summarising case notes; transcribing calls; coding suggestions	Reviews output before it is used	Low direct risk; errors caught downstream
1. Recommend	Scores and flags	Fraud and anomaly flags; provider steering suggestions; reserving estimates	Decides, with the score visible	The system narrows attention; a person acts
2. Approve within limits	Approves by explicit rule	Elective requests and chronic refills that pass eligibility, benefit, tariff and accreditation checks; clean low-value claims	Samples and audits; handles all failures	Approval of covered care harms no one; bounded value limits exposure
3. Never decide	None	Denials of authorisation or claims; emergencies; contested coverage; clinical judgements; sanctions against providers	Decides, with full reasons	Significant effects on the member; NDPA s.37; trust

Two principles follow from the matrix.

Approve by rule, never deny by rule. An automated system may approve a request that satisfies every rule. A request that fails any rule goes to a person, with the failing rule attached. This keeps the insurer's automated decisions on the side that cannot harm a member, and places every adverse decision with an accountable human. It also removes the asymmetric incentive described in section 2.2: the system can only accelerate payment, not withhold it.

Emergencies are outside automation. The NHIA already exempts emergencies from prior authorisation. No rule engine, model or queue should sit between an emergency and care.

5.2 Control two: fix the data before the model

AI inherits the quality of the data it runs on. In Nigerian health insurance, three data problems recur.

Inconsistent definitions. The same metric, the loss ratio being the most important, is often computed differently in different reports: on paid or incurred claims, on approved or submitted amounts, with or without unadjudicated claims. A model trained or evaluated on one definition and reported on another will mislead its users. Define each metric once, on an incurred basis where claims are concerned, and test that every report uses the same definition.

Uncoded claims. Claims recorded as free text or with incomplete diagnosis and procedure codes cannot be checked against rules or modelled reliably. The single largest one-time improvement available to most HMOs is moving to coded electronic claims with mandatory fields.

Unknown provenance. Every record used by an AI system should be traceable to where it came from and when. Seed, test and production data must never mix.

A practical test: before deploying any model, an HMO should be able to reproduce the model's key input figures from source records by hand, for a sample of cases.

5.3 Control three: record every decision

Every automated or AI-assisted decision should produce a record that contains:

- the request or claim as received, with its arrival timestamp;
- the data the system used, including the member's eligibility and benefit state at that moment;
- the rule or model, and its version, that produced the outcome;
- the outcome and the reason, expressed as a reason code and a plain-language sentence;
- any human review, with the reviewer's identity, decision and reason if they overrode the system;
- the time of the final response to the provider.

This record satisfies the NHIA's record-keeping requirement, supports the safeguards required by NDPA s.37(3), matches the logging the draft NITDA Code anti-

cipates, and is the evidence base for answering complaints. It is also the dataset that shows the HMO where its own rules are too tight or too loose.

Reasons should be written for the recipient. A provider told "PA-07: requested service exceeds annual physiotherapy limit of 12 sessions; 12 used" can act on it. A provider told "declined by system" cannot.

5.4 Control four: make human oversight real

Human oversight fails in two predictable ways. Reviewers without time rubber-stamp whatever the system recommends, a pattern known as automation bias. Reviewers without information cannot tell whether the system is right.

Oversight is real when:

- **reviewers have authority** to overrule the system without escalation for routine cases;
- **reviewers have time:** the queue is sized so that each exception receives genuine attention within the regulatory clock, which is precisely what automating the routine cases buys;
- **reviewers see the reasoning:** the failing rule, the score and its main drivers, and the relevant history;
- **override rates are measured.** An override rate near zero suggests rubber-stamping; a very high one suggests the rules are wrong. Both are signals;
- **there is an appeal.** Members and providers can contest a decision and reach a person who did not make the original one.

5.5 Control five: test, monitor and retire

Before deployment, test each system on historical data and in a limited live setting. Compare its decisions against experienced reviewers. Check outcomes across groups that should be treated equally: by state and region, urban and rural providers, age band, sex and plan tier. GAID article 43 requires testing in controlled, low-risk settings before deployment; a shadow mode, in which the system decides but people still act, is the simplest way to meet it.

After deployment, monitor weekly at first: approval rates, override rates, time to decision, complaints, and outcome differences across groups. Models drift as tariffs, provider behaviour and disease patterns change. Rules drift when benefits change and the rules are not updated.

