

GOVERNANCE WHITE PAPER

STORES

Sovereign Telecom Orchestration, Regulation and Enforcement System

> A reference architecture for AI-enabled, evidence-based and sovereign telecom regulatory oversight: from network monitoring to trusted, cross-sector infrastructure stewardship

Raghid W. Charara | Khaled M. Chatila

Version 1.0 | October 2026

About this white paper

Audience. Telecom regulators, policymakers and ministries responsible for digital infrastructure, together with operators, infrastructure owners and technology partners who work with them.

Purpose. To present a reference architecture, a governance model and a practical roadmap for modernising national telecom oversight with AI, verifiable evidence and cross-sector coordination, while protecting data sovereignty, competition and citizens' rights.

Scope. STORES is a technology-neutral reference model. It does not describe a specific product or deployment, and it is designed to complement, not replace, operator systems and existing regulatory tools.

Evidence. Market statistics are taken from the sources listed in the References and reflect the period stated in each source. Figures and scenarios labelled illustrative use hypothetical values to explain concepts.

Reading guide by audience

- **Ministers and board members (about 10 pages):** Executive Summary; Section 10 (costs); Section 12 (recommendations)
- **Regulator leadership:** add Sections 1 to 3 (the case for change), 6 (AI governance) and 11 (roadmap)
- **Technical and architecture teams:** Sections 4 and 5 (architecture, evidence model, threat model, NROC) and Section 8 (standards)
- **Legal and policy teams:** Section 7 (legal design, rights and due process) and Appendix A (model provisions)
- **Operators:** Sections 4.1 and 4.2 (nodes and evidence), 7.4 (due process), 10.3 (operator burden) and Appendix A.7 (node agreement)
- **Self-assessment:** Appendix B: where your organisation stands today

Suggested citation

Charara, R. W. and Chatila, K. M. (2026). STORES: Sovereign Telecom Orchestration, Regulation and Enforcement System. Governance white paper, version 1.0, October 2026.

Disclosure of interests

Raghid W. Charara is Vice President of Analitico International, which provides consulting and technology services to telecom regulators and operators, including device-identity (CEIR) and quality-of-service solutions. Khaled M. Chatila is Chief Executive Officer of UNIEQUIP, which delivers technology and turnkey projects, including in telecommunications. Both companies may compete for projects related to the topics discussed. STORES is a vendor-neutral reference model, not a product of either company: every component described can be built from open standards and sourced from multiple suppliers. The authors have developed an implementation based on this reference model; this paper presents the model and its research basis and makes no product claims.

Disclaimer

The views expressed in this paper are solely those of the authors and do not necessarily reflect the positions of any organisations with which they are affiliated. This publication is intended for informational and discussion purposes only and does not constitute legal, regulatory, financial or investment advice. The authors make no representation or warranty as to the completeness or accuracy of the information presented.

© 2026 Raghid W. Charara and Khaled M. Chatila. All rights reserved.

Contents

Executive Summary 5

1. The Paradigm Shift in National Regulatory Landscapes 7

2. Why Existing Telecom Governance Models Are No Longer Sufficient 9

2.1 The building blocks are already in use 9

3. The Regulatory Visibility Gap and the QoS/QoE Paradox 11

3.1 Network slicing and algorithmic neutrality 12

3.2 Information asymmetry and the trust deficit 12

4. Architectural Foundation: The STORES Reference Architecture 13

4.1 Layer 1: The sovereign data fabric 14

4.2 Layer 2: Trust and evidence 15

4.3 Layer 3: Intelligence and specification-aware AI 17

4.4 Layer 4: The NROC and controlled orchestration 18

4.5 Threat model 18

5. The National Regulatory Operations Centre (NROC) 20

5.1 Operational scenario: national 5G resilience during a regional power disruption 20

5.2 Power grid integration and energy resilience 21

5.3 Fibre plant: precise fault localisation and sensing 21

5.4 The National Infrastructure Digital Twin (NIDT) 22

6. AI Governance, Ethics and Human Oversight 23

6.1 Explainability and accountability 23

6.2 The STORES regulatory autonomy model 23

6.3 Security: defending the agentic regulator 24

7. Legal Design, Rights and Due Process 25

7.1 Legal basis, mandate and jurisdiction 25

7.2 Classification of the regulator's own AI 25

7.3 Data protection 26

7.4 Due process and admissibility of evidence 26

7.5 Net neutrality, emergency measures and liability 27

7.6 Competition safeguards 27

7.7 Preventing function creep 27

8. Alignment with Global Standards 28

9. Global Applicability Across Regulatory Environments 29

10. Costs, Resourcing and Funding 30

10.1 Classifying a market 31

10.2 People and skills 33

10.3 Operator burden and offsets 33

10.4 Funding models 33

10.5 Assessing value 33

11. Implementation Roadmap 34

11.1 Key risks and mitigations 35

12. Strategic Recommendations 36

Recommendation 1: Adopt standards-native design and open APIs 36

© 2026 R. W. Charara and K. M. Chatila

Recommendation 2: Institutionalise the National Infrastructure Digital Twin....36
Recommendation 3: Transition to intent-based performance standards, with published methods....36
Recommendation 4: Treat sovereignty as strategic control, not isolation....36
Recommendation 5: Invest in human judgment and AI literacy....36
Recommendation 6: Modernise the legal basis for data, evidence and AI....36

13. Conclusions: A New Era of National Infrastructure Stewardship....37

Appendix A. Illustrative Model Provisions....38

Appendix B. STORES Readiness Self-Assessment....40

Appendix C. Economic Modelling with Input-Output Methods....41

C.1 The Leontief input-output model....41
C.2 The dynamic extension....42

About the Authors....43

Glossary of Terms....44

References....45

© 2026 R. W. Charara and K. M. Chatila

Executive Summary

Connectivity has become the substrate on which payments, health care, logistics, energy and government services run. Yet most national regulators still oversee it with tools designed for a simpler era: periodic operator reports, sample-based drive tests and after-the-fact audits. As 5G Standalone (5G-SA), network slicing, software-defined infrastructure and AI-driven operations take hold, the gap between what regulators can see and what citizens and businesses actually experience is widening.

This white paper proposes **STORES, the Sovereign Telecom Orchestration, Regulation and Enforcement System**: a vendor-neutral, standards-aligned reference architecture that helps regulators move, in controlled steps, from monitoring to regulatory intelligence and, where the law allows, to controlled orchestration. STORES does not replace operator network operations centres (NOCs), OSS/BSS platforms or existing regulatory tools. It connects and extends them, and it is designed so that every finding can be verified and contested.

STORES in five ideas

- **See the experience, not only the network.** Correlate network KPIs with Quality of Experience (QoE) and slice-level performance, so that compliance reflects what users actually receive.
- **Keep data sovereign.** A federated data fabric lets national indicators and models learn from operator data without raw subscriber records leaving operator perimeters or national borders.
- **Make evidence verifiable and contestable.** Measurements used as evidence are signed at source, time-stamped and anchored in an append-only log, and the operator concerned can inspect and challenge them.
- **Use AI with bounded autonomy.** Specification-aware AI assistants test network behaviour against licence obligations under a five-level autonomy model (R0 to R4): only non-binding actions are automated, and every action with legal effect is approved by a named official.
- **Govern infrastructure as a system.** A National Regulatory Operations Centre (NROC) and a National Infrastructure Digital Twin (NIDT) coordinate, through formal agreements, with power, fibre and emergency authorities to anticipate cascading failures.

Why now. Ookla's Q1 2026 measurements show average peak-hour download speeds falling by 27% on 5G and 32% on 4G across ten high-5G European markets, and median speed falling by 66% in Spain[1], even where headline indicators look healthy. Data-centre electricity demand is projected to more than double by 2030[2], tightening the coupling between energy and digital infrastructure. Telecom revenues are forecast to grow below inflation[3], so oversight must improve without imposing disproportionate cost on the sector. Regulators such as Arcep, Ofcom and the FCC already use crowdsourced measurement and open challenge processes[4,5,6]; STORES integrates these established building blocks into a coherent, lawful, regulator-side architecture.

© 2026 R. W. Charara and K. M. Chatila

What this paper adds

- **A regulator-side reference architecture.** ETSI ZSM and TM Forum Autonomous Networks describe how operators automate their own networks; STORES describes how a regulator supervises automated networks.
- **The Regulatory Performance Paradox**, a precise statement of why compliant KPIs can coexist with poor user experience (Section 1).
- **The R0 to R4 regulatory autonomy ladder**, which ties each type of AI-assisted action to its legal effect and the human control it requires (Section 6).
- **The two-path evidence model**, which separates exact, contestable evidence from privacy-protected analytics (Section 4.2).
- **A legal-design toolkit and cost model**: rights safeguards, model licence provisions and indicative five-year costs by market size (Section 7, Section 10 and Appendix A).

What to do. Section 12 sets out six recommendations: adopt standards-native design and open APIs; institutionalise a national infrastructure digital twin; move toward intent-based, outcome-oriented performance standards with published methods; treat sovereignty as strategic control rather than isolation; invest in human judgment and AI literacy; and modernise the legal basis for data access, digital evidence and AI-assisted supervision. Indicatively, a five-year programme costs about USD 5 to 19 million in a small market and USD 9 to 37 million in a medium one (Section 10). The phased roadmap in Section 11 delivers visible value within nine months and reaches a fully operational NROC in about four years.

© 2026 R. W. Charara and K. M. Chatila

1. The Paradigm Shift in National Regulatory Landscapes

The global telecommunications landscape is undergoing a structural transformation as connectivity evolves from a standalone utility into a foundational layer of national digital infrastructure. The drivers are hyper-connectivity, the maturation of 5G-SA architectures, increasingly software-defined and AI-enabled networks and, in the longer term, the evolution toward 6G. Traditional oversight, characterised by periodic reporting, reactive monitoring, manual auditing and static compliance dashboards, is increasingly challenged by the scale, complexity and pace of modern networks.

Historically, oversight focused on the technical health of the network layer, measured through Quality of Service (QoS) parameters such as availability, packet loss, jitter and latency[7]. 5G-SA capabilities such as network slicing, where shared physical infrastructure supports multiple logical networks with differentiated performance requirements, reduce the sufficiency of aggregate indicators when they are used in isolation.

This paper calls the resulting discrepancy the **Regulatory Performance Paradox**: operators can demonstrate compliance with defined technical Key Performance Indicators (KPIs) while consumers and enterprises continue to experience service-quality problems at the application level. The paradox is not evidence of bad faith; it is a measurement problem. Averages hide peaks, network counters hide application behaviour, and national aggregates hide local and temporal gaps.



Figure 1. The evolution of regulatory oversight. STORES builds on existing reporting and monitoring and extends them toward service assurance, regulatory intelligence and controlled orchestration.

STORES is conceived from the perspective of the regulator rather than the network operator. It provides a unified architectural model through which regulators can integrate network and service observability, QoS/QoE intelligence, trusted evidence, AI-assisted analysis, infrastructure resilience and controlled orchestration. Table 1 summarises the shift.

Table 1. From the conventional regulatory model to the STORES paradigm

Dimension	Conventional regulatory model	STORES paradigm	Intended regulatory and national value
Governance mode	Reactive oversight; periodic and manual audit	Continuous intelligence, predictive oversight and controlled orchestration	Faster detection, assessment and response
Visibility scope	Mainly network-centric technical KPIs	Integrated network, service and QoS/QoE observability, per slice, place and time	Closer link between technical performance and user experience

© 2026 R. W. Charara and K. M. Chatila

Dimension	Conventional regulatory model	STORES paradigm	Intended regulatory and national value
Data architecture	Centralised reporting and operator-specific silos	Sovereign, interoperable, federated multi-operator data fabric	Cross-operator intelligence with data protection and residency
Trust and evidence	Operator self-reporting; post-event verification	Signed, time-stamped evidence that operators can inspect and contest	Stronger evidence integrity; less reconciliation friction
Infrastructure perspective	Telecom considered largely in isolation	Telecom coordinated with power, fibre, cloud and emergency authorities	Earlier warning of cross-sector cascading failures
Decision intelligence	Manual analysis and rule-based assessment	AI-assisted analytics and decision support with bounded autonomy	Earlier identification of emerging risks; better-targeted intervention
Economic and strategic planning	Compliance-focused assessment	Digital-twin and cost-benefit based policy analysis	Stronger evidence base for spectrum, coverage and investment policy

The shift becomes more pressing as networks serve not only people but also machine-to-machine services, autonomous systems and AI agents that initiate and consume network services through programmable interfaces such as GSMA Open Gateway and CAMARA APIs[8]. This paper refers to this emerging dimension as the **Agentic Experience (AX)**: the quality of service as experienced by software agents rather than by humans. AX sharpens the regulatory question. If an AI agent requests a guaranteed-latency session for an industrial-control or tele-health application, who verifies that it was delivered as promised?

Within STORES, specification-aware AI assists regulators by interpreting technical standards, licence conditions and regulatory policies, correlating them with observed network conditions, and identifying potential deviations for assessment, recommendation or authorised action. The framework does not assume autonomous enforcement. Autonomy grows progressively, according to legal authority, operational risk and explicit human-oversight requirements (Section 6).

© 2026 R. W. Charara and K. M. Chatila

2. Why Existing Telecom Governance Models Are No Longer Sufficient

From national telecom monitoring to sovereign regulatory orchestration

Existing governance models were designed for comparatively static, human-operated networks and for sector-by-sector oversight. OSS/BSS platforms, operator NOCs and periodic compliance audits remain essential, but on their own they struggle to govern software-defined, multi-slice infrastructure that is reconfigured in seconds and increasingly operated by AI. Five structural gaps stand out.

