Technology Stewardship

INSTITUTE • GOVERNING AUTONOMOUS TECHNOLOGY

INSTITUTE FOR TECHNOLOGY STEWARDSHIP

ISSUES PAPER No. 4

The Ghost in the Network

Shadow AI, Embedded Models and the Missing Inventory of Critical Infrastructure

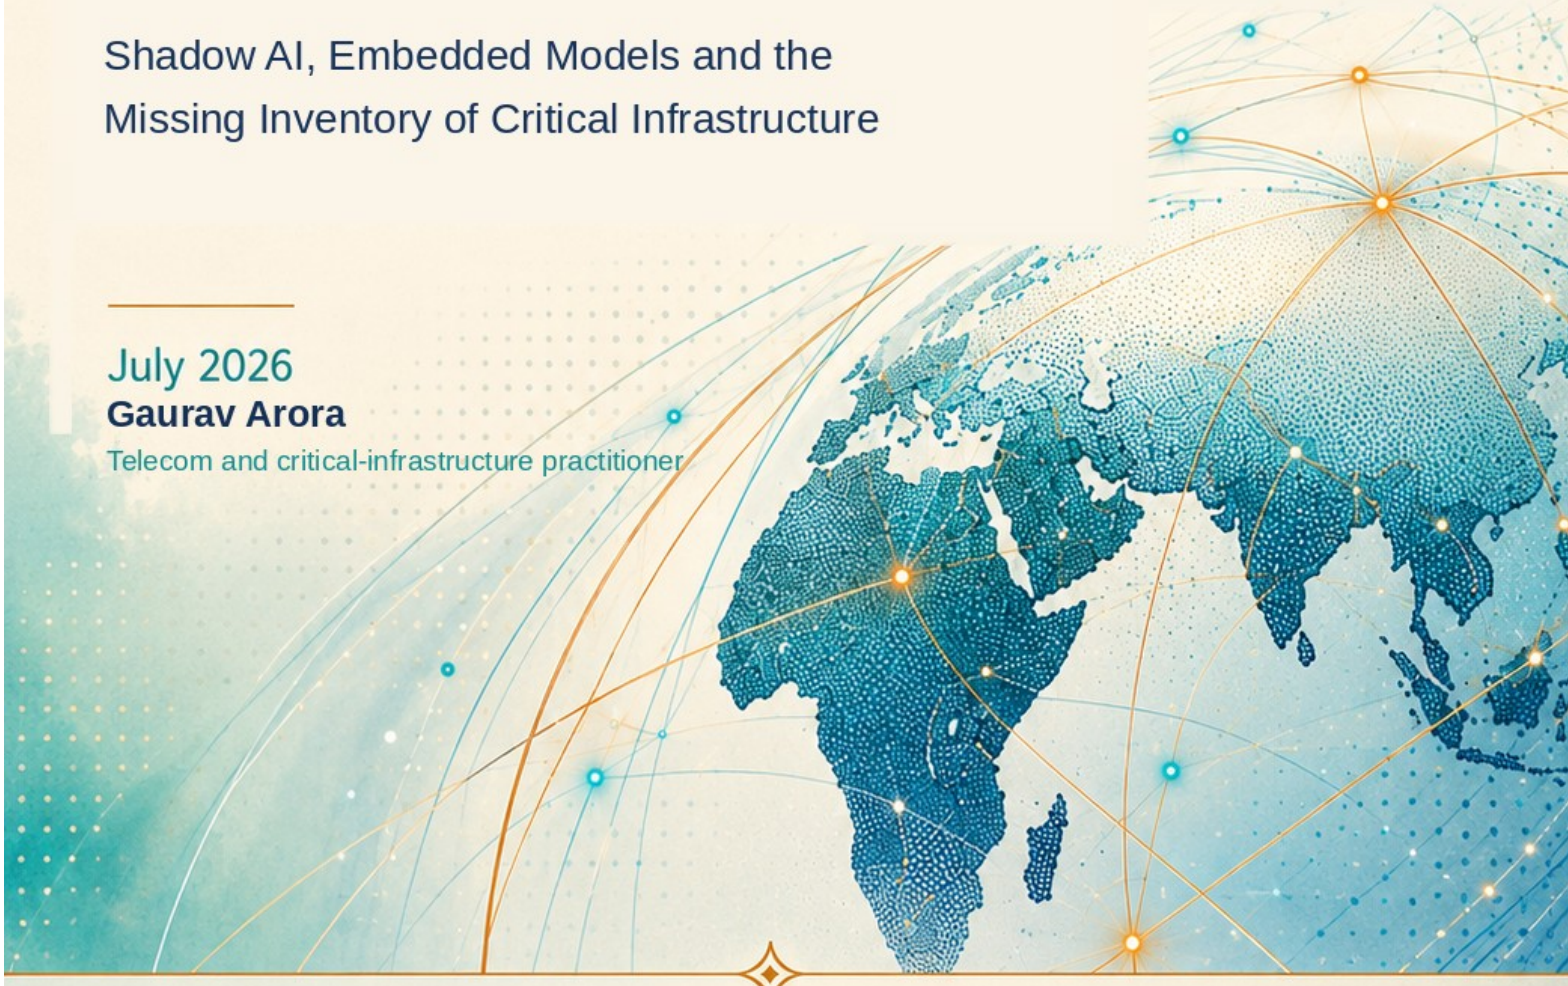

July 2026

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: telecommunications; shadow AI; AI inventory; Governed AI System Registry; workforce rights; critical infrastructure; Global South

Every AI governance architecture begins with inventory. NIST AI RMF 1.0 explicitly requires mechanisms to inventory AI systems under GOVERN 1.6, and the NIST Playbook describes an organised record of systems, documentation, responsible actors and supporting plans.[1,2] The problem is therefore not that governance frameworks ignore inventory. The problem is that conventional inventories are declaration-based, while modern AI enters telecom operations through employees, contractors, vendor platforms, software updates, local models, cloud services and autonomous agents.

IBM reported that one in five organisations studied experienced a breach involving shadow AI; organisations with high shadow-AI exposure incurred an average USD 670,000 in additional breach cost. IBM also reported that, among the relatively small subset of organisations experiencing an AI-model or application security incident, 97 per cent lacked proper AI access controls.[3] These figures are cross-sector and do not establish telecom prevalence. They nevertheless demonstrate that undeclared AI use is an operational security problem, not a speculative governance concern.

Telecom makes the discovery problem unusually consequential. Engineers handle call records, location telemetry, subscriber information, network topology, firmware details and live operational states. A public model may influence an alarm diagnosis, generate a configuration or process a photograph of network equipment even when no formal AI deployment has been approved. The model can therefore affect critical infrastructure through a human proxy while remaining absent from the official inventory.

This paper proposes a Governed AI System Registry: a continuously maintained, rights-respecting architecture that integrates four functions. The Enterprise Layer records declared systems. The Vendor Layer captures AI embedded in purchased products and managed services. The Operational Layer provides usable approved alternatives at the point of need. The Discovery Layer identifies undeclared or suspected use and measures the confidence and residual blind spots of the resulting inventory.

