


ENGINEERING / OCTOBER 2026

The AI telecom operating model

Make autonomy earn its place in network and service economics



THE DECISION IN BRIEF

Telecom operators should fund AI autonomy where measured network value survives service constraints, interacting control loops and the full cost of recovery.

- 01

Local optimization can shift cost or degrade service elsewhere. Establish the shared operating boundary before allowing energy, service and network agents to coordinate actions.
- 02

Observe, recommend, execute within limits and coordinate across domains are different delegations. Each requires evidence of useful outcomes, controllable interaction and recovery; they are not maturity rankings.
- 03

Keep electricity, emissions and financial return separate. Scale a domain only when service quality holds and the net-value account includes integration, computing, support and exception capacity.

THE MAYDEN VIEW


An energy saving in one cell is not a network saving until the wider service boundary is measured.

INSIDE THIS PERSPECTIVE

The argument at a glance	3
01 / Choose the network decision and its value	4
Exhibit 01 / Test the energy loop and the service loop together	5
02 / Bound autonomy and assess evidence	6
Exhibit 02 / Each expansion of autonomy creates a new proof obligation	7
03 / Protect service while pursuing efficiency	8
Exhibit 03 / Industry emissions progress does not identify an AI effect	9
04 / Run a coupled trial and decide transfer	10
The leadership agenda	11
Method and references	12

ONE DECISION, FOUR CONNECTED QUESTIONS

Bound the change to the network



CONTROLLED SIGNAL FIELD / CONCEPTUAL ARTWORK

01 Choose the network decision and its value

Prove the value across interacting control loops

02 Bound autonomy and assess evidence

Each operating mode needs different evidence

03 Protect service while pursuing efficiency

Energy is a real opportunity, but emissions are not an AI score

04 Run a coupled trial and decide transfer

Design the trial so neighboring cells cannot hide the cost

An energy saving in one cell is not a network saving until the wider service boundary is measured.

CHAPTER 01

01

Choose the network decision and its value

Prove the value across interacting control loops

An energy controller puts a radio resource to sleep while a service controller adds capacity nearby. Each can satisfy its own objective while the combined network becomes more expensive or less reliable. This illustrative conflict captures the investment challenge for AI in telecoms: value must survive interaction across the operating boundary, not merely appear inside one control loop.

Start with a domain where an accountable owner can connect an action to an observable outcome. Specify the baseline, controllable decision and service constraints. A prediction creates value only if someone can act on it in time; an automated action creates value only if its wider effects remain acceptable. These links should define the investment case.

The industry's standards work supports a systems perspective. ETSI's zero-touch network and service management programme addresses end-to-end automation across domains and layers. ^[1] It describes an architectural direction, not evidence that any particular operator has achieved reliable full autonomy. The CEO should therefore ask how the proposed investment changes an operating decision today and what additional conditions must hold before that decision can be delegated more broadly.

Separate three sources of economic value

Network efficiency concerns capacity, energy and fault response within service obligations. Define the domain and observation period, recording traffic, weather, configuration and service conditions. Lower energy expenditure may reflect tariffs or demand rather than better optimization. Test transfer to different equipment and operating conditions before multiplying a local result across the estate.

The second source is service productivity. AI may help staff diagnose issues, find policy information or coordinate a field visit. Measure resolution and repeat contact, not just the speed of the first response. A chatbot that reduces contact-centre volume by making escalation difficult can look efficient while worsening customer outcomes. Similarly, a faster field schedule can fail economically if it increases repeat visits or allocates the wrong skills. The complete service chain matters more than the local task metric.

The third source is revenue and investment quality. Better demand forecasting, enterprise solution design or retention targeting may improve commercial decisions. These claims need careful attribution because pricing, competition and product changes occur simultaneously. Where controlled trials are feasible and appropriate, use them. Where they are not, disclose the weaker inference. The enterprise should avoid combining speculative revenue uplift with well-measured cost savings in one headline figure that conceals the different confidence levels.

