Technology Stewardship

INSTITUTE • GOVERNING AUTONOMOUS TECHNOLOGY

INSTITUTE FOR TECHNOLOGY STEWARDSHIP

ISSUES PAPER No. 1

# The Unpriced Risk

## Telecom AI Governance and the Data Sovereignty Gap in the Global South

July 2026

**Gaurav Arora**

Telecom and critical-infrastructure practitioner

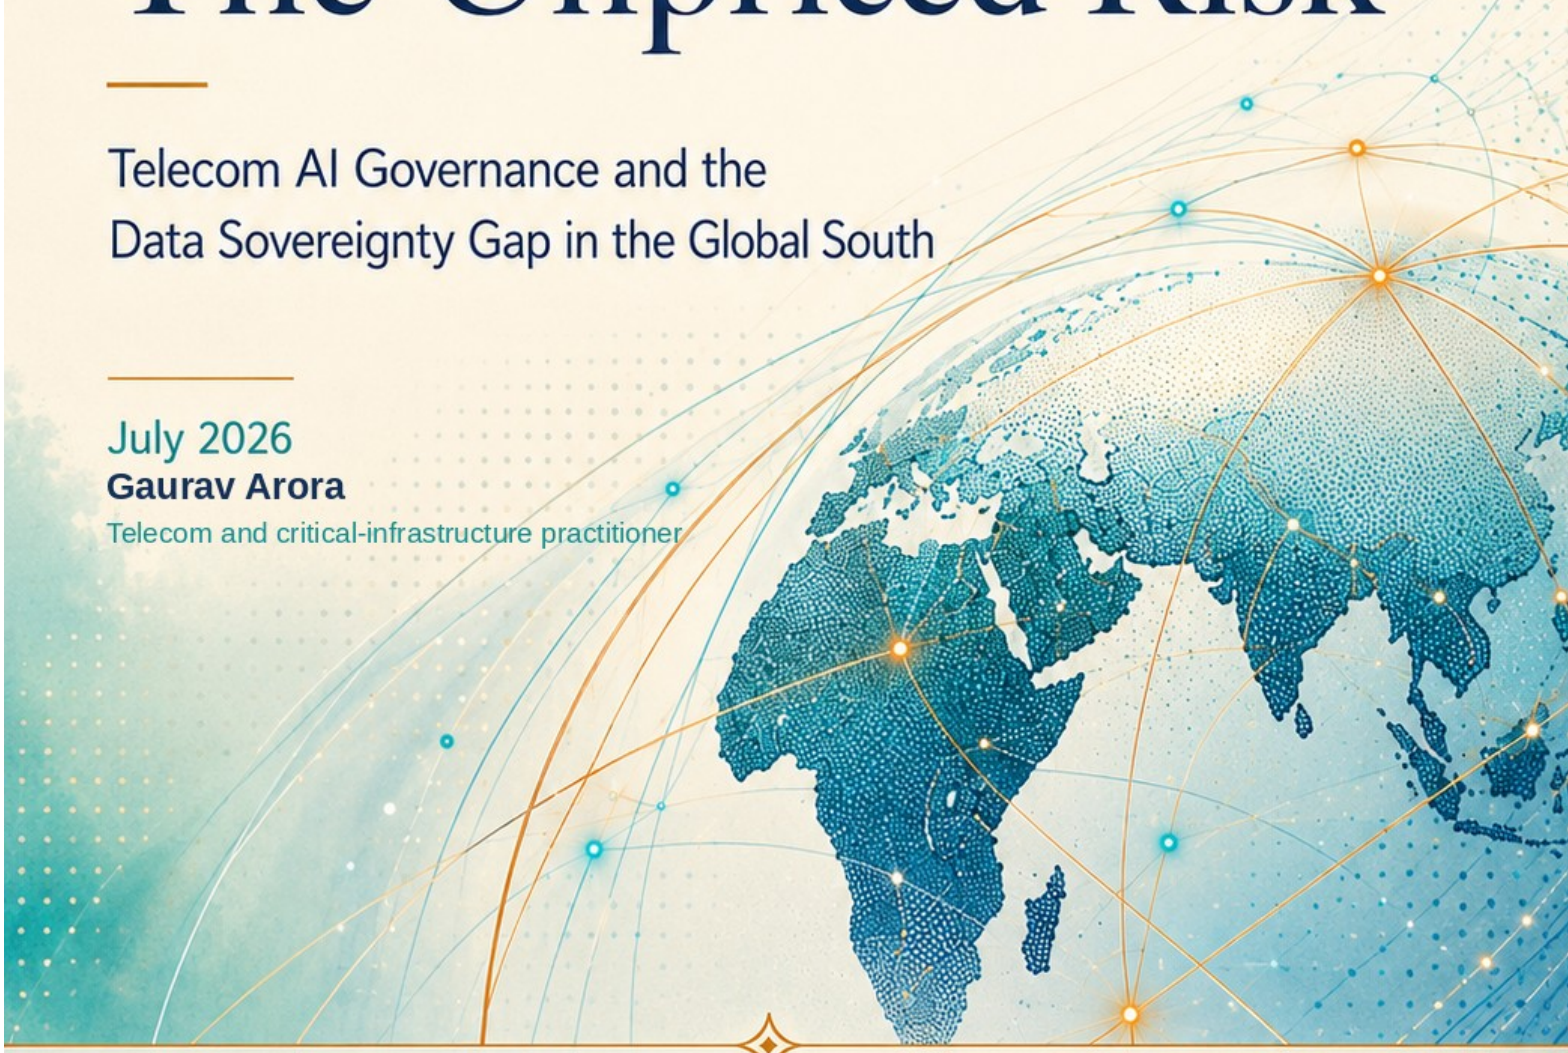

*Ensuring technology remains in service of humanity — not the other way around.*

## Executive Summary

*Author: Gaurav Arora | Institute for Technology Stewardship | Research cut-off: 18 July 2026*

**Keywords:** telecom AI governance; operational sovereignty; subscriber data; Global South; critical infrastructure

Telecommunications networks generate a uniquely revealing combination of communications metadata, location information, billing records and, in many markets, mobile financial data. Artificial intelligence is increasingly applied to this data for churn prediction, fraud control, credit scoring, personalised offers, service eligibility and network optimisation. The operational benefits can be substantial. The governance consequences can also be substantial because these systems can affect a person's connectivity, economic opportunity and access to services.[1, 2, 3]

The central risk is not the use of a third-party model by itself. The risk arises when an operator cannot demonstrate what data built and now feeds the system, where and under whose jurisdiction processing occurs, and how consequential decisions can be reconstructed, explained, challenged and reversed. In that situation, the operator may have acquired an AI capability without retaining operational sovereignty over it.

European law illustrates why precise classification matters. The EU AI Act does not make every telecom AI system high-risk. Annex III point 2 is limited to AI used as a safety component in the management and operation of critical digital infrastructure. The Digital Omnibus text approved by the European Parliament and the Council, but still pending publication in the Official Journal as of 18 July 2026, further clarifies the meaning of a safety function. Annex III point 5 directly reaches specified uses such as creditworthiness assessment and emergency-call classification. Other telecom systems may remain subject to the GDPR, ePrivacy, consumer-protection, communications and non-discrimination requirements. A system-by-system assessment is therefore essential.[4, 5, 26, 27]