The contribution is not a promise of complete visibility. Complete visibility is generally impossible. Personal devices, encrypted traffic, local models, contractor environments and opaque vendor platforms create unavoidable uncertainty. The proposal therefore treats inventory completeness as a measured claim. A credible registry states what is known, what is detected, what is suspected, which discovery methods were used, when the evidence was last verified and what remains outside coverage.

The Global South dimension is central. Many operators rely heavily on managed-service providers, shared devices, foreign cloud services and vendor-controlled network-management platforms while having less access to enterprise discovery tooling, local inference capacity and specialist assurance teams. An inventory standard must not become a procurement barrier that only the largest operators can satisfy. The paper therefore proposes modular implementation, shared assurance resources, open schemas and a minimum viable registry that is explicitly provisional rather than falsely complete.

Inventory also matters to people. If an AI system is absent from the registry, an operator may be unable to explain why service was restricted, identify which vendor influenced a decision, reconstruct an incident or provide a meaningful route to correction and remedy. Inventory is not administrative housekeeping. It is the first condition of explanation, contestability and accountability.

THREE INVENTORY QUESTIONS

1. What declared, vendor-embedded, detected and suspected AI systems interact with the operator's network, data or decisions?
2. How confident is the inventory, which discovery methods support it, and what residual blind spots remain?
3. When an incident or adverse decision occurs, can the operator reconstruct the role of every relevant AI system and identify a responsible route for review and remedy?

A registry that cannot state its uncertainty is not an assurance mechanism; it is a list.

1. Methodology, Scope and Definitions

This is a desk-based issues paper focused on AI systems used in or around telecommunications operations and subscriber-facing processes. Sources include primary legal texts, official government guidance, institutional frameworks, industry publications and the earlier ITS papers. Company and workforce scenarios are illustrative composites based on ordinary telecom workflows; they are not allegations against a particular operator or employee.