ILLUSTRATIVE DECISION TOOL

Test the energy loop and the service loop together

Proposed radio-cluster experiment; no operator results

Energy controller proposes sleep	Joint measurement boundary Target cell plus affected neighbors	Condition before expansion Service constraints hold; a local saving is insufficient
Service controller requests capacity	Joint measurement boundary Same window, demand pattern and equipment state	Condition before expansion Named coordinator resolves conflicting actions
Controlled intervention runs	Joint measurement boundary Energy, quality tails, handovers and added compute/support	Condition before expansion Record confounders and failures; retain adverse periods
Rollback or transfer	Joint measurement boundary Restored state across both loops; retained telemetry	Condition before expansion Prove recovery; retest new equipment or traffic conditions
THE MANAGEMENT IMPLICATION		
Expand only when the joint system meets service, energy and recovery conditions, not when one controller wins its local metric.		

Original mayden test design informed by closed-loop and resilience sources. No energy saving, ROI, universal threshold or verified vendor performance is claimed. Sources: [1], [2], [6], [8]

CHAPTER 02

Bound autonomy and assess evidence

02

Each operating mode needs different evidence

mayden proposes four operating modes. Observation establishes evidence and a baseline. Recommendation tests whether advice changes useful decisions. Bounded execution delegates a narrow action with exposure limits and rollback. Cross-domain coordination governs interactions among loops. Moving between them changes the proof obligation; it is not a prescribed climb toward maximum autonomy.

ETSI's January 2026 study of agents in autonomous networks distinguishes resource, policy and goal conflicts, and discusses delegation, escalation and arbitration. ^[8] This is a standards-group study, not evidence that multi-agent networks deliver a particular saving. Its practical contribution is to make coordination a separate design obligation: capable agents can still pursue incompatible objectives on the same network.

For the energy-and-service example, define which constraints prevail, who arbitrates conflicting proposals and how the combined action is reversed. Test that conflict before connecting autonomous domains. Measure the shared outcome, including neighboring cells: adding local improvements does not establish network value.

ETSI's 2026 report on closed-loop security examines risks across monitoring, analysis, decision and execution, including interactions among loops. ^[2] The 3GPP Release 18 description documents AI/ML and management-related work within the wider mobile-system evolution. ^[3] Neither source is a benchmark of commercial deployment quality. Their practical implication is architectural: procurement should examine interfaces,

lifecycle management and control interaction, not merely model accuracy or a vendor's claim to autonomy.

Keep standards, research and vendor evidence in their proper roles

Standards provide common concepts and interfaces. Research identifies possibilities and failure modes. A supplier demonstration shows what occurred under specified conditions. An operator's measured production trial provides evidence about its own environment. These categories complement one another, but they should not be substituted. A claim that a product follows a standard does not establish its financial return, and a strong trial result does not establish interoperability with every future system.


The ITU focus group on AI-native telecommunications networks was established to explore architectural requirements, challenges and opportunities. ^[4] Its existence is evidence that the architecture remains an active area of work. It is not proof that the industry has settled one complete operating model. Operators should preserve the ability to evolve components and avoid committing core control authority to interfaces they cannot inspect, test or replace. The cost of integration and exit belongs in the investment case.

Require vendors to identify the baseline, population, operating conditions and period behind performance claims. Separate simulation from live trials and customer reports from independent verification. Ask for logs and reproducible failure conditions. Where disclosure is constrained, narrow the claim and the delegation it can support.

ORIGINAL MAYDEN FRAMEWORK

Each expansion of autonomy creates a new proof obligation

mayden analytical model



THE MANAGEMENT IMPLICATION

Local model quality does not prove safe interaction between control loops.

Not an industry maturity score or certification. Sources: [1], [2], [3]

CHAPTER 03

03

Protect service while pursuing efficiency

Energy is a real opportunity, but emissions are not an AI score

The GSMA's 2026 report records a 13% reduction in operators' operational emissions between 2019 and 2024 while connections grew 10% and traffic more than quadrupled. Its coverage includes 116 operators representing 85% of mobile connections. ^[5] These figures describe an industry trend assembled from disclosures. They do not isolate the contribution of AI. Changes in energy sourcing, equipment, traffic mix and other operating practices can contribute to the observed result.