Selected jurisdictions across the Global South are not regulatory vacuums. India has enacted a data-protection regime with phased implementation; Saudi Arabia combines binding personal-data rules with AI ethics and assessment instruments; ASEAN has developed regional governance guides and a responsible-AI roadmap; and the African Union has adopted a Continental AI Strategy. What remains uneven is binding, sector-specific assurance for consequential telecom AI and a common evidence model that operators can use across jurisdictions.[10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20]

This paper proposes three operational sovereignty questions for telecom AI. It then translates them into a minimum assurance baseline for regulators, operators, vendors and standards bodies. The objective is not to require domestic ownership of every model or to reject cloud services. It is to ensure that the operator remains able to evidence control, accountability and lawful processing throughout the AI supply chain.

| KEY FINDINGS | |
|---|---|
| **1** | Telecom data is high-impact data infrastructure: communications, movement and financial signals become materially more revealing when combined. |
| **2** | AI Act exposure is use-case specific. Telecom status alone does not make an AI system high-risk, but credit scoring, emergency-call functions and genuine safety components may fall directly within Annex III. |
| **3** | Operational sovereignty is broader than data localisation. It includes provenance, jurisdictional control, decision traceability, human intervention and credible exit rights. |
| **4** | The lowest-cost point to build governance is before AI systems become deeply embedded in subscriber and network operations. |
| **5** | A common telecom assurance baseline can improve regulatory readiness without requiring every jurisdiction to adopt identical legislation. |

## Methodology and Scope

This is a desk-based issues paper. It reviews publicly available legal texts, regulatory guidance, government strategies, company announcements and industry materials available up to 18 July 2026. The jurisdictional discussion is selective

rather than exhaustive and is intended to identify governance patterns, not to provide legal advice for any particular deployment. Corporate announcements are treated as market signals, not as independently audited evidence of compliance or technical performance.