- **Visibility gap.** Regulators see reported network KPIs, not the end-to-end service outcome per slice, location and time of day.
- **Trust and information asymmetry.** Most evidence is generated, processed and reported by the regulated entity, which lengthens verification and dispute cycles.
- **Speed gap.** Quarterly or annual reporting cannot keep pace with networks that software and AI reconfigure continuously.
- **Silo gap.** Telecom failures increasingly originate in adjacent systems, such as power, fibre routes, cloud and data centres, which sit under other authorities.
- **Accountability gap.** As operators adopt AI-driven automation, regulators need ways to assess automated decisions, not only network outcomes.

STORES addresses these gaps with continuous, independent observability; intent-based performance standards; a federated and verifiable evidence base; bounded AI assistance; and cross-sector coordination through the NROC.

2.1 The building blocks are already in use

Most foundational capabilities used by STORES are already deployed individually in regulatory, telecom or critical-infrastructure settings. STORES' contribution is to integrate them into a coherent regulator-side architecture. Their integration at national scale, the specification-aware regulatory AI and broad NROC coordination are proposed here as an architecture, not presented as an operating model already demonstrated by a regulator. Table 2 gives examples of the building blocks in use.

© 2026 R. W. Charara and K. M. Chatila

Table 2. Regulators already using STORES building blocks

Regulator	Practice	STORES building block	Lesson
Arcep (France)	"Mon réseau mobile" combines operator coverage maps, annual field audits and, since 2022, about 200,000 crowdsourced measurements, roughly 100 times the volume of the annual audit; Arcep publishes the methodological limits[4]	Independent QoE measurement; data-driven regulation	Crowdsourcing scales coverage, but its limits must be published and controlled
Ofcom (UK)	Mobile Matters reports latency, video-streaming success and speeds by network and local area using crowdsourced device data[5]	Experience-based indicators by place and time	Experience metrics can be reported routinely and transparently
BEREC (EU)	Commissioned an open-source net-neutrality measurement tool, reference system and portal for consistent QoS measurement across Europe[9]	Common, standards-based measurement method	Shared open tools improve comparability and reduce lock-in
FCC (US)	Mobile challenge process: consumers and others challenge coverage claims with app-based speed tests against defined thresholds; providers have 60 days to concede or rebut[6]	Contestable evidence with due process	Clear thresholds and response rights make data-driven findings fair
DoT (India)	CEIR device-identity register enforced across all operators (see box below)[10]	Lawful cross-operator orchestration	Orchestration works with a mandate, standard interfaces and safeguards

Illustration: cross-operator orchestration already works at national scale

National device-identity registers, known as Central Equipment Identity Registers (CEIR), show how a regulator can move from passive monitoring to active, lawful orchestration across all operators. A CEIR links the device identifiers (IMEI) seen on every network with a national register, so that devices reported lost or stolen, and identifiers that are duplicated or invalid, can be blocked across all networks under a clear legal mandate.

India's CEIR, delivered through the Sanchar Saathi platform, shows the scale such systems reach. By December 2025 more than 4.2 million lost or stolen handsets had been blocked, about 2.6 million traced and more than 720,000 returned to their owners, while over 14 million fraudulent mobile connections had been disconnected following citizen reports[10].

The lesson for STORES: cross-operator orchestration succeeds when four conditions hold: a clear legal mandate, standard interfaces to every operator, auditable evidence, and transparent safeguards for citizens' rights. STORES generalises these conditions from device identity to service quality, resilience and AI governance. Unlike a CEIR, which acts on individual devices under its own specific legal mandate, STORES processes network and service performance data and takes no action against individual users or devices.

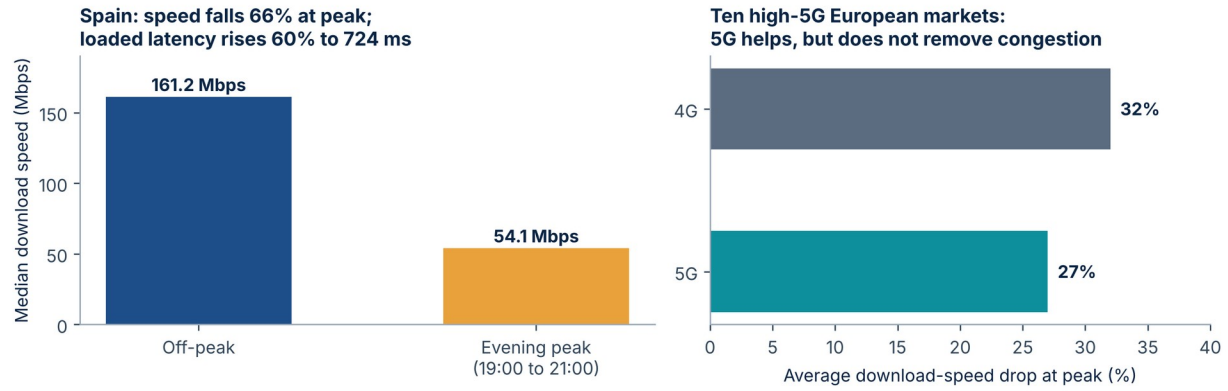
© 2026 R. W. Charara and K. M. Chatila

3. The Regulatory Visibility Gap and the QoS/QoE Paradox

A central challenge for modern telecom regulators is the gap between reported technical performance and the quality that users experience. Mobile broadband has largely operated on a best-effort basis, and regulatory monitoring has commonly focused on network-level indicators such as availability, throughput, latency, packet loss and coverage. As 5G-SA enables differentiated services, network slicing and demanding industrial, enterprise and consumer applications, these indicators remain essential but are insufficient when assessed in isolation.

The result is a regulatory visibility gap between network condition and service outcome. Network Management Systems (NMS) and Operations Support Systems (OSS) provide critical information on network health, faults, capacity and performance, but they do not give regulators an end-to-end view of how network conditions translate into application and user experience. European regulators have already adopted a common position on informing consumers about mobile coverage[11]; STORES extends that logic to multi-layer service intelligence, correlating network KPIs with service-level and QoE indicators where technically feasible and legally appropriate.

Recent European measurements illustrate the distinction. Ookla's analysis of Q1 2026 mobile performance identified a recurring evening trough between 19:00 and 21:00 local time. Across ten markets with significant 5G adoption, average download speed fell by about 32% on 4G and 27% on 5G at peak, and 5G loaded latency at peak was 12% to 44% lower than 4G in every market tested: 5G improves performance under load but does not eliminate congestion. Outcomes varied widely. In Spain, median download speed fell from 161.2 Mbps off-peak to 54.1 Mbps at peak while loaded latency rose by 60% to 724 ms; Luxembourg, by contrast, showed almost no peak-hour degradation[1]. A headline national average would conceal both the problem and its geography.



Source: Ookla, Q1 2026 European mobile peak-hour congestion analysis, as reported by ISPreview (May 2026).

Figure 2. Peak-hour congestion in European mobile networks, Q1 2026. Headline speeds can hide a large evening degradation.

© 2026 R. W. Charara and K. M. Chatila

Table 3. From network QoS to experience and intent

Metric category	Traditional focus (QoS)	Modern regulatory focus (QoE)	STORES target (intent)
Unit of measure	Throughput, latency, packet loss, jitter	Video quality (MOS)[12], page-load time, call set-up success	Intent fulfilment: share of sessions meeting the service's declared requirement
Visibility scope	Network layer (Layer 3)	Application layer (Layer 7)	End-to-end journey across radio, transport, core and application, per slice
Data source	Operator OSS/NMS	Drive tests, app-based and crowdsourced measurement	Federated operator nodes, independent probes and crowdsourcing, cross-validated
Cadence	Periodic, often manual	Near-real-time dashboards	Continuous, with statistically designed sampling
Typical trigger	Network alarms	Loaded latency, peak-hour degradation	Deviation from intent or obligation, with probable root cause

3.1 Network slicing and algorithmic neutrality

Slicing introduces multi-tenant complexity. A regulator may need to verify that a massive-IoT slice serving a smart grid meets its contracted reliability target, for example 99.9%, without being degraded by a consumer mobile broadband slice, and that premium slices are offered on fair, transparent and non-discriminatory terms. This raises a new supervisory question, which this paper calls **algorithmic net neutrality**: are resource-allocation algorithms isolating slices as promised, and are any services systematically starved of resources? Answering it requires slice-aware observability that draws on standardised analytics such as the 3GPP Network Data Analytics Function (NWDAF)[13] and intent-based slice management signalling such as ITU-T Q.5038[14], cross-checked by independent measurement. Only aggregated, slice-level outputs and evidence records without personal data leave the operator; per-subscriber analytics stay inside its perimeter.

3.2 Information asymmetry and the trust deficit

The most persistent pain point is the asymmetry between the regulator and the regulated. Regulators typically depend on data generated, interpreted and reported by operators, which complicates verification and invites dispute. Reconciliation of measurement disputes, inter-operator settlement and audit findings can take weeks or months, consuming technical and legal capacity on both sides. The industry itself has explored shared ledgers to compress reconciliation in wholesale roaming[15]. The asymmetry also runs the other way: operators often cannot see or test the measurements a regulator relies on. A sovereign framework should therefore provide a common, tamper-evident evidence base that both sides can inspect, re-run and contest.

© 2026 R. W. Charara and K. M. Chatila

4. Architectural Foundation: The STORES Reference Architecture

STORES is a modular, layered reference architecture. It separates data, trust, intelligence and orchestration so that each can be procured, governed and evolved independently, and it relies on open standards to avoid vendor lock-in. Two cross-cutting planes apply to every layer: governance with human oversight, and Zero Trust security with privacy by design.

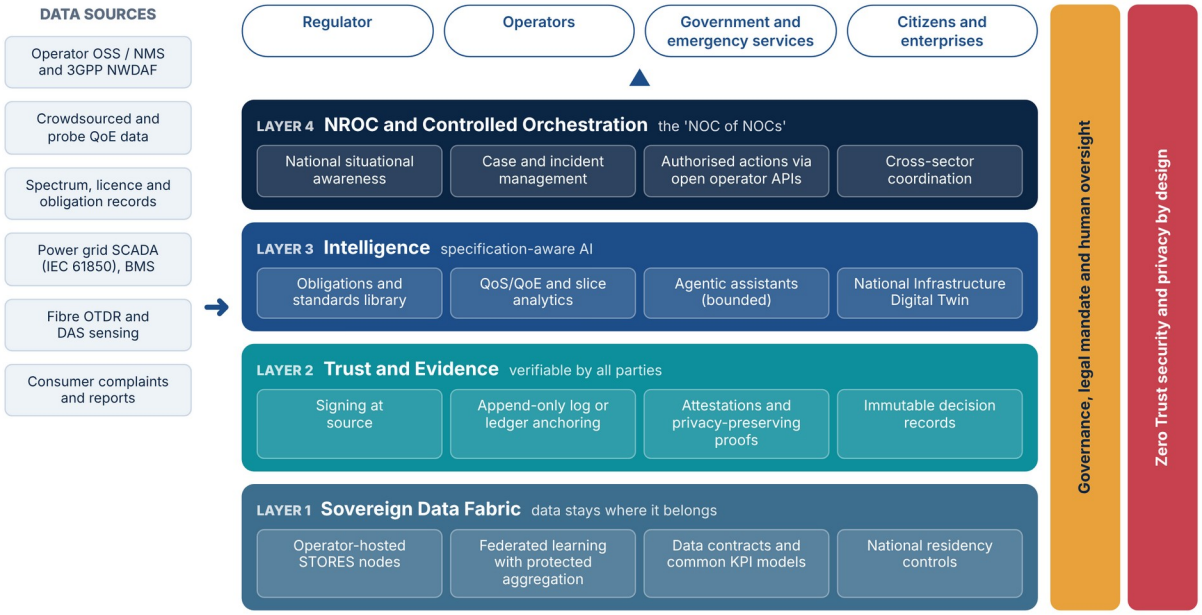


Figure 3. The STORES reference architecture: four layers, two cross-cutting planes, and data sources that remain under their owners' control.

Six design principles guide every implementation choice:

- **Regulator-centric.** Built around legal obligations and public-interest outcomes, not network optimisation.
- **Complement, do not replace.** Interoperate with operator NOCs, OSS/BSS and existing regulatory systems through standard interfaces.
- **Sovereign by design.** Data residency, local processing and national control of keys, models and evidence.
- **Evidence first, and contestable.** Every finding is traceable to signed, time-stamped data that the affected party can inspect.
- **Bounded autonomy.** Only non-binding, reversible actions are automated; anything with legal effect is approved by a named official.
- **Technology-generation neutral.** Applicable to 4G, 5G-SA, fixed and satellite networks today, and extensible to 6G.

© 2026 R. W. Charara and K. M. Chatila

What STORES is not

- **Not a surveillance system.** It processes network and service performance data, not communications content, and it does not track individuals. Purpose limitation is set in law (Section 7).
- **Not a remote control for operator networks.** STORES never writes to operator network elements; any action passes through the operator's own systems.
- **Not a single product or vendor platform.** It is a reference model that can be assembled from multiple suppliers and open-source components.
- **Not a replacement for regulatory judgment.** AI informs decisions; officials take them, and operators retain full rights of challenge and appeal.