For an operator, the relevant trial should distinguish electricity saved from emissions accounting and financial savings. Electricity is a physical quantity; the bill depends on tariffs and contracts; reported emissions depend on the accounting method and energy sources. A reduction in one does not establish the same percentage reduction in the others. Keep them in separate lines of the business case. Where a network intervention shifts demand to another site or increases computing elsewhere, include those effects in the chosen boundary.

mayden recommends an energy-and-service acceptance test. Specify the permitted change in service quality before running the optimization. Compare energy, coverage, latency, reliability and customer complaints over representative operating conditions. Include the additional inference and orchestration cost. An intervention that saves energy while degrading a critical service is not automatically a good trade. The operator must decide the trade-off explicitly, within its obligations, rather than letting an optimization objective make that decision invisibly.

Resilience is the condition for economic scale

More automation can reduce some human errors while increasing the speed and reach of a mistaken action. That does not justify avoiding automation; it changes the requirements for containment. A control loop should have bounded permissions, observable actions and a clear way to suspend or revert changes. The operator also needs staff who understand the fallback. A manual procedure that has not been exercised may not work when many sites or customers are affected at once.

Ofcom's network and service resilience guidance provides a UK-specific regulatory context for communications providers. ^[6] Its relevance here is the focus on resilient operation, not a universal rulebook for every telecom market. The NCSC's AI guidance separately reinforces lifecycle security. ^[7] mayden's operational extension is to test the full recovery path: detect anomalous behavior, stop further actions, restore an acceptable configuration and communicate with the affected service teams. A model rollback alone may not reverse changes already executed elsewhere.

Dependency concentration deserves board attention. Several domains may depend on the same data platform, orchestration layer or external model. Individually reasonable project choices can create an estate-wide dependency that no project owner sees. Maintain a portfolio view of shared services and correlated failure. Test degraded modes under constrained capacity, including the increased workload on people. The resilience case should describe what service remains available and for how long, rather than offer an undifferentiated claim that a backup exists.

REPORTED EMPIRICAL RESULT · DERIVED INDEX


Industry emissions progress does not identify an AI effect

Index: 2019=100

2019 baseline



2024 operational emissions



THE MANAGEMENT IMPLICATION

Use industry context to frame opportunity, then measure the intervention in the operator's own network.

mayden calculation from GSMA's reported 13% reduction 2019-2024. Operational emissions, not electricity use. No causal attribution to AI. The GSMA report draws on 116 operators covering 85% of connections. Sources: [5]

CHAPTER 04

Run a coupled trial and decide transfer 04

Design the trial so neighboring cells cannot hide the cost

In an illustrative radio-cluster energy trial, include affected neighboring cells in the measurement boundary. A sleeping resource can move traffic and energy use elsewhere. Define the cluster, adjacent resources, observation windows and service stop conditions before intervention. Otherwise the trial can reward cost displacement rather than improvement.

Make the interaction concrete: an energy controller requests that a resource sleep while a service-quality controller requests additional capacity in an adjacent cell. Neither local objective should silently override the other. A named coordinating function applies the agreed service constraints and permits, defers or reverses the action. Record both proposals, the governing rule and the observed result across the shared boundary. If coordination is unavailable, use the previously approved restricted mode rather than let competing loops improvise a new priority. This is an illustrative test condition, not a claim about a deployed operator architecture.

Where feasible, compare intervention and reference periods or clusters with similar traffic and equipment. Record events that break comparability, including maintenance, special events, software upgrades and unusual weather. A perfectly matched control may be unavailable, in which case the report should state the remaining attribution limits. Do not discard adverse periods merely because they complicate the result. Network leaders need to know whether an optimization is useful under the demand conditions that make their service difficult to operate.