The term Global South is used as a practical shorthand for a diverse set of countries with different legal systems, institutional capacities and market structures. The paper does not claim that these jurisdictions share a single regulatory position. It also does not claim that most telecom AI deployments are opaque, that third-party models are inherently ungovernable, or that local hosting alone creates sovereignty. Where evidence is limited, the paper frames the issue as a governance question rather than a universal factual claim.

**What this paper contributes**

A narrow operational proposition: a telecom operator should not claim sovereignty over an AI-enabled capability unless it can evidence the system's data provenance, processing jurisdiction and consequential decision record. The three questions introduced in Section 6 provide a minimum diagnostic, not a complete regulatory framework.

# 1. Telecom Data Is High-Impact Data Infrastructure

Telecom data is often described as metadata, operational data or customer data. Those labels can understate what the data reveals when aggregated. The Office of the UN High Commissioner for Human Rights has observed that communications metadata can support precise conclusions about behaviour, relationships, preferences and identity. Location histories can reveal places of residence and work, recurring routines, social associations and visits to sensitive locations. The governance issue is therefore not only the sensitivity of individual fields, but the inferential power created when multiple categories are linked.[1, 2]

## 1.1 Communications metadata

Call detail records and related network events can identify communicating parties, time, duration, routing and network elements. They do not ordinarily contain call content, but at scale they can map social and organisational relationships. AI systems can turn those records into behavioural segments, anomaly scores, churn predictions and relationship graphs. The resulting inferences may be more consequential than the original records because they convert historical activity into predictions about future behaviour.

## 1.2 Location and mobility data

Mobile networks continuously generate location-related signals through cell registration, handovers, roaming and service use. Accuracy varies by technology and context, but repeated observations can reveal patterns of movement. Location data should therefore be treated as a high-impact input, especially where an AI system links it with identity, communications or financial records.[2]

## 1.3 Billing, service and mobile financial data

Recharge behaviour, payment history, device type, service tier and consumption patterns can be used to infer income stability, price sensitivity and creditworthiness. Where operators or affiliates provide mobile money, lending or insurance services, the boundary between telecommunications data and financial decisioning becomes especially important. The GSMA has documented the use of telecom data to build predictive credit scorecards for mobile financial products.[3]

**Illustrative case: Kenya's mobile-money ecosystem**

Kenya illustrates why the boundary between telecom data and financial decisioning matters. The M-Shwari service historically used Safaricom-derived variables relating to airtime, airtime credit, M-Pesa usage and the length of the customer relationship to support credit scoring. This example is not presented as an allegation of regulatory failure, and the historical scorecard should not automatically be treated as AI under every legal definition. It demonstrates how data generated through telecommunications and mobile-money relationships can become an input to consequential financial decisions in a widely used national service. As machine-learning scoring becomes more sophisticated, provenance, contestability and decision evidence become correspondingly more important.[29]

### 1.4 The combined data layer

Individually, these data categories require protection. Combined, they create a high-resolution account of relationships, movement and economic behaviour. This paper uses the phrase surveillance-grade data as a descriptive governance term, not as a legal classification. It refers to data whose aggregation can support continuous observation and consequential inference even where the original records were collected for routine network or billing purposes.

## 2. What AI Is Doing with Telecom Data

Telecom AI is not one technology category. Predictive models, optimisation systems, generative models and autonomous agents create different risks and should not be governed as though they were interchangeable. A useful starting taxonomy is shown below.

| Use case | Typical inputs | Potential consequence | Primary governance concern |
|---|---|---|---|
| **Churn and retention** | Usage, recharge and service history | Targeted offers, contact or service changes | Profiling, fairness, purpose limitation and transparency |
| **Fraud and abuse controls** | CDRs, device, identity and transaction signals | Restriction, suspension or investigation | False positives, service continuity, contestability and incident evidence |
| **Creditworthiness and lending** | Recharge, device, usage and mobile-money data | Access to loans, limits or pricing | High-impact decisioning, discrimination, explainability and redress |
| **Dynamic pricing and eligibility** | Usage, location and inferred willingness to pay | Different prices, packages or subsidies | Consumer fairness, discrimination and transparency |
| **Network optimisation** | Traffic, congestion, topology and location density | Capacity allocation, routing and quality of service | Geographic inequity, resilience and emergency-service impact |
| **Generative and agentic support** | Logs, manuals, topology and operator prompts | Recommendations, scripts or autonomous actions | Data leakage, tool access, provenance, permissions and action accountability |

### 2.1 Predictive and decisioning systems

Churn, fraud, pricing and credit models generally use structured historical data. Their governance challenge is not model novelty but decision consequence. A modest statistical model can have a significant effect if it determines whether a person receives a loan, loses service or pays a different price. Risk classification should therefore focus on intended purpose, affected persons and practical consequence rather than model size alone.