4.1 Layer 1: The sovereign data fabric

To give the regulator a national view without compromising operator confidentiality or data residency, STORES uses a federated data fabric. Each operator hosts a **STORES node**: a sandboxed, read-only appliance that sits inside the operator's own perimeter and under its operational control. The node reads only from agreed data exports, maps them to a common KPI data model defined in data contracts, computes standard indicators and, where needed, trains shared models locally. It never writes to network elements, runs only signed code approved through a joint change board, proves its software state through remote attestation, and can be isolated by the operator in an emergency; every isolation is logged, reported to the regulator and covered in the meantime by independent probes. Evidence-signing keys are generated inside a regulator-controlled hardware security module housed in the node, under the regulator's national public-key infrastructure, and cannot be extracted by the operator. Even so, node outputs are treated as operator-originated data: they are cross-validated against independent probes and crowdsourcing, which remain the primary source of evidence for enforcement. Funding, key custody and liability for the node are set out in the data-access and node agreement (Appendix A.7).

For most regulatory indicators, **federated analytics** is enough: the node computes agreed aggregates, such as peak-hour throughput percentiles per district and slice, applies minimum-size thresholds so that no small group of users can be singled out, and shares only those results. Machine learning is reserved for problems that need it, such as congestion forecasting or detecting fraud patterns such as SIM-box activity; fraud models run inside each operator, and only aggregate indicators are shared.

For those problems, STORES applies **federated learning (FL)**: only model updates, never raw records, leave the operator. FL is not private by default; research has shown that shared gradients can, under some conditions, be inverted to reconstruct training data[16]. Most markets also have only three or four operators, so classical secure aggregation[17] hides little: each participant can infer much about the others' contributions from the combined result. STORES therefore combines three protections:

- **Subscriber-level differential privacy inside each operator.** Training bounds the combined contribution of all records belonging to any one subscriber (user-level differential privacy); updates are clipped and calibrated noise is added before anything leaves the node[18,19]. This places a mathematically bounded limit, set by the privacy budget, on what anyone can infer about an individual subscriber from the shared updates. The privacy budget (ε, δ) is set per use case by the data-protection governance board and published.

© 2026 R. W. Charara and K. M. Chatila

- **Protected aggregation for commercial confidentiality.** Where an operator's update could reveal commercially sensitive information, aggregation runs inside an attested confidential-computing enclave (a trusted execution environment)[20], so that neither the regulator nor other operators see individual updates.
- **Robust aggregation and output rules.** Updates are validated and combined with aggregation methods that resist poisoning by a faulty or dishonest participant. National models and indicators are released only after leakage checks, and never at a granularity that would expose a single operator's commercially sensitive position.

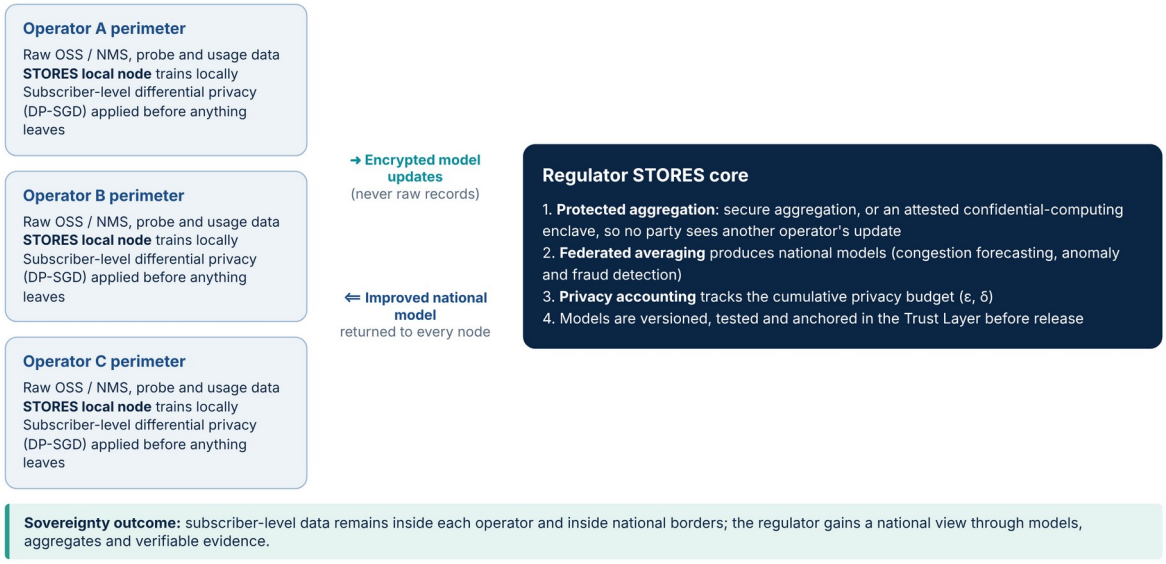


Figure 4. The federated learning loop. Raw data stays within each operator; only protected model updates travel, and no party sees another operator's individual update.

4.2 Layer 2: Trust and evidence

The Trust Layer turns measurements into evidence. Its central design choice is to separate two kinds of data flow (Figure 5). The **evidence path** carries exact, full-fidelity records that can support findings and disputes. The **analytics path** carries privacy-protected aggregates and models that support early warning and policy analysis. Differential-privacy noise is never applied to evidence, and analytics outputs are never used on their own as legal evidence: they can trigger an investigation, which then relies on the evidence path.

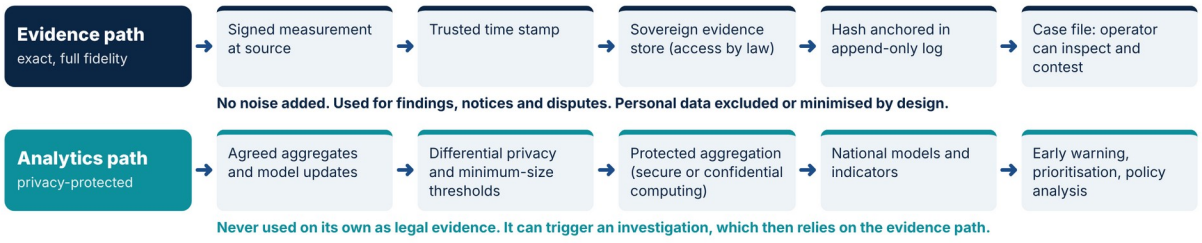


Figure 5. The two-path evidence model: exact, contestable evidence is kept separate from privacy-protected analytics.

© 2026 R. W. Charara and K. M. Chatila

On the evidence path, records are **signed at source** by the probe, crowdsourcing agent or operator node, and receive a **trusted time stamp** following RFC 3161[21] or, where available, a qualified electronic time stamp, which in the EU benefits from a legal presumption of accuracy and integrity[22]. Full records are kept in a sovereign evidence store, and only their cryptographic hashes and metadata are **anchored** in a tamper-evident log, so that any later alteration or deletion is detectable. Measurement equipment used for evidence is calibrated by laboratories accredited to ISO/IEC 17025[23]. Evidence records carry performance values, time and coarse location, such as a map grid cell; device identifiers and precise location are removed at source.

The trust chain must itself be sovereign. Time stamps should come from a national or nationally accredited time-stamping authority, signing keys should sit in nationally controlled hardware security modules, and enclave attestation services and cloud providers should be chosen with regard to jurisdiction, including exposure to foreign laws that permit extraterritorial access to data.

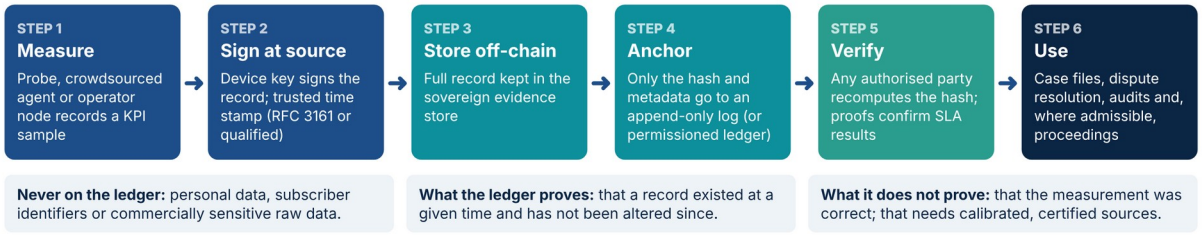


Figure 6. The STORES evidence chain: from a signed measurement to evidence that every authorised party can verify.

Is a blockchain needed? Not by default. Table 4 compares the options. For most regulators, signed records with trusted time stamps and an append-only transparency log, of the kind used to make web certificates publicly auditable[24], provide the required integrity at lower cost and complexity. A permissioned ledger such as Hyperledger Fabric becomes worthwhile only when several parties must write to and co-govern the same record, for example in inter-operator settlement or a multi-agency NROC. Attestations and privacy-preserving proofs, including zero-knowledge proofs where mature, can demonstrate a property such as "this slice met its latency target in 99.9% of intervals in May" without disclosing the underlying data; they are cryptographic techniques that work with any of the three options.

Table 4. Options for anchoring evidence

Option	What it provides	When it is sufficient	Relative cost and complexity
Signed records with trusted time stamps	Integrity and time proof for each record; legally recognised time stamps	Regulator is sole custodian; low dispute volume	Low
Append-only transparency log (STORES default)	Verifiable, append-only history; detects deletion or reordering; operators can audit it	Most regulators; operators need to verify that nothing was changed	Low to medium
Permissioned ledger	Shared governance and multi-party write; workflow automation	Several parties must write and co-govern records, such as settlement or multi-agency operations	Medium to high

© 2026 R. W. Charara and K. M. Chatila

Workflow automation serves procedure, not punishment. Once an authorised official has issued a notice, the system tracks the operator's response deadline, sends reminders and assembles the case file. It never issues notices or penalties by itself (Section 6). Personal data is excluded from the anchoring log by design, so immutability does not conflict with data-protection rights such as erasure (Section 7.3).

Table 5. Traditional audit compared with the STORES evidence model

Feature	Traditional audit	STORES evidence model	Benefit for regulator and operators
Data integrity	Manual reconciliation of operator reports	Signed at source; time-stamped; hash anchored	Disputes shift from data authenticity to interpretation
Coverage	Sample-based tests and periodic returns	Continuous measurement with statistically designed sampling	Earlier detection on a sounder statistical basis
Speed	Months or quarters	Near-real-time for anchored indicators	Faster, proportionate intervention
Contestability	Operator rarely sees the regulator's raw data	Operator can inspect records, methods and calibration certificates	Fairer process; fewer appeals on procedural grounds
Transparency	Limited audit trail	Complete, time-stamped trail of data, analysis and decisions	Stronger accountability and public legitimacy

4.3 Layer 3: Intelligence and specification-aware AI

The Intelligence Layer turns evidence into insight. Its foundation is an **obligations and standards library**: licence conditions, QoS regulations, coverage obligations and relevant standards expressed as machine-readable rules ("policy as code"), each linked to the KPI definitions and measurement methods on which it relies. On top of this library, analytics and AI models detect anomalies, forecast congestion, estimate QoE from network data and locate probable root causes. They build on standards such as ITU-T Y.3192, which specifies customer-oriented network quality optimisation with AI, including the case where users complain despite normal technical indicators[25].

STORES then uses **agentic AI** in a bounded form. Unlike a chatbot that answers questions, an agent can plan a sequence of steps and call tools: query the data fabric, run a digital-twin simulation, retrieve the relevant licence clause and draft a finding. A specification-aware agent knows which obligation applies to which service, area and period, and can explain why it believes a deviation occurred.

For example, if analytics detect a persistent degradation of video-streaming QoE in one governorate, an agent can correlate it with cell congestion, transport capacity and slice configuration, identify the licence obligations involved, assemble an evidence pack from the evidence path and draft a request for information. The draft goes to a named official, who decides whether to issue it (autonomy level R3). Penalties are decided collectively under due process (R4).

© 2026 R. W. Charara and K. M. Chatila

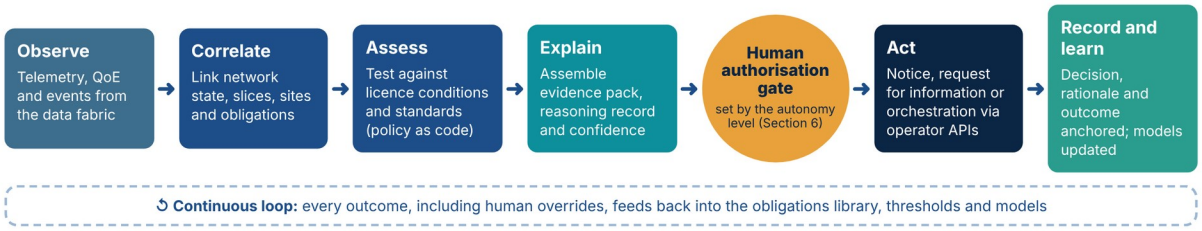


Figure 7. The specification-aware compliance loop, with a human authorisation gate set by the autonomy level of each action.

This mirrors the industry's own direction. Operators are moving toward TM Forum Level 4 ("highly autonomous") networks, in which systems take intent-driven, predictive and closed-loop decisions within defined domains[26], and ETSI ZSM specifies architectures for closed-loop automation[27]. As operators automate, regulators need comparable capability to supervise automated networks and to ask operators for evidence of how their own AI is governed.

4.4 Layer 4: The NROC and controlled orchestration

Layer 4 is the operational face of STORES: the National Regulatory Operations Centre described in Section 5. Orchestration in STORES means coordination through the owners of each system. The NROC requests or recommends actions through standard operator interfaces such as TM Forum Open APIs and GSMA Open Gateway or CAMARA APIs, and the operator executes them in its own network. In declared emergencies, pre-agreed playbooks allow faster action, still approved by a named official and still executed by the operator.

4.5 Threat model

A system that concentrates regulatory insight is itself an attractive target. Table 6 sets out the principal threats and the architectural controls that address them.