The analysis distinguishes three categories that are often collapsed. Declared AI is deliberately procured or developed and known to governance teams. Embedded AI is introduced through vendor software, network equipment, managed services or updates and may

be only partially disclosed. Shadow AI is used or operated outside the approved governance path, including personal accounts, local models, unauthorised plugins and undeclared agents. A system can move between categories as it is discovered, assessed, approved, prohibited or retired.

The paper also distinguishes inventory from discovery. An inventory is the maintained record. Discovery is the continuing process used to test whether that record reflects operational reality. The Governed AI System Registry joins the two and requires the organisation to state its confidence rather than claim perfect completeness.

Legal references are illustrative, not exhaustive legal advice. The EU AI Act does not classify every telecom AI system as high-risk; classification depends on intended purpose and actor role.[6] CIRCIA reporting requirements remain prospective until CISA's final rule becomes effective.[7] CISA and international partner guidance on agentic AI is non-binding guidance, although it may influence procurement, security practice and regulatory expectations.[4]

2. Existing Inventory Requirements - and the Discovery Gap

NIST AI RMF 1.0 already addresses inventory. GOVERN 1.6 states that mechanisms should be in place to inventory AI systems and resource that activity according to organisational risk priorities.[1] The Playbook describes an inventory as an organised database that can contain system documentation, incident-response plans, data dictionaries, code links and responsible actors.[2] Any ITS proposal must therefore begin by recognising this foundation rather than presenting AI inventory as a new concept.

The unresolved problem is operational discovery. Conventional governance records are usually populated through procurement, architecture review, model-development processes or self-declaration. Those channels identify systems that have entered the approved lifecycle. They do not reliably reveal an engineer using a personal account, a contractor running a local model, a vendor enabling an AI feature through an update, a managed-service agent operating under remote credentials or a foundation-model dependency that changed without operator approval.

This is the discovery gap: the difference between AI that the organisation has declared and AI that can influence its data, infrastructure or decisions. A conventional inventory is principally a declaration - a record of what the organisation believes it has. The Governed AI System Registry is also an estimator: it continuously tests that declaration against operational signals, identifies what may be missing and states the uncertainty surrounding its own coverage. A useful registry must therefore combine governance records with technical and organisational discovery signals, and preserve uncertainty when those signals are incomplete.

The ITS contribution is not the invention of AI inventory. It is the integration of declaration-based records, continuous discovery, vendor disclosure, operational enablement and calibrated confidence into one rights-respecting governance mechanism. A registry that cannot state its uncertainty is not an assurance mechanism; it is a list.

3. Shadow AI in Telecom Operations

3.1 The field engineer

A field engineer is troubleshooting equipment whose alarms do not match the current manual. Vendor support has a long callback window. The engineer photographs a panel and asks a public multimodal model to identify the fault. The model responds quickly and the immediate operational problem is solved.

The photograph may also expose equipment serial numbers, firmware versions, network or rack topology, and embedded geolocation metadata. The deeper issue is not only data leakage. The model may influence a diagnosis or configuration decision without system identification, approved data handling, independent verification or a reconstructable decision record.

3.2 The NOC engineer

A NOC engineer managing an alarm storm at night has access to an approved assistant, but the tool is separated from operational terminals and refuses to process topology-related data. A personal account returns a correlation in seconds. Redaction is improvised under time pressure, and operational identifiers or subscriber-linked fragments may remain.

The shadow tool has now become part of the decision chain. If the recommendation is followed, the operator may be unable to reconstruct which model influenced the response, what information left the environment or which limitations applied. This is ungoverned decision influence, even though a human remained formally responsible for the action.

3.3 Vendor-embedded and shadow agentic AI

The discovery challenge is not limited to conversational tools. Network-management products increasingly incorporate optimisation, remediation and agentic functions. A vendor may activate or materially change an AI capability through a software update, managed-service workflow or remote orchestration layer without a separate AI procurement event. Joint guidance on secure AI integration in

operational technology reinforces the need to identify AI-enabled components, their interfaces and their operational dependencies before they are trusted in consequential environments.[11] NIST's AI Agent Standards Initiative likewise identifies interoperable identity, authorisation and secure agent interactions as emerging standards priorities.[12]

The operator may therefore face a difficult combination: legal and operational accountability for a system whose model, update process, permissions or event records are controlled elsewhere. The Vendor Layer and Discovery Layer must treat embedded autonomy as a first-class inventory problem, not merely a procurement footnote.

4. Decision Influence and Human Consequences

AI use should be classified by the degree to which it can influence an operational or subscriber outcome. A web-like technical query does not create the same governance exposure as a generated configuration or an autonomous network action. The classification below is an internal ITS risk model, not a legal classification.