> **Illustrative consequence.** Consider a prepaid subscriber whose number is automatically suspended after a false fraud alert. If that number is also used for mobile-money authentication, loan repayments, government-service access and family remittances, the practical consequence extends beyond loss of telephone service. It may become temporary exclusion from essential economic and civic functions. The scenario is hypothetical, but it shows why governance intensity should follow the consequence of a decision, not merely the sophistication of the model.

### 2.2 Network optimisation systems

Network AI may appear operationally neutral because it acts on capacity, routing and energy use rather than directly on a named subscriber. Its effects can nevertheless be distributed unevenly. Optimisation objectives can influence which areas receive investment, which traffic is deprioritised during congestion and whether emergency or essential

communications remain available. These systems require resilience and public-interest analysis even where they fall outside a statutory high-risk category.

### 2.3 Generative AI and autonomous agents

Large language models and autonomous agents add a different supply chain. They may access external tools, retrieve documents, generate executable configurations or send data to remote inference services. Their governance therefore requires prompt and tool logging, permission controls, retrieval-source provenance and action records. These controls should not be confused with demands to disclose a model's hidden chain of thought, which may be unavailable or unreliable as an audit artifact.

## 3. Where Operational Sovereignty Can Disappear

Vendor participation is normal in telecommunications. Operators have always depended on network equipment providers, software suppliers, cloud platforms and managed-service partners. The governance problem begins when responsibility remains with the operator but the evidence needed to discharge that responsibility is distributed across contracts, platforms and jurisdictions that the operator cannot interrogate.

Four evidence gaps are especially important:

- Data supply chain: the operator cannot identify which data categories were used for training, fine-tuning, retrieval or ongoing inference, or whether customer data is retained for provider improvement.
- Model and software supply chain: the operator cannot identify the foundation model, material updates, embedded third-party components, evaluation results or known limitations relevant to the use case.
- Processing and jurisdiction chain: the operator lacks a verified map of inference, storage, logs, backups, subprocessors, remote administration and government-access exposure.
- Decision and accountability chain: the operator cannot reconstruct which system version acted, what information it relied on, what action was proposed or executed, who intervened and how an affected person can obtain review.

> **Operational sovereignty**
>
> For this paper, operational sovereignty means the demonstrable ability to understand, govern, intervene in and, where necessary, replace an AI-enabled capability without losing accountability for data, decisions or critical service continuity. It does not require domestic ownership of every model or physical server.

## 4. The Regulatory Landscape

### 4.1 European Union: classification is system-specific

The EU AI Act is relevant to telecom, but its application must be analysed by intended purpose. Annex III point 2 covers AI systems intended to be used as safety components in the management and operation of critical digital infrastructure. The Commission's draft classification guidelines, published on 19 May 2026, indicate that ordinary service-quality and operational optimisation may fall outside this category unless the system performs the required safety function.[4, 5]

The Digital Omnibus text approved by the European Parliament and the Council, but still pending publication in the Official Journal as of 18 July 2026, sharpens this boundary. It specifies that a component fulfils a safety function where its intended purpose is to prevent or mitigate risks to the health and safety of persons or property. The accompanying recitals distinguish this from systems intended solely for user assistance, performance optimisation, service efficiency, automation, convenience or non-safety quality control. The clarification reinforces the need to assess the actual intended function of each telecom AI system rather than treating every network-management system as high-risk.[26, 27]

Annex III point 5 directly covers specified decisioning uses, including evaluation of creditworthiness or establishment of a credit score, subject to the Act's exclusions, and AI used to evaluate and classify emergency calls or dispatch emergency response services. It also covers certain public-authority decisions concerning essential public assistance.

These provisions do not automatically classify private telecom churn, dynamic pricing or routine fraud controls as high-risk. Those systems require assessment under their actual purpose and other applicable law.[4]

Legal role also matters. High-risk provider obligations, including conformity assessment, are not identical to deployer obligations. An operator may be a deployer of a vendor system, a provider of an internally developed system, or may become a provider where it places the system under its own name, substantially modifies it or changes its intended purpose in circumstances covered by the Act. Governance inventories must therefore record both system classification and economic-actor role.[4]

The Council of the European Union approved the European Parliament's first-reading position and adopted the legislative act on 29 June 2026. The act was signed on 8 July but, as of 18 July 2026, remained pending publication in the Official Journal. It will enter into force on the third day following publication. The agreed application dates are 2 December 2027 for stand-alone Annex III high-risk systems and 2 August 2028 for high-risk systems embedded in regulated products. The Commission's draft high-risk classification guidelines remain under consultation until 23 July 2026.[5, 6, 26, 27]