Table 6. STORES threat model

Threat actor	Example attack	Key controls
External attacker	Compromise the NROC to issue false requests or leak data	Zero Trust access; no direct write path to networks; dual control on R3 approvals; segregated environments
Malicious or coerced insider at the regulator	Use data for surveillance or to favour an operator	Purpose limitation in law; no individual-level data held centrally; immutable access logs reviewed by an independent oversight board
Curious operator	Infer a competitor's performance or strategy from shared outputs	Confidential-computing aggregation; output thresholds; no cross-operator raw data sharing
Compromised or manipulated node	Falsify measurements or exfiltrate data	Remote attestation; signed code; cross-validation against independent probes and crowdsourcing
Measurement gaming and spoofing	Prioritise known probe traffic; submit fake crowdsourced results from emulated devices	Probe diversity; randomised schedules; app and device attestation; rate limits and outlier filtering; cross-checks between sources

© 2026 R. W. Charara and K. M. Chatila

Threat actor	Example attack	Key controls
Attacks on federated learning	Model poisoning by a participant; membership inference; side-channel attacks on enclaves	Update validation and robust aggregation; subscriber-level DP bounds membership inference; patched, attested enclaves; FL used only where federated analytics cannot answer the question
Adversarial content against AI	Prompt injection or poisoned reference documents	Policy engine outside the model; signed knowledge sources; red-teaming (Section 6.3)

© 2026 R. W. Charara and K. M. Chatila

5. The National Regulatory Operations Centre (NROC)

STORES culminates in the NROC, a cross-sector coordination hub rather than a monitoring room. The NROC brings operator data and third-party infrastructure data, including power grids, energy storage, fibre plant, cloud and emergency services, into one shared situational picture. This "NOC of NOCs" approach matters because many of the most damaging telecom outages originate in adjacent sectors: a substation failure, a cut fibre route, a cooling failure in a data centre.

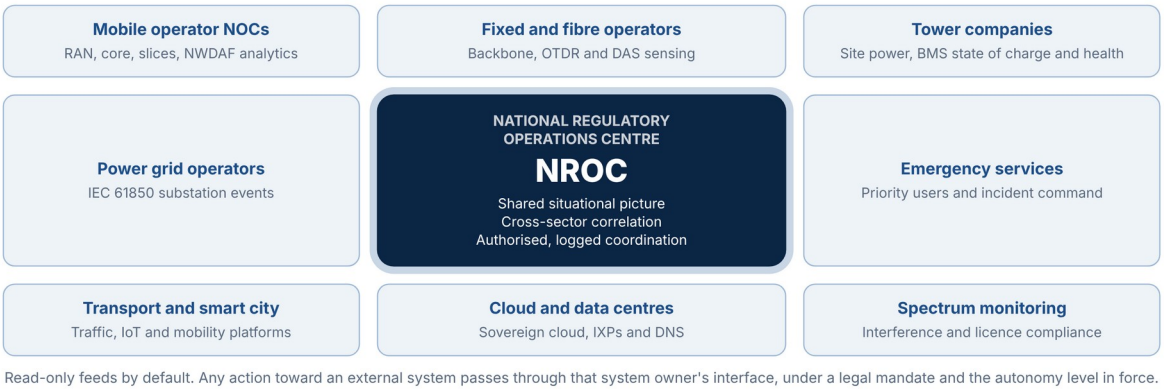


Figure 8. The NROC as a "NOC of NOCs": a shared picture across telecom and adjacent critical infrastructure.

A telecom regulator has no authority over power grids, transport or emergency services, and STORES does not assume any. Integration therefore follows four rules. Cross-sector participation rests on memoranda of understanding or legislation that define purpose, data, retention and roles. Feeds are read-only by default. Each authority keeps command of its own sector, and in an emergency the national emergency-management authority leads. Any action toward an external system passes through its owner's interfaces, with a logged decision.

5.1 Operational scenario: national 5G resilience during a regional power disruption

Consider a regional grid disruption affecting several metropolitan areas during the evening peak. In a conventional setting, each operator manages its own sites, information reaches the regulator through phone calls and emails, and nobody holds a joint picture of which hospitals, emergency services or critical enterprise slices are at risk.

With STORES, the NROC correlates the grid operator's substation events, received through an IEC 61850 gateway[28], with site power alarms and battery telemetry from operators and tower companies. The NIDT maps the affected area to the sites, slices and critical services at risk, and ranks sites by remaining battery autonomy and criticality. Within minutes the NROC proposes a priority plan to the emergency-management authority and the operators: protect emergency communications and hospital connectivity, dispatch generators where batteries are short, and reroute traffic over resilient transport paths. A named official authorises the regulator's part of the plan (R3), operators execute it through their own systems, and status is reported automatically. After the event, a complete, anchored record supports lessons learned and any compliance follow-up.

© 2026 R. W. Charara and K. M. Chatila

	T0	T+2 min	T+5 min	T+15 min	T+60 min	T+24 h
Power grid	Substation trips; regional feeders lost	Grid operator publishes event via IEC 61850 gateway	Restoration estimate shared: 3 to 4 hours		Partial restoration	Grid restored; post-event data shared
STORES / NROC	Correlates grid event with site power alarms	Digital twin maps 412 sites, 3 slices and 2 hospitals at risk	Ranks sites by battery autonomy (BMS) and criticality	Recommends priority plan; evidence pack generated	Tracks QoE and slice KPIs against plan	Incident report and lessons learned anchored
Operators	Sites switch to batteries	Receive risk map through API	Confirm generator dispatch and load shedding plans	Execute priority for emergency slice and rerouting	Report status automatically	Root-cause and compliance data submitted
Regulator and authorities	Alert raised to duty officer	Incident opened; emergency services informed		Authorises priority plan (human decision)	Public information update	Reviews obligations; decides on follow-up

Illustrative timeline. Site counts and durations are hypothetical and shown only to explain the sequence of events and decision points.

Figure 9. Illustrative incident timeline: who knows what, and who decides what, during a regional power disruption.

The scenario shows the shift from passive monitoring to active, accountable stewardship: faster shared awareness, better-targeted response and an evidence trail that serves both operational learning and regulation.

5.2 Power grid integration and energy resilience

5G sites, edge computing and future 6G services demand high power reliability. IEC 61850, the international standard for communication in electrical substations, provides a structured way to share substation events with authorised parties. When a transformer or feeder fails, the NROC can identify which sites depend on it and alert operators to activate backup procedures or prioritise critical slices while on battery power.

Battery Management System (BMS) telemetry, notably state of charge (SoC) and state of health (SoH), tells the NROC how long each site can survive and which batteries are degrading before the next incident exposes them. Regulators can use it to set and verify backup-power obligations for critical sites, for example a minimum number of hours of autonomy, instead of relying on declarations. Energy and digital infrastructure are converging fast: the International Energy Agency projects that global data-centre electricity consumption will more than double to around 945 TWh by 2030, slightly more than Japan's total electricity consumption today[2].

5.3 Fibre plant: precise fault localisation and sensing

Resilience ultimately depends on the physical integrity of the fibre plant. Optical Time Domain Reflectometry (OTDR) locates a break or degradation along a fibre to within metres, while Distributed Acoustic Sensing (DAS) turns existing fibre into a continuous sensor for vibration, ground movement and temperature. Integrated into the NROC, these signals shorten repair times, help to distinguish accidental damage, such as civil works, from deliberate interference, and reveal routes where several operators depend on the same physical path: a hidden single point of failure.

© 2026 R. W. Charara and K. M. Chatila

Table 7. Cross-sector integration in the NROC

Infrastructure sector	Integrated data stream	NROC role	Expected outcome
Power grid	IEC 61850 substation events	Map dependencies; recommend slice priority and backup activation	Fewer and shorter cascading site outages
Energy storage	BMS state of charge and health	Rank sites by autonomy; verify backup-power obligations	Targeted generator dispatch; evidence-based battery obligations
Passive fibre	OTDR and DAS	Localise faults; detect shared-route risks	Faster repair; fewer common-mode failures
Emergency services	Incident and priority-user data	Share telecom status; support the lead authority	Faster, better-coordinated response

5.4 The National Infrastructure Digital Twin (NIDT)

The NIDT is a living virtual replica of the nation's digital infrastructure, fed by the data fabric and the NROC. Unlike static coverage maps, it combines network topology, radio propagation models, traffic and mobility patterns, and dependencies on power, fibre and cloud. It enables "what-if" analysis for resilience and planning, for example simulating rerouting options after a subsea cable cut, or the effect of a new data centre on regional power capacity.

For spectrum policy, the twin can model coverage and interference before licences are awarded, for instance in 26 GHz millimetre-wave bands whose propagation is highly sensitive to buildings and terrain, and compare alternative obligations before they are written into licences. Several governments, including Saudi Arabia, are developing national or city-scale digital twins[29,30]. Fidelity is the key test for regulators: a twin that is not continuously calibrated against measured data becomes an attractive but misleading model.

© 2026 R. W. Charara and K. M. Chatila

6. AI Governance, Ethics and Human Oversight

A regulator that uses AI must meet a higher standard than the entities it supervises, because its decisions affect rights, markets and public trust. STORES therefore embeds AI governance in the architecture, aligned with the human-oversight requirements of the EU AI Act (Article 14)[31], the NIST AI Risk Management Framework[32] and the ISO/IEC 42001 AI management system standard[33]. Section 7.2 addresses how the AI Act's risk classes may apply.

6.1 Explainability and accountability

Every AI-assisted finding produces a **decision record**: the data used, with evidence-log references; the obligation tested; the model and its version; the confidence level; the alternative explanations considered; and the human decisions taken. Rather than relying on a model's free-text "reasoning", which may not faithfully reflect how a result was computed, STORES ties explanations to verifiable artefacts: evidence, rules and model documentation. Officials can reconstruct why a finding was made, and operators can challenge it on its merits.

6.2 The STORES regulatory autonomy model

Autonomy is set per type of action, not per system, and is proportionate to legal effect. The dividing line is simple: **anything that creates an obligation, starts a legal deadline or affects a network requires a named official's approval before it happens.** Figure 10 and Table 8 define five levels.

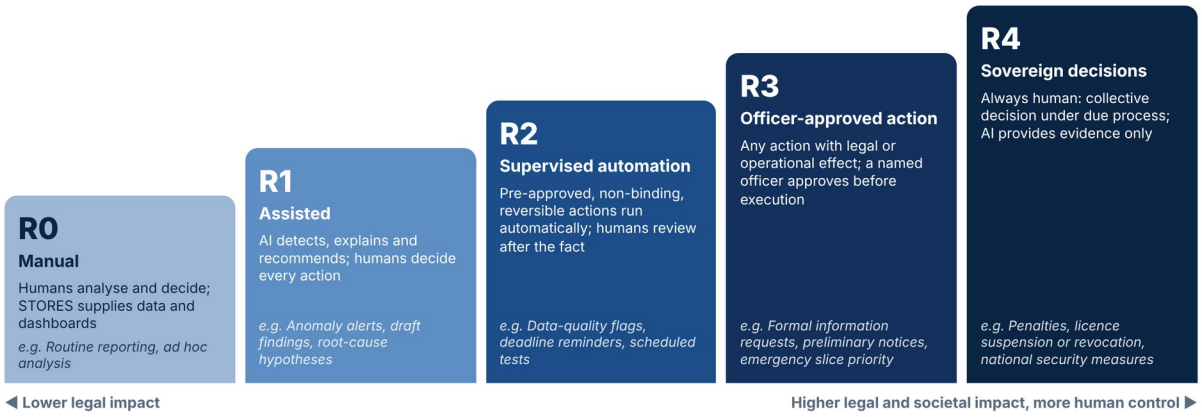


Figure 10. The STORES regulatory autonomy ladder (R0 to R4).

© 2026 R. W. Charara and K. M. Chatila

Table 8. Regulatory autonomy levels and safeguards

Level	Example actions	Human role	Safeguards
R0 Manual	Routine reporting, ad hoc analysis	Analyse and decide	Data-quality controls
R1 Assisted	Alerts, draft findings, root-cause hypotheses	Decide every action	Decision record; confidence and evidence shown
R2 Supervised automation	Non-binding actions only: data-quality flags, deadline reminders, scheduled tests	Review after the fact; can reverse	Pre-approved playbooks; reversibility; sample audits; no legal effect
R3 Officer-approved action	Formal requests for information, preliminary notices, emergency slice-priority requests	Named official approves before execution	Evidence pack; dual control for high impact; time-boxed measures; full logging
R4 Sovereign decisions	Penalties, licence suspension or revocation, national-security measures	Collective decision under due process	AI supplies evidence only; right to be heard and to appeal

The model should be reviewed every year. Moving a use case to a higher level requires evidence of model performance, error analysis (including false positives that would burden operators unfairly) and staff readiness. Automation complacency, the tendency to accept machine output without scrutiny, is a known risk; training and periodic challenge reviews of AI findings mitigate it.

