

ARTICLE 04

The Autonomous Operator

Intent, agents, digital twins and the governance of machine-operated networks



The Autonomous Operator

Intent, agents, digital twins and the governance of machine-operated networks

AUTHOR Moderion Advisory

SERIES The Total Connectivity Economy

EVIDENCE CURRENT TO 6 September 2026

EXECUTIVE PROPOSITION

Autonomy is not achieved by attaching a language model to operations. It requires a governed control architecture in which intent, data, models, digital twins, assurance, human authority and rollback work as one closed loop.

Executive summary

Autonomy is not achieved by attaching a language model to operations. It requires a governed control architecture in which intent, data, models, digital twins, assurance, human authority and rollback work as one closed loop.

This evidence-led thought article develops an executive point of view using standards, official data, company disclosures and Moderion Advisory analysis. It separates established facts from emerging direction and reported implementation.

EXECUTIVE QUESTION

How can operators move from bounded automation to production autonomy without surrendering safety, accountability or human authority?

Key insights at a glance

- TM Forum reported in 2026 that 75% of operators planned to increase autonomous-network investment and 81% aimed for Level 4 or above by 2030; these are intentions, not verified production outcomes.
- A TM Forum regional survey placed current maturity at 12% Level 0, 36% Level 1, 31% Level 2, 17% Level 3 and 4% Level 4, illustrating the distance between aspiration and present operating reality.
- Autonomy requires a closed governance loop: intent, context, planning, policy, authorization, pre-deployment verification, execution, assurance, learning and rollback.
- Incomplete inventory and inconsistent operational data remain first-order constraints. An agent cannot reliably reason about a network state that the organization cannot describe.
- Digital twins are most valuable as decision controls rather than presentation models. Their usefulness depends on state freshness, model fidelity, uncertainty bounds and correlation with live outcomes.
- Validation labels need careful interpretation. TM Forum states that solution validation does not independently certify that a live network performs at a claimed autonomy level without detailed investigation and testing.
- Brownfield evolution should use coexistence by design: legacy domains remain behind controlled adapters and domain services, decisions run in shadow mode before authority is delegated, and migration proceeds domain by domain with an exercised fallback path.

- TM Forum's maturity scale applies to bounded scenarios and domains; a Level 4 claim in one process does not establish enterprise-wide or end-to-end autonomy.
- O-RAN provides a useful RAN-domain intent path through SMO, Non-RT RIC, Near-RT RIC, rApps and xApps. Cross-domain operator autonomy additionally requires core, transport, cloud, service, security and BSS control loops.
- Conflict management is an explicit architectural function. Competing energy, capacity, mobility and quality policies require detection, precedence, arbitration, safe execution and post-action accountability.
- This article contributes to the Total Connectivity Economy by testing its enabling mechanism against universal reach, open interoperability, governed intelligence, trusted participation and sustainable economics.
- A 2025 IEEE Communications Surveys & Tutorials review supports intent-driven lifecycle management as a central 6G research direction, but identifies unresolved architecture, assurance and full-lifecycle challenges rather than a finished autonomous-network design.
- Peer-reviewed O-RAN experiments show that independently reasonable xApps can conflict and materially degrade performance. Conflict detection, sandboxing and policy arbitration are therefore safety functions, not optional optimization features.

Exhibit 1 | Autonomous-network maturity: current position and ambition

Level	Current survey	2026 target	2030 target
0	12%	0%	0%
1	36%	6%	0%
2	31%	25%	5%
3	17%	46%	10%
4	4%	23%	85%

Note: TM Forum regional guide survey. Targets are respondent ambitions, not audited forecasts.

Exhibit 2 | Brownfield migration protocol

State	Machine authority	Required controls	Exit evidence
1. Observe	Read and model only	Inventory, identity, telemetry, data quality	State is sufficiently complete and current
2. Shadow	Recommend; no production actuation	Decision comparison, twin/canary, conflict logging	Stable uplift with bounded error
3. Bounded execute	Reversible action in one domain	Policy, canary, rollback, human override	Production KPI and incident thresholds met
4. Federated autonomy	Cross-domain closed loop	Intent arbitration, end-to-end assurance, fallback domain	Sustained customer and risk outcomes

Note: Moderion Advisory synthesis informed by ETSI ZSM legacy-domain integration and progressive automation principles.

Exhibit 3 | Scenario-level autonomy progression

Capability	Level 2 - guided	Level 3 - conditional	Level 4 - bounded high autonomy
Intent	Human workflow target	Structured objective	Machine-testable objective and constraints
Decision	Rule or recommendation	Closed loop in defined cases	Predictive decision within signed policy
Scope	Single task	Single domain / scenario	Federated scenario with bounded authority
Recovery	Human correction	Automated rollback for known faults	Verified rollback plus exception escalation
Evidence	Task completion	Domain KPI	Outcome, side effects, confidence and audit trail

Note: Interpret maturity at scenario and domain level; do not extrapolate one validated use case to the full operator.