The postponement of the high-risk-system requirements does not generally delay Article 50 transparency obligations. Where their statutory conditions are met, relevant duties concerning direct interaction with AI systems, machine-readable marking of synthetic content, deepfake disclosure and certain AI-generated public-interest text apply from 2 August 2026. The Omnibus provides a targeted transition until 2 December 2026 for providers bringing pre-existing systems into conformity with the marking and detection obligation in Article 50(2); it does not move all Article 50 duties to 2027.[4, 26, 28]

## 4.2 GDPR, ePrivacy and meaningful information

The GDPR remains relevant whether or not a telecom AI system is classified as high-risk under the AI Act. Article 22 applies where a decision is based solely on automated processing and produces legal effects or similarly significant effects. Articles 13, 14 and 15 create transparency rights concerning automated decision-making where their conditions are met, while the ePrivacy framework applies sector-specific protections to electronic communications data.[7]

In February 2025, the Court of Justice of the European Union clarified that meaningful information about automated decision-making should explain, in a concise and understandable form, the procedure and principles actually applied to the person's data so that the result can be understood and challenged. This does not necessarily require disclosure of source code or a complete algorithm, but it does require more than a generic statement that a model was used.[8]

The European Data Protection Board's Opinion 28/2024 also cautions against assuming that an AI model is anonymous merely because personal data cannot be directly extracted through a simple query. Anonymity must be assessed case by case, taking account of the likelihood of identification or extraction and the measures applied throughout development and deployment.[9]

## 4.3 Selected Global South jurisdictions: progress without a common telecom assurance model

The selected jurisdictions below demonstrate why the issue should not be framed as Europe regulated and the Global South unregulated. Data-protection, ethics and AI-strategy instruments exist, but their legal status, sectoral reach and assurance mechanisms differ substantially.

| Jurisdiction / region | Existing instruments | Governance implication for telecom AI |
|---|---|---|
| **India** | Digital Personal Data Protection Act 2023; Digital Personal Data Protection Rules 2025 with phased commencement. | The cited instruments do not create an AI-specific telecom risk classification, conformity assessment or general right to an algorithmic explanation. Sectoral and consumer rules may still apply. |

| Jurisdiction / region | Existing instruments | Governance implication for telecom AI |
|---|---|---|
| **Saudi Arabia** | Personal Data Protection Law and implementing rules; cross-border transfer regulation; AI Ethics Principles, self-assessment and adoption instruments. | Saudi Arabia has binding personal-data and cross-border-transfer requirements alongside national AI ethics, assessment and adoption instruments. Its current architecture does not include a dedicated binding conformity-assessment mechanism for consequential telecom AI comparable to the EU AI Act. |
| **ASEAN** | ASEAN Guide on AI Governance and Ethics, Expanded Guide for Generative AI and Responsible AI Roadmap. | The regional instruments promote interoperability and responsible adoption but are primarily voluntary; binding obligations remain with member-state laws and sector regulators. |
| **African Union** | Continental AI Strategy and wider digital-governance initiatives, alongside diverse national data-protection and emerging AI policies. | There is no single binding continental telecom AI regime. National maturity and sectoral authority vary considerably, making regional capacity and evidence interoperability important. |

Sources for the table are the relevant official instruments and institutional publications.[10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20]

## 5. The Cost of Building Governance Late

Telecom has already experienced the operational cost of addressing strategic dependency after infrastructure has been deployed. In 2020, the United Kingdom required Huawei equipment to be removed from 5G networks by the end of 2027. The government estimated that the combined decisions could delay rollout by two to three years and create costs of up to GBP 2 billion. The exact policy and security context is different from AI, but the operational lesson is relevant: hidden dependency becomes expensive when the organisation must change architecture under an external deadline.[21]

AI is not a direct repetition of the 5G equipment case. Models can be replaced or reconfigured more rapidly than physical radio infrastructure in some circumstances. In other cases, however, model interfaces, data pipelines, vendor-specific agents, decision logic and workforce processes can become deeply embedded. The useful analogy is therefore not that every foreign AI provider will be removed. It is that operators should understand concentration, portability, evidence ownership and exit constraints before a regulatory or security event forces them to do so.