6.3 Security: defending the agentic regulator

AI agents expand the attack surface. Large language models can be manipulated through prompt injection, including instructions hidden in documents or data they process; retrieval sources can be poisoned; and agents with broad tool access can act beyond their intended scope. The OWASP Top 10 for LLM Applications catalogues these risks[34]. STORES responds with layered controls:

- **Policy enforcement outside the model.** A separate, deterministic policy engine checks every tool call against the agent's permissions, the autonomy level and parameter limits before execution. The model's own judgment is never the last line of defence.
- **Least privilege and segregation of duties.** Agents receive narrowly scoped, short-lived credentials; read and write paths are separated; no agent can both propose and approve an action.
- **Untrusted content and signed sources.** All retrieved content and model output are treated as untrusted and validated before reaching downstream systems; the obligations library and reference documents are version-controlled and signed against retrieval poisoning.
- **Zero Trust architecture.** Every user, device and service is authenticated and authorised continuously, following NIST SP 800-207[35].
- **Privacy controls.** Personally identifiable information is protected in line with ITU-T X.1058 (02/2026), which is aligned with ISO/IEC 29151[36], and with national data-protection law.
- **Model sovereignty, red-teaming and monitoring.** Prefer AI models that can be hosted in-country; where external general-purpose models are used, prompts, documents and outputs stay in national hosting and the provider's obligations are written into contracts. Agents and models are red-teamed regularly and monitored for drift and misuse.

© 2026 R. W. Charara and K. M. Chatila

7. Legal Design, Rights and Due Process

Technology alone does not create regulatory power, and a powerful supervisory system without strong legal design would rightly be resisted. "Legal design" means shaping the legal framework and the architecture together, so that STORES outputs are lawful, proportionate and usable in proceedings. This section sets out the questions every implementing country must answer; Appendix A offers illustrative model provisions. The analysis uses EU law as a well-documented reference point; other jurisdictions should map each item to their own legislation, such as the national data-protection laws now in force in several GCC states and many other countries, and seek local legal advice.

7.1 Legal basis, mandate and jurisdiction

STORES needs an explicit legal basis for: continuous access to specified operator data; hosting regulatory nodes and probes in operator networks; and processing measurement data for defined regulatory purposes. In the EU, Article 20 of the European Electronic Communications Code already empowers regulators to require information from undertakings, provided requests are proportionate and justified[37]; continuous, interface-based access and regulatory nodes may still need explicit provisions. Where existing telecom law covers only periodic reporting, licence conditions or secondary legislation should be updated (Recommendation 6). The mandate must also be bounded: STORES covers network and service performance, resilience and licence compliance; it excludes communications content and the monitoring of individuals.

Cross-sector functions need their own basis. A telecom regulator cannot direct grid operators, transport agencies or emergency services. NROC participation by those bodies should rest on memoranda of understanding or legislation that set out what data is shared, for which purposes, who leads in an emergency (normally the national emergency-management authority) and how costs and liabilities are shared.

7.2 Classification of the regulator's own AI

Under the EU AI Act, whether a STORES component is "high-risk" depends on its intended purpose. Annex III lists, among others, AI used as safety components in the management and operation of critical digital infrastructure (point 2)[38]. Its other categories mainly concern decisions about natural persons, which STORES does not make. Analytical assistants that draft findings for officials are unlikely, on their own, to be safety components. R3 emergency playbooks that influence the operation of networks during a crisis could, however, fall within point 2.

STORES therefore recommends a precautionary approach: apply high-risk-equivalent controls to all R2 and R3 functions (risk management, data governance, logging, transparency, human oversight, accuracy and robustness testing), whatever their formal classification. Under Article 27, public bodies deploying high-risk AI must carry out a fundamental rights impact assessment before use, although systems in the critical-infrastructure area of point 2 are excluded from that specific duty[39]. STORES recommends carrying out a rights impact assessment for every use case regardless, because the reputational and legal stakes for a regulator are high. A regulator that develops or commissions STORES AI and puts it into service under its own name may also count as a provider, not only a deployer. For any high-risk component it would then carry provider duties under the AI Act, such as a quality management system, technical documentation, conformity assessment and registration, in addition to deployer duties such as human oversight and monitoring[31]. Procurement contracts should allocate these roles explicitly. Countries outside the EU can use the same approach as a benchmark alongside their national AI principles.

© 2026 R. W. Charara and K. M. Chatila

7.3 Data protection

STORES is designed to minimise personal data, but some inputs, such as crowdsourced measurements, device identifiers or NWDAF analytics, can contain or be linked to personal data. The governing principles are:

- **Roles.** Operators remain controllers of their subscriber data. The regulator is controller for the data it processes for supervision. Cross-sector participants are independent controllers under the terms of their agreements. Roles are documented for each data flow.
- **Lawful basis.** The regulator's processing rests on legal obligation or the performance of a task in the public interest, as recognised in the EU GDPR[40] and in laws such as Qatar's Law No. 13 of 2016 on Personal Data Privacy Protection[41]. It does not depend on consent, so it cannot be undermined by consent being withdrawn.
- **Telecom-specific rules.** Traffic and location data are subject to stricter sector rules: in the EU, the ePrivacy Directive permits their processing only under narrow conditions, typically after anonymisation[42]. STORES therefore takes such data out of operator networks only in aggregated or anonymised form. National equivalents apply elsewhere.
- **Minimisation.** Per-subscriber analytics stay inside operators; only aggregates, protected model updates and evidence records that exclude personal data by design leave the node.
- **Crowdsourcing.** Taking part is voluntary: users choose to install and run the measurement app and are told clearly what it collects. Reading measurements from a user's device requires the user's consent under rules such as Article 5(3) of the EU ePrivacy Directive[42], which the app obtains at installation; the regulator's subsequent processing of those measurements rests on its public task. Location is coarsened before analysis, and published methods state the known limits of uncontrolled tests[4].
- **Retention.** Each data category has a defined retention period: for example, evidence records are kept for the limitation period of enforcement proceedings plus appeal, and raw crowdsourced data no longer than needed for validation.
- **Immutability and erasure.** No personal data, and no unsalted hash of personal data, is written to the anchoring log. Evidence records in the off-chain store follow normal retention and deletion rules.
- **Impact assessments.** A data-protection impact assessment precedes each use case and is reviewed when its purpose or data changes.

7.4 Due process and admissibility of evidence

Findings based on STORES evidence must survive challenge. Each case should give the operator: notice of the alleged deviation; access to the records, measurement methods, calibration certificates and model documentation relied upon; a reasonable period to respond, including by submitting its own measurements; independent re-measurement on request; and review by an official not involved in the original finding, followed by appeal. Investigation and adjudication should be separated: the team that builds a case does not decide it. R4 decisions are taken by the regulator's collegiate decision-making body, such as its board or commission, which states the standard of proof applied. Where penalties are severe, courts in many jurisdictions treat them as quasi-criminal, which in Council of Europe states engages the fair-trial guarantees of Article 6 of the European Convention on Human Rights; the procedure should be designed to that standard, including the possibility of suspending a sanction pending appeal. All decisions remain subject to judicial review, and operators can ask for any automated R2 action to be reviewed by an official. The FCC's mobile

© 2026 R. W. Charara and K. M. Chatila

challenge process, with published test thresholds and a 60-day period for providers to concede or rebut a challenge, shows how such rules can be made explicit[6].

On admissibility, trusted time stamps and hash anchoring strengthen chain of custody, but they prove only that a record existed at a given time and has not been changed. Accuracy must be shown separately, through calibrated equipment, published methods and statistical validity. In the EU, qualified electronic time stamps enjoy a presumption of accuracy of time and integrity of data, and electronic time stamps cannot be refused as evidence solely because they are electronic[22]. Other jurisdictions should confirm the evidential status of these records in law before relying on them in enforcement.

7.5 Net neutrality, emergency measures and liability

Prioritising emergency traffic is a form of traffic management, and open-internet rules constrain it. In the EU, Article 3(3) of Regulation 2015/2120 allows measures beyond reasonable traffic management only as necessary and for as long as necessary to comply with law or orders of public authorities, to preserve network integrity and security, or to prevent impending congestion and mitigate exceptional or temporary congestion[43]. STORES emergency playbooks should therefore be grounded in national emergency-telecommunications law, triggered only by a declared emergency, limited in time and scope, and reviewed afterwards. An operator that acts in good faith on an authorised request should be protected from liability for that action, while the authority's own liability follows public law. These allocations belong in legislation or licence conditions (Appendix A).

7.6 Competition safeguards

A regulator that sees every operator's performance holds commercially sensitive information. STORES must never become a channel for exchanging it. Information barriers separate case teams; comparative indicators are published only at levels of aggregation and delay that prevent inference of a competitor's network plans or commercial position; and operator-specific data shared with one operator never includes another's. Competition authorities should be consulted on publication rules.

7.7 Preventing function creep

The greatest long-term risk of any national data capability is gradual expansion beyond its original purpose. STORES purposes should be fixed in law, with any extension requiring the same legislative process. An independent oversight board, separate from the regulator's executive and including technical, legal and civil-society members, should review access logs and use cases and publish an annual transparency report. A periodic review clause, for example every five years, keeps the system under democratic control.

© 2026 R. W. Charara and K. M. Chatila

8. Alignment with Global Standards

STORES is designed to be standards-native: every KPI, interface and orchestration flow should map to an internationally recognised definition. This matters for legal defensibility, because a finding based on a standard method is harder to contest; for interoperability across operators and vendors; and for international comparability. Table 9 maps the principal standards to their role in STORES.

Table 9. Standards and frameworks underpinning STORES

Domain	Standard or framework	Role in STORES
QoS definitions	ITU-T E.800[7]	Common terminology for quality of service
Video QoE	ITU-T P.1203[12]	Objective estimation of streaming video quality (MOS)
AI-based quality optimisation	ITU-T Y.3192 (12/2025)[25]	Customer-oriented quality optimisation, including complaints despite normal KPIs
Intent-based slicing	ITU-T Q.5038 (04/2026)[14]	Signalling and protocol for intent-based management and orchestration of network slices
Tactile internet QoS	ITU-T Y.3149 (08/2025)[44]	QoS assurance requirements for ultra-low-latency, high-reliability services
Preventive maintenance	ITU-T M.3370 (03/2025)[45]	Overview of preventive maintenance tasks, supporting self-healing network behaviour
Network analytics	3GPP TS 23.288 (NWDAF)[13]	Standard analytics on load, QoS and user experience in the 5G core
Intent-driven management	3GPP TS 28.312[46]	Expressing and fulfilling intents in network management
Autonomous operations	ETSI ZSM[27]; TM Forum Autonomous Networks[26]	Closed-loop automation and autonomy levels
Open architecture and APIs	TM Forum ODA and Open APIs[47]; GSMA Open Gateway and CAMARA[8]	Vendor-neutral integration and network APIs
Power systems	IEC 61850[28]	Substation communication for grid integration
Evidence integrity	IETF RFC 3161[21]; RFC 9162[24]; ISO/IEC 17025[23]	Trusted time stamps; transparency logs; accredited calibration
Security and privacy	NIST SP 800-207[35]; ITU-T X.1058[36]	Zero Trust architecture; protection of personal data
AI governance	ISO/IEC 42001[33]; NIST AI RMF[32]; EU AI Act[31]	AI management system, risk management and human oversight

Building on open, consensus-based specifications keeps the architecture technically neutral, lowers integration costs for operators and reduces licensing and lock-in risk. Procurement rules should require open standards, open APIs and data portability, so that the state retains control over its supervisory capability.

© 2026 R. W. Charara and K. M. Chatila

9. Global Applicability Across Regulatory Environments

STORES is a reference architecture, not a product. Its modular design lets each country start where its needs and maturity are, and grow capabilities over time. Table 10 suggests typical starting points, and Appendix B provides a self-assessment checklist.

Table 10. Adopting STORES by market context

Market context	Typical starting point	Priority STORES capabilities	Typical first wins
Emerging digital economies	Operator self-reporting; limited independent measurement	Independent QoE and coverage measurement; obligations library; CEIR-style cross-operator interfaces	Transparent coverage maps; verified rural obligations
Developing 5G markets	Drive tests and probes; early 5G-SA	Slice-aware KPIs; data fabric; evidence anchoring	Peak-hour transparency; faster dispute resolution
Advanced 5G-SA markets	Rich operator analytics; mature regulator	Federated models; specification-aware AI; NIDT	Supervision of automated networks; outcome-based licensing
Smart-city and sovereign-cloud programmes	Strong central digital agenda	NROC cross-sector coordination; twin-based planning	Coordinated resilience across telecom, power and transport

Whatever the starting point, the same principles apply: open standards, sovereign control of data and keys, verifiable and contestable evidence, and bounded AI.

© 2026 R. W. Charara and K. M. Chatila

10. Costs, Resourcing and Funding

A credible programme needs a credible budget. Figure 11 and Table 11 give indicative five-year total cost of ownership by component and market size. They are order-of-magnitude estimates built bottom-up from typical public-sector ICT price levels, intended as illustrative scenario estimates to support early planning, not to replace a market consultation. Actual expenditure will depend on the number of participating operators, geographic coverage, existing regulatory infrastructure, measurement density, integration complexity, data sovereignty requirements, staffing arrangements, local labour costs and the scope of AI and digital-twin capabilities deployed.