Influence level	Description	Minimum governance response
Information retrieval	General technical information; no operator-specific data or action recommendation.	Basic usage record and approved-tool policy.
Interpretation assistance	Model interprets operational information; human forms an independent judgement.	Data classification, system identification and user responsibility.
Recommendation	Model proposes a specific root cause, service restriction or operational action.	Decision-event record, qualified human review and outcome monitoring.
Configuration generation	Model creates scripts, commands, rules or change content.	Mandatory testing, provenance tagging, approval and rollback.
Autonomous execution	Agent uses tools or permissions to change infrastructure or a consequential subscriber state.	Paper 2 authority controls, runtime enforcement, continuous monitoring and safe stop.

The affected-person consequence can be as important as the technical action. A false fraud decision may suspend a mobile number that is also used for mobile-money authentication, remittances, digital identity and public services. If the influencing system is absent from the registry, the operator may be unable to identify the model and version, explain the intervention, direct the complaint to the correct owner or determine whether a vendor participated. Discovery is therefore a condition of redress, not only cybersecurity.

RIGHTS-BASED PRINCIPLE

1. No person should be denied a meaningful review because the operator cannot identify the AI system that influenced the decision.
2. Every consequential system should have a responsible owner, decision record and route for correction.
3. Discovery records used for workforce governance must not become general productivity-surveillance data.

Inventory is the first condition of explanation, contestability and remedy.

5. The Governed AI System Registry

The Governed AI System Registry is a continuously maintained system of record and discovery process for AI that interacts with telecom infrastructure, data, employees, contractors or subscriber decisions. It is analogous to a CMDB only in part: unlike hardware inventory, its subject can be remotely updated, accessed through personal accounts, embedded in third-party services and changed at the model or orchestration layer without a visible physical asset. The registry should therefore operate as an AI-specific system of record integrated with, but not necessarily contained within, the enterprise CMDB. The CMDB remains authoritative for infrastructure assets and service relationships; the registry adds model versions, inference locations, vendor dependencies, decision influence, authority levels, discovery confidence and affected-person accountability. Appendix A provides the minimum registry record.

The registry uses a status taxonomy rather than a binary approved/unapproved field: approved; provisionally approved; vendor-embedded; detected shadow use; suspected or estimated; prohibited; under investigation; and retired. Status changes must be time-stamped and linked to the evidence supporting the change.

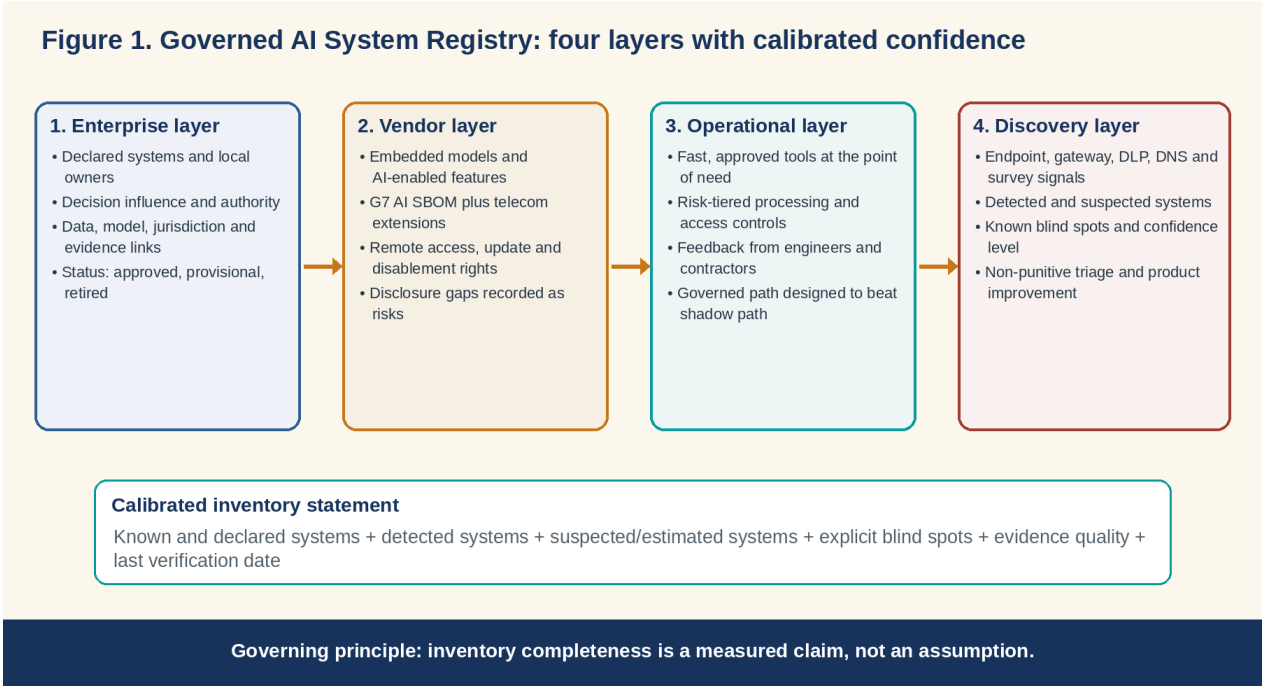

Figure 1. The four registry layers and the calibrated inventory statement.