The acceptance record should then separate three decisions. First, whether the energy effect

is credible within the tested boundary. Second, whether service performance remains acceptable, including at the distribution's extremes. Third, whether the financial return survives implementation, computing and support costs. Passing one test does not imply passing the others. Expansion to another equipment family or traffic pattern is a transfer question that requires additional evidence. This approach gives the CTO and CFO a shared way to challenge a vendor's savings claim while preserving a route to scale when the measured result is persuasive.

Tie the next delegation to the next investment

Some value may require simultaneous change across several domains. Preserve that strategic hypothesis while testing the next commitment. Prove local control, interfaces and arbitration before connecting their authority. Premature integration can make both attribution and recovery harder.

Choose one material service or cost problem and an accountable domain owner. Agree the baseline, intervention and acceptance criteria with finance, operations, security and customer teams. Where appropriate, start in recommendation mode and record rejected advice. Grant limited execution when the evidence supports it, rather than treating advisory accuracy as permission to act.

At the investment gate, reconcile observed benefit with implementation, computing, support, exceptions and recovery capacity. Test sensitivity to traffic, supplier price and equipment differences. Then scale, redesign, retain supervision or stop. The chosen delegation should follow demonstrated service and economic performance.

FROM INSIGHT TO ACTION

Decisions to take forward

NETWORK SERVICE OWNER

Agree the conditions that stop or restrict an intervention.

Evidence to require: Target/neighboring-cell boundary, traffic baseline and service constraints

CONTROL-SYSTEM AND RESILIENCE OWNERS

Withhold wider authority where controllers cannot preserve joint constraints.

Evidence to require: Conflicting-loop, coordination-loss and rollback exercises

CTO AND FINANCE OWNER

Expand within tested conditions; treat a new equipment or traffic regime as a fresh transfer decision.

Evidence to require: Whole-boundary energy, service and full operating-cost results

mayden connects senior judgement, research and engineering to turn strategic choices into working capabilities.

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Synthesis of eight opened standards, industry and regulator sources; cutoff 10 October 2026. Includes a transparent emissions index calculation and three original management exhibits.

Standards were used for architecture and lifecycle concepts, GSMA for descriptive industry context, and UK guidance for explicitly scoped resilience and security considerations.

No operator trial, causal estimate, vendor-product assessment or company ranking. GSMA disclosure-based emissions trends cannot identify AI's contribution. Proposed operating modes require local validation and jurisdiction-specific obligations.

Primary sources and author research

[1] Zero-touch network and Service Management

ETSI / Current programme page accessed 2026-10-10. Programme scope and end-to-end automation.

[Read the source](#)

[2] GR ZSM 017 V 1.1.1: Closed-Loop Automation Security aspects

ETSI / 2026-01. Scope and closed-loop security analysis; 37 pages.

[Read the source](#)

[3] Release 18 Description: Summary of Rel-18 Work Items

3GPP/ETSI / 2025-04. TR 21.918 V18.0.0, sections 11.2-11.3, printed pages 79-81.

[Read the source](#)

[4] Focus Group on Artificial Intelligence Native for Telecommunication Networks

ITU-T / Established 2024-07; page accessed 2026-10-10. Mandate and establishment.

[Read the source](#)

[5] Mobile Net Zero 2026: State of the Industry on Climate Action

GSMA / 2026-07. Executive summary; emissions and methodology.

[Read the source](#)

[6] Network and Service Resilience Guidance

Ofcom / 2024-09-05 statement; guidance updated 2026-06-16. Final statement and guidance.

[Read the source](#)

[7] Guidelines for secure AI system development

NCSC and partners / 2023-11-27; version 1.0. Deployment and operation.

[Read the source](#)

[8] ETSI GR ZSM 020 V1.1.1: Study on the Utilization of Agents in Autonomous Networks

ETSI / 2026-01. Clauses 6.2.2-6.2.4, printed pages 21-22; 7.1.5. Group Report rather than proof of interoperability, savings or operational certification.

[Read the source](#)