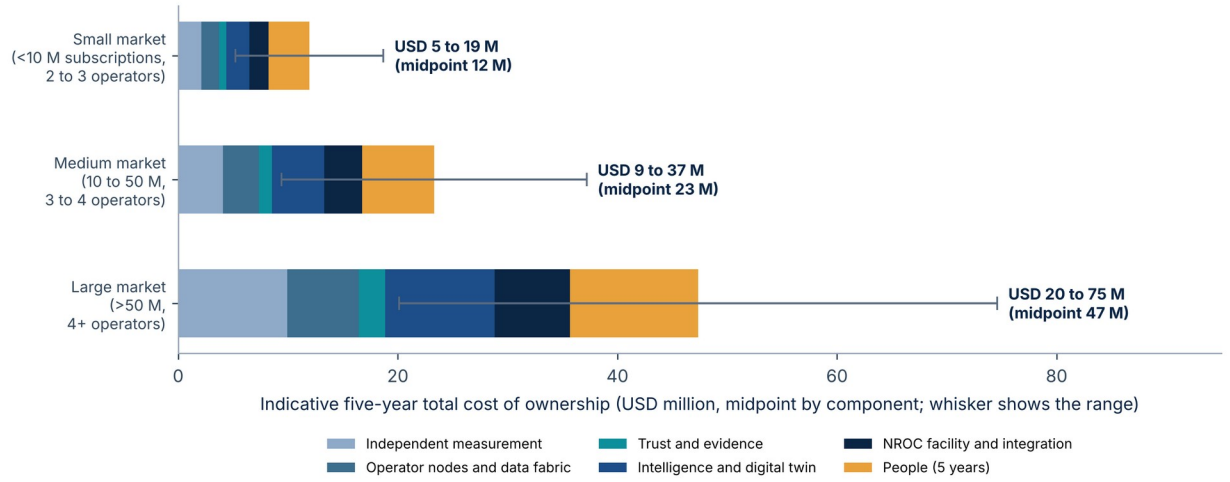


Figure 11. Indicative five-year total cost of ownership by market size and component.

Table 11. Indicative five-year cost ranges (USD million)

Component	Main cost drivers	Small market	Medium market	Large market
Independent measurement	Probes, drive-test kits, crowdsourcing app and data, calibration	0.9 to 3.3	1.7 to 6.5	4.3 to 15.6
Operator nodes and data fabric	Node appliances, integration with each operator, data contracts	0.7 to 2.6	1.3 to 5.2	2.6 to 10.4
Trust and evidence	Time-stamping, transparency log, evidence store, PKI	0.3 to 1.0	0.4 to 2.0	0.9 to 3.9
Intelligence and digital twin	Analytics platform, AI models, twin software, compute	0.9 to 3.3	1.7 to 7.8	4.3 to 15.6
NROC facility and integration	Operations room fit-out, case management, cross-sector interfaces	0.9 to 2.6	1.7 to 5.2	3.4 to 10.4
People (five years)	Regulator staff and specialist support	1.5 to 5.9	2.6 to 10.5	4.6 to 18.7
Total		**5 to 19**	**9 to 37**	**20 to 75**

© 2026 R. W. Charara and K. M. Chatila

Assumptions behind the estimates

Market sizes are defined in Table 13 and Section 10.1. Costs exclude land and buildings and assume reuse of existing regulator data-centre or sovereign-cloud capacity. People costs assume the staffing levels in Section 10.2, ramped up over the roadmap so that average staffing across the five years is 50% to 60% of full strength, at blended loaded costs of USD 40,000 to 80,000 per full-time equivalent (FTE) per year. In high-income markets people costs can be two to three times higher. The ranges are deliberately conservative: lower bounds were reduced by about 15% and upper bounds raised by about 30% from the central bottom-up estimate, so that budgets planned on the upper bound are unlikely to be exceeded.

Table 12. Basis of the cost estimates

Item	Assumption
Price basis	Constant 2026 US dollars; excluding taxes, customs duties and inflation, which should be added for the implementation years
Capital and operating expenditure	Roughly 40% capital and 60% operating over five years. Indicative capital share by component: measurement 60%, operator nodes 70%, trust and evidence 40%, intelligence and twin 50%, NROC facility 70%, people 0%
Hosting	Reuse of an existing regulator data centre or sovereign-cloud tenancy; compute sized for national indicators, models and evidence records, not for raw traffic
Salary baseline	Blended loaded cost of USD 40,000 to 80,000 per FTE-year, with a phased ramp-up averaging 50% to 60% of full staffing; two to three times higher in high-income markets
Contingency	Ranges include a contingency of about 20% and are widened on the conservative side (lower bounds about 15% below, upper bounds about 30% above the central estimate), as suits a pre-feasibility estimate
Exclusions	Land and buildings; operators' internal costs beyond integration; legislative drafting; communications campaigns
Validation	No public procurement covers this full scope, so ranges cannot be benchmarked against a single comparable programme. Regulators should validate them through a market consultation or request for information, and against their own past contracts for measurement systems, operations centres and data platforms

10.1 Classifying a market

The market size used in Table 11 is set mainly by the number of mobile subscriptions, because subscriptions drive traffic, measurement volume and the scale of analytics. Other factors then adjust the classification, since they drive particular cost components: the number of operators drives integration work, territory and radio sites drive measurement effort, and the scope adopted drives intelligence, NROC and staffing costs. Table 13 gives the criteria.

© 2026 R. W. Charara and K. M. Chatila

Table 13. Market size classification

Criterion	Small	Medium	Large
Mobile subscriptions (primary criterion)	Fewer than 10 million	10 to 50 million	More than 50 million
Mobile operators (typical)	2 to 3	3 to 4	4 or more
Radio sites, all operators (indicative)	Up to about 10,000	About 10,000 to 50,000	More than 50,000
Typical scope	Phases 1 to 3, NROC light	Phases 1 to 5	Phases 1 to 5 with extensive cross-sector feeds

How to classify. Start from the number of mobile subscriptions. Then check four adjustment factors:

- more operators than typical for that size;
- more radio sites than typical for that size;
- a large or difficult territory that needs extensive drive testing, for example more than about one million square kilometres, or mountainous or sparsely populated areas;
- fuller scope than typical, such as a 24-hour NROC with many cross-sector partners.

If one factor applies, plan on the upper half of the range for that size; if two or more apply, move up one size. If only Phases 1 and 2 are adopted, plan on the lower half of the range. Never move more than one size. For example, a market with 12 million subscriptions, two operators and a compact territory that adopts only Phases 1 and 2 is medium, in the lower half of its range. A market with 8 million subscriptions, four operators and a territory of 1.5 million square kilometres starts as small but has two upward factors, so it is planned as medium.

Because components scale with different drivers, regulators should adjust each cost line rather than only the total. Table 14 shows what drives each component.

Table 14. What drives each cost component

Component	Main scaling driver	How to adjust
Independent measurement	Territory, radio sites, population distribution, technologies monitored	Scale with sites and drive-test distance; add for each additional technology, such as fixed or satellite
Operator nodes and data fabric	Number of operators and other connected organisations	Roughly proportional to the number of connected organisations
Trust and evidence	Evidence volume	Largely fixed; grows slowly with volume
Intelligence and digital twin	Number of use cases; geographic extent of the twin	Mainly scope-driven, with a fixed platform core
NROC facility and integration	Operating hours; number of cross-sector partners	Fixed core plus a cost per partner integration
People	All of the above, especially round-the-clock operation	See Section 10.2

© 2026 R. W. Charara and K. M. Chatila

10.2 People and skills

Technology is the smaller part of the challenge. Indicatively, a small market needs about 15 to 25 FTE, a medium market 25 to 45 and a large market 45 to 80, phased in over the roadmap. Core roles are: NROC duty officers (round-the-clock cover needs about five people per seat); measurement and QoS engineers; data engineers; data scientists and AI specialists; security engineers; legal and policy analysts; and a product owner who manages the obligations library. Many regulators will combine a small permanent team with specialist support contracts that include knowledge transfer.

10.3 Operator burden and offsets

The programme funds node hardware and software. Each operator hosts the node (typically one or two racks or equivalent cloud capacity), integrates its data exports and supports data contracts. Indicatively this takes six to twelve person-months at onboarding and half to one FTE thereafter. In return, STORES should retire overlapping manual reporting, reduce ad hoc information requests, shorten disputes and give operators access to the same evidence the regulator uses. These offsets should be written into the programme's targets and measured.

10.4 Funding models

Options include the regulator's own budget, an earmarked share of licence or spectrum fees, or a managed-service contract under a public-private partnership in which the state retains ownership of data, keys and models. Direct funding of measurement by individual operators should be avoided, because it undermines independence.

10.5 Assessing value

Benefits should be estimated with a standard cost-benefit method: define the baseline; quantify costs over the programme life; estimate benefits for each value driver in Table 15; discount both; test sensitivity to key assumptions; and validate after the fact. Wider economic effects can be modelled with input-output methods, as explained in Appendix C.

Table 15. How STORES creates value, and how to measure it

Value driver	Main beneficiary	Mechanism	How to measure
Service quality	Citizens and enterprises	Outcome-based obligations; peak-hour transparency	Share of sessions meeting intent; peak-hour degradation by district
Regulatory efficiency	Regulator and operators	Automated evidence collection; fewer data disputes; retired manual reporting	Time from detection to resolution; disputes on data authenticity; reporting hours saved
Resilience	Government and society	Cross-sector early warning; verified backup power	Outage minutes avoided; restoration time for critical services
Investment efficiency	Operators and investors	Twin-based planning; visibility of shared risks	Capex per covered user; duplicate builds avoided
Economic growth	National economy	Better coverage and capacity decisions	Modelled output and GDP effects, validated after the fact

© 2026 R. W. Charara and K. M. Chatila

11. Implementation Roadmap

Moving to a fully integrated NROC requires a shift from purely ex-ante, rule-based regulation toward collaborative governance with operators and other infrastructure owners, including public-private partnership models where useful. The roadmap uses overlapping phases so that each one delivers value on its own, and legal instruments are prepared first. Timelines assume the legal changes in Section 7 can be adopted within about 18 to 36 months; where that takes longer, later phases proceed under interim, voluntary arrangements with reduced scope.

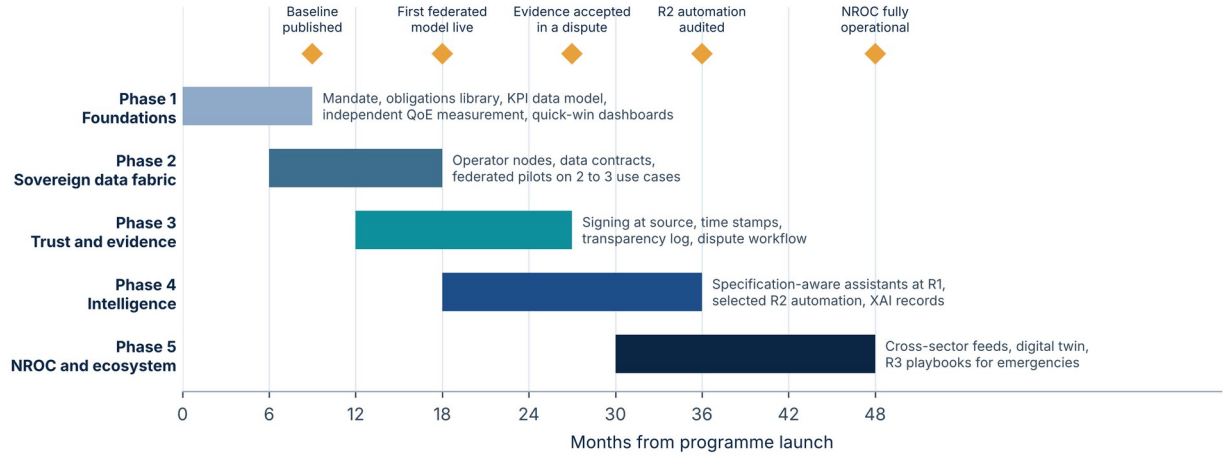


Figure 12. Indicative STORES implementation roadmap and milestones.

© 2026 R. W. Charara and K. M. Chatila

Table 16. Implementation phases and indicative exit criteria

Phase	Timeline	Key deliverables	Exit criteria (indicative)
1 Foundations	0 to 9 months	Legal review; draft instruments submitted (Appendix A); interim data-sharing MoUs under existing powers; oversight board; obligations library v1; common KPI data model; independent QoE measurement; baseline dashboards	Baseline published with methods; interim arrangements with all operators; legislative amendments submitted
2 Sovereign data fabric	6 to 18 months	Pilot nodes under interim arrangements; full roll-out with data-access and node agreements once legal amendments are adopted; operator nodes; data contracts; federated analytics; federated learning pilots on two or three use cases	National indicators computed without raw data transfer; privacy-budget governance in place
3 Trust and evidence	12 to 27 months	Signing at source; time stamps; transparency log; evidence store; dispute and challenge workflow	Anchored evidence used in at least one formal case and accepted after challenge, once the legal basis is in force
4 Intelligence	18 to 36 months	Specification-aware assistants at R1; selected R2 automation; decision records; model risk management	Measured precision of findings; independent audit of R2 playbooks
5 NROC and ecosystem	30 to 48 months	Cross-sector MoUs and feeds; NIDT; R3 emergency playbooks; joint exercises	NROC operational around the clock; first national resilience exercise completed

11.1 Key risks and mitigations

Table 17. Principal programme risks and mitigations

Risk	Mitigation
Operator resistance to data sharing	Federated design; clear legal basis; equal access to evidence; offsets such as retired reporting and faster disputes
Over-reliance on AI findings	Autonomy model; decision records; challenge reviews; training against automation complacency
Legal challenge to findings	Due-process rules; accredited calibration; published methods; independent re-measurement
Vendor lock-in	Open standards and APIs written into procurement; data portability; modular contracts
Cyber-attack on the platform	Threat model controls (Section 4.5); Zero Trust; national control of cryptographic keys
Privacy breach or function creep	Purpose limitation in law; minimisation; independent oversight board; transparency reports
Skills shortage and cost overrun	Phased capability building; knowledge-transfer clauses; stage gates tied to exit criteria

© 2026 R. W. Charara and K. M. Chatila

12. Strategic Recommendations

Recommendation 1: Adopt standards-native design and open APIs

Mandate standardised, open interfaces for the exchange between operators and STORES, for example TM Forum Open APIs within the ODA framework and, where relevant, GSMA Open Gateway and CAMARA. This prevents vendor lock-in, lowers integration cost for operators and keeps the Agentic Experience legible to the regulator's own AI assistants[47].