5.1 Enterprise Layer: declared systems

The Enterprise Layer records systems deliberately procured, developed or approved. At minimum it should identify the system and deployment instance; model and provider; operational owner; intended purpose; data categories; decision-influence level; autonomy and tool permissions; inference location and jurisdiction; human-oversight arrangement; decision-event logging; linked obligations; evidence sources; status; last verification date; and known limitations.

The registry should not require hidden chain-of-thought. The auditable object is the decision event: system and model version, relevant inputs and operational state, tools and permissions used, proposed and executed actions, policy checks, human intervention, outcome, rollback and escalation. This aligns Paper 4 with the vendor-neutral decision-event record developed in Paper 2.[14]

5.2 Vendor Layer: embedded AI and supply-chain evidence

The Vendor Layer records AI introduced through equipment, software, managed services and updates. The joint CISA and G7 partner guidance on Software Bill of Materials for AI establishes minimum elements across metadata, system properties, models, datasets, infrastructure, security properties and key performance indicators.[5] The registry should consume those elements rather than create a competing disclosure format.

Telecom-specific extensions remain necessary: operational domain; decision-influence and authority level; tool and permission access; inference and support jurisdiction; remote administrative access; model-update authority; agent-to-agent dependencies; consequence of model unavailability; operator disablement and rollback rights; and evidence-export capability. Missing vendor disclosure must be recorded as an explicit risk, not silently converted into a complete registry entry.

Telecom extension	Required disclosure
Composition	G7 AI SBOM elements: system, models, datasets, infrastructure, security and performance.
Operational role	Network domain, decisions influenced, authority and permitted tools.
Control and jurisdiction	Inference location, remote support, key ownership, update and rollback control.
Evidence rights	Event export, audit access, incident notification and retention.
Operator exit	Disablement, substitution, data return and service-continuity consequences.

5.3 Operational Layer: make the governed path usable

A registry that only records violations will not solve the cause of shadow AI. Approved tools must be available from NOC terminals, field devices and contractor workflows; support text, image, code and structured-data tasks; respond within operationally useful time; and route sensitive information to an appropriate processing environment. The governed path must be easier to use than the shadow path for legitimate work.

Processing tiers should be based on demonstrable controls rather than simplistic labels such as cloud, sovereign or on-premises. Local hosting does not by itself establish sovereignty. Relevant factors include jurisdiction, remote access, encryption-key control, replication, support location, software updates, logging, subprocessors and exit rights, as Paper 1 explains.[13] Data treatment should distinguish irreversible anonymisation from pseudonymisation, tokenisation and minimisation.

If an automated classifier routes prompts or operational data between processing tiers, that classifier is itself an AI or rules-based governance component. It must be registered, tested for false negatives, versioned, monitored and designed to fail safely.

5.4 Discovery Layer: continuous discovery with calibrated confidence

The Discovery Layer uses multiple signals because no single method can establish operational reality. Possible sources include endpoint inventory, browser and plugin records, API gateway logs, DLP alerts, DNS and traffic analysis, CASB records, vendor attestations, procurement review, confidential workforce surveys, self-reporting and incident forensics. Each method has blind spots and legal constraints.

Discovery source	What it can reveal	Known limitation
Endpoint and software inventory	Local models, plugins, AI applications and managed-browser use.	Misses personal and unmanaged devices; may miss browser-only use.
API, DNS and network signals	Connections to known AI services and unusual data flows.	Encryption, shared infrastructure, VPNs and new services reduce precision.
DLP and data-flow monitoring	Sensitive data moving toward AI-related destinations.	False positives; paraphrasing and novel formats may evade detection.
Vendor and contract discovery	Embedded features, managed-service agents and external inference.	Depends on disclosure quality and bargaining power.
Confidential surveys and self-report	Use cases that technical controls cannot see and capability gaps.	Subject to recall and response bias; requires credible non-punitive safeguards.
Incident and change review	Previously unknown AI influence revealed during investigations.	Reactive; evidence may be incomplete after the event.