Retire deliberately. Every system needs an owner who can switch it off, and a documented manual fallback. The GAID expects tools to be retooled, retested or discarded where risks are identified. That is only possible if someone is accountable for noticing.

5.6 Generative AI: specific uses and specific limits

Generative AI is different from rules and conventional models: it produces fluent text that may be wrong, and it is often consumed through third-party services hosted outside Nigeria.

Where it fits in an HMO today (Tiers 0 and 1):

- extracting structured fields from scanned claim forms, referral letters and invoices;
- summarising long case histories for a reviewer;
- drafting member and provider correspondence for human approval;
- answering general benefit questions in member channels, with escalation to a person for anything specific to an individual's cover;
- assisting staff to search policies, tariffs and guidelines.

Where it does not belong:

- determining coverage, medical necessity or claim outcomes;
- assigning final clinical codes without validation;
- sending any communication to a member or provider that has not been checked, where the content affects their entitlement.

Controls specific to generative AI:

- **Grounding.** Generated answers should cite the policy, tariff or record they rely on, so a reviewer can check them.
- **No personal data to unassessed services.** Member data should not be sent to a third-party model unless the DPIA covers it, the contract restricts the provider's use of the data, and the cross-border transfer has a lawful basis. Where possible, de-identify before processing.
- **Prompt and output logging.** Record what was asked and what was produced, as for any other decision support.
- **Fallback.** Staff must be able to do the task without the tool.

6. An adoption roadmap

The roadmap below assumes an HMO starting with largely manual processes. Each phase has entry conditions and measures of success.

PHASE	TIMING	SCOPE	ENTRY CONDITIONS	SUCCESS MEASURES
0. Foundations	Months 0 to 3	Inventory of all automated and AI systems in use; decision-rights matrix approved by the board; DPIA for each system; single timestamped intake for all authorisation requests; decision record standard	Executive sponsor; DPO in post	Every request recorded with arrival time; median time to decision known
1. Assist	Months 3 to 9	Generative AI for document extraction and summaries; fraud and anomaly flags as recommendations; shadow-mode rules for authorisation	Phase 0 complete; data definitions agreed	Reviewer time per case falls; flag precision measured; shadow decisions agree with reviewers on a high share of cases
2. Approve within limits	Months 9 to 18	Rule-based approval for bounded classes: elective requests and chronic refills within benefit and tariff; clean low-value claims	Shadow agreement at target; bias checks passed; appeal route live	Share decided inside one hour; share decided by rule; override and complaint rates stable or falling
3. Predict	Month 18 onwards	Predictive models for underwriting, chronic cost and provider risk, used as Tier 1 recommendations	Two years of clean, coded experience data; model risk policy	Model performance documented; outcomes across groups monitored; no Tier 3 decisions automated

The sequence matters. Organisations that start at phase 3 usually discover in phase 3 that they needed phase 0.

7. Risks and their controls

RISK	HOW IT SHOWS UP	PRIMARY CONTROL
Wrongful denial	A covered member is refused care or payment	Never deny by rule (5.1); human review with reasons (5.4)
Automation bias	Reviewers approve whatever the system suggests	Measured override rates; reviewer time protected (5.4)
Disparate outcomes	Rural providers, older members or particular states fare worse	Pre-deployment and ongoing group comparisons (5.5)
Data drift	Tariff changes or new fraud patterns degrade accuracy	Weekly monitoring; owner with authority to retire (5.5)
Hallucination	Generated text states wrong benefits or reasons	Grounding; human approval of outbound communication (5.6)
Data leakage	Member data processed by an unassessed foreign service	DPIA; contractual limits; lawful transfer basis; de-identification (5.6)
Fraud at speed	Automated payment processes fraudulent claims faster	Fraud detection as Tier 1 on all automated payment; value limits on Tier 2 (5.1)
Vendor opacity	The insurer cannot explain its vendor's decisions	Contractual right to decision records and model documentation (8)
Regulatory change	The NITDA Code or NHIA guidelines impose new duties	Build to the draft Code now; maintain the instrument map (4)

8. For boards and investors: a diligence checklist

The following questions distinguish an insurer, or an insurtech vendor, whose AI is governed from one whose AI is merely claimed.

1. Is there a written, board-approved statement of which decisions automated systems may take, and which they may not?
2. Can any automated decision from the last month be retrieved, with its inputs, rule or model version, reason and response time?
3. Are any adverse decisions, denials or reductions, made without a human?
4. What share of authorisation requests are answered within one hour, measured from the provider's submission?