Recommendation 2: Institutionalise the National Infrastructure Digital Twin

Use the NIDT as the primary sandbox for testing policy before adoption: spectrum award designs, coverage and resilience obligations, net-neutrality rules and disaster-response protocols. Fund its continuous calibration with measured data, so that the twin remains a trustworthy replica rather than an attractive model.

Recommendation 3: Transition to intent-based performance standards, with published methods

Complement network-level QoS obligations (for example packet loss at Layer 3) with outcome-based obligations defined at application level (Layer 7): the share of sessions in which a service achieves the performance it requires, by area and time of day. Publish measurement methods and thresholds in advance, give operators a right to challenge, and introduce the obligations gradually so that operators can adapt their investment plans[25].

Recommendation 4: Treat sovereignty as strategic control, not isolation

Define digital sovereignty as national control over data, keys, models and critical decisions, rather than as technological isolation. Federated architectures, confidential computing and sovereign cloud allow strict residency requirements to coexist with access to global innovation and scale, provided that keys, time-stamping and attestation remain under national control.

Recommendation 5: Invest in human judgment and AI literacy

The autonomy model is only as strong as the people who operate it. Invest in regulator and operator teams who can interrogate AI findings, recognise automation complacency and apply "two-factor judgment" to critical actions: an AI-supported finding plus an independent human assessment before any decision with legal effect.

Recommendation 6: Modernise the legal basis for data, evidence and AI

Update telecom law and licence conditions to provide explicit authority for continuous data access, regulatory nodes and probes, the evidential status of signed and time-stamped records, and the use of AI in supervision. Pair these powers with the safeguards in Section 7: purpose limitation, due process, competition safeguards and independent oversight. Appendix A offers illustrative drafting.

© 2026 R. W. Charara and K. M. Chatila

13. Conclusions: A New Era of National Infrastructure Stewardship

STORES and its NROC are more than a technology upgrade. They offer a practical way to modernise the regulatory mandate for the 5G-SA era and to prepare for 6G: from periodic, report-based oversight to continuous, evidence-based and sovereign stewardship of national digital infrastructure.

No country needs to build everything at once. Starting with the legal foundations, independent measurement, an obligations library and a common data model, a regulator can show value within months and grow toward an integrated NROC over several years, at a cost that is modest relative to the economic value of the networks it supervises. The outcome is a national digital backbone that is resilient, fair and high-performing, overseen by a regulator whose supervision is as precise, intelligent and transparent as the infrastructure it serves.

Looking ahead: the future of 6G governance

As networks evolve toward AI-native 6G, integrated sensing and communication, and machine-to-machine economic interaction, regulators will increasingly supervise systems that configure themselves. Governance will need to become programmable in the same way: obligations expressed as machine-readable intents, compliance verified continuously, and human authority preserved where it matters most. STORES is an early blueprint for that transition.

© 2026 R. W. Charara and K. M. Chatila

Appendix A. Illustrative Model Provisions

The provisions below illustrate how the safeguards in Section 7 could be expressed in licence conditions or secondary legislation. Provisions that grant emergency powers or limit liability (A.2, A.5 and A.7) normally require primary legislation and cannot rest on licence conditions alone. They are drafting aids only and must be adapted by legal counsel to national law and practice. Square brackets indicate values to be set nationally.

A.1 Regulatory data access

The Licensee shall make available to the Authority, through the interfaces and common data model specified by the Authority, the network performance and service quality data listed in Schedule [X], at the frequency and granularity specified therein. Personal data shall not be provided except where strictly necessary for a purpose specified in Schedule [X] and permitted under [national data-protection law].

A.2 Regulatory measurement nodes

The Licensee shall host a measurement node conforming to the Authority's published specification. The node shall have read-only access to agreed data exports, shall have no capability to modify network elements, and shall run only software approved through the joint change procedure. The Licensee may isolate the node where necessary to protect network integrity or security, and shall notify the Authority within [24] hours. The Licensee shall not be liable for loss caused by the node operating in accordance with the specification.

A.3 Digital evidence

Measurement records signed at source, time-stamped in accordance with [standard] and retained in the Authority's evidence store may be relied upon in proceedings under this Licence. Before any finding is made, the Licensee shall be given access to the records, methods, calibration certificates and model documentation relied upon, and a period of [30] days to respond, including by submitting its own measurements or requesting independent re-measurement.

A.4 Automated assistance and human review

No decision producing legal effects for the Licensee shall be taken solely on the basis of automated processing. Every such decision shall identify the official who approved it and the evidence relied upon. The Licensee may request review of a finding by an official not involved in making it, without prejudice to its right of appeal.

A.5 Emergency coordination

During an emergency declared under [law], the Authority may, in coordination with [the national emergency authority], request the Licensee to apply priority to designated emergency and critical services. Such measures shall be applied in accordance with [national open-internet rules, for example Article 3(3) of Regulation (EU) 2015/2120], limited to what is necessary, shall lapse after [72] hours unless renewed, and shall be reviewed after the emergency. A Licensee acting in good faith and within the scope of such a request shall, to the extent permitted by law, not be held liable for that action.

A.6 Confidentiality and competition

The Authority shall not disclose to any Licensee information concerning another Licensee that is commercially sensitive, and shall publish comparative indicators only at a level of aggregation and delay that prevents inference of a Licensee's network plans or commercial position.

© 2026 R. W. Charara and K. M. Chatila

A.7 Data-access and node agreement (key terms)

The agreement between the Authority and each Licensee shall cover at least: (a) the data in scope, its format and frequency; (b) ownership and funding of the node, with hardware and software funded by the Authority and integration effort borne by the Licensee; (c) custody of evidence-signing keys in a hardware security module controlled by the Authority; (d) security certification and patching of the node; (e) liability, under which each party is responsible for its own negligence, the Licensee is not liable for loss caused by the node operating to specification, and the Authority's liability follows [public law]; (f) incident notification within [24] hours; (g) audit rights for both parties; and (h) exit, including deletion or return of data.

A.8 Purpose limitation and oversight

Data and systems established under this [Regulation] shall be used only for the purposes listed in Schedule [Y]. An independent oversight board shall review access logs and use cases and publish an annual transparency report. This [Regulation] shall be reviewed every [five] years.

© 2026 R. W. Charara and K. M. Chatila

Appendix B. STORES Readiness Self-Assessment

Regulators can use this checklist to locate their current position and plan next steps. For each dimension, identify the level that best describes current practice; the next level is the natural target for the coming phase.

Table 18. STORES readiness self-assessment

Dimension	Level 1: Basic	Level 2: Developing	Level 3: Advanced	Level 4: Leading
Mandate and law	Periodic reporting duties only	Data-access powers in some licences	Continuous access, nodes and evidence rules in law	Full Section 7 safeguards, oversight board and review cycle
Measurement	Operator-reported KPIs	Periodic drive tests or probes	Continuous independent and crowdsourced QoE	Slice-aware, intent-based measurement with published methods
Data fabric	Spreadsheets and ad hoc files	Central database of operator returns	Common data model; operator nodes; federated analytics	Federated learning with privacy budgets and protected aggregation
Evidence	Unsigned reports	Some signed or archived data	Signed, time-stamped records with calibration	Transparency log; contestable case files; evidence accepted on appeal
Analytics and AI	Manual analysis	Dashboards and rule-based alerts	AI-assisted findings at R1 with decision records	Audited R2 automation; R3 playbooks; annual autonomy review
Cross-sector	No formal links	Informal contacts in emergencies	MoUs and read-only feeds with key sectors	NROC exercises with power, fibre and emergency authorities
People	Generalist staff	Small data team	Multidisciplinary team incl. legal and security	Full NROC staffing; continuous training; challenge reviews

© 2026 R. W. Charara and K. M. Chatila

Appendix C. Economic Modelling with Input-Output Methods

The economics of the sector make better oversight more important, not less. PwC forecasts global telecom revenue to grow by about 2.9% a year to 2028, below inflation[3], while investment needs for 5G-SA, fibre and resilience remain high. At the same time connectivity has large spillovers: ITU econometric work finds that a 10% increase in mobile broadband penetration is associated with an average 1.5% increase in GDP, with larger effects in less developed economies[48]. Capacity constraints carry real costs: a 2025 study by Accenture for the US wireless industry association CTIA estimated that a mid-band spectrum shortfall could cost the United States more than USD 1.4 trillion in lost economic growth by 2035[49]. Such figures are market-specific, and industry-sponsored estimates should be read as such, but they show why regulatory decisions on spectrum, coverage and resilience deserve quantitative support.

C.1 The Leontief input-output model

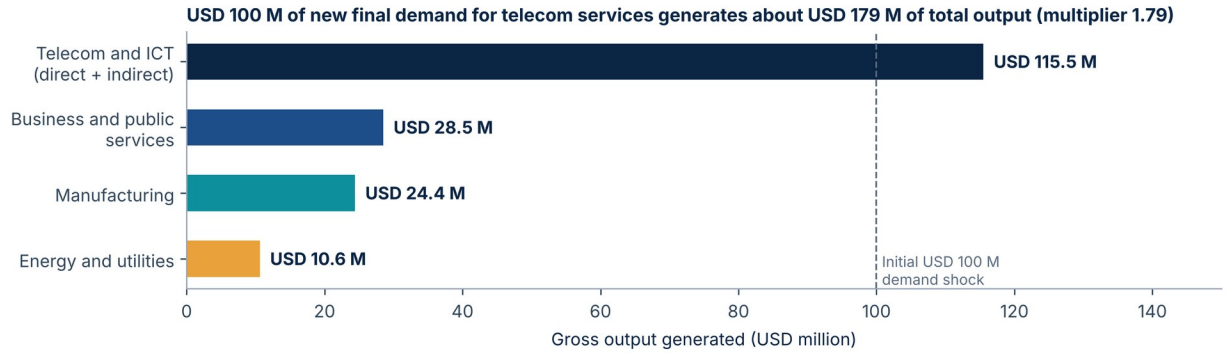
The Leontief input-output model captures how sectors buy from and sell to one another[50]. With n sectors, let **x** be the vector of total output, **A** the matrix of technical coefficients, where a_{ij} is the input from sector i required per unit of output of sector j, and **f** the vector of final demand. Total output must cover both intermediate use and final demand:

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{f} \quad \Longrightarrow \quad \mathbf{x} = (\mathbf{I} - \mathbf{A})^{-1}\,\mathbf{f} = \mathbf{L}\,\mathbf{f} \tag{1}$$

$\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$ is the Leontief inverse. Its element l_{ij} gives the total output required from sector i, directly and indirectly, per unit of final demand for sector j. The output multiplier of sector j is the column sum:

$$m_j = \sum_i l_{ij} \tag{2}$$

Figure 13 applies the method to a simplified four-sector economy with illustrative coefficients. An additional USD 100 million of final demand for telecom services generates about USD 179 million of total output, as telecom purchases equipment, services and energy, and those suppliers in turn purchase their own inputs. In practice, coefficients come from the national input-output tables published by statistics offices.



Illustrative four-sector model with hypothetical technical coefficients; for explaining the method only, not an estimate for any country.

Figure 13. Illustrative output multiplier of telecom investment in a four-sector economy.

© 2026 R. W. Charara and K. M. Chatila

C.2 The dynamic extension

The dynamic Leontief model adds a capital coefficients matrix **B** that links investment to the growth of production capacity between periods t and t+1:

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_t + \mathbf{B}\left(\mathbf{x}_{t+1} - \mathbf{x}_t\right) + \mathbf{f}_t \tag{3}$$

This lets planners simulate capacity over time, for example how fibre and 5G-SA investment today expands the productive capacity of digitally intensive sectors tomorrow. Embedded in the NIDT, such models help compare policy options by their estimated economic return and not only by their cost to operators. They estimate relationships under stated assumptions; they should inform, not replace, cost-benefit analysis and regulatory impact assessment.

© 2026 R. W. Charara and K. M. Chatila

About the Authors

Raghid W. Charara

Raghid W. Charara is a technology executive with more than 27 years of experience across telecommunications, technology and banking, including C-level roles in Middle East banks. As Vice President of Analitico International, he has led national-scale telecom programmes, including device-identity registers, quality-of-service monitoring and international voice gateways, and has advised telecom regulators on QoS, anti-fraud and big-data analytics, leading control-centre projects for telecom authorities. He is a member of faculty at INSEAD and IE Business School, where he teaches AI, innovation, digital transformation, strategy and FinTech, and has trained more than 2,000 executives in AI across governments, banks and investment funds. He holds a double major in Statistics and Computer Science from the American University of Beirut, more than a dozen INSEAD executive education certificates, including "Leading in a Transforming World" and "Innovation in the Age of Disruption", and the AI+ Chief AI Officer and AI+ Executive certifications from AI CERTs. He is pursuing a Doctorate in Business Administration (DBA) focused on AI and machine learning.

Contact: raghid.charara@insead.edu

Khaled M. Chatila

Khaled M. Chatila is an entrepreneur and senior international advisor with over 25 years of leadership experience spanning technology, telecommunications, transport, industrial operations and international business development. Throughout his career, he has led multinational operations involving more than 1,000 personnel across international markets, overseeing complex technology programmes, large-scale turnkey projects and strategic business initiatives for major organisations. As CEO of UNIEQUIP LTD, he provides strategic leadership and advisory services to organisations and governments, focusing on digital transformation, artificial intelligence, operational excellence and technology-driven growth. His expertise combines business strategy, commercial leadership, technology integration and the execution of complex, multidisciplinary programmes. He holds a B.Sc. in Computer Science from the American University of Beirut and an M.Sc. in Information Security from Royal Holloway, University of London.