The registry should not publish invented coverage percentages. Instead, each operator should document discovery scope, managed-device coverage, third-party coverage, known encrypted or unmanaged channels, survey calibration, false-positive and false-negative testing, last verification date and confidence by system class. Confidence should be expressed as supported, partially supported or uncertain, with the evidence basis stated.

CALIBRATED INVENTORY STATEMENT

1. Known and declared systems are listed with current owners and verification dates.
2. Detected systems are recorded with the signal and confidence supporting detection.
3. Suspected systems are retained as open discovery cases rather than discarded.
4. Blind spots are documented by workforce, device, vendor, network and jurisdictional coverage.
5. The organisation states what the registry cannot presently see.

Inventory completeness is a measured claim, not an assumption.

6. Workforce Rights and Non-Punitive Discovery

Endpoint, browser, DNS, gateway and DLP monitoring can become workforce-surveillance infrastructure regardless of the stated purpose. If engineers believe discovery data will be used for discipline or productivity scoring, they will move legitimate use to less visible channels and the registry will become less accurate. A rights-respecting architecture is therefore an operational requirement, not an optional ethics add-on.

The paper rejects performance-based suspicion. An engineer who resolves alarms faster than colleagues is not evidence of shadow AI use. Discovery should rely on system, endpoint, network, vendor or voluntary-disclosure signals, not unusual competence, protected characteristics or general behavioural scoring.

Safeguard	Minimum requirement
Purpose limitation	Use discovery data only for AI governance, security and explicitly defined incident investigation.
Notice and consultation	Inform employees and contractors; consult representatives where law or practice requires.
Data minimisation and retention	Collect the minimum necessary metadata; define short, risk-based retention.
Separation from performance management	Prohibit use for productivity scoring, routine appraisal or unrelated discipline.
Safe harbour	Good-faith self-reporting and first detected non-malicious use default to product improvement and education.
Due process	Security referrals require additional evidence, documented criteria, notice and a right to respond.
Independent oversight	Audit access, false-positive rates, referrals, retention and adherence to purpose limits.
Aggregate transparency	Publish internal statistics on detections, improvements, security referrals and unresolved gaps.

NASA's Aviation Safety Reporting System demonstrates the operational value of a confidential, voluntary and non-punitive reporting channel. The transferable lesson is not blanket immunity: it is that credible reporting protections can reveal system vulnerabilities and complement mandatory reporting and investigation rather than replace them.[16]

Safe harbour protects the reporter; it does not erase the event. A good-faith disclosure or first non-malicious detected use should not ordinarily be used for performance management or disciplinary action against the person reporting it. The operator must nevertheless assess the underlying data exposure, operational consequence and applicable notification duties, and complete any required remediation or reporting. These limits must be explained before workers are asked to rely on the safe-harbour process. Deliberate exfiltration, malicious conduct and intentional unsafe acts remain outside its protection.

Contractors and managed-service personnel should receive equivalent access to approved tools and proportionate safe-harbour protections. Contract terms alone are insufficient if the individual engineer never receives the approved tool, the data-handling guidance or the reporting channel.

7. Global South and Proportional Implementation

The discovery problem is global, but institutional capacity is uneven. Many operators in Africa, South Asia, Southeast Asia, the Middle East and small-island markets rely heavily on managed services, shared devices, foreign cloud infrastructure and network platforms whose AI functions are designed and updated elsewhere. They may also lack full EDR, CASB, DLP, model-assurance teams and local-language enterprise assistants.

These environments are not regulatory vacuums. Kenya's Data Protection Act 2019 establishes rights for data subjects and obligations for controllers and processors, and the accompanying regulations address access, rectification, objection and other operational rights.[8,9] The African Union Continental AI Strategy calls for Africa-centred, responsible and inclusive AI governance and capacity-building across member states.[10] A missing AI inventory can therefore undermine existing privacy, consumer and institutional accountability even where no comprehensive AI statute exists.

Capacity may determine how broadly an operator deploys discovery and governance, but it should not determine whether the operator states the resulting uncertainty honestly. A smaller operator may begin with fewer systems, narrower discovery coverage and shared regional assurance resources. It should not present partial coverage as complete or lower the reliability threshold for a consequential claim. Open schemas, pooled vendor assessments, shared test facilities and regulator-supported contract clauses can reduce implementation cost without weakening evidentiary honesty.