> ***The cost of governance debt does not disappear. It accumulates until regulation, litigation or operational failure makes it payable - usually when the operator has less time, less leverage and fewer choices.***

## 6. The Three Operational Sovereignty Questions

The following questions are intended as a minimum diagnostic. An operator may choose a cloud service, an external model or a managed platform and still answer all three. Conversely, domestic hosting can still fail the diagnostic if provenance, access and decision evidence remain opaque.

| | | |
|---|---|---|
| **1** | **What data built and now feeds this system?** | **Evidence should cover foundation-model provenance where available; training and fine-tuning data categories; retrieval sources; inference inputs; whether operator or subscriber data is retained for provider improvement; legal basis; retention; and material model updates.** |
| **2** | **Where, by whom and under whose jurisdiction is processing occurring?** | Evidence should include the locations of inference, storage, logs and backups; subprocessors; remote-administration paths; encryption and key ownership; cross-border transfer mechanisms; applicable legal-access regimes; service-continuity arrangements; and practical exit or substitution rights. |
| **3** | **Can the operator reconstruct and contest every consequential decision?** | The operator should be able to identify the system and model version, relevant inputs, retrieved sources, tools invoked, output or proposed action, policy checks, human intervention, executed action, outcome, rollback and review pathway. The objective is a reliable decision-event record, not disclosure of hidden chain of thought. |

### Sovereign AI is not synonymous with data localisation

Local inference can reduce some transfer and latency risks, but it does not by itself establish model provenance, lawful data use, effective human oversight, evidence retention or vendor substitutability. Operational sovereignty is an evidence and control condition, not a location label.

## 7. What Early Market Signals Show

Several operators are investing in infrastructure and partnerships explicitly associated with sovereign AI. These initiatives should not be treated as proof that every workload meets a particular legal standard. They do, however, show that control over infrastructure, data location and model choice is becoming a commercial proposition rather than only a compliance concern.

Deutsche Telekom launched its Industrial AI Cloud in February 2026. The company describes the Munich facility as sovereign computing infrastructure for Germany and Europe and reports 10,000 NVIDIA GPUs, with data remaining in Germany under German and European security standards.[22]

Singtel launched RE:AI in October 2024 and now describes it as a sovereign AI cloud business. In April 2026, RE:AI and Mistral AI announced a partnership to integrate open-weight models and AI infrastructure, while Singtel states that customers can retain control over where data is stored and processed.[23, 24]

SK Telecom used Mobile World Congress 2026 to present a Sovereign AI Package with telecom partners including Singtel, e& and NTT. It also presented A.X K1, a 519-billion-parameter foundation model developed within South Korea's sovereign AI programme. The relevant signal is the combination of telecom infrastructure, domestic or controlled model capability and cross-operator collaboration.[25]

The wider implication is not that every operator should build a 500-billion-parameter model or a national GPU cloud. It is that operators increasingly see infrastructure control, model choice and jurisdictional assurance as differentiators. Smaller operators can pursue the same governance objective through contract design, regional shared infrastructure, open-weight models, confidential computing, independently verifiable logging and credible exit plans.

## 8. A Minimum Telecom AI Assurance Baseline

The three sovereignty questions become useful when translated into repeatable evidence. The following baseline is deliberately technology-neutral and can be implemented before a jurisdiction adopts a dedicated telecom AI law.

| Control | Minimum evidence |
|---|---|
| **AI system inventory and role classification** | System owner, intended purpose, affected network domain or subscriber process, model/provider, operator role, legal risk classification and material dependencies. |
| **Data and model manifest** | Data categories, provenance, legal basis, training/fine-tuning/retrieval/inference use, retention, update process and known limitations. |
| **Processing and jurisdiction map** | Inference, storage, logging, backup, subprocessors, remote access, transfer safeguards, key control and service-continuity dependencies. |
| **Decision-event record** | Model/version, relevant inputs, retrieved sources, output/action, policy checks, human intervention, outcome, rollback and review status. |
| **Human review and subscriber redress** | Named accountable role, intervention authority, escalation conditions, understandable explanation process and correction or appeal route. |
| **Vendor assurance and exit controls** | Audit rights, incident notification, model-change notice, evidence portability, deletion/return obligations, substitution plan and continuity testing. |
| **Independent testing and monitoring** | Use-case-specific performance, subgroup and geographic impacts where relevant, drift, security, resilience, false-positive consequences and periodic review. |

### 8.1 Priorities for national telecom regulators

- Require operators to maintain an AI inventory covering subscriber-facing, network-operational, embedded vendor and agentic systems, with material changes reported through an appropriate risk-based process.
- Require enhanced impact assessment where telecom AI can restrict service, determine creditworthiness, influence essential connectivity, materially change pricing or affect emergency communications.
- Specify minimum records for data provenance, processing location, model changes, consequential decisions, human intervention and incident investigation.
- Require vendor contracts to preserve evidence access, incident notification, auditability and practical exit rights without mandating a single hosting or ownership model.
- Establish subscriber review and redress expectations proportionate to the seriousness of the decision.