Core analysis

Automation is not autonomy

Telecom networks already automate many tasks, but most automation executes predefined rules inside bounded domains. Autonomy begins when the system interprets an objective, plans actions across domains, learns from observed outcomes and adjusts without waiting for a human instruction at every step. That shift is attractive because network complexity, service diversity and real-time expectations are increasing faster than operations teams can scale.

It is also dangerous to overstate. A successful laboratory agent, optimization model or closed-loop use case does not prove carrier-wide autonomy. Production networks contain legacy equipment, incomplete inventories, conflicting objectives, vendor-specific controls and obligations that cannot be reduced to a single reward function.

The governed intent-to-action loop

A credible architecture starts with intent expressed as measurable objectives and constraints: service quality, energy, cost, security, customer impact and regulatory boundaries. A planning layer translates intent into candidate actions. Policy and authority controls determine which actions may be executed automatically, which require approval and which are prohibited. The system then tests material changes against a current digital twin or safe simulation environment before controlled deployment.

Assurance closes the loop by comparing expected and observed outcomes. Every action needs provenance: the intent, data, model version, confidence, approver, affected resources and rollback condition. Without this evidence, autonomy becomes operational opacity.

Conflict is the central engineering problem

Open and programmable networks enable several optimization applications to act on shared resources. Research on O-RAN conflict management shows that individually rational xApps can interfere and degrade system performance. The same problem will grow when agentic systems optimize radio, transport, cloud, energy and customer experience simultaneously. The network therefore needs an independent governance and verification plane able to detect conflicting objectives, rank authority and constrain unsafe combinations.

Language models can improve diagnosis, knowledge retrieval and plan generation, but they should not be confused with deterministic control. Their outputs must be grounded in approved data, constrained by policy and validated against network state. High-impact execution should remain bounded by verifiable tools, thresholds and human accountability.

The operating model changes

The autonomous operator requires product, network, IT, security, data and risk teams to share a common control model. Operations engineers move from repetitive intervention toward policy design, exception management, model supervision and continuous verification. Vendors must expose observability and control interfaces without making accountability disappear across contracts.

Progress should be measured by outcomes: fewer incidents, faster recovery, lower energy, reduced change failure, improved service quality and lower cost per managed element. Maturity labels are useful only when tied to audited capabilities and production evidence.

Moderion Advisory perspective

The autonomy control stack

Intent should be expressed as measurable objectives and constraints, not a natural-language aspiration. A service objective can include availability, latency, energy, cost, security and customer impact with explicit precedence rules. Planning agents generate candidate actions; policy determines permissible actions; a digital twin or canary environment tests material changes; and assurance compares observed outcomes with both the target and unintended consequences.

Authority must be granular. Routine, reversible actions can be automated within bounded risk. Cross-domain or customer-impacting changes may require approval. Safety-critical or legally constrained actions should remain prohibited. Every level needs identity, provenance, reason capture and a tested rollback path.

A credible progression model

The path to autonomy should be organized by closed-loop scope rather than a corporate maturity slogan. Start with high-volume, bounded use cases whose state and outcome are observable: energy optimization, ticket classification, parameter correction or capacity balancing. Expand only after the data, conflict handling and rollback controls perform consistently.

The business case should combine avoided incidents, mean time to detect and restore, resource utilization, energy, labor reallocation and customer effect. Project-reported catalyst results can reveal possibility - for example, a TM Forum 2026 project reported planning reductions from four weeks to two hours and restoration below 60 minutes - but they require replication in the operator's own environment.

A brownfield coexistence protocol

ETSI's ZSM reference architecture explicitly provides for coordination with legacy management domains through domain-control services or a cross-domain integration fabric. This supports a practical transition: establish a trusted inventory and telemetry contract; wrap legacy managers behind governed adapters; normalize service intent and assurance data; and keep the legacy execution path available while the new control loop proves itself.

Authority should then progress through four states. In observe, the new system only builds state and recommendations. In shadow, it generates decisions that are compared with human or existing-system actions. In bounded execute, it controls reversible use cases inside a defined domain with canary deployment and automatic rollback. In federated autonomy, cross-domain actions are permitted only after conflict, dependency and customer-impact controls have passed. Every state has an explicit exit criterion and fallback owner.

