Technology Stewardship

INSTITUTE • GOVERNING AUTONOMOUS TECHNOLOGY

INSTITUTE FOR TECHNOLOGY STEWARDSHIP

ISSUES PAPER No. 3

The Compliance Collision

Translating AI Governance into Regulator-Ready Evidence for Telecommunications

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; AI governance; compliance evidence; regulatory interoperability; NIST AI RMF; Global South; critical infrastructure

Telecommunications operators increasingly use the NIST Artificial Intelligence Risk Management Framework (AI RMF), sector guidance and internal assurance programmes to organise AI risk. This is rational. The AI RMF is voluntary, risk-oriented and designed to help organisations identify, measure and manage AI risks. It can generate inventories, assessments, test results, monitoring records and governance decisions.[1-3]

The compliance problem begins when those governance outputs are treated as proof that a specific legal, contractual, procurement or supervisory obligation has been satisfied. A framework does not determine whether an obligation applies to this system and actor, whether the underlying control operated, whether the evidence is current, or whether a regulator would regard the resulting claim as legally sufficient.

This paper proposes a Compliance Translation Layer (CTL): an operational architecture that maintains the traceable chain from governance outcome, to implemented control, to verifiable evidence, to a qualified compliance claim and external output. The CTL does not manufacture compliance. It maps, collects, validates, qualifies and packages evidence. Where an underlying control is absent or the evidence is insufficient, the layer must reject the claim and open a remediation gap.

The proposal builds directly on two emerging international foundations. The CISA and G7 partner guidance on Software Bill of Materials for AI provides a seven-cluster description of AI system composition, covering metadata, system-level properties, models, datasets, infrastructure, security properties and key performance indicators.[11] NIST OSCAL provides machine-readable formats for controls, implementations, mappings and assessment results.[4-6] The CTL fills the remaining operational gap: it links those system and control records to system-specific obligations, actor roles, evidence lineage, legal review, regulator-facing outputs and affected-person explanation and redress.

Harmonisation is already advancing through shared schemas, crosswalks, substantially similar reporting concepts and efforts to reduce overlapping requirements. The unresolved problem is evidence-level interoperability. A common schema does not prove that a control operated; a control mapping does not determine legal applicability; and a reporting exception is useful only where the operator can generate complete and trustworthy information in the first place.

The Global South dimension is central. Operators in Africa, South Asia, Southeast Asia and other emerging economies may deploy vendor-embedded AI across several markets while navigating local data-protection and telecom rules, multinational-group controls, foreign procurement expectations and voluntary international schemas. The objective is not to import the most complex foreign compliance stack. It is to ensure that local operators, regulators and affected communities can inspect the evidence relevant to local law, essential connectivity, financial access and public safety.

THREE COMPLIANCE-ARCHITECTURE QUESTIONS

1. Which binding, contractual and influential obligations apply to this system and actor?
2. Can every compliance claim be traced to a current control and verifiable evidence?
3. Who owns each unresolved gap, and by when must it be closed?

An operator that cannot answer all three has governance activity, but not a defensible compliance architecture.

Methodology, Scope and Terminology

This is a desk-based policy and operational analysis prepared in July 2026. It reviews primary regulatory texts, official agency guidance, NIST technical resources, regional AI-governance instruments and the companion ITS Issues Papers No. 1 and No. 2. The worked examples are illustrative and are not legal determinations about any named operator or product.

The paper is not legal advice and does not attempt an exhaustive jurisdiction-by-jurisdiction survey. Applicability depends on the system's intended purpose, deployment context, economic actor, location, licence conditions, contractual commitments and effective dates. Legal status is verified at the research cut-off and must be rechecked before operational reliance.

The term Global South is used as a practical shorthand for diverse countries with different legal systems, market structures and institutional capacities. It does not imply a single regulatory position or a lower assurance standard.

1. Legal Status Is the First Translation

The first function of a CTL is not document generation. It is classification. The same instrument may be binding law, an enacted requirement awaiting implementing rules, influential guidance, a voluntary governance framework or an ITS proposal. Treating these categories as equivalent creates false compliance claims before the evidence chain has even begun.