5. Has each AI system had a data protection impact assessment, and was it filed with the annual compliance return?
6. How are outcomes compared across states, age bands, sex and plan tiers, and when was this last done?
7. What is the human override rate on automated recommendations, and what is done when it changes?
8. Where is member data processed, by which third parties, and on what lawful basis for any transfer outside Nigeria?
9. Who can switch each system off, and what happens to the work when they do?
10. For vendors: does the contract give the insurer access to decision records and model documentation sufficient to answer a regulator?

An insurer that can answer all ten has auditable AI. One that cannot has a liability that will grow with its automation.

9. Conclusion

Nigerian health insurance is entering a period in which automation is not optional. Enrolment growth, the tariff reset, the one-hour authorisation rule and private provider price volatility together make manual operations too slow and too expensive. The question is not whether HMOs will use AI but whether the AI they use can be trusted by members, providers, regulators and investors.

The legal instruments to govern it already exist. What has been missing is their translation into operating practice: a decision-rights matrix, a data foundation, a decision record, real human oversight and disciplined testing. None of these is technically difficult. All of them require deciding, in advance and in writing, what the machine is for and what it is not.

An HMO that makes those decisions will find that auditability pays for itself: in faster decisions on routine care, fewer disputes, better fraud control and a credible answer to every regulator's question. One that does not will find its automation doing the opposite.

Annex A: Comparative note on Kenya, Ghana and South Africa

Nigeria's pattern, a social health insurance statute alongside a GDPR-style data protection law with an automated-decision clause, is shared by several African markets. The frameworks in this paper transfer with modest adjustment.

	KENYA	GHANA	SOUTH AFRICA
Health insurance structure	Social Health Authority under the Social Health Insurance Act 2023, replacing NHIF; private insurers alongside	National Health Insurance Authority under the National Health Insurance Act; National Health Insurance Scheme	Medical schemes regulated by the Council for Medical Schemes; private insurers
Data protection law	Data Protection Act 2019	Data Protection Act 2012 (Act 843)	Protection of Personal Information Act 2013 (POPIA)
Automated decision clause	s.35: right not to be subject to a decision based solely on automated processing, including profiling, with legal or significant effects	s.41: right to require that significant decisions are not based solely on automatic processing	s.71: no decision with legal consequences or substantial effect based solely on automated processing, subject to exceptions; where a contract exception applies, an opportunity to make representations and information about the underlying logic
National AI policy	National AI Strategy 2025 to 2030, launched March 2025; Insurance Regulatory Authority sandbox (BimaLab)	National AI Strategy 2023 to 2033 adopted	Policy work under way; no AI statute
Notable context	About 32 million registered with SHA by mid-2026, but only about 5 million actively contributing; Ksh 11 billion in fraudulent SHA claims reported to members of Parliament for October 2024 to April 2025	About 18.4 million active NHIS members (56 per cent) in 2025; 99.9 per cent of claims submitted and vetted digitally; one-time attendance codes piloted from 2026; biometric verification and analytics to flag duplicate claims	Most of the continent's current generative AI innovation in insurance, per McKinsey

Kenya offers the clearest warning. Rapid digitisation of a national scheme, with automated fraud flags but without matching controls on provider onboarding, enforcement and human follow-up, allowed large-scale fraud to be processed efficiently. The fraud types reported, fictitious procedures, implausible dependants and procedures at facilities lacking the means to perform them, are all detectable with the Tier 1 anomaly checks described in section 5.1, provided someone accountable acts on the flags.

Ghana offers a constructive example. Submitting and vetting nearly all claims digitally, piloting a unique one-time attendance code for each hospital visit, verifying identity biometrically in high-volume facilities, and using analytics across regions to flag duplicate claims are the kind of foundation controls described in sections 5.2 and 5.3. They precede, rather than follow, more advanced automation.

South Africa's POPIA section 71 is the most demanding of the four automated-decision clauses on one point: where a contract exception is relied on, the responsible party must give the individual sufficient information about the underlying logic of the processing to make representations. Nigerian insurers building to the reason-code standard in section 5.3 will meet that bar as well.

A full chapter on any of these markets would require local legal review. This annex is a comparative orientation, not a compliance assessment.

Annex B: Glossary

Adjudication. Deciding whether a claim is payable and for how much.

Automation bias. The tendency of people reviewing automated outputs to accept them without sufficient scrutiny.

CAR. Compliance Audit Return, filed annually with the Nigeria Data Protection Commission by qualifying organisations.