Contact: kc@uniequip.ltd

© 2026 R. W. Charara and K. M. Chatila

Glossary of Terms

Term	Meaning
3GPP	3rd Generation Partnership Project, the body that specifies mobile network standards (4G, 5G)
5G-SA	5G Standalone: 5G radio with a native 5G core, enabling slicing and service-based features
AX	Agentic Experience: service quality as experienced by AI agents and machines rather than humans
BMS	Battery Management System, reporting state of charge (SoC) and state of health (SoH)
CAMARA	Open-source project defining common network APIs, used by GSMA Open Gateway
CEIR	Central Equipment Identity Register: national register of device identifiers (IMEI)
DAS	Distributed Acoustic Sensing: using optical fibre as a continuous vibration and temperature sensor
DP / DP-SGD	Differential privacy; differentially private stochastic gradient descent
FTE	Full-time equivalent (staffing measure)
FL	Federated learning: training shared models without centralising raw data
MoU	Memorandum of understanding between authorities
MOS	Mean Opinion Score: a 1 to 5 scale of perceived quality
NIDT	National Infrastructure Digital Twin
NROC	National Regulatory Operations Centre: the cross-sector coordination hub of STORES
NWDAF	Network Data Analytics Function of the 5G core (3GPP TS 23.288)
ODA	TM Forum Open Digital Architecture
OSS / BSS / NMS	Operations Support, Business Support and Network Management Systems
OTDR	Optical Time Domain Reflectometry: locating faults along an optical fibre
QoS / QoE	Quality of Service (network performance) / Quality of Experience (user-perceived quality)
R0 to R4	The five STORES regulatory autonomy levels (Section 6.2)
RFC	Request for Comments: an Internet standard published by the IETF
SCADA	Supervisory Control and Data Acquisition systems used in power grids
STORES	Sovereign Telecom Orchestration, Regulation and Enforcement System
TEE	Trusted execution environment: hardware-protected enclave used in confidential computing
Transparency log	Append-only, cryptographically verifiable log that reveals any deletion or alteration
XAI	Explainable AI
ZKP	Zero-knowledge proof: proving a statement is true without revealing the underlying data

© 2026 R. W. Charara and K. M. Chatila

References

Each reference is labelled by source type so that readers can weigh its evidential value: standard, law, official regulator, government or intergovernmental source, peer-reviewed research, academic text, industry source or industry standards body, and secondary reporting.

[1] [Industry data, secondary reporting] Ookla (2026). Europe mobile peak-hour congestion analysis, Q1 2026. https://www.ookla.com/articles/europe-mobile-peak-hour-congestion-2026; reported by ISPreview, "Ookla examine impact of peak time UK and EU mobile network congestion", 4 May 2026. https://www.ispreview.co.uk/index.php/2026/05/ookla-examine-impact-of-peak-time-uk-and-eu-mobile-network-congestion.html

[2] [Intergovernmental source] International Energy Agency (2025). Energy and AI. Paris: IEA. https://www.iea.org/reports/energy-and-ai/executive-summary

[3] [Industry source] PwC (2025). Global telecoms industry revenue rose 4.3% in 2023 to hit US$1.1 trillion, but faces continuing headwinds (Global Telecoms Outlook), press release, 28 February 2025. https://www.pwc.com/gx/en/news-room/press-releases/2025/pwc-global-telecoms-outlook.html

[4] [Official regulator source] Arcep (2022). Mobile coverage and quality of service: crowdsourced measurements added to the "Mon réseau mobile" map, press release, 11 February 2022. https://en.arcep.fr/news/press-releases/view/n/mobile-coverage-110222.html

[5] [Official regulator source] Ofcom (2026). Mobile Matters: annual report on mobile data performance using crowdsourced measurements. https://www.ofcom.org.uk/phones-and-broadband/mobile-phones/mobile-matters

[6] [Official regulator source] Federal Communications Commission, Broadband Data Collection Help Center. What to Expect after Filing a Mobile Challenge. https://help.bdc.fcc.gov/hc/en-us/articles/10472720221211-What-to-Expect-after-Filing-a-Mobile-Challenge; FCC Speed Test App User Guide. https://www.fcc.gov/sites/default/files/FCC_Speed_Test_App_User_Guide.pdf

[7] [Standard] ITU-T Recommendation E.800 (09/2008). Definitions of terms related to quality of service. International Telecommunication Union. https://www.itu.int/rec/T-REC-E.800

[8] [Industry source] GSMA. Open Gateway: open network APIs for developers. https://www.gsma.com/solutions-and-impact/gsma-open-gateway/; see also the CAMARA project (Linux Foundation). https://camaraproject.org/

[9] [Official regulator source] BEREC (2018). Net neutrality measurement tool: result of the tender, 6 August 2018. https://www.berec.europa.eu/en/latest-news/net-neutrality-measurement-tool-result-of-the-tender

[10] [Official government source] Press Information Bureau, Government of India (2025). Sanchar Saathi App: Telecom Empowerment at Citizens' Fingertips, 2 December 2025. https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/dec/doc2025122715201.pdf

[11] [Official regulator source] BEREC (2018). BEREC Common Position on information to consumers on mobile coverage, BoR (18) 237, 6 December 2018. https://www.berec.europa.eu/en/document-categories/berec/regulatory-best-practices/common-approachespositions/berec-common-position-on-information-to-consumers-on-mobile-coverage

[12] [Standard] ITU-T Recommendation P.1203 (10/2017). Parametric bitstream-based quality assessment of progressive download and adaptive audiovisual streaming services over reliable transport. https://www.itu.int/rec/T-REC-P.1203

[13] [Standard] 3GPP TS 23.288. Architecture enhancements for 5G System (5GS) to support network data analytics services (NWDAF). https://www.3gpp.org/DynaReport/23288.htm

[14] [Standard] ITU-T Recommendation Q.5038 (04/2026). Signalling and protocol for intent-based management and orchestration of network slicing in IMT-2020 networks and beyond. https://www.itu.int/rec/T-REC-Q.5038/en

[15] [Industry source] GSMA (2021). Global Mobile Operators and GSMA Team up to Build Blockchain to Transform Back-Office Roaming, press release, June 2021. https://www.gsma.com/newsroom/press-release/global-mobile-operators-and-gsma-team-up-to-build-blockchain-to-transform-back-office-roaming/

© 2026 R. W. Charara and K. M. Chatila

[16] [Peer-reviewed research] Zhu, L., Liu, Z. and Han, S. (2019). Deep Leakage from Gradients. Advances in Neural Information Processing Systems (NeurIPS). https://arxiv.org/abs/1906.08935

[17] [Peer-reviewed research] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A. and Seth, K. (2017). Practical Secure Aggregation for Privacy-Preserving Machine Learning. Proceedings of ACM CCS 2017. https://eprint.iacr.org/2017/281

[18] [Peer-reviewed research] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K. and Zhang, L. (2016). Deep Learning with Differential Privacy. Proceedings of ACM CCS 2016. https://doi.org/10.1145/2976749.2978318

[19] [Peer-reviewed research] McMahan, H. B., Ramage, D., Talwar, K. and Zhang, L. (2018). Learning Differentially Private Recurrent Language Models. International Conference on Learning Representations (ICLR). https://arxiv.org/abs/1710.06963

[20] [Industry source] Confidential Computing Consortium (Linux Foundation). About confidential computing. https://confidentialcomputing.io/about/

[21] [Standard] Adams, C., Cain, P., Pinkas, D. and Zuccherato, R. (2001). Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP). IETF RFC 3161. https://www.rfc-editor.org/rfc/rfc3161

[22] [Law] Regulation (EU) No 910/2014 (eIDAS), Article 41: Legal effect of electronic time stamps. https://eur-lex.europa.eu/eli/reg/2014/910/oj

[23] [Standard] ISO/IEC 17025:2017. General requirements for the competence of testing and calibration laboratories. https://www.iso.org/standard/66912.html

[24] [Standard] Laurie, B., Messeri, E. and Stradling, R. (2021). Certificate Transparency Version 2.0. IETF RFC 9162, December 2021. https://www.rfc-editor.org/rfc/rfc9162

[25] [Standard] ITU-T Recommendation Y.3192 (12/2025). Requirements and functional framework for customer-oriented network quality auto-optimization with artificial intelligence. https://www.itu.int/rec/T-REC-Y.3192/en

[26] [Industry standards body] TM Forum. Autonomous Networks: Empowering digital transformation, evolving from Level 2/3 towards Level 4 (IG1326). https://www.tmforum.org/resources/standard/ig1326-autonomous-networks-empowering-digital-transformation-evolving-from-level-2-3-towards-level-4; TM Forum (2026). Autonomous Networks Leadership Forum spotlights Level 4 shift from cost-cutting to revenue growth. https://tmforum.org/news-insight/newsroom/autonomous-networks-leadership-forum-spotlights-level-4-shift

[27] [Standard] ETSI GS ZSM 002. Zero-touch network and Service Management (ZSM); Reference Architecture. European Telecommunications Standards Institute. https://www.etsi.org/deliver/etsi_gs/ZSM/001_099/002/01.01.01_60/gs_zsm002v010101p.pdf

[28] [Standard] IEC 61850 series. Communication networks and systems for power utility automation. International Electrotechnical Commission.

[29] [Secondary reporting] Hassan, N. (2024). Saudi Arabia announces national plan to stimulate urban policies by 2025 (national Digital Twin model). Arab News, 12 February 2024. https://www.arabnews.com/node/2458586/business-economy

[30] [Peer-reviewed research] Bahreldin, I., Samir, H., Maddah, R., Hammad, H. and Hegazy, I. (2025). Leveraging advanced digital technologies for smart city development in Saudi Arabia: opportunities, challenges, and strategic pathways. International Journal of Low-Carbon Technologies, 20, 834-847. https://doi.org/10.1093/ijlct/ctaf044

[31] [Law] Regulation (EU) 2024/1689 of 13 June 2024 (Artificial Intelligence Act), Article 14: Human oversight. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj; article text: https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-14

[32] [Official government framework] NIST (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0), 26 January 2023. https://www.nist.gov/itl/ai-risk-management-framework

[33] [Standard] ISO/IEC 42001:2023. Information technology, Artificial intelligence, Management system. International Organization for Standardization. https://www.iso.org/standard/42001

[34] [Industry source] OWASP Foundation (2025). OWASP Top 10 for Large Language Model Applications. https://genai.owasp.org/llm-top-10/

© 2026 R. W. Charara and K. M. Chatila

[35] [Official government standard] Rose, S., Borchert, O., Mitchell, S. and Connelly, S. (2020). Zero Trust Architecture. NIST Special Publication 800-207, August 2020. https://csrc.nist.gov/pubs/sp/800/207/final

[36] [Standard] ITU-T Recommendation X.1058 (02/2026). Controls, requirements and guidance for personally identifiable information protection (aligned with ISO/IEC 29151). https://www.itu.int/rec/T-REC-X.1058-202602-I/en

[37] [Law] Directive (EU) 2018/1972 establishing the European Electronic Communications Code, Article 20: Information request to undertakings. https://eur-lex.europa.eu/eli/dir/2018/1972/oj

[38] [Law] Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III: High-risk AI systems referred to in Article 6(2). https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3

[39] [Law] Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 27: Fundamental rights impact assessment for high-risk AI systems. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-27

[40] [Law] Regulation (EU) 2016/679 (General Data Protection Regulation). https://eur-lex.europa.eu/eli/reg/2016/679/oj

[41] [Law (secondary summary)] State of Qatar, Law No. (13) of 2016 on Personal Data Privacy Protection; summary in DLA Piper, Data Protection Laws of the World: Qatar. https://www.dlapiperdataprotection.com/index.html?c=QA&t=law

[42] [Law] Directive 2002/58/EC (Directive on privacy and electronic communications), Articles 6 and 9 on traffic and location data. https://eur-lex.europa.eu/eli/dir/2002/58/oj

[43] [Law] Regulation (EU) 2015/2120 laying down measures concerning open internet access, Article 3(3). https://eur-lex.europa.eu/eli/reg/2015/2120/oj

[44] [Standard] ITU-T Recommendation Y.3149 (08/2025). Quality of service assurance requirements for the tactile internet. https://www.itu.int/rec/T-REC-Y.3149-202508-I/en

[45] [Standard] ITU-T Recommendation M.3370 (03/2025), amended 03/2026. Telecommunication preventive maintenance task: Overview. https://www.itu.int/rec/T-REC-M.3370

[46] [Standard] 3GPP TS 28.312. Management and orchestration; Intent driven management services for mobile networks. https://www.3gpp.org/DynaReport/28312.htm

[47] [Industry standards body] TM Forum. Open Digital Architecture (ODA) and Open APIs. https://www.tmforum.org/open-digital-architecture/

[48] [Intergovernmental source] ITU (2018). The economic contribution of broadband, digitalization and ICT regulation. Geneva: International Telecommunication Union. https://www.itu.int/pub/D-PREF-EF.BDR-2018

[49] [Industry-sponsored study, secondary reporting] Accenture for CTIA (2025). Securing the Future of U.S. Wireless Networks: The Looming Spectrum Crisis; as reported by Telecompetitor, 28 March 2025. https://www.telecompetitor.com/spectrum-shortage-threatens-u-s-gdp-and-network-capacity-study/

[50] [Academic text] Miller, R. E. and Blair, P. D. (2009). Input-Output Analysis: Foundations and Extensions, 2nd edition. Cambridge University Press; Leontief, W. (1986). Input-Output Economics, 2nd edition. Oxford University Press.

© 2026 R. W. Charara and K. M. Chatila