Legal-status category	Meaning for the CTL	Illustrative sources
Binding and operative	Assess applicability and support the claim with legally sufficient evidence.	Applicable FCC equipment-authorisation rules; local data-protection and telecom requirements.
Binding but system- or actor-dependent	Applies only where classification, role and territorial conditions are met.	EU AI Act high-risk duties, including distinct provider and deployer obligations.[13]
Enacted but awaiting operative rules	Build evidence capability, but do not present proposed fields as final law.	CIRCIA reporting; final rule remained pending at the research cut-off.[7,8]
Non-binding but influential	May shape procurement, assurance, insurance and supervisory expectations.	CISA and international-partner agentic-AI guidance; G7 AI SBOM guidance.[11,12]
Voluntary governance framework	Organises risk management but does not by itself establish legal sufficiency.	NIST AI RMF and regional voluntary guidance. [1,17,18]
ITS proposal	Provides an operational design for testing and collaboration.	Compliance Translation Layer; Governed AI System Registry; Continuous Shadow Deployment.

This taxonomy is deliberately reusable. Every obligation record in the CTL should carry a status, effective date, verification date, jurisdiction, responsible actor and interpretation owner. A change in any of those fields should trigger reassessment of the associated evidence and claims.

2. Governance Outcomes, Controls, Evidence and Compliance Claims

The argument depends on distinguishing four objects that are often collapsed into one another.

Object	Question it answers	Telecom example
Governance outcome	What condition should the organisation achieve?	Third-party AI risks are identified and managed.
Implemented control	What operational mechanism produces that condition?	Supplier screening, component inventory, restricted deployment and update approval.
Evidence	What reliable record shows that the control existed and operated?	Time-stamped vendor attestation, AI/SBOM record, approval log and scan result.
Compliance claim	What proposition is made to a regulator or another endpoint?	This specific system and version satisfied the applicable supply-chain requirement on the stated date.

The AI RMF can support all four objects, but it does not collapse them. GOVERN 6, for example, addresses third-party risks. An operator may document a mature supplier-risk policy and still lack a complete component inventory, a current vendor attestation or evidence that the approved version is the version operating in production.

Compliance therefore requires an additional sequence:

1. Determine whether the obligation applies to the system, jurisdiction and actor.
2. Identify the control or combination of controls needed to satisfy it.
3. Collect evidence from authoritative operational sources.
4. Validate completeness, integrity, timing and legal sufficiency.
5. Generate the required artifact, submission, attestation or disclosure.
6. Retain the evidence and reassess it when the system, actor or law changes.

A crosswalk that maps one framework clause to another is helpful but insufficient. It shows conceptual alignment. It does not prove control operation, evidence freshness or legal sufficiency. That distinction is the foundation of the Compliance Translation Layer.

3. The Multi-Regime Telecom Evidence Problem

Telecom AI systems sit inside layered legal and institutional environments. A single deployment may involve national data-protection law, telecommunications licence conditions, cybersecurity reporting rules, group governance, procurement commitments, vendor contracts, foreign-market requirements and voluntary technical guidance. Different endpoints ask different questions of the same system:

- What hardware, software, model, dataset and supplier components are present?
- What incident occurred, when was it detected and what was its operational impact?
- What role did a human overseer play, and could the action be interrupted or reversed?
- Which agent or service identity invoked a tool, with what permission and under whose authority?
- What personal data was processed, transferred or used to produce a consequential decision?
- What explanation, review and remedy were available to an affected subscriber?

The practical objective is not four separate evidence universes. It is evidence reuse without legal flattening: collect reliable operational records once, preserve their lineage, and generate only those claims that the relevant jurisdiction and actor profile supports.

3.1 Harmonisation is advancing - but mostly above the evidence layer

The regulatory environment is not uniformly fragmented. NIST publishes crosswalks and OSCAL models; the G7 AI SBOM guidance develops a common system-description vocabulary; CIRCIA rulemaking includes mechanisms intended to reduce duplicative reporting; and the EU Digital Omnibus seeks to reduce overlap in parts of the AI and product-safety architecture.[3-6,8,11,14,15] These developments matter and should be used.

They do not eliminate the operational evidence problem. A shared AI SBOM does not establish that the deployed version matches the approved record. An OSCAL mapping does not decide whether an operator is legally a provider, deployer or another actor. A substantially similar reporting exception helps only if the organisation can produce the information in the first place. Harmonisation reduces friction among frameworks and endpoints; the CTL addresses the remaining chain from system fact to defensible claim.