### 8.2 Priorities for operators

- Begin with systems that can deny, restrict, price or materially influence a subscriber service, followed by network systems whose failure could affect emergency or essential communications.
- Classify the operator's legal and operational role for each system rather than assuming the vendor carries all provider responsibility.
- Test whether decision evidence can actually be produced during an incident or complaint, not merely whether a logging clause exists in a contract.
- Treat portability and evidence ownership as resilience controls. A technically accurate model can still create strategic risk if the operator cannot migrate or investigate it.

### 8.3 Priorities for standards bodies and international organisations

- Develop a telecom AI system-card and event-record profile that can map to NIST, ISO/IEC, ITU, regional data-protection and future AI-law requirements.
- Define common terminology for model provenance, inference jurisdiction, decision traceability, human intervention and material model change in telecom environments.
- Support regulatory capacity-building and shared testing infrastructure so that smaller and emerging-market operators can demonstrate assurance without constructing national-scale AI infrastructure.
- Create mechanisms for cross-border learning from telecom incidents without requiring disclosure of subscriber data or commercially sensitive network details.

## 9. Research Agenda and Call for Input

This paper identifies a minimum governance condition rather than resolving every policy question. The next stage of research should test the framework against real deployments, regulatory practice and the experience of affected subscribers. Priority questions include:

1. Which telecom AI decisions should be treated as consequential even where they do not meet the legal threshold for solely automated decisions or statutory high-risk classification?
2. What is the minimum event record that supports investigation, explanation and redress without exposing security-sensitive network information or demanding unreliable chain-of-thought disclosure?
3. How should operators measure geographic and socioeconomic distributional effects from network optimisation and personalised pricing?
4. Which vendor disclosures can be standardised without requiring disclosure of protected intellectual property?
5. How can regional or shared sovereign infrastructure serve smaller operators while maintaining competition and avoiding a new form of concentration risk?
6. What governance mechanisms are necessary where telecom data is used jointly for connectivity, identity, mobile money, insurance and public-service access?

The Institute for Technology Stewardship invites responses from telecom operators, regulators, vendors, standards bodies, civil-society organisations and researchers. ITS is particularly interested in operational case studies, counterexamples to the proposed diagnostic, evidence-field recommendations and perspectives from emerging economies. Substantive responses will inform the future development of the framework and future work on AI governance in telecommunications and critical infrastructure.

## 10. Conclusion

Telecom AI governance begins with a simple but demanding proposition: responsibility cannot be outsourced merely because the model, platform or inference service is outsourced. The operator remains the institution through which subscriber data is collected, services are delivered and operational decisions reach the public.

The EU AI Act demonstrates that legal classification is narrower and more use-case specific than broad critical-infrastructure rhetoric suggests. The selected Global South jurisdictions demonstrate that governance is advancing through data laws, ethics instruments, strategies and regional guidance, but without a common telecom assurance model. This creates a period in which operators can either build evidence and control deliberately or allow dependency to deepen until an external event makes remediation urgent.

The three questions in this paper provide a practical starting point. What data built and feeds the system? Where and under whose jurisdiction is processing occurring? Can every consequential decision be reconstructed and contested? An operator that can answer these questions with verifiable evidence has begun to establish operational sovereignty. An operator that cannot has purchased capability without securing accountability.

## About the Author

Gaurav Arora is the founder of the Institute for Technology Stewardship and a telecom and critical-infrastructure practitioner with more than two decades of experience in network operations, service delivery, technology programme leadership and cross-functional global teams. His research focuses on translating operational practices from telecommunications into evidence-based governance for AI and autonomous systems.

## About the Institute for Technology Stewardship

The Institute for Technology Stewardship (ITS) is an independent research initiative focused on the governance of autonomous and emerging technology systems. ITS bridges operational expertise with policy research, developing practical frameworks intended to keep technology deployment in critical infrastructure safe, transparent, accountable and human-centred.

ITS works from the premise that institutions with long experience of human-machine collaboration at scale - including telecom network operations, air traffic control and other critical infrastructure - hold lessons that should be more fully represented in AI governance and international policy discussions.

### Declarations

**Funding:** No external funding was received. **Competing interests:** The author declares no competing interests. **Data availability:** No original dataset was generated; the analysis is based on publicly available sources. **Ethical approval:** Not applicable.

### Suggested Citation

Arora, Gaurav. 2026. The Unpriced Risk: Telecom AI Governance and the Data Sovereignty Gap in the Global South. Institute for Technology Stewardship, Issues Paper No. 1, July 2026.

### Disclaimer

This issues paper provides independent policy analysis and does not constitute legal advice. It does not represent the position of any operator, regulator, standards body, company or government. References to organisations and initiatives are included for analytical purposes and do not imply endorsement or criticism. Legal and regulatory status should be verified for the jurisdiction and date of any actual deployment.