7.1 Minimum viable registry

A minimum viable registry is a provisional operating baseline, not a compliance certificate. It can be established through four steps:

1. Catalogue all declared systems and the highest-consequence vendor-embedded systems, with owner, purpose, data categories, decision influence and last verification date.
2. Use one or two feasible discovery methods - for example managed-device inventory and a confidential survey - and document their blind spots.
3. Identify the most common legitimate shadow use case and provide one approved alternative that is accessible at the point of work.
4. Publish an internal calibrated statement explaining what is covered, what remains uncertain and the plan for expanding discovery.

Regional regulators, operator associations and development partners can reduce cost through shared contract clauses, common registry schemas, pooled vendor assessments, regional test facilities and incident-support capacity. Open evidence exports are especially important where a local operator lacks leverage over a global vendor.

GLOBAL SOUTH PRINCIPLE

1. An inventory standard must not become a procurement barrier that only the largest operators can satisfy.
2. Vendor-controlled autonomy must not eliminate local operator control over discovery, logging and authority contraction.
3. Shared regional services can reduce implementation cost, but local accountability and affected-person rights remain with the responsible actor.

Reduce the scope of unsupported deployment - not the reliability of high-consequence claims.

8. Inventory, Explanation and Remedy

The registry must serve more than internal governance. Where AI materially influences service access, fraud restriction, credit, pricing, emergency communications or another consequential outcome, an affected person may need an explanation, human review, correction or remedy. Paper 3 identifies an affected-person evidence interface as a distinct output of the Compliance Translation Layer.[15] Paper 4 supplies the system identity and ownership needed to make that interface real.

A rights-oriented registry should support the same affected-person evidence and redress interface defined in Paper 3, with Paper 4 supplying the system identity, ownership, vendor participation and discovery-confidence fields. Where legally and operationally appropriate, the package should include: the system and decision identifier; responsible operator function; principal data categories; system and model version; whether a vendor participated; human-review status; available challenge route; correction status; and any disclosure limitation required for security or third-party rights.

The registry should also record when no explanation can be generated because the system was unknown, the vendor did not provide evidence or the decision record was not retained. Such failure is not merely an administrative gap. It is an unresolved accountability incident requiring remediation and, where applicable, notification to compliance or supervisory functions.

9. The Four-Paper ITS Architecture

The three questions presented in the Executive Summary constitute the Paper 4 diagnostic. Their answers provide the discovery input required by the system-control, authority and compliance layers described in Papers 1-3. The four papers are numbered as publications but operate in a different analytical sequence: discovery comes first, followed by system control, delegated authority and compliance evidence.

Layer	ITS mechanism	Source	Core question
Discovery	Governed AI System Registry	Paper 4	What AI exists, how was it discovered and how complete is the inventory?
System and data control	Operational Sovereignty Diagnostic	Paper 1	What data built and feeds the system, where does processing occur and can decisions be reconstructed?
Delegated authority	Continuous Shadow Deployment	Paper 2	Was authority earned through realistic evidence, and can humans intervene in time?
Compliance evidence	Compliance Translation Layer	Paper 3	Which requirements apply, what evidence supports each claim and who owns unresolved gaps?

The sequence is intentionally dependent. Paper 1 cannot establish control over a system that has not been discovered. Paper 2 cannot evidence authority without a system identity and decision record. Paper 3 cannot translate governance into compliance when the contributing system, vendor or evidence source is unknown. Paper 4 is therefore the discovery foundation of the complete architecture.

10. Recommendations

10.1 For telecom operators

- Establish a Governed AI System Registry with explicit declared, embedded, detected, suspected, prohibited and retired statuses.
- Replace claims of complete inventory with calibrated confidence statements, verification dates and known blind spots.
- Require decision-event records for recommendation, configuration-generation and autonomous-execution systems.
- Provide fast, multimodal approved tools at the point of operational need and treat legitimate shadow use as product feedback.
- Create workforce purpose limits, safe harbour, due process and separation from performance management.
- Integrate the registry with CMDB, procurement, vendor management, incident response, change management and the Compliance Translation Layer.

10.2 For vendors and managed-service providers

- Provide G7 AI SBOM-aligned disclosures with telecom-specific operational and authority extensions.
- Disclose model, inference, remote-access, update, agent, permission and evidence-export arrangements before deployment and after material change.
- Provide operator-controlled disablement, rollback and local event export for consequential functions.
- Extend approved tools, training and reporting channels to contractor and managed-service personnel.

10.3 For regulators, standards bodies and international organisations

- Define reasonable discovery expectations rather than impossible absolute-completeness duties.
- Promote open, vendor-neutral registry schemas and map them to AI SBOM, control and incident-reporting standards.
- Create shared assurance and capacity-building mechanisms for smaller and emerging-market operators.
- Require affected-person review routes where an unidentified or vendor-controlled AI system materially influences a consequential decision.
- Consult workers, civil society and operational practitioners when designing discovery and monitoring requirements.