Mapping guided automation to bounded Level 4 autonomy

At Level 2, systems execute selected tasks while humans define workflows and handle most cross-domain decisions. At Level 3, closed loops can analyze and act within defined scenarios, but human supervision and exception management remain material. Level 4 introduces intent-driven, predictive and continuously learning closed loops within a governed domain. The unit of assessment is therefore the scenario and its scope, not the corporate brand.

Progression requires evidence at each layer: machine-testable intent; complete and timely state; deterministic policy precedence; verified action plans; bounded authority; observable outcomes; automated rollback; and accountable exception ownership. A business target such as 'latency below 10

milliseconds for a factory service' is incomplete until the service boundary, measurement point, confidence interval, competing objectives and authorized remedies are defined.

Intent enforcement across RAN and operator domains

In an O-RAN example, the service intent reaches the SMO and Non-RT RIC, where policy and rApps can optimize over longer intervals. The Near-RT RIC uses xApps for faster control through the E2 interface. Candidate actions must still respect radio safety, admission control and vendor constraints. The resulting service outcome also depends on transport, user-plane placement, core policy, edge compute, security and enterprise application behavior.

The operator therefore needs a federated intent contract: every domain receives a measurable sub-intent, declares feasibility and confidence, exposes planned actions and returns assurance evidence. The parent orchestrator resolves infeasible combinations before execution. This avoids presenting one RIC or one AI agent as the control plane for an entire service chain.

A conflict protocol for autonomous applications

Conflicts may be direct, when two applications change the same parameter; indirect, when separate actions consume the same resource; or objective-based, when energy savings degrade a protected quality target. The control loop first detects overlapping scope, then evaluates signed policy precedence and predicted outcome. Low-risk conflicts can be arbitrated automatically; uncertain or high-impact conflicts move to shadow mode or human approval.

Every decision should retain the competing intents, models, evidence, selected action, rejected alternative and rollback condition. O-RAN's work on conflict mitigation is important precisely because multi-app intelligence does not self-coordinate. Governance is a runtime engineering capability, not a committee added after deployment.

Intelligence must be governed before it becomes autonomous

The North Star requires governed intelligence: systems that can interpret intent, act within bounded authority, explain material decisions, preserve evidence and recover safely. Autonomous-network maturity is therefore measured by scenario, control boundary and verified outcome - never by a general corporate claim that a network is autonomous.

From maturity claims to reproducible assurance

Industry validations demonstrate progress in bounded domains, while academic evidence remains concentrated in architectures, simulations and testbeds. Neither evidence class alone establishes safe end-to-end autonomy across a heterogeneous operator estate.

Publication-grade autonomy claims should name the domain, scenario, delegated authority, observation period, fallback path and measured business outcome. A level label without those fields is not comparable evidence.

Exhibit 2 | Autonomy deployment gates

Gate	Pass condition	Stop condition
State	Inventory and telemetry are current	Unknown or conflicting topology
Intent	Objective and constraints are machine-testable	Ambiguous priority or ownership
Verification	Twin/canary outcome within bounds	Unquantified uncertainty
Authority	Action is within signed policy	Material cross-domain impact without approval

Gate	Pass condition	Stop condition
Recovery	Rollback is tested	Irreversible change or missing baseline

Note: Moderion Advisory governance framework.

Executive action agenda

- Define the enterprise intent model, including hard constraints, competing objectives and accountable owners.
- Build a trusted data and topology foundation, wrap legacy domains through governed adapters, and use observe-shadow-bounded-execute gates before expanding autonomous scope.
- Require digital-twin testing, canary deployment, observability and automated rollback for material changes.
- Create an independent conflict-management and assurance function across agents and vendor domains.
- Report production outcomes and incidents, not only maturity levels and demonstration counts.

Conclusion

The autonomous operator is an institutional design as much as a technical design. Machine speed can improve resilience and economics only if authority, evidence and accountability travel with every action. The winning operators will not be those that remove humans fastest, but those that make human judgment scalable through governed autonomy. In the wider series, this capability is treated as an enabler of the Total Connectivity Economy rather than as an end state in itself.