# References


1. **Office of the United Nations High Commissioner for Human Rights.** *The Right to Privacy in the Digital Age, A/HRC/27/37*, 30 June 2014. Source
2. **European Data Protection Board.** *Guidelines 01/2020 on Processing Personal Data in the Context of Connected Vehicles and Mobility Related Applications, Version 2.0*, 9 March 2021. Source
3. **GSMA Mobile for Development.** *Leveraging Telco Data for Predictive Credit Scorecards in Mobile Money*, 13 November 2014. Source
4. **European Union.** *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, 13 June 2024. See in particular Articles 3, 6, 14, 16, 25 and 26, and Annex III. Source
5. **European Commission.** *Draft Commission Guidelines on the Classification of High-Risk AI Systems and Targeted Consultation*, 19 May 2026. Consultation deadline extended to 23 July 2026. Source
6. **Council of the European Union.** *Artificial Intelligence: Council Gives Final Green Light to Simplify and Streamline Rules*, 29 June 2026. Source
7. **European Union.** *Regulation (EU) 2016/679 (General Data Protection Regulation)*, 27 April 2016. See Articles 13, 14, 15 and 22. Source
8. **Court of Justice of the European Union.** *Judgment in Case C-203/22, CK v Magistrat der Stadt Wien and Dun & Bradstreet Austria*, 27 February 2025. Source
9. **European Data Protection Board.** *Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models*, 17 December 2024. Source
10. **Government of India, Ministry of Electronics and Information Technology.** *Digital Personal Data Protection Act, 2023*, 11 August 2023. Source
11. **Government of India, Ministry of Electronics and Information Technology.** *Digital Personal Data Protection Rules, 2025*, 14 November 2025. Source
12. **Government of India, Ministry of Electronics and Information Technology.** *Notification bringing provisions of the Digital Personal Data Protection Act into force on a phased basis*, 14 November 2025. Source

13. **Saudi Data and AI Authority.** *Personal Data Protection Law*, 2023 consolidated English version. Source
14. **Saudi Data and AI Authority.** *Regulation on Personal Data Transfer Outside the Kingdom, Version 2.0*, August 2024. Source
15. **Saudi Data and AI Authority.** *AI Ethics Principles*, September 2023. Source
16. **Saudi Data and AI Authority.** *Artificial Intelligence Adoption Framework*, 16 November 2025. See also SDAIA AI Ethics Assessment services. Source
17. **ASEAN Secretariat.** *ASEAN Guide on AI Governance and Ethics*, February 2024. Source
18. **ASEAN Secretariat.** *Expanded ASEAN Guide on AI Governance and Ethics - Generative AI*, January 2025. Source
19. **ASEAN Secretariat.** *ASEAN Responsible AI Roadmap*, 2025. Source
20. **African Union.** *Continental Artificial Intelligence Strategy: Harnessing AI for Africa's Development and Prosperity*, July 2024. Source
21. **Government of the United Kingdom.** *Digital, Culture, Media and Sport Secretary's Statement on Telecoms*, 14 July 2020. The statement estimated cumulative delays of two to three years and costs of up to GBP 2 billion. Source
22. **Deutsche Telekom.** *Germany's First AI Factory for Industry Officially Goes into Operation*, 4 February 2026. See also the 2 March 2026 Industrial AI Cloud update. Source
23. **Singtel.** *Singtel Launches RE:AI, an AI Cloud Service*, 10 October 2024. Source
24. **Singtel.** *RE:AI and Mistral AI Partner to Strengthen Singapore's Industry AI Capabilities*, 27 April 2026. Source
25. **SK Telecom.** *The Cathedral of AI at MWC26 and SKT's Bold Path to Change*, 13 March 2026. See also SK Telecom's 22 February 2026 MWC preview on A.X K1. Source
26. European Parliament and Council of the European Union. PE-CONS 30/26: adopted text of the Digital Omnibus on AI, 18 June 2026. Source See the amended definition of safety component and the transitional provisions for Article 50.
27. EUR-Lex. Procedure 2025/0359/COD: Digital Omnibus on AI, status checked 18 July 2026. Source Records Council adoption on 29 June 2026, signature on 8 July 2026 and the procedure as ongoing pending Official Journal publication.
28. European Commission, AI Office. Signing the Code of Practice on Transparency of AI-Generated Content, updated 8 July 2026. Source Explains the 2 August 2026 application date and the targeted Article 50(2) transition for pre-existing systems.
29. Consultative Group to Assist the Poor (CGAP). How M-Shwari Works: The Story So Far, April 2015. Source Documents the use of Safaricom variables relating to airtime, airtime credit, M-Pesa and customer tenure in credit scoring.