4. Four Compliance Collisions

A compliance collision occurs when an organisation can point to a governance outcome or policy but cannot produce the system-specific evidence needed to support the external claim. Each collision below follows the same pattern: governance outcome, implemented control, required evidence, validation gap and consequence.

4.1 System provenance and supply-chain evidence

The FCC's March 2026 Covered List update added specified foreign-produced consumer-grade routers, subject to approval exceptions and conditions. It is origin- and authorisation-based and does not prohibit equipment merely because it contains an AI model.[9,10] The earlier framing of this issue as 'NIST-compliant, FCC-prohibited' would therefore be a category error.

The real collision is more useful. An operator may have a complete GOVERN 6.1 third-party-risk policy yet be unable to answer which deployed network elements contain embedded models, which supplier and production location apply, which model and firmware version is running, which datasets or external services it depends on, and when those facts

were last verified. The policy outcome exists; the system-level provenance needed for screening, procurement or regulatory assurance does not.

The G7 AI SBOM guidance directly addresses part of this gap by defining minimum elements across metadata, system properties, models, datasets, infrastructure, security and performance.[11] The CTL must correlate those component records to live deployment instances, validate vendor assertions, track changes and apply the correct jurisdiction profile before generating a supply-chain claim.

4.2 Incident process versus time-bound evidence

CIRCIA was enacted to establish reporting obligations for covered cyber incidents and ransomware payments, but CISA's final implementing rule remained pending at the research cut-off.[7,8] The paper therefore does not claim that CIRCIA already mandates AI-specific fields.

The collision arises when an operator has an incident communication process but cannot rapidly assemble a trustworthy, system-specific package. An AI-related incident may require the system and model version, triggering event, tools and permissions used, decision-event record, human-intervention status, affected services, blast radius, vendor participation, containment and recovery. These are ITS-proposed operational fields, not a prediction of the final CIRCIA form.

The reporting clock can begin before the organisation has determined which system acted, whether the event was cyber-related, which vendor logs are available or whether an existing filing qualifies as substantially similar. A process document does not solve that evidentiary race. A CTL pre-positions evidence collectors, mappings, validation rules and output adapters before an incident occurs.

4.3 Human-oversight policy versus operational demonstration

The EU AI Act is binding, but telecom AI is not categorically high-risk. Annex III point 2 concerns AI used as a safety component in critical digital infrastructure; classification depends on intended purpose and the consequences of failure. Other telecom use cases may fall within other categories or outside the high-risk regime.[13] Actor roles also matter: provider and deployer duties are not interchangeable.

Where a system is in scope, a written human-oversight policy is not the same as operational evidence. A defensible record may need to show what the overseer could see, what authority the system possessed, how long intervention remained possible, whether override or safe stop worked, what action was actually executed and how the outcome was reviewed. Paper 2 develops the decision-event and earned-authority evidence needed for this purpose.[26]

The Digital Omnibus received Council final approval on 29 June 2026, with latest application dates of 2 December 2027 for stand-alone Annex III systems and 2 August 2028 for product-related Annex I systems, subject to Official Journal publication and entry into force.[14,15] Legislative status was verified at the research cut-off and should be rechecked before reliance.

4.4 Agent shutdown policy versus identity and authority evidence

CISA and international partners recommend mitigations for agent identity, permissions, tool access, orchestration and multi-agent behaviour.[12] The guidance is not binding law. NIST is also developing complementary agent identity, authorisation and interoperability work through its AI Agent Standards Initiative.[27]

The collision remains concrete. A disengagement policy cannot prove which agent invoked a tool, which credential it used, what scope was granted, whether another agent influenced the action, whether the credential was short-lived or whether authority changed during the event. A shutdown procedure is an intended outcome; identity, permission and event records are the evidence.

The CTL should therefore treat agent identity and authority records as reusable operational evidence. Different jurisdictions and contracts may attach different legal meaning, but the underlying event should remain reconstructable through vendor-neutral identifiers, signed records and time-synchronised logs.

5. Worked Examples: One Evidence Core, Different Endpoints

The following examples are illustrative rather than legal determinations. They show how common operational evidence can support different outputs without assuming that every jurisdiction imposes the same rights, classifications or reporting duties.