DPIA. Data protection impact assessment.

GAID. General Application and Implementation Directive, issued by the NDPC in 2025.

Generative AI. AI that produces new text, images or other content, such as large language models.

Incurred basis. Counting claims by the date the service was provided rather than the date they were paid, so that submitted but unpaid claims are included.

Loss ratio. The share of premium income paid out in claims.

NDPA. Nigeria Data Protection Act 2023.

NHIA. National Health Insurance Authority.

Override rate. The share of automated recommendations that a human reviewer changes.

Shadow mode. Running a system on live cases without acting on its outputs, to compare its decisions with people's.

References

1. National Health Insurance Authority Act 2022; KPMG Nigeria, commentary on the National Health Insurance Authority Act.
2. Channels TV, "Approve treatment within one hour, NHIA directs HMOs", 7 April 2025.
3. 21st Century Chronicle, "NHIA sanctions 96 health providers, HMOs", 30 June 2025.
4. The Guardian (Nigeria), "NHIA enrolment hits 23m as cancer, HIV/TB coverage expands", 20 September 2026.
5. The Guardian (Nigeria), "Health insurance: how NHIA is serving and protecting Nigerians", August 2026 (16.2 million in Q4 2023; 22,025,342 in Q1 2026).
6. Nairametrics, "Rising drug prices force sharp increases in health insurance premiums across Nigeria", 14 November 2025 (premium rises of 8 to 59 per cent by tier).
7. Nigeria Data Protection Act 2023; KPMG Nigeria, "The Nigeria Data Protection Act 2023".

8. Nigeria Data Protection Commission, General Application and Implementation Directive 2025, articles 18 and 43; G Elias, "GAID 2025 unpacked (series II)", Mondaq, November 2025.
9. Federal Ministry of Communications, Innovation and Digital Economy, National Artificial Intelligence Strategy, September 2025; OECD.AI policy navigator entry.
10. NITDA, Draft Code of Practice for Artificial Intelligence, 2025, as summarised by Regulations.ai (checked September 2026); primary draft text not reviewed.
11. News Agency of Nigeria, "FG strengthens digital infrastructure for AI in health", 13 November 2025.
12. McKinsey & Company, "Leading, not lagging: Africa's gen AI opportunity", May 2025.
13. The Star (Kenya), "InsurTech Forum Nairobi launches Execution Advantage AI challenge", 22 June 2026.
14. Blueprint, "NHIA reluctant as Nigerians groan over near-death experiences with HMOs", 2026.
15. Citizen Digital, "Revealed: how Kenyans lost Ksh 11 billion in fraudulent SHA claims", 28 January 2026; Kenyans.co.ke, "CS Duale reveals gaps leading to missing Ksh 11B SHA funds", 2026.
16. Kenya Data Protection Act 2019, section 35; Kenya Social Health Insurance Act 2023.
17. Capital FM, "Only 1 in 6 SHA members actively contributing despite 32m registration", 2026; Bowmans, "Kenya: unveiling of the National AI Strategy 2025 to 2030", 2025.
18. Ghana Data Protection Act 2012 (Act 843), section 41; Regulations.ai, Ghana AI regulation overview.
19. National Health Insurance Authority (Ghana), "NHIS registers over 18 million active members", 2025; MyJoyOnline, "NHIA begins 50% NHIS tariff increase, introduces digital claims verification", August 2026.
20. Protection of Personal Information Act 2013 (South Africa), section 71; Bowmans, "South Africa: AI in the workplace (2 of 6): POPIA considerations", August 2026.

About the author

Moyo Awoyokun is a physician, health insurance specialist and health IT analyst working across the United Kingdom and Nigeria. He holds an MB ChB, an MBA from Warwick Business School, a Fellowship of the Institute of Health Insurance and Managed Care of Nigeria and Membership of the Chartered Insurance Institute (UK). He has held operational and board-level roles in Nigerian health maintenance organisations, has led electronic health record build work in NHS programmes, and is Product Visionary at Ajé Intelligent InsurTech, where he is responsible for data protection and AI governance.

THIS PAPER IS A PRACTITIONER'S SYNTHESIS OF PUBLIC SOURCES AND DOES NOT CONSTITUTE LEGAL, ACTUARIAL OR INVESTMENT ADVICE. VIEWS ARE THE AUTHOR'S OWN. REPRODUCTION WITH ATTRIBUTION IS WELCOMED.