11. Limitations, Research Agenda and Call for Collaboration

This paper does not provide validated detection percentages, universal retention periods or fixed implementation timelines. Such values depend on operator architecture, workforce model, jurisdiction, risk appetite and available tooling. The proposed fields and controls require empirical testing with operators, vendors, workers, regulators and civil-society organisations.

Priority research questions include: how to measure false negatives in shadow-AI discovery; how local models and encrypted channels can be governed without disproportionate surveillance; which AI SBOM elements reliably predict operational risk; how contractors and managed-service providers can be included across jurisdictions; how affected-person explanations can remain meaningful without exposing security-sensitive information; and how shared assurance facilities can support smaller operators.

The Institute for Technology Stewardship invites telecom operators, regulators, vendors, engineers, worker representatives, researchers and civil-society organisations to test the Governed AI System Registry and contribute anonymised lessons. Responses may be submitted through technologystewardship.org.

Appendix A. Minimum Registry Record

The following schema is a minimum record for a consequential AI system. It is deliberately technology-neutral and should be implemented in a machine-readable format where possible. The record is cross-consistent with Paper 2's decision-event schema and Paper 3's evidence-lineage and affected-person interfaces: Paper 4 identifies the system and its discovery confidence, Paper 2 records consequential decisions and actions, and Paper 3 traces current evidence to qualified compliance claims.

Record group	Minimum fields
Identity and status	Unique system/deployment ID; declared, embedded, detected, suspected, prohibited or retired; owner and last verified date.
Purpose and context	Intended purpose, network or subscriber domain, users, jurisdictions and

Record group	Minimum fields
	consequence of failure.
Model and vendor	Provider, model and version, foundation-model dependencies, AI SBOM link, update authority and support access.
Data and processing	Data categories, training/fine-tuning use, inference location, storage, replication, key control and subprocessors.
Decision influence	Information, interpretation, recommendation, configuration generation or autonomous execution.
Authority and tools	Permissions, tools, limits, escalation conditions, human oversight, safe stop and rollback.
Evidence and lineage	Decision-event record, change record, vendor attestation, discovery source, integrity protection and retention.
Obligations and rights	Applicable legal, contractual and influential requirements; affected-person notice, review and remedy route.
Confidence and blind spots	Evidence quality, discovery methods, known unmanaged channels, unresolved vendor opacity and next review date.

About the Institute for Technology Stewardship

The Institute for Technology Stewardship (ITS) is an independent research initiative focused on the operational governance of autonomous and emerging technologies in telecommunications and critical infrastructure. ITS develops practical mechanisms that connect system discovery, data control, delegated authority, compliance evidence and human accountability.

Author: Gaurav Arora, 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 and technology programme leadership.

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 Ghost in the Network: Shadow AI, Embedded Models and the Missing Inventory of Critical Infrastructure. Institute for Technology Stewardship, Issues Paper No. 4, July 2026.

Disclaimer

This issues paper reflects independent analysis and does not represent the position of any operator, regulator, standards body or government. Legal references are included for analytical purposes and should not be treated as legal advice.

References

1. NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023. Source

2. NIST AI Resource Center, AI RMF Playbook - GOVERN 1.6, AI system inventory. Source

3. IBM and Ponemon Institute, Cost of a Data Breach Report 2025: The AI Oversight Gap, 2025. Source

4. CISA and international partners, Careful Adoption of Agentic AI Services, 1 May 2026. Source

5. CISA and G7 international partners, Software Bill of Materials for AI - Minimum Elements, 12 May 2026. Source

6. European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act). Source

7. CISA, Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA) overview and FAQs. Source

8. Republic of Kenya, Data Protection Act No. 24 of 2019. Source

9. Republic of Kenya, Data Protection (General) Regulations, 2021. Source

10. African Union, Continental Artificial Intelligence Strategy, endorsed July 2024. Source

11. CISA and international partners, Principles for the Secure Integration of Artificial Intelligence in Operational Technology, December 2025. Source

12. NIST, AI Agent Standards Initiative, February 2026. Source

13. Arora, Gaurav. The Unpriced Risk, Issues Paper No. 1, July 2026. Source

14. Arora, Gaurav. Beyond the Sandbox, Issues Paper No. 2, July 2026. Source

15. Arora, Gaurav. The Compliance Collision, Issues Paper No. 3, July 2026. Source

16. NASA, Aviation Safety Reporting System - confidential, voluntary and non-punitive safety reporting. Source