System and consequence	Common evidence core	Possible external outputs	Key caution
Vendor-embedded autonomous RAN remediation agent used by a regional operator	Agent identity; tool calls; authority level; network domain; change record; canary evidence; human intervention; vendor/model version.	Local telecom or cybersecurity report; group assurance package; procurement evidence; EU high-risk dossier where legally applicable.	The operator remains accountable for local network outcomes even where the vendor controls the model or orchestrator.
Subscriber fraud model that can suspend a number linked to mobile money and digital public services	Input categories; model/version; fraud signal; human review; action taken; duration; notification; correction and appeal outcome.	Data-protection record; telecom consumer complaint file; affected-person explanation and redress package; internal fairness review.	A common schema must not assume that every adverse decision attracts identical legal rights in every market.
AI-enabled customer-premises equipment or network component	Manufacturer; production location; authorisation status; firmware and AI component inventory; supplier attestation; update history; scan result.	FCC or procurement screening where applicable; local supply-chain assurance; insurer or customer due-diligence package.	The March 2026 FCC router action is origin- and approval-based, not a general prohibition on AI components. [9,10]

The examples make the paper's central design choice concrete: capture facts from authoritative systems once, preserve their lineage, and generate claims only through jurisdiction- and actor-specific profiles. Evidence reuse is efficient; legal equivalence must never be assumed.

6. What the Compliance Translation Layer Is - and Is Not

A Compliance Translation Layer is an evidence mediation architecture. It sits between operational governance systems and external endpoints that require proof. It does not replace legal analysis, regulator judgement, human assurance, incident response or the controls themselves.

GOVERNING PRINCIPLE

1. Map obligations to systems, actors and controls.
2. Collect evidence from authoritative operational sources.
3. Validate lineage, integrity, timing and sufficiency.
4. Generate the claim only where the evidence supports it.

An absent control cannot be translated into compliance.

The telecom analogy is a mediation platform. Network systems generate alarms, counters and events in internal formats. External consumers expect different schemas, timing, authentication and semantics. The operator does not rebuild the network-management system for every interface. It creates a controlled mediation layer while preserving the source record.

The analogy has a limit: compliance is not only data transformation. Applicability and sufficiency require professional judgement. The CTL may automate data collection, mapping checks, freshness tests and output generation, but a human legal or assurance function must own interpretations, exceptions, material qualifications and disputed claims.