Selected evidence and references

- T019 | Multiple authors. A Comprehensive Survey of Large AI Models for Future Communications: Foundations, Applications, and Challenges. 2026. IEEE Communications Surveys & Tutorials. <https://ieeexplore.ieee.org/document/11370829/>
- T020 | Zou et al.. Large language models in 6G from standard to on-device networks. 2026. Nature Reviews Electrical Engineering. <https://www.nature.com/articles/s44287-025-00239-6>
- T021 | Multiple authors. AI-Native O-RAN Architectures for 6G: Toward Real-Time Adaptation, Conflict Resolution, and Efficient Resource Management. 2026. IEEE Transactions on Network and Service Management. <https://ieeexplore.ieee.org/document/11483133/>
- T022 | Multiple authors. PACIFISTA: Conflict Evaluation and Management in Open RAN. 2025. IEEE Transactions on Mobile Computing. <https://ieeexplore.ieee.org/document/11007770/>
- T023 | Multiple authors. X5G: An Open, Programmable, Multi-Vendor, End-to-End, Private 5G O-RAN Testbed. 2025. IEEE Transactions on Mobile Computing. <https://ieeexplore.ieee.org/document/11039666>
- T024 | Multiple authors. A Survey on Intent-Driven End-to-End 6G Mobile Communication System. 2025. IEEE Communications Surveys & Tutorials. <https://ieeexplore.ieee.org/document/11017620/>
- T034 | Multiple authors. Research on Digital Twin Technology for Future 6G. 2025. Intelligent and Converged Networks. <https://ieeexplore.ieee.org/document/10949695/>
- T042 | Boston Consulting Group. The AI-First Telco. 2026. BCG. <https://www.bcg.com/publications/2026/the-ai-first-telecommunications-strategy>
- T043 | McKinsey & Company. AI-driven telecom networks. 2026. McKinsey. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/issue-brief-ai-driven-telecom-networks>
- P001 | IEEE. A Comprehensive Survey of Large AI Models for Future Communications: Foundations, Applications, and Challenges. 2026. 10.1109/COMST.2026.3660844. <https://ieeexplore.ieee.org/document/11370829/>
- P002 | IEEE. AI-Native O-RAN Architectures for 6G: Toward Real-Time Adaptation, Conflict Resolution, and Efficient Resource Management. 2026. 10.1109/TNSM.2026.3684880. <https://ieeexplore.ieee.org/document/11483133/>
- P003 | IEEE. PACIFISTA: Conflict Evaluation and Management in Open RAN. 2025. 10.1109/TMC.2025.3570632. <https://ieeexplore.ieee.org/document/11007770/>
- TM Forum. Level 4 autonomous-network shift. <https://www.tmforum.org/news-insight/newsroom/autonomous-networks-leadership-forum-spotlights-level-4-shift>
- TM Forum. A regional guide to autonomous networks progress. https://info.tmforum.org/rs/021-WLD-815/images/TM_Forum_-_A_regional_guide_to_autonomous_networks_progress.pdf?version=0
- TM Forum. Living Network Phase III. <https://www.tmforum.org/catalysts/projects/C26.0.913/living-network-phase-iii>
- TM Forum. Autonomous Networks Level solution packs and validation caveat. <https://www.tmforum.org/solutions/anl-solution-packs>
- ETSI. GS ZSM 002: Zero-touch network and Service Management Reference Architecture. https://www.etsi.org/deliver/etsi_gs/ZSM/001_099/002/01.01.01_60/gs_ZSM002v010101p.pdf
- ETSI. GR ZSM 005: Means of Automation, including brownfield and classic services. https://www.etsi.org/deliver/etsi_gr/ZSM/001_099/005/01.01.01_60/gr_zsm005v010101p.pdf
- ETSI. Zero-touch network and Service Management. <https://www.etsi.org/technical-groups/zsm/>
- TM Forum. Autonomous Networks Mission and maturity taxonomy. <https://www.tmforum.org/missions/autonomous-networks>

- TM Forum. IG1251 Autonomous Networks Reference Architecture. <https://www.tmforum.org/resources/introductory-guide/ig1251-autonomous-networks-reference-architecture-v1-0-0/>
- O-RAN Alliance. O-RAN Conflict Mitigation v1.00 overview. <https://www.o-ran.org/blog/74-new-or-updated-o-ran-technical-documents-released-since-july-2024>
- O-RAN Alliance. dApp use cases and requirements v2.21. <https://mediastorage.o-ran.org/ngrg-rr/nGRG-RR-2024-10-dApp%20use%20cases%20and%20requirements.pdf>
- PR010 | A Survey on Intent-Driven End-to-End 6G Mobile Communication System. IEEE Communications Surveys & Tutorials 28, 882-915 (2025). <https://doi.org/10.1109/COMST.2025.3575041>
- PR011 | PACIFISTA: Conflict Evaluation and Management in Open RAN. IEEE Transactions on Mobile Computing 24(10), 10590-10605 (2025). <https://doi.org/10.1109/TMC.2025.3570632>
- PR012 | Zou, H. et al. Large language models in 6G from standard to on-device networks. Nature Reviews Electrical Engineering 3, 123-134 (2026). <https://doi.org/10.1038/s44287-025-00239-6>
- PR013 | TM Forum. Scaling to AN Level 4 to unlock network value (2026). <https://inform.tmforum.org/research-and-analysis/reports/scaling-to-an-level-4-to-unlock-network-value>