7. Compliance Translation Layer Architecture

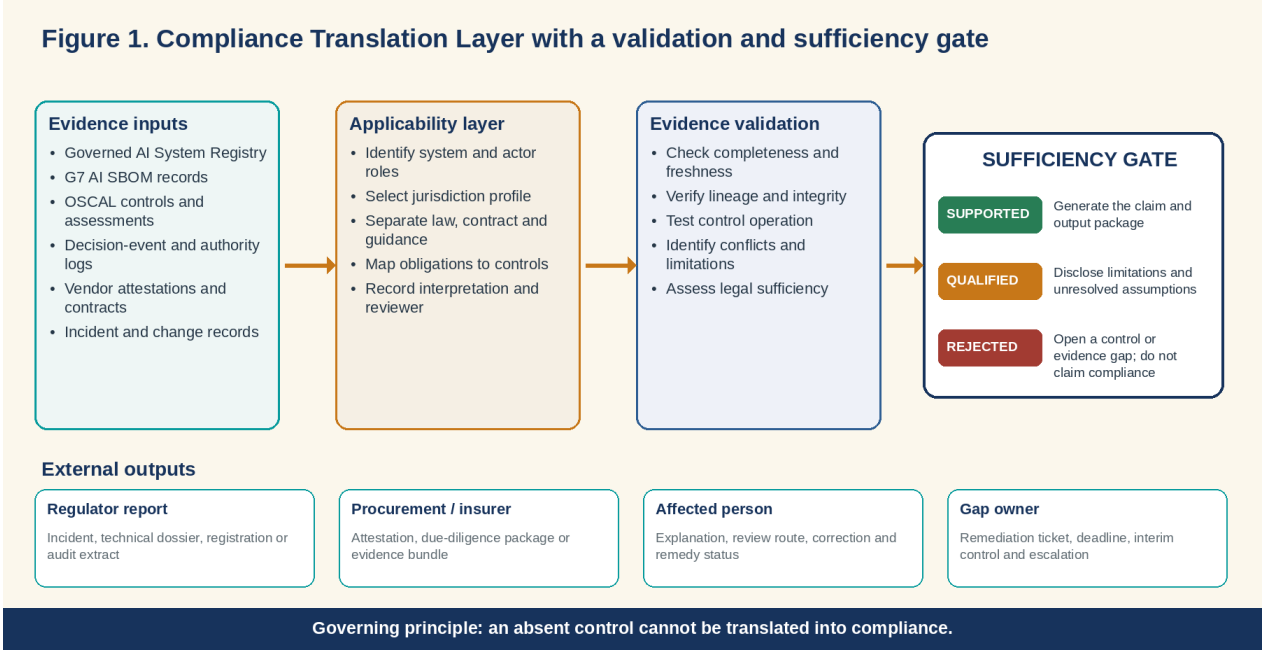

Figure 1. Compliance Translation Layer: the validation gate distinguishes supported, qualified and rejected claims.

7.1 Governed AI System Registry and actor registry

The layer begins with a reliable inventory. Each record should identify the system and version, intended purpose, deployment instances, jurisdictions, provider and deployer roles, data categories, autonomy level, vendor dependencies and links to Papers 1 and 2 evidence. The registry must include approved systems, embedded vendor AI, discovered shadow AI, unclassified systems and retired systems.

Paper 4 develops this inventory and discovery architecture as the Governed AI System Registry. The ITS papers are ordered analytically rather than by publication chronology: Paper 3 specifies the compliance evidence pressure; Paper 4 supplies the inventory foundation without which no system-specific claim can be reliable.

7.2 Obligation registry and jurisdiction profiles

Each obligation record should contain the source, legal status, jurisdiction, responsible actor, system trigger, effective date, deadline, output format, retention expectation, enforcement consequence, interpretation owner and verification date. A jurisdiction profile is a scoped view of the obligations relevant to a specific operator, system and transaction. It is not a hierarchy in which EU or US rules automatically override local law.

7.3 Control-to-obligation mapping

The mapping connects an external obligation to the operational control that supports it. It should identify the control owner, implementation location, evidence source, test frequency, residual assumptions and review status. OSCAL's Control Mapping Model can represent structured relationships among controls and requirements, including overlaps and gaps.[5] The CTL adds system applicability, actor role, operational evidence lineage and legal-sufficiency review.

7.4 Common evidence core: G7 AI SBOM, OSCAL and ITS extensions

The G7 AI SBOM guidance and OSCAL provide the clearest current precedents for a common evidence core. The AI SBOM describes what an AI system contains and depends upon across seven clusters: metadata, system-level properties, models, dataset properties, infrastructure, security properties and key performance indicators.[11] OSCAL represents controls, implementations, mappings, assessment plans and assessment results in machine-readable form.[4-6]

Neither instrument is designed to determine which telecom law applies to which actor, whether a particular item of evidence is legally sufficient, or what explanation and remedy an affected subscriber should receive. The ITS proposal

therefore adds obligation profiles, system-to-deployment correlation, decision and authority records, evidence lineage, legal validation, output adapters and redress interfaces.

Function	G7 AI SBOM	OSCAL	ITS CTL contribution
Describe system composition	Primary function	Limited	Consume, correlate and validate against live deployments
Model and dataset lineage	Strong	Extensible	Link provenance to obligations, decisions and claims
Represent controls and mappings	Limited	Primary function	Add actor, jurisdiction and system applicability
Represent assessment results	Limited	Strong	Evaluate evidence freshness, integrity and sufficiency
Determine legal applicability	Not a primary purpose	Not a primary purpose	Core human-governed CTL function
Generate regulator-specific artifacts	No	Supports automation	Core output-adapter function
Affected-person explanation and redress	No	Not a primary purpose	Dedicated CTL interface

This positioning converts the CTL from an isolated architecture into a defined contribution between two named international instruments. It also exposes a Global South risk: a G7-endorsed voluntary schema may become a procurement condition before smaller operators and regulators have the capacity to generate, validate or challenge the records it expects.

7.5 Evidence lineage, collectors and the sufficiency gate

Evidence should be collected from authoritative sources: the registry, decision-event records, identity and security logs, vendor-management systems, change platforms, component inventories, incident systems, human-oversight records and affected-person case files. Each item should carry a source, owner, collection time, system version, integrity mechanism, transformation history, retention period and confidentiality classification.

Where evidence comes from a vendor-controlled platform, the operator should distinguish vendor assertion from independently observed fact. Contracts should require exportable records, stable schemas, audit rights, update notification, time synchronisation and retention compatible with local obligations.

Before a claim is generated, the sufficiency gate should test completeness, freshness, integrity, actor applicability, approvals and consistency across sources. The result must be one of three states: supported, qualified or rejected. A qualified claim discloses limitations and unresolved assumptions. A rejected claim opens a control or evidence gap and must not be transformed into a compliance statement.

7.6 Regulatory, procurement and assurance output adapters

Adapters transform validated evidence into the form expected by an endpoint. Outputs may include regulatory reports, technical dossiers, registration records, procurement attestations, insurer evidence, audit extracts and executive dashboards. The adapter should preserve a link to the source evidence, validation result, interpretation owner and generation time.

7.7 Affected-person evidence and redress interface

Compliance evidence is not only sent upward to regulators. Where AI materially influences access to essential connectivity, mobile money, credit or public services, some evidence must be translated outward to the affected person. The CTL should support a package containing, where legally appropriate:

- notice that AI materially influenced the action;
- a decision and system identifier;
- the principal data categories and relevant decision factors at an appropriate level;
- human-review status and the responsible operator contact;
- a challenge, correction or appeal route;
- the outcome and remedy status;

- necessary restrictions protecting security, third-party rights and confidential information.

This interface is central to IGF alignment. It connects infrastructure governance to human rights, inclusion, meaningful participation and remedy. It also prevents a compliance architecture from becoming a one-way reporting system that serves institutions while remaining unintelligible to the person who experienced the decision.

7.8 Gap, change and remediation management

Every unresolved control or evidence gap should have an owner, severity, deadline and interim treatment. Changes in model version, purpose, data, vendor, jurisdiction, autonomy, procurement terms or law should trigger reassessment. The CTL should distinguish a missing control from missing evidence of a control: both prevent an unqualified claim, but they require different remediation.

8. Three Compliance-Architecture Questions

8.1 Which obligations apply to this system and actor?

The operator should enumerate binding law, licence conditions, contractual duties, procurement requirements and influential guidance separately. The answer should identify the jurisdiction, actor role, trigger, effective date, enforcement route and interpretation owner. "We follow the AI RMF" is not an applicability analysis.

8.2 Can every claim be traced to current, verifiable evidence?

For each claim, the operator should identify the implemented control, source system, evidence owner, collection time, integrity mechanism, validation result, material limitations and output artifact. A general policy or conceptual crosswalk cannot substitute for that chain.

8.3 Who owns the gap and the deadline?

No single function naturally owns the full chain. Security owns identity and logging; the CTO owns network platforms; the AI function owns model governance; Legal interprets obligations; consumer teams own complaints and remedy. The CTL requires cross-functional accountability for evidence production, not only risk discussion.

MINIMUM DIAGNOSTIC

1. Can the operator produce a system-specific obligation profile?
2. Can each external claim be traced to evidence that is current and appropriately reviewed?
3. Can every missing control or artifact be assigned to a named owner with a time-bound remediation plan?

You select a governance framework. A compliance obligation selects the system, actor and evidence it applies to.

9. Integration Across the ITS Issues Paper Series

Paper 3 is the connective layer in the ITS architecture. It depends on the inventory foundation developed in Paper 4, the operational-sovereignty analysis in Paper 1 and the delegated-authority evidence in Paper 2. The sequence is logical: a translation layer cannot produce reliable claims for an unknown system; it cannot establish sovereignty where data and processing control are absent; and it cannot demonstrate earned authority where consequential decisions are not reconstructable.

Layer	Question	Primary ITS contribution
Inventory	What AI systems, embedded models and agents actually exist?	Governed AI System Registry - Issues Paper No. 4.
Operational sovereignty	Who controls data, processing and reconstructable decisions?	The Unpriced Risk - Issues Paper No. 1.[25]
Delegated authority	What authority was earned, evidenced and kept safely reversible?	Beyond the Sandbox - Issues Paper No. 2.[26]
Compliance translation	Which obligations apply, what evidence supports each claim and who owns the gap?	The Compliance Collision - this paper.

10. Proportional Implementation and the Global South

Regulatory interoperability is a capacity and participation issue. The African Union's Continental AI Strategy calls for governance adapted to local context, capacity and development priorities.[16] ASEAN instruments emphasise practical and interoperable guidance while remaining largely voluntary.[17,18] India's data-protection obligations are maturing on a different timetable from AI-specific conformity regimes.[19,20]

Kenya illustrates the layered environment likely to face many operators. The Data Protection Act 2019 and its 2021 regulations establish rights and obligations for personal-data processing, access, objection, rectification and complaints.[21,22] The Communications Authority of Kenya maintains consumer redress mechanisms and a regulatory sandbox that expressly includes AI-driven ICT and telecommunications services.[23,24] A local operator may therefore need to distinguish national law and licence expectations from group policy, vendor contract terms, foreign-market requirements and voluntary G7 or NIST schemas.

A CTL designed only for large European or North American groups would reproduce the asymmetry it seeks to solve. Smaller operators and regulators may lack legal-engineering teams, continuous assurance tooling and bargaining power over global vendors. Yet subscribers should not receive weaker evidence merely because their operator has fewer resources.

Proportionality should apply to implementation cost, not to the reliability of high-consequence claims. This does not require a mid-size operator to replicate a multinational compliance platform. It may deploy fewer autonomous systems, start with one high-consequence use case, use shared assurance services, retain low-risk evidence for shorter periods or adopt narrower jurisdiction profiles. What it should not do is lower the reliability threshold for a claim that an emergency-affecting system is controlled, a subscriber restriction is explainable or an incident report is complete. The scope of deployment may be proportional to capacity; the truthfulness of the consequential claim cannot be.

A proportionate Global South implementation should combine:

1. A common evidence core using open, vendor-neutral records for systems, components, decisions, incidents, authority and human intervention.
2. Modular jurisdiction profiles maintained and reviewed locally, distinguishing law from guidance and imported contract expectations.
3. Shared capacity mechanisms such as regional test facilities, pooled expertise, open-source validation tools, regulator cooperation and South-South learning.
4. Vendor obligations to export evidence in stable formats and to support local audit, explanation and remedy requirements.
5. Affected-person interfaces that make evidence meaningful to communities, not only to compliance teams.

The objective is not to import the most complex foreign stack. It is to prevent international interoperability from becoming an evidentiary requirement that only large operators can afford, while preserving local legal meaning and meaningful avenues of challenge and remedy.

11. Recommendations

11.1 For telecom operators

- Build the obligation registry at system level, not only at enterprise-policy level.
- Adopt or request G7 AI SBOM-compatible component records and correlate them to deployed instances.
- Use OSCAL-compatible control and assessment representations where they reduce manual mapping.
- Require vendors to export identity, decision, update, incident and authority evidence in stable formats.
- Pilot the CTL on one high-consequence system and include affected-person outputs from the start.

11.2 For regulators and governments

- Publish machine-readable obligation profiles and minimum evidence fields where feasible.
- Distinguish binding requirements from influential guidance and voluntary assurance practices.
- Design reporting schemas that allow evidence reuse without erasing local legal differences.
- Support shared assurance capacity and vendor-neutral tools for smaller operators and emerging economies.
- Include civil society and affected-user perspectives in explanation, complaint and remedy requirements.

11.3 For NIST, CISA, G7 partners and standards bodies

- Extend conceptual crosswalks toward evidence, artifact and validation profiles.
- Map the G7 AI SBOM clusters to OSCAL controls, implementations and assessment results.
- Standardise vendor-neutral decision-event, agent-identity and delegated-authority records.
- Develop open jurisdiction-profile templates that can be locally governed and maintained.
- Support regional capacity building, interoperability pilots and South-South cooperation.

12. Limitations, Research Agenda and Call for Collaboration

This paper proposes an architecture, not a validated production platform. It has not yet been tested against a complete operator control environment, vendor evidence export or regulator submission process. Legal profiles will require continuous maintenance, and some evidence cannot be automated without risking false certainty.

Priority research questions are:

1. Which minimum evidence fields can be standardised across jurisdictions without erasing local legal meaning?
2. How should evidence sufficiency and freshness be scored for continuously changing AI systems?
3. What governance should apply to the legal rules and mappings encoded in a CTL?
4. How can operators prove evidence integrity when critical logs remain under vendor control?
5. What minimum affected-person package supports meaningful explanation, challenge and remedy?
6. Which shared regional facilities would most reduce compliance-engineering costs for smaller operators and regulators?

The Institute for Technology Stewardship invites collaboration from telecom operators, regulators, standards bodies, legal and assurance professionals, civil society organisations and researchers. ITS is particularly interested in pilot partners willing to test the CTL against one consequential telecom AI system and one live obligation profile.

Responses may be submitted through technologystewardship.org. Substantive feedback will inform the planned ITS synthesis paper on governing autonomous networks.

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 experience with policy research, developing practical frameworks that keep technology accountable to people, institutions and the public interest.

Author: Gaurav Arora, Founder, Institute for Technology Stewardship. He brings more than two decades of experience in telecommunications, network operations, service delivery and technology programme leadership.

Core belief: Ensuring technology remains in service of humanity - not the other way around.

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 Compliance Collision: Translating AI Governance into Regulator-Ready Evidence for Telecommunications. Institute for Technology Stewardship, Issues Paper No. 3, July 2026.

Disclaimer

This issues paper provides independent policy and operational analysis. It is not legal advice and does not represent the position of any operator, regulator, standards body or government. References to organisations and instruments are included for analytical purposes and do not imply endorsement or criticism.

References

1. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023. Source
2. National Institute of Standards and Technology. Concept Note: AI RMF Profile on Trustworthy AI in Critical Infrastructure, April 2026. Source
3. NIST AI Resource Center. Crosswalk Documents for the NIST AI RMF. Source
4. National Institute of Standards and Technology. Open Security Controls Assessment Language (OSCAL). Source
5. National Institute of Standards and Technology. OSCAL Control Mapping Model. Source
6. National Institute of Standards and Technology. OSCAL Assessment Results Model. Source
7. Cybersecurity and Infrastructure Security Agency. Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA) - implementation information. Source
8. Cybersecurity and Infrastructure Security Agency. CIRCIA FAQs and rulemaking resources. Source
9. Federal Communications Commission. FCC Updates Covered List to Include Foreign-Made Consumer Routers, 23 March 2026. Source
10. Federal Communications Commission. FAQs on Recent Updates to the Covered List Regarding Routers Produced in Foreign Countries, 12 May 2026. Source
11. Cybersecurity and Infrastructure Security Agency and G7 international partners. Software Bill of Materials for AI - Minimum Elements, 12 May 2026. Source
12. Cybersecurity and Infrastructure Security Agency and international partners. Careful Adoption of Agentic AI Services, May 2026. Source
13. European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Source
14. Council of the European Union. Artificial Intelligence: Council Gives Final Green Light to Simplify and Streamline Rules, 29 June 2026. Source
15. Council of the European Union. PE-CONS 30/26, adopted text of the Digital Omnibus on AI, June 2026. Source
16. African Union. Continental Artificial Intelligence Strategy, July 2024. Source
17. ASEAN Secretariat. ASEAN Guide on AI Governance and Ethics, 2024. Source
18. ASEAN Secretariat. Expanded ASEAN Guide on AI Governance and Ethics - Generative AI, 2025. Source
19. Government of India, Ministry of Electronics and Information Technology. Digital Personal Data Protection Act, 2023. Source
20. Government of India, Ministry of Electronics and Information Technology. Digital Personal Data Protection Rules, 2025. Source
21. Republic of Kenya. Data Protection Act, No. 24 of 2019. Source
22. Republic of Kenya. Data Protection (General) Regulations, 2021. Source
23. Communications Authority of Kenya. Consumer protection and redress mechanisms, March 2026. Source
24. Communications Authority of Kenya. Regulatory Sandbox, including AI-driven ICT and telecommunications services. Source
25. Arora, Gaurav. The Unpriced Risk, Issues Paper No. 1, July 2026. Source
26. Arora, Gaurav. Beyond the Sandbox, Issues Paper No. 2, July 2026. Source
27. National Institute of Standards and Technology. AI Agent Standards Initiative, 2026. Source